Erschienen in: Journal of Peace Research. 2017, 54(2), pp. 125-140
https://doi.org/10.1177/0022343316630359

Using night light emissions for the prediction of local wealth

Nils B Weidmann

Department of Politics and Public Administration, University of Konstanz

Sebastian Schutte

Zukunftskolleg & Department of Politics and Public Administration, University of Konstanz

Abstract

Nighttime illumination can serve as a proxy for economic variables in particular in developing countries, where data are often not available or of poor quality. Existing research has demonstrated this for coarse levels of analytical resolution, such as countries, administrative units or large grid cells. In this article, we conduct the first fine-grained analysis of night lights and wealth in developing countries. The use of large-scale, geo-referenced data from the Demographic and Health Surveys allows us to cover 39 less developed, mostly non-democratic countries with a total sample of more than 34,000 observations at the level of villages or neighborhoods. We show that light emissions are highly accurate predictors of economic wealth estimates even with simple statistical models, both when predicting new locations in a known country and when generating predictions for previously unobserved countries.

Keywords

economic data, night lights, spatial prediction

Introduction

In many developing countries, official economic statistics are either not available or of poor quality (Jerven, 2013). This may create problems for cross-national research on political violence that has established strong links between economic conditions and civil war (Hegre & Sambanis, 2006). These limitations, however, may be much more serious in disaggregated analyses at the subnational level, since data requirements are considerably higher. Are wealthier regions predominantly affected by conflict? What is the economic damage of violence across the regions of a country, or the local impact of development aid on economic recovery? To reach solid conclusions on these important questions, it is necessary that empirical researchers base their analyses on reliable economic datasets with high temporal and spatial resolution. These statistics are available for many developed countries – in Europe, for example, via the GEOSTAT project (EUROSTAT, 2015) – but typically not for those countries that conflict researchers are most interested in.

In order to overcome this problem, researchers have resorted to alternative data sources. When it comes to economic data, nighttime illumination observed from satellites has been proposed as an alternative measurement approach. So far, however, these attempts have mostly been limited to coarse resolutions; we know, for example, that night light patterns track economic growth at the national level. How far can we increase the resolution of this approach? Is it possible to predict economic wealth at the village or the neighborhood level solely from the level of nighttime illumination? In this article, our aim is to show that it is. Our approach is different in several ways from the majority of works on conflict prediction, and therefore also from most of the other contributions to this special issue. First, rather than predicting violence or war, we predict a covariate that is frequently required in conflict studies: economic

Corresponding author:
nils.weidmann@uni konstanz.de

wealth. Therefore, we see prediction as an improvement of the data collection process for conflict studies, before any analyses of the actual outcome of interest are conducted. Second, and more importantly, our innovation as regards prediction is the use of an alternative data source, not a new prediction model or methodology. The latter typically constitutes the focus of current research in prediction (see for example Beger, Dorff & Ward, 2014, or other contributions to this special issue). In contrast, we show that progress can also be made by using new data in conjunction with simple models.

In our analysis below, we find night lights to be good predictors of wealth at the local level. Across the countries we analyzed, the correlation between night light emission and wealth is on average 0.73, and can be as high as 0.87. In order to test the accuracy of our predictions out-of-sample, we set up two prediction tasks. The first is *within-country*: given a training set with data on wealth and night lights for a number of locations in a country, can we predict the wealth of unseen locations based only on their night light illumination? Predictive performance is very high, indicating that night lights data have great potential for subnational analyses. The second, *cross-national* task is to predict subnational wealth for new countries, training the models on data for other countries. This task is more difficult, because absolute levels of night light emissions can vary considerably across countries (Kyba et al., 2014). However, we show that with a simple normalization procedure, we can also reach a reasonable predictive performance in this second task.

Our article starts by briefly reviewing the literature on the use of night lights in the social sciences, before describing our data and methodology. We then present our results, starting with simple descriptives and moving on to our two prediction tasks, within-country and across-country. Last, we summarize our findings and outline ways in which they can be used in actual research.

Existing work

The use of night lights as a proxy for economic variables typically assumes that nighttime illumination corresponds to wealth through one of at least three channels. First, access to the power grid (or a power generator) requires financial investment, which is likely to be made by people with the necessary resources. Second, night lights indicate economic activity, which can lead to higher levels of wealth for the people involved (Henderson, Storeygard & Weil, 2011). Third, nighttime illumination (street lamps) can be a result of preferential treatment by the state for certain societal groups (Hodler & Raschky, 2014). Whatever mechanism we assume, high light emissions should be correlated with high levels of wealth. This assumption, however, is not unproblematic. For example, economic activity may not benefit the people living at the location where it occurs – bright commercial centers in cities may be inhabited by poor people. Also, the amount of light emitted by economic activity may not scale directly with the benefits it generates for the local population: oil refineries, for example, are typically illuminated at night, but require few staff and do not coincide with residential areas.

Therefore, the question of whether night lights correlate with wealth is an empirical one. Existing analyses have tried to assess the use of these data at different levels of analysis. Early work conducted at the country level reveals a clear correlation between the area illuminated at night and economic output (Elvidge et al., 1997). This correlation alone, however, does not tell us whether night lights track economic activity, since the correlation could simply be due to country size – on average, larger countries have larger economies, but also emit more light at night. This is why subsequent work has examined this relationship further. Henderson, Storeygard & Weil (2011) show that *changes* in night lights track economic *growth*, which provides strong support for the night lights–wealth relationship.

Similar results hold for economic output of subnational units of analysis (states or provinces), which have been approximated using night light patterns. An article by Sutton, Elvidge & Ghosh (2007) conducts such an analysis for four countries (China, India, Turkey, and the USA), showing that night lights track economic output also at this finer level of resolution. More recently, Chen & Nordhaus (2011) present a global study that compares night light emissions to economic output measured at the level of 1-degree (approx. 100 km by 100 km) grid cells. One of the most detailed analyses so far was carried out by Mellander et al. (2013) for Sweden using fine-grained official socio-economic data on businesses and individuals. This study clearly shows the limitations of night lights-based analyses, in particular in developed countries. In these countries, a key problem of night light measurement – top-coding (see below) – makes these data much less useful as a predictor of wealth.

Because of the number of promising results, night lights data have seen some adoption in the social sciences and conflict research in order to approximate economic variables. Shortland, Christopoulou & Makatsoris (2013) use fine-grained data on night light emission to estimate the economic impacts of violence in Somalia. A number of works on inequality rely on night lights data, due to the fact that economic data from other sources are not available at a global scale. Most of these works combine the night lights data with geographic data on ethnic settlement regions, making it possible to link illumination to particular groups. For example, Alesina, Michalopoulos & Papaioannou (forthcoming) construct an indicator for inequality between ethnic groups based on these data. Cederman, Weidmann & Bormann (2015) employ a similar approach, but triangulate the night lights data with other sources in order to compensate for some of their weaknesses. Kuhn & Weidmann (2015) demonstrate how this approach can be adapted to yield estimates of intragroup inequality. Finally, Hodler & Raschky (2014) use nighttime illumination to study ethnic favoritism.

Despite the adoption of night light data as an alternative data source at the subnational level, there have been few attempts to validate these data at high levels of analytical resolution. While many of the above works assume that night lights are good predictors of wealth at the local level, this has not been thoroughly tested. An article by Min et al. (2013) describes an attempt to infer electrification at the local level from nighttime lights, using a survey conducted in Senegal and Mali. The results show that night lights seem to track electric infrastructure reasonably well, but cannot tell whether the frequent assumption that they proxy wealth actually holds. An approach close to our analysis below is presented in Michalopoulos & Papaioannou (2013). They show that night light emissions closely correlate with wealth as measured by the Demographic and Health Surveys (DHS), which we also use below. Their analysis is limited to four countries. Most importantly, however, they do not examine performance when predicting out of sample. Below, we expand our set of countries significantly and show that simple changes to the computation of the night lights-based indicator compared to Michalopoulos & Papaioannou (2013) lead to considerable improvements in its predictive power.

Data and prediction models

We use times-series data on night light emissions from the Defense Meteorological Satellite Program's Operational Linescan System (DMSP-OLS), provided by the US National Oceanic and Atmospheric Administration. The data are provided as annual rasters with a resolution of 30 arc seconds, which corresponds to approximately 1 km. We use the 'stable lights' version of the data, which has non-stable light sources such as forest fires removed (National Geophysical Data Center, 2014a). For each raster point, the dataset provides a so-called 'digital number' (DN) between 0 and 63 that encodes the level of radiation. Since our units of observation have different sizes (see below), we compute the radiation per square km emitted from these units. Note that when comparing night light emissions across time, it is necessary to calibrate the data in order to eliminate variation due to atmospheric noise or sensor drift (Wu et al., 2013). However, since our comparisons are exclusively cross-sectional, we use uncalibrated data.

Night lights data suffer from a number of potential problems that we can discuss only briefly here. One of these problems is 'overglow', or the fact that light can spill over from one cell to adjacent cells (Doll, 2008). This leads to a loss of precision in the night lights data – for example, it is difficult to detect a dark neighborhood in a city due to the spillover of light from bright neighborhoods nearby. In our application, however, this measurement error works against us. If we achieve high prediction accuracy *despite* overglow, this suggests that accuracy should be even higher if we were able to correct this problem. Two other related problems concern the limited scale that is used for encoding the light emissions. As stated above, the DN is limited to values between 0 and 63. This can result in top-coding (different high values end up at the top of the scale) or bottom-coding (low emissions are recorded as 0). This problem can only be addressed by using a radiance-calibrated version of the DMSP-OLS dataset with a larger value range (National Geophysical Data Center, 2014b). We repeat some parts of our analysis with this dataset. However, since the radiance-calibrated dataset includes non-stationary light sources that have been removed in our default dataset, it may not be an optimal choice for our analysis.

The validation data for our prediction analysis are taken from the Demographic and Health Surveys (DHS), a comprehensive survey project conducted in developing countries.[1] Using a standardized data collection methodology across different countries and years, these surveys provide detailed information on individual and household characteristics. Importantly for our purpose, the surveys include a wealth index for the socio-economic status of a household. Since the estimation of wealth using only financial income is difficult, the DHS wealth index is based on assets, for example bikes or radios, house construction materials, and water access (Rutstein

[1] See http://dhsprogram.com.

& Johnson, 2004). The wealth index at the household level is reported as a number from 1 to 5, which correspond to survey quintiles (with 1 corresponding to the poorest 20%, etc.).

While the asset-based approach of estimating wealth can circumvent problems of applicability and availability of conventional measures, it is not without problems. In particular, asset-based indices seem to capture household expenditure only imperfectly (Howe et al., 2009). Also, another problem seems to be that the concept of a 'household', and therefore the assets it owns, can vary, which of course affects the quality of the index (Randall & Coast, 2015). Still, due to the fact that the DHS index is essentially the only one available across a large sample of developing countries and with high spatial resolution, it is the best choice for our study.

The surveyed households are randomly sampled starting from a survey cluster location (typically, a village or a neighborhood of a city), which is reported with longitude/latitude coordinates. For the vast majority of clusters, the cluster location is measured using GPS, which guarantees a high level of precision. For some surveys, however, the location is based on information from gazetteers, which is not precise enough for our purpose. Gazetteers typically only include a single point for each village or city, and are therefore unable to distinguish between different neighborhoods. We believe that the considerable uncertainty caused by the use of gazetteers prevents the use of gazetteer-based cluster locations for a fine-grained analysis such as ours. For that reason, we retain only those clusters with coordinates measured by GPS for our analysis. In order to preserve anonymity, the cluster locations in the DHS data are randomly distorted with an artificial error of 2 km in urban locations, and 5 km in rural locations. The type of cluster (urban/rural) is known, so we can incorporate this locational uncertainty into our analysis. The cluster location allows us to geo-reference the survey results to a particular location on the globe and compare it with night light emissions at that location from the same year the survey was conducted. Since we do not have location information about the households, we conduct our analysis at the level of clusters and assign each cluster point the average wealth level of the households that belong to it.

The amount of night light radiation for that cluster is calculated by drawing a circular buffer around the cluster centroid whose radius corresponds to the artificial error (2 km or 5 km), and by computing the *average light radiation per square km* within that buffer for the year the survey was conducted. This ensures that the real cluster location is within the buffer. In order to test whether larger buffer sizes as used, for example, in Michalopoulos & Papaioannou (2013) improve predictive accuracy, we also used alternative buffer radii of 5 km, 10 km, and 20 km. If, however, night light emissions at the local level are a good predictor of local wealth, predictive accuracy should decrease as we increase the buffer size.

We model the relationship between night light emissions and wealth using three simple statistical models. While more complex models can possibly achieve better performance, we believe that simplicity is an asset: it reduces the number of free parameters and therefore alleviates the risk of overfitting. In addition, the simple models we employ are readily implemented in all statistical toolkits, so users can generate similar predictions without the need for specialized software. As we have stated above, the unit of analysis in all three models is the survey cluster, using the level of light emissions (per square km) as the predictor variable, and the average level of household wealth in the cluster as the response. Our first model is a simple linear OLS regression model. In our second model, following earlier work such as Chen & Nordhaus (2011) or Michalopoulos & Papaioannou (2013), we log-transform the night lights variable. The third model is a generalized additive model (GAM), fitted using the standard settings from the *mgcv* package in R. The GAM computes the optimal functional form of how night lights relate to wealth from the data, which in many cases improves prediction beyond standard parametric approaches.

Results

Our analysis is limited to the years 2003 to 2012 and includes only those surveys where both geographic information and the wealth index were provided. As mentioned above, we only use clusters where the location was measured with a GPS, since alternative sources of geographic coordinates are not precise enough. Our sample for this analysis consists of a total of $N = 34{,}047$ cases (survey clusters) from 56 surveys in 39 countries (see Appendix A for details).

Descriptives

We start with a simple example. Figure 1 shows an image of the 2008 night lights data for Pakistan (left panel). On the right panel, we enlarge the area around Hyderabad and add the corresponding survey data. In total, there are five sampling clusters in the area. The two clusters in the downtown area of the city (two values at the bottom) have bright light values, and at the same time have high wealth scores close to the upper end of the 1–5 scale. As

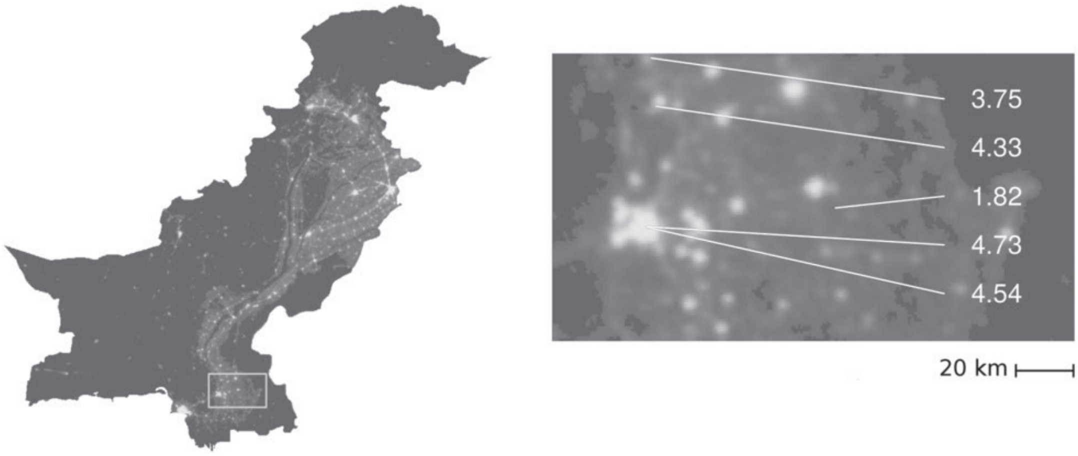

Figure 1. Night light emissions for Pakistan and Hyderabad (2008)

Left panel: Night light emissions for Pakistan. The rectangle shows the location of the Hyderabad area. Right panel: Zoom in on the Hyderabad area. The city corresponds to the bright spot on the left. The lines indicate the locations of the survey clusters and their respective average household wealth values.

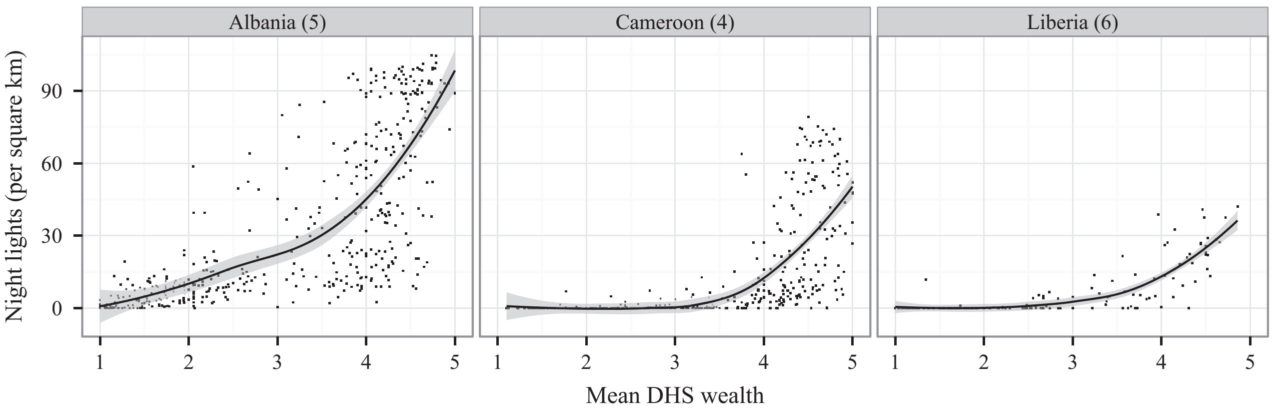

Figure 2. Scatterplot of night lights and wealth for three countries/surveys

The lines show the smoothed Loess fits.

the poor cluster (1.82) shows, dark areas correspond to low wealth values. The topmost point indicates that night lights seemingly capture also intermediate levels of wealth – a moderately bright spot has an intermediate level of wealth (3.75) according to the survey data. Therefore, at least in this initial example, there seems to be a high correlation between night light emissions and wealth at the local level.

In order to examine whether this relationship holds more generally, we provide scatter plots for night light emissions (using our default buffer sizes of 2 km for urban and 5 km for rural areas) and wealth for each survey. Figure 2 shows these plots for a selection of three countries (Albania, Cameroon, and Liberia). Complete results are available in Appendix B. The figure provides two important insights. First, the relation between night lights and wealth seems to have a characteristic 'hockey stick' shape, where locations with little to no emissions are generally in the poorest two quintiles of the distribution. The plots in the appendix show that this holds across the vast majority of all countries/surveys examined in our analysis. Second, while the general shape of the relationship seems to be comparable across countries, the absolute levels of night lights are not. For example, in Albania the richest household clusters emit roughly three times the amount of night lights as those in Liberia. This is something to keep in mind when predicting wealth *across* countries, as we need to take these differences in magnitude into account.

Before we proceed to derive precise wealth values from night lights data, we start with a simple bivariate test. In particular, we analyze how well we can rank-

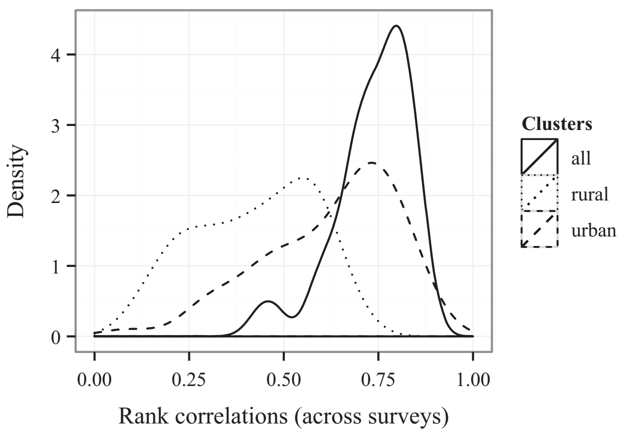

Figure 3. Distributions of rank correlation coefficients across the 56 surveys

Solid line: all survey clusters; dashed line: only urban clusters; dotted line: only rural clusters.

order survey clusters with respect to their wealth when using the night lights they emit. To this end, we compute rank correlations by survey between the night lights and wealth values. Figure 3 shows the distribution of correlations across the 56 surveys in our sample, for all clusters (solid line), and separately for urban (dashed line) and rural clusters (dotted line).

Figure 3 indicates that night light emissions correlate highly with wealth at the local level. Using all clusters for validation, the correlation is on average 0.73 and reaches values up to 0.87 (solid line). In Appendix C, we list the rank correlations by survey wave. While it is difficult to discern a clear pattern, some of the smaller countries in the sample (Rwanda, Burundi) achieve low correlations, which could suggest that night lights-based prediction performs worse in small countries due to denser settlement patterns. Differences between urban and rural locations account for much, but not all, of the correlation between night lights and wealth: rural locations typically have higher levels of poverty but will at the same time have lower levels of light. Therefore, when analyzing the sample of urban or rural clusters separately, the correlation between night lights and wealth should be lower as compared to the full sample, which is what we see in Figure 3 (dashed and dotted lines). For urban locations, the correlation is on average only 0.62, whereas it decreases even further for rural clusters (0.42). Together, these results suggest that simply rank-ordering locations in a country according to their night light emissions can give researchers a good idea of their relative wealth status. But can we do better? In the next section, we proceed to modeling wealth as a function of night lights using different types of regression models.

Predicting wealth within countries

Our first task is to predict the wealth of household clusters within countries based on a simple model fitted on data from the same country. Since this task operates exclusively with data from the same country for a single round of prediction, we avoid the problem of different magnitudes in night light emissions that we describe above. Our prediction exercise proceeds as follows. For each survey in our sample, we conduct a ten-fold cross validation by randomly splitting the survey's data points into ten bins. We then perform ten runs where each time, a single bin is left out for testing purposes, and the model is estimated on the remaining bins and evaluated on the test bin. For each of the 56 surveys, we compute the mean absolute error between the model predictions and the actual values. Alternative measures of predictive accuracy are of course possible, but we prefer the mean absolute error due to its intuitive interpretation.[2] We use the three models described above – the linear (LM), log-linear (LLM), and generalized additive model (GAM) – as well as the four different buffer sizes for the night lights (our default choice of 2/5 km, 5 km, 10 km, 20 km). For comparison, we also include a linear model that includes only the urban/rural predictor as a dummy variable. Figure 4 shows the results.

Comparing the different models, the results in Figure 4 clearly demonstrate that the logarithmic specification and the semi-parametric one outperform the linear model. As the figure shows, both achieve consistently lower prediction errors. For our default buffer size of 2 km (urban areas) and 5 km (rural areas), the average prediction error decreases by about 20% (from 0.70 to 0.57 for the log-linear model, and 0.55 for the GAM). In other words, this means that we predict the wealth quintile of a household (cluster) using night lights with a low error of about half a quintile overall. Across all surveys examined, in the worst case our predictions would be about one quintile off for a linear model (0.97), but only by about 25% less (0.73) when using a GAM. This analysis shows that night light emissions do not scale linearly with wealth, since other specifications (log-linear and GAM) provide much better results.

Buffer size seems to matter a lot. As we increase the buffer size from the minimum radius of 2/5 km to 20 km, we see a continuous decrease in predictive performance (an increase in the prediction error). This is a

[2] For deterministic predictions as in our case, the continuous rank probability score (CRPS) that is used elsewhere in this special issue is equivalent to the mean absolute error that we use.

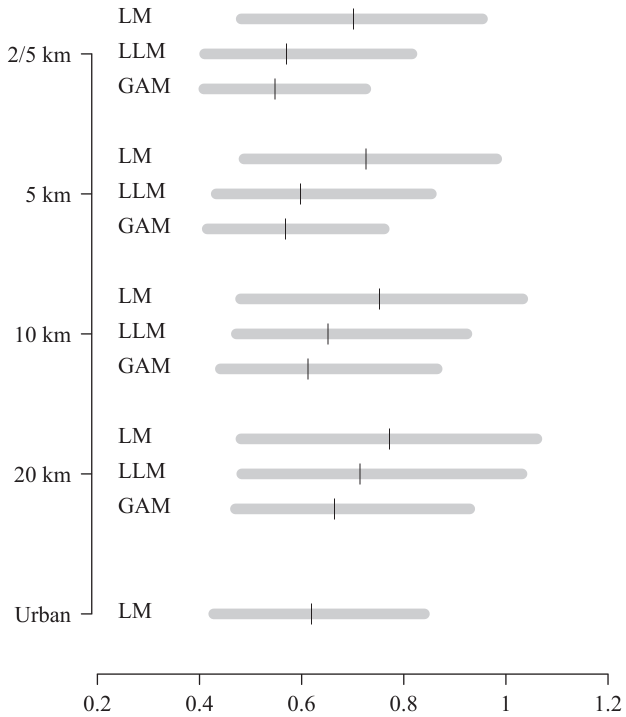

Figure 4. Predicting wealth from night light emissions, within country

The figure shows the averages (black lines) and ranges (minimum maximum, grey bars) of the mean absolute prediction errors across the 56 surveys in our sample. Lower values indicate better performance.

clear indication that the local levels of night lights – and not the emissions across a region – seem to matter for prediction. This also shows that the arbitrary choice of a 10 km or 20 km buffer size as in Michalopoulos & Papaioannou (2013) may not be optimal; in fact, we can increase the resolution considerably and obtain much better results. Our decision to use different buffer sizes for urban and rural areas turns out to be a good one – the 2/5 km buffers perform best for prediction if we use either a log-linear model or a GAM.

Is the quality of our predictions entirely dependent on night lights picking up differences between urban and rural areas? To find out, we conduct our prediction exercise also with a simple linear model that only includes the urban/rural dummy from the DHS (Figure 4, last row). Compared to the non-linear models based on night lights, the results are consistently worse. For example, the mean prediction error across the surveys increases to 0.62, which is about a 15% increase from the GAM. Also, a separate prediction analysis of only the urban clusters reveals the considerable predictive power of the night lights approach: if we repeat the above analysis for urban clusters only, we obtain consistently *lower* error rates as compared to the overall sample (see results in Table D.I in the appendix). Thus, we can conclude that there is a significant added value when predicting wealth from night light emissions, and that night lights do not simply reflect urban/rural divides.

How much does the problem of the top- and bottom-coded night light values in the standard DMSP-OLS data affect our predictions? We repeat our within-country prediction with the radiance-calibrated night lights dataset (National Geophysical Data Center, 2014b). Due to the more limited temporal coverage of these data, this reduces our sample to 35 surveys. Table D.II in the appendix reports the results. For comparison, we also report the results when using the standard night lights time series on the reduced sample. We see that there is a slight improvement for the radiance-calibrated data, which achieve slightly lower prediction errors for all three models. Therefore, if the reduced temporal coverage of these data is not an issue, these data can be a better choice in a within-country prediction task.

Predicting wealth across countries

Our first prediction task represents a situation where for a given country, we have high-resolution data both for night lights and wealth and can use these data to train a model and predict wealth for new locations in this country. For a prediction exercise of this kind, the data requirements are still considerable, as we need at least a certain amount of fine-grained data on wealth for the country we are interested in. In many cases, this will not be possible. For that reason, we now explore whether we can use a prediction model trained on other countries to predict wealth in a previously unseen country. As we have seen above, the greatest obstacle to this is the different magnitudes of night light emissions we observe across countries, where the amount of light emitted by people of a certain socio-economic status can differ by up to a factor of three. For a new country, we usually do not know what the typical amount of emissions would be for households at the center or the upper end of the wealth scale.

Therefore, when generating predictions for new countries, we can address this question in two ways. The first and more complex way would be to model the country-specific magnitude of the night lights emissions explicitly. This could be done with a simple multilevel model, where the intercept or the effect of night lights on wealth is a function of country attributes. We choose a second and, we believe, more straightforward way by normalizing the night light scores from a country. Rather than

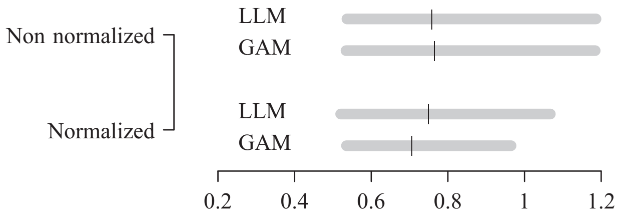

Figure 5. Predicting wealth from night light emissions, across countries

The figure shows the averages (black lines) and ranges (minimum maximum, grey bars) of the mean absolute prediction errors across the 56 surveys, both for non normalized and normalized night lights values as predictors. As above, lower values indicate better performance.

using raw night lights values as above, we use the *estimated percentiles of the respective values in the distribution of all night light values in the survey*. In other words, for our normalization we take the entire distribution of night light emission in a survey, and assign each cluster its percentile (0–100) in this distribution. This is effectively a rank-ordering of the night light values, which, as we have seen above, tracks the wealth distribution closely.

In our above analysis, we have seen that the log-linear and the GAM model using the 2/5 km night light buffers work best, which is why we restrict this analysis to these models and buffer sizes. We set up a cross-validation task similar to the one above. We leave out each survey in turn and estimate the model on the remaining data, before predicting the wealth values for the survey that is left out. This way, we generate predicted values for each cluster and compute again the mean absolute error for these values by survey.

Figure 5 shows some interesting results. First, as we expected, the cross-national prediction task is much harder than within-country prediction. As we have shown above, for the latter we can achieve absolute errors of around 0.5 on the wealth scale. For the former, the absolute error is 0.7 and higher, regardless of what model we choose. Second, our proposed night lights normalization seems to work. Prediction using non-normalized values (upper pair of bars) results in considerably higher error as compared to normalized values (lower pair). This is also reflected in the mean values. For the log-linear model, there is only a very small reduction in the mean error across surveys, but normalization can reduce the worst outcomes, reducing the maximum error from 1.19 to 1.07 (a 10% decrease). For the GAM, the decrease is more pronounced (the average goes down by about 8% from 0.77 to 0.71, the maximum from 1.18 to 0.97, about 20%). In sum, this exercise shows that predicting wealth values in previously unseen countries is possible with mean errors of around 1 in the worst cases, which means that our predictions can be up to a quintile off. Whether this error margin is acceptable or too high depends on the respective application.

We also repeat the cross-national prediction with the radiance-calibrated data. As Table E.I in the appendix shows, these data reach consistently lower performance as compared to the standard night lights data. One reason for this may be that the former do not exclude non-stationary lights, which could be a potential source of error in our analysis.

Conclusion

In this article, we made one of the first attempts to predict wealth at the local level from satellite imagery of night light emissions in developing countries. While earlier research has been able to demonstrate that night lights can serve as a useful proxy for economic variables at the level of large spatial entities (countries or first-tier administrative units), it has not been tested systematically whether this approach can also be applied at a higher geographic resolution. Basing our analysis on the freely-available DMSP-OLS night lights data, we examine how the amount of light emitted from a village or a city neighborhood relates to its wealth. Using wealth estimations available from survey data for validation, we find a considerable correlation between the level of night light emissions and household wealth. Some of this is due to night lights picking up differences between urban (bright) and rural (dark) locations, but we have also shown that night lights add significant predictive value even when controlling for these differences. Thus, in short, our analysis demonstrates that local wealth can be predicted with reasonable accuracy if only the household's location is known. This works both for predicting new cases in a country where training data is available and, to a lesser extent, for new countries, training the prediction model only on data from others.

Our results have at least two implications. First, it is possible – with a small margin of error – to predict the wealth of a household if only the household location is known. When no alternative data sources are available (as is often the case in developing countries), this is a powerful and widely applicable approach. For example, micro-level research on political violence often relies on events with spatial coordinates, which can be used to compute night lights emissions as a proxy for wealth. Also, night lights may prove useful as more social science data are being collected via mobile devices, where the device's location is known (through GSM tower triangulation or from GPS). Second, our results underline two important caveats when using night lights: cross-national differences and functional form. The level of night lights emissions from 'rich' locations varies tremendously across countries. This is an issue that cannot be ignored in cross-national studies. We have shown that a simple normalization can alleviate this problem in the cross-national prediction task, but alternative ways of dealing with it are certainly possible. Also, wealth does not scale linearly with night lights; rather, we find that a log-linear model provides the best fit and is able to achieve much higher predictive accuracy compared to a linear model.

In this study, however, we have only begun to move the analysis of night lights to a high resolution. Future research will have to expand upon this work in at least two directions. The first, and most important, question is about selecting locations for measurement. Our results do not imply that one can choose any location on the globe and use its level of light emissions to gauge economic wealth. For obvious reasons, this would be impossible, as we do not know whether low emissions indicate low population or low economic wealth (or both). Therefore, in order to develop an indicator with global applicability, we would have to have precise population density estimates that allow us to compute a per capita level of emissions. Existing data sources such as Gridded Population of the World (Center for International Earth Science Information Network (CIESIN) & Centro Internacional de Agricultura Tropical (CIAT), 2005) or the LandScan database (Oak Ridge National Laboratory, 2008) are either too coarse or have been computed partly using night lights data, in which case they do not give us much independent information. A second question arises regarding the predictive methodology we use. Can we improve predictions by using more powerful models? Our aim in this study was to start simple and see what is possible with standard models, but one can likely push these results even further. In general, however, our study has shown that the use of night lights data need not be limited to large geographic units, and holds tremendous value also for local-level analyses.

Acknowledgements

The authors would like to thank Lukas Kawerau for research assistance, and Philipp Hunziker, Espen Rød, Sabine Otto, and Peter Selb, as well as two anonymous reviewers and the special issue editors for comments.

Funding

Financial support from the Alexander von Humboldt Foundation (Sofja Kovalevskaja Award) and the EU FP7 Marie Curie Zukunftskolleg Incoming Fellowship Program (Grant #291784) is gratefully acknowledged.

References

Alesina, Alberto; Stelios Michalopoulos & Elias Papaioannou (forthcoming) Ethnic inequality. *Journal of Political Economy*.

Beger, Andreas; Cassy L Dorff & Michael D Ward (2014) Ensemble forecasting of irregular leadership change. *Research & Politics* 1(3) (http://rap.sagepub.com/content/1/3/2053168014557511).

Cederman, Lars Erik; Nils B Weidmann & Nils Christian Bormann (2015) Triangulating horizontal inequality: Toward improved conflict analysis. *Journal of Peace Research* 52(6): 806 821.

Center for International Earth Science Information Network (CIESIN) & Centro Internacional de Agricultura Tropical (CIAT) (2005) Gridded population of the world v3 (GPWv3). Palisades, NY: CIESIN, Columbia University (http://sedac.ciesin.columbia.edu/gpw/).

Chen, Xi & William D Nordhaus (2011) Using luminosity data as a proxy for economic statistics. *Proceedings of the National Academy of Sciences* 108(21): 8589 8594.

Doll, Christopher (2008) *CIESIN Thematic Guide to Night-Time Light Remote Sensing and its Applications*. Center for International Earth Science Information Network (http://sedac.ciesin.org/binaries/web/sedac/thematic guides/ciesin nl tg.pdf).

Elvidge, Christopher D; Kimberley E Baugh, Eric A Kihn, Herbert W Kroehl, Ethan R Davis & Chris W Davis (1997) Relation between satellite observed visible near infrared emissions, population, economic activity and electric power consumption. *International Journal of Remote Sensing* 18(6): 1373 1379.

EUROSTAT (2015) The GEOSTAT project (http://web.ornl.gov/sci/landscan/).

Hegre, Håvard & Nicholas Sambanis (2006) Sensitivity analysis of empirical results on civil war onset. *Journal of Conflict Resolution* 50(4): 508 535.

Henderson, Vernon; Adam Storeygard & David N Weil (2011) A bright idea for measuring economic growth. *American Economic Review* 101(3): 194 199.

Hodler, Roland & Paul A Raschky (2014) Regional favoritism. *Quarterly Journal of Economics* 129(2): 995 1033.

Howe, Laura D; James R Hargreaves, Sabine Gabrysch & Sharon RA Huttly (2009) Is the wealth index a proxy for consumption expenditure? A systematic review. *Journal of Epidemiology and Community Health* 63(11): 871 877.

Jerven, Morten (2013) *Poor Numbers: How We Are Misled by African Development Statistics and What To Do About It*. Ithaca, NY: Cornell University Press.

Kuhn, Patrick & Nils B Weidmann (2015) Unequal we fight: Between and within group inequality and ethnic civil war. *Political Science Research and Methods* 3(3): 543 568.

Kyba, Christopher CM; Stefanie Garz, Helga Kuechly, Alejandro Sánchez de Miguel, Jaime Zamorano, Jürgen Fischer & Franz Hölker (2014) High resolution imagery of earth at night: New sources, opportunities and challenges. *Remote Sensing* 7(1): 1 23.

Mellander, Charlotta; Kevin Stolarick, Zara Matheson & José Lobo (2013) Night time light data: A good proxy measure for economic activity? CESIS Electronic Working Paper Series number 315. Royal Institute of Technology (https://static.sys.kth.se/itm/wp/cesis/cesiswp315.pdf).

Michalopoulos, Stelios & Elias Papaioannou (2013) Pre colonial ethnic institutions and contemporary African development. *Econometrica* 81(1): 113 152.

Min, Brian; Kwawu Mensan Gaba, Ousmane Fall Sarr & Alsassane Agalassou (2013) Detection of rural electrification in Africa using DMSP OLS night lights imagery. *International Journal of Remote Sensing* 34(22): 8118 8141.

National Geophysical Data Center (2014a) DMSP OLS nighttime lights time series, version 4 (http://ngdc.noaa.gov/eog/dmsp/downloadV4composites.html).

National Geophysical Data Center (2014b) Global radiance calibrated nighttime lights (http://ngdc.noaa.gov/eog/dmsp/download radcal.html).

Oak Ridge National Laboratory (2008) Landscan global population database (http://www.efgs.info/geostat/).

Randall, Sara & Ernestina Coast (2015) Poverty in African households: The limits of survey and census representations. *Journal of Development Studies* 51(2): 162 177.

Rutstein, Shea Oscar & Kiersten Johnson (2004) The DHS wealth index (http://dhsprogram.com/pubs/pdf/CR6/CR6.pdf).

Shortland, Anja; Katerina Christopoulou & Charalampos Makatsoris (2013) War and famine, peace and light? The economic dynamics of conflict in Somalia 1993 2009. *Journal of Peace Research* 50(5): 545 561.

Sutton, Paul C; Christopher D Elvidge & Tilottama Ghosh (2007) Estimation of gross domestic product at sub national scales using nighttime satellite imagery. *International Journal of Ecological Economics & Statistics* 8(S07): 5 21.

Wu, Jiansheng; Shengbin He, Jian Peng, Weifeng Li & Xiaohong Zhong (2013) Intercalibration of DMSP OLS night time light data by the invariant region method. *International Journal of Remote Sensing* 34(20): 7356 7368.

Appendix

(A) List of countries and surveys

Table A.I shows a list of countries and surveys used for the analysis, along with the number of survey clusters, the total number of households and their average cluster wealth, and its standard deviation included in each survey. In total, our analysis includes 56 surveys from 39 countries, with 821,857 households surveyed in 34,047 clusters.

Table A.I. List of countries and surveys included in the analysis

Country	*DHS wave*	*No. of clusters*	*No. of households*	*Cluster wealth (mean)*	*Cluster wealth (SD)*
Albania	5	450	7,999	3.125	1.249
Angola	6	110	3,708	3.344	1.265
Bangladesh	4	355	10,324	3.050	0.948
Bangladesh	5	361	10,400	3.140	1.023
Bangladesh	6	598	17,086	3.034	0.994
Bolivia	5	998	19,526	3.033	1.246
Burkina Faso	4	390	8,863	2.997	1.132
Burkina Faso	6	505	12,706	3.136	1.120
Burundi	6	376	8,596	3.048	0.980
Cambodia	5	548	14,017	2.942	1.066
Cameroon	4	464	10,413	3.185	1.150
Cameroon	6	577	14,189	3.109	1.217
Colombia	6	4,416	45,249	2.674	1.278
Congo Democratic Republic	5	289	8,563	2.925	1.207

(continued)

Table A.I. (continued)

Country	*DHS wave*	*No. of clusters*	*No. of households*	*Cluster wealth (mean)*	*Cluster wealth (SD)*
Dominican Republic	5	1,414	32,110	2.581	1.059
Egypt	5	2,540	39,440	3.003	1.132
Ethiopia	5	528	13,550	3.171	1.322
Ghana	4	410	6,221	3.053	1.259
Ghana	5	404	11,574	2.982	1.230
Guinea	5	291	6,191	3.108	1.195
Guyana	5	284	4,934	2.761	0.996
Haiti	5	332	9,773	2.993	1.195
Honduras	6	1,127	20,967	2.860	1.192
Indonesia	4	1,319	31,393	2.730	1.177
Kenya	4	399	8,542	3.301	1.267
Kenya	5	397	9,033	3.280	1.305
Lesotho	4	378	8,037	2.991	1.119
Liberia	5	425	17,661	2.783	0.949
Liberia	6	148	4,110	2.645	1.178
Madagascar	5	585	17,578	3.053	1.243
Madagascar	6	266	8,066	3.171	1.307
Malawi	4	520	13,649	2.839	0.817
Mali	5	390	12,476	3.129	1.075
Moldova	5	399	11,066	3.324	1.208
Morocco	4	480	11,513	2.967	1.287
Mozambique	5	270	6,097	3.342	1.149
Mozambique	6	609	13,899	3.310	1.219
Namibia	5	491	9,036	3.087	1.156
Nepal	5	260	8,707	3.050	1.182
Nepal	6	289	10,826	3.119	1.228
Nigeria	4	360	7,187	3.143	1.243
Nigeria	5	886	34,070	2.940	1.219
Nigeria	6	239	5,895	3.236	1.244
Pakistan	5	762	7,4659	3.043	1.132
Peru	5	704	19,278	2.180	0.857
Philippines	4	656	10,161	2.988	1.019
Rwanda	5	702	30,139	3.063	0.681
Rwanda	6	492	12,540	2.942	0.938
Senegal	4	366	7,224	2.871	1.139
Senegal	5	267	7,763	2.597	1.200
Senegal	6	358	7,236	2.553	1.141
Swaziland	5	270	4,756	3.391	1.084
Tanzania	4	315	5,927	2.988	1.170
Tanzania	5	466	8,350	3.137	1.091
Tanzania	6	1,023	29,908	3.049	0.914
Zimbabwe	5	396	9,234	3.023	1.201
Zimbabwe	6	393	9,442	3.224	1.210

(B) Scatter plots of night lights and wealth

The three plots in Figures B.1a, B.1b, and B.1c show scatter plots of night lights and wealth, by survey. Each survey is identified by the country and the wave (in parentheses).

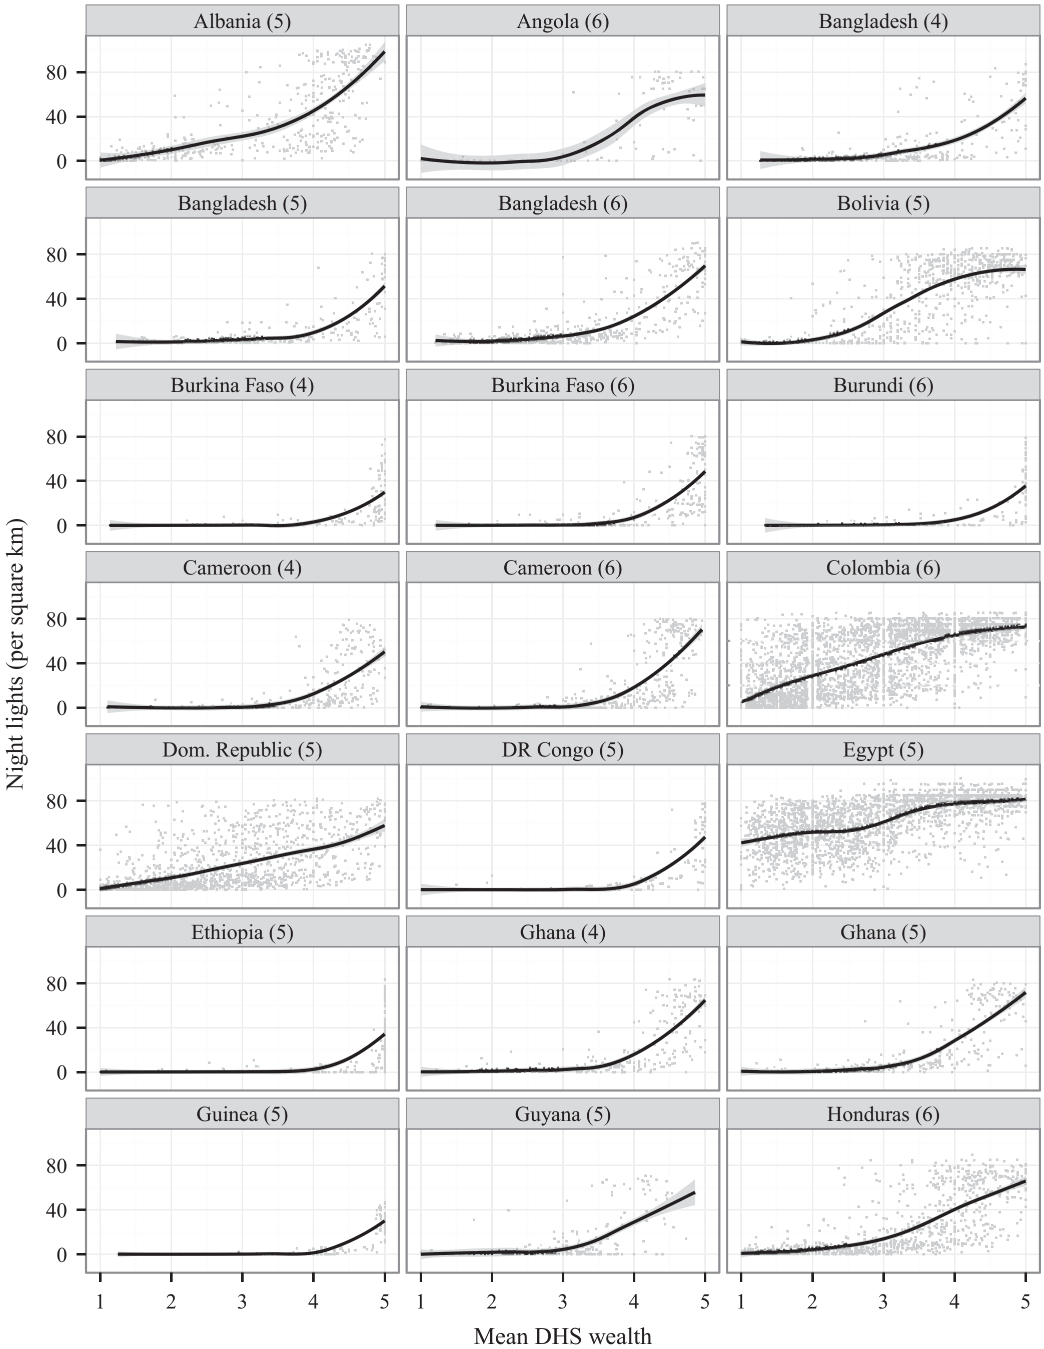

Figure B.1a. Scatter plots of night lights and wealth (part 1)

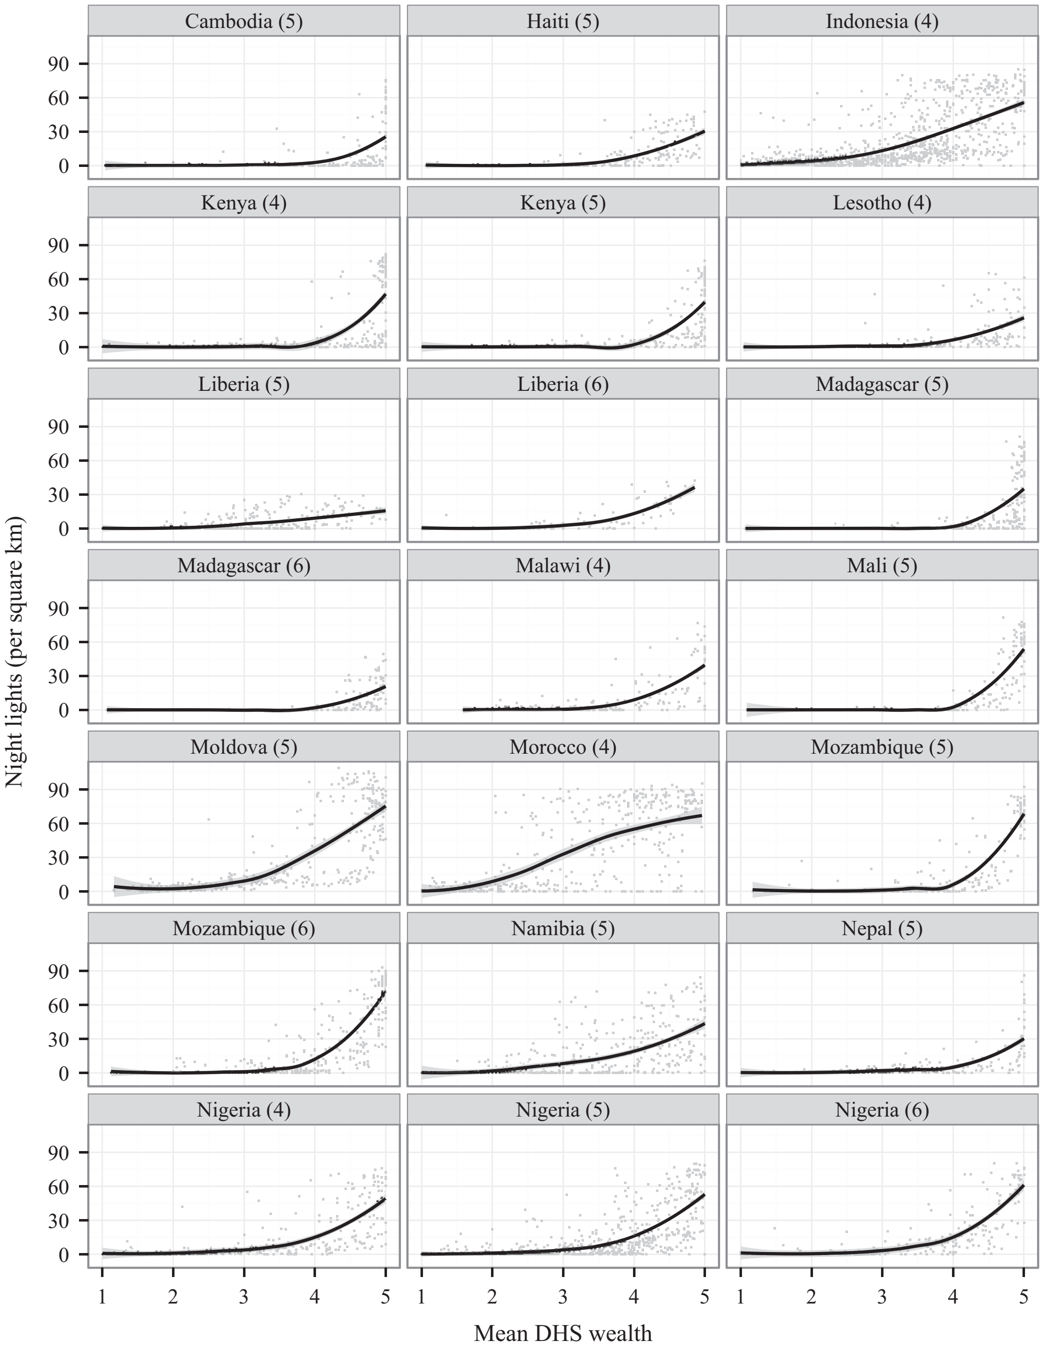

Figure B.1b. Scatter plots of night lights and wealth (part 2)

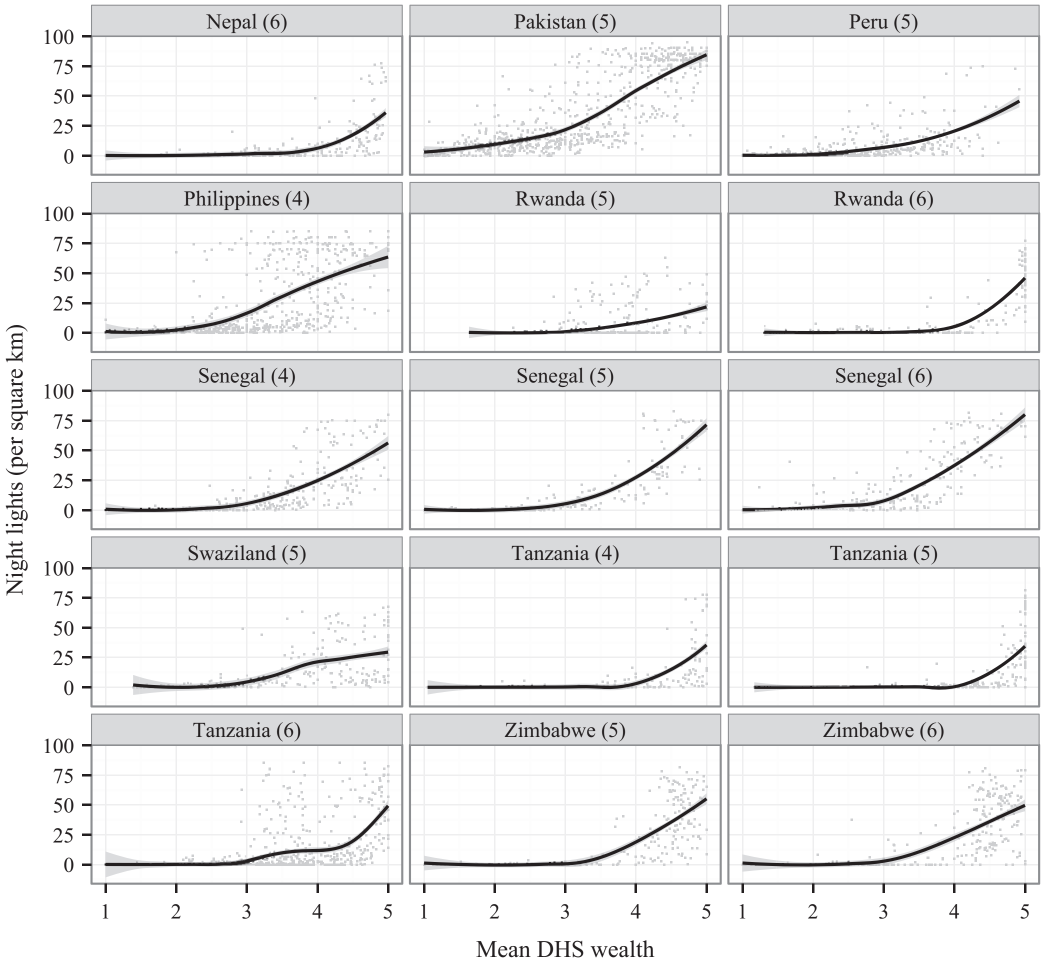

Figure B.1c. Scatter plots of night lights and wealth (part 3)

(C) Rank correlations: Full list of results

Table C.I. List of rank correlations between night lights and survey based estimates across all surveys

Country	*DHS Wave*	*Rank Correlation*
Rwanda	5	0.43
Rwanda	6	0.46
Malawi	4	0.48
Cambodia	5	0.58
Guyana	5	0.59
Liberia	5	0.60
Burundi	6	0.62
Tanzania	6	0.62
Morocco	4	0.65
Egypt	5	0.65
Namibia	5	0.66
Peru	5	0.68
Dominican Republic	5	0.68
Bangladesh	4	0.69

(continued)

Table C.I. (continued)

Country	*DHS Wave*	*Rank Correlation*
Burkina Faso	4	0.70
Tanzania	4	0.70
Madagascar	5	0.70
Ethiopia	5	0.70
Tanzania	5	0.70
Bangladesh	5	0.71
Madagascar	6	0.72
Haiti	5	0.72
Nepal	5	0.73
Mali	5	0.73
Burkina Faso	6	0.74
Kenya	4	0.74
Bangladesh	6	0.74
Lesotho	4	0.75
Congo Democratic Republic	5	0.76

(continued)

Table C.I. (continued)

Country	*DHS Wave*	*Rank Correlation*
Swaziland	5	0.77
Nigeria	5	0.77
Philippines	4	0.77
Nepal	6	0.77
Kenya	5	0.78
Liberia	6	0.78
Zimbabwe	5	0.79
Indonesia	4	0.79
Guinea	5	0.79
Colombia	6	0.80
Angola	6	0.80
Zimbabwe	6	0.80
Nigeria	4	0.81
Cameroon	4	0.81
Pakistan	5	0.82
Cameroon	6	0.82
Bolivia	5	0.82
Honduras	6	0.82
Mozambique	5	0.82
Albania	5	0.83
Nigeria	6	0.83
Senegal	6	0.84
Ghana	4	0.84
Senegal	5	0.86
Mozambique	6	0.86
Ghana	5	0.86
Moldova	5	0.86
Senegal	4	0.87

(D) Within-country prediction: Additional results

This section repeats the within-country prediction analysis for urban clusters only. Table D.I shows the averages, minimum, and maximum values of the mean average prediction error across all surveys. We use again three models (LM, LLM, and GAM) and 2/5 km buffers around the cluster centroids. For comparison, we also show the values for all clusters that are displayed visually in Figure 4, top three bars.

Table D.I. Mean absolute prediction error for all clusters (top panel) and urban clusters only (bottom panel)

	LM	*LLM*	*GAM*
All clusters			
Min.	0.4806	0.4097	0.4080
Mean	0.7022	0.5717	0.5493
Max.	0.9560	0.8166	0.7258
Only urban clusters			
Min.	0.1879	0.1666	0.1296
Mean	0.4614	0.4173	0.4117
Max.	0.8119	0.7431	0.6833

We also present results using radiance-calibrated night lights data that remove the problem of top- and bottom-coding in the standard DMSP-OLS time series (for all clusters). Table D.II shows again the results using the standard DMSP-OLS data (top panel) for comparison, computed on the same set of surveys for which we have radiance-calibrated data. The bottom panel reports the results for the Global Radiance-Calibrated Nighttime Lights time series.

Table D.II. Mean absolute within country prediction error using the standard DMSP OLS stable lights (top panel) and the radiance calibrated lights (bottom panel)

	LM	*LLM*	*GAM*
DMSP-OLS stable lights			
Min.	0.4836	0.4172	0.4133
Mean	0.6988	0.5696	0.5496
Max.	0.9572	0.8014	0.7274
Radiance-calibrated nighttime lights			
Min.	0.2522	0.2415	0.2371
Mean	0.5825	0.4840	0.4656
Max.	0.8549	0.6700	0.6684

(E) Cross-national prediction: Additional results

This section repeats the cross-national prediction analysis using radiance-calibrated imagery. Table E.I shows the results similar to those presented in Figure 5 in the main paper, but only for the subset of surveys where we have radiance-calibrated data (top panel). The bottom panel shows the prediction errors when using the radiance-calibrated imagery.

Table E.I. Mean absolute prediction error using the standard DMSP OLS stable lights time series (top panel) and the radiance calibrated version (bottom panel)

	LLM (raw values)	*LLM (normalized)*	*GAM (raw values)*	*GAM (normalized)*
DMSP-OLS stable lights				
Min.	0.5373	0.5536	0.5243	0.5587
Mean	0.7446	0.7354	0.7542	0.6989
Max.	1.2297	0.9973	1.2431	0.8831
Radiance-calibrated nighttime lights				
Min.	0.5379	0.4838	0.5297	0.4825
Mean	0.9109	0.7989	0.9022	0.7983
Max.	1.6740	1.2587	1.6911	1.2595

NILS B WEIDMANN, b. 1976, PhD in Political Science (ETH Zurich, 2009); Professor of Political Science, University of Konstanz, Germany (2012) and head of the Communication, Networks and Contention research group; research interests: violent and nonviolent conflict, with a particular focus on the impact of communication and information technology.

SEBASTIAN SCHUTTE, b. 1979, PhD in Political Science (ETH Zurich, 2013); Marie Curie Fellow, University of Konstanz, Germany (2014); research interests: conflict processes in civil wars and communal violence with a focus on reactive dynamics.