

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 165

The Resource Curse Exorcised: Evidence from a Panel of Countries

Brock Smith

OxCarre

The Resource Curse Exorcised: Evidence from a Panel of Countries

Brock Smith

Oxford Centre for the Analysis of Resource Rich Economies

Department of Economics

Manor Road Building

Manor Road

Oxford, UK

OX1 3UQ

ABSTRACT

This paper evaluates the impact of major natural resource discoveries since 1950 on GDP per capita and its proximate causes. Using panel fixed-effects estimation and resource discoveries in countries that were not previously resource-rich as a plausibly exogenous source of variation, I find a positive effect on GDP per capita levels following resource exploitation that persists in the long term. Results vary significantly between OECD and non-OECD treatment countries, with effects concentrated within the non-OECD group. I further test GDP effects with synthetic control analysis on each individual treated country, yielding results consistent with the average effects found with the fixed-effects model. Productivity, capital formation and education were also positively affected by resource discovery, while growth accounting analysis suggests productivity gains were a major distinguishing factor in GDP effects.

I would like to thank Giovanni Peri, Ann Stevens, Hilary Hoynes, Christopher Meissner, Douglas Miller, and Alan Taylor for helpful feedback. Support from the BP funded Oxford Centre for the Analysis of Resource Rich Economies is gratefully acknowledged.

Keywords: Natural resource curse; economic growth; growth regressions; growth accounting; oil.

1 Introduction

Following the seminal work of Sachs & Warner (1995), a near-consensus formed supporting the existence of a “resource curse”, the counter-intuitive finding that countries rich in natural resources tend to experience slower growth. Sachs and Warner used a simple cross-sectional design to find that countries with a higher ratio of commodity exports to GDP in 1970 saw slower average growth over the next 20 years.

Although recent studies have called into question the existence of a resource curse (see the following section, and van der Ploeg, 2011 for a survey), much of the literature on natural resources and growth has more or less taken the Sachs & Warner (1995) result as given and extended their design in an attempt to pinpoint the mechanisms through which natural resources harm growth, and to find which factors cause a resource curse or blessing to materialize. One commonly cited culprit is the so-called “Dutch disease” (a term coined in 1977 after the natural gas boom in the Netherlands), whereby resource exports increase exchange rates, reducing the competitiveness of exporters in the manufacturing sector (Sachs & Warner 1995, Gylfason, et al 1999, Sala-i-Martin & Subramanian 2003). Others have explored the link between natural resources and quality of institutions. One form of the institutions-driven resource curse is that resource discovery subsequently weakens institutions and thus growth (Ross 2001, Leite & Weidmann 2002). Another form treats institutions as exogenous to resource wealth, and the interaction between resources and institutions explains the divergent outcomes of resource-rich countries (Robinson et al 2006, Mehlum et al 2006, Sarr et al 2011). Caselli & Michaels (2013) find that oil-rich municipalities in Brazil report significantly higher revenues and spending, but with little to no benefit to the wider population, suggesting corruption by municipal officials¹. Other papers have argued that low levels of human capital (Gylfason 2001, Papyrakis & Gerlagh 2004, Ortega & Gregorio 2005), lack of investment (Atkinson & Hamilton 2004), and increased risk of civil war (Collier & Hoeffler 1998) also play a role.

The majority of the empirical literature on the resource curse suffers from two significant identification flaws. First, the most commonly used measures of “resource wealth” are more accurately described as resource dependence. Sachs & Warner (1995) and its various extensions model resource wealth as resource exports’ share of GDP. However, as Brunnschweiler & Bulte (2008) and Alexeev & Conrad (2009) point out, using resource dependence creates an endogeneity problem; poor growth resulting from structural factors independent of resource wealth will cause a lower GDP, and thus a higher share of resources in GDP. This creates a possible omitted variable bias, where whatever unobserved structural factors cause

¹Other papers studying the corruption channel include Torvik (2002) and Papyrakis (2004).

high dependence will likely impact subsequent performance. A more appropriate measure is what Brunnschweiler & Bulte (2008) call resource abundance, which measures resource wealth per capita, independent of overall GDP. To the extent that population is exogenous to growth, this measure should be free of the kind of endogeneity described above (although the use of resource abundance in a cross-sectional setting is still problematic, as discussed in the following section).

The second flaw is the use of cross-sectional data. One of the more intuitive a priori explanations of the resource curse is the historical path-dependence story of the sort espoused by Acemoglu, et al (2001). By this reasoning, early institutions were partially determined by known resource abundance, as pernicious extractive regimes were installed in resource-abundant colonies, and these bad institutions then persisted into the post-colonial era. If this mechanism is present, then a cross-sectional sample of modern data, like that used in Sachs & Warner (1995), cannot separate such “colonial baggage” from the present-day effects of resource wealth. Presumably, this mechanism has not been explored in the literature because it does not lend itself to empirical testing. Resource wealth cannot be used as an instrument for institutions because it also directly affects present-day growth and is therefore non-excludable.

In this paper, I eliminate colonial baggage from estimates of the resource effect by only considering relatively recent discoveries. There have been sufficient discoveries of oil (and one discovery each of diamonds and natural gas) in previously non-resource-rich countries in the past six decades to implement a quasi-experimental design with plausibly exogenous resource shocks and a clearly defined treatment group. It is advantageous that nearly all recent discoveries have been in oil, the commodity most commonly associated with the curse. Wick & Bulte (2009) argue that oil’s high spatial concentration makes it susceptible to control by one party, promoting inequality and civil strife.

This paper estimates the effect of the transition to resource abundance from non-resource abundance on subsequent growth in GDP per capita. Using panel difference-in-differences, event study and synthetic control designs, I compare countries that have become resource-rich since 1950 with countries that have remained resource-poor (countries that were resource-rich already in 1950 are dropped from the analysis—see Appendix C for the list of sample countries). One concern with this design is the possibility that resource discovery is not exogenous; David & Wright (1997) and Bohn & Deacon (2000) have argued that discovery may be more likely in more democratic countries or those with better institutions. I test this proposition empirically in Section 3 and find no evidence that it is true. However, even if it were the case, the difference-in-differences specification controls for structural differences in institutions, so any institutional bias would have to arise from institutions independently

changing after discovery, and such that the direction of change is correlated with being a discovery country.

I find that newly resource-rich countries on average experience a large short-term boost in GDP growth and non-negative long-run effects on growth, resulting in a large positive effect on long-run GDP levels. I find little to no pre-exploitation trends in the outcomes, further supporting the exogeneity of the treatment. The finding that resource discovery appears to have a long-run level effect on GDP, and no long-run growth effect, is consistent with a simple Solow model in which there is a temporary shock to productivity growth. This attracts additional investment, which further enhances growth during a transitional period until the capital stock per worker settles at a higher level and normal growth resumes. In an endogenous growth setting, the result could be thought of as analogous to the model proposed in Jones (1995), in which an increase in the share of output in R & D (which could be thought of as drilling infrastructure in this case) results in a permanent level effect but no long-term growth effect.

Further, I find that the positive GDP effects are concentrated in developing countries, with small and insignificant effects for developed countries. This runs counter to much of the literature, which argues that countries with better institutions experience a more positive effect from natural resources. The main reason I find no effects for developed countries is that I am making within-region comparisons (by including region-year fixed effects in the main specification); while developed treatment countries have not performed poorly over the period studied, this is likely due to many factors besides natural resources, as their regional counterparts have performed similarly well.

Moving beyond simple difference-in-differences analysis, I use the synthetic control methodology developed by Abadie & Gardeazabal (2003) and Abadie et al (2010) for each discovery country. This method uses a data-driven algorithm to find a weighted combination of control countries that best replicates the pre-treatment behavior of a single treatment country. This is a useful extension of the analysis in this paper for two reasons: it provides a further robustness check by evaluating performance against an alternative counterfactual, and also reveals the heterogeneity of treatment outcomes by country, rather than just a single average effect. While the synthetic control results do reveal a fairly wide range of individual outcomes, they are consistent with the average positive effects found with the difference-in-differences model, and also with the differing outcomes between developed and developing countries.

Finally, I evaluate the effect of resource discoveries on proximate causes of growth. I find that capital formation, productivity, and education were all positively affected by resources. I perform a growth accounting exercise to estimate the contribution of each

of these three factors, and compare them to control countries during the relevant period. For developing countries, the major distinguishing factor is productivity growth, which was strong in treated countries and negative for the average non-treated country, implying that resource discoveries provided a high-return sector during a period when other productivity-enhancing opportunities were scarce.

The finding of positive growth effects from resource discoveries does not necessarily contradict the cross-sectional Sachs-Warner result. Because I am focusing on countries that became resource-rich since 1950, if both results were assumed to be valid it would imply that the negative Sachs-Warner result is driven primarily by countries that were resource-rich stretching back to the colonial era, whereas for post-colonial discoveries the curse has been lifted (at least in terms of GDP per capita).

This paper contributes to the literature in the following ways: first, it is to my knowledge the first paper to use a quasi-experimental, treatment-control approach to the resource curse question in a cross-country setting, and provides a more plausible test of causality for the effect of natural resources on GDP per capita than has been heretofore performed. Second, apart from Mideksa (2013), which focuses on Norway, this paper is also the first to my knowledge to study the resource curse using the synthetic control method, which allows for causal analysis for many individual countries. Third, it is the first to empirically evaluate by direct observation both the short and long-run effects of resource discoveries on growth (the closest to my knowledge is Collier & Goderis 2012, which uses an error correction model to estimate the long-run effect of resource price changes). This is especially important since many of the proposed resource curse mechanisms, such as deteriorating institutional quality, could take many years to materialize. Fourth, it is the first to my knowledge to carry out a growth accounting analysis focusing on resource-rich economies, and to evaluate the impact of resources on different components of GDP growth (capital, education, TFP).

The rest of this paper proceeds as follows: the following section briefly reviews recent advances in the empirical resource curse literature. Section 3 gives a brief historical overview of the oil industry and exploration. Section 4 outlines my empirical design. Section 5 presents and discusses the main results. Section 6 evaluates components of GDP growth and presents the growth accounting exercise. Section 7 concludes.

2 Recent Literature

A number of recent studies have challenged the finding that resources harm growth, primarily by using alternative measures of resource abundance rather than the resource share of GDP. Brunnschweiler & Bulte (2008) examines the relationship between 1970-2000 average growth

rates and “subsoil assets” per capita measured in 1994 and 2000, and finds a positive effect². Alexeev & Conrad (2009) find a positive association between hydrocarbon deposits per capita in 1993 (or alternatively, the value of oil production per capita in 2000) and the level of GDP in 2000. These papers rely on the argument that natural resource endowments are exogenous, geographic variables. While this is compelling, van der Ploeg & Poelhekke (2010) point out that the available resource abundance measures are closely associated with current resource rents and thus endogenous to growth and income. They further critique the World Bank’s estimates of subsoil wealth and argue that it is more of a one-off estimate of natural capital and net adjusted saving, but not a suitable measure of actual subsoil wealth. A related argument is that what is truly being measured is *known* resource endowments (or an estimate based on known endowments), which depend on how thoroughly a given country has been prospected, which in turn may be affected by the country’s wealth and institutions³. While similar concerns could be raised for the initial discovery of resources as this paper uses, the difference-in-differences design controls for time-invariant factors present before and after discovery.

A few recent studies have also incorporated oil discoveries into their specifications. Cotet & Tsui (2013) argue that for most oil-producing countries, the most significant oil discoveries are concentrated over a few years. They evaluate the relationship between growth and health measures and estimated oil endowments over different periods of time after this “peak discovery period”, and find positive effects. However, this method faces the same causal uncertainty as described above resulting from estimated oil endowments. Tsui (2011) uses a similar analysis and finds that countries that discover more oil (with oil discovered instrumented by estimated endowments) become less democratic in the following decades. Cotet & Tsui (2013) additionally exploit data on the number of exploratory wells dug in a given year and find that civil conflict is largely uncorrelated with oil wealth per capita.

A number of papers have used panel methods to study the relationship between resources and political outcomes. Bruckner et al (2012) and Caselli and Tesei (2011) both use panel data to estimate the effect of income shocks driven by commodity price fluctuations on democratic institutions in commodity-exporting countries (reaching different conclusions)⁴. To my knowledge, few other papers have used panel data to examine the relationship between growth and natural resources. Collier & Goderis (2009) use a panel cointegration approach to estimate a specified long-run equilibrium relationship between growth and resource-export

²Lederman & Maloney (2007) take a similar approach, though using different measures of abundance, and also find positive effects.

³Michaels (2010) uses a similar approach to study long-run outcomes of United States counties. This paper makes a convincing causal argument since the US has been extensively prospected.

⁴See also Aslaksen (2010) and Haber & Menaldo (2011) for panel studies on oil and democracy.

prices, finding a negative long-run effect of price increases. Cotet & Tsui (2012) includes a panel specification that evaluates the effect of changes in oil rents on different outcomes over 5-year periods, finding no significant effect on income but positive effects on health measures. Michaels & Lei (2011) examine whether giant oil field discoveries (defined as containing 500 million barrels of recoverable reserves) leads to armed conflict.

Michaels & Lei (2011) is perhaps closest to this paper's approach in terms of source of variation, but differs in two important respects: first, it uses every giant oil field discovery a country experiences, whereas I use only the first discovery that makes a country resource-rich. Field discoveries subsequent to the first one are less plausibly exogenous, since the initial discovery typically leads to enhanced exploration, and also may not be expected to have the same effect as the initial discovery since it is already known that the country has oil. Second, Michaels & Lei (2011) is primarily focused on the effects on civil conflict, while this paper is focused on economic indicators.

In summary, recent work on the resource curse has employed more convincing empirical designs and challenged the old consensus that resources harm GDP growth. However, these designs still face significant potential endogeneity problems. While this paper adds to the dissent on the existence of the resource curse, I argue that the quasi-experimental, treatment-control approach using initial discoveries as a plausibly exogenous source of variation is a more compelling test of causality. Firstly because structural characteristics present before and after discovery are controlled for in the fixed effects design. Second, it is shown in the following section that, with the exception of population, several important initial country characteristics do not predict resource discovery. Third, there is no significant difference in GDP trends between treatments and controls before resource exploitation or discovery, particularly in the case of synthetic control analysis, where the counterfactual is explicitly constructed such that pre-treatment levels and trends are approximately equal.

3 Background of Oil Discovery

This section provides a brief history of the oil industry, with emphasis on how and when production spread geographically, and what factors drove further exploration. I argue qualitatively that new discoveries were primarily driven by global factors exogenous to any one country. I then test if any of several country characteristics are able to predict oil/gas discovery since 1950.

The modern oil industry is typically said to have started in 1859 when Edwin Drake struck oil in Pennsylvania with the first well that was drilled for the sole purpose of finding

oil.⁵ In subsequent decades the oil industry was thoroughly dominated by the United States, though by the turn of the century Russia and the Dutch East Indies (present-day Indonesia) also had significant production. World War 1 made it clear that military might would hinge on access to oil, which, along with the rise of the automobile, led to significant expansion in exploration activities around the world.

Advances in exploration and drilling technology have been and remain a constant theme in the spread of oil discoveries. Initially, oil fields were found simply through seeps to the surface. Prior to World War 1, exploration was based on “surface geology”, in recognition of the fact that oil seeps often occurred in specific types of rock formations. However, limiting exploration to geology associated with surface seeps was not suitable for the vast majority of later-discovered fields, which required no specific surface rock formations. It was not until the invention of the seismograph in the early 1920s that sub-surface structures could be plotted. This and other technologies (aerial surface plotting, micropaleontology) led to an explosion in discoveries in the United States. Still, these methods had a long way yet to go. A British 1926 geological report declared that Saudi Arabia appeared “devoid of all prospects for oil”.

Although oil production had spread to many parts of the globe, at the eve of World War 2 the global market was completely dominated by just a handful of countries. As of 1938, just 8 countries⁶ accounted for 94% of world oil production, and the U.S. alone accounted for almost two thirds. Using UN Commodities data, I calculate the share of the top eight countries in 1950 (the start of this paper’s sample period) to be 92%. Between 1950 and the present day oil production would become far more distributed; in 2008 the figure is 55%.

A convergence of factors led to a flurry of discoveries following World War 2. First, since each theater of the war depended so critically on access to oil, which was arguably the determining factor for the allied victory, governments were ever more eager to secure access to reserves, for military rather than commercial purposes. One clear consequence of this dynamic was the push into Africa in the 1950s. France, which was dependent on imports for its oil supply, began a drive under Charles de Gaulle to develop oil production within its empire. It thus began exploration in its African colonies. Even as the colonial era was winding down, ties to these countries remained strong and would provide a dependable source of oil. Africa up to that point remained largely unexplored, partly due to remoteness and lack of infrastructure, but also because prospects were thought to be sparse⁷. But France’s

⁵This section borrows heavily from the canonical book on the history of oil *The Prize* by Daniel Yergin.

⁶USA, Mexico, Russia, Indonesia, Romania, Iraq, Iran, Venezuela.

⁷In another sign that oil prospecting was still a highly imperfect science, shortly prior to the Algerian discovery a prominent professor of geology at Sorbonne announced that he was “so sure that there was no oil in the Sahara that he would happily drink any drops of oil that happened to be found there”.

push led to discoveries in Gabon and Algeria in the 1950s. In 1956 oil was also found in Nigeria (a British colony). Following these finds, Africa was seen as the “new frontier” of oil, and many companies began exploration across the continent, leading major discoveries in Libya and the Republic of Congo, and smaller ones elsewhere.

A second factor was that the decades following the war was a period of breakneck growth in commercial demand for oil. Driven by rapidly rising incomes, the spread of automobiles and the expanding plastics industry, between 1949 and 1972 oil demand increased by more than five and a half times. A third and related factor was an explosion in competition among producers. By 1970, the old order of the “seven sisters”, the seven giant companies that controlled almost all oil production, had given way to a much more distributed industry. From 1953-1972, over 350 companies entered the non-US oil industry or significantly expanded participation. This surge in competition was itself driven by several factors. Witnessing the benefits being derived in spite of foreign companies controlling operations in countries like Iran and Saudi Arabia, potential producing countries increasingly adopted favorable concessionary policies to encourage exploration. Changes in the U.S. tax code were made to encourage foreign investment. Improvements in transportation and communications made all parts of the world more accessible. Finally, exploration and drilling technology continued to improve and diffuse, reducing risk and barrier to entry. Among the important advances made over the 20th century were satellite imaging, sedimentology, geochemistry, and computing, the last of which helped geologists process large amounts of seismographic data.

A particularly important advance that led to several discoveries in the latter part of the century was deepwater drilling. Offshore drilling dates back to the late 19th century, but was long confined to shallow waters near the coast. In 1947 a milestone was reached when a rig was built 18 miles off the coast of Louisiana, albeit still in shallow waters. The first semi-submersible drilling rig was built in 1961. When the huge (onshore) Groningen gas field was discovered in the Netherlands,⁸ geologists realized that the North Sea floor had similar geology, and exploration into the sea yielded its first discovery in 1970. Up to that point drilling at the depths involved in North Sea drilling had never even been attempted, but this discovery fortuitously coincided with a new generation of offshore technology that made it viable. Major offshore discoveries in previously non-oil-rich nations were made in Malaysia, the United Kingdom, Norway, Denmark, and Equatorial Guinea.

To summarize, major oil discoveries in previously non-producing nations have been driven to a great extent by global factors exogenous to any one country, particularly tech-

⁸Drilling efforts in Western Europe, which dated back to at least the 1920s, had proved mostly unsuccessful. But efforts were renewed following the Suez crisis of 1956, eventually leading to the Groningen discovery.

nology advance and enormous growth in global demand (along with, of course, geographic luck of the draw). As will be shown in the following section, oil prices do not appear to have been a factor in driving exploration in countries without previous discoveries, as most of the major initial discoveries occurred during a time when oil prices remained relatively stable and low, before the price spike of the 1970s.

This is not to say the distribution of discoveries is completely random. Africa was under-explored entering the post-war period at least partially due to lack of infrastructure, but to the extent that this was a region-wide phenomenon, the region-year fixed effects in this paper's regressions control for it. Also, the timing of some African discoveries (and possibly others) had a geopolitical element, as the French colonies were explored earlier due to the French push for access. Hence there may be some caveats to the design of this paper, but there does not appear to be an obvious mechanism that would systematically bias results.

Can the data tell us anything about the likelihood of oil discovery? I use regression analysis to check for whether several initial observable characteristics that may affect future growth are able to predict oil discovery. Each characteristic has been used in past empirical growth literature as a predictor of growth, and several appear in the commonly used specification of Barro & Lee (1991). Each characteristic is observed at 1950, except for Democracy score and investment/GDP, which is observed in 1960 due to data limitations. I run cross-sectional linear probability regressions with having experienced an oil discovery since 1950, conditional on not being resource-rich prior to 1950, as the dependent variable (or having experienced a discovery since 1960 in the cases mentioned above). This indicator is equal to one for all countries with such a discovery, including those not in the treatment group because subsequent production was insignificant.⁹

The results are shown in Appendix Table A1. In the univariate regressions, initial levels of log GDP per capita, democracy level, log of average years schooling, investment/GDP ratio and ethnic fragmentation are all insignificant. Only initial log of population is a significant predictor of discovery. One may guess this is because population is correlated with geographic land area, and countries with large area have more opportunity to discover oil. However, even when controlling for land area (which is predictive in a univariate regression), population is still strongly significant. Another possible explanation is the fact that oil is more likely to be found under softer soil, which is also better able to accommodate larger populations. In any case, any resulting bias in the GDP per capita growth estimates is likely to be downward, since oil wealth is being spread among more people.

When I combine all predictors into one joint regression, I lose all but 40 observations due to data limitations, but the results are largely the same, except that ethnic fragmentation

⁹There are 39 discovery countries by this definition, compared to 78 non-discovery countries.

is positive and significant at a 10% level. Similarly to population, if conditionally more fragmented countries are more likely to discover oil, this would likely cause a downward bias in growth estimates, as fragmentation has been widely found to hinder growth. Further, the country fixed effects in the main regression specifications (which use panel data, rather than a cross section as in Table A1) should largely control for any population and fragmentation effects, since relative population and fragmentation levels are fairly stable over time.

In the regressions shown in Table A1 I am assuming that the size of the discovery is independent of a discovery being made, so even small discoveries are included. If I relax this assumption and run the same regressions with being a treatment country (defined below) as the dependent variable, all coefficients are insignificant, including the one for population.

4 Empirical design

The average effect of resource discovery on post-exploitation outcomes is estimated with the following equation:

$$Y_{crt} = Post_{ct}\delta + \alpha_c + \gamma_{rt} + \epsilon_{ct} \quad (1)$$

Where Y_{crt} is an outcome of interest for country c in region r in year t , $Post_{ct}$ is a country-specific indicator for being after the exploitation event, α_c is country fixed effects, and γ_{rt} is a set of regional year dummies, which control for any common shocks experienced across a region. Regions are assigned according to World Bank country groups where applicable¹⁰.

Effects are also estimated using an event study specification, allowing the treatment effect to vary over time:

$$Y_{crt} = E_{ct}\delta + \alpha_c + \gamma_{rt} + \epsilon_{ct} \quad (2)$$

Where E_{ct} is a vector of indicator dummies for being within some specified 3-year period before or after the exploitation event, and δ is a vector of coefficients corresponding to each 3-year period. In this specification, identification comes from comparing the outcome variable for treatment countries during a given event-time period to the omitted period of 1-3 years before the event. Treatment observations are trimmed in this specification so that each event-time coefficient is estimated with the same number of treatment observations.

¹⁰One difficult case is treatment country New Zealand, which does not naturally fit in any of the listed regions. If I created an Oceania region, New Zealand would be the only country, because Australia is dropped as an initially resource rich country, and other countries are too small or lack data. Therefore I include New Zealand in the Northern Europe region. While obviously not a match geographically, as one of the “neo-Europes” New Zealand has similar culture, institutions, and wealth as Northern European nations.

This is done so that differences in the treatment effect over time are not driven by different compositions of treatment countries identifying each event-time coefficient, an especially important consideration given the small number of treatment countries. Hence the sample is not identical to that used in the baseline specification of equation (1).

Although each event-time coefficient is estimated with a small number of observations relative to the baseline difference-in-differences design, this method has two significant advantages. First, it checks for the existence of pre-existing trends that could lead to spurious difference-in-differences results. Second, it reveals the temporal pattern of the treatment effect, rather than just a post-event average. This advantage becomes increasingly acute to the extent that the treatment effect over time deviates from a simple step function. Most importantly, I can identify differences in short-run versus long-run effects.

It is not obvious how to define the treatment group or the event in question. The purpose is to identify countries that began the 1950-2008 sample with negligible resource production and subsequently achieved substantial resource production on a per capita basis. For oil and gas (or hydrocarbon) discoveries, which make up the entire treatment group except Botswana, a country is included if annual oil and gas production per capita in 1950 was less than one oil barrel energy equivalent¹¹ (henceforth referred to as barrels) per capita, and subsequently passed 10 barrels per capita for a sustained period. Countries that produced more than one barrel at the start of the period, or already had significant mineral wealth are dropped from the sample as unsuitable comparison countries. 27 countries are excluded for this reason (see Appendix C).¹² Thus the regressions compare countries that started resource-poor and became resource rich with countries that remained resource-poor throughout.


These are somewhat arbitrary thresholds, but they satisfactorily uphold the purpose of the treatment group. One barrel per capita generates trivial wealth for the country, whereas 10 barrels generate anywhere from \$100 to over \$800, depending on oil and gas prices in a given year. Further, most countries that pass 10 barrels per capita do so in the early stages of exploitation after a major discovery and go on to produce at much higher levels. In other words, the threshold is effective at separating low-level producers from high-level ones. This is illustrated in Figure 1, a histogram showing the maximum level of annual barrel production per capita achieved over the entire sample period, and only includes sample countries that achieved some non-zero production level. The vertical line represents the threshold to be

¹¹Natural gas production is converted to its oil barrel equivalent in terms of energy generation using the conversion rate of 0.00586152 oil barrels per terajoule, since the raw natural gas production data is given in terajoules.

¹²Former Soviet nations are also excluded, since they lack GDP data before the fall of the Soviet Union, and anyways have obvious confounding factors. Countries with populations of less than 200,000 as of 2007 are also dropped. These exclusions do not meaningfully change the results.

included in the treatment group. The sensitivity of this threshold is tested for the main GDP per capita regression by alternatively setting it to five barrels and 20 barrels (see Appendix Table A3).

Figure 1: Maximum Barrel Production Histogram



There are six countries that matched the above definition in terms of hydrocarbon production but are not included in the treatment group. Four of these countries already generated significant wealth from some other mineral commodities (Suriname, Angola, Australia, and Bolivia). Israel is a unique case in that it only maintained production over 10 barrels per capita for a six year period, then fell to nearly zero from 1976 on, and so cannot be considered to be resource rich. Additionally, Abu Dhabi of what is now the United Arab Emirates discovered oil in 1962, nearly a decade before the emirates were combined into a single nation, so a before-after comparison is neither feasible nor appropriate and the UAE is dropped from the analysis.

The one non-oil and gas country is Botswana (the Netherlands is also a unique case in that it almost exclusively produces natural gas rather than oil), which has yielded tremendous wealth from diamonds on par with the oil-extracting countries in the treatment group. To my knowledge, there are no other non-oil extracting countries appropriate for this treatment group, as nearly all major mineral producers discovered their mineral wealth long before the period studied here.

Table 1 lists the 17 treatment countries, along with event year and first non-zero production year, which are defined and discussed below. The treatment group, while somewhat small, represents a reasonably representative geographic spread, and a variety of economic and political backgrounds. Appendix Table A2 presents summary statistics separately for treatment and control countries as of 1970. I choose 1970 because this is the first year that data for most of the attributes shown are available for all countries, and is still generally

before or shortly after the event years, thus providing a reasonable comparison snapshot of treatment and control countries. The averages are generally similar between the two groups. Control countries do have a significantly higher average population, but this average is skewed by a few very large countries, which the treatment group lacks. The treatment group actually has a slightly larger median population than the control group.

Table 1: Treatment Countries

Country	Event Year	Initial Discovery	1st Production Year	Production Lag	Event Lag
Algeria	1959	1956	1958	2	4
Gabon	1959	1957	1959	2	2
Libya	1961	1958	1961	3	3
Oman	1966	1963	1966	3	3
Netherlands	1966	1959	1963	4	7
Syria	1968	1959	1968	9	9
Nigeria	1969	1956	1957	1	13
Botswana (diamonds)	1971	1967	1971	4	4
Malaysia	1971	1963	1970	7	8
Ecuador	1972	1967	1972	5	5
Republic of Congo	1972	1951	1960	9	21
Norway	1972	1967	1971	4	5
New Zealand	1976	1959	1970	11	17
United Kingdom	1976	1970	1975	5	6
Denmark	1982	1966	1972	6	16
Yemen	1991	1984	1986	2	7
Equatorial Guinea	1992	1984	1992	8	8


One possible way to define the event year is the year of discovery, but this does not make sense for a growth regression, since GDP is not directly affected by the discovery of resources, but rather their extraction¹³. Further, the initial discovery is not always the one that makes a country a major oil producer. For example, the first oil field discovered in the Republic of Congo was Point Indienne in 1951, but this was a minor field and the next one was not discovered until 1969, and production did not take off until 1972. Therefore I define the event to be the year that resource production begins to surge upwards. In more concrete terms, the event year is the first year that growth in oil and gas production increases by 0.5 barrels per capita. All treatment countries have such a year, all of which mark the first year in a surge of production. One exception to this rule is Nigeria, which saw production drop to nearly zero shortly after the event year as defined above (1965), so in this case I assign the second such year (1969), after which production proceeds to surge upwards (this pattern

¹³It is possible that GDP is indirectly affected before extraction by countries borrowing against future windfalls. However, the mostly flat trend prior to extraction shown in the event study graph of Figure 3 suggests that, on average, this is not a major factor.

is likely associated with the Nigerian Civil War that lasted from 1967-1970). For Botswana I assign 1971 as the event year, as this is the first year of operation for the Orapa diamond mine. While the 0.5 barrels threshold is arbitrary by necessity, it successfully captures the point in time that oil and gas production takes off. This is demonstrated in Appendix Figure A1, which shows, for each treatment country besides Botswana, a graph of barrel production over time, with a vertical line denoting the event year.

Defining the event year in this way raises the concern of endogeneity of timing. One argument is that countries, upon making an initial discovery, will not undertake the investment in drilling infrastructure until oil prices are suitably high. However, the timing of exploitation does not typically coincide with high prices. Figure 2 shows the time series of benchmark world oil prices, measured in constant 2005 U.S. Dollars, along with vertical lines indicating event years for oil-producing countries (bold lines indicate two events in the same year). The majority of exploitation events were made in the pre-1970s period of low and stable prices. Two more were in 1988 and 1992, another low-price era. Only two events occurred during the price spike of the 1970s (Denmark, New Zealand), and while we cannot rule out timing endogeneity for these cases, it would be surprising if none of the events fell into this roughly 10-year window, even if the timing of events was completely random.

Figure 2: Oil Price and Event Years



Another concern is that lesser-developed countries will take longer to develop drilling infrastructure, so that the lag between discovery and exploitation somehow induces endogeneity. Here it is useful to consider a third date (in addition to discovery year and event year): the first year of non-zero production. This may differ from the event year if a country initially produces a very small amount of oil, but is a good indicator when at least some drilling infrastructure was in place. Column 4 of Table 1 shows the lag between discovery of the first oil field and the first year of non-zero production. The average lag is five years, with

a minimum of two and maximum of eleven. While there is some variation, it is encouraging that there are no exceptionally long lag times, and even in a hypothetical world where all nations had similar levels of development and institutions, we would expect variation based on geography (how close the country is to a pipeline network) and how accessible the oil is (how deep in the ground, type of soil, remoteness of field, offshore fields, etc.). However, to address the possibility of endogenous variation in production lag, as a robustness check I run a specification with the years between discovery and the event year omitted, so that I am only comparing pre-discovery periods with post-exploitation periods (see Appendix Table A3).

4.2 Synthetic Controls

An alternative way to measure the effect of resource discovery is the synthetic control methodology developed in Abadie and Gardazabal (2003) and Abadie, Diamond, and Hainmueller (2010). Designed for cases where the treatment in question only applies to a single unit, the idea is to construct, through a data-driven algorithm, a weighted combination of control units that matches the pre-treatment outcome behavior of the treated unit, thus creating a post-treatment counterfactual, or “synthetic control”. I apply this method individually for each treatment country, essentially performing 16 different case studies.¹⁴ This both serves as an additional robustness check for the fixed effects model results, and gives greater context to the findings, as we can examine the effect on each individual country, rather than an average effect.

A brief outline of the procedure follows (for more detail, see the aforementioned papers by Abadie, et al). For each treatment country, the pool of possible controls is restricted to countries in its own region, and which neither start the period resource rich nor become resource rich. Suppose there are J control countries and K predictors.¹⁵ Then control country weights are found through an optimization procedure minimizing the following function:

$$(X_1 - X_0W)'V(X_1 - X_0W)$$

Where X_1 is a $(k \times 1)$ vector of predictors for the treatment country, X_0 is a $(K \times J)$ matrix of pre-event predictors for the control countries, W is a $(J \times 1)$ vector of time-invariant weights assigned to control countries which sum to one, and V is a $(K \times K)$ diagonal matrix

¹⁴Gabon is excluded for reasons discussed in section 5.3.

¹⁵For this procedure, a “predictor” can be any linear combination of a pre-treatment variable, including the outcome variable. For example, population one year before the event year could be one predictor, and average population from 2-5 years before the event year could be another.

with the diagonal elements representing the importance of each predictor.¹⁶ Given these weights, the treatment effect in a given post-event period t is:

$$Y_1t - \sum_{j=2}^{J+1} w_j^* Y_jt$$

Where Y_1 is the outcome variable for the treatment country, Y_j is the outcome for control country j and w_j^* is the optimized weight assigned to country j . The main output of the procedure is a simple graph of the outcome variable over time for both the treatment and the synthetic control. Ideally, before treatment the two curves largely overlap, and then diverge after treatment if there is a causal effect.

5 Empirical Results

5.1 Difference-in-Differences

Table 2 presents the regression results for the main specification of equation (1). In the full sample, treatment countries saw a statistically significant average effect of approximately .35 on the log of GDP per capita. This result is economically significant, as it implies that GDP per capita was on average over 40 percentage points higher than the no-discovery counterfactual in the post-exploitation period.

Column 2 shows the results of the main specification if only non-OECD countries are included, and Column 3 if only OECD countries are included¹⁷. They reveal a striking difference in effects between the two groups. The effect for non-OECD treatments is considerably larger than the overall average effect, while that of the OECD countries is actually negative, though small and insignificant. This is not to say that OECD treatments performed badly, but their fellow Northern European control countries likewise experienced steady, robust growth during the sample period, and the relative magnitude of resource wealth is simply too small to have a major effect-this is more clearly illustrated in the synthetic control results discussed below. As for the large effect on non-OECD treatments, in one sense this is not surprising; the non-OECD treatments are much poorer, so an oil discovery can have a greater impact on GDP. However, it would seem to contradict the theory that countries with better institutions upon discovery are better able to avoid a resource curse¹⁸.

¹⁶the V matrix is found through a nested optimization procedure such that the mean squared prediction error of the pre-treatment outcome variable is minimized.

¹⁷Five of the 17 treatment countries are in the OECD: Denmark, Netherlands, New Zealand, Norway,

Table 2: Difference-in-Differences: GDP/capita

	(1) Full Sample	(2) Non-OECD	(3) OECD only
Post	0.350* (0.157)	0.540** (0.199)	-0.102 (0.105)
N	6195	4956	1239
R^2	0.684	0.620	0.962

Notes: The dependent variable is the natural log of real GDP/capita. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.

5.2 Event Study

Table 3 shows the results of the event study specification of equation (2). For treatment countries, only observations from nine years before to 17 years after are included to obtain a balanced (by event-time) panel. In the full sample of Column 1, there are no significant effects on GDP for any time before resource exploitation, but rather dramatic positive effects in the years following, reaching a coefficient of .43 by the end of the time frame studied. The effect on growth appears to subside after about 10 years, leaving no long-term growth effects but a persistent and large level effect. The same pattern, to a greater degree, is followed for the sample with non-OECD countries only. With only OECD countries included, there is a slight negative downward trend before the event year and no long-term effect. The graphical representation of this table is shown in Figure 3¹⁹.

Even if we hypothesized a resource curse, we might expect the years immediately following exploitation to see positive growth effects, as the direct contribution of resource extraction to GDP is growing rapidly, while the negative mechanisms could take time to manifest. While there does not seem to be a long-run negative growth effect from the event study specification, it is still possible that negative effects begin even farther into the future. To test this, I extend the event-time period analyzed out to 30 years after exploitation. To keep a balanced panel in this case, I only need to drop two treatment countries from the analysis (Equatorial Guinea and Yemen). The graphical result of this specification is shown

United Kingdom.

¹⁸This is somewhat consistent with Davis, 2013, which replicated the Sachs and Warner result that oil-rich countries with poor institutions performed worse, but found that the result was sample-dependent and driven by a few outliers.

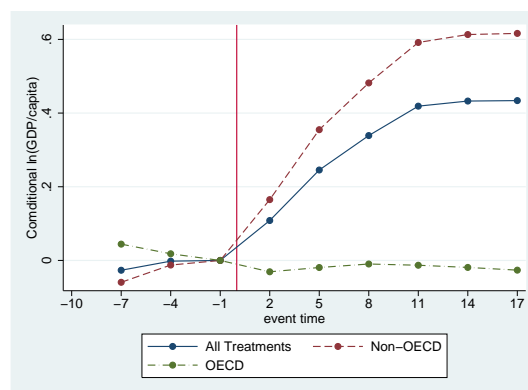
¹⁹In this and all subsequent event study graphs, each point corresponds to the event-time coefficient representing observations from the previous three event-time years. For example, in Figure 3 the point shown at event-time negative seven represents the coefficient for “Exploitation Year - 7-9”. Hence the graph actually represents a period going back to nine years before the exploitation year.

Table 3: Event Study: GDP/capita

	(1) Full Sample	(2) non-OECD Treatments	(3) OECD Treatments
Exploitation year - 7-9	-0.027 (0.033)	-0.059 (0.045)	0.044* (0.019)
Exploitation year - 4-6	-0.002 (0.027)	-0.013 (0.037)	0.018 (0.013)
Exploitation year + 0-2	0.108** (0.037)	0.165*** (0.045)	-0.031 (0.020)
Exploitation year + 3-5	0.245** (0.090)	0.355** (0.114)	-0.019 (0.039)
Exploitation year + 6-8	0.339** (0.121)	0.482** (0.154)	-0.010 (0.044)
Exploitation year + 9-11	0.419** (0.145)	0.592** (0.183)	-0.013 (0.057)
Exploitation year + 12-14	0.433** (0.163)	0.613** (0.204)	-0.019 (0.076)
Exploitation year + 15-17	0.434** (0.162)	0.616** (0.205)	-0.026 (0.085)
<i>N</i>	5650	4571	1079
<i>R</i> ²	0.701	0.632	0.966


Notes: The dependent variable is the natural log of real GDP/capita. The omitted category is 1-3 years before exploitation or never experiencing an exploitation event. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.

Figure 3: Event Study, GDP per capita



is Figure 4. Conditional GDP per capita remains roughly flat from years 10-30 (note that the magnitude of the effect is smaller due to the exclusion of Equatorial Guinea, which experienced extremely high growth rates following exploitation). Although there is a slight downward trend at the end of the period, there is scant evidence of a long-term curse.

Figure 4: Event Study, GDP per capita, Long Panel



5.3 Synthetic Controls

As a robustness check and to show the variation of effects within the treatment group, I next run synthetic control analysis for each treated country.²⁰ For the effect on GDP per capita, I use the following six predictor variables to construct each synthetic control: ethnic fragmentation, population one year before the event, and GDP per capita one, three, five and seven years before the event. The weights making up each country's synthetic control for this analysis are shown in Appendix E.


The graphical results for each individual treatment country are shown in Appendix D. Each graph shows the time series of GDP per capita for each treated unit and its corresponding synthetic control over the entire period from 1950-2008. The results are largely consistent with the difference-in-differences results, in that we see a positive average effect in the short and long term. However, there is an interesting variety of outcomes. There are five countries (Botswana, Republic of Congo, Equatorial Guinea, Nigeria, and Oman) that perform significantly better than their synthetic counterpart (although Nigeria's advantage was nearly gone before the oil price surge in the 2000s). There are three countries (Algeria, New Zealand, Yemen) that do noticeably worse in the long term (although the pre-trend of New Zealand is not especially well replicated, as New Zealand was one of the world's richest countries at the start of the sample period). There is generally little to no effect on OECD

²⁰There is one country, Gabon, where the pre-event level and trend of GDP per capita is not well replicated by its synthetic control. This is because at the onset of oil exploitation, Gabon was already the wealthiest country in the sample of sub-Saharan African countries. Abadie et al (2010) states that the method may not be appropriate if the predictors of the treatment unit do not lay within the convex hull of those of the control units. As it turns out, For Gabon the method gives 100% weight to the second richest pre-event control country, Mauritius. As this does not adequately reproduce Gabon's pre-treatment behavior and is not a credible counterfactual, Gabon is excluded from this part of the analysis. Similarly, Oman's synthetic control is 100% Egypt, but the GDP per capita levels in the years preceding the event are reasonably well-replicated, so Oman is included.

countries, as steady growth is matched by their synthetic counterparts. For a few countries a striking spurt of growth following the event year is followed by a sharp drop, particularly in the case of Libya, in which all of the gains are lost. In these cases the surge and subsequent fall closely correspond to similar patterns in production levels, indicating that these countries in particular failed to develop the non-hydrocarbon economy. Overall, the synthetic control results portray positive or non-negative short-run results for most treatment countries, but a more mixed record in the long-run, particularly in lesser-developed regions.

Figure 5 shows the synthetic control results for a representative sample of five countries in a single graph. The selected countries are intended to illustrate the different types of cases discussed in the preceding paragraph. Each line in Figure 5 represents the results for one country, and is the difference between the log of GDP per capita of the treatment country and that of the synthetic control for each year of event-time.

Figure 5: Synthetic Control Results, Selected Countries



Note: each line represents the difference in the log of GDP per capita between the country and its synthetic control.

5.4 Heterogeneous Effects

The synthetic control results show that although the average effect of discoveries is positive, outcomes vary widely by individual country. Are there characteristics at the start of the sample period that can predict a large or small treatment effect? To attempt to answer this question I take the specification in Equation 1 and add interaction terms between the post-exploitation variable and various initial characteristics that may affect growth and the effect of resources on growth.²¹

²¹Initial log population, log GDP per capita and log of average years of education are measured in 1950.

Column 1 of Table 4 shows the results for the full sample and all interactions included. Since there is no education data for three treatment countries (Equatorial Guinea, Oman, and Republic of Congo), these countries are not included in this specification. Column 2 shows the results when the education interaction is dropped and thus all treatment countries are included. In both cases the interactions with initial GDP per capita and population are negative and significant, with the intuitive implication that a natural resource boom has a greater impact on growth in countries with smaller starting economies and fewer people to “spread” the wealth between. Consistent with Hodler (2006), higher ethnic fragmentation has a negative effect, but the estimate is only significant at a 10% level in the first specification. The initial infant mortality interaction has a negative but insignificant effect, while the initial education interaction has a positive effect (consistent with Ortega & Gregorio, 2005 and Gylfason, 2001), indicating that countries with higher overall levels of development, after controlling for GDP per capita, receive greater benefits from resource discoveries. The estimate for infant mortality increases considerably in magnitude when education is dropped, as the two are strongly correlated.

Because the positive overall growth results are driven by the non-OECD treatments, and since those groups of countries differ in ways that may not be fully captured with the variables used here, I run the same specifications dropping OECD treatments. The results, shown in Columns 3 and 4 of Table 4, are similar to that of the full sample, except that the infant mortality interaction loses significance. Overall, only the population and GDP per capita interactions are robustly significant.

5.5 Robustness

In this section I run several robustness checks, for which all results are shown in Appendix A. First, Since inclusion in the treatment group involves a somewhat arbitrary cutoff (a maximum production level of at least 10 barrels of oil or oil-equivalent gas during the period studied), I test the sensitivity of increasing and decreasing this cutoff. In Column 1 of Appendix Table A3, Panel A I increase the cutoff to 20 barrels, which eliminates five treatment countries.²² The treatment effect with this reduced treatment group increases considerably, as would be expected given the higher intensity of treatment. In Column 2 I decrease the cutoff to five barrels, which adds six countries.²³ The effect is slightly smaller, but still statistically significant.

In Columns 3 and 4 I perform robustness checks against the endogeneity of production

Infant mortality is measured in 1955. Fragmentation is only measured once per country, but relative fragmentation levels are assumed to be largely stable over time.

²²Ecuador, New Zealand, Nigeria, Syria, and Yemen.

²³Albania, Cameroon, Egypt, Hungary, Indonesia, and Tunisia.

Table 4: Heterogeneous Treatment effects: Non-OECD Treatments

	(1)	(2)	(3)	(4)
	All countries	All countries	non-OECD Treatments	non-OECD Treatments
Post	5.47*** (0.72)	5.71*** (0.57)	6.94*** (1.03)	7.08*** (0.72)
Post*(log pop.)	-0.11 ⁺ (0.062)	-0.17*** (0.046)	-0.20** (0.065)	-0.22*** (0.027)
Post*(log GDP/cap)	-0.62*** (0.11)	-0.67*** (0.075)	-0.62*** (0.069)	-0.67*** (0.071)
Post*(log fragmentation)	-0.15 ⁺ (0.081)	-0.018 (0.056)	-0.19 (0.12)	0.037 (0.039)
Post*(log inf. mortality)	-0.058 (0.21)	-0.41*** (0.11)	0.39 (0.45)	0.12 (0.27)
Post*(log avg. yrs school)	0.18 ⁺ (0.096)		0.24 ⁺ (0.13)	
<i>N</i>	6018	6195	4779	4956
<i>R</i> ²	0.730	0.734	0.658	0.675

Notes: The dependent variable is the natural log of real GDP/capita. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.

lag. Column 3 excludes treatment country observations between the first recorded oil field discovery and the first year of non-zero production, so that only pre-discovery and post-production outcomes are compared. Column 4 excludes the observations between the first recorded discovery and the actual event year. In both cases the results are similar to the main specification, but the estimates are slightly larger than for the full sample.

In Column 1 of Panel B I run the main specification of equation (1) using Penn World Table 7.0 GDP data. As discussed in Appendix B, PWT data does not have complete coverage going back to 1950 for many countries, and as a result five treatment countries do not have data before the event year (Algeria, Gabon, Libya, Oman, and Yemen) and thus cannot contribute to identification of the treatment effect. With these countries dropped, PWT data yields a similar point estimate to that of the full Maddison sample, but with larger standard errors due to the reduction in treatment observations. In Column 6 I run the same specification with the same observations as in Column 2, but with Maddison GDP data. Hence the difference in the estimated effect is due solely to differences in GDP measurement, rather than sample differences. This estimate is actually slightly smaller than the PWT estimate, suggesting that, if anything, Maddison data underestimate the treatment effect.

In Column 3 of Panel B I run the main specification using GDP measured in constant 2000 US Dollars as the dependent variable. One possible concern about the GDP results is that the PPP adjustments used in Maddison and Penn World Tables does not sufficiently

reflect the higher price differences found in resource-rich economies (if, for example, the adjustments are made using a basket of goods that is not representative). To test for this I use a third GDP data source, World Development Indicators (WDI), that provides GDP in constant 2000 US Dollars (non-PPP adjusted). WDI has similar data limitations as Penn World Tables, and four treatment countries are omitted since they do not have pre-treatment data (Algeria, Gabon, Libya, Yemen). The estimate for constant 2000 US Dollars is again similar but slightly smaller than the main result. However, as shown in Column 4, the estimate using Maddison GDP for the equivalent sample is very similar. This suggests that erroneous PPP adjustments are not inflating the estimated effects.

Another possible concern is the non-stationarity of GDP per capita. Given that the sample has a large number of time periods, if residuals are non-stationary even after controlling for year fixed effects, this could lead to inconsistent standard error estimates. To address this I use two alternative specifications that mitigate non-stationarity. First, I include country-specific time trends (which also controls for the possibility that results are driven by differing long-term trends between treatment and control countries). Second, I insert the GDP per capita growth rate (specifically, the year-on-year difference in the natural log of GDP per capita) as the dependent variable, and include the lagged level of GDP per capita on the right hand side. The results are given in Appendix Table A4. For the specification including country trends, the estimate is only slightly lower and actually more precisely measured. Using an augmented Dickey-Fuller unit root test on the residuals of this regression rejects the null hypothesis that all panels contain unit roots at a 5% level, so this specification is successful in mitigating non-stationarity. The growth rate estimate in Column 2 is also significant at a 5% level, and implies a 2.1 percentage point effect on growth rates. However, as suggested in the GDP level event study in Figure 3, the growth effect is not permanent. This is likewise borne out in a growth rate event study specification, which is shown in Appendix Figure A2. The first graph shows the effects on growth rates for the full sample from 9 years before exploitation to 17 years after, while the second graph shows effects for the longer panel, where Equatorial Guinea and Yemen are dropped (this is analogous to Figure 4). As expected, after the initial spike in growth following discovery, effects are close to zero in the long-term.

6 GDP Components

6.1 Non-hydrocarbon GDP

While the results thus far have established a positive average effect of resource discovery on GDP per capita, they have been silent on the mechanisms of growth. In a broad sense,

there are two possible mechanisms: first is the obvious one of resource production directly adding to GDP; second is the indirect effect of reinvesting part of the windfall for future growth. A simple way of positing this is to imagine a simple Solow model where GDP is augmented by an exogenous resource shock in a given period. The output from this shock can either be consumed or invested to increase future output. Additionally, oil revenues can either be reinvested back into the oil sector to support further exploration and drilling infrastructure, or used to support other industries, such as manufacturing, in an effort to diversify the economy.

While we don't have sector-specific investment data, we can adapt the empirical design of this paper to explore the impact of resource discovery on the non-oil sector by constructing a non-resource-generated GDP per capita variable, which I insert as the dependent variable in the main specification of Equation (1). I derive resource value, measured in current U.S. dollars, by combining the UNINDCOM data on oil and natural gas production levels, oil price data from *UNCTADstat* online²⁴ and U.S. natural gas wellhead prices from the Energy Information Administration. This value is converted to a real value and subtracted off the real GDP level from Penn World Tables.²⁵ As mentioned above, because of less extensive GDP data coverage in Penn World Tables, there are five treatment countries that do not have pre-exploitation data (Algeria, Gabon, Libya, Oman, and Yemen), and thus cannot contribute to identifying a treatment effect and are dropped from the sample. In addition, Botswana is not included in this part of the analysis since diamond prices vary significantly by individual diamond, so price indices are not available. Finally, there are three years in which the estimated value of oil and gas extracted in Equatorial Guinea exceeds the GDP given in Penn World Tables. This may indicate an overstatement of production, or that the price indices used exceed the prices Equatorial Guinea received. For this reason Equatorial Guinea is also dropped from this analysis. This leaves only 10 treatment countries, so the following results should be viewed with heightened caution, but the results are still suggestive.

Column 1 of Table 5 shows the results for the full sample. The effect on the log of non-hydrocarbon GDP per capita was approximately -.18, or roughly -20 percentage points. To check whether developed countries were more effective in diversification, I run the same regression for the sub-sample that includes non-OECD countries only, and likewise for OECD countries. While the negative impact is slightly larger for non-OECD countries, the effect is similar for both groups. These results suggests that the estimated GDP per capita gains come exclusively from the direct contribution of oil and gas revenues, and that other sectors

²⁴This is an average of equally weighted Dubai, Brent, and Texas crude oil prices.

²⁵PWT is used instead of Maddison here because Maddison only provides real PPP-adjusted GDP levels, without showing the deflators used, so I cannot convert current resource value to a comparable real value.

of the economy may have been crowded out.

Table 5: Difference-in-Differences: Non-hydrocarbon GDP/capita

	(1)	(2)	(3)
	All countries	non-OECD Only	OECD Only
Post	-0.18** (0.060)	-0.20* (0.091)	-0.15+ (0.084)
<i>N</i>	5993	4617	1376
<i>R</i> ²	0.715	0.609	0.959

Notes: The dependent variable is the log of real non-hydrocarbon generated GDP. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.

6.2 Growth Accounting


In this section I evaluate the effect of resource exploitation on proximate causes of GDP growth, and estimate the contribution of each factor to treatment country growth in the years following exploitation. Data on capital, labor force, and TFP are taken from the UNIDO World Productivity Database (WPD), while education data is taken from the Barro-Lee Educational Attainment data set.

I first run the event study specification of equation (2) for four proximate causes of GDP growth: capital formation, labor force size, human capital, and Total Factor Productivity (TFP). This analysis is carried out for non-OECD treatment countries only, since this is the group for which resources were found to have a large effect on GDP. Due to coverage limitations of both the WPD and Barro-Lee data sets, some treatment countries are missing²⁶ and except for education the panels are not balanced, so these results should be viewed with caution. Still, the results shown in Figure 6 are suggestive that each component of growth was impacted by resources. Capital Stock saw large positive effects, while TFP experienced a smaller but still substantial gain. There was an upward trend in capital stock prior to exploitation, which likely reflects the investment in drilling infrastructure in the period between discovery and exploitation. This likely helps explain the downward trend seen in TFP prior to exploitation, as during this period large amounts of extra capital are being formed but not actually producing significant oil output yet.

There are also positive effects on both the size and quality of the labor force. The

²⁶For the capital, TFP and labor force regressions, Libya, Oman and Yemen are excluded due to lack of data. For the education regression, Equatorial Guinea, Nigeria and Oman are excluded.

Figure 6: Proximate Causes of Growth, Non-OECD countries



increase in labor force size may reflect an influx of migrant workers following the resource boom. In any case, the increase of workers mirrors an increase in overall population in treatment countries, as there is no effect on the labor utilization rate (regression not shown), so this does not contribute to the effect on GDP per capita. However, the population also became more educated²⁷, with an effect of about 1 additional year of average schooling 20 years after exploitation. This may be a result of an influx of more educated migrants, or an increase in public investment in education resulting from oil revenues, or some combination thereof.

To find the contributions of these factors to treatment country growth in GDP per worker²⁸, I use a conventional growth accounting framework (also used in Cho & Tien, 2014) with a Cobb-Douglas production function in which labor is augmented by education:

$$Y = AK^\alpha (Le^{\gamma s})^\beta \quad (3)$$

²⁷Since education is measured only every five years, each event time coefficient covers a five year window, which only includes one observation per country occurring at some point within the window.

²⁸Data on income per worker is also taken from the World Productivity Database. Of course, this is not exactly equal to GDP per capita, but it should be a close substitute as far as growth rates are concerned, since there is generally very little change in labor utilization rates in treatment countries.

Where A is equal to TFP, K is total capital stock, L is the size of the labor force, γ is a parameter that determines the returns to education, s is average years schooling, α is the share of capital in the economy, β is the share of labor in the economy. Dividing by L , taking the natural log of both sides and differencing yields the following:

$$\Delta \ln \frac{Y}{L} = \Delta \ln(A) + \alpha \Delta \ln \frac{K}{L} + \beta \gamma \Delta s \quad (4)$$

This is the growth accounting equation I use to decompose sources of growth. I assume constant returns to scale (ie $\alpha + \beta = 1$), and follow the convention of setting α equal to $1/3$ and β equal to $2/3$. The return to schooling is $\beta\gamma$. I follow Bernanke & Gurkaynak (2002) in setting the return to schooling at seven percent (hence setting $\gamma = .105$). These choices for α , β and γ are merely conventions, and studies have shown that there is significant variation in these values between countries (Oduor, 2010 and Uwaifo, 2006 for example), so the results should be viewed as rough estimates. Growth in TFP is calculated using the residual method, in which growth not accounted for directly by inputs is assumed to be a result of productivity growth.

Equation (4) identifies three components of growth in GDP per worker: growth in TFP, growth in capital per worker, and growth in average years schooling. Plugging in the data for these inputs²⁹ and using the parameter values described above, growth in each component and its share in overall GDP per worker growth is calculated. These shares are reported for all treatment countries with the necessary data in Tables 6 and 7. The analysis is carried out over two different periods: the extremely high growth period of 0-8 years following the exploitation event³⁰, and the steadier period of 9-17 years after the event. Additionally, for both OECD and non-OECD countries, I calculate the average growth rates and component shares for non-treatment countries that experienced positive annual GDP per worker growth of at least .05%³¹ over relevant time frames: 1960-90 for non-OECD controls, and 1970-90 for OECD controls. These time frames correspond to the years where the most treatment countries were between 0-17 years after exploitation. So while it is not a direct comparison, it is suggestive of differences between the two groups.

As might be expected from the previous results in this paper, the OECD treatments' growth profiles generally do not differ radically from the average OECD control country (set-

²⁹Since schooling data is only given every five years, for the growth accounting exercise I fill in the missing years by linear interpolation.

³⁰There are two minor exceptions to this. Because the event year for Algeria and Gabon is 1959, one year before the WPD data starts, for these two countries the growth accounting results shown in Table 6 are for the 2-8 years following the event.

³¹This restriction is made to avoid results driven by outlier component shares caused by near-zero growth, and to compare mechanisms with non-resource driven countries that also experienced positive growth.

ting aside outlier observations, such as New Zealand, which experienced near-zero growth in income per worker and thus hugely inflated “shares” of growth). However, among non-OECD countries there is a striking difference between treatments and controls, beyond the difference in GDP growth rates. Non-OECD control countries on average relied almost exclusively on increases in capital per worker (or “capital deepening”) for growth in GDP per worker. Education increases played a minor role, while TFP actually decreased. Negative TFP growth is consistent with past growth accounting studies focusing on Sub-Saharan Africa (Cho & Tien 2004, Bosworth & Collins 2003, and Tahari, et al 2004.). On the other hand, non-OECD treatments, while generally also experiencing robust capital deepening and moderate education growth, also saw substantial contributions from TFP, particularly in the first 8 years after the exploitation event (two countries do have negative TFP growth in the 9-17 years post-exploitation period). The difference is all the more remarkable when considering that the control countries in this analysis are limited to those that had experienced at least moderate positive growth over the relevant period.

This result offers the intuitive interpretation that resource discovery provided an extremely high yield sector during a period when other productivity-enhancing opportunities were generally not available to developing countries. This naturally also led to greater levels of capital investment and high growth derived from both investment and productivity gains, while other developing countries relied almost solely on capital accumulation for much more modest growth.

Table 6: Growth Accounting: 0-8 years After Event

Country	Y/L growth	K/L share	Education share	TFP share
Non-OECD				
Algeria	1.6%	-45.4%	22.3%	123.0%
Botswana	11.6%	44.7%	2.3%	53.0%
Congo, Republic of	5.4%	21.1%	8.1%	70.8%
Ecuador	6.7%	27.3%	3.6%	69.1%
Gabon	10.1%	18.9%	4.4%	76.7%
Malaysia	3.8%	58.9%	5.9%	35.2%
Syria	5.1%	44.7%	6.7%	48.7%
Non-treated 1960-90 avg.	2.2%	94.5%	13.3%	-7.8%
OECD				
Denmark	1.2%	27.9%	1.6%	70.5%
Netherlands	3.1%	62.6%	4.5%	32.9%
New Zealand	-0.1%	-266.8%	7.8%	359.0%
Norway	2.5%	50.4%	2.1%	47.5%
United Kingdom	1.8%	65.3%	2.7%	32.1%
Non-treated 1970-90 avg.	2.4%	50.1%	4.6%	45.3%

Table 7: Growth Accounting: 9-17 years After Event

Country	GDP per worker growth	Cap. Per worker share	Education share	TFP share
Non-OECD				
Algeria	3.1%	38.4%	15.5%	46.1%
Botswana	4.9%	42.5%	11.5%	46.0%
Republic of Congo	5.4%	22.7%	3.6%	73.7%
Ecuador	-2.3%	13.5%	-4.7%	91.2%
Gabon	5.6%	55.9%	7.6%	36.5%
Malaysia	2.2%	92.1%	9.3%	-1.4%
Syria	1.3%	162.7%	20.0%	-82.6%
Non-treated 1960-90 avg.	2.2%	94.5%	13.3%	-7.8%
OECD				
Denmark	1.9%	28.3%	0.6%	71.0%
Netherlands	-0.1%	-660.7%	-96.8%	857.5%
New Zealand	-0.7%	-1.1%	-0.4%	101.5%
Norway	1.9%	44.2%	5.3%	50.5%
United Kingdom	1.5%	33.8%	2.3%	63.8%
Non-treated 1970-90 avg.	2.4%	50.1%	4.6%	45.3%

7 Conclusion

This paper takes a novel approach to estimating the impact of natural resources, using modern discoveries, longitudinal data, and sophisticated empirical methods to provide a more rigorous test of the existence of the resource curse than has been heretofore performed. I find positive short-term GDP per capita growth effects and positive long-term level effects in lesser developed countries, and no effect for developed countries. In evaluating the proximate causes of growth, I find positive effects on capital formation, productivity and education. These results are consistent with the predictions of a simple Solow model in which there is a temporary shock to TFP growth, or with the semi-endogenous R & D based growth model proposed by Jones (1995). The fact that treated developing countries saw robust TFP growth while the average non-treated developing country saw negative TFP growth strongly suggests that high-return growth opportunities were scarce in the developing world over the time frame studied. This is further supported by the finding of negative effects on the non-hydrocarbon sector, which may reflect the presence of negative spillovers in the form of dutch disease.

There is little evidence in this study to support the presence of other common resource curse channels, such as harm to political and economic institutions. These types of effects would be expected to be felt some number of years after the beginning of exploitation, but

this study finds no evidence of negative long-run growth effects. It is possible that these channels take an even longer time to manifest than is studied here. It is also possible that the usual negative association between institutions and resource wealth applies mainly to countries that were known to be resource-rich in the colonial era and were thus given extractive institutions. In any case, further research is needed on the link between resources and institutions, as even the negative impact on democracy has come under recent challenge in the literature.

Perhaps the most pressing research question going forward is how equitably resource-driven growth is distributed within countries. Equatorial Guinea, with one of the highest levels of GDP per capita in the region and yet one of the highest poverty rates, is a stark demonstration of the perils of using GDP per capita as an overall measure of welfare. Assessing the impact of resource exploitation on inequality would be a useful extension of the empirical designs used in this paper, but the demands of panel inequality data is a difficulty that must be overcome.

References

- Acemoglu, D., Johnson, S., and Robinson, J. (2001). The colonial origins of comparative development: An empirical investigation. *American Economic Review*, 91(5):1369–1401.
- Acemoglu, D., Johnson, S., and Robinson, J. (2002). Reversal of Fortune: Geography and Institutions in the Making of the Modern World Income Distribution*. *Quarterly Journal of Economics*, 117(4):1231–1294.
- Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., and Wacziarg, R. (2003). Fractionalization. *Journal of Economic Growth*, 8(2):155–194.
- Alexeev, M. and Conrad, R. (2009). The elusive curse of oil. *The Review of Economics and Statistics*, 91(3):586–598.
- Aslaksen, S. (2007). Corruption and oil: evidence from panel data. *Unpublished Manuscript*.
- Atkinson, G. and Hamilton, K. (2003). Savings, growth and the resource curse hypothesis. *World Development*, 31(11):1793–1807.
- Barro, R. (1999). Determinants of democracy. *Journal of Political economy*, 107(S6):158–183.
- Barro, R. and Sala-I-Martin, X. (1995). Economic growth. *Economic Growth Theory*.
- Bernanke, B. S. and Gürkaynak, R. S. (2002). Is growth exogenous? taking mankiw, romer, and weil seriously. In *NBER Macroeconomics Annual 2001, Volume 16*, pages 11–72. MIT Press.
- Bohn, H. and Deacon, R. T. (2000). Ownership risk, investment, and the use of natural resources. *American Economic Review*, pages 526–549.
- Bosworth, B. and Collins, S. M. (2003). The empirics of growth: An update. *Brookings papers on economic activity*, 2003(2):113–206.
- Bravo-Ortega, C. and De Gregorio, J. (2005). The Relative Richness of the Poor? Natural Resources, Human Capital, and Economic Growth. *World*.
- Brückner, M., Ciccone, A., and Tesei, A. (2012). Oil price shocks, income, and democracy. *Review of Economics and Statistics*, 94(2):389–399.
- Brunnschweiler, C. and Bulte, E. (2008). The resource curse revisited and revised: A tale of paradoxes and red herrings. *Journal of Environmental Economics and Management*, 55(3):248–264.
- Caselli, F. and Michaels, G. (2013). Do oil windfalls improve living standards? evidence from brazil. *American Economic Journal: Applied Economics*, 5(1):1–31.
- Caselli, F. and Tesei, A. (2011). Resource windfalls, political regimes, and political stability. Technical report, National Bureau of Economic Research.
- Cho, Y. and Tien, B. (2014). Sub-saharan africa’s recent growth spurt: an analysis of the sources of growth. *World Bank Policy Research Working Paper*, (6862).

- Collier, P. and Goderis, B. (2012). Commodity prices and growth: An empirical investigation. *European Economic Review*, 56(6):1241–1260.
- Collier, P. and Hoeffler, A. (1998). On economic causes of civil war. *Oxford economic papers*, 50(4):563.
- Collier, P. and Rohner, D. (2008). Democracy, development, and conflict. *Journal of the European Economic Association*, 6(2-3):531–540.
- Cotet, A. M. and Tsui, K. K. (2013a). Oil and conflict: What does the cross country evidence really show? *American Economic Journal: Macroeconomics*, 5(1):49–80.
- Cotet, A. M. and Tsui, K. K. (2013b). Oil, growth, and health: What does the cross-country evidence really show? *The Scandinavian Journal of Economics*, 115(4):1107–1137.
- David, A. P. and Wright, G. (1997). Increasing returns and the genesis of american resource abundance. *Industrial and Corporate Change*, 6(2):203–245.
- Davis, G. A. (2011). Replicating sachs and warner: The 1997 working paper.
- Davis, G. A. (2013). Replicating sachs and warner’s working papers on the resource curse. *The Journal of Development Studies*, (ahead-of-print):1–16.
- Gylfason, T. (2001). Natural resources, education, and economic development. *European Economic Review*, 45(4-6):847–859.
- Gylfason, T., Herbertsson, T., and Zoega, G. (1999). A mixed blessing. *Macroeconomic Dynamics*, 3(02):204–225.
- Haber, S. and Menaldo, V. (2011). Do natural resources fuel authoritarianism? a reappraisal of the resource curse. *American Political Science Review*, 105(1):1–26.
- Hodler, R. (2006). The curse of natural resources in fractionalized countries. *European Economic Review*, 50(6):1367–1386.
- Jones, C. I. (1995). R & d-based models of economic growth. *Journal of political Economy*, pages 759–784.
- Lederman, D. and Maloney, W. (2003). Trade structure and growth. *World Bank Policy Research Working Paper*, (3025).
- Lei, Y.-H. and Michaels, G. (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics*.
- Leite, C. and Weidmann, J. (2002). Does mother nature corrupt? Natural resources, corruption, and economic growth. *Governance, corruption, and economic performance*, pages 159–196.
- Mehlum, H., Moene, K., and Torvik, R. (2006). Institutions and the Resource Curse*. *The Economic Journal*, 116(508):1–20.
- Michaels, G. (2011). The long term consequences of resource-based specialisation*. *The Economic Journal*, 121(551):31–57.

- Michaels, G. and Lei, Y.-H. (2011). Do giant oilfield discoveries fuel internal armed conflicts? Technical report, mimeo, LSE.
- Mideksa, T. K. (2012). The economic impact of natural resources. *Journal of Environmental Economics and Management*.
- Oduor, J. (2010). Are prior restrictions on factor shares appropriate in growth accounting estimations? *Economic Modelling*, 27(2):595–604.
- Papyrakis, E. and Gerlagh, R. (2004). The resource curse hypothesis and its transmission channels. *Journal of Comparative Economics*, 32(1):181–193.
- Robinson, J., Torvik, R., and Verdier, T. (2006). Political foundations of the resource curse. *Journal of Development Economics*, 79(2):447–468.
- Ross, M. (2001). Does oil hinder democracy? *World Politics*, 53(3):325–361.
- Sachs, J. and Warner, A. (1995). Natural resource abundance and economic growth. *NBER working paper*.
- Sachs, J. D. and Warner, A. M. (2001). The curse of natural resources. *European economic review*, 45(4):827–838.
- Sala-i Martin, X. and Subramanian, A. (2003). Addressing the natural resource curse: an illustration from Nigeria. *NBER working paper*.
- Sarr, M., Bulte, E., Meissner, C., and Swanson, T. (2011). On the looting of nations. *Public Choice*, 148(3):353–380.
- Sokoloff, K. and Engerman, S. (2000). History lessons: Institutions, factors endowments, and paths of development in the new world. *The Journal of Economic Perspectives*, 14(3):217–232.
- Tahari, A. (2004). *Sources of Growth in Sub-Saharan Africa*, volume 4. International Monetary Fund.
- Torvik, R. (2002). Natural resources, rent seeking and welfare. *Journal of Development Economics*, 67(2):455–470.
- Tsui, K. K. (2011). More oil, less democracy: Evidence from worldwide crude oil discoveries. *The Economic Journal*, 121(551):89–115.
- Uwaifo, R. O. (2006). *Three Essays on the Return to Education Enigma in Nigeria*. ProQuest.
- Van der Ploeg, F. and Poelhekke, S. (2010). The pungent smell of red herrings: Subsoil assets, rents, volatility and the resource curse. *Journal of Environmental Economics and Management*, 60(1):44–55.
- Wick, K. and Bulte, E. (2009). The curse of natural resources. *Annu. Rev. Resour. Econ.*, 1(1):139–156.
- Yergin, D. (2008). *The prize: The epic quest for oil, money, & power*. Free Pr.

Appendix A: Additional Graphs and Tables

Figure A1: Hydrocarbon Production and Event Years

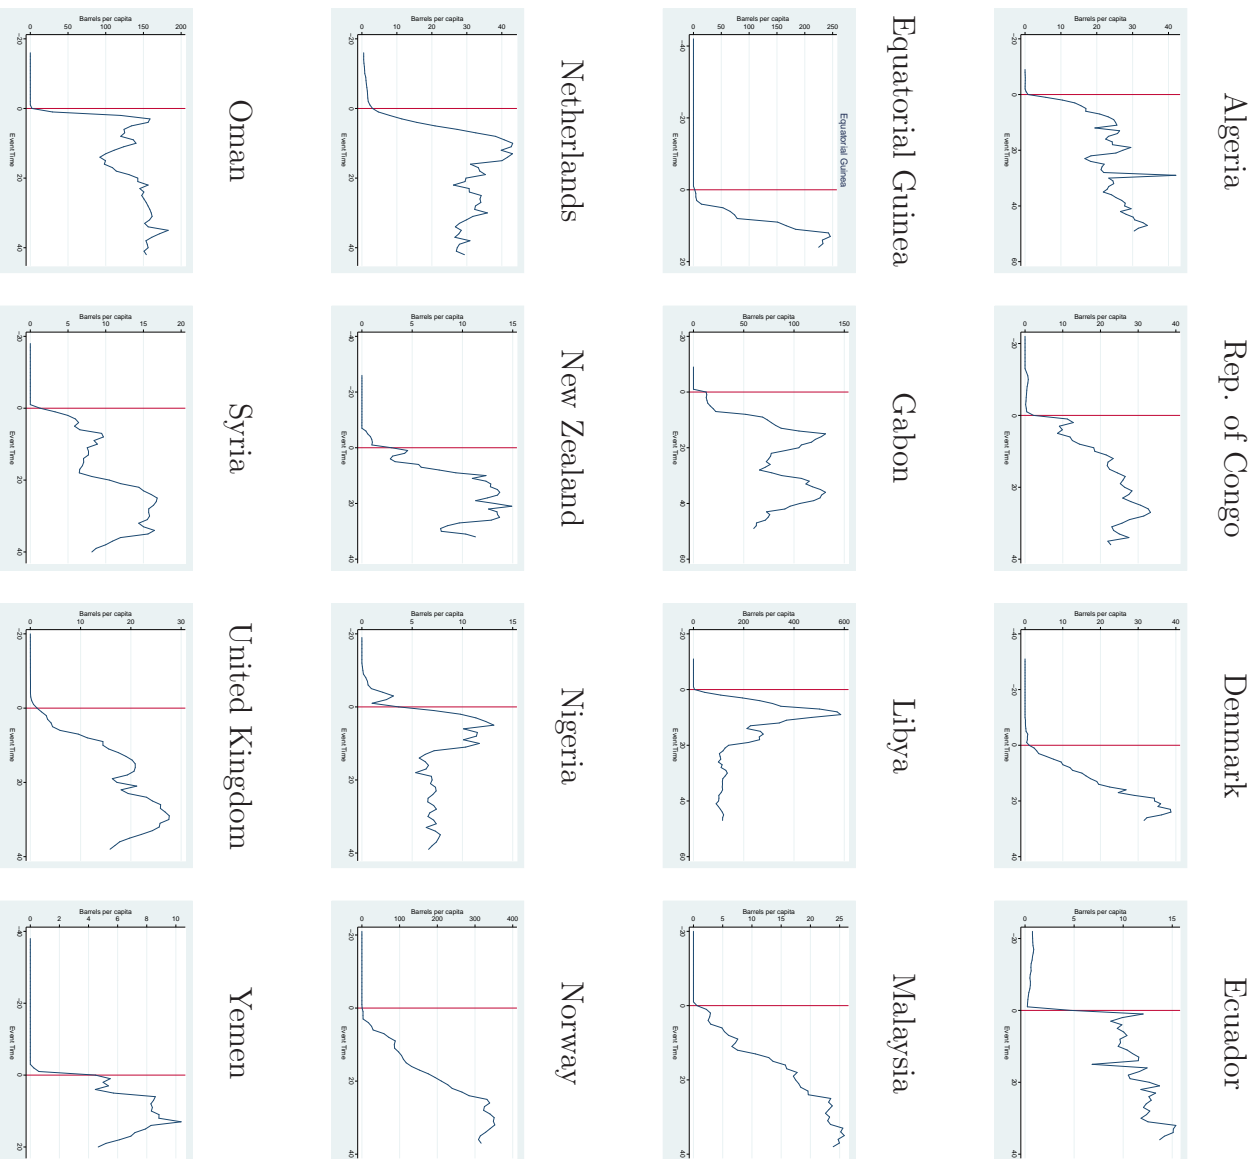


Table A1: Initial Characteristics as Predictors of Discovery

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Initial Pop.	0.109*** (0.021)	0.091** (0.029)						0.254* (0.096)
Area		0.025 (0.025)						0.016 (0.063)
Initial GDP/capita			0.072 (0.103)					0.116 (0.263)
Initial Democ. Score				-0.014 (0.019)				0.006 (0.031)
Initial Avg. Schooling					-0.021 (0.030)			0.036 (0.070)
Initial investment/GDP						-0.001 (0.003)		0.006 (0.006)
Fragmentation							0.152 (0.208)	1.044+ (0.536)
N	117	114	93	62	94	73	113	40
r2	0.27	0.29	0.15	0.14	0.19	0.09	0.14	0.50

Notes: The dependent variable is an indicator for making an initial oil discovery since 1950, or since 1960 in columns 4, 6 and 8. All covariates are measured at 1950, or at 1960 in columns 4, 6 and 8. All regressions include region fixed effects. White-Robust standard errors are reported. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table A2: 1970 Summary Statistics

	Treatment	Control
Average Real GDP/capita (1970)	8710 (9787)	7611 (12604)
Average Population (000s) (1970)	11088 (17293)	22000 (82431)
Average Democracy Score (1970)	3.1 (4.4)	3.2 (4.0)
Average Ethnic Fractionalization	0.43 (0.32)	0.45 (0.25)
Average years schooling (1970)	3.7 (3.1)	4.5 (2.6)
Average openness (1970)	.33 (0.16)	0.26 (0.19)
N	17	88

Table A3: Robustness, Panel A

	(1) Reduced T Group	(2) Increased T Group	(3) Production Lag Omit	(4) Event Lag Omit
Post	0.506* (0.206)	0.308* (0.126)	0.390* (0.177)	0.445* (0.203)
N	5900	6077	6112	6065
R^2	0.690	0.672	0.685	0.687

Robustness, Panel B

	(1) PWT sample	(2) Madd. with PWT sample	(3) Non-PPP GDP	(4) Madd. with WDI sample
Post	0.301 (0.196)	0.245 (0.178)	0.319 ⁺ (0.183)	0.307 ⁺ (0.168)
N	5353	5353	4375	4375
R^2	0.668	0.719	0.648	0.688

Notes: The dependent variable is the natural log of real GDP/capita. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.


Table A4: Robustness, Obtaining Stationary Residuals

	(1) Country Trends	(2) Growth Rate
Post	0.311** (0.111)	0.021* (0.010)
$\ln(\text{GDP/cap})_{t-1}$		-0.013* (0.006)
N	6195	6090
R^2	0.911	0.144


Notes: The dependent variable is the natural log of real GDP/capita in Column 1, and the year-on-year difference in log GDP/capita in column 2. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. +, *, **, *** represent significance at 10%, 5%, 1%, .1%, respectively.

Figure A2

Event Study, GDP Growth Rate



Event Study, GDP Growth Rate, Long Panel



Appendix B: Data Sources

Resource production data comes from UN Industrial Commodities Statistics, which includes production quantities of commodities for all countries and years from 1950-2001. GDP and population data covering the years 1950-2007 comes from Maddison Historical Statistics, which measures GDP in 1990 International Geary-Khamis dollars. I use Maddison in favor of Penn World Tables because the latter is missing data from 1950-1970 for many less-developed countries, including some in my treatment group, yielding a lack of pre-event data. However, I use Penn World Tables for purchasing power parity conversion factors (needed to find non-hydrocarbon GDP used in section 6.1) and investment to GDP ratio, which are not found in the Maddison data. I also use World Development Indicators to obtain GDP measured in constant 2000 US Dollars for the regression shown in Appendix Table A3, Panel B.

Oil discovery dates were found using the 2007 and 1994 editions of the Oil and Gas Journal Data Book, which lists all oil fields along with their discovery dates for each country. The discovery date used in this paper is the earliest given field discovery date. However, this method is not 100% reliable, as when comparing these dates with the UN production data, some countries (three from the treatment group) begin producing oil before the initial discovery. The most likely reason is that fields that have been shut down do not appear in the Oil and Gas Journal. It is also possible that especially small fields do not appear, since in all such cases, the amount produced is trivial until sometime after the first listed field is discovered. I have attempted to confirm discovery dates for each country in my treatment group with external sources, and just two adjustments have been made from the method described above.³²

Data on TFP, capital stock, labor force, and income per worker are drawn from the UN World Productivity Database, which provides data for a global sample of countries going back to 1960.

Education data is drawn from the Barro-Lee (2010) data set, which is a balanced panel of 145 countries, with data on several educational attainment variables measured every fifth year from 1950-2010. The variable of interest in this study is average years of schooling. The degree of democracy comes from the Polity IV index, a simple measure that ranges from 0 (hereditary monarchy) to 10 (consolidated democracy) for all countries from 1800-2009. For

³²In the United Kingdom, North sea oil was not discovered until 1970. A negligible amount of inland oil was produced before that, but since the North Sea bonanza was what made the UK a relevant producer, 1970 is a more appropriate date. Similarly, Ecuador produced a negligible amount until a major discovery in 1967 made it a major producer. See Figure 2 for illustrations of these two cases. Additionally, I adjusted the first non-zero production year in Algeria to 1958, even though the UN production data has Algeria producing trivial amounts of oil before then, before surging up in 1958. The oil history book “The Prize” pinpoints the discovery date as 1956, which is consistent with the Oil and Gas Journal.

ethnic fragmentation I use the data set compiled by Alesina, et al (2003). Their formula for fractionalization is the Herfindahl index, which ranges from zero (completely homogeneous) to 1 (every citizen is a different ethnic group).

Appendix C: List of Sample Countries by Region

Treatment countries are in bold.

East Asia

Cambodia, China, Hong Kong, Indonesia, Japan, Korea, Republic of, Laos, **Malaysia**, Mongolia, Philippines, Singapore, Taiwan, Thailand, Vietnam.

Eastern Europe

Albania, Bulgaria, Czech Republic, Hungary, Poland.

Latin America and the Caribbean

Costa Rica, Cuba, Dominican Republic, **Ecuador**, El Salvador, Guatemala, Honduras, Jamaica, Nicaragua, Panama, Paraguay, Puerto Rico, Uruguay.

Middle East and North Africa

Algeria, Djibouti, Egypt, Israel, Jordan, Lebanon, **Libya**, Morocco, **Oman**, **Syria**, Tunisia, Turkey, **Yemen**.

Northern Europe

Belgium, **Denmark**, Finland, France, Germany, Ireland, **Netherlands**, **New Zealand**, **Norway**, Sweden, Switzerland, **United Kingdom**.

Southern Europe

Greece, Italy, Portugal, Spain.

South Asia

Afghanistan, Bangladesh, India, Nepal, Pakistan, Sri Lanka.

Sub-Saharan Africa


Benin, **Botswana**, Burkina Faso, Burundi, Cameroon, Cape Verde, Central African Republic, Chad, **Republic of Congo**, Cote d'Ivoire, **Equatorial Guinea**, **Gabon**, Gambia, Ghana, Guinea, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritania, Mauritius, Mozambique, Namibia, Niger, **Nigeria**, Rwanda, Senegal, Somalia, Sudan, Swaziland, Tanzania, Togo, Uganda, Zambia, Zimbabwe.

Countries dropped for being resource rich before sample period


Angola, Argentina, Australia, Austria, Bahrain, Bolivia, Brazil, Brunei, Canada, Chile, Colombia, Dem. Rep. of Congo, Iran, Iraq, Kuwait, Mexico, Papua New Guinea, Peru, Qatar, Romania, Saudi Arabia, Sierra Leone, South Africa, Suriname, Trinidad & Tobago, United States, Venezuela.

Appendix D: Synthetic Control GDP/capita Results


Algeria




Botswana




Rep. of Congo




Denmark




Ecuador




Equatorial Guinea




Libya




Malaysia




Netherlands



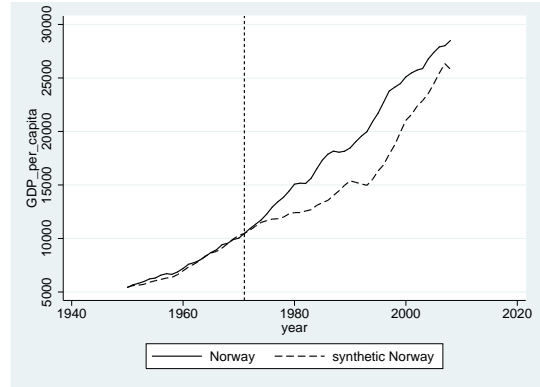
New Zealand



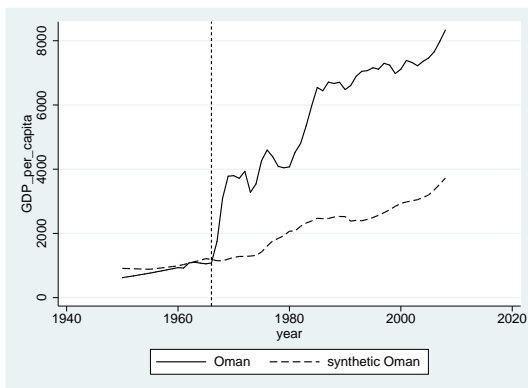
Nigeria




Norway




Oman



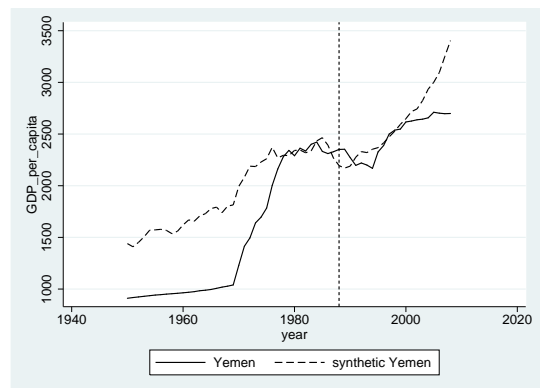
Syria



United Kingdom



Yemen



Appendix E: Synthetic Control Weights by Country

The following tables show the weights given to each control country in the synthetic control analysis for GDP per capita in Appendix Figure A1. These are not the same weights assigned for other outcomes analyzed with synthetic controls.

Algeria

Egypt	0.483
Israel	0.096
Jordan	0.29
Tunisia	0.13

Botswana

Burundi	0.277
Malawi	0.354
Rwanda	0.369

Denmark

Belgium	0.074
France	0.563
Sweden	0.25
Switzerland	0.113

Ecuador

Cuba	0.347
Dominican Republic	0.035
Guatemala	0.489
Nicaragua	0.008
Uruguay	0.121

Equatorial Guinea

Gambia	0.144
Lesotho	0.474
Liberia	0.048

Mauritania	0.071
Mauritius	0.025
Swaziland	0.237

Libya

Egypt	0.652
Jordan	0.348

Malaysia

Hong Kong	0.16
Indonesia	0.509
Philippines	0.216
Singapore	0.022
Thailand	0.091

Netherlands

Belgium	0.637
Germany	0.094
Sweden	0.166
Switzerland	0.103

Nigeria

Chad	0.83
Mauritius	0.012
Namibia	0.086
Sudan	0.072

Norway

Ireland	0.401
Sweden	0.599

New Zealand

Ireland	0.282
Sweden	0.589
Switzerland	0.129

Oman

Egypt	1
-------	---

Rep. of Congo

Liberia	0.779
Mozambique	0.114
Namibia	0.1
Swaziland	0.006

Syria

Djibouti	0.4
Israel	0.362
Lebanon	0.238

United Kingdom

Ireland	0.526
Switzerland	0.474

Yemen

Djibouti	0.419
Egypt	0.175
Lebanon	0.152
Tunisia	0.254