Left in the Dark? Oil and Rural Poverty

Brock Smith and Samuel Wills*

October 2015

Abstract

How are oil revenues shared throughout society? We combine high-resolution geo-coded data on night-time lights and population to construct global measures of rural poverty from 2000-2013. We find that oil booms, due either to high prices or new discoveries, increase light intensity and GDP. However, the increase in output is limited to cities and towns, with no evidence that it benefits the rural poor. We also find that while urbanization is occurring throughout the developing world, there is no evidence that it is hastened by oil wealth.

*Oxford Center for the Analysis of Resource Rich Economies, Department of Economics, University of Oxford, UK. Samuel Wills would like to thank the Economic and Social Research Council for financial support [grant number ES/K009303/1]. Corresponding email: brock.smith@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 work employing cross-country regressions (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). Most of this research is concerned with oil's impact on aggregate economic activity. However, if we are concerned with living standards, then it is important to look beyond country-level measures to study how oil income is shared throughout society. This has largely been ignored by the literature. 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 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 30% of countries in any given year (see Figure 1.1), 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".

This paper presents a global panel study of inequality and extreme poverty rates by constructing a novel measure of rural poverty. We do this using two detailed and geographically disaggregated datasets, on night-time lights and population (see Figure 1.2). The first records the amount of light emitted at night around the globe at a 1km^2 resolution, which is a useful geographic proxy for economic activity (Henderson et al., 2011; 2012). This has been used in studies covering institutions (Michalopoulos and Papaioannou, 2013), political favoritism (Hodler and Raschky, 2014), and infrastructure investment (Jedwab and Moradi, 2015; Jedwab et al., 2015) amongst others. The second is from LandScan and uses several spatial data inputs to measure population, also at a 1km^2 resolution. After aggregating the data to 10km^2, our final sample includes information from 1.04 million cells each year for 2000-2013. The LandScan data has not received much attention in economics research, but it reveals an important point: that many people live in areas not illuminated at night (see for example Figure 1.3). They are the rural poor.

Our poverty measure is constructed by recording the population share of each country living in unlit rural areas. While this is a somewhat crude proxy for rural poverty, it correlates well with World Bank poverty estimates (shown in Section 3.1). As an alternative we also create a calibrated poverty rate that uses the relationship between lights and poverty rates at the national level to assign poverty rates at local levels (adapting the work of Elvidge et al., 2009). This yields what is to our knowledge the first global balanced panel data set of severe poverty at the national (and sub-national) level.

Why is illumination related to poverty? Illumination is essentially a measure of electrification. Poverty is the state of being unable to meet one's basic needs. As we move

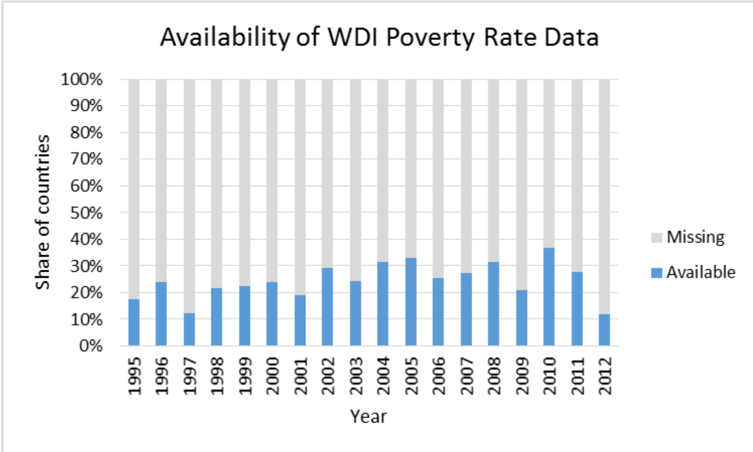

Figure 1.1: Share of countries with data available on people living below US$2 per day (World Bank WDI).

up Maslow's hierarchy from food and shelter we soon need electrification. Not only does it increase productivity during the working day, it extends the day, improving access to healthcare, safety and education. Thus, while illumination is only one element of the poor's consumption bundle, there are reasons for it to be an accurate proxy for poverty.

There are two mechanisms for improving rural illumination: public or private investment. Public investment involves connecting rural areas to a wider electrical grid, or providing off-grid electricity through diesel, solar or hydro generators. Private investment is mainly spent on off-grid solutions. The path from typically government-owned oil revenues to illumination therefore depends heavily on policy. If the government prioritises rural over urban infrastructure then rural illumination will improve. If the government instead provides direct cash transfers then private investment might achieve the same, while tax cuts would instead benefit wealthier urban areas. So, by investigating how oil booms affect the rural poor we are by proxy studying government spending priorities in oil rich countries.

To start we outline a few stylised facts about rural poverty and urbanisation around the world. First, GDP is steadily rising in both oil dependent and non-dependent countries. Second, the share of people living in unlit rural areas is steadily falling. This occurs because both lights are being switched on and people are moving to towns and cities.

We then assess how oil booms affect the rural poor and to do so we need a well-identified oil shock. We approach this in two ways, covering both prices and quantities. The first uses the sharp rise in oil prices after 2003 as an exogenous, demand-driven shock to oil revenues. According to Kilian (2009), "the surge in the price of oil after 2003 was driven primarily by the cumulative effects of positive global demand shocks". However, illumination during this period might also be due to higher global demand for all goods which could possibly affect oil-dependent countries differentially. To overcome this, and to distinguish between a price and a production shock, we also use data on exogenous giant oil and gas field discoveries (which has also been used as a well-identified income shock by Lei and Michaels, 2014; Arezki et al., 2015; Smith, 2015 and Wills, 2015).

Our analysis offers three main results. First, the decade of high oil prices from 2003 saw night-time illumination grow significantly faster in oil dependent countries than non-

i. ii.

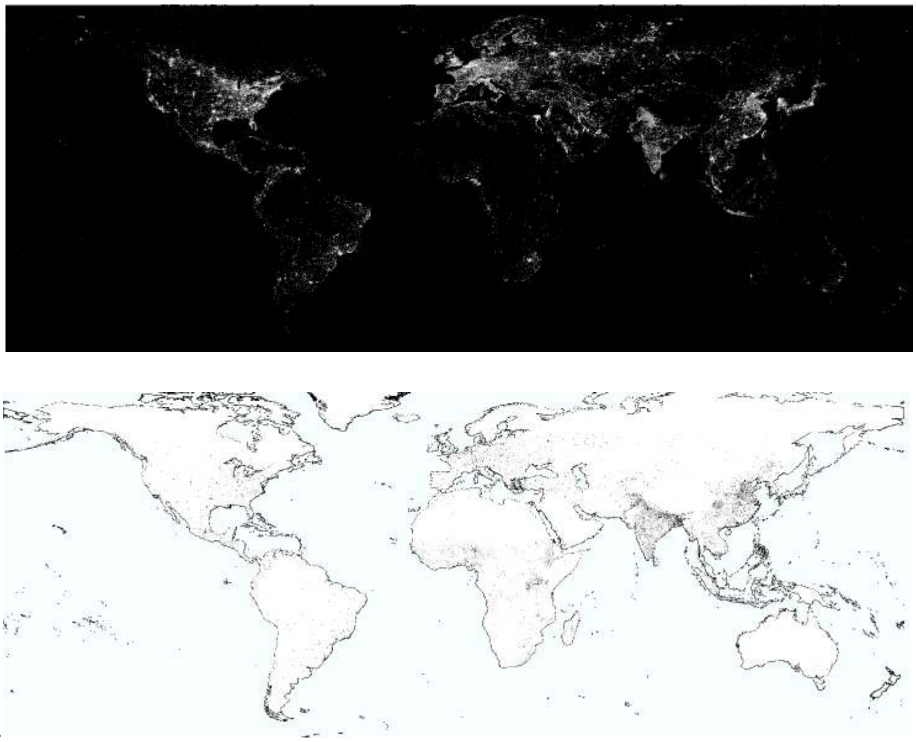

Figure 1.2: i) Night-time lights and ii) population around the world.

dependent ones. By the end of the decade the price boom was responsible for approximately a .29 log point increase in total illumination, with similar effects for GDP. Furthermore, illumination grew significantly in countries that had made giant oil discoveries, starting with a six year lag (consistent with Arezki et al., 2015). We find that an oil discovery with a net present value worth 100% of GDP increases illumination by .07 log points after ten years.

Second, economic growth from oil booms is not shared with the rural poor. While rural poverty is falling around the world, this is not hastened by oil booms. Price booms have an insignificant effect on the rural poverty rates of oil dependent countries, while giant discoveries cause a small (1 percentage point) but significant fall in rural poverty after ten years. For both the period of high prices and the giant oil discoveries, the entire additional rise in illumination occurred in towns and cities where lights already existed. There is no evidence of unlit rural areas becoming illuminated at a higher rate in either case. Declining rural poverty after a giant discovery is due to the population of unlit rural areas falling, and that of towns rising, most likely due to migration. This invites a different interpretation to that of Aragon 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 close to mines.

Third, oil booms do not change the rate of urbanization, countering the view that natural resources promote urbanization without industrialization (Gollin et al., 2013; Cavalcanti et al., 2014). To determine why total illumination increases but poverty remains unaffected, we divide our data into three areas: cities (or urban areas), towns and unlit but inhabited rural areas. Consistent with Gollin et al (2013), non-OECD countries with oil have higher levels of urbanization than those without. Urbanization also increases for both groups during the sample period. However, we find no evidence that either the rise in oil prices or giant oil discoveries hastened the pace of urbanization. This differs from the work by Gollin et al. (2013) because we explicitly identify exogenous oil shocks. The oil price shock also did not alter population shares in unlit rural areas or towns, while oil discoveries encourage some relocation from the former to the latter.

To further understand these main results, we zoom in to the grid cell level and attempt to identify the mechanisms behind falling poverty around the world. Estimating the hazard rate of unlit cells switching on during our sample, we find that cells are more likely to switch on if they are adjacent to other lit cells (and therefore the existing grid), are near the capital city or have higher population density. This is true for both dependent and non-dependent countries. Furthermore, cells in countries that have benefited from higher overall growth in illumination in the past are more likely to switch on in the present. However, the effect of past growth on illumination is significantly dampened in oil-dependent countries, further implying that the extra growth from oil shocks does not benefit the poor.

This paper contributes to a number of strands of literature. The first is on the measurement of income, poverty and inequality, which has a long and distinguished history (for example Kuznets 1937, 1941, 1953; Stone 1959, 1961; Atkinson, 1970; Deaton 1980, 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 (2014) use night-time lights data to reconcile this conflict, finding in favour of aggregate measures. The advantage of surveys is that they offer 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), within-household distributions (Deaton and Muellbauer, 1986) and a host of other factors. Compared to this we offer a relatively crude measure of poverty, though it covers the entire world, at fine resolution and regular intervals, and is relatively cheap to collect.

The second is the literature on how growth generally affects poverty and inequality. Kuznets (1955) famously observed that inequality widened as workers moved from agriculture to industry, up to a point where the trend reversed. Recent groundbreaking work on top income shares shows that they peaked in the inter-war years, fell until the 1970s-90s, and started rising again since (see survey by Atkinson et al., 2009). They were initially dominated by capital income, but in recent years labour income dominates. These studies focus only on the top of the income distribution, but this is correlated with relative poverty at the bottom (Leigh, 2009). Collectively this body of work uses up to three centuries of top income data, covering over 20 mostly developed countries. In contrast our study focuses directly on the bottom of the income distribution. We are limited by satellite data availability to 2000-2013, though we cover the whole world at high resolution.

The third is the resource curse. This literature has generally been concerned with how resource wealth affects aggregate income and growth. Little has been done on the distribution of income, poverty and inequality. One of the few studies which does notes that "the empirical literature on the inequality and resource boom connection is relatively thin" (Bhattacharyya and Williamson, 2013). This study analyses the effect of

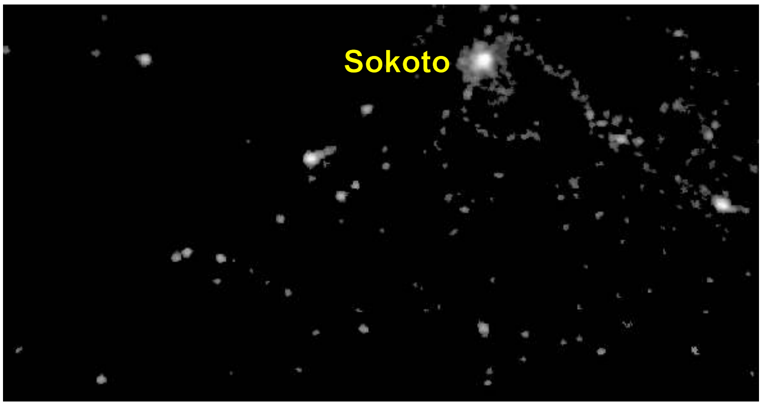

i.

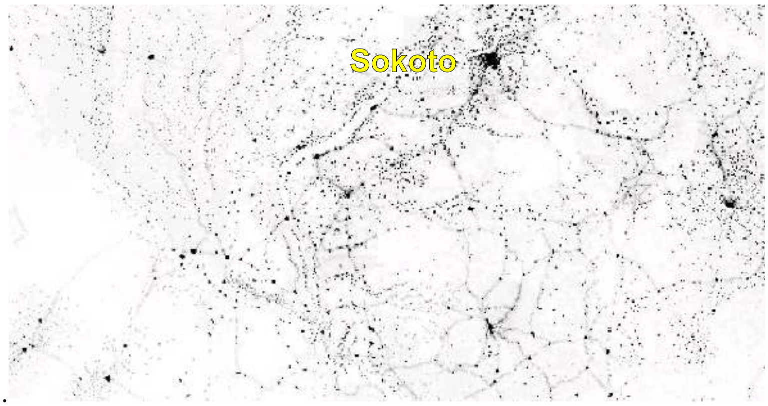

ii.

Figure 1.3: i) Night-time lights and ii) population for the region surrounding Sokoto in north-west Nigeria.

commodity price shocks on top income shares in Australia, building on data collected by Atkinson and Leigh (2007). They find that the richest benefit disproportionately from resource booms, but not agricultural booms, and attribute this to the initial distribution of assets. Resource booms have also been found to increase inequality using cross-section data (Gylfason and Zoega, 2003), a 90 country panel (Goderis and Malone, 2011) and inter-regional variation in Russia (Buccellato and Mickiewicz, 2009).

The paper proceeds as follows. Section 2 introduces our data and Section 3 the methodology. Section 4 presents and discusses our findings and a range of robustness checks. Section 5 concludes.

2 Data

2.1 Night-time lights

Satellites from the Defense 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 resolution of 30x30 arcseconds, or about 1 square kilometer near the equator, and ranges from 0 to 63. The data is constructed by

i) 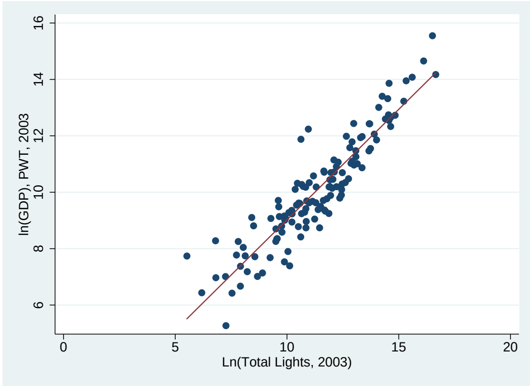

Figure 2.1: PPP-adjusted real GDP vs Night-time lights (in logs)

overlaying all daily images over the course of a year, discarding those that are obfuscated by cloud cover, lightning, aurora, etc. for a given pixel.

The pioneering work of Henderson et al (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 .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 measures (described in Section 3.1) by combining it with spatial population data and identifying populated areas with no light activity.

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, such that we cannot distinguish levels of economic activity above this value, and occurs in especially dense or economically active areas. According to Michalopoulos & Papaioannou (2014) this problem is far more prevalent in developed countries, but it does create difficulties in estimating urban poverty rates (see Section 3.1). Lights data also include significant luminosity readings resulting from gas flares, which do not reflect comparable economic activity. This is an especially important issue for this paper since our treatment group mostly includes significant gas-producing countries. While this problem is mitigated to the extent that the flares are similar over time by the inclusion of country fixed effects, in all analysis we drop grid cells that include gas flare activity according to spatial data provided by the Earth Observation Group (which also oversees the lights data).

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 30x30 arcsecond resolution.[1] This is similar to NASA's Socioeconomic Data and Applications Center (SEDAC),

[1]This is the same resolution as the lights data, although the pixels are not aligned. The grid cells described in Section 3.1 are aligned with the lights rasters but not the population rasters. The Zonal

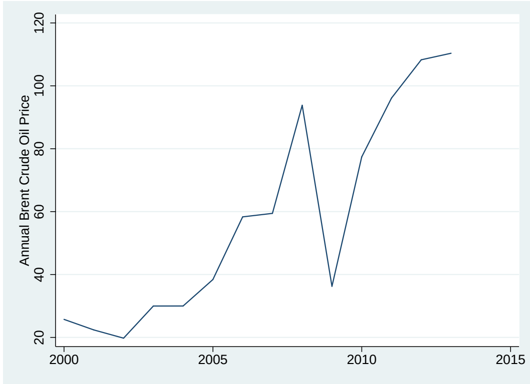

Figure 2.2: Annual Brent crude oil prices (1990-2013).

which also measures population at a 30x30 arcsecond 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.[2]

2.3 Urban and Rural Classifications

SEDAC also provides an "Urban Extents Grid", which uses 1995 population count estimates to classify each square of a 30x30 arcsecond 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.

2.4 Oil Prices and Discoveries

We study two types of oil shocks: to prices and to the quantity of ultimately recoverable reserves. The oil price shock exploits the period of high prices from 2003 and 2013, and its differential effect on oil dependent and non-dependent countries (see Figure 2.2).

The classification of oil dependence comes from Baunsgaard et al (2012), which categorises countries as resource-dependent 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.[3]

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.

[2]For further detail http://web.ornl.gov/sci/landscan/landscan_documentation.shtml)

[3]As mentioned below in Section 3.4, Iraq and Syria are dependent countries that are excluded from this study, since both experienced devastating wars during the sample period.

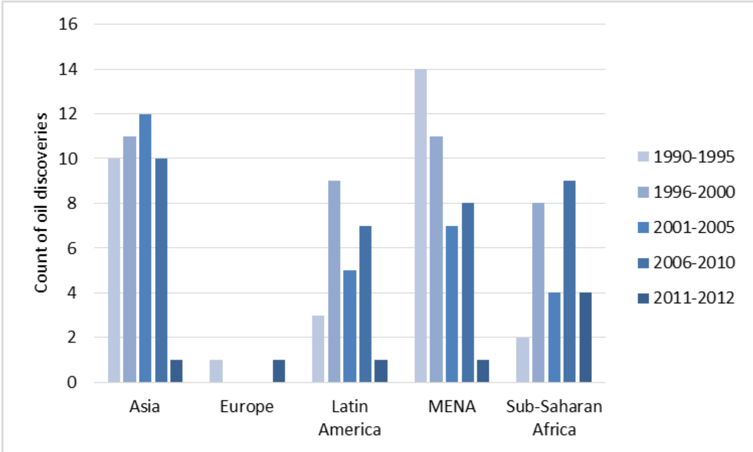

Figure 2.3: Count of giant oil discoveries by region included in our sample (excluding OECD countries).

The oil quantity shock uses data on giant oil 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), and it has been used in studies by Lei and Michaels (2014), and Arezki et al. (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 139 unique discovery-years in our sample (see Figure 2.3).

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., 2015),

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

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. (2015).

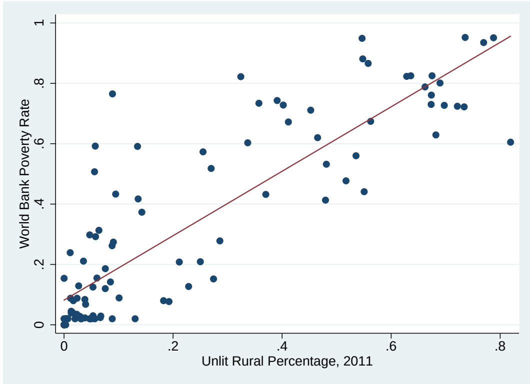

Figure 3.1: Less than $2 per day percentage (World Bank) vs 2011 Unlit rural percentage

3 Empirical Methodology

3.1 Constructing Poverty Measures

We construct and analyze a novel measure for rural poverty by combining the spatial lights and population data. For robustness we also use a second, calibrated measure. Both calculate poverty at a local level and can then be aggregated to a national level.

The first step is to create a global grid consisting of 10x10 pixel cells, which is approximately 10km x 10km near the equator. Cells are divided at the border so no cell occupies more than one country. Therefore not all cells are 10km x 10km, though the vast majority are.

Our primary measure is the unlit rural percentage. It is constructed by summing the population that that live in cells with zero light readings anywhere in the cell, and dividing by the total country population (excluding those cells dropped due to gas flare activity). We reason that people in cells containing moderate population density with no light activity, indicating little economic activity beyond agriculture, are among the extreme poor. Of course, there are presumably cells with non-poor residents with sufficiently low population density to render lights readings of zero. By definition these cells will contribute little to the overall unlit percentage.

Figure 3.1 plots the unlit rural percentage by country in 2011 against the World Bank's most updated estimates of percentage of people living on less than $2 per day.[4] While there are some countries for which the unlit percentage is a poor predictor,[5] there is a strong correlation, and the corresponding regression yields an adjusted r-squared of .74.

We also construct a second poverty measure from a calibration procedure adapted from Elvidge et al. (2009). That paper constructed a global map of poverty for the year 2008,

[4]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.

[5]Figure 3.1 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 Phillipines, Swaziland, Bangladesh, India and Pakistan. With the exception of Swaziland, these are high-density poor countries with presumably high rates of urban poverty, which the unlit rural percentage does not capture. The calibrated poverty measure described below performs much better for these particular countries.

but did not study its determinants as we do. In short, this procedure relates observed lights per capita to world bank poverty rate estimates at the country level, then applies this relation to the grid cell level to obtain a poverty percentage estimate for each cell, then making various adjustments to address issues with the spatial data. Since we use the World Bank's measure of people living on less than $2 per day, we end up with an estimate of this at the cell level, which of course can then be re-aggregated to a national or sub-national level.

The calibration procedure works as follows: for a given year, we calculate 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.[6] We then apply these regression estimates to the cell level based on each cell's lights per capita to obtain a cell-level poverty rate estimate.

We then make adjustments to certain types of cells. First, because the calibration procedure does not apply to cells with zero lights, we make the assumption that 90% of the population of such cells is living on under $2 per day. This is higher than the lowest decile (which is 76% in 2011), but specifying 100% as in Elvidge et al (2009) consistently overestimates poverty in very poor countries. Second, we specify that poverty rates are zero if the average light reading in a cell exceeds 50 (recall that the max value for a pixel is 63). In such cells the lights per capita value can be significantly biased downwards because of top-coding as discussed in the previous section, and while top-coding is rare in developing countries in terms of pixel percentage, it occurs in densely populated areas so can be significant in terms of population affected. Hence without this adjustment we consistently overestimate poverty in highly urbanized countries. Third, since the previous step does not completely solve the problem of overestimating poverty in more urbanized countries, we specify that for cells coded as urban the final poverty rate is 60% of the rate specified by the calibration procedure. Again, this is likely necessary due to top-coding, even if the problem is not as severe as in the case of mean luminosity exceeding 50.

Following all these steps we have a cell-level poverty measure we call the Calibrated Poverty Rate (CPR). While the adjustments described above involve somewhat arbitrary assumptions about poverty rates, they yield a close relationship between World Bank poverty rates and national CPR, as shown in Figure 3.2. The regression line is close to the ideal of a 45 degree line passing through the origin, and the adjusted r-squared is .86.

Because the CPR is more accurate but involves arbitrary assumptions, particularly in urban areas, we use both measures to estimate the effects of oil shocks, with a preference for the unlit rural percentage due to its simplicity, transparency and intuitiveness.

3.2 Urban and Rural Classifications

Since we are interested in the differential effects of oil booms on cities, towns and rural areas, we need to classify each cell as such. We base our classification on SEDAC's Urban

[6] The non-parametric decile-based approach was found to perform better than a linear specification as used in Elvidge et al (2009).

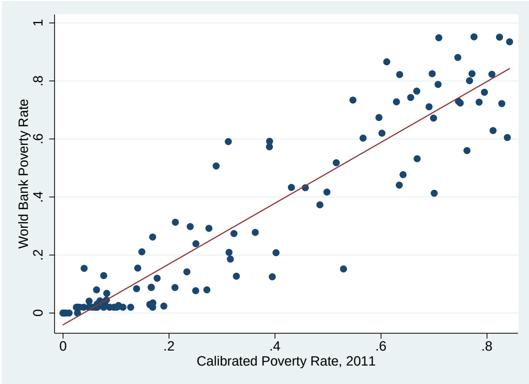

Figure 3.2: Less than $2 per day percentage (World Bank) vs 2011 Calibrated Poverty Rate (CPR)

Extents Grid, which as described above classifies each pixel as urban or non-urban. Since cells will then contain a mix of urban and non-urban pixels, we must specify a cutoff to be considered an urban cell. We classify a cell as urban if at least 33% of the pixels contained are urban. While necessarily arbitrary, this definition yields a global urban population very close to that estimated by the World Bank as of 2000. Cities are then 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.

3.3 Oil Prices and Discoveries

To study the effect of oil booms on rural poverty we exploit exogenous variation in oil prices and giant oil discoveries. Kilian (2009) argues that the increase in oil prices from 2003 was almost exclusively due to an exogenous global demand shock. This is done in a VAR framework using a structural decomposition to identify shocks to global oil supply, global industrial demand and oil-specific precautionary demand. Based on this we treat the increase in oil prices from 2003 as shock to global demand that is exogenous to oil-dependent countries. For robustness we also exclude OPEC countries and find that our results do not appreciably change.

The second type of shock is the discovery of giant oil fields, which we also treat as exogenous after controlling for time and country fixed-effects. Oil discoveries are a type of quasi-natural experiment. Some might argue that these discoveries are endogenous, depending on past discoveries or the quality of political institutions in a country. However, we are concerned specifically with the timing of the discovery, which is less predictable. We also only focus on giant discoveries, which are less predictable again. Finally, we control for time-invariant country-specific characteristics using country fixed-effects.

3.4 Estimating Equations

The effects of oil price shocks on spatial outcomes aggregated to the country level are estimated as follows:

$$Y_{i,t} = \alpha + \sum_{t=2000}^{2013} \beta_t(\lambda_t \times D_i) + \sum_{t=2000}^{2013} (\lambda_t \times Y_{i,2000}) + \lambda_t + \varphi_i + +region_i * t + \epsilon_{i,t} \quad (3.1)$$

where Y_{it} is the outcome of interest for country i in year t, D_i is an indicator variable 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, $region_i * t$ are regional linear trends,[7] λ_t is year fixed effects and Φ_i is country fixed effects. Each coefficient β_t 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 as follows:

$$Y_{i,t} = \alpha + \sum_{j=0}^{10} \beta_j Size_j + \sum_{t=2000}^{2013} (\lambda_t \times Y_{i,2000}) + \lambda_t + \varphi_i + +region_i * t + \epsilon_{i,t} \quad (3.2)$$

Where $Size_j$ is the NPV relative to GDP of a discovery made in year $t - j$. Thus β_j is the effect of a discovery made j years ago equal to 100% of GDP. This specification allows us to measure the dynamic effect of discoveries over time and to analyze multiple discoveries made within countries. The delay between discovery and production is typically 4-6 years, so we hypothesize positive effects on economic activity following this lag, but there may also be anticipation effects observed during the lag period. See Arezki et al. (2015) for a full discussion of these effects.

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%[8] as of 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 Results

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. As our data is geographically disaggregated we are able to isolate the regions that benefit from each type of boom. Overall we find that both price and quantity booms stimulate economic growth in the short to medium term. In both instances these booms do not benefit the rural poor. There is some evidence that oil discoveries cause a small amount of migration from unlit rural areas to towns, though not to cities.

[7] Regional 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.

[8] Results are robust to a threshold of 10%.

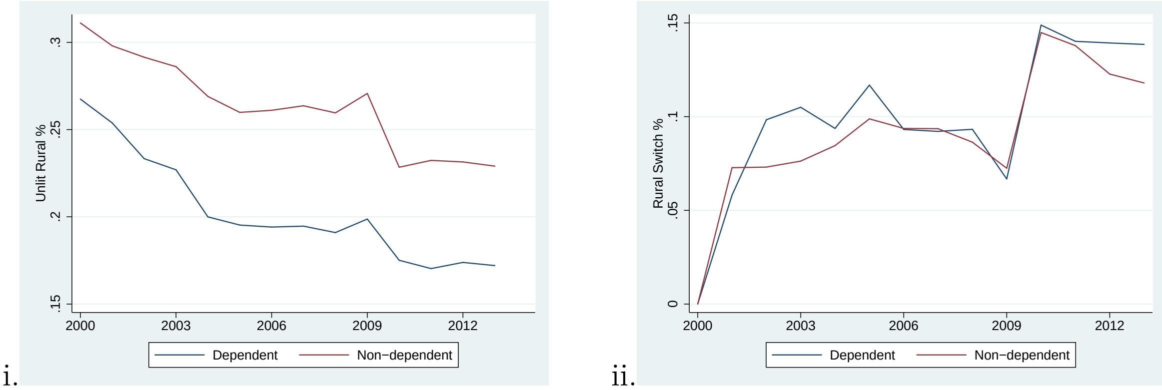

Figure 4.1: Unconditional trends in the i) unlit rural percentage and ii) switch-on percentage.

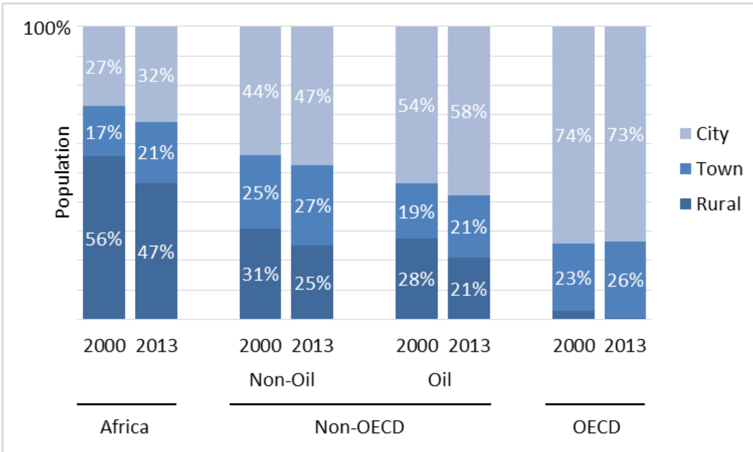

Figure 4.2: Breakdown of population share by area classification as of 2000.

4.1 Broad trends

To begin with let us examine the trends in the raw data. The first headline is that GDP is growing in both oil dependent and non-dependent countries. We know this from a broad range of sources, but it is best not to use the unconditional lights data to illustrate it because of year-to-year changes in light sensitivity and the introduction of new satellites in 2000, 2004 and 2010. In the main analysis we control for these changes using time fixed effects.

The second headline is that our poverty measure - the share of people living in unlit rural areas - is generally falling in both oil dependent and non-dependent countries (Figure 4.1). This is consistent with broad evidence that progress is being made in the fight against global poverty (e.g. Pinkovskiy and Sala-i-Martin, 2009; Chen and Ravallion, 2010). There are two reasons why the unlit rural percentage may fall: unlit areas are becoming illuminated or people are leaving these areas for towns and cities. We find evidence for both. Figure 4.1 introduces the “Rural switch percentage”, which records whether areas that were populated but unlit in 2000 became lit in subsequent years. As such it is a cumulative measure of rising economic activity in rural areas. We see that initially rural areas are steadily becoming illuminated during our sample, with the jumps in the data

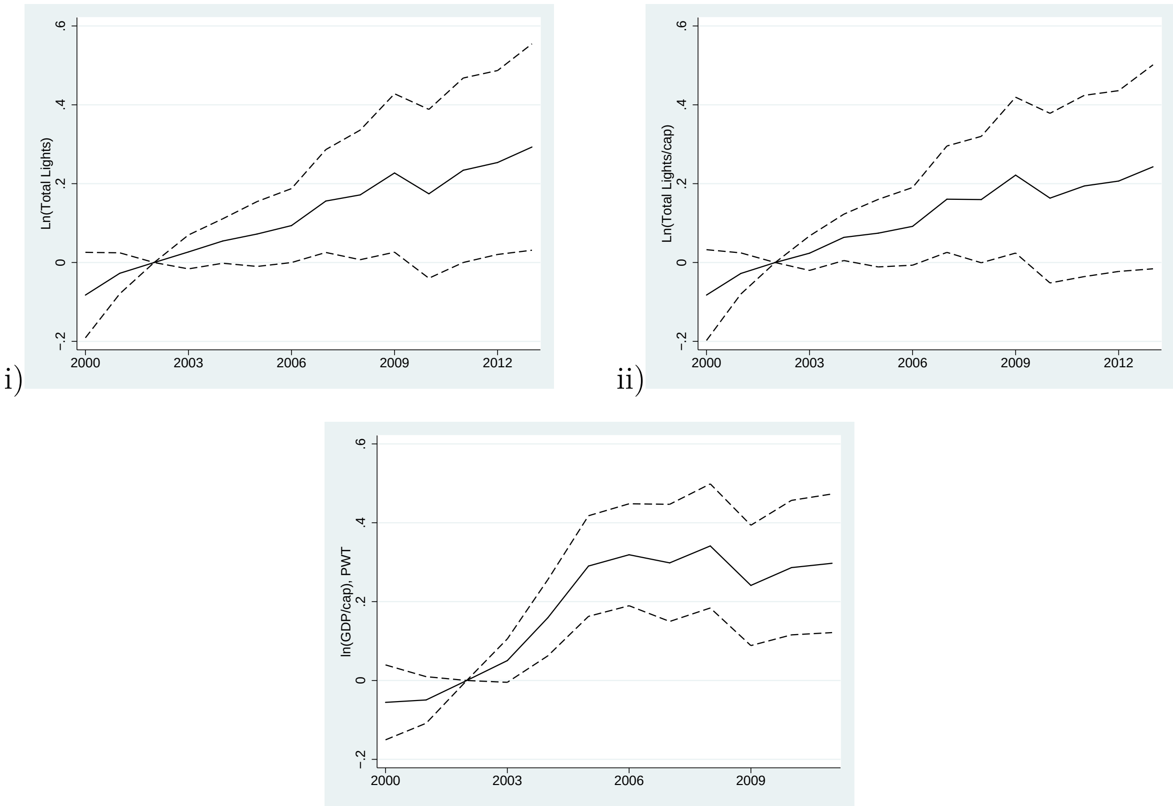

Figure 4.3: Effect of the 2000s oil price boom on i) aggregate lights (estimates and 95% confidence bands), ii) aggregate lights per capita and iii) PPP-adjusted real GDP/capita.

partially attributed to changes in satellite sensitivity. Figure 4.2 shows that there is also a general trend around the world for people to leave unlit rural areas for towns and cities.

4.2 Price booms: the 2000's

4.2.1 High oil prices stimulate economic growth in oil-dependent countries

During the period of high oil prices in the 2000s, night-time illumination in non-OECD oil-dependent countries increased significantly relative to non-dependent countries. This is illustrated in Figure 4.3, which shows the log difference in aggregate night-time lights between oil dependent and non-dependent countries, relative to the omitted year of 2002 (coefficient β_t in equation 3.1). The results show that this difference increased steadily during the oil boom, and by the end of the sample the effect on lights in dependent countries is .29 log points. The effect on lights per capita[9] is slightly smaller but still statistically significant, though only at a 10% level at the end of the sample. There was a similar and strongly significant effect on PPP-adjusted real GDP/capita (expenditure based).

[9] Since the elasticity between lights growth and actual economic activity is unknown, comparing lights growth to known population growth may give misleading estimates to the extent that the elasticity differs from 1. We present the results for both lights and lights per capita for completeness.

i. ii.

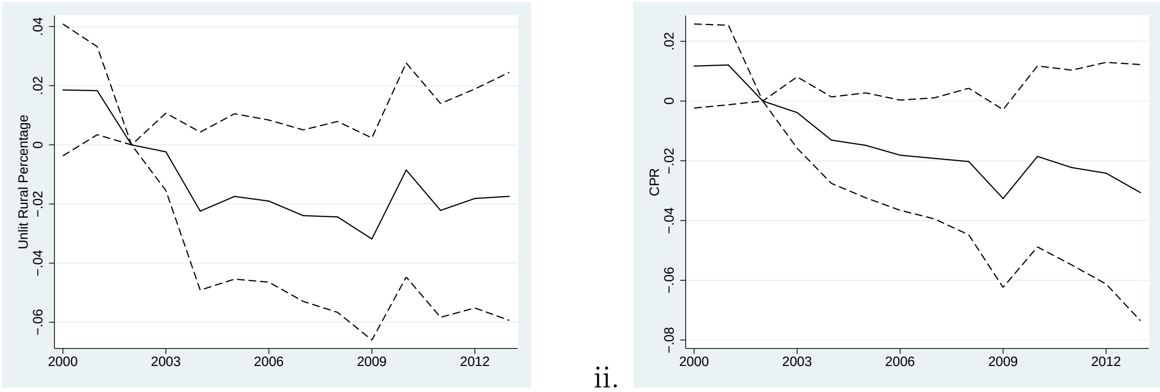

Figure 4.4: Effect of the 2000s price boom on i) the population share living in unlit rural areas, and ii) the Calibrated Poverty Rate.

4.2.2 Economic growth is not shared with the rural poor

While aggregate night-time illumination increased in oil dependent countries relative to non-dependent ones, there is no evidence that it benefitted the rural poor. This is illustrated in Figure 4.4, which shows the share of oil-dependent countries' population living in unlit rural areas, relative to the 2002 base year. We find no significant change or trend in the share of rural poor. A slightly larger though insignificant result is found when using the Calibrated Poverty Rate (which, recall, attempts to measure poverty in non-rural areas as well). However, these results may mask some illumination and de-illumination of cells and movement of people between them, which we address next.

The results shown in Figures 4.3 and 4.4 raise the question: how did light intensity increase so much without significantly reducing rural poverty? To answer this question we define three types of grid cells: A cell is classified as rural if it has a non-zero population but no lights activity as of the initial year 2000. A cell is classified as a town if it is coded as non-urban but has non-zero lights activity as of 2000, and as urban according to the definition from Section 2.3. We then aggregate the lights and population of these three types to the country level and analyze them with specification 3.1.

By the construction of our poverty measure, poverty can be reduced in two ways: lights spreading to previously unlit cells, and population shifts from unlit to lit cells. We investigate both mechanisms. Figure 4.5 shows that lights in initially rural cells (as of 2000) did not grow any faster for oil-dependent than non-dependent countries, but those in towns and cities did. To illustrate this we use two different measures. The first is the "Rural switch percentage", described in Section 4.1. We find no evidence that the oil price boom caused lights to be turned on in rural areas, further confirming that the rural poor did not benefit. The second is the log difference in town and city lights between oil dependent and non-dependent countries, which is similar to the aggregate measure in Figure 4.3. We find that economic activity in both the towns and the cities of oil dependent countries grew faster than in their non-dependent counterparts during this period, with an effect of .32 log points in 2013. In cities the difference was a more modest .14. For both towns and cities the estimates at the end of the period are significant at a 10% level, though when using lights per capita they are significant at a 5% level (not shown).

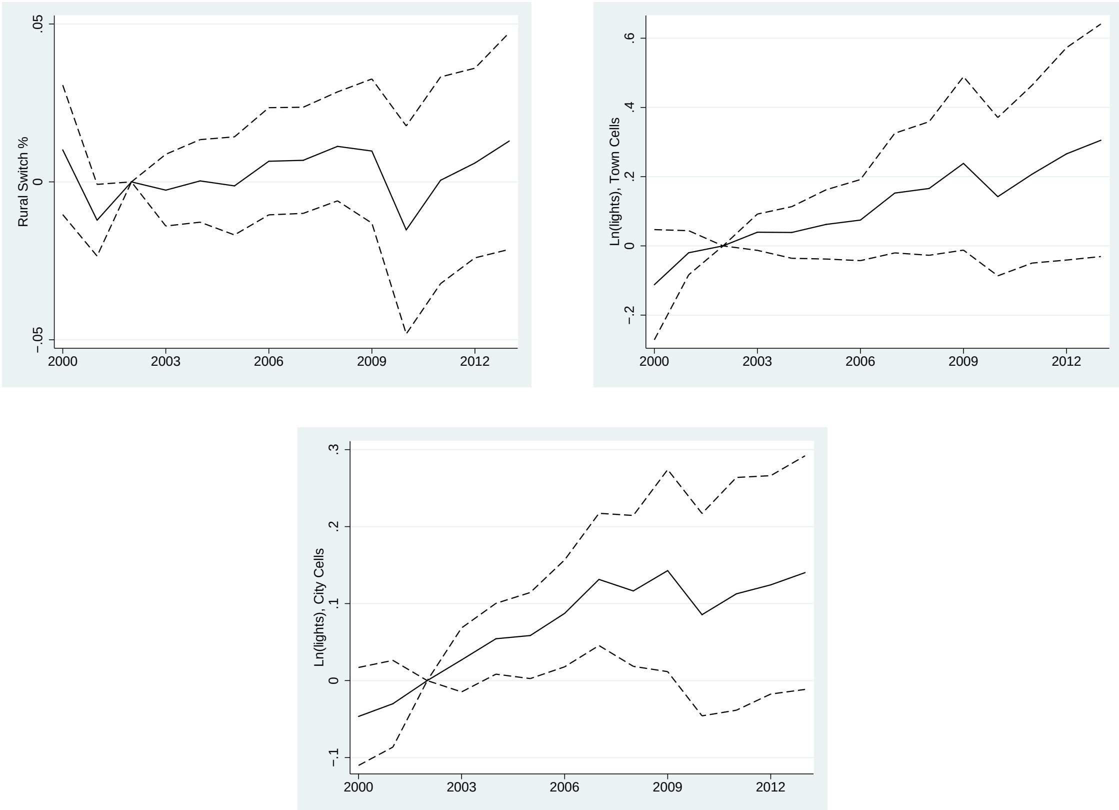

Figure 4.5: Effects of the 2000s oil price boom on illumination in rural areas, towns and cities, based on their 2000 classification.

4.2.3 Oil price spikes do not affect migration

We now consider how the 2000s oil price boom influenced migration between rural and urban areas, and we find no evidence of an effect. While there has been a general 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 4.6 illustrates the difference in the population share in rural areas, towns and cities, between oil dependent and non-dependent countries[10]. In all three instances this difference did not significantly change.

4.3 Quantity booms: giant oil discoveries

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

When a giant oil field is discovered, aggregate night-time lights in the country increase relative to other non-OECD countries after about six years. This is illustrated in Figure 4.7, 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

[10]For population share regressions we do no control for initial population shares since these are not theoretically subject to convergence as output measures are.

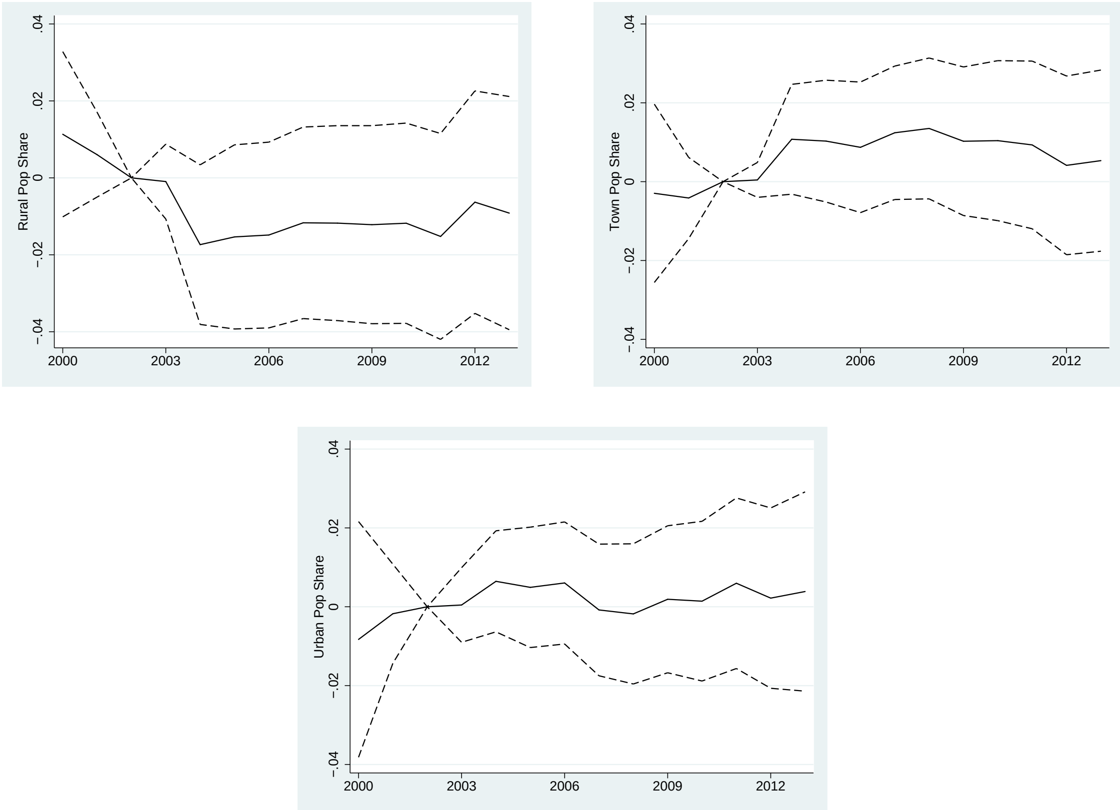

Figure 4.6: Effect of the 2000s oil price boom on the population share in rural areas, towns and cities, based on their 2000 classification.

with net present value of 100% of GDP increases lights and lights per capita by .06 log points, with a similar effect on GDP/capita. The pattern of a slight decrease in lights following discovery following a sharp rise roughly around the time production begins is consistent with the GDP results found in Arezki et al (2015).

4.3.2 This growth is not shared with the rural poor

Figure 4.8 illustrates the population share living in unlit rural areas after an oil discovery (like Figure 4.4). Unlike the 2000s price boom we find that oil discoveries do reduce the share of the population living in unlit rural areas by around 1%, with a lag of nine years. Discoveries have a similar effect on the CPR. However it is unclear whether this is because previously unlit areas gain lights, or because people leave the unlit areas for better opportunities elsewhere. We disentangle this next.

Figure 4.9 shows that while cities and towns gain illumination after an oil discovery, unlit rural areas do not. We start by again classifying cells according to their city/town/rural type as of the year 2000. For rural areas, we find no evidence that they become illuminated after an oil discovery. If anything, rural lights switch on more in countries that don't have oil discoveries, though the effect is generally insignificant. It is important to note that this result does not comment on the local effects of oil wells (see the review by Cust and Poelhekke, 2014), as it is making cross-country rather than within-country comparisons.

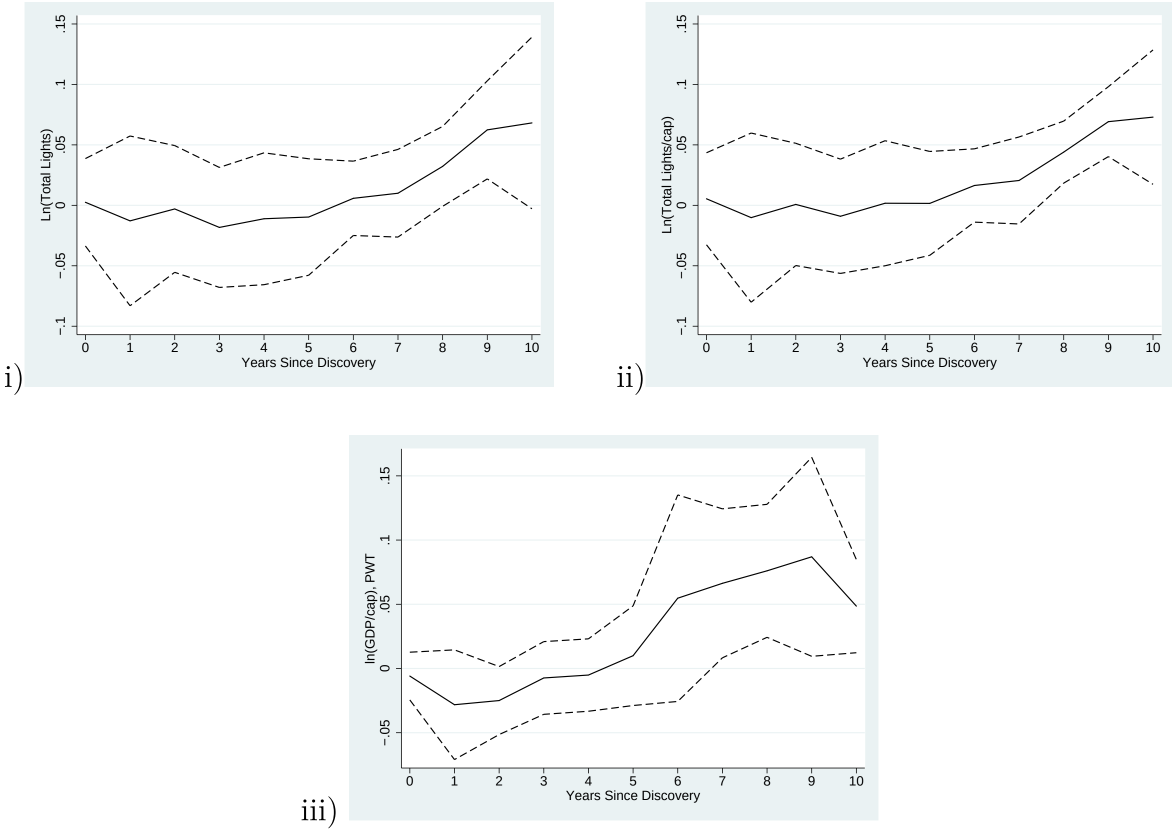

Figure 4.7: Effect of oil discoveries on i) aggregate (log) lights, ii) log aggregate lights/capita and iii) PPP-adjusted real GDP/capita.

For towns and cities, both see illumination trend upwards six years after oil is discovered. This trend becomes significant after eight years (though not in year 10 due to large increase in the standard errors, although the trend remains consistent), confirming that the gain in economic activity from oil discoveries is concentrated in towns and cities.

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

We have just established that the share of oil dependent countries' populations living in unlit rural areas falls after an oil discovery, but this is not due to unlit rural areas becoming illuminated. The implication is that it must be because people migrate away from rural areas. Figure 4.10 confirms this. It shows that the rural population share in oil dependent countries trends downwards (relative to non-dependent countries) after an oil discovery, and is approximately 0.6 percentage points lower after ten years. It also shows that this migration leads to a similar increase in the population of towns, but not cities. As with the price boom we don't find evidence that quantity booms encourage urbanization.

4.4 Illumination Mechanisms

The results so far have been for unlit rural areas in aggregate. We have shown that while rural areas are illuminating around the world, they are not doing so any faster in oil-

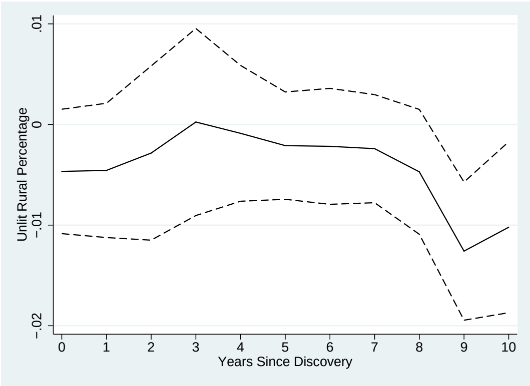

i.

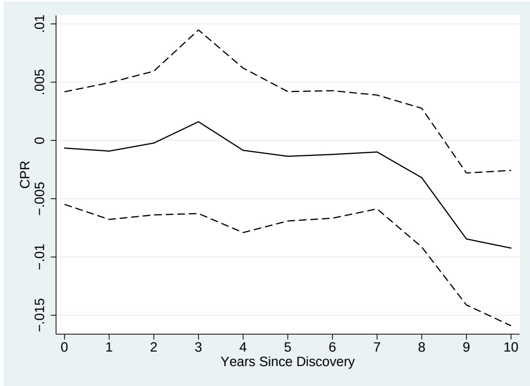

ii.

Figure 4.8: Effect of oil discoveries on i) the population share living in unlit rural areas and ii) the Calibrated Poverty Rate (CPR).

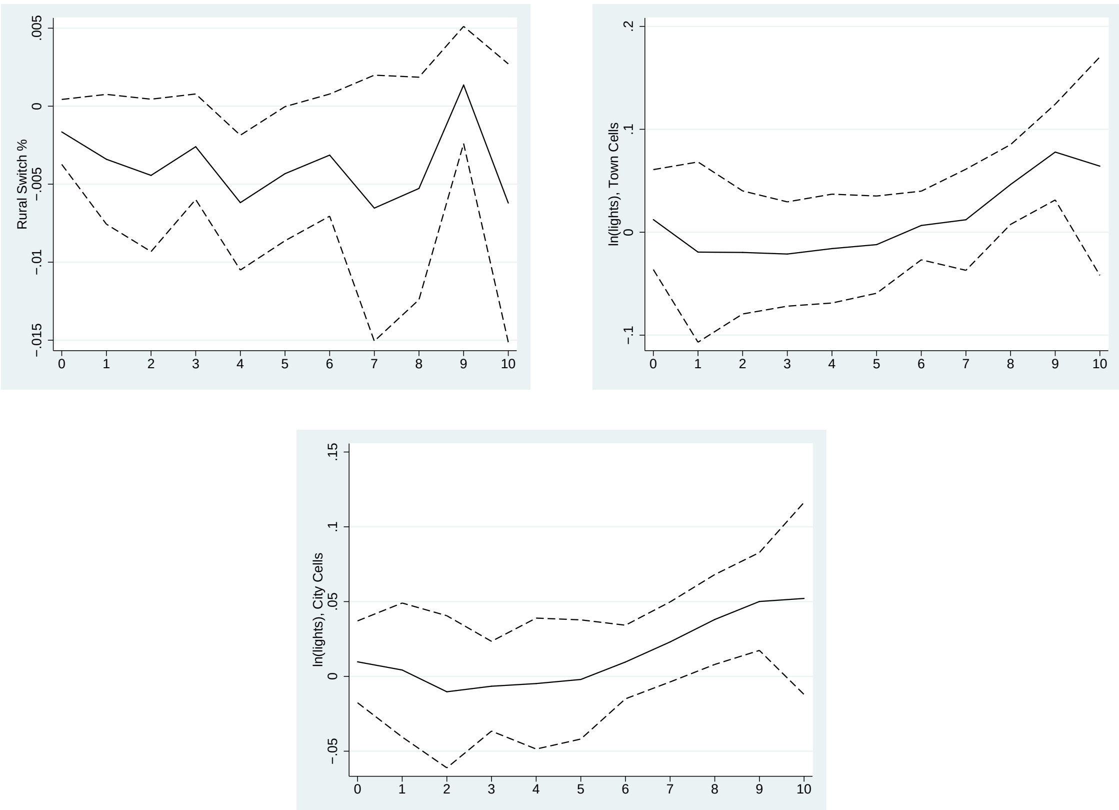

Figure 4.9: Effect of giant oil discoveries on the illumination of rural areas, towns and cities, based on their 2000 classification.

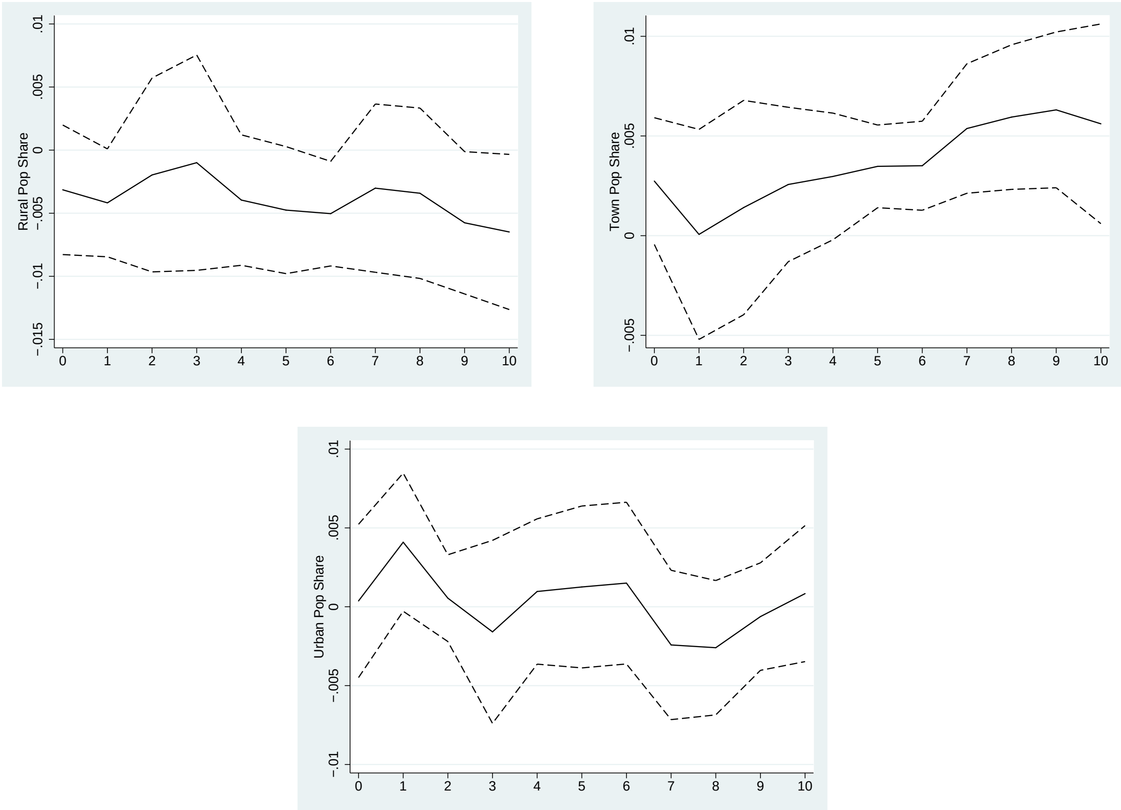

Figure 4.10: Effect of giant oil discoveries on the population share in rural areas, towns and cities, based on their 2000 classification.

dependent countries. However, these average results might be hiding some details in the darkness. To understand these details we use a hazard model to study what causes an unlit but inhabited rural cell to become lit in any given year. 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 evidence that the first three of these mechanisms are more active in oil-dependent countries. We do find that oil-dependent countries are less efficient at converting growth to poverty reduction.

We model the probability of unlit rural cells becoming lit in a given year using the following hazard rate specification,

$$I^{RS}_{cit} = \alpha + \beta_1 X_{cit} + \beta_2 D_i X_{cit} + \lambda_t + \varphi_i + \epsilon_{c,i,t} \tag{4.1}$$

where I^{RS}_{cit} 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

	(1)
	switchon
Non-adjacent	-0.031***
	(0.005)
Non-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%, 1%, .1%, respectively.

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

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.

We hypothesize that being next to a lit cell generally represents being near an existing electricity network, and such cells switching on represents public investment in the grid. Therefore, being next to a lit cell should increase the probability of illumination, because other cells would presumably require more expensive off-grid generators. We find that being next to a lit cell does increase the probability of an unlit rural cell switching on by 3 percentage points. Appendix Figure C.1 shows the cumulative effect of this mechanism (not controlling for the other three in equation 4.1).[11] Over the decade from 2003 adjacent cells became ~15 percentage points more illuminated than non-adjacent cells. 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.

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 C.2). 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 C.3). 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.

4.5 Robustness

To check the robustness of our main results we try four alternative specifications, with the results shown in Appendix B. Each focuses on the price boom specification in equation

[11]Graphs 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.

3.1. 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 B.1). The estimates are more or less the same, but with larger standard errors since 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 either measure of poverty.

In the main specification we use a dependence dummy indicator because it yields a simple and transparent graphical representation of effects, and does not rely on functional forms of dependence. However, some information is lost in the binary classification. In Appendix Figure B.2, we replace the dependence indicator variable D_i in equation 3.1 with the continuous variable $Rents_i$, which measures average oil rents as a share of GDP from 2000-2012. This specification implies a linear relationship with dependence. We also do the same but drop Equatorial Guinea (EG), an outlier where rents account for 75% of GDP during the period and may skew the results. As shown in Figure B.2, if oil rents account for 100% of GDP, then the boom would have caused lights to increase by 1.34 log points (with EG) or .92 log points (without), with smaller but significant effects for lights per capita. Our poverty measures fall by .13 log points (again if oil is 100% of GDP) in this specification, driven largely by Equatorial Guinea. Excluding that country we find moderate but insignificant reductions in poverty. This reduction in turn is driven almost entirely by Gabon and the Republic of Congo, two other very high-dependence countries that saw substantial reductions in unlit percentage due to migration away from rural areas in the early part of the sample, after which the unlit rural percentage is nearly trendless. Continuous oil dependence has a small and insignificant impact on rural cells becoming lit, with or without Equatorial Guinea (not shown), again indicating that oil revenues are not reaching rural areas but may in somce cases encourage migration (similarly to the results using discoveries).

The fourth robustness check replaces the year fixed effects and regional trends in equation 3.1 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.

5 Conclusion

This paper offers an annual global panel study of extreme rural poverty. We construct a world-wide measure of rural poverty using two high-resolution spatial datasets: on night-time illumination as a proxy for economic activity, and population. We combine this with data on oil prices and giant oil discoveries to answer the question, does the income from oil booms benefit the very poor?

We start out by describing some stylised facts about rural poverty and urbanization. The share of people living in rural darkness is falling world-wide, both because previously unlit areas are gaining electricity and because people are moving from rural areas to existing towns and cities. Electrification occurs faster in areas that are close to the existing grid, close to the capital city, and have higher population.

Then we present three main results, which are similar for both price and quantity booms. First, oil booms do increase aggregate economic activity - proxied by night-time lights. For periods of high oil prices this happens immediately, while for oil discoveries it begins after a six year lag, corresponding to the typical lag between discovery and production. Second, this activity is confined to cities and towns, and does not benefit the rural poor. We find no evidence that unlit rural areas become more illuminated, whether close to the grid or the capital, while lights in cities and towns rise significantly. Third, oil discoveries encourage a small amount of migration from unlit rural areas to towns (though not cities), but oil price booms do not. Thus, the evidence suggests that the only way the rural poor benefit from oil discoveries is to leave these regions to pursue opportunities elsewhere.

This work suggests a wide range of extensions, exploiting both our new measure of rural poverty and the global panel method we employ. Such work may include studying other determinants of extreme poverty and the success of programs designed to alleviate it.

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 (2015). News shocks in open economies: Evidence from giant oil discoveries. Technical report, National Bureau of Economic Research.

[5] Atkinson, A. B. (1970). On the measurement of inequality. *Journal of economic theory 2*(3), 244–263.

[6] Atkinson, A. B. and A. Leigh (2007). The distribution of top incomes in australia*. *Economic Record 83*(262), 247–261.

[7] Atkinson, A. B., T. Piketty, and E. Saez (2009). Top incomes in the long run of history. Technical report, National Bureau of Economic Research.

[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] 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.

[11] Buccellato, T. and T. Mickiewicz (2009). Oil and gas: a blessing for the few. hydrocarbons and inequality within regions in russia. *Europe-Asia Studies 61*(3), 385–407.

[12] Cavalcanti, T., D. Da Mata, and F. Toscani (2014). Winning the oil lottery: The impact of natural resource extraction on growth. *Unpublished manuscript, 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] Collier, P. and B. Goderis (2012). Commodity prices and growth: An empirical investigation. *European Economic Review 56*(6), 1241–1260.

[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 microeconometric 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., J.-P. Muller, and J. G. Morley (2006). Mapping regional economic activity from night-time light satellite imagery. *Ecological Economics 57*(1), 75–92.

[22] Elvidge, C. D., K. E. Baugh, E. A. Kihn, H. W. Kroehl, and E. R. Davis (1997). Mapping city lights with nighttime data from the dmsp operational linescan system. *Photogrammetric Engineering and Remote Sensing 63*(6), 727–734.

[23] Feenstra, R. C., R. Inklaar, and M. Timmer (2013). The next generation of the penn world table. Technical report, National Bureau of Economic Research.

[24] Goderis, B. and S. W. Malone (2011). Natural resource booms and inequality: Theory and evidence. *The Scandinavian Journal of Economics 113*(2), 388–417.

[25] Gollin, D., R. Jedwab, and D. Vollrath (2013). Urbanization with and without industrialization. Technical report.

[26] Gylfason, T. and G. Zoega (2003). Education, social equality and economic growth: a view of the landscape. *CESifo Economic Studies 49*(4), 557–579.

[27] 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.

[28] Henderson, J. V., A. Storeygard, and D. N. Weil (2011). A bright idea for measuring economic growth. *The American economic review 101*(3), 194.

[29] Henderson, J. V., A. Storeygard, and D. N. Weil (2012). Measuring economic growth from outer space. *The American Economic Review*, 994–1028.

[30] Hodler, R. and P. A. Raschky (2014). Regional favoritism*. *The Quarterly Journal of Economics*, qju004.

[31] 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.

[32] 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.

[33] James, A. (2015). The resource curse: A statistical mirage? *Journal of Development Economics 114*, 55–63.

[34] 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.*

[35] Jedwab, R. and A. Moradi (2015). The permanent effects of transportation revolutions in poor countries: Evidence from africa. *Review of Economics and Statistics.*

[36] 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.

[37] Kuznets, S. (1937). Appendices to" national income and capital formation, 1919-1935". In *National Income and Capital Formation, 1919-1935*, pp. 61–90. NBER.

[38] Kuznets, S. (1953). International differences in income levels: Reflections on their causes. *Economic Development and Cultural Change*, 3–26.

[39] Kuznets, S. (1955). Economic growth and income inequality. *The American economic review*, 1–28.

[40] Kuznets, S., L. Epstein, and E. Jenks (1941). National income and its composition, 1919-1938, volume i. *NBER Books.*

[41] 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.

[42] Lei, Y.-H. and G. Michaels (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics 110*, 139–157.

[43] Leigh, A. (2009). Top incomes. *The Oxford Handbook of Economic Inequality*, 150–176.

[44] Michalopoulos, S. and E. Papaioannou (2013). Pre-colonial ethnic institutions and contemporary african development. *Econometrica 81*(1), 113–152.

[45] Pfeffermann, G. and R. Webb (1983). Poverty and income distribution in brazil. *Review of Income and Wealth 29*(2), 101–124.

[46] Pinkovskiy, M. and X. Sala-i Martin (2009). Parametric estimations of the world distribution of income. Technical report, National Bureau of Economic Research.

[47] Pinkovskiy, M. and X. Sala-i Martin (2014). Lights, camera,... income!: Estimating poverty using national accounts, survey means, and lights. Technical report, National Bureau of Economic Research.

[48] Ravallion, M. (2003). The debate on globalization, poverty and inequality: why measurement matters. *International Affairs 79*(4), 739–753.

[49] Ross, M. L. (2007). How mineral-rich states can reduce inequality. *Escaping the resource curse 23775*.

[50] Sachs, J. D. and A. M. Warner (1995). Natural resource abundance and economic growth. Technical report, National Bureau of Economic Research.

[51] Sachs, J. D. and A. M. Warner (2001). The curse of natural resources. *European economic review 45*(4), 827–838.

[52] Sala-i Martin, X. (2006). The world distribution of income: falling poverty and... convergence, period. *The Quarterly Journal of Economics*, 351–397.

[53] Smith, B. (2015). The resource curse exorcised: Evidence from a panel of countries. *Journal of Development Economics 116*, 57–73.

[54] Stone, R. and G. Croft-Murray (1959). *Social accounting and economic models*. Bowes & Bowes London.

[55] Stone, R. and O. E. de Cooperació Econòmica (1961). *Input-output and national accounts*. Organisation for European Economic Co-operation Paris.

[56] Subramanian, S. and A. Deaton (1996). The demand for food and calories. *Journal of political economy*, 133–162.

[57] Van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature*, 366–420.

[58] Wills, S. (2015). Optimal monetary responses to oil discoveries. *OxCarre Research Paper 121*.

[59] World Bank (2013). Global monitoring report 2013: Rural-urban dynamics and the millennium development goals. Technical report, The World Bank Group.

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 Robustness

Figure B.1: Effect of 2000s oil price boom on key variables, excluding OPEC.

C Illumination Mechanisms

In addition to the hazard rate model described in Section 4.4, 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,

$$RuralSwitch\%_{c,i,t} = \alpha + \sum_{t=2000}^{2013} \beta_t(\lambda_t \times X_c) + \lambda_t + \varphi_i + \epsilon_{c,i,t} \tag{C.1}$$

Where X_c is one of the dependent variables described in Section 4.4. 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.[12]

[12]This takes, for example, the difference in rural illumination between cells that are/are not adjacent

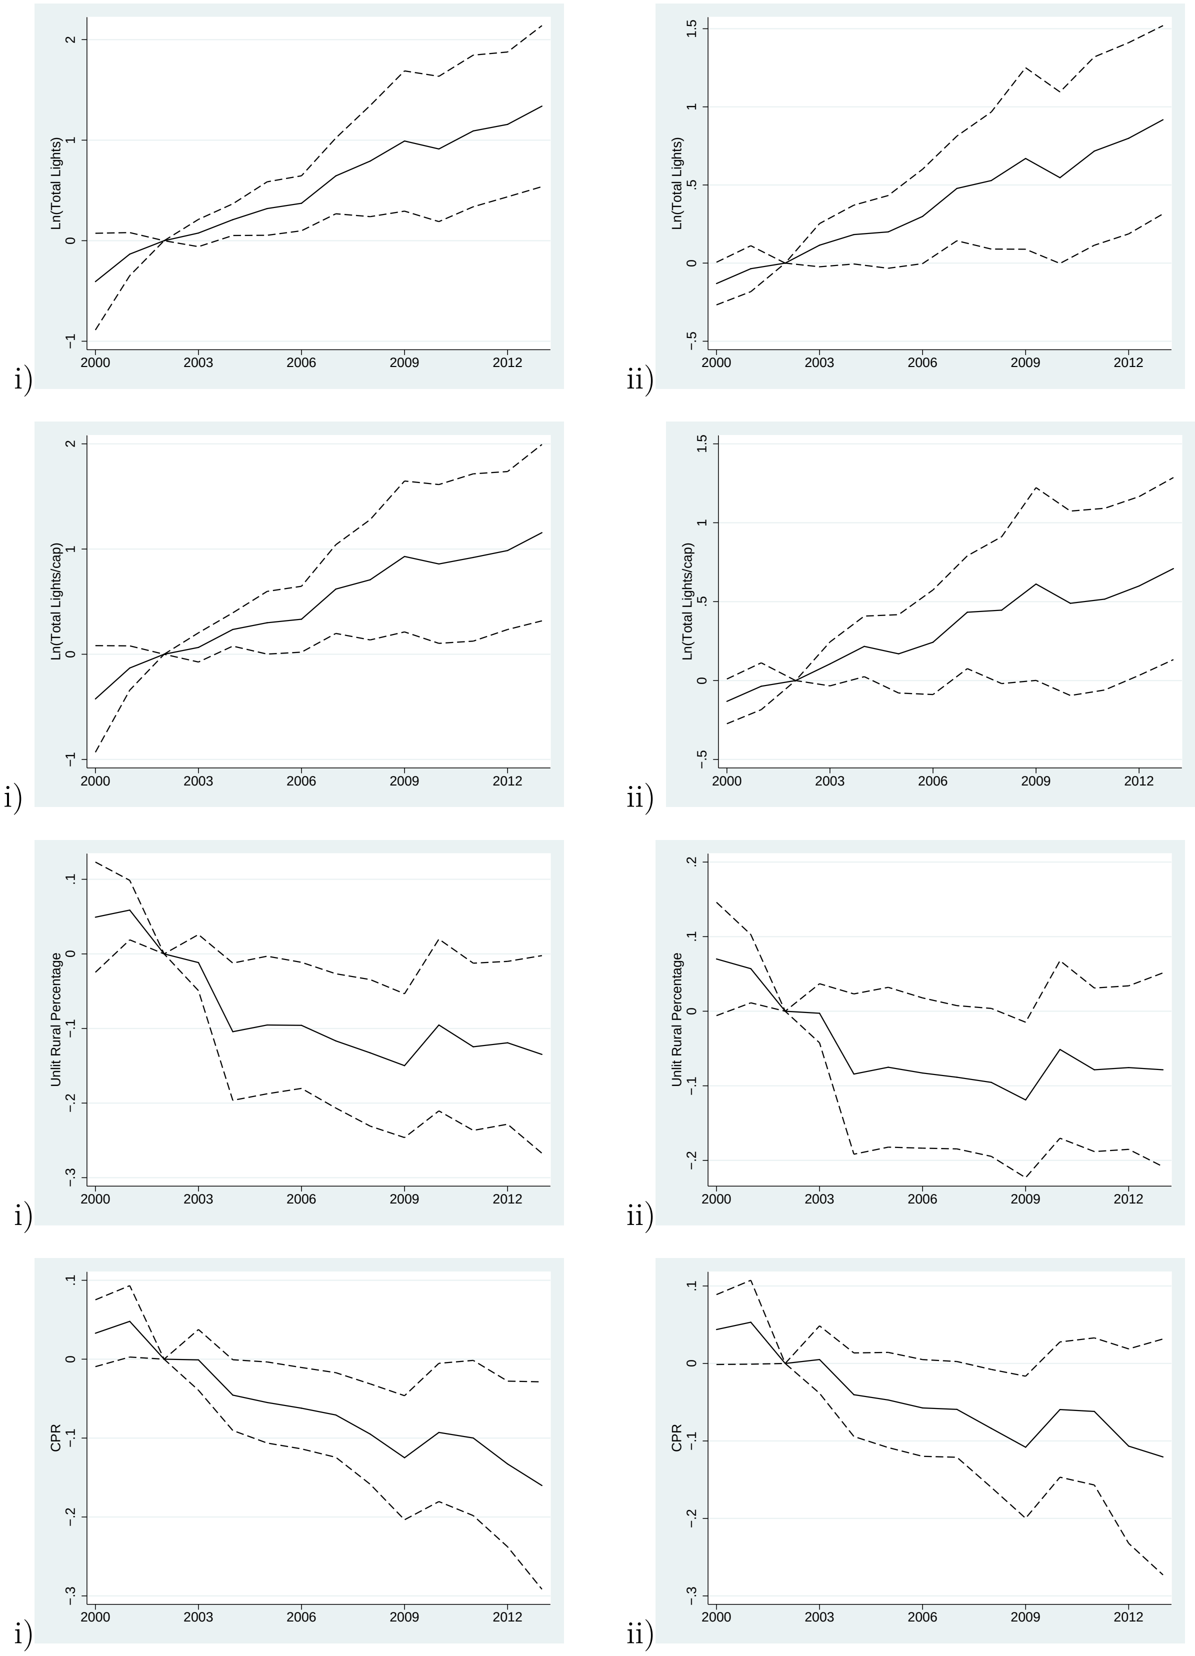

Figure B.2: Effect of 2000s price boom on key variables, using a continuous measure of resource dependence i) with and ii) without Equatorial Guinea.

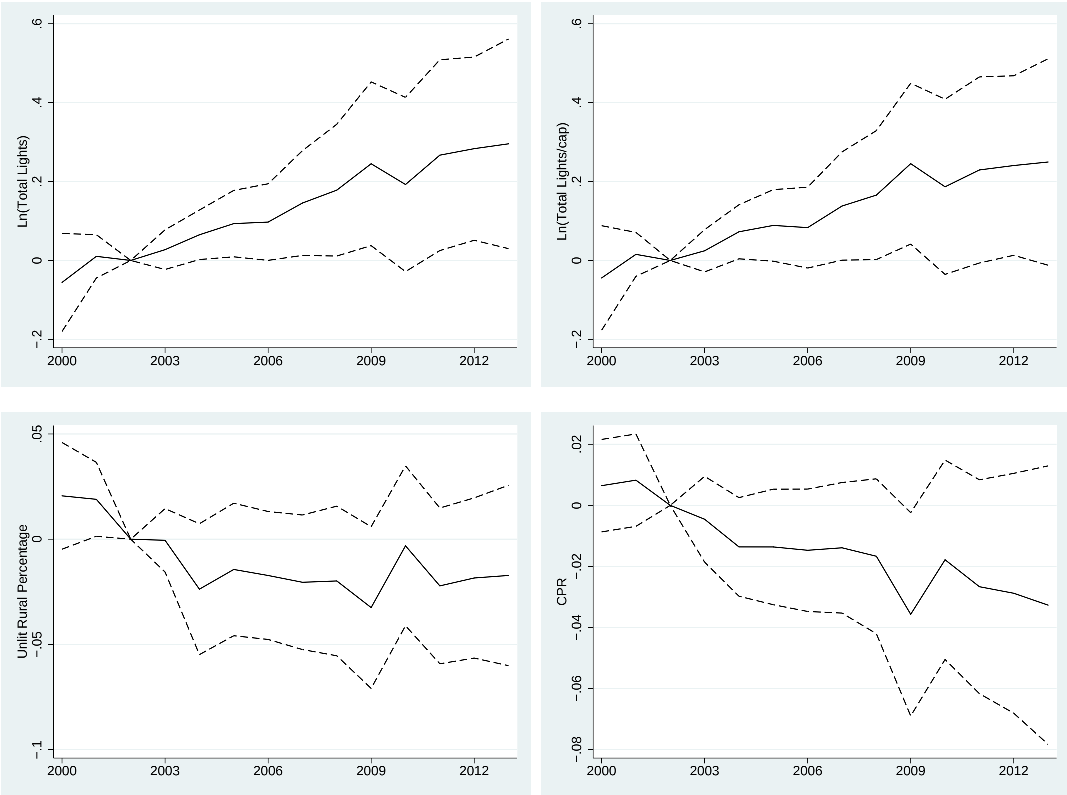

Figure B.3: Effect of 2000s oil price boom on key variables, using region-year instead of year fixed effects.

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 coeffiecients are shown in Figure C.1, iii.

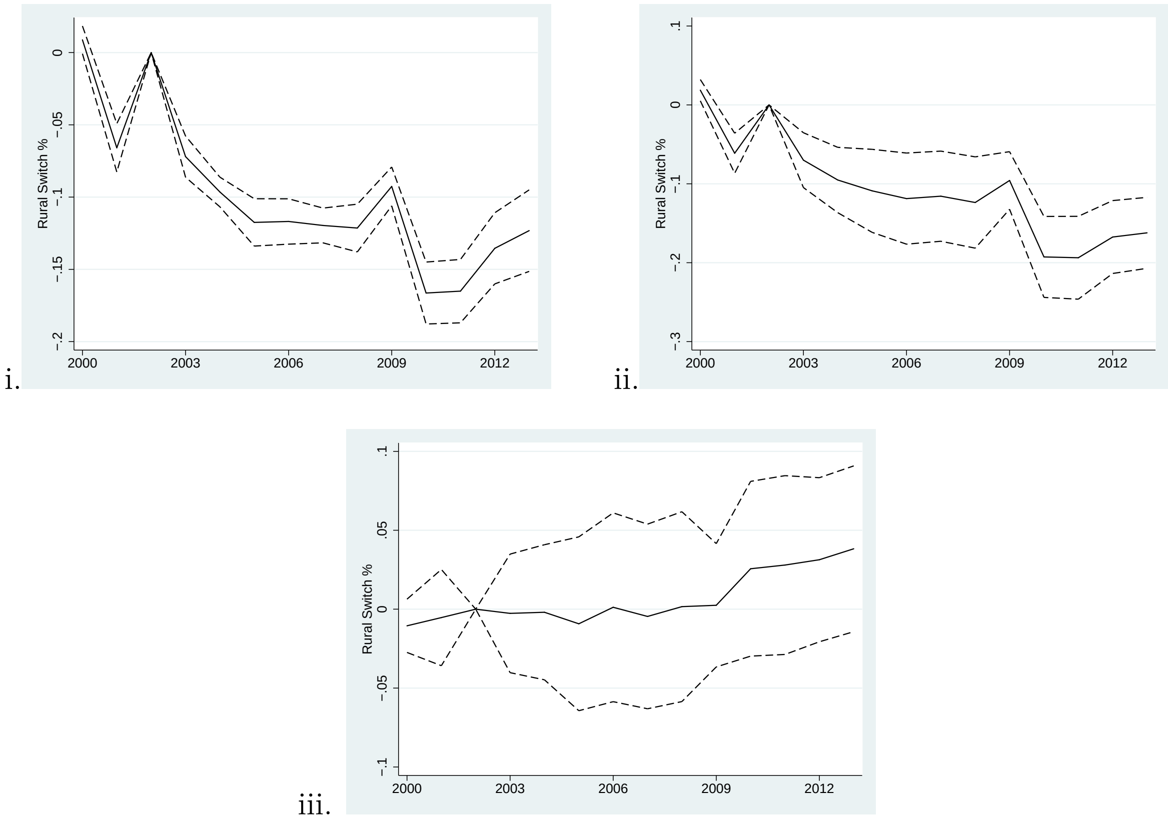

Figure C.1: The difference in switch-on percentage for rural areas that are adjacent vs not-adjacent to lit squares density in i) oil-dependent countries, ii) non-dependent countries and iii) the difference between i) and ii).

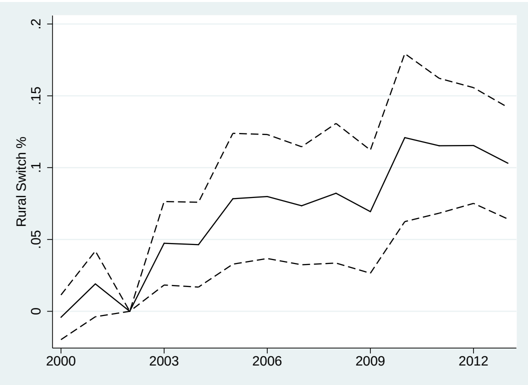

i.

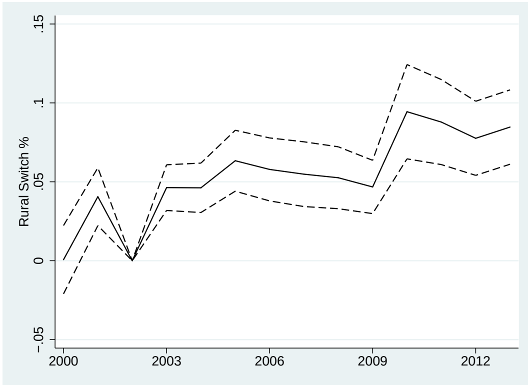

ii.

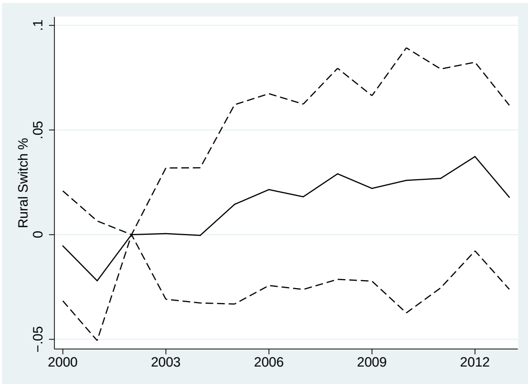

iii.

Figure C.2: The difference in switch-on percentage between rural areas that are 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).

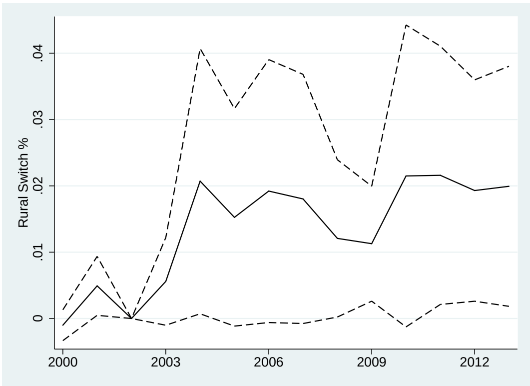

i.

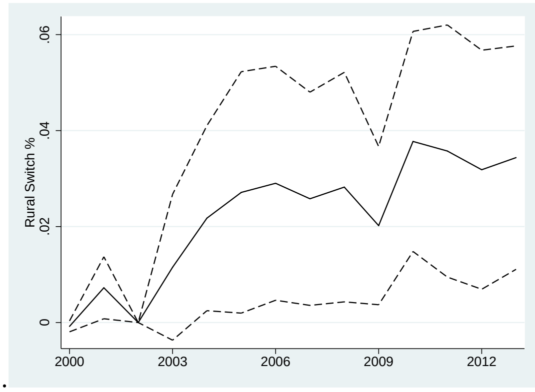

ii.

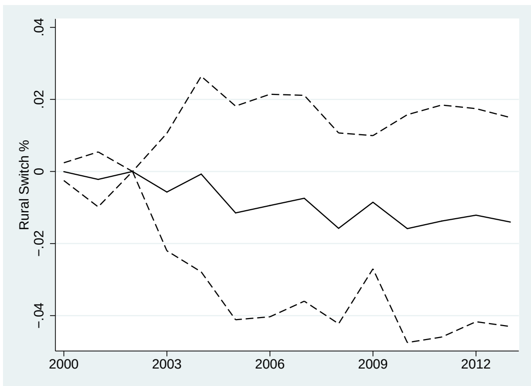

iii.

Figure C.3: 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).