Left in the Dark? Oil and Rural Poverty

Brock Smith, Samuel Wills

Abstract: Do oil booms reduce rural poverty and inequality? To study this we measure rural poverty by counting people who live in darkness at night: combining high-resolution global satellite data on night-time lights and population from 2000 to 2013. We develop a measure that accurately identifies 74% of households as above or below the extreme poverty line when compared to over 600,000 household surveys. We find that both high oil prices and new discoveries increase illumination and GDP nationally. However, they also promote regional inequality because the increases are limited to towns and cities with no evidence that they benefit the rural poor.

JEL Codes: D31, E01, O11, O13, O47, Q32, Q33, Q43

Keywords: Night-time lights, Oil, Poverty measurement, Regional inequality, Rural poverty, Urbanization

THE VIEW THAT OIL AND OTHER NATURAL RESOURCES "curse" the countries that own them is widely held in academic and policy circles. This literature has focused mainly on how resources affect aggregate economic activity (Sachs and Warner 1995, 2001; Brunnschweiler and Bulte 2008; Alexeev and Conrad 2009; Smith 2015; James 2015; survey by van der Ploeg 2011). However, to properly understand their

Brock Smith is in the Department of Economics, Montana State University, USA (brock.smith1@montana.edu). Samuel Wills is in the Department of Economics and Oxford Center for the Analysis of Resource Rich Economies, University of Oxford, UK; and School of Economics, University of Sydney and CAMA, ANU, Australia. Samuel Wills would like to thank the Economic and Social Research Council for financial support, grant number ES/K009303/1. We are grateful for research assistance in producing poverty maps from Thomas McGregor, information on the DHS comparative wealth index from Shea Rutstein, and for helpful comments from Tiho Ancev, Vernon Henderson, Rick van der Ploeg, Simon Quinn, Gerhard Toews, Tony Venables, and seminar participants at SURED 2016, University of California, Merced, University of Colorado, Denver, Hebrew University of Jerusalem, Montana State University, and Oxford University.

Received May 25, 2017; Accepted November 29, 2017; Published online August 15, 2018.

JAERE, volume 5, number 4. 2333-5955/2018/0504-0006$10.00 http://dx.doi.org/10.1086/698512

effect on welfare we must know how they affect the income distribution. There has been very little empirical research on how natural resources affect inequality and poverty. We aim to fill that gap.

The challenge with studying inequality and poverty in developing countries is mostly one of data. According to the World Bank's definition, 76% of the world's poor live in rural areas (World Bank 2013). Data on income and wealth in rural areas are collected infrequently, if at all. When the data are collected they are expensive and time consuming, relying on household surveys that are rarely comparable across countries. The global standard for poverty data comes from the World Bank, which has done a remarkable job collecting and aggregating a huge number of detailed surveys and national accounts (Chen and Ravallion 2010). However, the coverage of these data remains below one-third of countries in any given year, making cross-country causal analysis and urban/rural comparisons difficult. Furthermore, in the words of Ross (2007, 238), "Surprisingly little is known about the relationship between mineral wealth and vertical income inequality . . . Data on income inequality are missing for most of the world's oil-dependent countries. In fact . . . there is a strong negative relationship between a country's dependence on mineral rents and the amount of data we have about its inequality levels."

To overcome these data issues we measure rural poverty by counting people who live in darkness at night. We do this using two detailed and geographically disaggregated data sets on night-time lights and population (see app. A; apps. A–G are available online). The first records the amount of light emitted at night around the globe at approximately a 1 km^2 resolution from 1992 to 2013. This has been used by many recent studies as a useful geographic proxy for economic activity (Henderson et al. 2011, 2012), covering institutions (Michalopoulos and Papaioannou 2013), political favoritism (Hodler and Raschky 2014), and infrastructure investment (Jedwab and Moradi 2015; Jedwab et al. 2015), among others. However, looking for poverty under lights is much like looking for lost keys under a street lamp. Instead, we focus on darkness. This is possible using a second data set from LandScan that estimates global population at a 1 km^2 resolution. After aggregating the data to approximately 100 km^2 cells, our final sample covers 1.04 million observations each year for 2000–2013. It reveals that 30% of people outside the Organization for Economic Cooperation and Development (OECD) live in darkness at night, rising to 50% in Africa (see, e.g., app. A).

People living in darkness is admittedly not a perfect measure of poverty. It only captures rural, rather than urban, poverty, so this paper is only about the former. It only considers one component of a household's consumption bundle, light, and ignores the many other unmet needs that characterize poverty. It will also pick up different levels of poverty around the world, as the provision of light depends on how effectively central governments provide electricity grids (see sec. 2).

However, using darkness to measure rural poverty also has a number of strengths. First, it is simple and rooted in basic needs. Lighting has very high returns, and people tend to switch from kerosene to electric lighting (so appear in our data) soon after leaving extreme poverty (see sec. 2.1). Second, it uses existing satellites and so is cheaper and faster to collect than household surveys. Third, it covers the whole world at fine resolution and regular intervals, producing a globally balanced panel data set at ~1 km^2 resolution that is well suited to empirical analysis. Fourth, it is strongly positively correlated with national-level estimates of extreme poverty. Fifth, when compared to a sample of 636,448 individual Demographic and Health Surveys (DHS) household surveys from 36 countries it accurately identifies 74% of households as being above or below the poverty line using ~100 km^2 cells, rising to 83% using 1 km^2 cells. The type I error (nonpoor living in unlit cells) is 5%, and the type II error (urban poor living in lit cells, outside the scope of this study) is 21%.

Darkness lets us study how oil and gas booms affect poverty and regional inequality in two experiments. The first exploits the sharp rise in oil prices after 2003, which is treated as exogenous to oil producers because it "was driven primarily by the cumulative effects of positive global demand shocks" (Kilian 2009, 1053). The second exploits giant oil and gas discoveries, to distinguish between price and quantity booms. Both show that oil and gas booms increased aggregate economic activity, but only in lit areas, and they did not reduce the share of people living in darkness at night.

We find that both oil price and quantity booms stimulated aggregate economic activity. High oil prices from 2003 to 2013 saw night-time illumination grow significantly faster in oil-dependent countries than in nondependent ones. After 10 years the price boom was responsible for approximately a 40% increase in total illumination, with similar effects for GDP per capita. Illumination also grew significantly in countries that made giant oil discoveries, with a 6-year lag (consistent with Arezki et al. 2016). We find that a giant oil or gas discovery with a net present value worth 100% of GDP increased illumination by 19%, and purchasing-power-parity-(PPP)-adjusted GDP by 8%, after 10 years.

The economic growth from oil and gas booms accrued to cities and towns but was not shared with the rural poor.[1] By the end of the 2000s oil price boom, economic activity in cities in oil-dependent countries was 15% higher than in their nondependent counterparts, and in towns was 38% higher. In contrast, the share of population living in darkness at night did not significantly change. There is no evidence that rural areas became lit, or that people moved to cities and towns, any faster in oil-dependent than nondependent countries. Ten years after a giant oil discovery the economic activity of cities and towns grew by 15% and 22%, respectively, relative to the control group, starting with a 4–6-year lag. Once again, there is no evidence that the population living in rural darkness changed significantly. Rural areas did not become lit,

1. Cities are defined as urban cells (as defined in sec. 2.2), towns as lit nonurban cells, and rural areas as populated, unlit nonurban cells.

but there may have been a small (1%) reallocation of people from unlit rural areas to towns. These results suggest that regional inequality grew during resource booms.[2]

These results are important as they provide the first global panel evidence that natural resources rents increase regional inequality within countries and do not benefit the rural poor. We find that oil booms did increase aggregate economic activity relative to the counterfactual, so rents were accrued domestically. However, these rents only accrued to areas of existing economic activity. They were not shared with the rural poor in a way that expanded their access to night-time lights. Furthermore, we find no evidence that people in unlit rural areas moved to urban areas with higher activity. This implies that public policy in resource-rich countries has failed to share resource rents through redistribution, rural development programs, and investment in public infrastructure.

To further understand the mechanisms behind unlit rural areas becoming lit we estimate a hazard model for individual cells. We find that unlit cells are more likely to illuminate if they have high population density, are adjacent to a lit cell (and thus the existing grid), or are near the capital city (consistent with Pinkovskiy 2013; Michalopoulos and Papaioannou 2014). This is true for both oil-dependent and nondependent countries. Cells in countries that have seen high light growth in the past are more likely to switch on in the present, but significantly less so in oil-dependent countries. This suggests that oil-dependent countries are less effective at converting economic growth to poverty reduction than nondependent ones.

Why don't oil booms reduce the share of people living in darkness at night? It ultimately comes down to the behavior of governments, and thus institutional quality, for two reasons. First, point source resources like oil and gas initially accrue to the government (van der Ploeg and Venables 2012). Second, the main means of expanding the reach of electric lighting is through government investment in the electricity grid (see sec. 2.1). Therefore, the quality of the institutions that dictate government behavior, be they democracy, courts, the rule of law, or other checks and balances, will be crucial.[3] Our findings are consistent with earlier work which finds that investment is inadequate and inequitable in resource-dependent countries for economic and institutional reasons (for more detail, see the discussion in secs. 3.4.2 and 3.4.6).

Our work contributes to two main strands of literature. The first is measuring income, poverty and inequality, which has a long and distinguished history (e.g., Kuznets

2. Our results invite a different interpretation to Aragón and Rud (2013), who find that the benefits from a Peruvian gold mine are shared evenly across the income distribution, though they focus only on areas near mines. Note that we do not directly calculate a spatial Gini coefficient because Elvidge et al. (2012) tried and found no correlation with existing Gini estimates due to urban poverty.

3. The main private channel for the rural poor to benefit from oil and gas windfalls—through employment in the industry—will be limited in this study. This is because oil and gas wells employ a small number of high skilled workers, typically from elsewhere.

1937, 1953, 1955; Kuznets et al. 1941; Stone and Croft-Murray 1959; Stone and Organització Europea de Cooperació Econòmica 1961; Atkinson 1970; Deaton 1985, 1997). Much of this combines data from national accounts and household surveys to approximate income distributions (Sala-i-Martin 2006; Pinkovskiy and Sala-i-Martin 2009; Chen and Ravallion 2010). There is a conflict between household and aggregate measures of income (Ravallion 2003), partly because aggregate measures exclude nonmarket services (Deaton 2005). Pinkovskiy and Sala-i-Martin (2016) use night-time lights data to reconcile this conflict, finding in favor of aggregate measures. Surveys have the advantage of offering detailed, targeted measures of poverty that take into account consumption bundles and price levels, quality (Deaton 1988), calorific demands (Subramanian and Deaton 1996), life expectancy (Pfeffermann and Webb 1983), and within-household distributions (Deaton and Muellbauer 1986), among other factors. The top income share is a proxy for the income distribution (see survey by Atkinson et al. 2009) and is correlated with poverty (Leigh 2009). Almås et al. (2014) assume that people living in darkness is a proxy for poverty in China. Other recent efforts to measure poverty with large data sets include using night-time lights to spatially distribute World Bank poverty rates (Elvidge et al. 2009; which we adapt in app. D), using mobile phone records in Rwanda (Blumenstock et al. 2015), and using machine learning on daytime satellite images in five African countries (trained with night-time lights; Jean et al. 2016). To our knowledge this is the first paper to test darkness as a poverty measure against existing methods like household surveys, and the first to use it in a causal study.

To our knowledge this paper is also the first global panel study of how oil booms affect poverty and regional inequality. Existing literature has focused on how natural resources affect aggregate income and growth. Little has been done on poverty, inequality, or the distribution of income: regional or otherwise. Gylfason and Zoega (2003) and Goderis and Malone (2011) provide evidence that resource booms increase survey-based Gini indices but note issues in comparing data across countries. Parcero and Papyrakis (2015) attribute this finding to underreported data. Bhattacharyya and Williamson (2013) show that commodity price shocks increase Australian top income shares. Van der Ploeg and Poelhekke (2016, 214) survey the empirical resource curse literature and find that "much more work on the impact of natural resources on . . . inequality . . . is needed."

The paper proceeds as follows. Section 1 introduces the data. Section 2 introduces and tests a new measure of extreme poverty: the number of people living in darkness at night. Section 3 uses the measure to investigate how oil price and quantity booms affect rural poverty and regional inequality. Section 4 concludes.

1. DATA

1.1. Night-Time Lights

Satellites from the Defence Meteorological Satellite Program's Operational Linescan System (DMSP-OLS) have recorded average annual night-time light intensity around the world since 1992. The data are provided at a pixel resolution of 30 × 30 arcseconds, or 0.86 square kilometers near the equator, and range from 0 to 63. The data are constructed by overlaying all daily images over the course of a year and discarding those that are obfuscated by cloud cover, lightning, aurora, and so on for a given pixel. Our analysis will focus on "aggregate lights": the sum of all light intensity values in a given country; and unlit areas: cells where total light intensity is zero.

There is a strong link between growth in country-level GDP and mean light intensity (Henderson et al. 2011, 2012). Doll et al. (2006) and Michalopoulos and Papaioannou (2014) have also performed cross-validation work for GDP levels. While we refer to those papers for a more detailed analysis, in figure 1 we plot the log of the sum of light readings by country against the log of PPP-adjusted real GDP (expenditure based) in 2003. The corresponding regression yields an adjusted r-squared of 0.82. Given their high resolution, lights data have been used in several studies for subnational analysis of GDP levels and growth rates. While this study is primarily at the country level, we leverage the fine spatial nature of the data to construct our poverty measure (see sec. 2.2).

Lights data are subject to some confounding issues. First, "top-coding" refers to pixels with a max value of 63, beyond which we cannot distinguish levels of economic activity. This occurs in especially dense or economically active areas in developed countries (Michalopoulos and Papaioannou 2014), so estimating urban poverty rates is difficult. This is not a major concern for our study because we focus on rural poverty in the least lit areas of developing countries. Second, gas flares appear in the lights data but do not reflect comparable economic activity. This is important for our paper as gas flares happen at oil wells when gas is not captured for sale. To control for this we drop

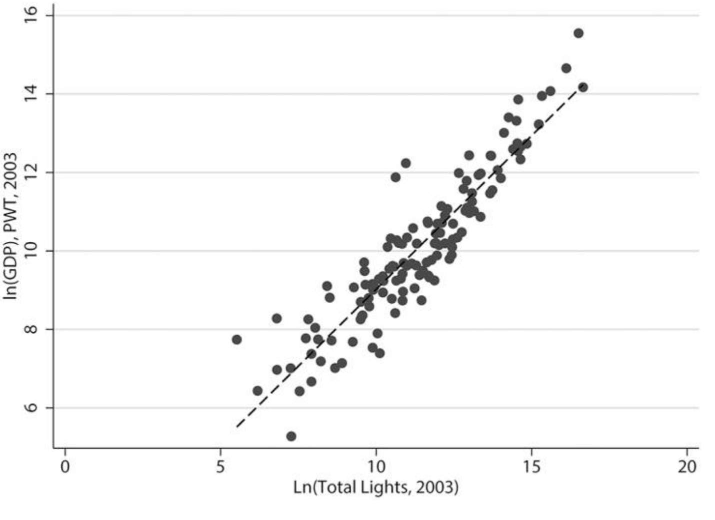

Figure 1. PPP-adjusted real GDP versus night-time lights (in logs)

all cells with gas flare activity according to the lights data provider (the Earth Observation Group). This gas flare shape file was made in 2009, so new flare activity since then is missed, but the aggregate lights results are nevertheless consistent with results using GDP data.[4] Third, overglow (or "blooming") occurs where light is recorded in pixels away from its origin and is magnified over terrain like water and snow (Doll 2008). Small et al. (2005) and McGregor and Wills (2017) find that overglow is linearly proportional to lit area, which is consistent with a physical model for atmospheric scattering. Pinkovskiy (2013) uses unlit wastelands to show that overglow on land is statistically insignificant more than 10 km from the light's origin. As we compute national-level estimates, use ~100 km^2 cells and focus on people living in areas far from cities; it is unlikely that this source of error will be sufficiently large or correlated with our outcome variables that it will confound our analysis (as in Pinkovskiy and Sala-i-Martin 2016). Fourth, the satellites used to construct the data change in 2000, 2004, and 2010, and the effectiveness of the sensors diminishes over time, which we control for using time fixed effects. Fifth, some cells at the margin flicker between unlit and lit, introducing measurement error in the dependent variable, which reduces precision. To address this in section 3.7 we also analyze "never-been-lit" cells (light intensity is zero for all previous years) and "barely lit" cells (light intensity is zero or below the bottom 10th percentile of lights per capita among lit cells globally in a given year), which are respectively more and less conservative measures of darkness than our main "unlit" cells measure (light intensity is zero).

1.2. Population Data

The Oak Ridge National Laboratory produces the LandScan data set, which estimates spatial population at a 30 × 30 arcsecond resolution annually from 2000 to 2013.[5] It estimates "ambient" average population over 24 hours, rather than just where people sleep. It is based on second-order administrative unit data, which are compiled by the International Programs Center of the US Bureau of Census. Population is then distributed throughout the grid according to a likelihood model that uses inputs including elevation, land cover, roads, coastlines, settlements, and high resolution satellite imagery, among others. It is similar to the Gridded Population of the World data

4. Excluding gas flare cells means we cannot capture local investment around oil and gas wells within gas flare zones. However, evidence from oil-rich Brazilian municipalities shows that oil windfalls do not improve local public good provision or living standards (Caselli and Michaels 2013). There is more evidence for positive local effects around mines, which have a greater role for cheap unskilled labor in extraction (Aragón and Rud 2013; Lippert et al. 2014; Cust and Poelhekke 2015).

5. This is the same resolution as the lights data, although the pixels are not perfectly aligned before 2010. The grid cells described in section 2.2 are aligned with the lights rasters but not the pre-2010 population rasters. The Zonal Statistics tool used in ArcGIS to find grid cell light and population counts addresses this by internally resampling the raster files so that they are aligned.

from NASA's Socioeconomic Data and Applications Center (SEDAC; see Dell 2010; Alesina et al. 2015), which is only available at 5-year intervals, and uses areal weighting to estimate population within census units as a function of land area only. We use the LandScan data because multiple studies have found likelihood models to provide more accurate population estimates than areal weighting (for an overview, see Lloyd et al. 2017), it has been found to be an effective small area estimator of rural population shares (using Burundi as a test case; Albert 2012), and it does not use night-lights data as an input.[6] To further reduce possible measurement error across space and time we aggregate population into three large subnational totals (cities, towns, and unlit rural areas) and focus on changes over 13 years to include multiple censuses. For robustness we use also Penn World Tables data (v8.1) on national population, and World Bank/UN data on urbanization rates.[7]

1.3. Household Survey Data

The Demographic and Health Surveys (DHS) Program collects individual- and household-level data on several questions about living conditions, possessions, health, and so on in over 90 countries. We use the most recent standard DHS survey from 36 countries, covering 636,448 households from 2003 to 2013. Each DHS survey is designed to be a nationally and regionally representative sample, with any deviations corrected for using sampling weights. Surveys are collected in clusters of 25–30 rural households, and 20–25 urban ones (ICF 2012). For confidentiality the location of each cluster is randomly displaced by up to 2 km in urban areas, and 5 km in rural areas with 1% of rural clusters displaced by 10 km (though all remain within their country and region, www.dhsprogram.com/faq.cfm). Following DHS recommendations, we manage this by basing our analysis on ~100 km^2 cells.

1.4. Urban and Rural Classifications

SEDAC also produces data on an "urban extents grid" from its Global Rural-Urban Mapping Project (GRUMP), which uses 1995 population estimates to classify each square of a 30 × 30 arcsecond global grid as either urban or nonurban. The classification is based on contiguous lit squares (as of 1995) and settlements known to hold at least 5,000 people, and agrees with urban extents based on DHS surveys (Dorélien

6. While early LandScan releases used night-lights data as an input, as does the Global Rural-Urban Mapping Project (GRUMP) data set, more recent releases, including ours, do not. Lights have been replaced with machine learning and computer vision techniques applied to daytime satellite images that are available at much higher resolutions (Vijayaraj et al. 2008; Rose and Bright 2014). For further detail, see http://web.ornl.gov/sci/landscan/landscan_documentation.shtml.

7. This is calculated using World Bank population estimates and urban ratios from the United Nations World Urbanization Prospects. For further detail, see http://data.worldbank.org/indicator/SP.URB.TOTL.IN.ZS.

et al. 2013). As it is based on 1995 data, this is a conservative definition of urban areas for our purpose.

1.5. Oil Prices and Discoveries

We study two types of oil and gas shocks: to prices and to the quantity of ultimately recoverable reserves. We use annual Brent crude prices from 1990 to 2013, which rose from $20 per barrel in 2002 to over $110 per barrel in 2013. Countries are classified as oil dependent or not following Baunsgaard et al. (2012), which is based on resource exports and revenues as a percentage of GDP for the years 2006–10. We include countries that are classified as resource dependent and where oil and/or gas is listed as the main commodity in appendix 1 of Baunsgaard et al. (2012). See appendix A for the list of dependent countries. Summary statistics for treatment and control groups as of 2000 (shortly before the price boom and the beginning of the sample for most outcomes studied), along with the *p*-values from a *t*-test of mean differences, are given in appendix table E.1 (tables B.1, C.1, C.2, E.1–E.10 are available online). They show that the treatment group was somewhat brighter and more urbanized than the control group. This is one reason why it is important to control for convergence effects, as we do in our main specifications in section 3. Further, despite any differences in levels our results generally support the assumption of parallel preexisting trends, as required by the difference-in-differences model.

The quantity shock uses data on giant oil and gas discoveries from the American Association of Petroleum Geologists. This is an updated version of the data set by Horn (2003, 2004), which builds on Halbouty et al. (1970). The data record the field name, location, date of discovery, type, and estimates of ultimate recoverable reserves—which must exceed 500 million barrels of oil equivalent (MMOBE) to be considered a "giant" discovery. In total the data cover 1,019 discoveries, 245 of which occur between 1990 and 2013. We take discoveries back to 1990 to study their effect on lights up to 10 years after the discovery (as population data begin in 2000). Dropping the OECD and aggregating multiple discoveries for a country in a given year leaves 98 unique discovery years in our sample: 51 for oil and 52 for gas (some country-years had both; see appendix fig. B.1).

We account for the size of each discovery by constructing a measure of its net present value divided by GDP (following Arezki et al. 2016),

$$Size_{i,t} = \frac{\sum_{j=5}^{J} q_{i,t+j}\, oilprice_t (1 + r_i)^j}{GDP_{i,t}} \times 100, \tag{1}$$

where $Size_{i,t}$ is the discovery size for country i at the time of discovery, t. This is the discounted sum of total revenue based on an approximate production profile, $q_{i,t+j}$,[8]

8. Production is also affected by institutions and conflict (e.g., van der Ploeg and Rohner 2012).

from the fifth year after discovery to the exhaustion date, J, valued at the oil price at the time of discovery, $oilprice_t$ (assuming prices follow a random walk). We therefore exploit variation in the size of the discovery, as well as its timing via the oil price. Revenue is discounted using country-specific, risk-adjusted discount rates to account for differences in political risk. This assumes a riskless rate of 5%, and predicted premia based on the past relationship between bond spreads (41 countries in the Emerging Markets Bond Index) and political risk (133 countries in the International Country Risk Guide), to account for the limited data on bond spreads. For more detail, see Arezki et al. (2016).

2. USING DARKNESS TO MEASURE RURAL POVERTY

We measure extreme rural poverty by counting people who live in darkness at night. This allows us to focus on the bottom of the income distribution, in contrast to studies using light to proxy economic activity. As the measure is constructed from satellite data it is faster, cheaper, and has better coverage than existing survey measures. For robustness we also conduct the analysis in section 3 using a different (but related) measure that uses lights to allocate World Bank poverty rates to individual cells (based on Elvidge et al. 2009) and find similar results as described in appendix D.

2.1. Motivation

Lighting is linked to poverty because it is a basic need. It extends the working day and increases productivity, safety, and education. It is also a normal good and is consumed more when incomes rise (Lee 2013). When a household's income rises it changes its lighting source according to an "energy ladder" (Leach 1992; Lee 2013). This involves using inefficient and polluting sources (e.g., biomass) first, more efficient transitional fuels like kerosene next, and clean but expensive sources such as electricity last (Mills 2003; Bacon et al. 2010). Electricity can be used either by connection to a public grid or, particularly in rural areas, using private off-grid diesel, solar, or wind generators.

Areas appear in our night-time light data after their households switch from kerosene to electric lighting. A typical kerosene lamp delivers from 1 to 6 lux of useful light; a standard 60 W bulb delivers 100 lux at useful distances; and the western standard for tasks such as reading is 300 lux (Mills 2003). Therefore, there is a considerable jump in illumination when a household switches from kerosene to electric lighting. Elvidge et al. (2011) find that the number of people living in areas with detectable lighting in our data set is highly correlated with reported electrification rates. Similarly, we find from the DHS surveys that 71% of households in lit cells have electricity, while only 15% in unlit cells do.

Households adopt electric lighting soon after crossing the World Bank poverty line, though the exact level differs by country and urbanization. Lee (2013) studies energy use in Uganda using a cross-section of 6,775 urban and rural households in 2009–10. It finds that kerosene is only used for lighting and has an inverse U-shaped relationship with income, consistent with the energy ladder theory. Kerosene usage peaks at a per capita income of UGX 0.83 million per year, equivalent to \$2.45 per day in 2005 PPP (Penn World Tables 8.1). Beyond this point, electricity starts to be used for lighting, while kerosene is retained for reliability. Bacon et al. (2010) shows that Uganda provides an upper bound for electricity adoption, as it has relatively poor electrical infrastructure (see fig. C.1). Other countries with denser populations and better infrastructure adopt electricity earlier: the lowest quintile of rural households in India, Pakistan, and Vietnam have average incomes of \$0.80, \$1.32, and \$1.42 per person per day (2005 PPP), but electricity is the main source of energy in 33%, 69%, and 90% of households, respectively. Urban electrification is again higher, though not universal (Lee et al. 2014). These thresholds can be compared to the World Bank's poverty line of \$1.25 per person per day (2005 PPP). While electric lighting is just one element of a household's consumption bundle it is a simple, consistent, intuitive, and readily observable measure of when a household leaves the most abject state of poverty.

2.2. Measurement

We measure poverty using the "unlit rural percentage" (URP) by combining spatial data on lights and population. To evaluate the measure we compare it to two definitions of extreme poverty: using the Comparative Wealth Index (CWI) and Unmet Basic Needs (UBN), constructed from 636,448 DHS household surveys from 36 countries.

To calculate the URP we first create a global grid of cells comprising 12 × 12 pixels each. The vast majority are approximately 100 km^2 near the equator, though cells are divided at national borders. We drop cells with zero population. Ideally we would analyze each 1 × 1 pixel, but with 100 million observations per year this is computationally impractical. Aggregating pixels like this speeds up the analysis without losing accuracy for our purposes (see sec. 2.3.2) and has become standard in the literature (e.g., Michalopoulos and Papaioannou 2013, 2014; Jedwab and Moradi 2015).

Second, we sum the population living in cells with zero lights everywhere in the cell, and divide by the total country population (excluding cells dropped due to gas flares). We reason that people in cells with moderate population density and no lights are among the world's extreme poor. This provides a binary measure of poverty that is comparable to the World Bank Poverty Rate and Unmet Basic Needs. Some cells have a nonpoor population at a sufficiently low density to record no night-time lights. By definition these cells will contribute little to the overall unlit percentage. There will also be poor living in lit cells who are not captured, which we study using a different measure in appendix D.

Third, we classify lit cells as urban or rural, according to the GRUMP Urban Extents Grid definition in 2000 (see sec. 1.4). Each cell contains a mix of urban and nonurban pixels, so we classify it as urban if at least 33% of its pixels are urban. This definition yields an urbanization percentage with a high correlation with World Bank national urbanization rates of .78 in 2000, dropping only to .76 by 2013. Therefore, nearly all of the increase in urbanization is due to migration to existing urban areas and our static definition of urban areas is not a major issue. Cities are defined as urban cells, towns as nonurban cells with lights, and unlit rural areas as cells without lights but with nonzero population. The number of unlit urban cells is negligible.

The Comparative Wealth Index measures wealth at a household level. Each DHS survey includes a household wealth index which is normalized to have a mean of zero and standard deviation of one. It is calculated in two stages: first, a survey-wide index is calculated using items common to urban and rural areas; then it is adjusted for items unique to urban and rural areas (including controlling for lower rural electrification rates; see Rutstein 2008). To make the normalized index comparable across surveys we follow Rutstein and Staveteig (2014) and calculate a CWI for each household.[9] The threshold for extreme poverty is a CWI < -0.27, which is approximately the median for households with two or more unmet basic needs.

Unmet Basic Needs directly measures whether households lack the four basic needs used in the CWI: poor housing construction, overcrowding, poor sanitation, or high economic dependency. This assumes that poverty is an inability to meet one's basic needs and is based on the UBN index developed by the Economic Commission for Latin America and the Caribbean.[10] Extreme poverty is defined as having at least two (from four) unmet basic needs.

We focus on the CWI rather than the UBN as it is a more complete, and less noisy, measure of poverty. The UBN assumes that people will satisfy their basic needs before purchasing luxuries, which is not true in practice. Of the people in extreme poverty according to the UBN ≥ 2 definition, 33% own a television, 20% a fridge, 9% a phone, and 5% a car. The CWI acknowledges that these households may have been able to afford to meet their basic needs but chose not to. The CWI is also less noisy because it is continuous so has more variation than the five UBN categories.

9. This uses survey-level coefficients generously provided by Rutstein and Staveteig (2014). In brief, these coefficients are constructed by comparing household wealth indices at common "anchor points" of asset-based wealth that can be found across all surveys. At the bottom of the wealth distribution these anchors are whether households meet four basic needs: poor housing construction, overcrowding, poor sanitation, and high economic dependency. At the top they are whether the household owns a television, refrigerator, car/truck, and fixed telephone. The survey-level coefficients are then calculated by linear regression that relates wealth indices at a given survey's anchor points to those of a baseline survey (the Vietnam 2002 survey). See Rutstein and Staveteig (2014) for a complete explanation.

10. For details, see Rutstein and Staveteig (2014). In addition we allow "good sanitation" to include households sharing a toilet, as otherwise over 95% of households have poor sanitation, which is uninformative. When a household is missing information we determine whether this creates uncertainty about its UBN score, given the other criteria. If uncertainty exists, we drop the household from the UBN analysis, which amounts to less than 1% of the sample.

2.3. Evaluation

2.3.1. National Level

The unlit rural percentage is strongly positively correlated with existing poverty measures at a national level. Figure 2 plots the 2011 URP by country against poverty rates from the World Bank, the CWI and the UBN. The World Bank's rate is based on its most recent estimates of people living on less than \$2 per day.[11] The CWI rate is based on the −0.27 threshold, and the UBN rate is based on having two or more unmet basic needs. While there are some countries for which the unlit percentage is a poor predictor,[12] there is a strong correlation for all three regressions. The corresponding regressions yield an adjusted *r*-squared of 0.74 (World Bank), 0.56 (CWI), and 0.60 (UBN).

2.3.2. Household Level Globally

We find that darkness accurately identifies 74% of households in our sample as being above or below the CWI threshold for extreme poverty at 12×12 pixel resolution (33% are poor and live in unlit cells, 41% are not poor and live in lit cells; see fig. 3). This rises to 83% at 1×1 pixel resolution,[13] or 66% using the UBN ≥ 2 definition at 12×12 (see table C.1). We proceed using the 12×12 resolution because global analysis at the individual pixel level is computationally impractical, we suffer only minor loss of accuracy, and overglow is less of an issue.

The poverty share in unlit cells is 2.6 times higher than in lit cells at 12×12 pixel resolution (88% vs. 33%), or 4.4 times higher at 1×1 resolution. The type I error at 12×12 resolution (nonpoor living in unlit cells) is 5%, and the type II error (poor living in lit cells) is 21% (again, see fig. 3). The type I error is more relevant for our analysis because the type II error mostly captures the urban poor, which darkness does not identify and which we explicitly do not study.

Tightening the definition of poverty does not greatly alter the measure's accuracy (see fig. C.2). The CWI definition classifies approximately half the households in our sample as poor. Accuracy remains above 70% if the definition is tightened and only falls below 70% when the definition is loosened to count 60% or more of our sample as poor. As the definition is tightened the number of households above the poverty

---

11. These rates are not all from 2011—they are taken in various years. We use only World Bank estimates that have been made since 2005.

12. Figure 2*A* shows a cluster of countries in the top left of the graph—that is, countries with low unlit rural percentages and high World Bank poverty estimates. These six countries are Indonesia, the Philippines, Swaziland, Bangladesh, India, and Pakistan. With the exception of Swaziland, these are high-density poor countries with high rates of urban poverty, which the unlit rural percentage does not capture. The calibrated poverty measure described in app. D performs much better for these particular countries.

13. The 1×1 figures are based on 25 surveys from 2010–13 because of a grid alignment issue before then.

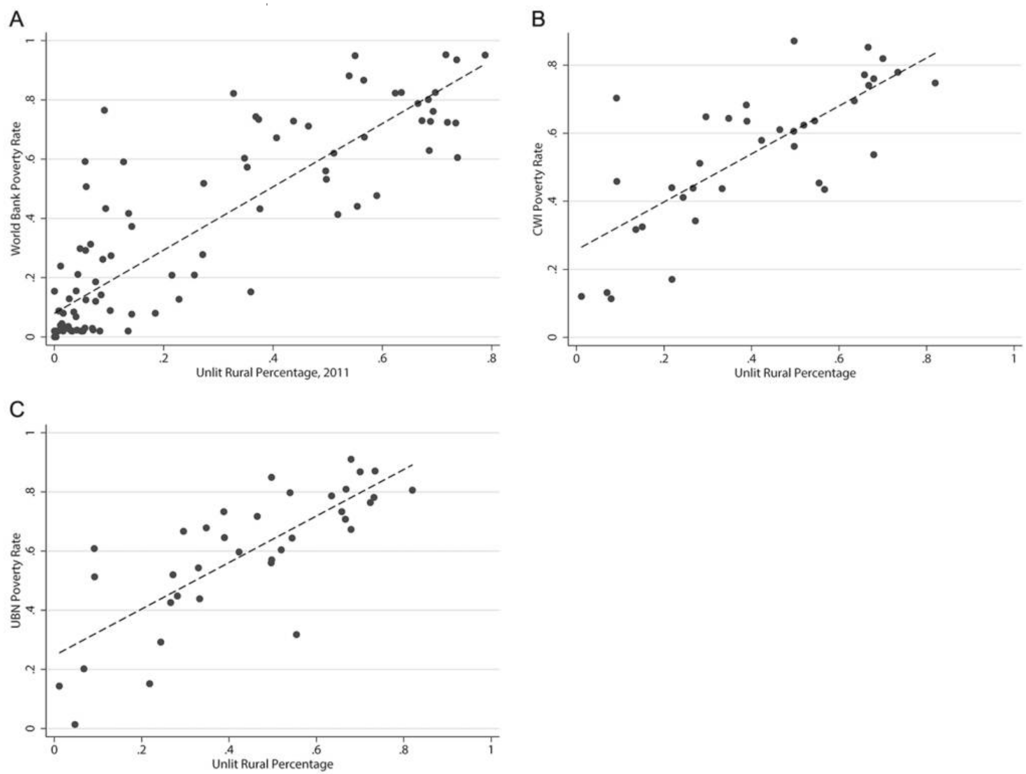

Figure 2. Unlit rural percentage versus three poverty measures: less than \$2 per day percentage (*A*) (World Bank 2015), CWI extreme poverty rate (*B*), UBN $\geq$ 2 poverty rate (*C*).

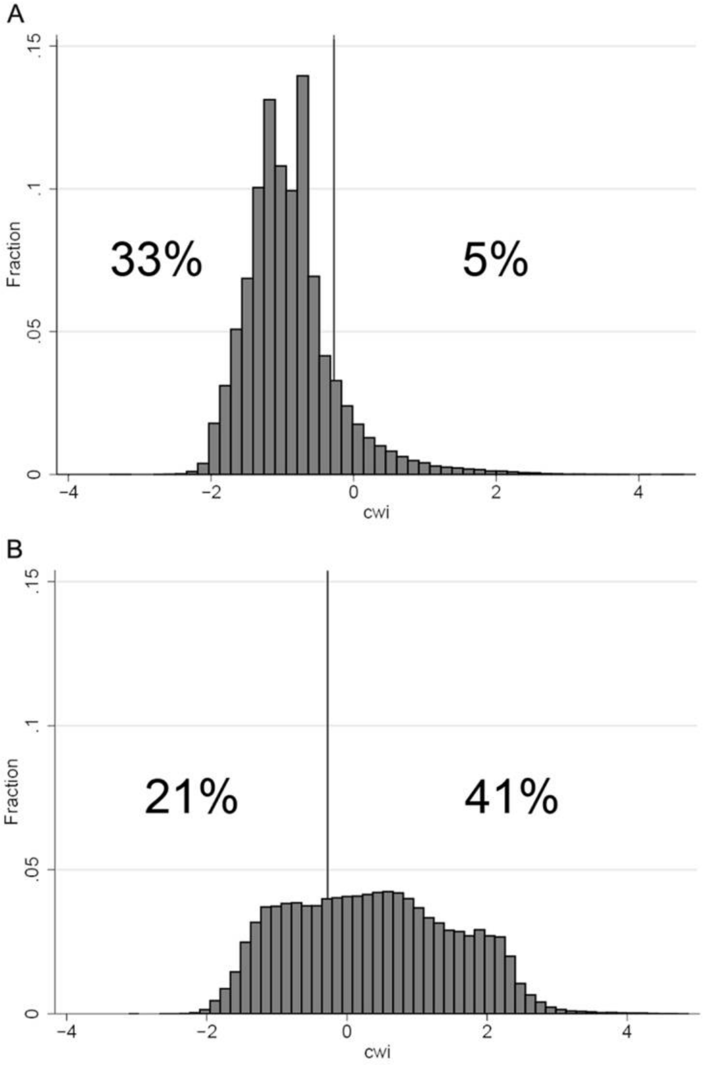

Figure 3. Distribution of wealth in unlit (*A*) and lit (*B*) cells at 12 × 12 pixel resolution (CWI extreme poverty threshold indicated by the vertical line).

line in lit areas rises, while the number below the line in unlit areas falls. Type I error increases and the type II error falls.

2.3.3. Household Level by Country

Darkness's effectiveness as a poverty measure is relatively consistent across countries, accurately identifying over 60% of households at 12 × 12 pixel resolution in all but one of our surveys (see table C.2). The other, Bangladesh, is unique in its high population density, electrification, and poverty rates. Increasing the resolution to 1×1 pixel improves the accuracy in Bangladesh to 62%, and in Ethiopia, Rwanda, and Congo DRC to over 90%. If we exclude Bangladesh, then the correlation of accuracy at 12×12 resolution with the number of lit cells, CWI, urbanization, population density, and electrification is insignificant at the 10% level, suggesting that the measure is not systematically biased.[14]

3. THE EFFECT OF OIL BOOMS ON POVERTY AND INEQUALITY

To understand how oil booms affect rural poverty we study exogenous increases in both prices, using the high oil prices of the 2000s, and quantities, using data on giant oil discoveries. We isolate the regions that benefit from each type of boom using darkness as a spatial measure of rural poverty. Overall we find that both price and quantity booms stimulate aggregate economic growth in the short to medium term. In both instances this growth promotes regional inequality as it accrues to cities and towns, but not the rural poor. High oil prices do not affect the population distribution, but oil discoveries cause a small reallocation from unlit rural areas to towns (though not cities).

3.1. Identification: Oil Prices and Discoveries

The first experiment exploits the period of high prices between 2003 and 2013, and its differential effect on oil-dependent and nondependent countries. We treat the rise in prices from 2003 as a shock to global demand that is exogenous to oil-dependent countries, following the evidence in Kilian (2009) that uses a structural decomposition in a VAR framework. We control for global demand shocks affecting our treatment countries through channels other than oil by using nondependent countries as the counterfactual. To control for any endogenous supply responses we also exclude OPEC countries and find that our results do not appreciably change. To control for other common unobserved characteristics of oil exporters we compare the periods before and during the oil price boom, although this still leaves the possibility of other unobserved contemporaneous effects impacting oil producers. Our second experiment is less vulnerable to this possibility.

The second experiment exploits the discovery of giant oil and gas fields, which we treat as exogenous after controlling for time and country fixed effects, as do Lei and Michaels (2014), Smith (2015), Wills (2015), and Arezki et al. (2016). Giant oil and gas discoveries are a type of quasi-natural experiment, occurring in only 2% of all wells drilled. Countries have limited ability to affect this. Toews and Vezina (2016) find that a country

14. The average accuracy for the five oil-rich countries in our DHS sample is roughly 5 percentage points higher than for non-oil-rich countries, a significant difference according to a *t*-test. It is unclear why accuracy could be systematically higher in oil-rich countries, and the difference may be explained by differences in the composition of treatment and control countries in the DHS sample versus the full sample.

must double drilling activity to increase the probability of a giant discovery by one percentage point. Anecdotally, a giant Norwegian discovery in 2010 was made only three meters from where drilling failed to find oil in 1971 (Kavanagh 2013). We are concerned specifically with the timing of discoveries, which we argue is exogenous after including fixed effects for time (controlling for the price of oil and drilling equipment) and country (controlling for past discoveries and political institutions). The treatment is the size of the discovery relative to GDP (eq. [1]), which incorporates variation in both the quantity of oil and the price when it is discovered. The control group is countries that do not make a giant discovery and observations for countries that do but before a discovery is made.

3.2. Estimating Equations

We estimate the effects of oil price shocks on spatial outcomes aggregated to the country level using a difference in difference specification:

$$Y_{i,t} = \sum_{s=2000}^{2013} \beta_{1s}(\lambda_s D_i) + \sum_{s=2000}^{2013} \beta_{2s}(\lambda_s Y_{i,2000}) + \Lambda_t + \Phi_i + region_i \times t + \varepsilon_{i,t}, \quad (2)$$

where Y_{it} is the outcome of interest for country i in year $t = 2000, \ldots, 2013$, λ_s is a year indicator equal to one if $s = t$ and zero otherwise, D_i is an indicator equal to one if classified as an oil or gas-dependent country and zero otherwise, $Y_{i,2000}$ is the outcome variable at the beginning of the sample to control for convergence effects (except in the population regressions), Λ_t is year fixed effects, Φ_i is country fixed effects, and $region_i * t$ is regional linear trends.[15] Each coefficient β_{1t} then measures the average conditional difference in Y between dependent and nondependent countries in year t relative to the difference in the reference year 2002, which is the year before oil prices began to rise. Standard errors are clustered at the country level.

To estimate the effect of giant oil and gas field discoveries we use a distributed lag specification:

$$Y_{i,t} = \alpha + \sum_{j=0}^{22} \beta_{1j} Size_{i,t-j} + \sum_{s=2000}^{2012} \beta_{2s}(\lambda_s Y_{i,2000}) + \Lambda_t + \Phi_i + region_i \times t + \varepsilon_{i,t}, \quad (3)$$

where $Size_{i,t-j}$ is the net present value (NPV) relative to GDP of a discovery made in year $t - j$, and other variables are defined as above. In this case β_{1j} is the effect of a discovery made j years ago equal to 100% of GDP. This specification lets us measure the dynamic effect of discoveries over time and analyze multiple discoveries within countries. Because our data on the size of discovery only covers discoveries made on or

15. Regional classifications are taken from World Bank classifications. The main sample includes the following regions: Central Asia, East Asia and Pacific, Eastern Europe, Latin America and the Caribbean, Middle East and North Africa, South Asia, and sub-Saharan Africa.

before 2012, these regressions cover the years 2000–2012, rather than 2013 as in the price shock specification above. We focus on the 10 years after discovery due to data limitations, so use all discoveries since 1990 (since the population data begin in 2000). However, we include lags up to 22 years after discovery to ensure that our counterfactual for lags 0–10 are limited to countries that do not discover oil, and those that do in the years before the discovery is made. The delay between discovery and production is typically 4–6 years, so we hypothesize positive effects on economic activity following this lag plus any anticipation effects (see Arezki et al. 2016) for a full discussion). Standard errors are again clustered at the country level.

Zooming in to the cell level, we model how various spatial mechanisms affect the probability of an unlit rural cell becoming lit in a given year using the following hazard model specification (using 12 × 12 pixel cells, as used to construct the URP),

$$I_{cit}^{RS} = \alpha + \beta_1 X_{cit} + \beta_2 D_i X_{cit} + \Lambda_t + \Phi_i + \varepsilon_{c,i,t}, \tag{4}$$

where I_{cit}^{RS} is an indicator for whether a particular cell c switches on in year t (and is dropped from the sample thereafter); X_{cit} is a vector of independent variables describing adjacency to lit cells (indicator, measured as of 2000), being <100 km from the capital (indicator), population density (standard deviation units), and aggregate national light growth since 2000 (continuous); D_i is an indicator for being in an oil-dependent country; and Λ_t and Φ_i are year and country fixed effects. Standard errors are clustered at the country level. The coefficients in β_2 therefore estimate the differential effect of each variable on dependent relative to nondependent countries. The sample is restricted to cells that are unlit but inhabited as of 2000.

Since we are interested in extreme rural poverty and are thus focused on the developing world, we drop all OECD countries as well as countries with an unlit rural percentage of less than 5%[16] in 2000 from all specifications. We also drop three countries that experienced large wars during the sample period, including two potential treatment countries: Iraq, Syria, and Afghanistan. This leaves 105 countries in the main sample.

3.3. Results: Broad Trends

Examining the raw data, we find that GDP is growing and poverty is falling around the world. We know that GDP is growing from a variety of sources.[17] Global poverty, or the unlit rural percentage, is generally falling, which is consistent with broader evidence (fig. 4*A*; see also Pinkovskiy and Sala-i-Martin 2009; Chen and Ravallion

16. Results are robust to a threshold of 10%.

17. Unconditional illumination data do not illustrate this well because of annual changes in light sensitivity and new satellites in 2000, 2004, and 2010, which we control for using time fixed effects.

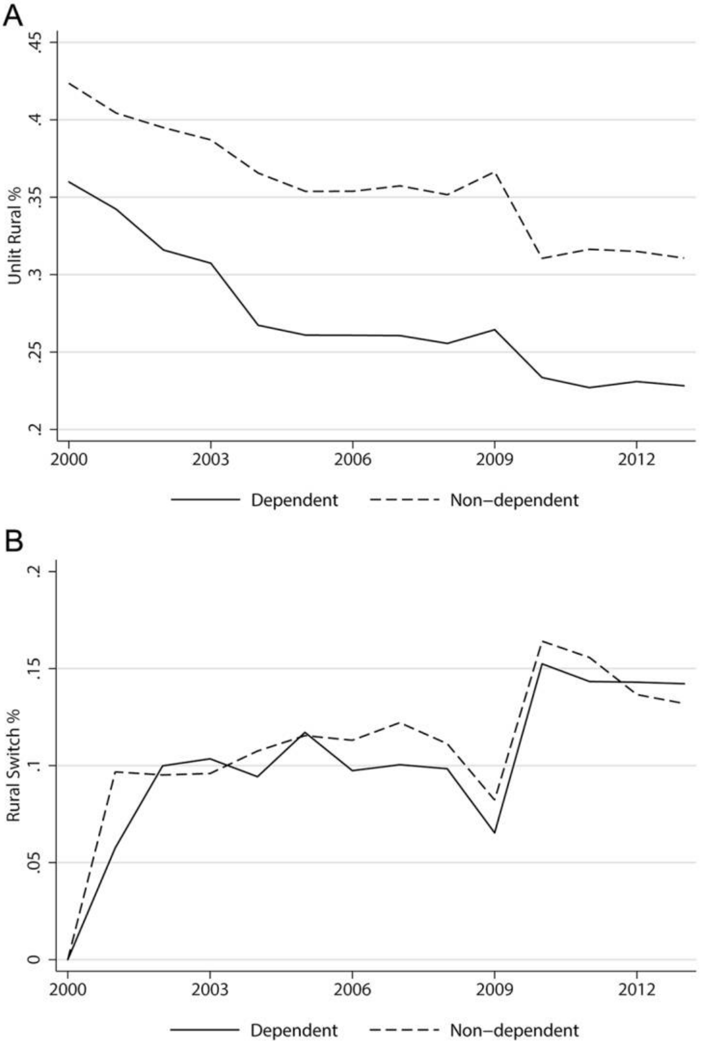

Figure 4. Unconditional trends in the unlit rural percentage (*A*) and switch-on percentage (*B*)

2010). The unlit rural percentage can fall for two reasons: unlit areas become illuminated, or people leave unlit areas for towns and cities. We find evidence for both. The first is measured using the "rural switch percentage," which records whether unlit, populated areas in 2000 subsequently became lit. Figure 4*B* shows that unlit rural areas steadily become illuminated during our sample, with jumps in part due to satellite sensitivity. The second is illustrated in figure 5, which shows the general trend for people to leave unlit rural areas for towns and cities.

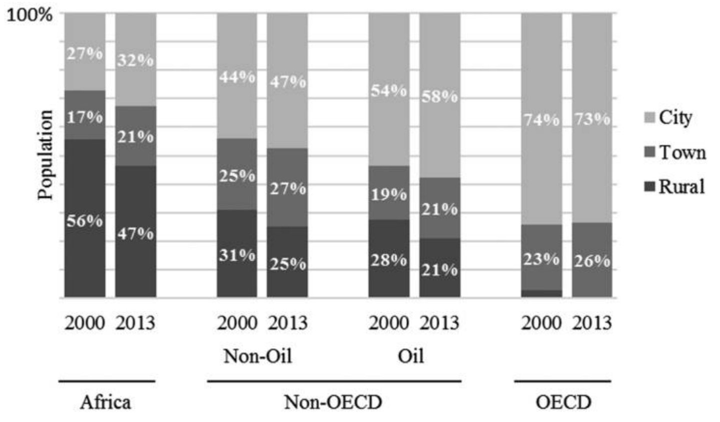

Figure 5. Breakdown of population share by area classification as of 2000

3.4. Results: The Oil Price Boom from 2003 to 2013

3.4.1. High Oil Prices Stimulated Growth in Oil-Dependent Countries

Aggregate illumination (the sum of all light intensity values in a given country) grew in oil- and gas-dependent countries relative to nondependent ones during the 2003–13 oil price boom. Figure 6[18] shows that in the 10 years before oil prices began to rise in 2003, the difference in log aggregate night-time lights between dependent and nondependent countries remained steady.[19] However, by 2013 log aggregate night-time lights had risen by 40% (0.34 log points) in dependent relative to nondependent countries, compared to 2002 (coefficient β_{1t} in eq. [2]). The point estimates show a clear trend throughout the 2003–13 boom. Log lights per capita and PPP-adjusted real GDP per capita (expenditure based) also grew steadily over the period, with statistically significant effects at a 5% level.

3.4.2. Growth from High Oil Prices Was Not Shared with the Rural Poor

The aggregate growth in illumination masked an increase in regional inequality as it was confined to cities and towns. By 2013, lights in cities had risen by 15% (0.14 log points) in oil- and gas-dependent countries relative to nondependent ones; and in towns they had risen by 38% (0.32 log points) (fig. 7*A*, 7*B*).[20] These estimates are significant at

18. These results and all results in the main text are shown in table form in app. E.2.

19. In the late 1990s light in oil- and gas-dependent countries began to rise relative to nondependent ones. This coincides with a mini-boom in oil prices, which rose from $10 per barrel in 1998 to over $30 per barrel in 2000 before briefly stabilizing.

20. The results in figures 7 and 8 start in 2000, when the LandScan data begin. The results in figure 6 start in 1992 to exploit all available night-time lights data and more fully observe pretreatment trends, using population data from the Penn World Tables (8.1).

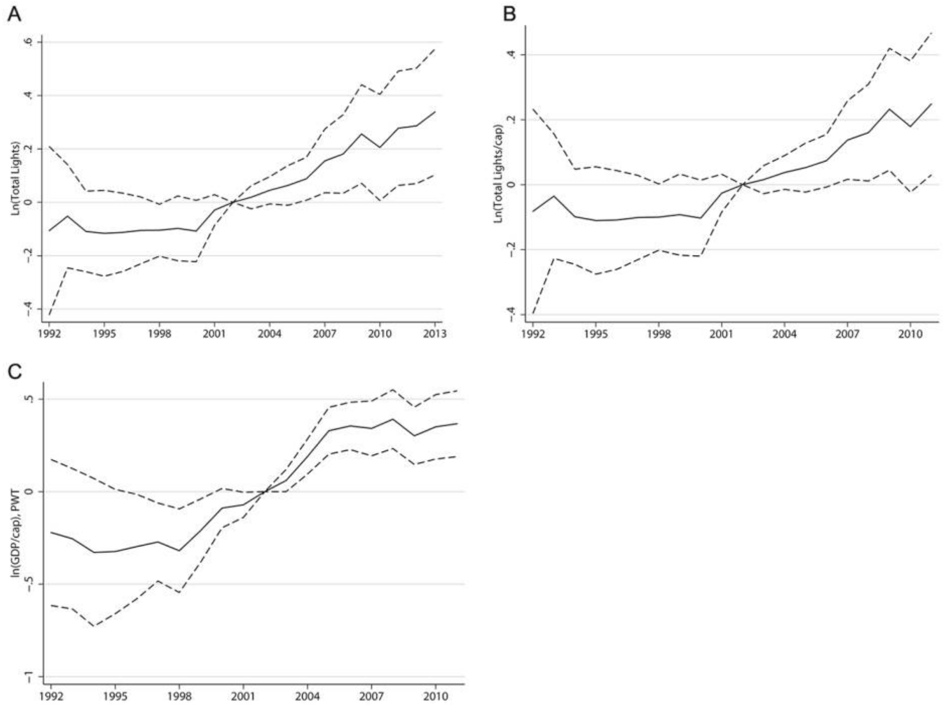

Figure 6. Effect of the 2003–13 oil price boom on (log) aggregate lights (*A*), (log) aggregate lights per capita (*B*), and PPP-adjusted (log) real GDP per capita (*C*) (95% confidence bands in dashes). Each point represents the conditional difference between oil- and gas-dependant and nondependent countries, compared to the conditional difference in 2002. The specification controls for convergence and includes year and country fixed effects and regional linear trends (see eq. [2]). Tables of these estimates, along with control variables, are found in appendix E.2.

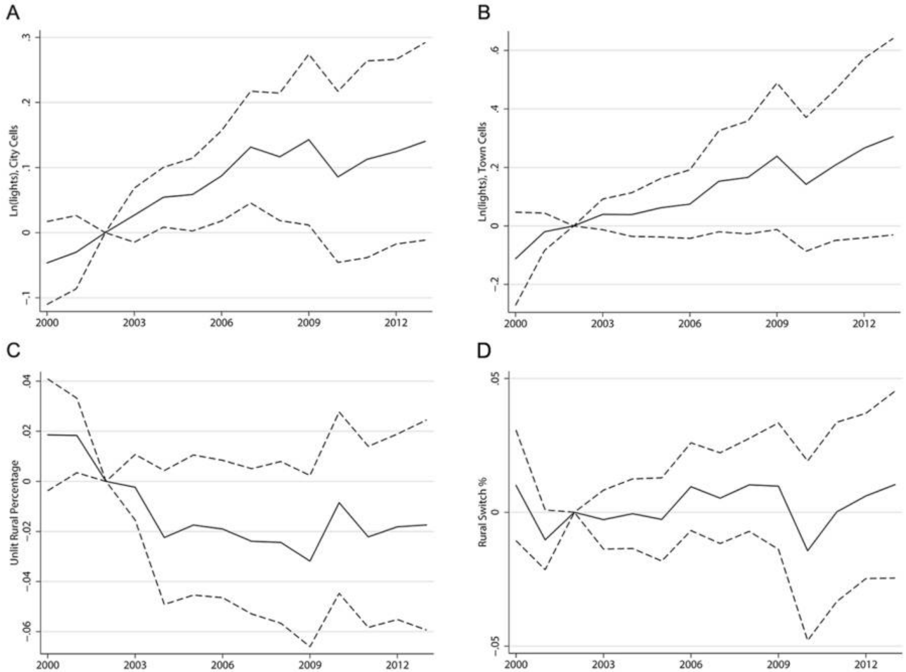

Figure 7. Effect of the 2003–13 oil price boom on illumination in cities (*A*), illuminaion in towns (*B*), the share of national population living in unlit rural areas (*C*), and the share of unlit rural areas that become lit (*D*) (classification in 2000). Each point represents the conditional difference between oil- and gas-dependent and nondependent countries, compared to the conditional difference in 2002 (95% confidence bands in dashes). The specification controls for convergence and includes year and country fixed effects and regional linear trends (see eq. [2]). The 2009 Great Recession may have affected these results if it had an asymmetric effect on oil-dependent versus nondependent economies. Tables of these estimates, along with control variables, are found in appendix E.2.

a 10% level, and at a 5% level when using lights per capita (not shown). They also display a clear upward trend from 2002 to 2013. This means that light intensity in cells that were already lit at the beginning of our sample became brighter in oil- and gas-dependent countries relative to nondependent ones.

There is no evidence that the oil boom benefited the rural poor. Despite the increase in aggregate illumination, there is no evidence that the share of people living in unlit rural areas fell in oil- and gas-dependent countries, relative to nondependent ones (fig. 7C). In 2013 the unlit rural percentage in oil- and gas-dependent countries had not changed relative to nondependent countries since the 2002 base year. There was also no evidence of any trend between 2003 and 2013.

This finding, that an oil price boom increases lights within the grid, rather than by extending the grid, is consistent with earlier findings that resource-dependent countries do not invest enough. The key policy challenge for these countries is to convert below- to above-ground assets to maintain permanently higher spending (Friedman 1957; Hartwick 1977). In developed countries this requires a sovereign wealth fund (van den Bremer et al. 2016), while in developing ones with high domestic returns it requires investing locally (Collier et al. 2010; van der Ploeg and Venables 2011; Venables and Wills 2016). However, evidence from the genuine savings literature finds that this domestic investment is not happening, in favor of public consumption (Collier et al. 2010; van der Ploeg 2011; Bhattacharyya and Collier 2013). Our results are consistent with inadequate investment because an increase in lighting at the extensive margin is likely to need investment in the grid; while an increase at the intensive margin can be driven by public consumption, which may cut taxes for formal sectors or employ additional public sector workers, both of which disproportionately benefit areas with existing lighting. In general, higher consumption can increase light intensity in existing lit areas, for example, through restaurants and houses leaving lights on longer and cars staying on the road later.

The literature on why investment in resource-dependent countries is too low has found both economic and institutional reasons. Economic reasons include population growth and anticipation of better times (van der Ploeg 2011) and absorption constraints as governments must develop the capacity to evaluate projects, construction workers must be trained, and new projects must be sequenced ("investing in investing," Collier 2010; van der Ploeg 2012; van der Ploeg and Venables 2013). Institutional reasons include the fact that resources encourage politicians to neglect productive investment in favor of rent seeking (Gelb 1988; Lane and Tornell 1996; Tornell and Lane 1999; Auty 2001; Ross 2001; Caselli and Cunningham 2009), and corruption (Dietz et al. 2007; Vicente 2010; Caselli and Michaels 2013; Sala-i-Martin and Subramanian 2013; Andersen et al. 2017), particularly in countries with poor initial institutions (Mehlum et al. 2006; Robinson et al. 2006; Boschini et al. 2007; Collier and Goderis 2008; Bhattacharyya and Hodler 2010; Libman 2013). The larger pie created by resources intensifies political competition, which reduces politicians' discount rates and encourages them to underinvest (Alesina and Tabellini 1989; Caselli 2006). Furthermore, point-source resources encourage conflict by increasing the returns to fighting (Fearon and Laitin 2003; Collier and Hoeffler 2004, 2005; Ross 2004; Fearon 2005; Humphreys 2005; Lujala 2010; Dube and Vargas 2013; Lei and Michaels 2014), which worsens institutions and discourages investment in productive capital (van der Ploeg 2011).

The finding in figure 7*C* that the unlit rural percentage did not change might mask lights turning on in unlit cells, or people moving from unlit to lit cells. The rural switch percentage in figure 7*D* shows that the oil price boom did not cause lights in unlit cells to turn on. We address the question of whether people moved from unlit to lit cells next.

3.4.3. High Oil Prices Did Not Affect the Population Distribution

To understand whether people in unlit areas may have benefited from the 2000s oil boom by moving to lit areas, we now turn our attention to urbanization. There is no evidence that high oil prices caused the population in oil-dependent countries to reallocate between cities, towns, and unlit rural areas. While there has been a global trend for people to leave unlit rural areas for towns and cities (see fig. 5), this has happened at the same pace in both oil-dependent and nondependent countries. Figure 8 shows that from 2002 to 2013 the population share of cities, towns, and rural areas in oil- and gas-dependent countries did not significantly change relative to nondependent ones. This differs from the result in Gollin et al. (2016), who argue that natural resource rents drive urbanization. Using World Bank data on urbanization rates, they find that non-OECD countries with oil are more urbanized than those without, which we confirm in figure 5. They also find a small effect of resource exports in a panel setting from 1960 to 2010. Our results can be reconciled if their effects on urbanization have diminished in recent years.

This result is robust to annual interpolation in the LandScan data as it covers a period of 13 years, which includes multiple censuses for each country. It is also robust to cell-by-cell measurement error as it is based on broad national aggregates of areas classed as cities, towns, or unlit rural. Furthermore, replicating the analysis using World Bank data on urbanization rates we find the same result (see fig. G.1i).

3.5. Results: Quantity Booms from Giant Oil and Gas Discoveries

3.5.1. Giant Oil Discoveries Stimulated Economic Growth at a 6-Year Lag

Aggregate illumination increased approximately 6 years after countries discovered a giant oil field, relative to the control group (of other non-OECD countries that have not yet discovered oil). This is illustrated in figure 9, which shows the effect on nighttime lights for countries that discovered oil t years ago, scaled for the size of the discovery relative to GDP. Nine years after countries discovered oil worth 100% of GDP, lights and lights per capita were 19% higher (0.17 log points), and GDP per capita was

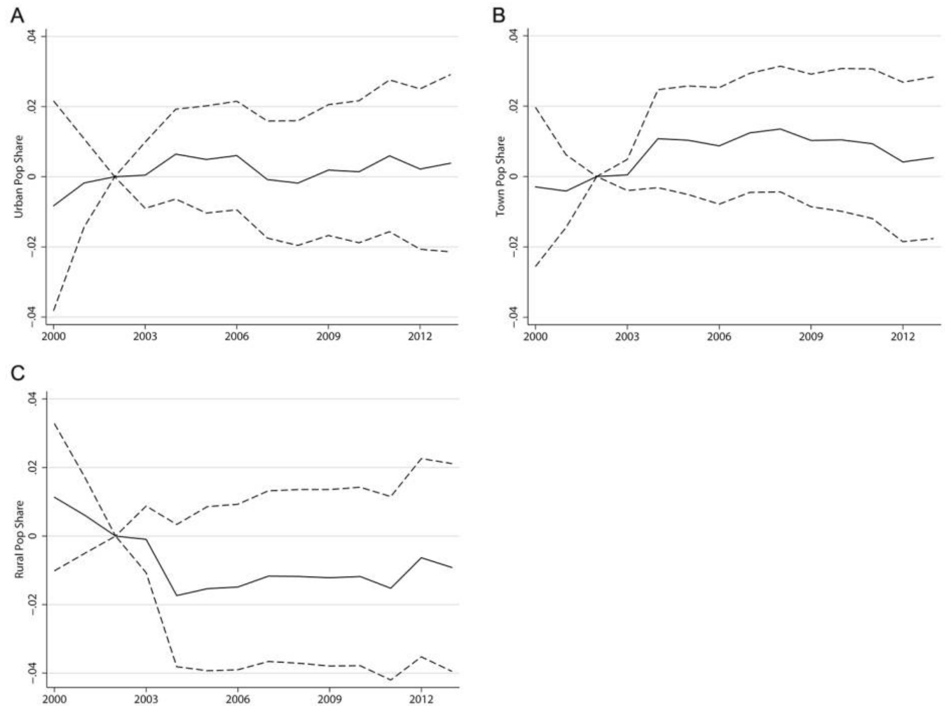

Figure 8. Effect of the 2003–13 oil price boom on the national share of population living in cities (*A*), towns (*B*), and rural areas (*C*) (classification in 2000). Each point represents the conditional difference between oil- and gas-dependent and nondependent countries, compared to the conditional difference in 2002 (95% confidence bands in dashes). The specification includes year and country fixed effects and regional linear trends (see eq. [2]). Tables of these estimates, along with control variables, are found in appendix E.2.

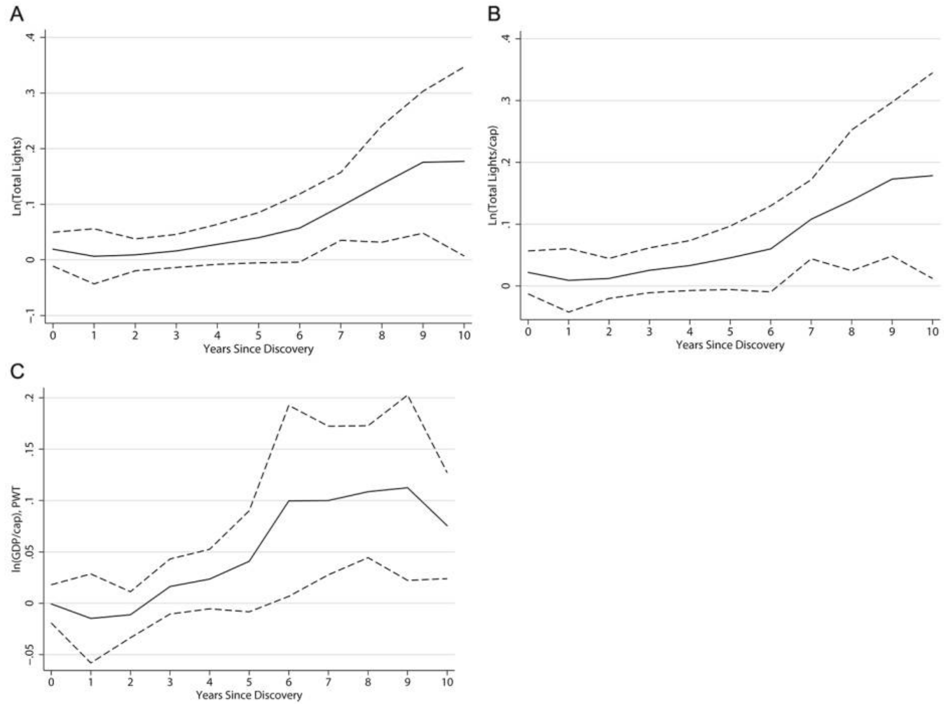

Figure 9. The difference in aggregate (log) lights (*A*), log aggregate lights per capita (*B*), and PPP-adjusted real GDP per capita (C), after a country discovers a giant oil field worth 100% of GDP, relative to countries that have not yet discovered oil (95% confidence bands in dashes). The specification controls for convergence and includes year and country fixed effects and regional linear trends (see eq. [3]). Tables of these estimates, along with control variables, are found in appendix E.2.

8% higher (0.08 log points), than they were in 2002 relative to the control group. The delay between discovery and production is consistent with oil discoveries being a form of news shock, confirming the results of Arezki et al. (2016).[21]

3.5.2. Growth from Oil Discoveries Was Not Shared with the Rural Poor

Once again, higher aggregate illumination masked a rise in regional inequality as it was confined to cities and towns. Illumination in the cities and towns of countries that discovered oil grew by 15% and 22% (0.14 and 0.20 log points), respectively, over our sample, relative to the control group. This effect was significant at the 5% level after 4 years in cities and 6 years in towns, consistent with a delay between discovery and production. This again means that light intensity in cells that were already lit at the beginning of our sample rose in oil-dependent countries relative to nondependent ones.

There is no evidence that giant oil discoveries benefited the rural poor. Up to 10 years after a discovery the unlit rural percentage only fell by 1 percentage point relative to the control group, and this is not significantly different from zero (fig. 10C). There is no evidence that previously unlit rural areas became lit due to a discovery; if anything, unlit areas in the control group illuminated faster (fig. 10*D*). This is consistent with the findings from the 2003–13 price boom discussed in section 3.4.2. Note that this result does not comment on the local effects of oil wells (see the review by Cust and Poelhekke 2015), as it is making cross-country rather than within-country comparisons.

3.5.3. Oil Discoveries Cause Some Population to Reallocate from Rural Areas to Towns but Not Cities

There is some evidence that oil discoveries led to a minor reallocation of population from unlit rural areas to towns. Figure 11 shows that the unlit rural population share was approximately 1 percentage point lower 10 years after an oil discovery worth 100% of GDP, relative to the control group. It also slightly trended downward in the intervening years. Towns saw a similar increase in their population share. These effects are of a similar order of magnitude to the (insignificant) fall in the unlit rural percentage in figure 10C. There is no discernible effect on the population of cities, which we confirm using World Bank data on urbanization rates (see fig. G.1ii, noting the very small magnitude of effects). As with the results in section 3.4.3, these results are robust to interpolation and measurement error in the LandScan data because they

21. These results are presented in table form in app. E.2. It shows that the estimated effect peaks at the thirteenth lag and then declines to zero by the twenty-second lag. This is consistent with the typical production profile of an oil field, which peaks approximately 7 years after production starts and declines to nearly zero after approximately 20 years (Höök et al. 2009). The effects of any subsequent discoveries are captured by earlier lags. However, we focus on lags 0–10 because beyond that the estimates are based on a shrinking set of discoveries, so are only included to ensure that lags 0–10 are estimated against the correct counterfactual.

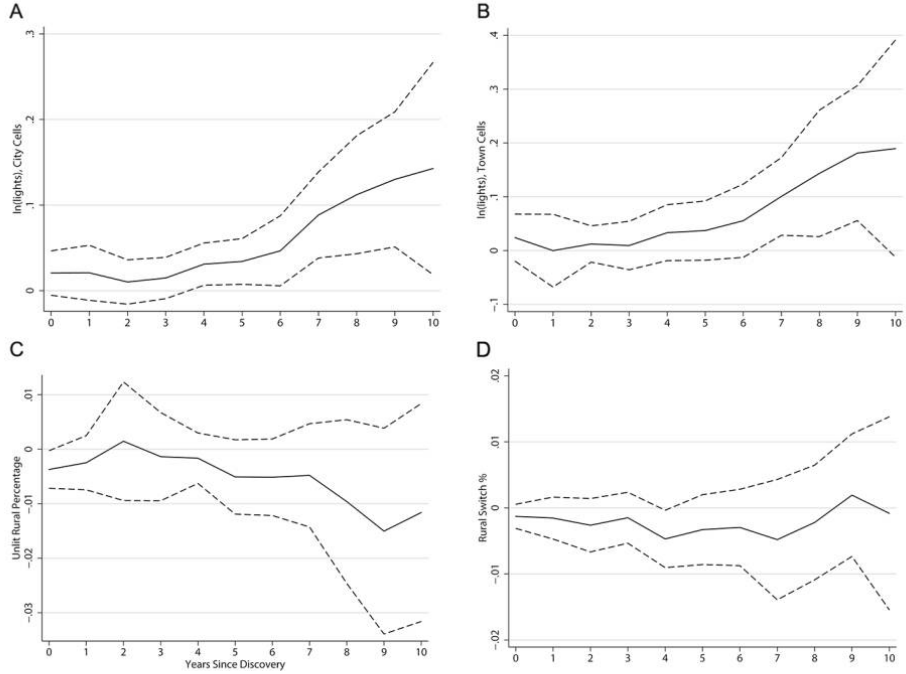

Figure 10. The difference in illumination in cities (*A*), illumination in towns (*B*), the share of national population living in unlit rural areas (C), and the share of unlit rural areas that become lit (*D*) (classification in 2000), after a country discovers a giant oil field worth 100% of GDP, relative to countries that have not yet discovered oil (95% confidence bands in dashes). The specification controls for convergence and includes year and country fixed effects and regional linear trends (see eq. [3]). Tables of these estimates, along with control variables, are found in appendix E.2.

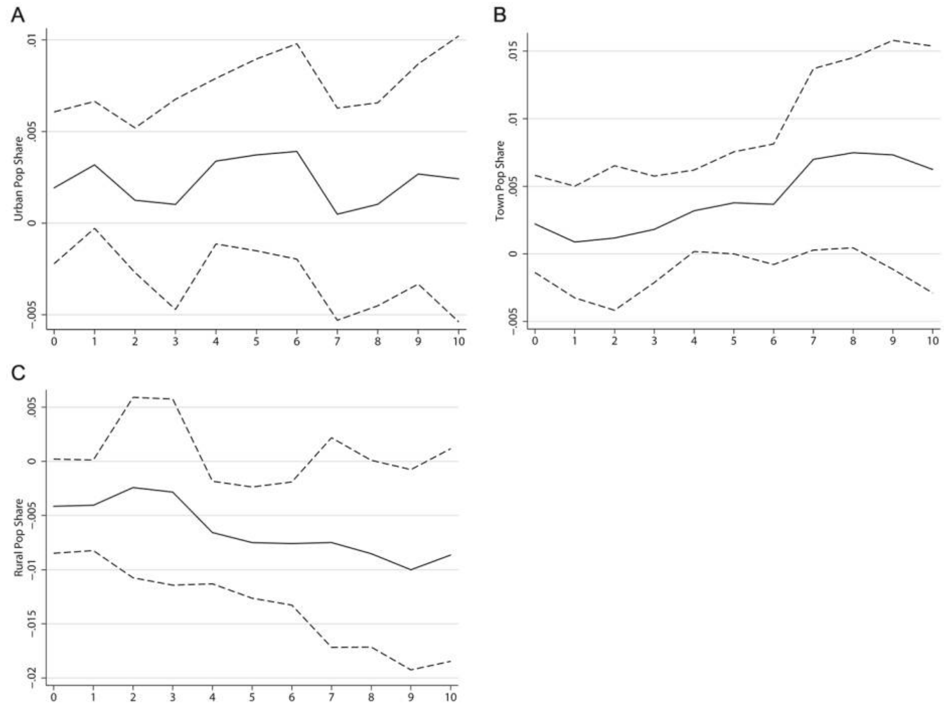

Figure 11. Difference in the population share of cities (*A*), towns (*B*), and rural areas (C) (classification in 2000), after a country discovers a giant oil field worth 100% of GDP, relative to countries that have not yet discovered oil (95% confidence bands in dashes). The specification includes year and country fixed effects and regional linear trends (see eq. [3]). Tables of these estimates, along with control variables, are found in appendix E.2.

cover an extensive time period and are based on broad, nationally aggregated area classifications.

3.6. Mechanisms

The aggregated results above show that while rural areas are illuminating around the world, oil booms do not hasten the process. However, these aggregate results might hide some details in the darkness. We now use the hazard model in equation (4) to understand what causes an unlit and inhabited rural cell to light up. We find that the probability of illumination is increased by (i) being adjacent to existing lit cells, (ii) being close to the capital, (iii) having a high population density, and (iv) being in a country with high aggregate light growth since the start of the sample, as shown in table 1. However, we do not find that the first three mechanisms are more or less active in oil-dependent countries during a price boom. Instead, we find that oil-dependent

Table 1. Results for a Model of the Hazard Rate of Unlit Rural Areas Switching On in a Given Year

	Switch On (1)
Adjacent	.031***
	(.005)
Adjacent × oil/gas dependence	.000
	(.005)
<100 km from capital	.006**
	(.002)
(<100 km from capital) × oil/gas dependence	.004
	(.004)
Population density	.004*
	(.002)
Population density × oil/gas dependence	−.002
	(.002)
Lights growth since 2000	.019***
	(.005)
(Lights growth since 2000) × oil/gas dependence	−.007*
	(.003)
N	7,479,895
R^2	.026

Note. The dependent variable is an indicator for the cell switching on in a given year. Regression includes country and year fixed effects. Robust standard errors clustered at the country level are reported in parentheses.

* Significant at 10%.

** Significant at 5%.

*** Significant at 1%.

countries are less efficient at converting growth to poverty reduction than their nondependent counterparts.

Being next to a lit cell suggests being near an existing electricity network. We hypothesize that this would increase the probability of illumination, allowing public grid investment rather than expensive private off-grid generators. We find that being next to a lit cell increases the probability of an unlit cell switching on by 3 percentage points. Between 2003 and 2013 adjacent cells were ~15 percentage points more likely to be illuminated than nonadjacent cells (see fig. F.1A, which shows the cumulative effect of this mechanism not controlling for the other three in eq. [4]).[22] However, adjacency did not increase the probability of illumination in oil-dependent relative to nondependent countries, suggesting that oil revenues were not systematically invested in larger electricity networks.

Being within 100 km of the capital, after controlling for adjacency to lit cells, can be interpreted as a proxy for geographic and political connections. We find that proximity to the capital raises the probability of becoming illuminated by 0.6 percentage points. These cells cumulatively gained ~10 percentage points more illumination from 2003 to 2013 (fig. F.1B). Again, being near to the capital did not increase the chance of illumination in oil-dependent versus nondependent countries, which suggests that oil revenues did not change the connectedness of these areas.

Cells with a higher population density will benefit more from electrical infrastructure and may be more politically organized. We see that a 1 standard deviation increase in population density increases the probability of illumination by 0.4 percentage points. Over the decade of high oil prices from 2003, cells with 1 standard deviation more population were 2 percentage points more likely to become lit (fig. F.2). However, as with the previous mechanisms, there is no evidence that oil-dependent countries invested more in electrical infrastructure in areas with high population density, which is consistent with more autocratic regimes (Min 2010). This further suggests that oil wealth is not benefiting areas with the highest density of rural poverty.

Finally, growth in aggregate national lights can be interpreted as a proxy for broader economic growth since 2000. We find that a 1 log point increase in past aggregate light growth raises the probability of a rural cell switching on by 1.9 percentage points. This indicates that, in general, aggregate economic growth does reduce rural poverty. However this effect is 0.7 percentage points smaller in oil-dependent countries, again implying that they are less effective at converting growth into rural poverty reduction than other countries, even after controlling for spatial mechanisms at the grid-cell level.

---

22. Graphs i. and ii. of figure F.1A show the results of regressing the cumulative switch-on percentage on indicators for adjacency to lit cells interacted with year fixed effects (also controlling for year and country fixed effects). Graph iii shows the results of a triple-difference specification that evaluates if the effect differs between dependent and nondependent countries.

These results, particularly the effect of being within 100 km of the capital, are consistent with earlier findings that any surge in investment in developing countries is often inequitable. "White elephants"—partisan projects with negative social surplus—can arise if politicians cannot credibly commit to their supporters (Robinson and Torvik 2005); and investments may focus on the homelands of rulers (Hodler and Raschky 2014). Mineral wealth can further erode the institutions that ensure equitable investments (Bulte et al. 2005; Tsui 2011) by disenfranchising the middle classes (Bourguignon and Verdier 2000), allowing governments to pacify dissent (Isham et al. 2005; Robinson et al. 2006) or buy off challengers (Acemoglu et al. 2004), and encouraging politicians to neglect them (Acemoglu and Robinson 2006). Inequitable investments further compound the effects on the poor of insufficient investment in resource-dependent countries, discussed in section 3.4.2. These are all reasons for why an oil price boom would not lead to the electricity grid expanding, increasing the probability of an unlit cell becoming lit and reducing the share of people living in darkness at night.

3.7. Robustness

To test the robustness of our main results we try five alternative specifications, with the results in appendix G. The first three focus on the price boom specification in equation (2), the fourth on oil discoveries in equation (3), and the fifth on addressing measurement error in the unlit rural percentage.

The first uses the same specification but drops OPEC countries from the treatment group, to further reduce the possibility that the oil price shock is endogenous due to supply disruptions (fig. G.3i). The estimates are more or less the same, but with larger standard errors because we drop a large proportion of our treatment group. Total lights grow by the same amount and the results are significant at the 10% (but not 5%) level. The oil boom has no effect on poverty.

The second replaces the year fixed effects and regional trends in equation (2) with region-year fixed effects. This controls for common shocks at the regional, rather than global level. This is a more restrictive specification as identification is based strictly on within-region comparisons, sometimes with small numbers of countries in a given region. Still, results are similar to the main specification as seen in figure G.3ii.

The third replaces the binary indicator for dependence, D_i, with a continuous variable, $Rents_i$. In the main specification we use a dependence indicator because it is simple and yields transparent graphical results. However, some information is lost in the binary classification. To account for the degree of resource dependence we use a continuous variable, $Rents_i$, which measures average oil rents as a share of GDP from 2000 to 2012 (implying a linear relationship). We find that if oil rents account for 10% of GDP, then the five times increase in prices over our sample would have increased lights by 14% (0.13 log points), or 9% (0.09 log points) excluding Equatorial Guinea (a high-growth outlier where rents account for 75% of GDP, see fig. G.4).

The unlit rural percentage drops by 1.3% (0.013 log points), driven primarily by migration in Equatorial Guinea, Gabon, and the Republic of Congo. Continuous oil dependence again has a small and insignificant impact on rural cells becoming lit, with or without Equatorial Guinea (not shown).

The fourth tests whether the type of resource discovery (oil/gas, onshore/offshore) matters. Figure G.5 shows that giant oil discoveries have a much larger effect on national lights than gas discoveries within 10 years of discovery. This might be because gas discoveries require more infrastructure, such as pipelines and liquification plants, which delays their effect. Offshore discoveries have a larger effect than onshore discoveries, and they disproportionately involve oil (40/65 of offshore discoveries are oil, compared to 11/35 of onshore discoveries).

The fifth controls for "flickering" lights as a possible source of measurement error. In the main analysis rural poverty is measured using the share of national population living in cells that are unlit (the unlit rural percentage). Some of these marginal cells flicker between lit and unlit from year to year, due in part to changes in satellite sensitivity. This may induce measurement error in the unlit rural percentage, reducing the precision of regressions where it is the dependent variable. To address this we rerun the analysis replacing the "unlit rural percentage" with two measures. The first is the "never-been-lit rural percentage," which is the population share of cells that had never been lit in the year of observation. This measure counts marginal, flickering cells as lit and therefore defines a deeper level of poverty that is not subject to flickering. The second is the "barely lit percentage," which is the population share of cells that are either unlit or in the bottom 10% of lit cells globally in a given year. This measure counts marginal cells as unlit, and therefore defines a broader level of poverty. Figure G.6 shows that in both cases the results are consistent with the main specification: oil and gas booms do not reduce the population share living in never-been-lit or barely lit cells, relative to countries that do not experience the boom.

4. CONCLUSION

This paper attempts to answer the question, Do oil booms reduce inequality and rural poverty? To do this we measure rural poverty by counting people who live in darkness at night, and then conduct two experiments by exploiting exogenous variation in oil prices and quantities. We find that oil booms stimulate aggregate economic activity; however, this promotes regional inequality as it is confined to cities and towns but does not benefit the rural poor.

We construct an annual, globally balanced panel of rural poverty by counting the people who live in darkness at night. This is done by combining two high-resolution spatial data sets: on night-time illumination as a proxy for economic activity, and on population. Evaluating the effectiveness of our measure against 636,448 household DHS surveys from 36 countries, we find that it accurately identifies 74% of people as being above or below the threshold for extreme poverty at 100 km^2 resolution and more at higher resolutions.

We then conduct two experiments using the rise in oil prices from 2003 to 2013 and giant oil discoveries. We argue that the oil price shock was due to an exogenous increase in global oil demand and further control for endogeneity by excluding OPEC countries in a robustness test. We also argue that giant oil discoveries are well identified after controlling for time and country fixed effects.

Our results provide the first evidence that oil booms promote regional inequality by benefiting those in towns and cities but not the rural poor. There is growing agreement that the natural resource curse does not affect all countries but is conditional on institutions that permit rent seeking and corruption. If this is true, then the curse should be associated with inequality and poverty, as resource wealth accrues to a small privileged few. We find that high oil prices and new discoveries stimulate economic activity (proxied by lights) in countries with oil and gas relative to those without. However, this new activity is restricted to towns and cities. Both types of oil boom have no effect on the rural poor. There is no evidence that oil booms cause unlit rural areas to become illuminated, and only a small number of people (1%) leave unlit areas for towns and cities after a giant oil or gas discovery. In general, illumination is more likely to happen in areas that are close to lit areas, close to the capital, have high population density, and are in a country that is growing quickly. However, none of these areas benefit any more during an oil price boom; if anything, booming countries are less efficient at converting growth to poverty reduction. These results suggest that oil booms increase regional inequality.

This work lends itself to a wide range of extensions. Darkness may be used to study other determinants of rural poverty and policies designed to alleviate it, as it provides geographically disaggregated and globally balanced panel data. Aid and humanitarian interventions may use poverty maps of unlit rural areas, like those in the appendix, to better direct their efforts. Finally, this research suggests that there is extensive scope for research into policies that can share resource revenues more equitably with the rural poor.

REFERENCES

Acemoglu, Daron, and James A. Robinson. 2006. De facto political power and institutional persistence. *American Economic Review* 96:325–30.

Acemoglu, Daron, Thierry Verdier, and James A. Robinson. 2004. Kleptocracy and divide-and-rule: A model of personal rule. *Journal of the European Economic Association* 2:162–92.

Albert, Donald P. 2012. *Geospatial technologies and advancing geographic decision making: Issues and trends*. Hershey, PA: IGI Global.

Alesina, Alberto F., Stelios Michalopoulos, and Elias Papaioannou. 2015. Ethnic inequality. *Journal of Political Economy* 124 (2): 428–88.

Alesina, Alberto, and Guido Tabellini. 1989. External debt, capital flight and political risk. *Journal of International Economics* 27:199–220.

Alexeev, Michael, and Robert Conrad. 2009. The elusive curse of oil. *Review of Economics and Statistics* 91:586–98.

Almås, Ingvild, Åshild Auglænd Johnsen, and Andreas Kotsadam. 2014. Poverty in China seen from outer space. Technical report, Memorandum, University of Oslo, Department of Economics.

Andersen, Jørgen Juel, Niels Johannesen, David Dreyer Lassen, and Elena Paltseva. 2017. Petro rents, political institutions, and hidden wealth: Evidence from offshore bank accounts. *Journal of the European Economic Association* 15 (4): jvw019.

Aragón, Fernando M., and Juan Pablo Rud. 2013. Natural resources and local communities: Evidence from a Peruvian gold mine. *American Economic Journal: Economic Policy* 5:1–25.

Arezki, Rabah, Valerie A. Ramey, and Liugang Sheng. 2016. News shocks in open economies: Evidence from giant oil discoveries. *Quarterly Journal of Economics* 132 (1): 103–55.

Atkinson, Anthony B. 1970. On the measurement of inequality. *Journal of Economic Theory* 2:244–63.

Atkinson, Anthony B., Thomas Piketty, and Emmanuel Saez. 2009. Top incomes in the long run of history. Technical report, National Bureau of Economic Research, Cambridge, MA.

Auty, Richard M. 2001. The political economy of resource-driven growth. *European Economic Review* 45:839–46.

Bacon, Robert, Soma Bhattacharya, and Masami Kojima. 2010. Expenditure of low-income households on energy: Evidence from Africa and Asia. World Bank Oil, Gas, and Mining Policy Division (June).

Baunsgaard, Thomas, Marcos Poplawski-Ribeiro, Mauricio Villafuerte, and Christine Richmond. 2012. Fiscal frameworks for resource rich developing countries. Technical report, International Monetary Fund, Washington, DC.

Bhattacharyya, Sambit, and Paul Collier. 2013. Public capital in resource rich economies: Is there a curse? *Oxford Economic Papers* 66:1–24.

Bhattacharyya, Sambit, and Roland Hodler. 2010. Natural resources, democracy and corruption. *European Economic Review* 54:608–21.

Bhattacharyya, Sambit, and Jeffrey G. Williamson. 2013. Distributional impact of commodity price shocks: Australia over a century. OxCarre Research Paper 117, Oxford.

Blumenstock, Joshua, Gabriel Cadamuro, and Robert On. 2015. Predicting poverty and wealth from mobile phone metadata. *Science* 350:1073–76.

Boschini, Anne D., Jan Pettersson, and Jesper Roine. 2007. Resource curse or not: A question of appropriability. *Scandinavian Journal of Economics* 109:593–617.

Bourguignon, François, and Thierry Verdier. 2000. Oligarchy, democracy, inequality and growth. *Journal of Development Economics* 62:285–313.

Brunnschweiler, Christa N., and Erwin H. Bulte. 2008. The resource curse revisited and revised: A tale of paradoxes and red herrings. *Journal of Environmental Economics and Management* 55:248–64.

Bulte, Erwin H., Richard Damania, and Robert T. Deacon. 2005. Resource intensity, institutions, and development. *World Development* 33:1029–44.

Caselli, Francesco. 2006. Power struggles and the natural resource curse. Unpublished manuscript, London School of Economics.

Caselli, Francesco, and Tom Cunningham. 2009. Leader behaviour and the natural resource curse. *Oxford Economic Papers* 61:628–50.

Caselli, Francesco, and Guy Michaels. 2013. Do oil windfalls improve living standards? Evidence from Brazil. *American Economic Journal: Applied Economics* 5:208–38.

Chen, Shaohua, and Martin Ravallion. 2010. The developing world is poorer than we thought, but no less successful in the fight against poverty. *Quarterly Journal of Economics* 125:1577–1625.

Collier, Paul. 2010. *The plundered planet: Why we must—and how we can—manage nature for global prosperity*. Oxford: Oxford University Press.

Collier, Paul, and Benedikt Goderis. 2008. Commodity prices, growth, and the natural resource curse: Reconciling a conundrum. *European Economic Review* 56:1241–60.

Collier, Paul, and Anke Hoeffler. 2004. Greed and grievance in civil war. *Oxford Economic Papers* 56:563–95.

———. 2005. Resource rents, governance, and conflict. *Journal of Conflict Resolution* 49:625–33.

Collier, Paul, Rick van der Ploeg, Michael Spence, and Anthony J. Venables. 2010. Managing resource revenues in developing economies. *IMF Staff Papers* 57:84–118.

Cust, James, and Steven Poelhekke. 2015. The local economic impacts of natural resource extraction. *Annual Review of Resource Economics* 7:251–68.

Deaton, Angus. 1985. The measurement of welfare: Theory and practical guidelines. Development Research Department, World Bank, Washington, DC.

——— . 1988. Quality, quantity, and spatial variation of price. *American Economic Review* 78 (3): 418–30.

———. 1997. *The analysis of household surveys: A microeconometric approach to development policy*. Washington, DC: World Bank Publications.

———. 2005. Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics* 87:1–19.

Deaton, Angus S., and John Muellbauer. 1986. On measuring child costs: With applications to poor countries. *Journal of Political Economy* 94 (4): 720–44.

Dell, Melissa. 2010. The persistent effects of Peru's mining *mita*. *Econometrica* 78:1863–1903.

Dietz, Simon, Eric Neumayer, and Indra De Soysa. 2007. Corruption, the resource curse and genuine saving. *Environment and Development Economics* 12:33–53.

Doll, Christopher N. H. 2008. CIESIN thematic guide to night-time light remote sensing and its applications. Center for International Earth Science Information Network of Columbia University, Palisades, NY.

Doll, Christopher N. H., Jan-Peter Muller, and Jeremy G. Morley. 2006. Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics* 57:75–92.

Dorélien, Audrey, Deborah Balk, and Megan Todd. 2013. What is urban? Comparing a satellite view with the Demographic and Health Surveys. *Population and Development Review* 39:413–39.

Dube, Oeindrila, and Juan F. Vargas. 2013. Commodity price shocks and civil conflict: Evidence from Colombia. *Review of Economic Studies* 80:1384–1421.

Elvidge, Christopher D., Kimberly E. Baugh, S. J. Anderson, P. C. Sutton, and T. Ghosh. 2012. The Night Light Development Index (NLDI): A spatially explicit measure of human development from satellite data. *Social Geography* 7:23–35.

Elvidge, Christopher D., Kimberly E. Baugh, Paul C. Sutton, Budhendra Bhaduri, Benjamin T. Tuttle, Tilotamma Ghosh, Daniel Ziskin, and Edward H. Erwin. 2011. Who's in the dark? Satellite based estimates of electrification rates. In *Urban remote sensing: Monitoring, synthesis and modeling in the urban environment*, ed. Xiaojun Yang. 211–24. Hoboken, NJ: Wiley.

Elvidge, Christopher D., Paul C. Sutton, Tilottama Ghosh, Benjamin T. Tuttle, Kimberly E. Baugh, Budhendra Bhaduri, and Edward Bright. 2009. A global poverty map derived from satellite data. *Computers and Geosciences* 35:1652–60.

Fearon, James D. 2005. Primary commodity exports and civil war. *Journal of Conflict Resolution* 49:483–507.

Fearon, James D., and David D. Laitin. 2003. Ethnicity, insurgency, and civil war. *American Political Science Review* 97:75–90.

Friedman, Milton. 1957. The permanent income hypothesis. In *A theory of the consumption function*, 20–37. Princeton, NJ: Princeton University Press.

Gelb, Alan H. 1988. *Oil windfalls: Blessing or curse?* Oxford: Oxford University Press.

Goderis, Benedikt, and Samuel W. Malone. 2011. Natural resource booms and inequality: Theory and evidence. *Scandinavian Journal of Economics* 113:388–417.

Gollin, Douglas, Remi Jedwab, and Dietrich Vollrath. 2016. Urbanization with and without industrialization. *Journal of Economic Growth* 21:35–70.

Gylfason, Thorvaldur, and Gylfi Zoega. 2003. Inequality and economic growth: Do natural resources matter? In *Inequality and growth: Theory and policy implications*, ed. Theo S. Eicher and Stephen J. Turnovsky, 255–92. Cambridge, MA: MIT Press.

Halbouty, Michel T., Arthur Augustus Meyerhoff, Robert E. King, Robert H. Dott Sr., H. Douglas Klemme, and Theodore Shabad. 1970. In *World's giant oil and gas fields, geologic factors affecting their formation, and basin classification. Part I: Giant oil and gas fields*, ed. Michel T. Halbouty, 502–28. McLean, VA: GeoScienceWorld.

Hartwick, John M. 1977. Intergenerational equity and the investing of rents from exhaustible resources. *American Economic Review* 67:972–74.

Henderson, J. Vernon, Adam Storeygard, and David N. Weil. 2011. A bright idea for measuring economic growth. *American Economic Review* 101:194.

———. 2012. Measuring economic growth from outer space. *American Economic Review* 102:994–1028.

Hodler, Roland, and Paul A. Raschky. 2014. Regional favoritism. *Quarterly Journal of Economics* 129 (2): qju004.

Höök, Mikael, Bengt Söderbergh, Kristofer Jakobsson, and Kjell Aleklett. 2009. The evolution of giant oil field production behavior. *Natural Resources Research* 18:39–56.

Horn, Myron K. 2003. Giant fields, 1868–2003 (CD-ROM). In *Giant oil and gas fields of the decade, 1990–1999*, ed. M. K. Halbouty. Houston: AAPG Memoir 78.

———. 2004. Giant fields, 1868–2004 (CD-ROM, revision to 2003 version). In *Giant oil and gas fields of the decade, 1990–1999*, ed. M. K. Halbouty. Houston: AAPG/ Datapages Miscellaneous Data Series. Version 1.2, 2004.

Humphreys, Macartan. 2005. Natural resources, conflict, and conflict resolution: Uncovering the mechanisms. *Journal of Conflict Resolution* 49:508–37.

ICF. 2012. Demographic and health survey: Sampling and household listing manual. Technical report, MEASURE DHS/ICF International.

Isham, Jonathan, Michael Woolcock, Lant Pritchett, and Gwen Busby. 2005. The varieties of resource experience: Natural resource export structures and the political economy of economic growth. *World Bank Economic Review* 19:141–74.

James, Alexander. 2015. The resource curse: A statistical mirage? *Journal of Development Economics* 114:55–63.

Jean, Neal, Marshall Burke, Michael Xie, W. Matthew Davis, David B. Lobell, and Stefano Ermon. 2016. Combining satellite imagery and machine learning to predict poverty. *Science* 353:790–94.

Jedwab, Remi, Edward Kerby, Alexander Moradi, et al. 2015. History, path dependence and development: Evidence from colonial railroads, settlers and cities in Kenya. *Economic Journal* 127:1467–94.

Jedwab, Remi, and Alexander Moradi. 2015. The permanent effects of transportation revolutions in poor countries: Evidence from Africa. *Review of Economics and Statistics* 98 (2): 268–84.

Kavanagh, Michael. 2013. Abandoned oil discoveries enjoy revival. *Financial Times* (April 8). https://www.ft.com/content/2b41056a-a054-11e2-88b6-00144feabdc0.

Kilian, Lutz. 2009. Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99:1053–69.

Kuznets, Simon. 1937. Appendices to "National income and capital formation, 1919–1935." In *National income and capital formation, 1919–1935*, 61–90. Cambridge, MA: National Bureau of Economic Research.

———. 1953. International differences in income levels: Reflections on their causes. *Economic Development and Cultural Change* 2 (1): 3–26.

———. 1955. Economic growth and income inequality. *American Economic Review* 45 (1): 1–28.

Kuznets, Simon, Lillian Epstein, and Elizabeth Jenks. 1941. *National income and its composition, 1919–1938*. NBER Books, vol. 1. Cambridge, MA: National Bureau of Economic Research.

Lane, Philip R., and Aaron Tornell. 1996. Power, growth, and the voracity effect. *Journal of Economic Growth* 1:213–41.

Leach, Gerald. 1992. The energy transition. *Energy Policy* 20:116–23.

Lee, Kenneth, Eric Brewer, Carson Christiano, Francis Meyo, Edward Miguel, Matthew Podolsky, Javier Rosa, and Catherine Wolfram. 2014. Barriers to electrification for "under grid" households in rural Kenya. Technical report, National Bureau of Economic Research, Cambridge, MA.

Lee, Lisa Yu-Ting. 2013. Household energy mix in Uganda. *Energy Economics* 39:252–61.

Lei, Yu-Hsiang, and Guy Michaels. 2014. Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics* 110:139–57.

Leigh, Andrew. 2009. Top incomes. In *Oxford handbook of economic inequality*, 150–76. Oxford: Oxford University Press.

Libman, Alexander. 2013. Natural resources and sub-national economic performance: Does sub-national democracy matter? *Energy Economics* 37:82–99.

Lippert, Alexander, et al. 2014. Spill-overs of a resource boom: Evidence from Zambian copper mines. Technical report, University of Oxford, Oxford Centre for the Analysis of Resource Rich Economies.

Lloyd, Christopher T., Alessandro Sorichetta, and Andrew J. Tatem. 2017. High resolution global gridded data for use in population studies. *Scientific Data*, vol. 4, article 170001.

Lujala, Päivi. 2010. The spoils of nature: Armed civil conflict and rebel access to natural resources. *Journal of Peace Research* 47:15–28.

McGregor, Thomas, and Samuel Wills. 2017. Surfing a wave of economic growth. OxCarre Research Paper 170, Oxford.

Mehlum, Halvor, Karl Moene, and Ragnar Torvik. 2006. Institutions and the resource curse. *Economic Journal* 116:1–20.

Michalopoulos, Stelios, and Elias Papaioannou. 2013. Pre-colonial ethnic institutions and contemporary African development. *Econometrica* 81:113–52.

———. 2014. National institutions and subnational development in Africa. *Quarterly Journal of Economics* 129:151–213.

Mills, Evan. 2003. Technical and economic performance analysis of kerosene lamps and alternative approaches to illumination in developing countries. Lawrence Berkeley National Laboratory Report.

Min, Brian Kyung-Hue. 2010. Democracy and light: Public service provision in the developing world. PhD diss., University of California, Los Angeles.

Parcero, O. J., and Elissaios Papyrakis. 2015. Income inequality and the resource curse. *Resource and Energy Economics* 45:159–77.

Pfeffermann, Guy, and Richard Webb. 1983. Poverty and income distribution in Brazil. *Review of Income and Wealth* 29:101–24.

Pinkovskiy, Maxim L. 2013. Economic discontinuities at borders: Evidence from satellite data on lights at night. Unpublished manuscript, MIT.

Pinkovskiy, Maxim, and Xavier Sala-i-Martin. 2009. Parametric estimations of the world distribution of income. Technical report, National Bureau of Economic Research, Cambridge MA.

———. 2016. Lights, camera, . . . income! Illuminating the national accounts-household surveys debate. *Quarterly Journal of Economics* 131:579–631.

Ravallion, Martin. 2003. The debate on globalization, poverty and inequality: Why measurement matters. *International Affairs* 79:739–53.

Robinson, James A., and Ragnar Torvik. 2005. White elephants. *Journal of Public Economics* 89:197–210.

Robinson, James A., Ragnar Torvik, and Thierry Verdier. 2006. Political foundations of the resource curse. *Journal of Development Economics* 79:447–68.

Rose, Amy N., and Eddie A. Bright. 2014. The LandScan Global Population Distribution Project: Current state of the art and prospective innovation. Technical report, Oak Ridge National Laboratory (ORNL).

Ross, Michael L. 2001. Does oil hinder democracy? *World Politics* 53:325–61.

———. 2004. What do we know about natural resources and civil war? *Journal of Peace Research* 41:337–56.

———. 2007. How mineral-rich states can reduce inequality. *Escaping the Resource Curse* 23775:237–55.

Rutstein, Shea O. 2008. The DHS Wealth Index: Approaches for rural and urban areas. DHS Working Paper no. 60. Demographic and Health Surveys, Washington, DC.

Rutstein, Shea Oscar, and Sarah Staveteig. 2014. Making the Demographic and Health Surveys wealth index comparable. Rockville: ICF International.

Sachs, Jeffrey D., and Andrew M. Warner. 1995. Natural resource abundance and economic growth. Technical report, National Bureau of Economic Research, Cambridge, MA.

———. 2001. The curse of natural resources. *European Economic Review* 45:827–38.

Sala-i-Martin, Xavier. 2006. The world distribution of income: Falling poverty and . . . convergence, period. *Quarterly Journal of Economics* 121 (2): 351–97.

Sala-i-Martin, Xavier, and Arvind Subramanian. 2013. Addressing the natural resource curse: An illustration from Nigeria. *Journal of African Economies* 22:570–615.

Small, Christopher, Francesca Pozzi, and Christopher D. Elvidge. 2005. Spatial analysis of global urban extent from DMSP-OLS night lights. *Remote Sensing of Environment* 96:277–91.

Smith, Brock. 2015. The resource curse exorcised: Evidence from a panel of countries. *Journal of Development Economics* 116:57–73.

Stone, Richard, and Organització Europea de Cooperació Econòmica. 1961. *Input-output and national accounts*. Paris: Organisation for European Economic Co-operation.

Stone, Richard, and Giovanna Croft-Murray. 1959. *Social accounting and economic models*. London: Bowes & Bowes.

Subramanian, Shankar, and Angus Deaton. 1996. The demand for food and calories. *Journal of Political Economy* 104 (1): 133–62.

Toews, Gerhard, and Pierre-Louis Vezina. 2016. Resource discoveries and FDI bonanzas. Typescript, University of Oxford, Department of Economics.

Tornell, Aaron, and Philip R. Lane. 1999. The voracity effect. *American Economic Review* 89 (1): 22–46.

Tsui, Kevin K. 2011. More oil, less democracy: Evidence from worldwide crude oil discoveries. *Economic Journal* 121:89–115.

Van Den Bremer, Ton, Frederick van der Ploeg, and Samuel Wills. 2016. The elephant in the ground: Managing oil and sovereign wealth. *European Economic Review* 82:113–31.

Van der Ploeg, Frederick. 2011. Natural resources: Curse or blessing? *Journal of Economic Literature* 49 (2): 366–420.

———. 2012. Bottlenecks in ramping up public investment. *International Tax and Public Finance* 19:509–38.

Van der Ploeg, Frederick, and Steven Poelhekke. 2016. The impact of natural resources: Survey of recent quantitative evidence. *Journal of Development Studies* 53 (2): 1–12.

Van der Ploeg, Frederick, and Dominic Rohner. 2012. War and natural resource exploitation. *European Economic Review* 56:1714–29.

Van der Ploeg, Frederick, and Anthony J. Venables. 2011. Harnessing windfall revenues: Optimal policies for resource-rich developing economies. *Economic Journal* 121:1–30.

———. 2012. Natural resource wealth: The challenge of managing a windfall. *Annual Review of Economics* 4:315–37.

———. 2013. Absorbing a windfall of foreign exchange: Dutch disease dynamics. *Journal of Development Economics* 103:229–43.

Venables, Anthony J., and Samuel E. Wills. 2016. Resource funds: Stabilising, parking, and inter-generational transfer. *Journal of African Economies* 25:ii20–ii40.

Vicente, Pedro C. 2010. Does oil corrupt? Evidence from a natural experiment in West Africa. *Journal of Development Economics* 92:28–38.

Vijayaraj, Veeraraghavan, Anil M. Cheriyadat, Phil Sallee, Brian Colder, Ranga Raju Vatsavai, Eddie A. Bright, and Budhendra L. Bhaduri. 2008. Overhead image statistics. In *Applied imagery pattern recognition workshop*, pp. 1–8, Applied Imagery Pattern Recognition '08. 37th Institute of Electrical and Electronics Engineers.

Wills, Samuel. 2015. Optimal monetary responses to oil discoveries. OxCarre Research Paper 121, Oxford.

World Bank. 2013. Global monitoring report 2013: Rural-urban dynamics and the millennium development goals. Technical report, World Bank Group, Washington, DC.

———. 2015. *World development indicators*. Washington, DC: World Bank (producer and distributor). https://data.worldbank.org/products/wdi.