Using Satellite Imagery of Night Lights to Study Patronage and Politics in Africa: A Research Proposal

Elizabeth Carlson, Brian Min and Daniel Posner
UCLA

12 May 2008

Introduction

One of the greatest impediments to a deeper understanding of the link between patronage and politics in Africa (and elsewhere) is the paucity of high quality, cross-national, longitudinal data on the distribution of valued patronage goods. We propose to fill this gap by employing satellite images of the earth at night to identify changes in the extent and coverage of light output at the district level. Since lights visible from space at night are highly correlated with state investments in electrification (Elvidge et al. 1997; Min 2008) and since electrification is highly desired by communities that are not yet electrified, we interpret changes in light output as being indicative of changes in the distribution of a valuable patronage good. To test for the political sources of these changes, we propose to study the relationship between changes in electrification rates and electoral returns in 541 administrative districts in seven African countries.[1] To maximize our leverage, we will focus (at least initially) on countries that have experienced changes in their presidential leadership between 1992 to 2006, where the incentives for governments to favor different constituencies (and thus administrative districts) are likely to have changed over time.[2]

This memo describes the data sources we propose to employ, outlines the theoretical questions that we hope to be able to answer, and raises several possible limitations of our approach. We are at a very early stage in this project and are grateful for feedback and advice of any kind.

Measuring Electrification Using Satellite Data

Across the world, approximately 1.6 billion people lack electricity, with a third of this total living in Africa (International Energy Agency, 2002). More than simply a modern convenience, access to electricity is a life-altering transformation that improves welfare and enables development. Electric light extends a day's productive hours, allowing children to study after the sun has set and enhancing safety at night. Refrigeration allows for the preservation of food and medicines. Electric power enables the development of industries and job creation. Powered water pumps reduce the effort needed to collect clean water. Electrical cooking stoves reduce the amount of time needed to gather wood and other biomass fuels. Electrification is thus a highly desired good. And with so many Africans lacking access to electricity, the question of who gets electrification and who does not is potentially highly instructive about patterns of

[1] The administrative maps onto which we superimpose the satellite nighttime lights data define two different administrative levels (I and II). Aggregation at the level of the administrative II level generates 541 units in our seven countries; aggregation at the administrative I level generates only 58 units.
[2] These countries are Benin, Ghana, Kenya, Malawi, Mozambique, Senegal, and Zambia. In later research, we plan to expand the analyses to include additional cases.

governmental favoritism. The question is: how might we determine which communities have received this highly valued good and which have not?

Various estimates of access to electricity have been made available by international agencies.[3] However, these data suffer from multiple deficiencies, including incomplete geographic coverage, uncertain accuracy, reporting bias, questionable comparability across cases, and availability at levels of analysis that are too aggregated to be useful for studies of sub-national patterns of electricity distribution. Several of these problems are especially severe for Africa.

To address these limitations, we will generate new estimates of national and sub-national variations in electrification with an unbiased procedure using satellite images of the earth at night. Images come from the Defense Meteorological Satellite Program (DMSP), a series of weather satellites that capture high resolution images of night lights across the globe every night.[4] These images are overlaid and processed into a series of annual composite images identifying time stable night lights in each year from 1992 to 2003 (Imhoff et al. 1997; Elvidge et al. 2001).[5] The resolution is 30 arc-seconds (approximately 1km^2 at the Equator) with each pixel encoded with a relative measure of its annual average brightness on a 6-bit scale from 0 to 63. We superimpose the satellite images onto digitized administrative maps to identify all electrified and unelectrified administrative districts across Africa, as well as the amount of light emitted from each district.

Figure 1 shows a reverse-color DMSP image of 2003 nighttime lights in Africa with darker dots indicating more brightly lit areas and white areas on the page indicating darkness. The image, which includes portions of Europe and the Middle East for comparison, reveals that much of Africa is unlit at night. While Sub-Saharan Africa is home to 12% of the global population, it produces only 2% of total light output. Some of the variation across Africa can be explained by geography and settlement patterns: the Sahara and Kalahari deserts as well as the Congolese jungles are predictably dark. States with larger economies have more lit settlements, as evidenced across South Africa, Nigeria, Ghana, and Côte d'Ivoire. The overall picture, however, reveals that most of Africa lives in the dark.

The lack of light across large swaths of Africa reveals a potentially significant limitation of the use of night-lights imagery to detect electrification. Because of limitations in the sensitivity of the DMSP sensors, not all dimly lit regions may be detectable in satellite images. In theory, the DMSP sensors are capable of detecting radiances as low as 10^{-9} watts/cm^2/sr/μm and field checks have revealed that lights from U.S. towns as small as 120 people are detectable (Elvidge et al., 2001). However, because DMSP annual composite images are produced through image processing algorithms designed to remove ephemeral light sources like lightning and fires, it is possible that some very dimly lit (or irregularly lit) areas also get blacked out. Thus annual

[3] See e.g., the United Nations' Energy Statistics Database and the International Energy Agency's *World Energy Outlook* series.
[4] The Defense Meteorological Satellite Program's Operational Linescan System (DMSP-OLS) has been flying in polar orbit since the 1970s, capturing high resolution images of the entire earth each night between 20:00 and 21:30 local time. Captured at an altitude of 830 km above the earth, images reveal concentrations of outdoor lights, fires, and gas flares at a fine resolution of 0.56 km and a smoothed resolution of 2.7 km. Digitized images are available only from 1992.
[5] Annual images from 2004 to 2006 should be available very soon.

composite DMSP images may not unambiguously detect the electrification of all small settlements. We intend to compare our night lights data against specific household surveys with information on electrification to evaluate the severity of this limitation. That said, satellite images capture substantial changes over time and identify many newly lit settlements. Figure 2 superimposes lights from 1992 and 2003, identifying newly lit areas for Ghana and Benin.

We use the DMSP data to construct three separate annualized measures of electrification at the level of the administrative district.[6]

1) *Total light output* sums the light radiance values of each pixel within the administrative district.
2) *Percent area lit* counts the proportion of populated pixels in the administrative district that are lit (we ignore how much light is being generated in each pixel and focus here on the 0-1 distinction between "any" light and "none").
3) *Lit population* sums the population count in all lit pixels within the administrative district. Comparing this figure with the total population of the district provides an estimate of the share of the population living in lit areas, or *proportion lit*.

We use these annual measures to generate estimates of change over time within the administrative district in total light output, percent area lit, and population living in a lit area. We interpret positive changes as being indicative of the district having received a patronage good.

Matching Lights with Election Data

The question of the project is how politics drives the distribution of this important patronage good. Is it provided as a reward for political support? Or as a means of generating new support where it presently is weak? These questions map onto the larger question in the political science literature regarding whether politicians reward core or swing voters: evidence that electrification tends to go to areas that have voted for the ruling party or president would provide empirical support for the first perspective; evidence that it goes to areas that have voted for the opposition would provide support for the latter (particularly if we can establish that these voters, once rewarded with electrification, switch their allegiance to the ruling party or president). To get at these questions, we compare district-level changes in electrification and district-level changes in electoral outcomes over time.

The primary challenge in conducting these (otherwise rather straightforward) analyses lies in merging the nighttime lights data and the electoral data at the same level of analysis. As noted, the satellite data is aggregated at the (administrative II) district level. The challenge is to aggregate electoral data at the same level—and, indeed, it may not be possible to produce electoral results at the administrative II district level in all countries. We believe such

[6] For the measures that require information about the population of the pixel or district, we use the high resolution population map produced by the LandScan project, which we project onto the same 30 arc-second grid as our DMSP data.

disaggregated electoral data are available for the seven countries we will focus on first, but moving beyond this set may be difficult for this reason.

The second problem is that, even if data are broken down below the regional level, they are likely to be reported at the level of electoral constituencies rather than at the administrative II district level. If constituencies are simply sub-sets of the districts, aggregation will not be a major obstacle. Similarly, if we are able to obtain GIS-ready maps of the boundaries of electoral constituencies, merging them with our administrative districts should be doable. However, if constituencies cross-cut district boundaries or can not be readily mapped, it may be possible to aggregate the electoral data only at the regional (administrative I) level. This would have the disadvantage of reducing our sample size, raising the likelihood of ecological inference problems, and (potentially) threatening the comparability of our analyses across countries by forcing us to compare different levels of analyses in each.

There is also the issue of whether doing district-level analysis actually *is* the best strategy given the data. Because the majority of Africa is poorly lit, aggregating at the district level may mean that many of our units will have no light at all (or no changes in light output, % area lit, or % population with lighting). Aggregating to the regional (administrative I) level will make it more likely that there will be meaningful variation on levels of light output between observations. Once we have put the data together we should be in a better position to assess the degree of variation over time at each level of analysis.

Proposed Analyses to Identify Political Patterns in Electricity Distribution

We are in the process of collecting district-level election data, as described above, for all of the listed presidential elections in the seven countries in Table 1 below.[7] For each district, we are then calculating the change in total light output, change in percent area lit and change in population (and proportion of population) living in a lit area that has taken place during each inter-election period. We will then ask: does the level of support for the president in a given district in one election predict changes in the level of total light output, percent area lit, or percent of population living in a lit area at the time of the next election? We will also ask the similar question: does the change in level of support for the president between one election and the next predict changes in the three dependent variables in the subsequent period (essentially, a "does change generate change" rather than "does level generate change" question)? Evidence in the affirmative would be consistent with a model of patronage distribution in which politicians reward their supporters. As noted, the fact that our seven countries include so many instances of presidential turnover (including turnover across parties—indicated in bold in Table 1) should provide large variation in levels of support for the incumbent, and thus plenty of leverage for tackling this question.

[7] These are the presidential elections that took place in a year for which satellite lights data are available. For the Zambia and Benin elections, which took place in October and March of 1991, we will use lights data from 1992.

Table 1: Countries and Elections

	Elections	Change in leader?*
Benin	1991, 1996, 2001, 2006	1991 **(Kerekou to Soglo)**; **1996 (Soglo to Kerekou)**, 2006 (Kerekou to Boni)
Ghana	1992, 1996, 2000, 2004	2000 (Rawlings to Kuofor)
Kenya	1992, 1997, 2002	**2002 (Moi to Kibaki)**
Malawi	1994, 1999, 2004	2004 (Muluzi to Mutharika)
Mozambique	1994, 1999, 2004	2004 (Chissano to Guebuza)
Senegal	1993, 2000, 2007	**2000 (Diouf to Wade)**
Zambia	1991, 1996, 2001, 2006	**1991 (Kaunda to Chiluba)**; 2001 (Chiluba to Mwanawasa)

* **Bold** type indicates that the leadership change entailed a change from one party (and, likely, support coalition) to another.

We will also ask whether districts that experienced the greatest positive change between one election and the next in our various measures of light output and access also see the biggest increase in support for the incumbent in the second contest. Evidence in the affirmative would confirm that voters do in fact reward politicians who have provided them with valuable benefits. These are just two of the many analyses that our data would put us in a position to undertake. We would be grateful for suggestions of others, and for thoughts on the controls we would need to include to make the findings more plausibly reflective of our underlying concerns.

A different set of questions that our data will put us in a position to shed light on (pardon the pun!) involve the distinction between party-based favoritism and ethnic-based favoritism. One way to get at this issue would be to look at cases where the ruling party remains the same but the president changes (these are the cases of change noted in Table 1 that are not in bold), particularly if the new president comes from a different region than the old president (for example, Mwanawasa and Chiluba in Zambia or Guebuza and Chissano in Mozambique). We can then see whether the pattern of electricity distribution remains the same (which would indicate party-based favoritism) or changes to favor members of the president's ethnic group (which would indicate ethnic-based favoritism).

Yet another analysis we are contemplating involves ascertaining whether lights intensity changes in election years, which might provide evidence of a pork electoral cycle. If we find evidence for such a cycle, it might indicate a) that leaders spend more money on electrification or fuel for generators in election years and/or b) that they make sure that the electric grid is up and running in election years. The specification to run would look at year on year changes in the various measures of nighttime lights and have "election year" as a dummy. If we get a positive sign on the election year dummy, then we have evidence for an electoral cycle…

As noted, these are but a few of the interesting analyses that these data, once collected and merged, will make possible. Suggestions for others and/or for ways of improving the simple specifications we have proposed would be most welcome.

Figure 1. Nighttime Lights in Africa, 2003

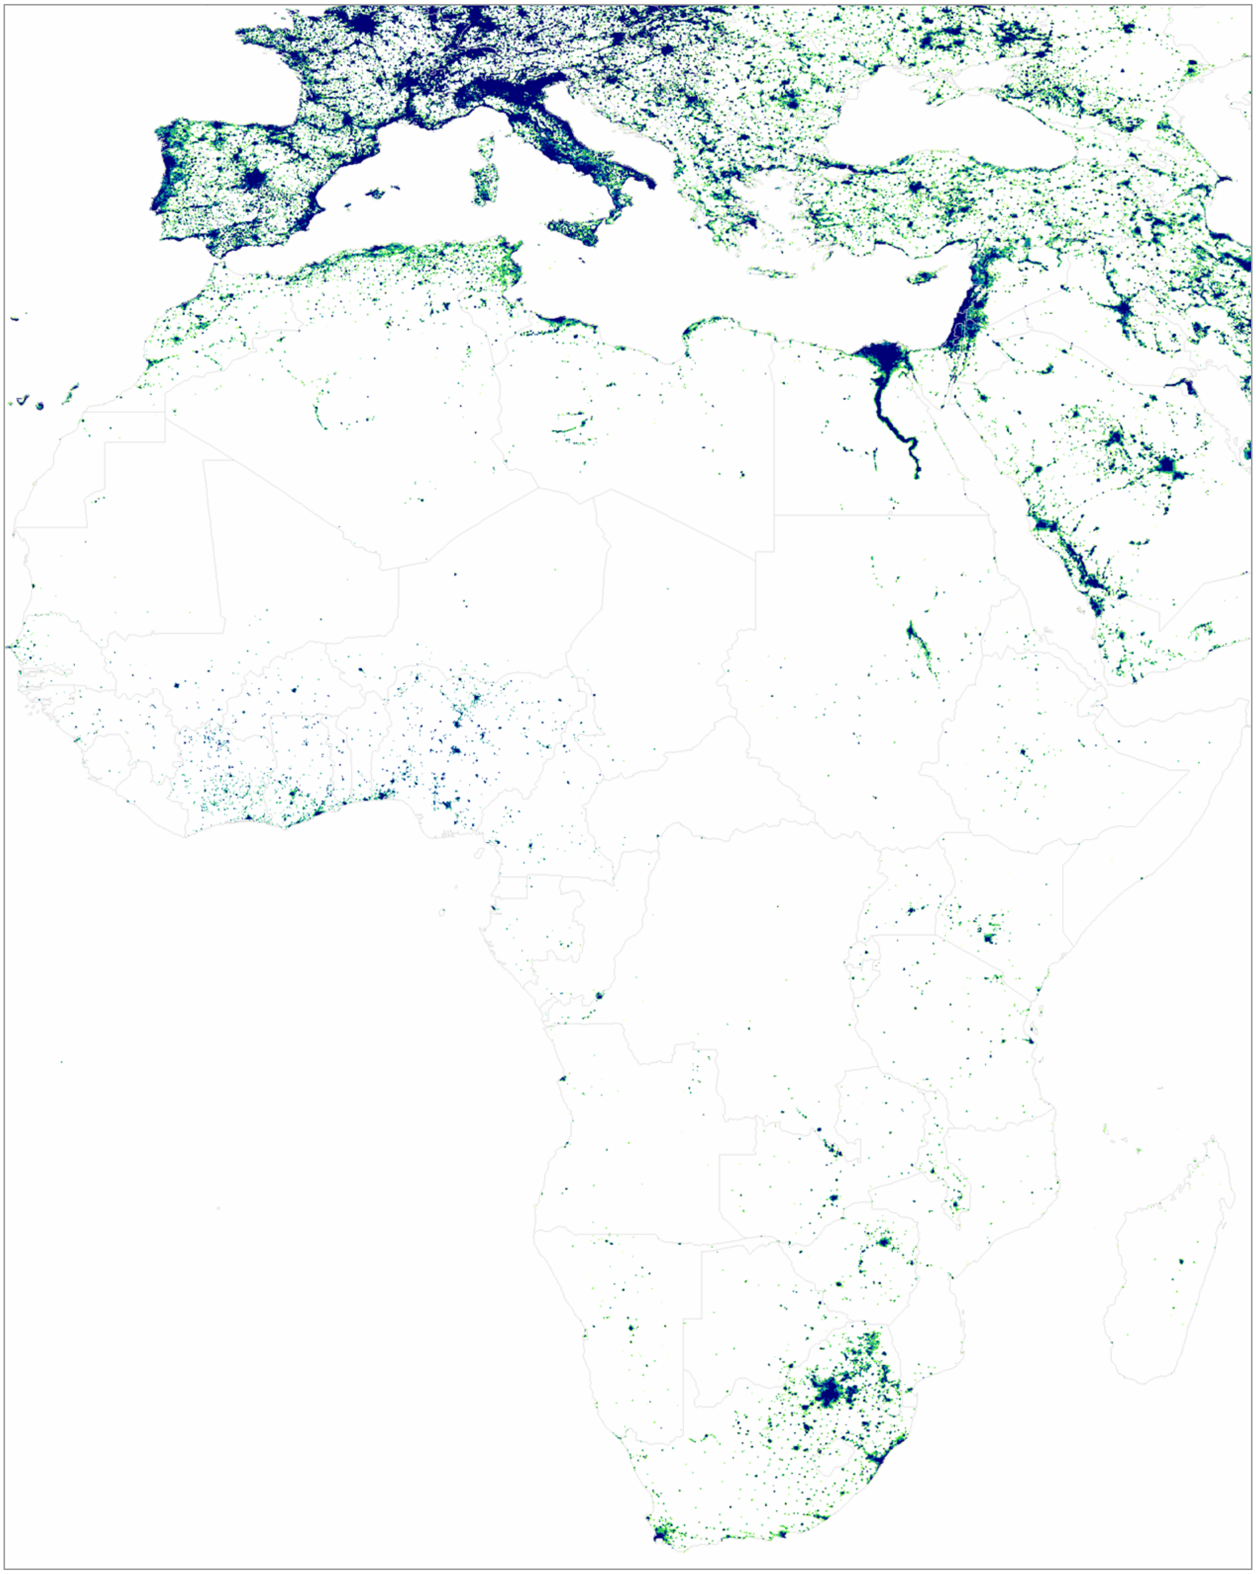

Source: DMSP-OLS F152003

Figure 2. Change in Lit Areas, 1992-2003

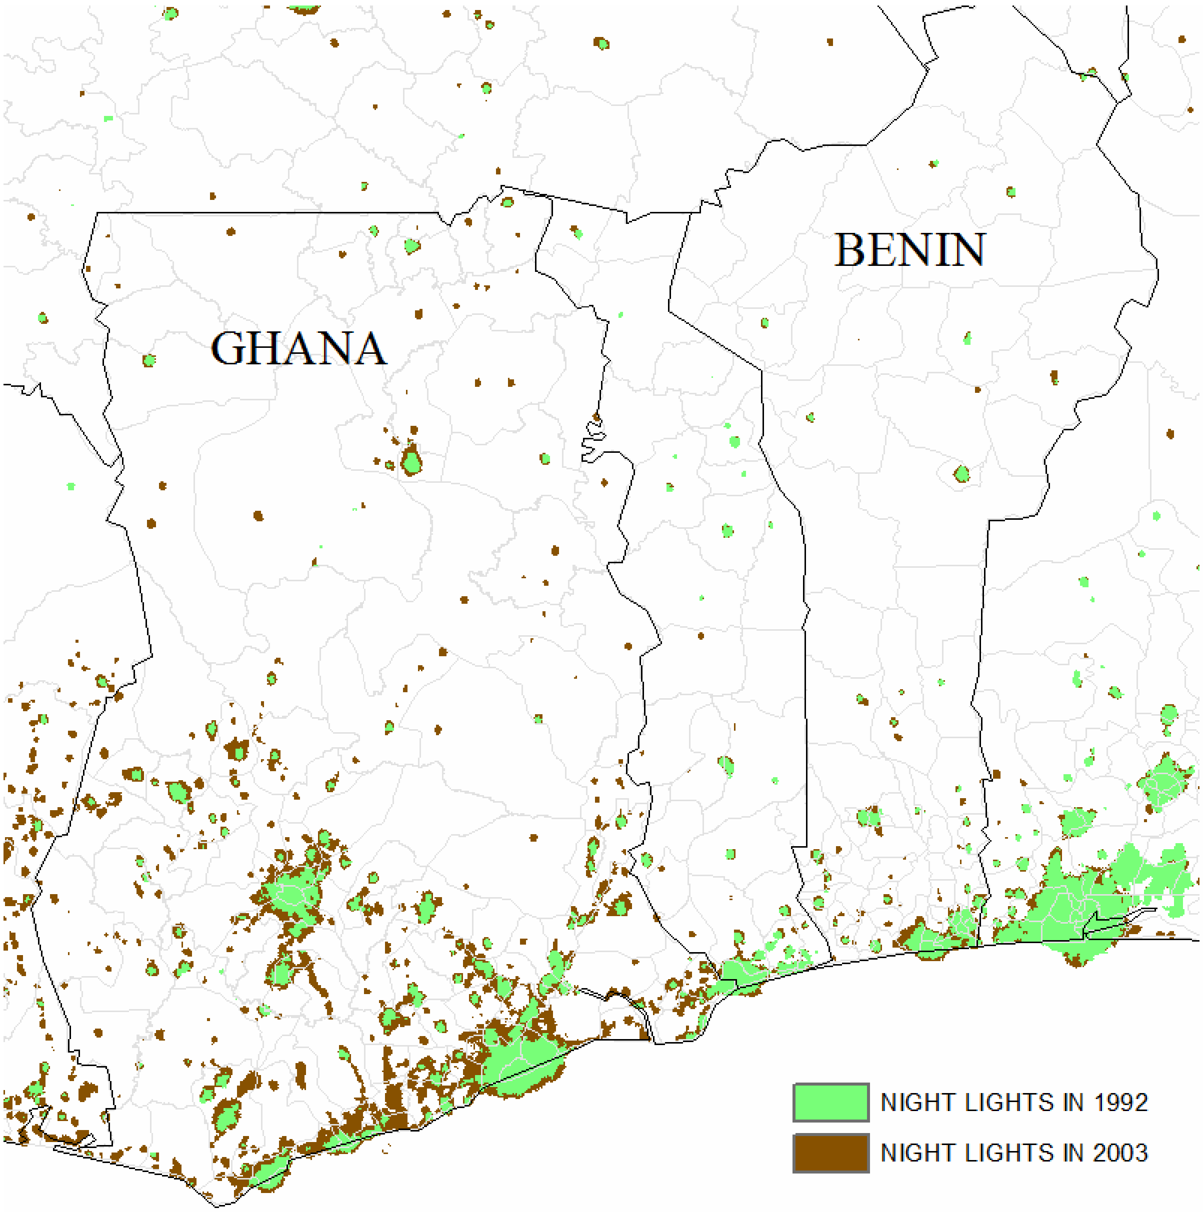

Note: Grey boundaries reflect second level administrative units.

References

Elvidge, Christopher et al. 1997. “Relation between satellite observed visible-near infrared emissions, population, economic activity, and power consumption.” *International Journal of Remote Sensing* 18(6):1373–1379.

Elvidge, Christopher et al. 2001. “Night-time lights of the world: 1994–1995.” *ISPRS Journal of Photogrammetry & Remote Sensing* 56:81–99.

Imhoff, Mark L. et al. 1997. “A technique for using composite DMSP/OLS “city lights” satellite data to map urban area.” *Remote Sensing of Environment* 61(3):361–370.

International Energy Agency. 2002. *World Energy Outlook 2002*. Paris, France: Organisation for Economic Cooperation and Development (OECD)/IEA.

Min, Brian. 2008. “Democracy and Light: Electoral Accountability and the Provision of Public Goods.” Unpublished manuscript. Available at http://bmin.bol.ucla.edu/demlight.html.