Published in final edited form as:

Am Econ Rev. 2011 May ; 101(3): 194–199. doi:10.1257/aer.101.3.194.

A Bright Idea for Measuring Economic Growth

J. Vernon Henderson,

Brown University and NBER

Adam Storeygard, and

Brown University

David N. Weil

Brown University and NBER

Most readers of this article are already familiar with images of the “Earth at Night,” which are really montages showing satellite views of visible lights.¹ Looking at such images, it is easy to see that the brightness of visible lights is strongly related to both population density and income per capita. The effect of population density shows up in the visibility of urban agglomerations, coastlines, river valleys, and even the higher density of population along the paths of highways and railroads. Similarly, the Northeastern United States is far brighter than the mountain West, even though they have similarly levels of income. The effect of income is similarly easy to see. Sharp economic boundaries, such as the inter-Korean border show up plainly. India, with a population density slightly higher than Japan, is significantly dimmer,

The observation that income per capita is one of the determinants of visible light suggests that one might potentially use visible light to help *measure* income. In this paper, we explore the usefulness of these night lights data to economists studying issues related to economic growth and the national or local level. In particular, we ask the extent to which lights data can augment or entirely replace data on Gross Domestic Product.²

There are a number of obstacles to using lights visible from space to measure income, broadly falling into two categories. First, the relationship between economic activity and the true amount of light emanating from the Earth’s surface is not constant across time and space. Light visible from space is a byproduct of consumption and production activities, but light is hardly produced in a fixed ratio to output. Some productive activities such as steelmaking produce a lot of light; others such as software design produce very little. Countries or regions may differ in the fraction of their economic activity that takes place after dark (for example, Las Vegas vs. Salt Lake City). Patterns of settlement (for example, whether people live in multi-storey buildings), the availability of hydropower, etc. may all affect how much light is visible from space for a given level of income and population. Second, true light is imperfectly measured by the satellites. Humidity, reflectivity, and time periods excluded from coverage because of sunlight, moonlight, and cloud cover

¹Because of limitations on print quality, we do not include any such pictures here. A good online version is at http://visibleearth.nasa.gov/view_rec.php?id=1438.

²Henderson, Storeygard, and Weil (2009) and Chen and Nordhaus (2010) present more detailed explorations of this same issue.

contamination all differ across the globe. Sensor scale for recording lights in practice varies across satellites and as satellites age over time. If one is measuring economic growth, some of the location-specific factors are differenced out. The rest we treat as measurement error in the relationship between GDP growth and light growth.

We note that while we focus on measuring total income growth of regions, one might be also interested in growth in income per capita. We don't pursue that here although we have explored the issue by using population growth numbers to supplement lights growth. Population is both conceptually and practically easier to measure than is income, so it is also feasible to alternatively focus on income growth per capita.

Despite problems, the night lights data have several advantages. The first has to do with the nature of measurement error in lights data. A number of studies have recently reminded economists that much of the data on GDP growth in developing countries is plagued by serious measurement error. For example, Johnson et al. (2009) study revisions in the Penn World Tables, a standard measure of GDP. They find that comparing version 6.1 of that data, released in 2002, with version 6.2 released in 2006, the standard deviation of the change in countries' average growth over the period 1970–1999 was 1.1% per year (which is very large relative to the average growth rate of 1.56% per year). This indicates that at least one of the measures, and presumably both, contained a great deal of error. We do not claim that the measurement error in the night lights data is smaller than that in conventional data, but we do think that the two forms of measurement error are poorly or not correlated. It is well known that combining two problematic measures can produce a composite with smaller measurement error than either of them.

A second advantage of the night lights data is that they are available for regions for which standard GDP measures are not. These include sub-national units such as cities and regions, as well as entities that cross national borders, like biomes. Economic analysis of growth and of the impacts of policies and events on cities and regions of many countries is hindered by a complete absence of any regular measure of local economic activity. Much of the interesting variation in economic growth takes place within rather than between countries. However, for the vast majority of economics research, “empirical analysis of growth” has become synonymous with use of national accounts data. The night lights data offer a tool that allows the researcher to set aside this limitation.

The rest of this paper is organized as follows. In Section I, we briefly discuss the night lights data. Section II outlines our statistical approach to combining data from lights with conventional measures of GDP, and presents results of an exercise applying our technique to a set of low and middle income countries. Section III discusses an application beyond the measurement of national income. We refer the reader to our companion paper, Henderson, Storeygard, and Weil (2009), for more detailed discussion of many of the technical and statistical issues raised here.

I Night Light Data

Our starting point is data on visible light emanating from earth as captured by a series of US Air Force weather satellites that observe every location sometime between 8:30 and 10:00 pm local time. Daily observations are excluded if the light visible is influenced by moonlight, auroral activity, forest fires, or summer months when the sun is still visible, and also exclude days in which cloud cover blocks light from reaching space. Data from all orbits of a given satellite in a given year are averaged over all valid nights to produce a satellite-year dataset. Datasets currently exist for 30 satellite-years covering the years 1992 to 2008. Data are reported for grid cells measuring 30 arc-seconds per side. This corresponds to a grid cell area of approximately 0.86 square kilometers at the equator. For a given satellite-year, we have data on approximately 181 million grid squares that fall on land. Light emanation is measured by a number on a linear scale between zero and 63, which corresponds closely but not exactly to the quantity of light (as measured by true radiance) reaching the satellite. A small fraction of observations (less than 0.1%) are top coded.³

To give a sense of what these data look like, Table 1 presents data for a few representative countries. For reference, we also include data on GDP per capita at PPP, population density, and the percentage of the population living in urban areas (all taken from the World Development Indicators).

Figure 1 shows a simple scatter plot of the change in log measured total GDP and the change in log measured light. Specifically, we look at the difference between the average for the first two years of our sample (1992 and 1993) and the last two (2005 and 2006). There is clearly a strong relationship between lights growth and measured income growth, although there are a number of interesting outliers.

II Statistical Model and National Income Results

Let y be the growth in true GDP, z the growth of GDP as measured in national income accounts, and x the growth of observed light. The variance of true income growth is σ_y^2 . For country j , we assume that there is classical measurement error in GDP growth as recorded in national income accounts:

$$z_j = y_j + \varepsilon_{z,j}, \quad (1)$$

where the variance of ε_z is denoted σ_z^2 . Later we allow for the variance of the measurement error in national income data, σ_z^2 , to vary among country groups.

The relationship between growth of lights and growth of true income is given by

³For an overview of this data, see Christopher Elvidge, et al. (2003).

$$x_j = \beta y_j + \varepsilon_{x,j}, \quad (2)$$

where the variance of ε_x is denoted σ_x^2 . The error term in this equation represents all the ways in which growth of measured light is not proportional to growth of income. Because we don't think measurement error in GDP is related in any consistent fashion to the error in the equation determining observable light, we assume that $\text{Cov}(\varepsilon_x, \varepsilon_z) = 0$.

While equation (2) specifies a production relationship between income and lights, in most applications we are concerned with using lights growth to predict income growth. As such, for predictive purposes, we want a regression of growth of income on growth of lights, or

$$z_j = \hat{\psi} x_j + e_j, \quad (3)$$

Even though the OLS parameter $\hat{\psi}$ is a biased estimate of the inverse of the elasticity of lights with respect to income, the fitted values from this regression (which we call $\hat{z}_j$) are best-fit proxies for income.

Fitted values of income growth based on lights growth can be created for sub-national units such as cities as well as for countries in which there are no available income data. When income data are available, the income growth proxy from satellite data can be combined with a national account measure to arrive at a composite estimate of income growth which will have lower error than either one separately.

$$\hat{y}_j = \lambda z_j + (1 - \lambda) \hat{z}_j, \quad (4)$$

As we show in Henderson, Storeygard and Weil (2010), the weight λ which minimizes variance of the gap between true GDP and our proxy $\hat{y}$ will depend on the variances of measurement error in national income accounts and satellite data, the variance of true GDP growth, and slope coefficient β in the lights-income structural equation, none of which are directly observed. In the data, only three moments for $\text{var}(x)$, $\text{var}(z)$, and $\text{cov}(x, z)$ are observed, which are not sufficient to pin down these four parameters.

To supplement these moments, we can use information on variations in the degree of measurement error in conventional GDP data among country groups. For example, if we have a set of "good data" countries in which there is zero measurement error in conventional GDP data, and a set of "bad data" countries in which this is not the case, then we can use the good data countries to pin down measurement error in bad data countries, and then estimate the other parameters required to do our optimal weighting for the bad countries (the optimal weight on the satellite data for good data countries is zero in this case). Alternatively, if we are willing to specify a "signal to noise" ratio in the good data countries, then the data will

tell us the corresponding ratio in bad data countries as well as optimal weights for both good and bad data countries.

We implement this approach examining a panel of 118 low and middle income countries where we have ratings of their statistical capacity on a scale of 0–10 by the World Bank (2002). We designate the 30 countries with ratings 0–3 as bad data countries and the rest as good data countries. We assume that the country groups differ in the quality of their conventional GDP measures, but not in the degree of measurement error in lights as a proxy for GDP nor in the economic relationship between lights and income.

2.1 Predicting income growth

To start, we fit income and lights data to get a prediction of income from lights. We implement three statistical models. The first looks at income growth and annual fluctuations, with country and time fixed effects to control for geographic and cultural differences, changes in scaling of satellite sensor readings, and the like. The second adds a country time trend, to focus on growth fluctuations about a country's growth path. The third is a long difference to focus on long term economic growth.

The top panel of Table 2 shows results. Note the estimated coefficient $\hat{\psi}$ is almost the same across the three formulations. The bottom panel of the table reports a regression of the squared residuals from the basic regression on a dummy variable for good data countries. The coefficient is negative, indicating lower measurement error and a tighter fit for good data countries, and is statistically significant except in the last column where the sample is small. If we add a dummy for very good data countries with a World Bank rating of 6–10 then it is negative and significant in all cases.

2.2 Improving true income growth estimates

Next we turn to improving income growth estimates following our model for the 30 poor data countries, with World Bank ratings of 3 or less. We examine the long term growth case. If we assume that GDP growth is perfectly measured in the national income accounts of the good data countries, the data imply that the optimal value of λ for bad data countries is 0.56. In other words, the optimal estimate of GDP growth is a combination of 56% growth as measured in the national income accounts and 44% growth as measured by satellite data.⁴ If we assume that national income accounts data is not perfect even in the good data countries, then the implied weight on satellite data is even higher. For example, if we assume that ratio of signal to total variance in then good country data is 0.9, then the implied weight on satellites data in forming an optimal measure of growth is 15% in good data countries and 52% in bad data countries.

To give an example of how satellite data can be used to improve national accounts estimates of GDP growth, we focus on the 30 bad data countries, and use the assumption just mentioned that the signal to total variance ratio in good countries is 0.9. The average absolute value of the gap between growth over the period 1992/3 and 2005/6 as measured by

⁴Under the assumption of zero measurement error in national income accounts data for the good data countries, we can also solve for the value of the structural coefficient β that relates lights growth to income growth in equation (2). The value we get is close to one.

national accounts and growth over the same period as measured by our optimal proxy is 0.75% per year. Some of the gaps are extremely large. For example, in Myanmar, national accounts data say that GDP grew at an annual rate of 10%, while our optimally weighted estimate is 6.5% per year; in Burundi, the WDI data say that annual GDP growth was -0.71% while our optimally weighted estimate is 1.1% per year. In both these cases, there is reason, beyond the night lights data, to suspect that GDP is particularly poorly measured in the WDI. Burundi experienced civil war and reconstruction for much of the period for which we have satellite data, while the economy in Myanmar was largely autarkic and non-market, with a governing regime that would not be averse to exaggerating GDP growth.

3. Other Applications

We have done a variety of applications at the sub-national level focused on sub-Saharan Africa. We note some examples here. In one we look at the debate about whether coastal as opposed to hinterland location is critical to economic growth, as advocated by Gallup, Sachs, and Mellinger, 1999. We follow Gallup, Sachs, and Mellinger (1999) in defining coastal region as being within 100 kilometers of an ocean or ocean navigable river and look at the period 1992/93 to 2007/08. This is a period of enormous expansion of world trade and rising oil prices, in a context with poor hinterland transport infrastructure. We sum over all pixels in coastal versus hinterland areas of sub-Saharan Africa for each of the two years. We find that inland lights grew by 0.133 log points *more* than coastal areas. Using the $\hat{\psi}$ coefficient of 0.329 from the long difference estimation in column 3 of Table 2, the lights data imply that the increase in total GDP inland was 4.4% greater than on the coast – a difference of 1/3 of a percent per year. While we cannot say anything about the long-run benefits over centuries of being on the coast, during a period of rapidly growing trade, coastal areas in Africa grew more slowly than non-coastal areas. There may be a number of competing explanations for this, but the supposed inherent advantage of coastal location for growth in this period in sub-Saharan Africa does not dominate other forces that were at work.

In another application based on Storeygard's work, we looked at 14 sub-Saharan countries with a coastal primate city and asked whether cities in their hinterland that are further from the coast suffer a greater disadvantage as oil prices and hence the costs of trade rise during the 1992–2008 time period. For these 14 countries there are 275 cities with population over 20,000 for which Storeygard defines polygon lit areas and the best road links to the coastal primate city. During this time period oil prices rose from about \$20/ barrel in the 1990s to over \$90/barrel. In estimating the growth in lights with city and time fixed effects and a growth trend for all primate versus non-primite cities, Storeygard finds that, by interacting oil prices with distance to the coastal primate city, an increase in oil prices from \$20 to \$90 per barrel leads to 12% lower lights for a city one standard deviation further from the coastal primate city. At a maximal distance he finds a 60% reduction in lights

In another exercise, we looked at whether rainfall helped economic prosperity in African cities. The presumption is that the exogenous changes in rainfall don't affect urban productivity per se. Rather they improve productivity in the immediate agricultural hinterland of the city. Farmers then spend money in their immediate urban area for both

consumption and production purposes, which improves urban income levels. The exercise is a test of whether cities in sub-Saharan Africa benefit from growth of their hinterlands (as opposed to cities just augmenting rural growth). We find that increased rainfall in the current and up to three years previously leads to increased city lights in a sample of 541 cities, allowing for city and time fixed effects and country growth trends. The precise specification is

$$\ln(x_{jkt}) = \sum_{i=0}^k \beta_i r_{jk,t-i} + \alpha_j + \lambda_t + \gamma_l t + \varepsilon_{jkt}, \quad (5)$$

where j is city, t is time and l is country and there is a lag structure to annual rainfall, r . Rainfall is measured for a 30 km buffer around each city, where cities are lit polygon areas of sub-Saharan cities with recorded population in the time period. Given the lag structure and rainfall data for 1995 to 2007 we look at city light growth over 9 years.

Results are in Table 3, where in column 2, increased rainfall from up to three years prior raises this year's city income. The insignificant coefficient for a 4th year back is shown. In column 3, as a counterfactual, increases in rainfall in the next year do not affect this year's lights. As an interpretation of results, a one standard deviation increase in rainfall in the current year leads to a 10% increase in city lights, a testament to the role of agricultural in city growth in less developed countries.

References

- Chen, Xi, Nordhaus, William D. The Value of Luminosity Data as a Proxy for Economic Statistics. NBER Working Paper. 2010 Aug.:16317.
- Elvidge, Christopher D., Hobson, Ruth, Nelson, Ingrid L., Safran, Jeffrey M., Tuttle, Benjamin T., Dietz, John B., Baugh, Kimberley E. Overview of DMSP-OLS and Scope of Applications. In: Mesev, Victor, editor. Remotely Sensed Cities. London: Taylor & Francis; 2003. p. 281-99.
- Gallup, John Luke, Sachs, Jeffrey, Mellinger, Andrew. Geography and Economic Development. International Regional Science Review. 1999; 22(2):179-232.
- Henderson, Vernon, Storeygard, Adam, Weil, David N. NBER Working Paper 15199. Jul. 2009 Measuring Economic Growth from Outer Space.
- Johnson, Simon, Larson, William, Papageorgiou, Chris, Subramanian, Arvind. NBER Working Paper 15455. 2009. Is Newer Better? Penn World Table Revisions and Their Impact on Growth Estimates.
- World Bank. Building Statistical Capacity to Monitor Development Progress. Washington: 2002.

Figure 1.
GDP versus lights: long differences

Table 1

Night Light Data for Representative Countries

	USA	France	India	Brazil	Ghana	Niger
Average Visible Light	4.66	8.63	2.89	0.634	0.494	0.0146
Gini for Visible Light	0.847	0.602	0.723	0.969	0.969	0.999
Percentage Cells Unlit	66.2	29.1	55.8	94.3	94.2	99.8
Percentage Cells Top Coded	0.5801	0.4835	0.0954	0.058	0.031	.00048
Population Density (per sq. km)	30.7	107.8	341.0	20.6	86.0	8.90
Percentage Population Urban	79.0	75.8	27.7	81.0	44.0	16.1
GDP per Capita, PPP (2005 \$)	37,953	28,458	1,816	8,046	1,078	593

Table 2

Results for rated low-middle income countries; growth in real GDP (local currency units)

	Fixed effects	Country time trend	Long difference
	(1)	(2)	(3)
ln(lights/area)	0.308 *** [0.037]	0.270 *** [0.043]	0.329 *** [0.046]
Constant	n/a	n/a	0.365 *** [0.028]
Observations	1953	1953	113
Number of Countries	118	118	113
(Within-country) R-sq	0.780	0.903	0.301
Country fixed effects	Yes	Yes	No
Year fixed effects	Yes	Yes	No
Country time trend	No	Yes	No
Regression of squared residuals:			
Good data dummy	-0.0055 *** [0.0017]	-0.0017 * [0.0010]	-0.029 [0.018]

Robust standard errors in brackets

p<0.01,**
p<0.05,*
p<0.1

In column 3, long differences are formed by averaging the first and last two years of levels data

Table 3

Hinterland rainfall and economic growth of cities

	(1) FE plus trend	(2) FE plus trend	(3) Counter - factual
rain(t)	.071 ** [.031]	.105 ** [.044]	.143 ** [.056]
rain(t-1)		.081 ** [.040]	.127 ** [.050]
rain(t-2)		.129 *** [.041]	.109 ** [.052]
rain(t-3)		.097 *** [.037]	.100 ** [.041]
rain(t-4)		{.015} { [.043]}	
rain(t+1)			.030 [.051]
N	4869	3787	2705
Cities	541	541	541
Within R-sq	0.13	0.11	0.086

p<0.01,**
p<0.05,*
p<0.1