

DEPARTMENT OF ECONOMICS

OxCarre

Oxford Centre for the Analysis of Resource Rich Economies

Manor Road Building, Manor Road, Oxford OX1 3UQ

Tel: +44(0)1865 281281 Fax: +44(0)1865 271094

oxcarre@economics.ox.ac.uk www.oxcarre.ox.ac.uk



OxCarre Research Paper 164

Left in the Dark? Oil and Rural Poverty

Brock Smith

OxCarre, Department of Economics

&

Samuel Wills

OxCarre, Department of Economics

Left in the Dark? Oil and Rural Poverty

Brock Smith and Samuel Wills*

June 2016

Abstract

Do oil booms reduce poverty and inequality? To study this we propose a new measure of rural poverty: counting people that live in darkness at night. We do this by combining high-resolution satellite data on night-time lights and population globally from 2000-2013. This measure accurately identifies up to 83% of households as above or below the poverty line when compared to over 600,000 surveys. We find that both high oil prices and new discoveries increase illumination and GDP nationally, but promote inequality because the increases are limited to towns and cities with no evidence that they benefit the rural poor.

Keywords: oil, rural poverty, poverty measurement, inequality, night-time lights, urbanization.

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

*Department of Economics and Oxford Center for the Analysis of Resource Rich Economies, University of Oxford, UK. First version October 2015. Samuel Wills would like to thank the Economic and Social Research Council for financial support [grant number ES/K009303/1]. We are grateful for information on the DHS comparative wealth index from Shea Rutstein, and for helpful comments from Tiho Ancey, 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. Websites: www.brockdsmith.com; www.samuelwills.wordpress.com. Corresponding email: samuel.wills@economics.ox.ac.uk.

1 Introduction

The view that oil and other natural resources “curse” the countries that own them is widely held in academic and policy circles. It comes from early cross-country work that focused on aggregate economic activity (Sachs and Warner, 1995; 2001; survey by van der Ploeg, 2011). More recent studies have challenged this view (Brunnschweiler and Bulte, 2008; Alexeev and Conrad, 2009; Smith, 2015; James, 2015). However, there has been 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 is collected infrequently, if at all. When the data is collected it is 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 this 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), “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 study how oil booms affect poverty and inequality we propose a novel measure of rural poverty: the number of people that live in darkness at night. We do this using two detailed and geographically disaggregated datasets, on night-time lights and population (see Figure 1.1). The first records the amount of light emitted at night around the globe at approximately a 1km^2 resolution from 1992-2013. 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 favouritism (Hodler and Raschky, 2014), and infrastructure investment (Jedwab and Moradi, 2015; Jedwab et al., 2015) amongst 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 dataset on global population at a 1km^2 resolution from LandScan, which has received little attention in the economics literature. After aggregating the data to approximately 100km^2 cells, our final sample covers 1.04 million observations each year for 2000-2013. It reveals an important point: that many people live in darkness at night (see for example Figure 1.2). They are the rural poor.

People living in darkness is an admittedly crude proxy for poverty. It only captures rural poverty and tells us nothing about urban poverty rates. It only considers one component of a household’s consumption bundle, light, and ignores the many other unmet needs that characterise 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 Section 3).

However, measuring poverty using darkness also has a number of strengths. First, it is simple and intuitive. Lighting is an important basic need with very high returns, and

people tend to switch from kerosene to electric lighting (so appear in our data) soon after leaving extreme poverty. Second, it uses existing satellites so is cheap and quick to collect compared to household surveys. Third, it produces a globally balanced panel data set at $\sim 1\text{km}^2$ resolution, which 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 DHS household surveys from 36 countries it accurately identifies 74% of households as being above or below the poverty line at 100km^2 resolution, rising to 83% at 1km^2 . The poverty share in unlit areas is 2.6 times higher than lit areas at 100km^2 resolution, or 4.4 times higher at 1km^2 . To our knowledge this is the first comparison of night-time light data to household surveys (see Section 3).

This new measure lets us study how oil booms affect poverty and inequality using two experiments: high oil prices and giant oil discoveries. We do not have a strong prior. Oil booms may reduce poverty and inequality if the rents are shared evenly, even more if they are spent on redistribution and public investment as in Botswana and Malaysia (van der Ploeg, 2011). Alternatively, they may exacerbate it if the rents are captured by economic and political elites as in Nigeria. Our first experiment uses the sharp rise in oil prices after 2003, which we treat as exogenous to oil producers because it “was driven primarily by the cumulative effects of positive global demand shocks” (Kilian, 2009). Our second experiment uses giant oil and gas discoveries, so we can distinguish between price and quantity booms.

Our baseline results show that approximately 30% of people outside the OECD lived in rural darkness in 2000, rising to over 50% in Africa. Since then economic activity around the world has risen and rural poverty has fallen, both because lights have been switched on and people have moved to cities and towns.

We find that both oil price and quantity booms stimulate aggregate economic activity. The decade of high oil prices from 2003 saw night-time illumination grow significantly faster in oil dependent countries than non-dependent ones. By the end of the decade the price boom was responsible for approximately a 34% increase in total illumination, with similar effects for GDP per capita. Furthermore, illumination grew significantly in countries that had made giant oil discoveries, starting with a six year lag (consistent with Arezki et al., 2016). We find that an oil discovery with a net present value worth 100% of GDP increases illumination by 19%, and PPP-adjusted GDP by 8%, after ten years.

The economic growth from oil booms accrues to cities and towns, but is not shared with the rural poor. By the end of the 2000s oil price boom the economic activity in cities in oil dependent countries was 15% higher than their non-dependent counterparts, and in towns was 38% higher. In contrast, the population living in rural darkness 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 non-dependent 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 four to six year lag. Once again, there is no evidence that the population living in rural darkness changed significantly. Rural areas did not become lit, but there may have been a small (1%) reallocation of people from unlit rural areas to towns. Together these results imply that inequality grows during oil booms.¹

¹Our results invite a different interpretation to Aragon and Rud (2013), who find that the benefits




Figure 1.1: i. Night-time lights and ii. population around the world.

To understand the mechanisms behind unlit rural areas becoming lit we estimate a hazard model. 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 dependent and non-dependent countries. Cells in countries that have seen high light growth in the past are more likely to switch on in the present, except in oil dependent countries which suggests they are less effective at converting economic growth to poverty reduction.

Our work contributes to two strands of literature. This is the first paper we are aware of that measures poverty by counting the people who live in darkness at night. Measuring income, poverty and inequality has a long and distinguished history (for example Kuznets 1937, 1941, 1953, 1955; Stone 1959, 1961; Atkinson, 1970; Deaton 1985, 1997). Much of this combines aggregate national accounts data with household survey data 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 services that are not exchanged in a market (Deaton, 2005). Pinkovskiy and Sala-i-Martin (2016) use night-time lights data to reconcile this conflict, finding in favour of aggregate measures. Surveys have the advantage of offering detailed, targeted measures of poverty that take into ac-

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.




Figure 1.2: i. Night-time lights and ii. population for the region surrounding Sokoto, Nigeria.

count consumption bundles and price levels, quality (Deaton, 1988), calorific demands (Subramanian and Deaton, 1996), life expectancy (Pfeffermann and Webb, 1983), within-household distributions (Deaton and Muellbauer, 1986) and a host of other factors. Other studies use the top income share as a proxy for the income distribution (see survey by Atkinson et al.), which is correlated with poverty (Leigh, 2009). In comparison we offer a relatively crude poverty measure, though it covers the entire world, at fine resolution and regular intervals, using free existing data, and is thus well-suited for empirical work. Other recent efforts to measure poverty with large datasets include using night-time lights to spatially distribute World Bank poverty rates (Elvidge et al., 2009; which we adapt in Appendix C), using mobile phone records in Rwanda (Blumenstock et al., 2015), and using machine learning on daytime satellite images in Uganda (trained with night-time lights, Xie et al., 2015).

To our knowledge this paper also offers the first global panel evidence on how oil booms affect poverty. The resource curse literature has focused on how commodities affects aggregate income and growth. Little has been done on poverty, inequality and the distribution of income. Evidence that resource booms increase inequality has been provided by Gylfason and Zoega (2003) who find that national resource dependence is positively correlated with the Gini index, and Goderis and Malone (2011) who use oil price variation in a cross-country panel; while Parcero and Papyrakis (2015) find that inequality is lower in




Figure 2.1: PPP-adjusted real GDP vs Night-time lights (in logs)

resource-rich countries after controlling for under-reported data. Bhattacharyya and Williamson (2013) show that commodity price shocks increase Australian top income shares and note that “the empirical literature on the inequality and resource boom connection is relatively thin”. Van der Ploeg and Poelhekke (2016) survey the empirical literature on the resource curse and find that “much more work on the impact of natural resources on... inequality... is needed”.

The paper proceeds as follows. Section 2 introduces the data. Section 3 introduces and tests a new measure of extreme poverty: the number of people living in darkness at night. Section 4 uses the measure to investigate how oil price and quantity booms affect poverty and inequality. Section 5 concludes.

2 Data

2.1 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 is provided at a pixel resolution of 30×30 arcseconds, or 0.86 square kilometres near the equator, and ranges from 0 to 63. The data is constructed by overlaying all daily images over the course of a year and discarding those that are obfuscated by cloud cover, lightning, aurora, etc. for a given pixel.

The pioneering work of Henderson et al. (2011, 2012) established a strong link between country-level GDP growth and growth in mean light intensity. Doll et al. (2006) and Michalopoulos & Papaioannou (2014) have also performed cross-validation work for GDP levels. While we refer to those papers for a more detailed analysis, in Figure 2.1 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 its high resolution, lights data has been used in several studies for sub-national 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 (described in Section 3.2).

Lights data are subject to a few confounding issues important to this study. First, “top-coding” refers to pixels assigned a max-value of 63, beyond which we cannot distinguish levels of economic activity. This occurs in especially dense or economically active areas, making it difficult to estimate urban poverty rates, but is a far more prevalent problem in developed countries (see Michalopoulos and Papaioannou, 2014). Second, light data includes significant luminosity readings from gas flares, which do not reflect comparable economic activity. This is an important issue for this paper as our treatment group includes many gas-producing countries. To control for this we drop all cells that include gas flare activity according to the provider of the lights data (the Earth Observation Group). Third, light data is affected by overglow (or “blooming”), 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) 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 10km from the light’s origin. As we compute national-level estimates, use $\sim 100\text{km}^2$ cells, and focus on people living in areas far from cities, it is unlikely that these sources of error will be sufficiently large or correlated with our variables of interest that they 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. To control for this we use time fixed effects.

2.2 Population Data

The Oak Ridge National Laboratory produces a data set called LandScan covering each year from 2000-2013, which provides spatial population counts at a 30×30 arc second resolution.² This is similar to NASA’s Socioeconomic Data and Applications Center (SEDAC), which also measures population at a 30×30 arc second resolution and has been used by Dell (2010) and Alesina et al. (2015) amongst others. However, the SEDAC data only covers the years 1990, 1995 and 2000. LandScan provides estimates of “ambient” population, which is a count for a given area over a 24-hour average, rather than just where people sleep. The data are generated by distributing known national and sub-national population counts throughout the grid according to a likelihood model that uses inputs including land cover data, roads data, and high resolution satellite imagery, among other sources.³

2.3 Household Survey Data

The Demographic and Health Surveys (DHS) Program collects individual- and household-level data on several questions about living conditions, possessions, health, etc., in over 90 countries. We use the most recent standard DHS surveys from 36 countries, covering

²This is the same resolution as the lights data, although the pixels are not aligned before 2010. The grid cells described in Section 3.2 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.

³For further detail http://web.ornl.gov/sci/landscan/landscan_documentation.shtml

636,448 households from 2003-2013. 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 2km in urban areas, and 5km in rural areas with 1% of rural clusters displaced by 10km (though all remain within their country and region, www.dhsprogram.com/faq.cfm). Following DHS recommendations we manage this by basing our analysis on $\sim 100\text{km}^2$ cells.

2.4 Urban and Rural Classifications

SEDAC also provides an “Urban Extents Grid”, which uses 1995 population count estimates to classify each square of a 30x30 arc second global grid as either urban or non-urban. The classification is based on contiguous lighted squares (as of 1995) and squares known to hold at least 5000 people.

2.5 Oil Prices and Discoveries

We study two types of oil shocks: to prices and to the quantity of ultimately recoverable reserves. We use annual Brent crude prices from 1990-2013, 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 2006-2010. 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.

The oil quantity shock uses data on giant oil and gas discoveries from the American Association of Petroleum Geologists. This is an updated version of the dataset by Horn (2003, 2004), which builds on that of Halbouty et al. (1970). These giant discoveries have also been used as a well-identified income shock by Lei and Michaels (2014), Arezki et al. (2016), and Smith (2015) amongst others. The data records 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 covers 1019 discoveries, 245 of which occur between 1990-2013. We take discoveries as far back as 1990 to study their effect on lights up to 10 years after the discovery (as population data begins in 2000). Dropping the OECD and aggregating multiple discoveries for particular countries 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 Figure 2.2).

We account for the size of each discovery by constructing a measure of its Net Present Value (NPV) divided by GDP (following Arezki et al., 2016),

$$NPV_{i,t} = \frac{\sum_{j=5}^J q_{i,t+j} \text{oilprice}_t (1 + r_i)^j}{GDP_{i,t}} \times 100$$




Figure 2.2: Count of giant oil discoveries by region included in our sample (excluding OECD countries).

where the NPV of country i at the time of discovery, t is the discounted sum of total revenue based on an approximate production profile, $q_{i,t+j}$, from the fifth year after discovery to the exhaustion date, J , valued at the oil price at the time of discovery. Revenue is discounted using country-specific, risk-adjusted discount rates to account for differences in political risk. This assumes a riskless rate of 5 per cent, 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).

3 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 4 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 C.

3.1 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) on the lowest rungs, more efficient transitional fuels like kerosene in the middle and clean but expensive sources such as electricity at the top (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.




Figure 3.1: Share of rural and urban households using electricity as main source of lighting (Bacon et al., 2010; own calculations).

Areas appear in night-time light data after their households switch from kerosene to electric lighting. A typical kerosene lamp delivers from 1-6 lux of useful light; a standard 60W 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 dataset is highly correlated with reported electrification rates. The DHS surveys show 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 urbanisation. Lee (2013) studies energy use in Uganda using a cross-section of 6775 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 (PWT 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 (Figure 3.1). 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.

3.2 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 100km^2 near the equator, though cells are divided at national borders. We drop cells with zero population. Aggregating pixels in this way speeds up the analysis with little loss of accuracy for our purposes, and has become standard practice in the literature.

Second, we sum the population living in cells with zero light readings 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. Of course, there are cells with a non-poor population at a sufficiently low density to record a light reading of zero. 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 C.

Third, we classify lit cells as urban or rural, according to the SEDAC definition in 2000 (see Section 2.4). Each cell will contain a mix of urban and non-urban pixels, so we classify it as urban if at least 33% of its pixels are. While necessarily arbitrary, this definition yields a global urban population very close to that estimated by the World Bank in 2000. Cities are defined as urban cells, towns as non-urban cells with lights, and unlit rural areas as non-urban cells without lights but with non-zero population.

The CWI measures wealth at a household level. Each DHS survey includes a household-level wealth index which is normalised 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 (see Rutstein, 2008). To make the index comparable across surveys we calculate a CWI for each household⁴. This maps the wealth index from each survey to a common baseline by comparing common “anchor points” at the bottom of the wealth distribution (the number of unmet basic needs), and at the top (ownership of a television, refrigerator, car/truck, and fixed telephone). The threshold for extreme poverty is a $CWI < -0.27$, which is approximately the median for households with two or more unmet basic needs.

The UBN directly measures whether households lack four basic needs: 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.⁵ These UBNs are also part of the CWI. 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 $UBN \geq 2$ 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.

3.3 Evaluation

3.3.1 National level

The unlit rural percentage is strongly positively correlated with existing poverty measures at a national level. Figure 3.2 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

⁴This uses survey-level coefficients generously provided by Rutstein and Staveteig (2014).

⁵For 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.

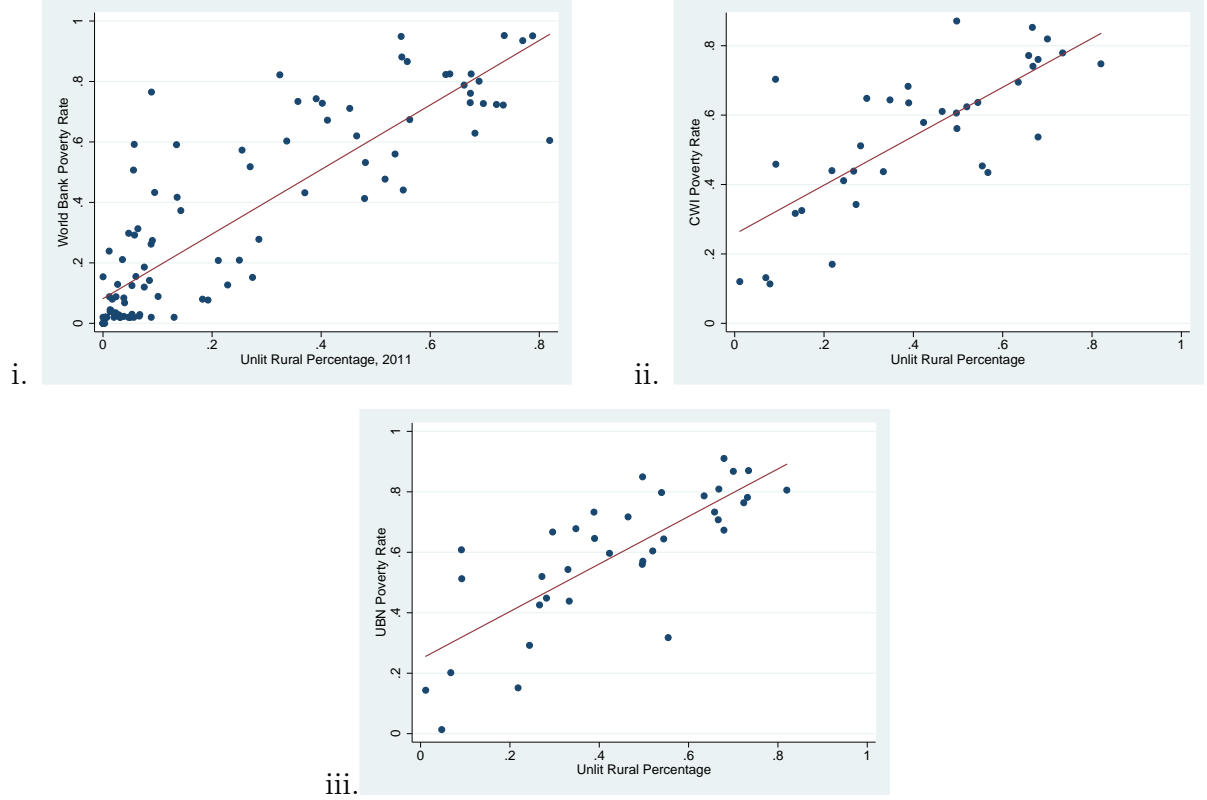


Figure 3.2: Unlit rural percentage vs three poverty measures: i. Less than \$2 per day percentage (World Bank, 2011), ii. CWI extreme poverty rate, iii. UBN=2 poverty rate.

recent estimates of people living on less than \$2 per day.⁶ The CWI rate is based on the -0.27 threshold, and the UBN rate based on having two or more unmet basic needs. While there are some countries for which the unlit percentage is a poor predictor,⁷ 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).

3.3.2 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 Figure 3.3). This rises to 83% at 1×1 pixel resolution⁸, or 66% using the $UBN \geq 2$ definition at 12×12 (see Table B.2). 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.

⁶These rates are not all from 2011-they are taken in various years. We use only World Bank estimates that have been made since 2005.

⁷Figure 3.2i. shows a cluster of countries in the top left of the graph - i.e. 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 Appendix C performs much better for these particular countries.

⁸The 1×1 figures are based on 25 surveys from 2010-2013 because of a grid alignment issue before then.




Figure 3.3: Distribution of wealth in i. unlit and ii. lit cells at 12×12 pixel resolution (CWI extreme poverty threshold in red).




Figure 3.4: The effect of varying the CWI poverty threshold on accuracy at 12×12 pixel resolution.

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 (non-poor living in unlit cells) is 5%, and the Type II error (poor living in lit cells) is 21%. Most of the error therefore comes from darkness not identifying the urban poor, which we acknowledge is outside the scope of this measure.

Tightening the definition of poverty does not greatly alter the measure's accuracy. 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 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 (see Figure 3.4).

3.3.3 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 B.1). 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, urbanisation, population density and electrification is insignificant at the 10% level. This suggests that there is not a major systematic bias in the measure.

4 The effect of oil booms on poverty and inequality

To understand how oil booms affect rural poverty we study both increases in prices, using the high oil prices of the 2000s, and increases in 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 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 inequality as it accrues to cities and towns, but not the rural poor. High oil prices do not affect migration, but oil discoveries cause a small movement from unlit rural areas to towns (though not cities).

4.1 Identification: Oil Prices and Discoveries

We conduct two experiments to understand how oil booms affect poverty and inequality. These use exogenous shocks from oil prices, and from giant oil discoveries.

The first experiment exploits the period of high prices between 2003 and 2013, and its differential effect on oil dependent and non-dependent 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) using a structural decomposition in a VAR framework. We control for global demand shocks affecting our treatment countries through channels other than oil by using non-dependent 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 the possibility that the global demand shock affected non-oil demand in oil dependent and non-dependent countries differently we employ a second experiment.

The second experiment exploits the discovery of giant oil fields, which we treat as exogenous after controlling for time and country fixed-effects. Giant oil discoveries are a type of quasi-natural experiment, occurring in only 2% of all wells drilled. Countries only have limited ability to affect this. Toews and Vezina (2016) find that a country 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 metres 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 control group is countries that don't make a giant discovery, and those that do before the discovery is made.

4.2 Estimating Equations

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

$$Y_{i,t} = \alpha + \sum_{s=2000}^{2013} \beta_{1s}(\lambda_s \times D_i) + \sum_{s=2000}^{2013} \beta_{2s}(\lambda_s \times Y_{i,2000}) + \Lambda_t + \Phi_i + region_i * t + \epsilon_{i,t} \quad (4.1)$$

where Y_{it} is the outcome of interest for country i in year $t = 2000, \dots, 2013$, 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, Λ_t is year fixed effects, Φ_i is country fixed effects, and $region_i * t$ is regional linear trends.⁹ Each coefficient β_{1t} then measures the average conditional difference in Y between dependent and non-dependent countries in year t relative to the difference in the reference year 2002, which is the year before oil prices began to rise.

We estimate the effect of giant oil and gas field discoveries using a distributed lag model:

$$Y_{i,t} = \alpha + \sum_{j=0}^{23} \beta_{1j} Size_j + \sum_{s=2000}^{2013} \beta_{2s}(\lambda_s \times Y_{i,2000}) + \Lambda_t + \Phi_i + region_i * t + \epsilon_{i,t} \quad (4.2)$$

Where $Size_j$ is the 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 analyse multiple discoveries within countries. We focus on the ten years after discovery due to data limitations, so use all discoveries since 1990 (since the population data begins in 2000). However, we include lags up to 23 years after discovery to ensure that our counterfactual is limited to countries that don't 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).

⁹Regional classifications are adapted 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.




Figure 4.1: Unconditional trends in the i. unlit rural percentage and ii. switch-on percentage.

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,

$$I_{cit}^{RS} = \alpha + \beta_1 X_{cit} + \beta_2 D_i X_{cit} + \Lambda_t + \Phi_i + \epsilon_{c,i,t} \quad (4.3)$$

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 <100km 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. The coefficients in β_2 therefore estimate the differential effect of each variable on dependent relative to non-dependent countries. The sample is restricted to cells that are unlit but inhabited as of 2000.

Since we are interested in effects on 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%¹⁰ in 2000 from all specifications. We also drop three countries that experienced large-scale wars during the sample period, including two that would be treatment countries: Iraq, Syria and Afghanistan. This leaves a sample of 105 countries in the main specification.

4.3 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.¹¹ Global poverty, or the unlit rural percentage, is generally falling which is consistent with broader evidence (Figure 4.1i., see also Pinkovskiy and Sala-i-Martin, 2009; Chen and Ravallion, 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

¹⁰Results are robust to a threshold of 10%.

¹¹Unconditional illumination data does 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 4.2: Breakdown of population share by area classification as of 2000.

measured using the “rural switch percentage”, which records whether unlit, populated areas in 2000 subsequently became lit. Figure 4.1ii. 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 4.2, which shows the general trend for people to leave unlit rural areas for towns and cities.

4.4 Price booms: the 2000’s

4.4.1 High oil prices stimulate growth in oil-dependent countries

Aggregate illumination grew in oil dependent relative to non-dependent countries from 2003-2013. This is illustrated in Figure 4.3, which shows the log difference in aggregate night-time lights between non-OECD oil dependent and non-dependent countries, relative to the omitted year of 2002 (coefficient β_{1t} in equation 4.1). The difference increased steadily during the oil boom, and by the end of the sample the effect on lights in dependent countries is 34% (0.29 log points). The effect on lights per capita is slightly smaller but still statistically significant at a 10% level. There was a similar and strongly significant effect on PPP-adjusted real GDP per capita (expenditure based). These results are in sharp contrast to the stability of these variables during the period of low and relatively stable oil prices in the 1990s (see Figure E.1). Although the effects on lights and GDP begin rising in the late 1990s, this is plausibly due to a mini-boom in which the Brent Crude Index rose from approximately \$13 to \$28 from 1998-2000.

4.4.2 Growth from high oil prices is not shared with the rural poor

Aggregate growth increased inequality by being confined to cities and towns, with no evidence that it benefited the rural poor. Economic activity in both the cities and towns of oil dependent countries grew faster than in their non-dependent counterparts from 2003-2013, with an effect of 15% and 38% (0.14 and 0.32 log points) respectively (Figure 4.4i.-ii.). These estimates are significant at a 10% level, and at a 5% level when using lights per capita (not shown).

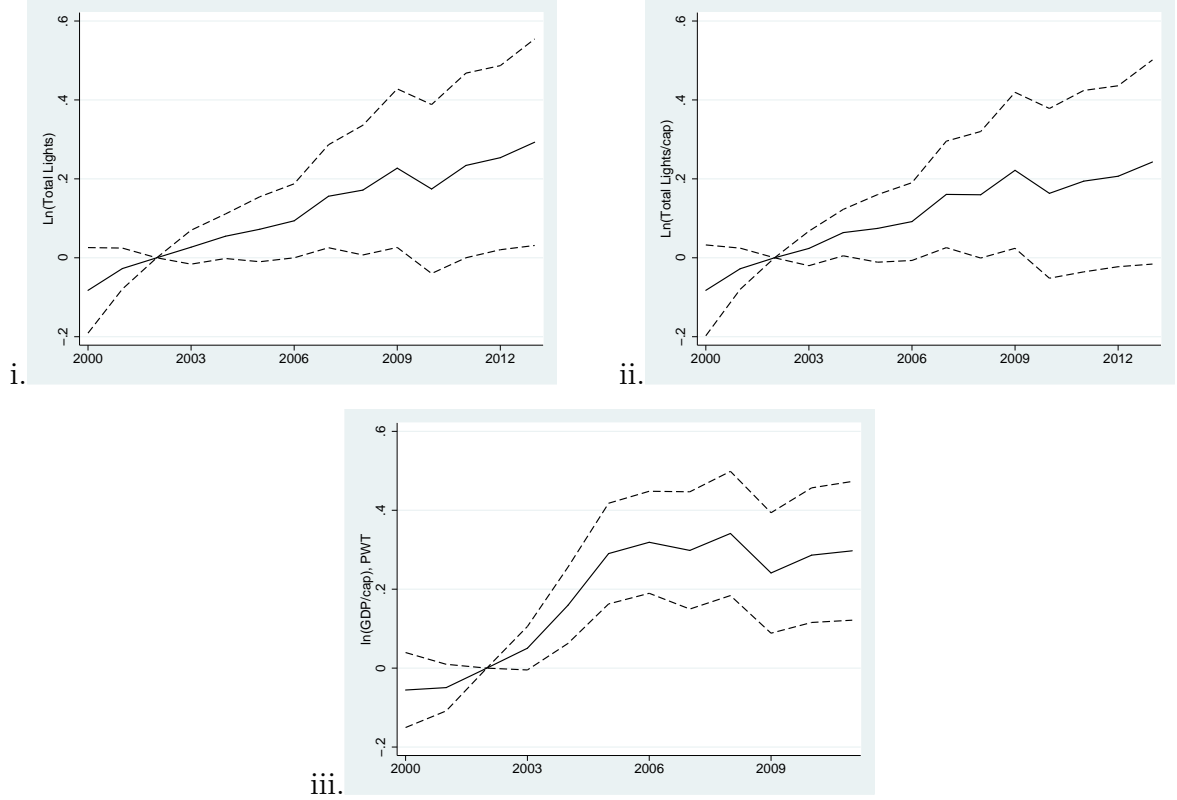


Figure 4.3: Effect of the 2000s oil price boom on i. aggregate lights (estimates and 95% confidence bands), ii. aggregate lights per capita and iii. PPP-adjusted real GDP per capita.

There is no evidence that the oil boom benefited the rural poor. Figure 4.4iii. shows the unlit rural percentage in oil-dependent countries, relative to non-dependent countries in the 2002 base year. There is no significant change or trend. The unlit rural percentage could fall by lights turning on in unlit cells or by population moving from unlit to lit cells, which Figure 4.4iii. might mask. The rural switch percentage in Figure 4.4iv. 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.

4.4.3 High oil prices do not affect migration

There is no evidence that high oil prices induced migration 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 Figure 4.2), this has happened at the same pace in both oil dependent and non-dependent countries. Figure 4.5 illustrates the difference in the population share in cities, towns and rural areas between oil dependent and non-dependent countries¹². In all three instances this difference did not significantly change. The same result is found using data on urbanization rates from the World Bank (not shown). This is in contrast to some existing work arguing that natural resource rents drive urbanization.¹³

¹²Our population share regressions do not control for initial levels, because they are not theoretically subject to convergence like output.

¹³See Cavalcanti et al. (2014) and Gollin et al. (2016). The latter finds that non-OECD countries with oil have higher levels of urbanization than those without, which we confirm in Figure 4.2. However after explicitly identifying exogenous oil shocks we find no evidence that they hasten the pace of urbanization.

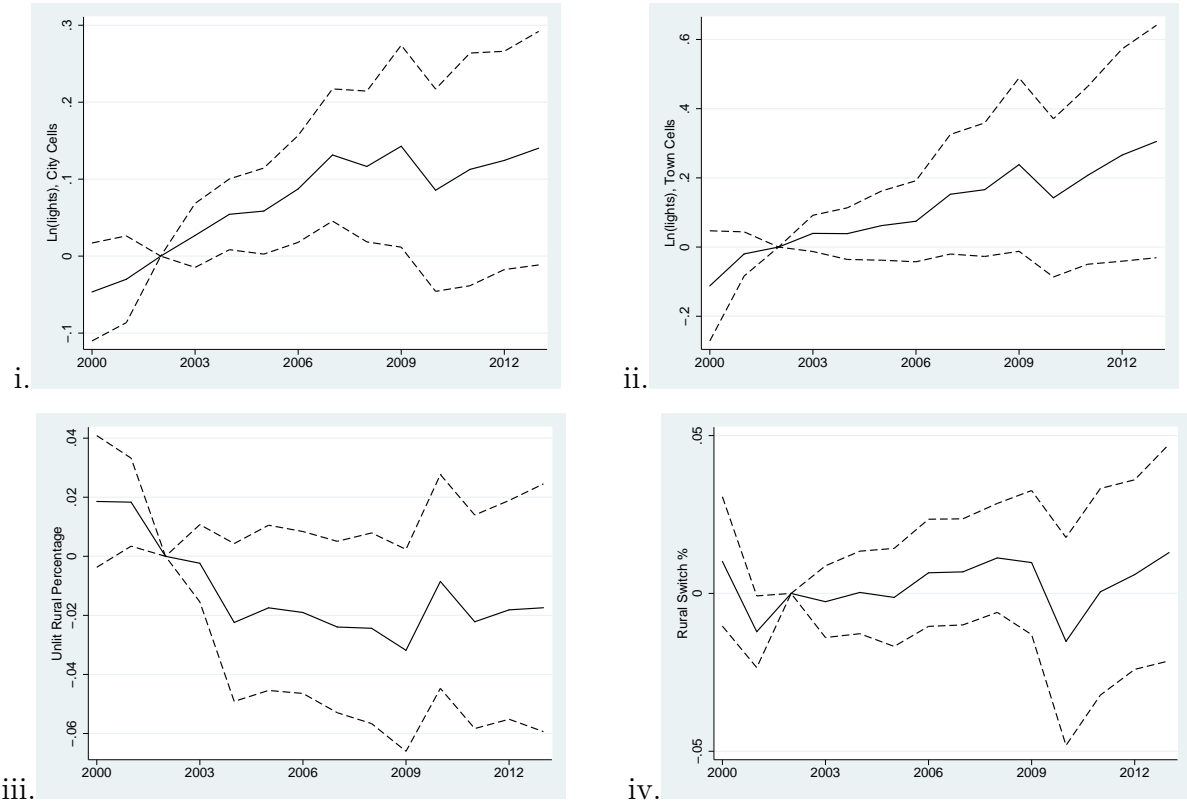


Figure 4.4: Effects of the 2000s oil price boom on i. illumination in cities, ii. illumination in towns, iii. the unlit rural percentage, iv. the share of unlit rural areas that become lit (classification in 2000).

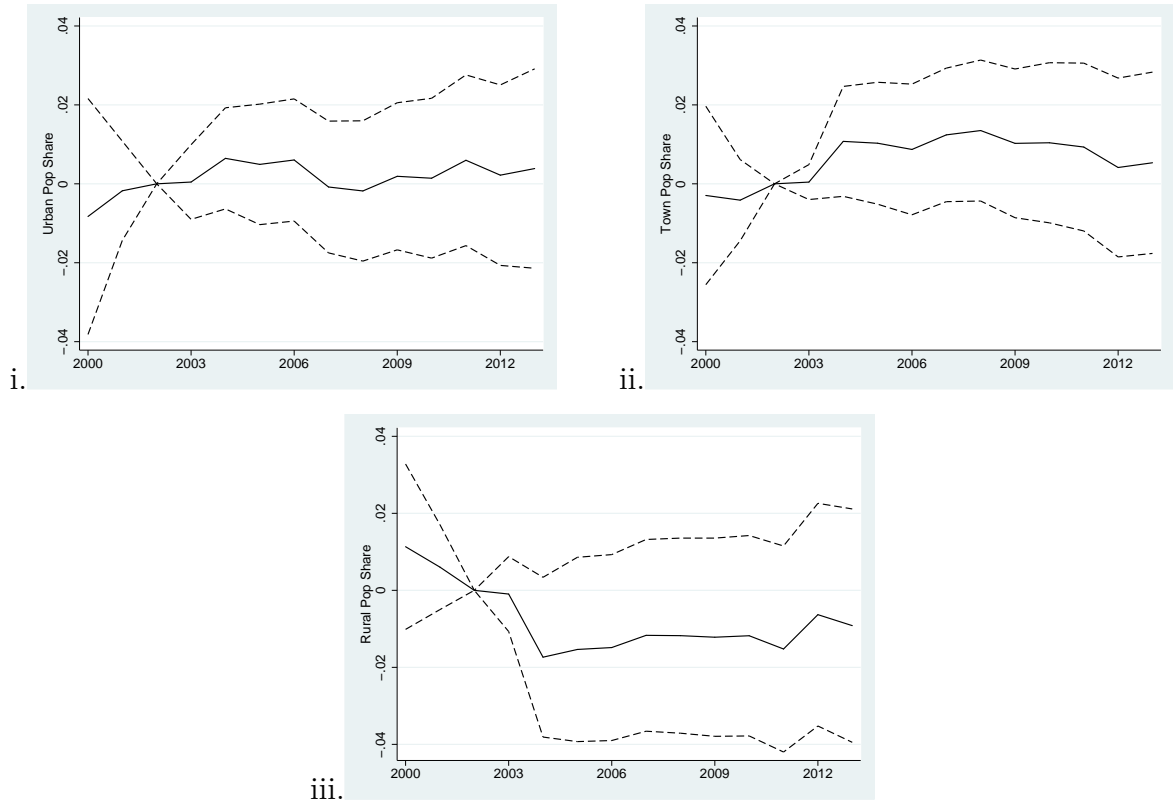


Figure 4.5: Effect of the 2000s price boom on population shares in i. cities, ii. towns and iii. rural areas (classification in 2000).

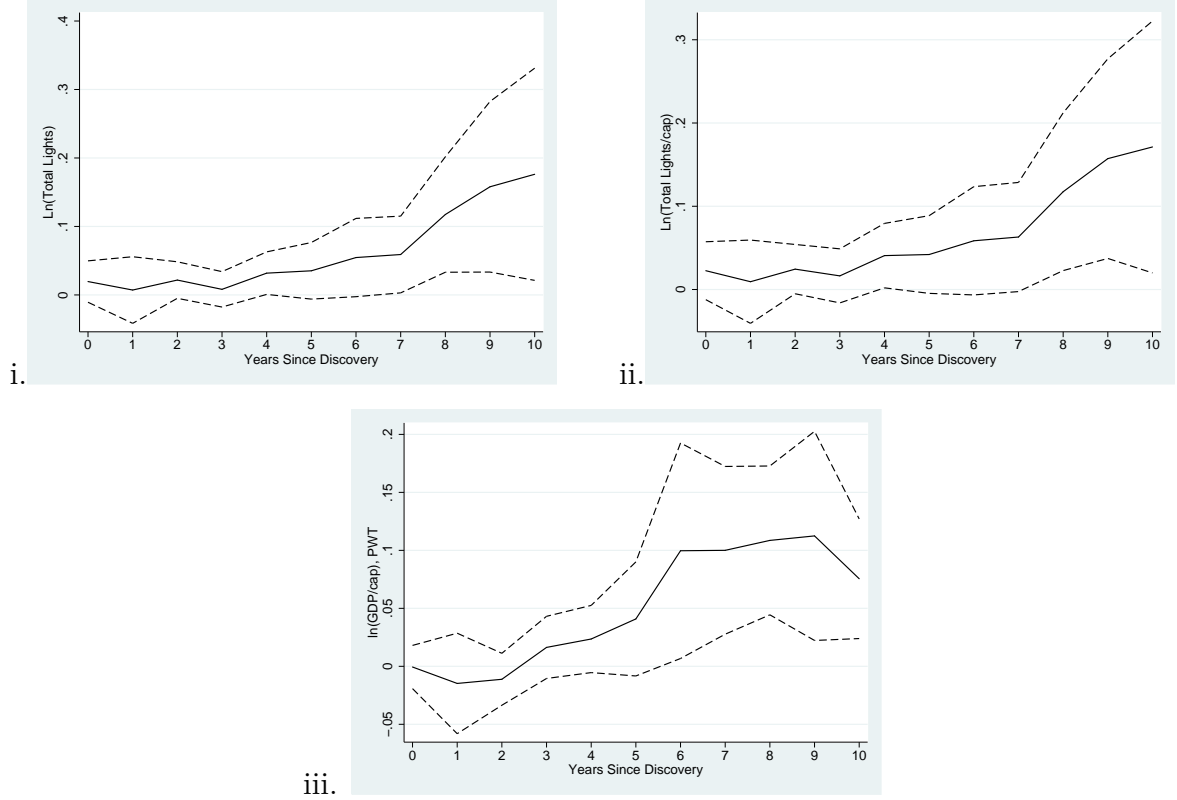


Figure 4.6: Effect of oil discoveries on i. aggregate (log) lights, ii. log aggregate lights per capita and iii. PPP-adjusted real GDP per capita.

4.5 Quantity booms: giant oil discoveries

4.5.1 Giant oil discoveries stimulate economic growth at a six year lag

Aggregate illumination increases about six years after countries discover 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 4.6, which shows the effect on night-time lights for countries that discovered oil t years ago, scaled for the size of the discovery relative to GDP. After nine years, discovering oil with a net present value of 100% of GDP increases lights and lights per capita by 19% (0.17 log points), and GDP per capita by 8% (0.08 log points). It is consistent with oil discoveries being a form of news shock, because there is a delay between discovery and production (see Arezki et al., 2016).

4.5.2 Growth from oil discoveries is not shared with the rural poor

Oil discoveries increase inequality because aggregate national growth is confined to cities and towns but does not benefit the rural poor. 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 is significant at the 5% level after 4 years in cities and 6 years in towns, consistent with a delay between discovery and production.

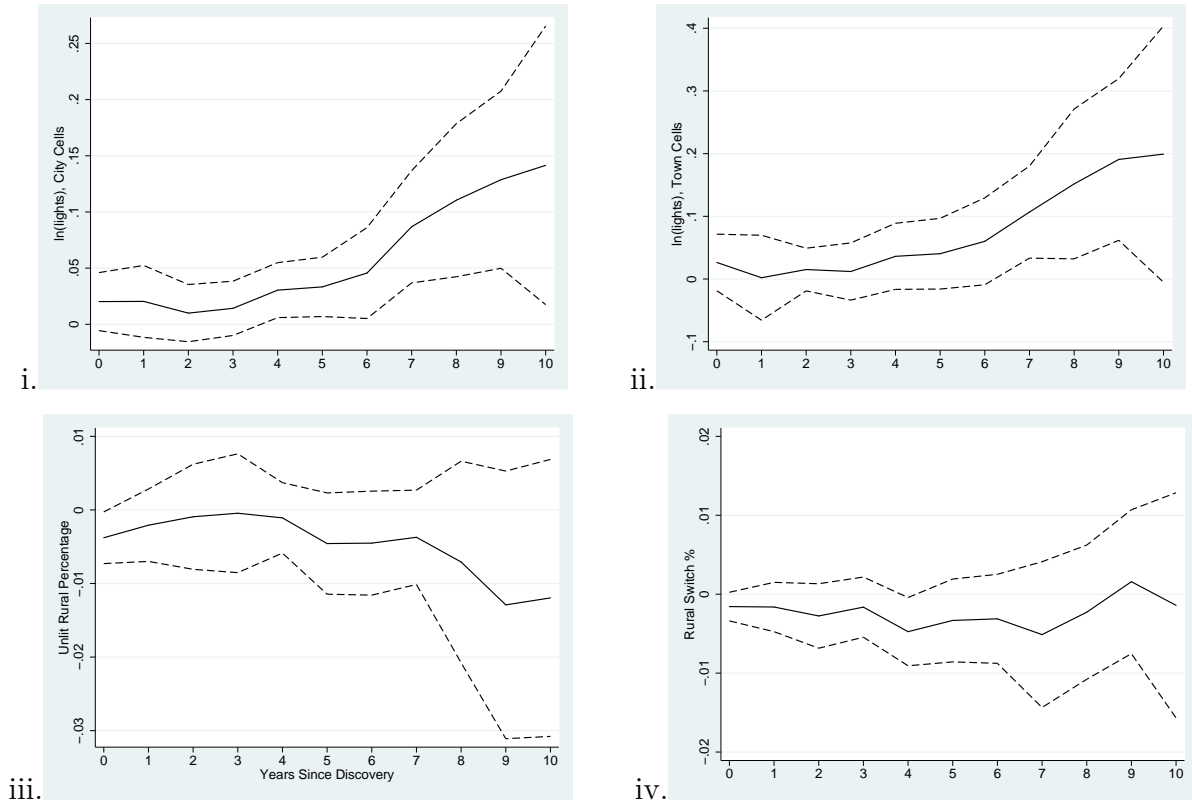


Figure 4.7: Effects of giant oil discoveries on i. illumination in cities, ii. illumination in towns, iii. the unlit rural percentage, iv. the share of unlit rural areas that become lit(classification in 2000).

There is no evidence that giant oil discoveries benefit the rural poor. Up to ten years after a discovery the unlit rural percentage only falls by 1% (0.01 log points) relative to the control group, and this is not significantly different from zero (Figure 4.7iii.). There is no evidence that previously unlit rural areas became lit due to a discovery; if anything, unlit areas in the control group illuminate faster (4.7iv.). 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.

4.5.3 Oil discoveries cause some rural poor to move to towns but not cities

There is some evidence that people may leave unlit rural areas for towns after an oil discovery. Figure 4.8 shows that the rural population share trends downwards after an oil discovery (relative to the control group), and is approximately 1% (0.01 log points) lower after ten years. Towns see 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 4.7iii. There is no discernible effect on the population of cities, which we confirm using World Bank data on urbanisation rates (not shown).

4.6 Mechanisms

The previous two sections analysed how oil price and quantity booms affect unlit rural areas in aggregate. We have seen that rural areas are illuminating around the world, but

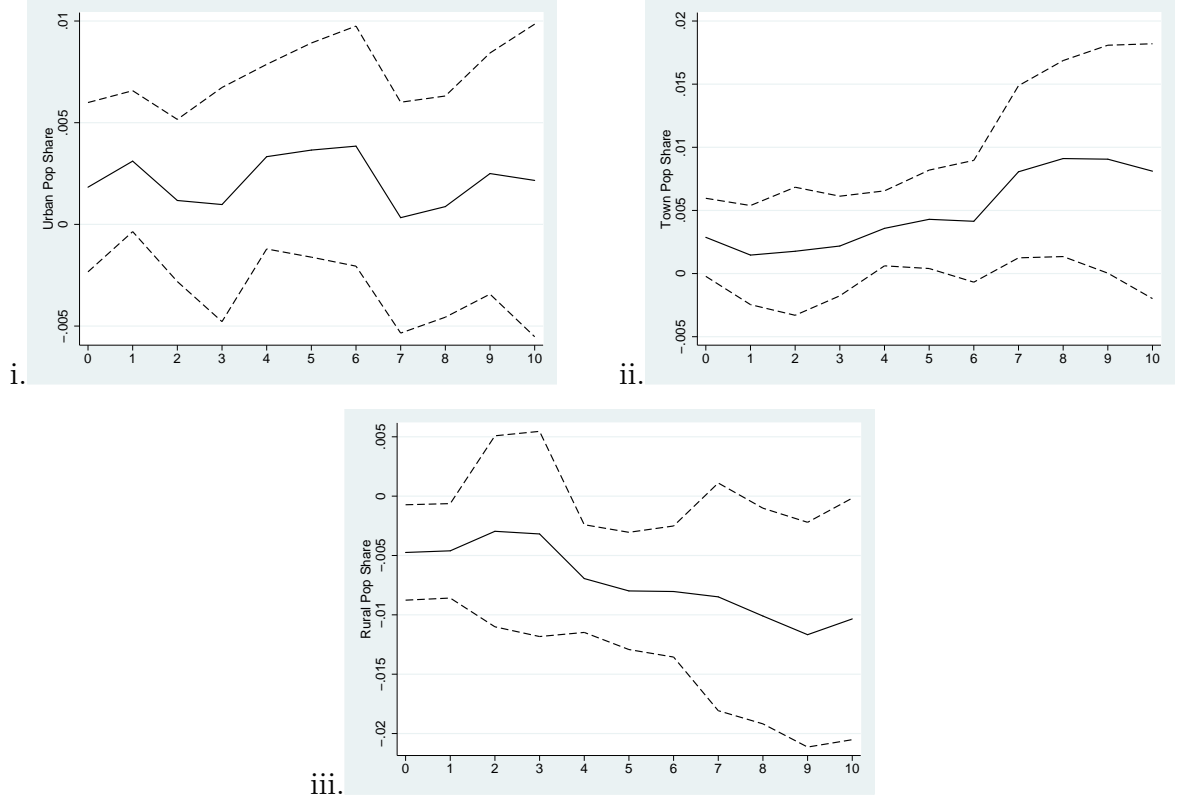


Figure 4.8: Effect of giant oil discoveries on the population share in cities, towns and rural areas (classification in 2000).

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 4.3 to understand what causes an unlit and inhabited rural cell to light up. We find that the probability of illuminating 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 4.1. 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 countries are less efficient at converting growth to poverty reduction.

Being next to a lit cell suggests being near an existing electricity network. We hypothesise 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 2003-2013 adjacent cells became ~15 percentage points more illuminated than non-adjacent cells (see Figure D.1A., which shows the cumulative effect of this mechanism not controlling for the other three in equation 4.3¹⁴). However, adjacency did not increase the probability of illumination in oil dependent relative to non-dependent countries, suggesting that oil revenues were not systematically invested in larger electricity networks.

¹⁴Graphs i. and ii. of this figure 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 non-dependent countries.

Hazard Rate of Unlit Rural Areas Switching On

	(1)
	switchon
Adjacent	0.031*** (0.005)
Adjacent*dep	0.000 (0.005)
<100km from capital	0.006** (0.002)
(<100km from capital)*dep	0.004 (0.004)
Pop. Density	0.004* (0.002)
Pop. Density*dep	-0.002 (0.002)
Lights growth since 2000	0.019*** (0.005)
(Lights growth since 2000)*dep	-0.007* (0.003)
N	7479895
R^2	0.026

Notes: 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 parenthesis. +, *, **, *** represent significance at 10%, 5%, and 1%, respectively.

Table 4.1: Results for a model of the hazard rate of unlit rural areas switching on in a given year.

Being within 100km 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 2003-2013 (Figure D.1B.). Again, being near to the capital did not increase the chance of illumination in oil dependent vs. non-dependent 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 organised. 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 (Figure D.2). As with the previous mechanisms there was no evidence of oil-dependent countries investing more in areas with high population density. This suggests that governments were not using oil revenues to “pork barrel” high-population areas.

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.

4.7 Robustness

To test the robustness of our main results we try four alternative specifications, with the results in Appendix E. The first three focus on the price boom specification in equation 4.1, the fourth on oil discoveries in equation 4.2.

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 (Figure E.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 4.1 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 E.3ii..

The third replaces the binary indicator for dependence, D_i , with a continuous variable, $Rents_i$. In the main specification we use an dependence indicator because it does not rely on a functional form and yields simple and 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 2000-2012 (implying a linear relationship). We find that if oil rents account for 100% of GDP, then the five times increase in prices over our sample would have increased lights by 280% (1.34 log points), or 150% (0.92 log points) excluding Equatorial Guinea (a high-growth outlier where rents account for 75% of GDP, see Figure E.4). The unlit rural percentage drops by 14% (0.13 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 E.5 shows that giant oil discoveries have much larger effect on national lights than gas discoveries within ten 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).

5 Conclusion

This paper attempts to answer the question, do oil booms reduce inequality and rural poverty? To do this we propose a novel way to measure rural poverty: counting the people who live in darkness at night, and then conduct two experiments: exploiting exogenous variation in oil prices and quantities. We find that oil booms stimulate aggregate economic activity, however this promotes 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 datasets: on night-time illumination as a proxy for economic activity, and on population. Evaluating the effectiveness of our measure against 636,448 DHS household surveys from 36 countries we find that it accurately identifies 74% of people as being above or below the threshold for extreme poverty at 100km² resolution, and more at higher resolutions.

We then conduct two experiments using the rise in oil prices from 2003-2013, 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 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 (Mehlum et al., 2006). If this is true then the curse should be associated with inequality and poverty. We find that high oil prices and new discoveries stimulate economic activity (proxied by lights) in countries with oil 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, or people to leave unlit areas for towns and cities. 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. Taken together these results suggest that oil booms increase inequality.

This work lends itself to a wide range of extensions. Darkness is well suited to studying many other determinants of rural poverty, and policies designed to alleviate it, as it provides geographically disaggregated and globally balanced panel data. Practically, darkness may also be a useful way to direct aid and humanitarian interventions.

References

- [1] Alesina, A. F., S. Michalopoulos, and E. Papaioannou (2015). Ethnic inequality. *Journal of Political Economy* (Forthcoming).
- [2] Alexeev, M. and R. Conrad (2009). The elusive curse of oil. *The Review of Economics and Statistics* 91(3), 586–598.
- [3] Aragón, F. M. and J. P. Rud (2013). Natural resources and local communities: evidence from a peruvian gold mine. *American Economic Journal: Economic Policy* 5(2), 1–25.
- [4] Arezki, R., V. A. Ramey, and L. Sheng (2016). News shocks in open economies: Evidence from giant oil discoveries. *Quarterly Journal of Economics* forthcoming.
- [5] Atkinson, A. B. (1970). On the measurement of inequality. *Journal of Economic Theory* 2(3), 244–263.
- [6] Atkinson, A. B., T. Piketty, and E. Saez (2009). Top incomes in the long run of history. Technical report, National Bureau of Economic Research.
- [7] Bacon, R., S. Bhattacharya, and M. Kojima (2010). Expenditure of low-income households on energy: Evidence from africa and asia. *World Bank Extractive Industries for Development Series* 16.
- [8] Baunsgaard, T., M. Poplawski-Ribeiro, M. Villafuerte, and C. Richmond (2012). Fiscal frameworks for resource rich developing countries. Technical report, International Monetary Fund.
- [9] Bhattacharyya, S. and J. G. Williamson (2013). Distributional impact of commodity price shocks: Australia over a century.
- [10] Blumenstock, J., G. Cadamuro, and R. On (2015). Predicting poverty and wealth from mobile phone metadata. *Science* 350(6264), 1073–1076.
- [11] Brunnschweiler, C. N. and E. H. Bulte (2008). The resource curse revisited and revised: A tale of paradoxes and red herrings. *Journal of environmental economics and management* 55(3), 248–264.

- [12] Cavalcanti, T., D. Da Mata, and F. Toscani (2014). Winning the oil lottery: The impact of natural resource extraction on growth. *Mimeo, Department of Economics, University of Cambridge*.
- [13] Chen, S. and M. Ravallion (2010). The developing world is poorer than we thought, but no less successful in the fight against poverty. *The Quarterly Journal of Economics* 125(4), 1577–1625.
- [14] Cust, J. and S. Poelhekke (2015). The local economic impacts of natural resource extraction. *Annu. Rev. Resour. Econ.* 7(1), 251–268.
- [15] Deaton, A. (1985). *The measurement of welfare: theory and practical guidelines*. World Bank, Development Research Department.
- [16] Deaton, A. (1988). Quality, quantity, and spatial variation of price. *The American Economic Review*, 418–430.
- [17] Deaton, A. (1997). *The analysis of household surveys: a microeconomic approach to development policy*. World Bank Publications.
- [18] Deaton, A. (2005). Measuring poverty in a growing world (or measuring growth in a poor world). *Review of Economics and Statistics* 87(1), 1–19.
- [19] Deaton, A. S. and J. Muellbauer (1986). On measuring child costs: With applications to poor countries. *The Journal of Political Economy*, 720–744.
- [20] Dell, M. (2010). The persistent effects of peru’s mining mita. *Econometrica* 78(6), 1863–1903.
- [21] Doll, C. N. (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*.
- [22] Doll, C. N., J.-P. Muller, and J. G. Morley (2006). Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics* 57(1), 75–92.
- [23] Elvidge, C. D., K. E. Baugh, S. Anderson, P. Sutton, and T. Ghosh (2012). The night light development index (nldi): a spatially explicit measure of human development from satellite data. *Social Geography* 7(1), 23–35.
- [24] Elvidge, C. D., K. E. Baugh, P. C. Sutton, B. Bhaduri, B. T. Tuttle, T. Ghosh, D. Ziskin, and E. H. Erwin (2011). Who’s in the dark. satellite based estimates of electrification rates. *Urban Remote Sensing: Monitoring, Synthesis and Modeling in the Urban Environment*, 211–224.
- [25] Elvidge, C. D., P. C. Sutton, T. Ghosh, B. T. Tuttle, K. E. Baugh, B. Bhaduri, and E. Bright (2009). A global poverty map derived from satellite data. *Computers & Geosciences* 35(8), 1652–1660.
- [26] Goderis, B. and S. W. Malone (2011). Natural resource booms and inequality: Theory and evidence. *The Scandinavian Journal of Economics* 113(2), 388–417.
- [27] Gollin, D., R. Jedwab, and D. Vollrath (2016). Urbanization with and without industrialization. *Journal of Economic Growth* forthcoming.

- [28] Gylfason, T. and G. Zoega (2003). Inequality and economic growth: Do natural resources matter?
- [29] Halbouty, M. T., A. A. Meyerhoff, R. E. King, R. H. Dott Sr, H. D. Klemme, and T. Shabad (1970). World's giant oil and gas fields, geologic factors affecting their formation, and basin classification: Part i: Giant oil and gas fields.
- [30] Henderson, J. V., A. Storeygard, and D. N. Weil (2011). A bright idea for measuring economic growth. *The American Economic Review* 101(3), 194.
- [31] Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring economic growth from outer space. *The American Economic Review*, 994–1028.
- [32] Hodler, R. and P. A. Raschky (2014). Regional favoritism*. *The Quarterly Journal of Economics*, qju004.
- [33] Horn, M. K. (2003). Giant fields 1868-2003 (cd-rom). In *Halbouty, M.K., ed., Giant Oil and Gas Fields of the Decade, 1990-1999*. Houston: AAPG Memoir 78, 2003.
- [34] Horn, M. K. (2004). Giant fields 1868-2004 (cd-rom - revision to 2003 version). In *Halbouty, M.K., ed., Giant Oil and Gas Fields of the Decade, 1990-1999*. Houston: AAPG/ Datapages Miscellaneous Data Series. Version 1.2,2004.
- [35] ICF (2012). Demographic and health survey: Sampling and household listing manual. Technical report, MEASURE DHS/ICF International.
- [36] James, A. (2015). The resource curse: A statistical mirage? *Journal of Development Economics* 114, 55–63.
- [37] Jedwab, R., E. Kerby, A. Moradi, et al. (2015). History, path dependence and development: Evidence from colonial railroads, settlers and cities in kenya. *The Economic Journal forthcoming*.
- [38] Jedwab, R. and A. Moradi (2015). The permanent effects of transportation revolutions in poor countries: Evidence from africa. *Review of Economics and Statistics*.
- [39] Kavanagh, M. (2013). Abandoned oil discoveries enjoy revival. *The Financial Times April 8*.
- [40] Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053–69.
- [41] Kuznets, S. (1937). Appendices to "national income and capital formation, 1919-1935". In *National Income and Capital Formation, 1919-1935*, pp. 61–90. NBER.
- [42] Kuznets, S. (1953). International differences in income levels: Reflections on their causes. *Economic Development and Cultural Change*, 3–26.
- [43] Kuznets, S. (1955). Economic growth and income inequality. *The American economic review*, 1–28.
- [44] Kuznets, S., L. Epstein, and E. Jenks (1941). National income and its composition, 1919-1938, volume i. *NBER Books*.

- [45] Leach, G. (1992). The energy transition. *Energy policy* 20(2), 116–123.
- [46] Lee, K., E. Brewer, C. Christiano, F. Meyo, E. Miguel, M. Podolsky, J. Rosa, and C. Wolfram (2014). Barriers to electrification for "under grid" households in rural kenya. Technical report, National Bureau of Economic Research.
- [47] Lee, L. Y.-T. (2013). Household energy mix in uganda. *Energy Economics* 39, 252–261.
- [48] Lei, Y.-H. and G. Michaels (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics* 110, 139–157.
- [49] Leigh, A. (2009). Top incomes. *The Oxford Handbook of Economic Inequality*, 150–176.
- [50] Mehlum, H., K. Moene, and R. Torvik (2006). Institutions and the resource curse*. *The economic journal* 116(508), 1–20.
- [51] Michalopoulos, S. and E. Papaioannou (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica* 81(1), 113–152.
- [52] Michalopoulos, S. and E. Papaioannou (2014). National institutions and subnational development in africa. *The Quarterly Journal of Economics* 129(1), 151–213.
- [53] Mills, E. (2003). Technical and economic performance analysis of kerosene lamps and alternative approaches to illumination in developing countries. *Lawrence Berkeley National Laboratory Report*.
- [54] Parcerro, O. and E. PAPYRAKIS (2015). Income inequality and the resource curse. *Journal of Comparative Economics*, forthcoming.
- [55] Pfeffermann, G. and R. Webb (1983). Poverty and income distribution in brazil. *Review of Income and Wealth* 29(2), 101–124.
- [56] Pinkovskiy, M. and X. Sala-i Martin (2009). Parametric estimations of the world distribution of income. Technical report, National Bureau of Economic Research.
- [57] Pinkovskiy, M. and X. Sala-i Martin (2016). Lights, camera,... income!: Estimating poverty using national accounts, survey means, and lights. *The Quarterly Journal of Economics* forthcoming.
- [58] Pinkovskiy, M. L. (2013). Economic discontinuities at borders: Evidence from satellite data on lights at night.
- [59] Ravallion, M. (2003). The debate on globalization, poverty and inequality: why measurement matters. *International Affairs* 79(4), 739–753.
- [60] Ross, M. L. (2007). How mineral-rich states can reduce inequality. *Escaping the resource curse* 23775.
- [61] Rutstein, S. O. (2008). The dhs wealth index: Approaches for rural and urban areas. *DHS Working Paper No. 60*.

- [62] Rutstein, S. O. and S. Staveteig (2014). Making the demographic and health surveys wealth index comparable. *Rockville: ICF International*.
- [63] Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. Technical report, National Bureau of Economic Research.
- [64] Sachs, J. D. and A. M. Warner (2001). The curse of natural resources. *European economic review* 45(4), 827–838.
- [65] Sala-i Martin, X. (2006). The world distribution of income: falling poverty and... convergence, period. *The Quarterly Journal of Economics*, 351–397.
- [66] Small, C., F. Pozzi, and C. D. Elvidge (2005). Spatial analysis of global urban extent from dmsp-ols night lights. *Remote Sensing of Environment* 96(3), 277–291.
- [67] Smith, B. (2015). The resource curse exorcised: Evidence from a panel of countries. *Journal of Development Economics* 116, 57–73.
- [68] Smith, B. and S. Wills (2016). Darkness and poverty. *MIMEO, Department of Economics, University of Oxford*.
- [69] Stone, R. and G. Croft-Murray (1959). *Social accounting and economic models*. Bowes & Bowes London.
- [70] Stone, R. and O. E. de Cooperació Econòmica (1961). *Input-output and national accounts*. Organisation for European Economic Co-operation Paris.
- [71] Subramanian, S. and A. Deaton (1996). The demand for food and calories. *Journal of political economy*, 133–162.
- [72] Toews, G. and P.-L. Vezina (2016). Resource discoveries and fdi bonanzas. *MIMEO, Department of Economics, University of Oxford*.
- [73] Van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature*, 366–420.
- [74] Van Der Ploeg, F. and S. Poelhekke (2016). The impact of natural resources: Survey of recent quantitative evidence. *The Journal of Development Studies*, 1–12.
- [75] Wills, S. (2015). Optimal monetary responses to oil discoveries. *OxCarre Research Paper* 121.
- [76] World Bank (2013). Global monitoring report 2013: Rural-urban dynamics and the millennium development goals. Technical report, The World Bank Group.
- [77] Xie, M., N. Jean, M. Burke, D. Lobell, and S. Ermon (2015). Transfer learning from deep features for remote sensing and poverty mapping. *arXiv preprint arXiv:1510.00098*.

Appendix

A Oil-dependent countries

Country	Resource Exports/ Total Exports (%, 2006-2010)	Commodity Revenues/ GDP (%, 2006-2010)
Algeria	98	30
Angola	95	35
Azerbaijan	94	26
Bahrain	81	23
Bolivia	5	11
Brunei	96	45
Cameroon	47	6
Chad	89	15
Congo, Republic of	94	3
Ecuador	55	7
Equatorial Guinea	99	31
Gabon	83	18
Indonesia	10	5
Iran	79	17
Kazakhstan	60	11
Kuwait	93	62
Libya	97	56
Malaysia	8	8
Mexico	15	8
Nigeria	97	22
Oman	73	37
Papua New Guinea	80	10
Qatar	88	23
Russia	50	11
Saudi Arabia	87	42
Sudan	97	11
Trinidad & Tobago	38	17
Turkmenistan	91	11
United Arab Emirates	41	24
Venezuela	93	19
Vietnam	14	6
Yemen	82	22

Table A.1: List of oil dependent countries (Baunsgaard et al., 2013).

B Evaluating darkness as a poverty measure

Country	Survey Year	Poor		Rich		Accuracy	Poverty Share	
		Unlit	Lit	Unlit	Lit		Unlit	Lit
Angola	2011	28%	15%	6%	50%	78%	82%	23%
Bangladesh	2011	10%	61%	1%	29%	39%	93%	68%
Benin	2012	30%	38%	2%	30%	60%	94%	56%
Bolivia	2008	23%	11%	5%	61%	84%	82%	15%
Cambodia	2010	32%	29%	7%	32%	64%	82%	48%
Cameroon	2011	44%	12%	7%	37%	82%	87%	24%
Colombia	2010	5%	8%	4%	83%	88%	56%	9%
Congo, Dem. Rep.	2013	74%	8%	3%	15%	90%	97%	33%
Dominican Republic	2013	1%	11%	1%	87%	88%	59%	11%
Ethiopia	2011	60%	18%	2%	20%	80%	97%	47%
Gabon	2012	33%	11%	8%	48%	81%	81%	18%
Ghana	2008	22%	29%	3%	46%	69%	89%	38%
Guinea	2012	49%	5%	13%	33%	82%	79%	13%
Guyana	2009	14%	3%	24%	59%	72%	36%	5%
Haiti	2012	30%	35%	3%	32%	62%	91%	52%
Honduras	2012	10%	36%	2%	52%	63%	86%	41%
Kenya	2009	42%	20%	5%	32%	74%	89%	39%
Lesotho	2009	41%	23%	6%	31%	71%	88%	43%
Madagascar	2009	69%	6%	7%	18%	87%	90%	25%
Malawi	2010	50%	37%	3%	10%	60%	95%	79%
Mali	2013	59%	15%	2%	24%	83%	96%	38%
Morocco	2003	10%	22%	7%	62%	71%	59%	26%
Mozambique	2011	47%	22%	2%	29%	76%	96%	44%
Namibia	2013	21%	23%	9%	47%	67%	69%	33%
Nepal	2011	44%	17%	8%	31%	75%	84%	35%
Nigeria	2008	37%	21%	6%	36%	74%	87%	36%
Peru	2009	21%	23%	3%	53%	74%	87%	30%
Philippines	2008	13%	20%	8%	60%	72%	62%	25%
Rwanda	2010	65%	20%	3%	12%	77%	96%	63%
Senegal	2011	35%	29%	2%	33%	68%	94%	47%
Sierra Leone	2013	64%	14%	2%	20%	84%	96%	40%
Swaziland	2006	12%	29%	4%	55%	67%	74%	35%
Tajikistan	2012	5%	7%	7%	81%	86%	40%	7%
Tanzania	2010	53%	23%	3%	21%	74%	94%	52%
Zambia	2013	45%	19%	5%	31%	76%	90%	37%
Zimbabwe	2010	38%	7%	16%	38%	76%	70%	16%

Table B.1: Effectiveness of darkness as a poverty measure by country at 12×12 pixel resolution, using the CWI definition of extreme poverty from 636,448 household DHS surveys.

		CWI Poor (1×1)					UBN ≥ 2 Poor (12×12)				
		Yes	No	Total			Yes	No	Total		
i.	Unlit	Yes	44%	7%	52%	ii.	Unlit	Yes	36%	7%	44%
		No	9%	39%	48%			No	26%	30%	56%
		Total	54%	46%	100%			Total	63%	37%	100%

Table B.2: Summary of how accurately darkness measures extreme poverty according to i. the CWI definition at 1×1 pixel resolution and the ii. $UBN \geq 2$ definition at 12×12 pixel resolution.

C Calibrated Poverty Rate

For robustness we also calculate a second calibrated poverty measure, building on Elvidge et al. (2009). We first relate observed lights per capita to World Bank \$2 per day poverty rates at a national level. Then we apply this relation at a cell level. Finally, we correct for issues with the spatial data to produce a Calibrated Poverty Rate (CPR) for each cell. We then use the CPR in robustness exercises in Appendix E.

First we use a non-parametric, decile-based approach to link observed lights per capita to World Bank \$2 per day poverty rates at a national level each year. This involves, each year, calculating the ratio of total lights to total population for each country that has a World Bank poverty rate estimate made in or after 2005 (103 countries). We then rank the countries and assign them to a decile. We then regress the World Bank poverty rate, which is the percentage of people living on less than \$2 per day, on the indicator for each decile, effectively calculating the mean World Bank poverty rate for each decile. We find this non-parametric decile-based approach performs better than the linear specification used in Elvidge et al. (2009).

Then, we apply the national-level relationship between lights per capita and the World Bank poverty rate to each cell (12×12 pixels). This produces an estimate of the World Bank poverty rate for each cell, which can readily be re-aggregated to a sub-national or national level. The output is a globally balanced panel from 2000-2013 at a 12×12 pixel resolution.

Finally, we make adjustments to correct some issues with the spatial data. In cells with zero lights we assume that 90% of the population lives below the poverty line. This is more than the lowest decile (which is 76% in 2011), but less than the 100% assumed by Elvidge et al. (2009) which causes the CPR to overestimate poverty in very poor countries. In cells with a lot of light we assume that the poverty rate is zero if the average light reading exceeds 50 (out of 63). This addresses top-coding which biases lights per capita downwards, and overestimates poverty in highly urbanised countries. While top-coding does not affect many pixels in developing countries it may affect a lot of people, as those pixels will be in densely populated areas. We also assume that the final poverty rate in urban cells is 60% of that given by the calibration procedure. This further helps our measure fit national World Bank poverty rates in highly urbanised countries. Overall the CPR involves some arbitrary assumptions but has a close relationship with World Bank poverty rates at a national level, as shown in Figure C.1. The regression line is close to the ideal of a 45 degree line passing through the origin, and the adjusted r-squared is 0.86. Because the CPR involves these assumptions we have a preference for the unlit




Figure C.1: Less than \$2 per day percentage (World Bank) vs 2011 Calibrated Poverty Rate (CPR)

rural percentage due to its simplicity, transparency and intuitiveness, and use the CPR for robustness.

D Mechanisms

In addition to the hazard model described in Section 4.6, which studies the marginal effect of various variables on the probability of a rural cell becoming lit, we also study the cumulative effect of each variable individually. To do this we run the following specification at the grid-cell level,

$$RuralSwitch\%_{c,i,t} = \alpha + \sum_{s=2000}^{2013} \beta_s (\lambda_s \times X_c) + \Lambda_t + \Phi_i + \epsilon_{c,i,t} \quad (D.1)$$

Where X_c is one of the dependent variables described in Section 4.6, and the other variables are defined in Section 4.2. The specification therefore estimates, for example, the likelihood of unlit rural cells that are adjacent to lit cells switching on relative to other unlit cells, controlling for country-level averages. Standard errors are again clustered at the country level. We run this specification separately for dependent and non-dependent countries, and then assess whether the effect differs between the two using a triple-difference specification.¹⁵

¹⁵This takes, for example, the difference in rural illumination between cells that are/are not adjacent to lit cells, relative to their difference in 2000, and assesses how much this differs between oil dependent and non-dependent countries. The specification is run for all sample countries and includes dummies for being in a dependent country and being adjacent to a lit cell, dependence-by-year interactions, adjacency-by-year interactions, and triple interactions for which the coefficients are shown in Figure D.1, iii.

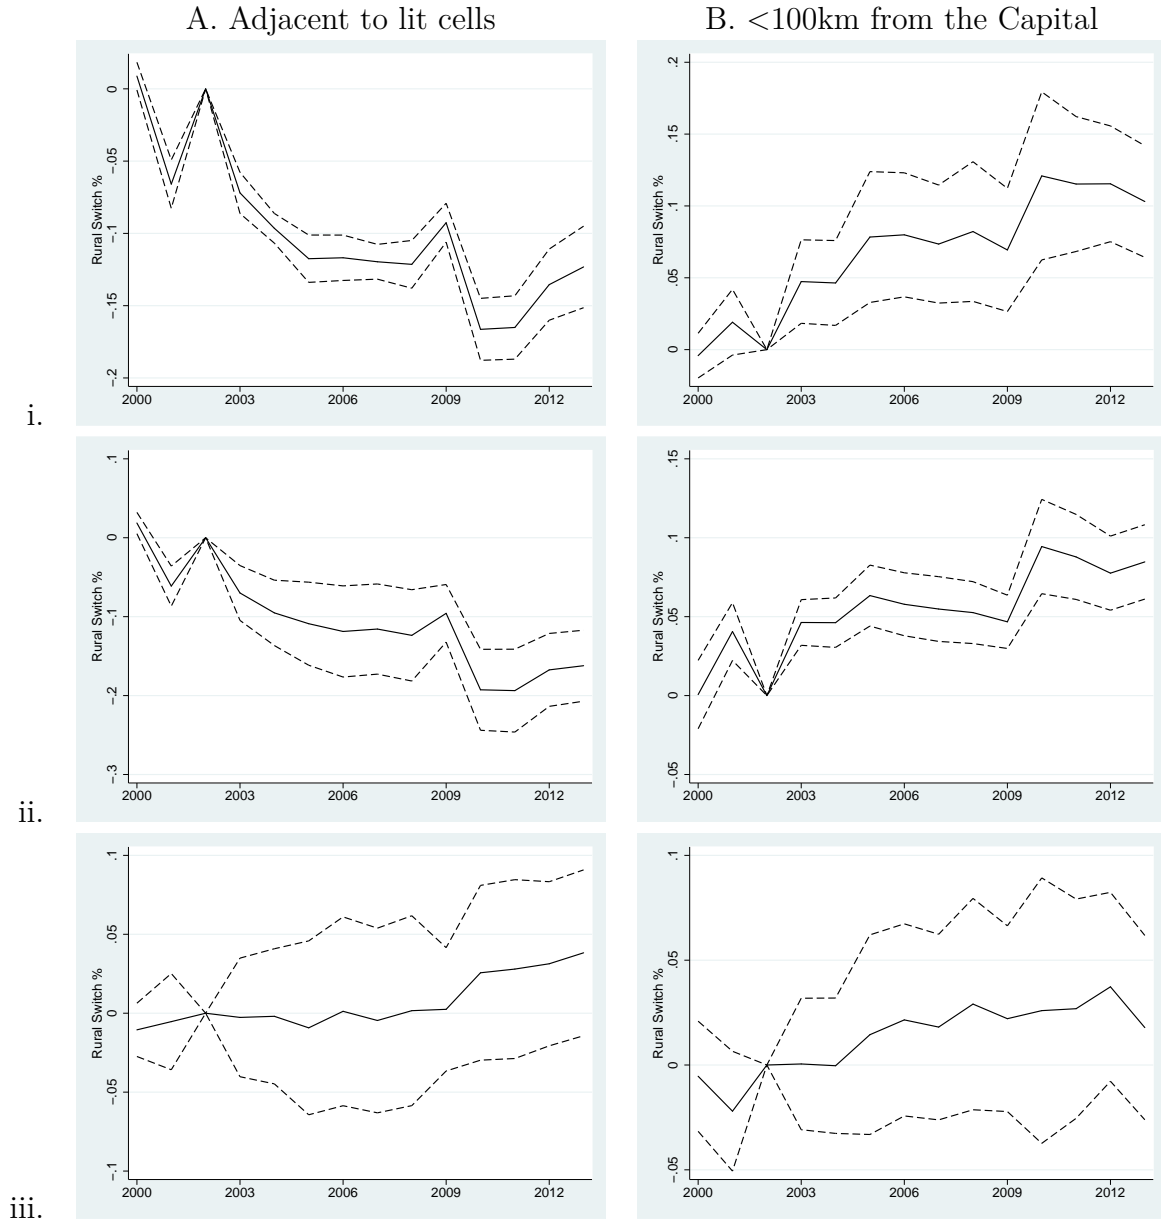


Figure D.1: The difference in switch-on percentage for rural areas that are A. adjacent vs not-adjacent to lit squares and B. near (<100km) vs far (>100km) from the capital city, in i. oil-dependent countries, ii. non-dependent countries and iii. the difference between i. and ii..

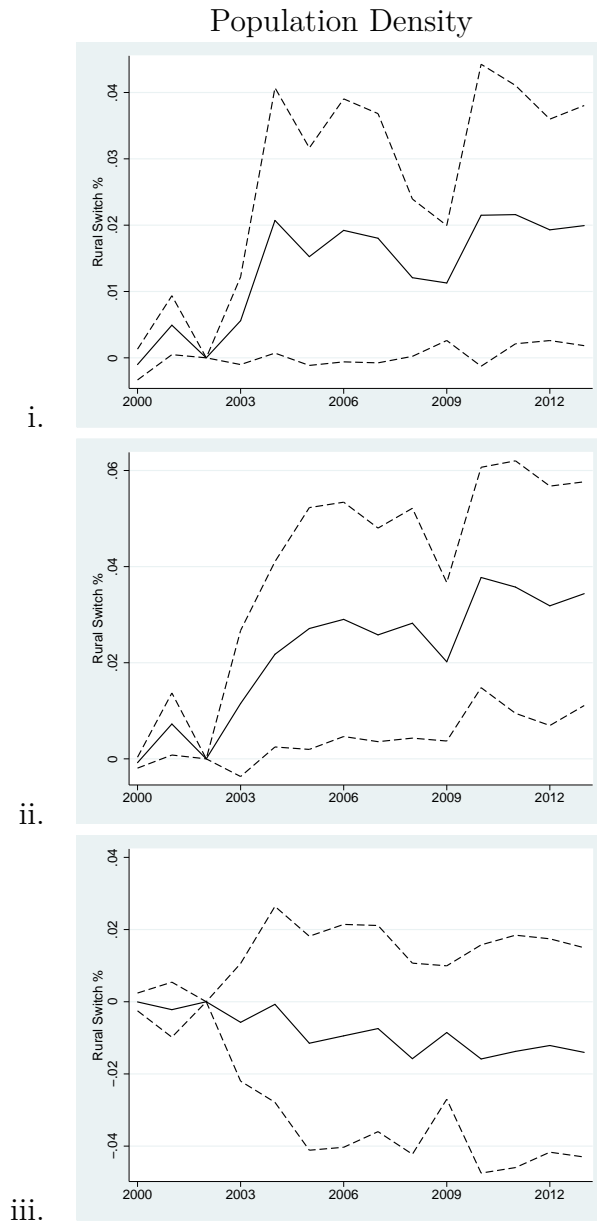


Figure D.2: The difference in rural switch-on percentage for a one standard-deviation increase in population density in i. oil-dependent countries, ii. non-dependent countries and iii. the difference between i. and ii..

E Robustness




Figure E.1: Effect of the 2000s oil price boom on i. aggregate lights (estimates and 95% confidence bands), and ii. PPP-adjusted real GDP per capita, extended time axis.




Figure E.2: Changes in the Calibrated Poverty Rate due to i. the 2000s oil price boom and ii. giant oil discoveries.

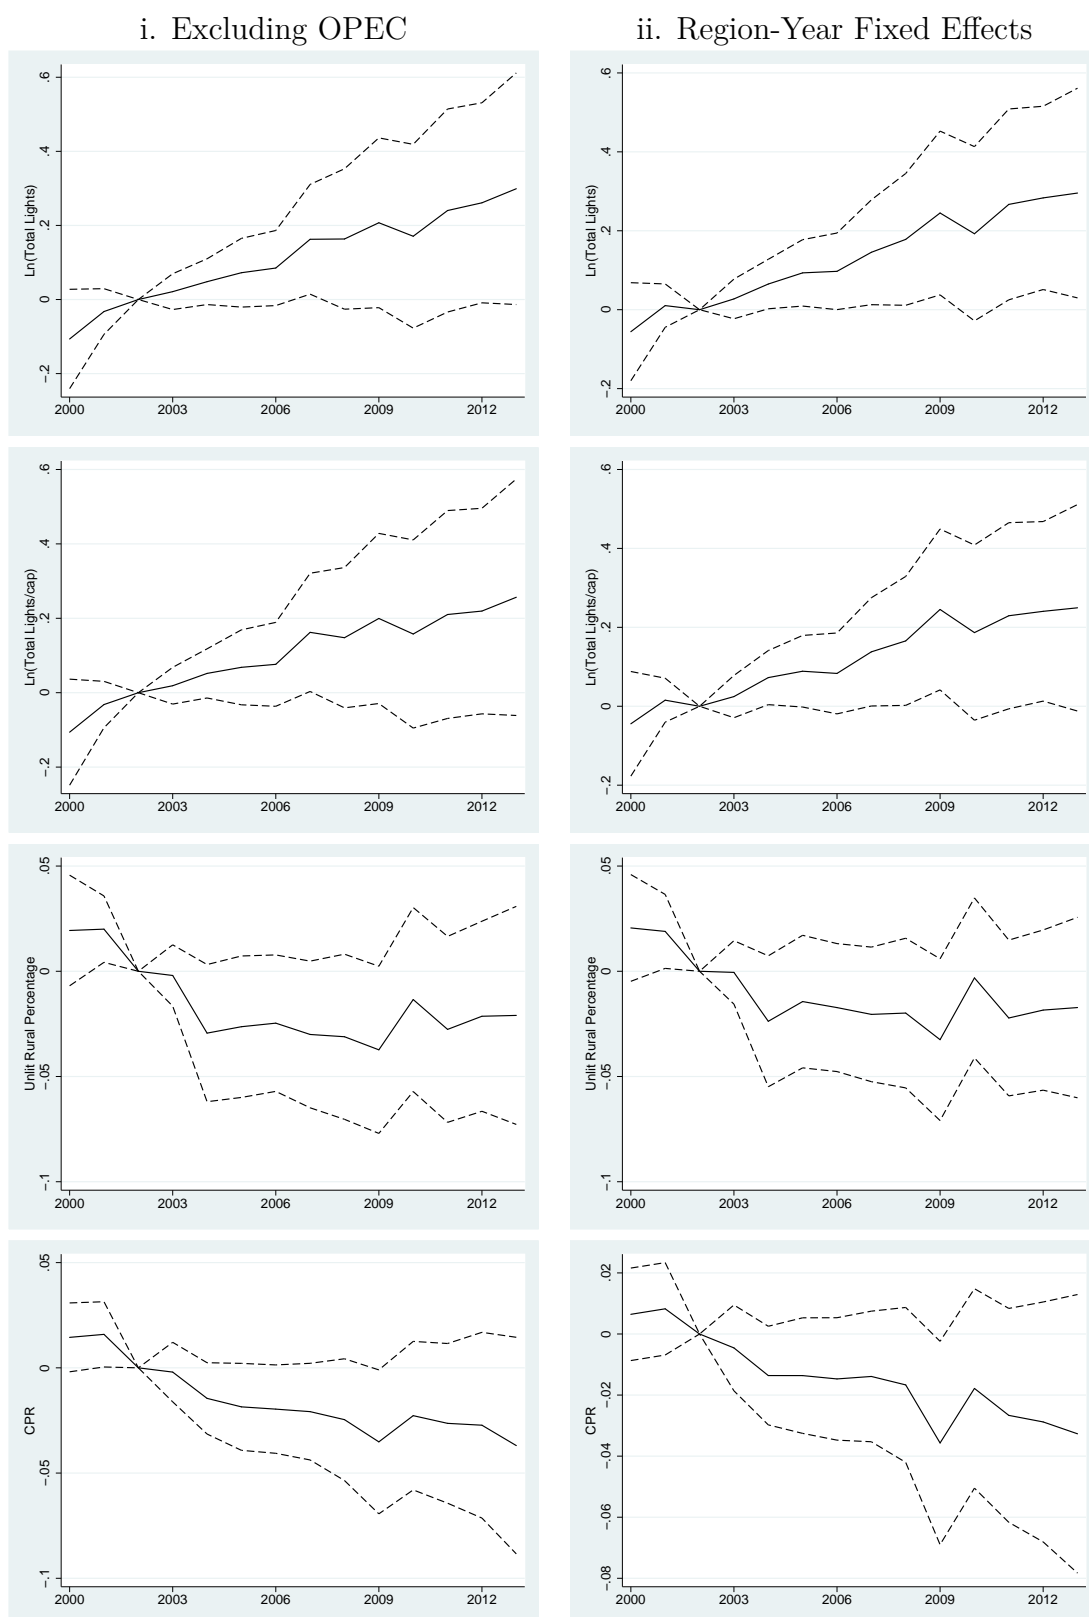


Figure E.3: Effect of 2000s oil price boom on key variables i. excluding OPEC and ii. using region-year instead of year fixed effects.

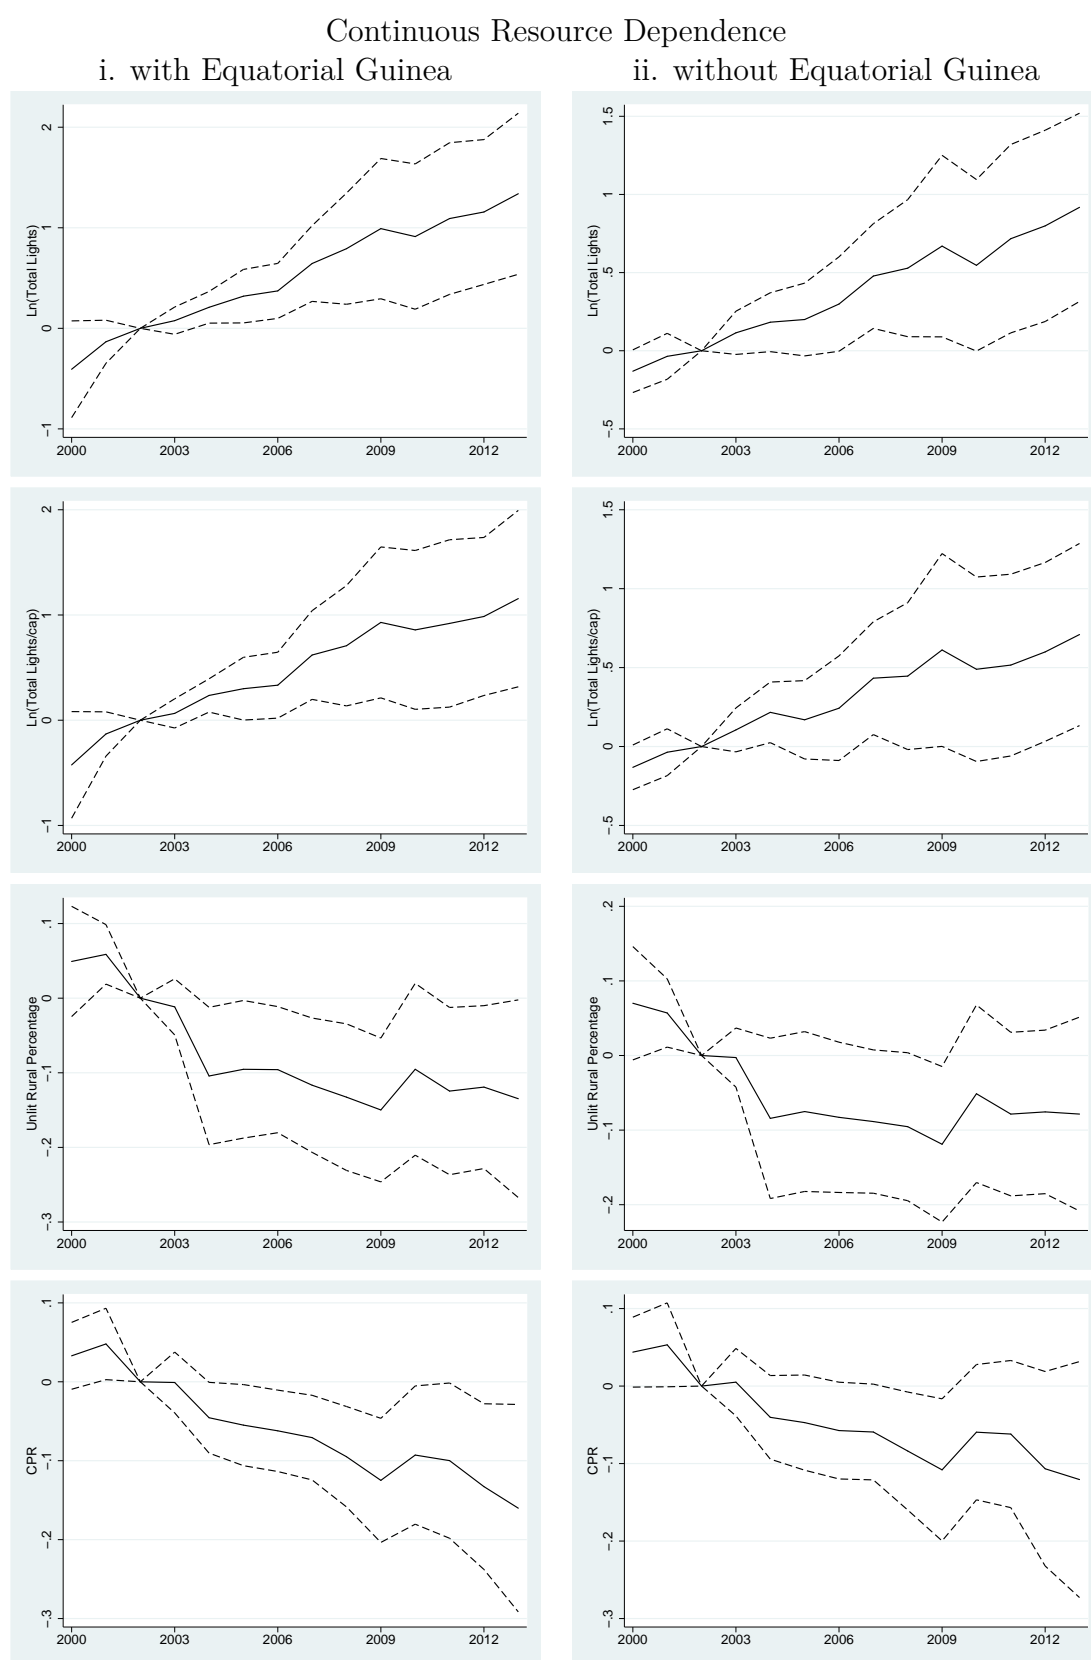


Figure E.4: Effect of 2000s price boom on key variables, using a continuous measure of resource dependence i. with and ii. without Equatorial Guinea.

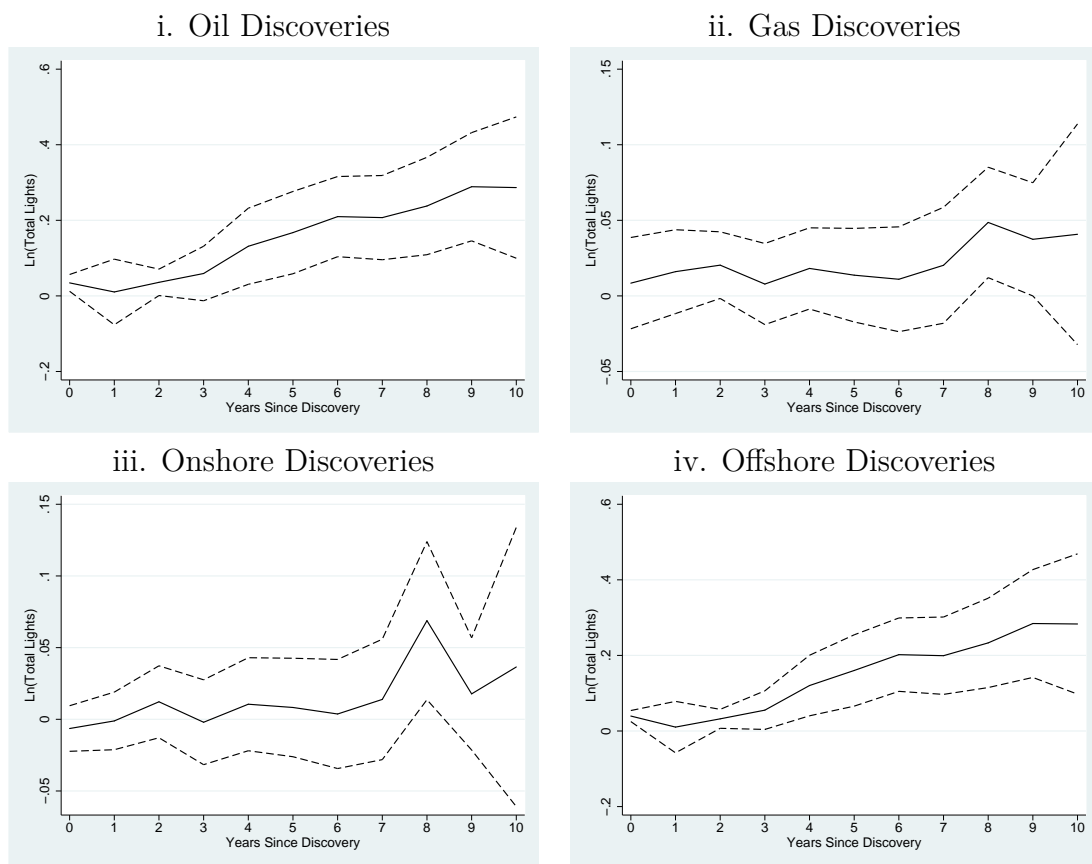


Figure E.5: The effect on aggregate (log) lights of i. oil, ii. gas, iii. onshore, iv. offshore discoveries.