



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 OX1 3UQ, UK

ARTICLE INFO

Article history:

Received 4 February 2014

Received in revised form 28 March 2015

Accepted 1 April 2015

Available online 10 April 2015

Keywords:

Natural resource curse

GDP regressions

Synthetic controls

Oil discoveries

ABSTRACT

This paper evaluates the impact of major natural resource discoveries since 1950 on GDP per capita. 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.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Following the seminal work of [Sachs and Warner \(1995\)](#), a large literature has developed analyzing the existence and causes of the “resource curse”, the counter-intuitive finding that countries rich in natural resources tend to experience slower growth. Much of this literature has attempted to identify the market or institutional failures that might cause a resource curse to materialize. One commonly cited culprit is the so-called “Dutch disease”, whereby resource exports increase exchange rates, reducing the competitiveness of other exporting sectors ([Gylfason et al., 1999](#); [Sachs and Warner, 1995](#); [Sala-i Martin and Subramanian, 2003](#)). Others have argued that resource discovery subsequently weakens institutions and thus growth ([Leite and Weidmann, 2002](#); [Ross, 2001](#)), while another strand treats institutions as exogenous to resource wealth, and the interaction between resources and institutions explains divergent outcomes of resource-rich countries ([Mehlum et al., 2006](#); [Robinson et al., 2006](#); [Sarr et al., 2011](#)). [Caselli and 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 officials.¹ Other papers have argued that low levels of human capital ([Bravo-Ortega and De Gregorio, 2005](#); [Gylfason, 2001](#); [Papyrakis and Gerlagh, 2004](#)), lack

of investment ([Atkinson and Hamilton, 2003](#)), and increased risk of civil war ([Collier and Hoeffler, 1998](#)) also play a role.²

In recent years the cross-sectional design of [Sachs and Warner \(1995\)](#) has come under increased scrutiny,³ and more sophisticated designs have called the existence of the resource curse into question, although finding convincing exogenous variation remains a challenge. In this paper, I use major resource discoveries in the post-colonial period to evaluate the link between resources and GDP per capita. 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. 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). An objection to this design is the possibility that resource discovery is not in fact exogenous; [David and Wright \(1997\)](#) and [Bohn and 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 2](#) 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 would like to thank Giovanni Peri, Ann Stevens, Hilary Hoynes, Christopher Meissner, Douglas Miller, Alan Taylor and Steve Bond for their helpful feedback. I also thank the editor Gerard Padro i Miquel and two anonymous referees. Part of this work was completed while employed at the BP funded Centre for the Analysis of Resource Rich Economies at the University of Oxford.

¹ Other papers studying the corruption channel include [Torvik \(2002\)](#) and [Papyrakis and Gerlagh \(2004\)](#).

² For an extensive survey see [van der Ploeg \(2011\)](#).

³ For example [Alexeev and Conrad \(2009\)](#) and [Brunnschweiler and Bulte \(2008\)](#).

I find that newly resource-rich countries on average experience a large increase in GDP per capita levels that persists into the long-term. I find little to no pre-exploitation trends in GDP. Further, I find that the positive GDP effects are concentrated in developing countries, with small and insignificant effects when the sample is limited to OECD countries. This runs counter to much of the literature, which argues that countries with better institutions benefit more from natural resources. The reason is that while developed treatment countries have not performed poorly over the period studied, this is likely due to many factors besides natural resources, as their OECD counterparts have performed similarly well.

Extending the analysis beyond difference-in-differences, I use the synthetic control methodology developed by [Abadie and 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 for two reasons: it provides an additional 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 analyze the impact of resources on proximate causes of GDP. Although data limitations for these outcomes make precise estimation difficult, I find mixed evidence of long-run positive effects on capital stock, total factor productivity, labor force and human capital accumulation.

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.

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 as in [Sachs and Warner \(1995\)](#). [Brunnschweiler and Bulte \(2008\)](#) examine the relationship between growth rates and “subsoil assets” per capita measured in 1994 and 2000, and find a positive effect.⁴ [Alexeev and 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 and Poelhekke \(2010\)](#) point out that the available resource abundance measures are closely associated with current resource rents and thus endogenous to growth and income, and function more as 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 institutions.⁵ While similar concerns could be raised for

the initial discovery of resources as this paper uses, the fixed-effects design controls for time-invariant factors present before and after discovery.

A few recent studies have also incorporated oil discoveries into their specifications. [Cotet and Tsui \(2013b\)](#) argue that for most oil-producing countries, the most significant discoveries are concentrated over a few years. They evaluate the relationship between income and health measures and estimated oil endowments over different periods of time after this “peak discovery period”, and find positive effects.⁶ However, this method faces the same causal uncertainty as described above resulting from estimated oil endowments. [Cotet and Tsui \(2013a\)](#) 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 also used panel designs to study the relationship between resources and political outcomes. [Brückner 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).⁷ To my knowledge, few other papers have used panel data to examine the relationship between GDP and natural resources. [Collier and Goderis \(2012\)](#) use an error correction approach to estimate a specified long-run equilibrium relationship between growth and resource-export prices, finding a negative long-run effect of price increases. [Cotet and Tsui \(2013b\)](#) 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.

[Lei and Michaels \(2014\)](#) examine whether “giant” oil field discoveries leads to armed conflict, and 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, [Lei and Michaels \(2014\)](#) are primarily focused on the effects on civil conflict, while this paper focuses on economic indicators.

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. 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 evaluate the impact of resources on proximate causes of GDP (capital, TFP, labor force, education).

The rest of this paper proceeds as follows: the following section gives a brief historical overview of oil discovery and tests for endogenous discovery. [Section 3](#) describes the main data sources and defines

⁴ [Lederman and Maloney \(2003\)](#) take a similar approach, though using different measures of abundance, and also find positive effects.

⁵ [Michaels \(2011\)](#) uses a similar approach to study long-run outcomes of United States countries. This paper makes a convincing causal argument since the US has been extensively prospected.

⁶ [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.

⁷ See also [Aslaksen \(2007\)](#) and [Haber and Menaldo \(2011\)](#) for panel studies on oil and democracy.

the treatment group. Section 4 outlines the empirical design. Section 5 presents and discusses the results. Section 6 concludes.

2. Background of oil discovery

On the eve of World War 2, the global oil market was dominated by just a handful of countries and companies, with the U.S. alone accounting for almost two thirds of production.⁸ Following the war a convergence of factors led to a flurry of discoveries and a far more distributed industry. Governments scrambled to secure reserves after access to oil had proven critical in the allied victory. The spread of automobiles and the expanding plastic industry drove breakneck growth in commercial demand for oil. On the supply side, the industry structure became much less concentrated as barriers to entry fell. 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. Several important exploration advances were made over the 20th century, but perhaps the most important was deepwater drilling, which led to the North Sea boom and several major finds elsewhere.

To summarize, major oil discoveries in previously non-producing nations, while not completely random, have been driven to a great extent by global factors exogenous to any one country, particularly technological advances 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.

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 (1989). 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 effects 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 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 fixed effects in the main regression specifications 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.

3. Data and treatment assignment

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 and an imbalanced panel. However, I use Penn World Tables GDP data as a robustness check. Resource production data comes from UN Industrial Commodities Statistics, which includes production quantities of oil and gas for all countries and years from 1950–2001.¹⁰

The purpose of the treatment group 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 discoveries, 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 per capita 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 Fig. 1, a histogram showing the maximum level of annual

⁸ This section borrows from the canonical book on the history of oil *The Prize: The epic quest for oil, money & power* by Daniel Yergin (2011).

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

¹⁰ See Appendix B for a description of all other data sources used.

¹¹ Natural 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.

¹² Former 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.




Fig. 1. Maximum barrel production histogram.

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 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 A2).

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. Finally, 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, 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

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

Table 2
Summary statistics.

	(1) Treatment (full sample)	(2) Control (full sample)	(3) Treatment (non-OECD)	(4) Control (non-OECD)
Real GDP/capita (1950)	7.52 (1.01) [17]	7.08 (0.77) [88]	6.98 (0.64) [12]	6.84 (0.61) [72]
Population (000s) (1950)	7.90 (1.54) [17]	8.32 (1.59) [88]	7.52 (1.53) [12]	8.09 (1.57) [72]
Democracy score (1970)	3.9 (4.6) [16]	2.9 (3.9) [77]	1.2 (2.1) [11]	2.1 (3.2) [61]
Years schooling (1950)	3.3 (3.27) [14]	2.6 (2.16) [79]	1.08 (0.79) [9]	1.86 (1.44) [63]
Infant mortality (1955)	.13 (.07) [17]	.14 (.06) [87]	.17 (.04) [12]	.16 (.05) [71]
Ethnic fragmentation	.43 (.29) [17]	.45 (.27) [87]	.55 (.26) [12]	.51 (.26) [71]

Notes: Standard deviations are shown in parentheses and sample counts are shown in brackets.

all major mineral producers discovered their mineral wealth long before the period studied here.

Table 1 lists the 17 treatment countries. While somewhat small, the treatment group represents a reasonably representative geographic spread, and a variety of economic and political backgrounds. Table 2 presents summary statistics separately for treatment and control countries, both for the full sample and the non-OECD sample. When possible, statistics are shown for 1950, the start of the sample period. Infant mortality data begins in 1955, and democracy data coverage is poor until 1970. Ethnic fragmentation is only measured once per country at various times, but is presumably relatively stable over time. In the full sample, treatment countries had somewhat higher GDP per capita, schooling and democracy due to the prevalence of highly developed North Sea countries. For non-OECD countries (which is where positive GDP effects are found) GDP per capita is well-balanced, while democracy and schooling are actually lower in the treatment group, although these means are very low for both groups. Control countries do have a significantly higher average population, but this is skewed by a few very large countries, which the treatment group lacks. The treatment group in fact has a slightly larger median population than the control group.

The research design also requires an appropriate country-specific event year. One possible definition is the year of discovery, but GDP is not directly affected by the discovery of resources, but rather their extraction.¹³ Further, the initial discovery is not always the one that makes a country a major oil producer.¹⁴ 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

¹³ It 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 Fig. 3 suggests that, on average, this is not a major factor.

¹⁴ 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.




Fig. 2. Oil price and event years.

which production proceeds to surge upwards.¹⁵ 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 Fig. 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. Fig. 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.

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.

4. Empirical design

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

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

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 applicable.¹⁶

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

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

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. 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 Eq. (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.

4.1. Synthetic controls

An alternative way to measure the effect of resource discovery is the synthetic control methodology developed in Abadie and Gardeazabal (2003) and Abadie et al. (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.¹⁷ This both serves as an additional robustness check for the fixed effects model results, and gives greater

¹⁵ This pattern is likely associated with the Nigerian Civil War that lasted from 1967–1970.

¹⁶ One 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.

¹⁷ Gabon is excluded for reasons discussed in Section 5.3.

Table 3
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.

- + Significant at 10%.
* Significant at 5%.
** Significant at 1%.
*** Significant at .1%.

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 paper by *Abadie et al. (2010)*). 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.¹⁸ 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 with the diagonal elements representing the importance of each predictor. 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 3 presents the regression results for the main specification of Eq. (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.

¹⁸ For 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.

Table 4
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)
N	5650	4571	1079
R 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.

- + Significant at 10%.
* Significant at 5%.
** Significant at 1%.
*** Significant at .1%.

Column 2 shows the results of the main specification if only non-OECD countries are included, and column 3 if only OECD countries are included.¹⁹ 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 does not imply 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 in *Section 5.3*. 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 curse.²⁰

5.2. Event study

Table 4 shows the results of the event study specification of Eq. (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 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

¹⁹ Five of the 17 treatment countries are in the OECD: Denmark, Netherlands, New Zealand, Norway, United Kingdom.

²⁰ This is somewhat consistent with *Davis (2013)*, which replicated the Sachs & 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.




Fig. 3. Event study, GDP per capita.


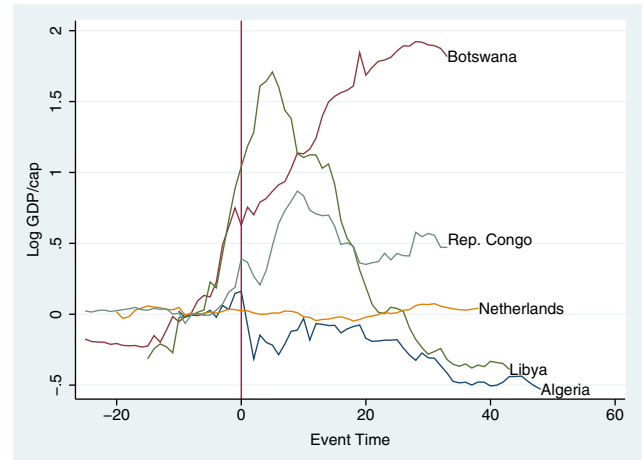


Fig. 4. Event study, GDP per capita, long panel.

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 Fig. 3.²¹

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 main 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 in Fig. 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

²¹ In 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 Fig. 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.



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

Fig. 5. Synthetic control results, selected countries.

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.

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.²² 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

²² There 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.

Table 5
Heterogeneous treatment effects: non-OECD treatments.

	(1) All countries	(2) All countries	(3) Non-OECD treatments	(4) 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)	
N	6018	6195	4779	4956
R 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.

+ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

to no effect on OECD 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 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.

Fig. 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 Fig. 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.

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 Eq. (1) and add interaction terms between the post-exploitation variable and various initial characteristics that may affect GDP and the effect of resources on GDP.²³

Column 1 of Table 5 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

Table 6
Obtaining stationary residuals.

	(1) Country trends	(2) Error correction model
Post	0.311** (0.111)	0.020* (0.009)
ln(GDP/cap) _{t-1}		-0.022*** (0.005)
Growth _{t-1}		0.163*** (0.039)
Growth _{t-2}		0.081** (0.025)
Growth _{t-3}		0.052* (0.026)
LR effect		(0.897)+ (0.513)
N	6195	5775
R ²	0.911	0.182

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.

+ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

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 Bravo-Ortega and De 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 5, 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. Obtaining stationary residuals

A possible concern with this paper's main results 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, in column 1 of Table 6 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. This 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.

²³ Initial log population, log GDP per capita and log of average years of education are measured in 1950. 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.

The second way to address non-stationarity is to substitute the GDP per capita growth rate as the dependent variable. I run the following “error correction” specification, which includes the lagged level of GDP per capita, and three lags of GDP growth to control for persistence in growth rates:

$$\Delta \ln(Y_{ct}) = \delta Post_{ct} + \theta Y_{t-1} + \sum_{j=1}^3 \beta_j \Delta Y_{t-j} + \alpha_c + \gamma_{rt} + \epsilon_{ct}. \quad (3)$$

This specification implies that in the long-run steady state, the effect of resource exploitation on the level of GDP per capita is $-\delta/\theta$.²⁴ Like the event study specification of Eq. (2), this specification allows analysis of short and long run effects on GDP, albeit in a more parametric way. The results are given in column 2 of Table 6. The “short run” effect of exploitation is 0.02 and significant at a 5% level. The long-run effect is 0.897 and significant at a 10% level.²⁵ While this point estimate is larger than the long-run level effects implied by the event study results in Fig. 3, the estimates are in fact not statistically different from each other based on a Students t-test.

As suggested by the event study results in Figs. 3 and 4, the observed growth effect is not permanent. This is likewise borne out in a growth rate event study specification, which is shown in Appendix Fig. A2 (these specifications also include the lagged GDP level and three lags of GDP growth.). 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 Fig. 4). Consistent with the level regressions, after the initial spike in growth following discovery, effects are close to zero in the long-term.

5.6. Robustness

In this section I run several robustness checks, for which all results are shown in Appendix Table A2. 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 A2, panel A I increase the cutoff to 20 barrels, which eliminates five treatment countries.²⁶ 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.²⁷ The effect is slightly smaller, but still statistically significant.

In columns 3 and 4 I perform robustness checks against the endogeneity of production 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 Eq. (1) using Penn World Table 7.0 GDP data. As mentioned in Section 3, PWT data does not have complete coverage going back to 1950

Table 7
Difference-in-differences: GDP proximate causes, non-OECD.

	(1)	(2)	(3)	(4)
	Capital stock	TFP	Labor force	Avg. years schooling
Post	0.50 ⁺ (0.30)	0.16 (0.16)	0.043 (0.032)	0.87 ⁺ (0.44)
N	2849	2870	3369	1040
R ²	0.827	0.167	0.925	0.881

Notes: The dependent variable is the natural log of the outcome given in the column header. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis.

- * Significant at 5%.
- ** Significant at 1%.
- *** Significant at .1%.
- + Significant at 10%.

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), which 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.

5.7. GDP proximate causes

As a final exercise, to shed light on the mechanisms of the positive effects found I evaluate the effect of resource exploitation on proximate causes of GDP that are typically included in traditional growth models. This analysis is carried out for non-OECD countries only, since this is the group for which resources were found to have a significant effect on GDP.

Table 7 reports the difference-in-differences results for capital, Total Factor Productivity (TFP), labor force and average years schooling. Due to data coverage limitations, some treatment countries are missing²⁸ and except for education the panels are not balanced, so

²⁴ See Collier and Goderis (2012) for a more detailed discussion of a similar specification.

²⁵ The standard error for the long-run effect $-\delta/\theta$ is found using the post-estimation command nlcom, which is based on the “delta method”, in Stata 13.

²⁶ Ecuador, New Zealand, Nigeria, Syria, and Yemen.

²⁷ Albania, Cameroon, Egypt, Hungary, Indonesia, and Tunisia.

²⁸ For 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.

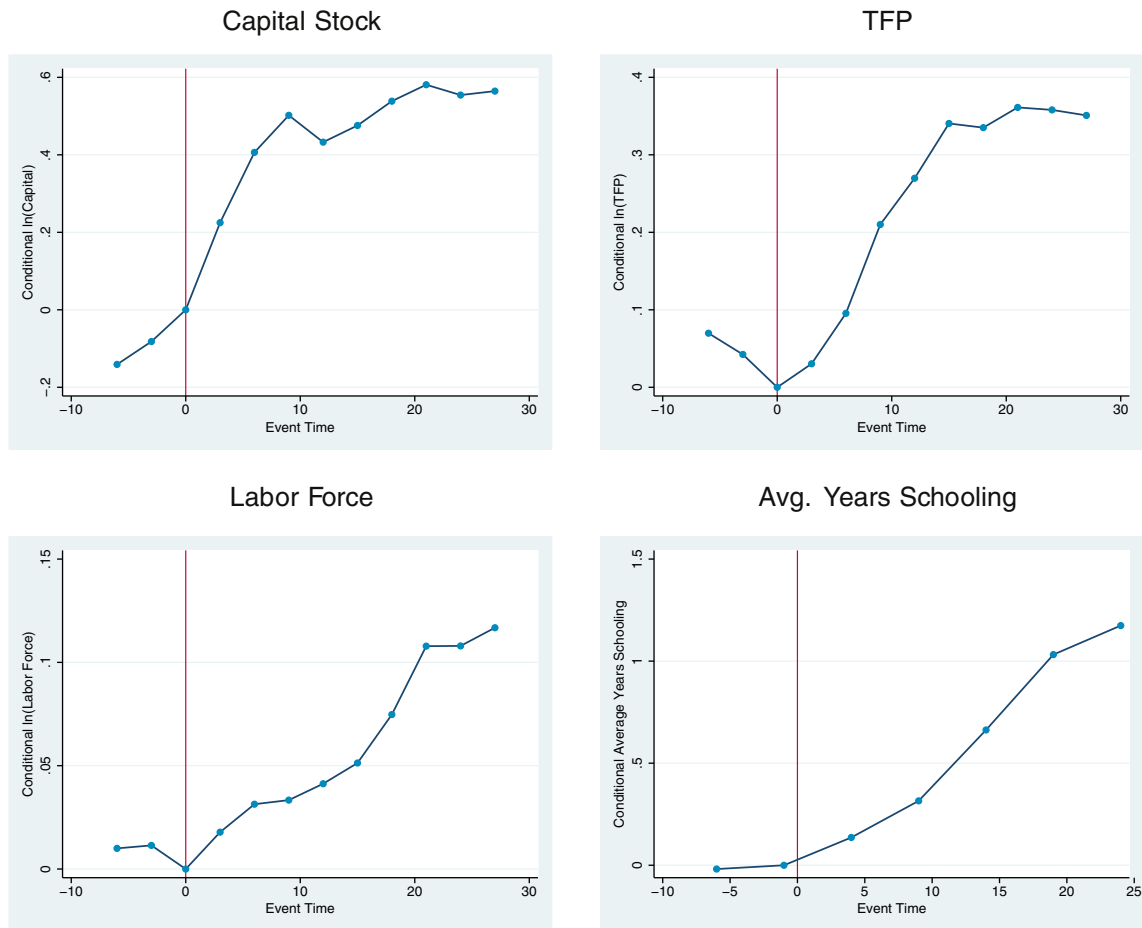


Fig. 6. Event study, GDP proximate causes, non-OECD.

these results should be viewed with caution. While all four outcomes see qualitatively significant effects, the standard errors are large and only the estimates for capital and schooling are significant at a 10% level. However, the graphical event study specification results shown in Fig. 6 are more suggestive of strong impacts. The first three graphs show effects from 9 years before exploitation to 26 years after (this range was chosen to obtain the longest time frame for which data coverage among the treatment group is acceptable, but again the panel is not perfectly balanced). 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 may 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 significant increases in both the size and quality of the labor force. The increase in labor force size may reflect an influx of migrant workers following the resource boom. In any case, the rise in workers mirrors a rise 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 educated,²⁹

²⁹ Since 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.

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. The exact event study estimates and standard errors are shown in Appendix Tables A3 and A4, and show that effects on long run levels relative to the reference period just before exploitation are generally statistically significant for all four outcomes.

6. 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 effects on GDP per capita that persist in the long term for developing countries, and no effect for developed countries. In evaluating the proximate causes of GDP, I find mixed evidence of positive long-run effects on capital formation, productivity, labor force 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).

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 and poverty would be a useful extension of the empirical designs used in this paper, but the demands of panel income distribution data is a difficulty that must be overcome.

Appendix A. Additional graphs and tables

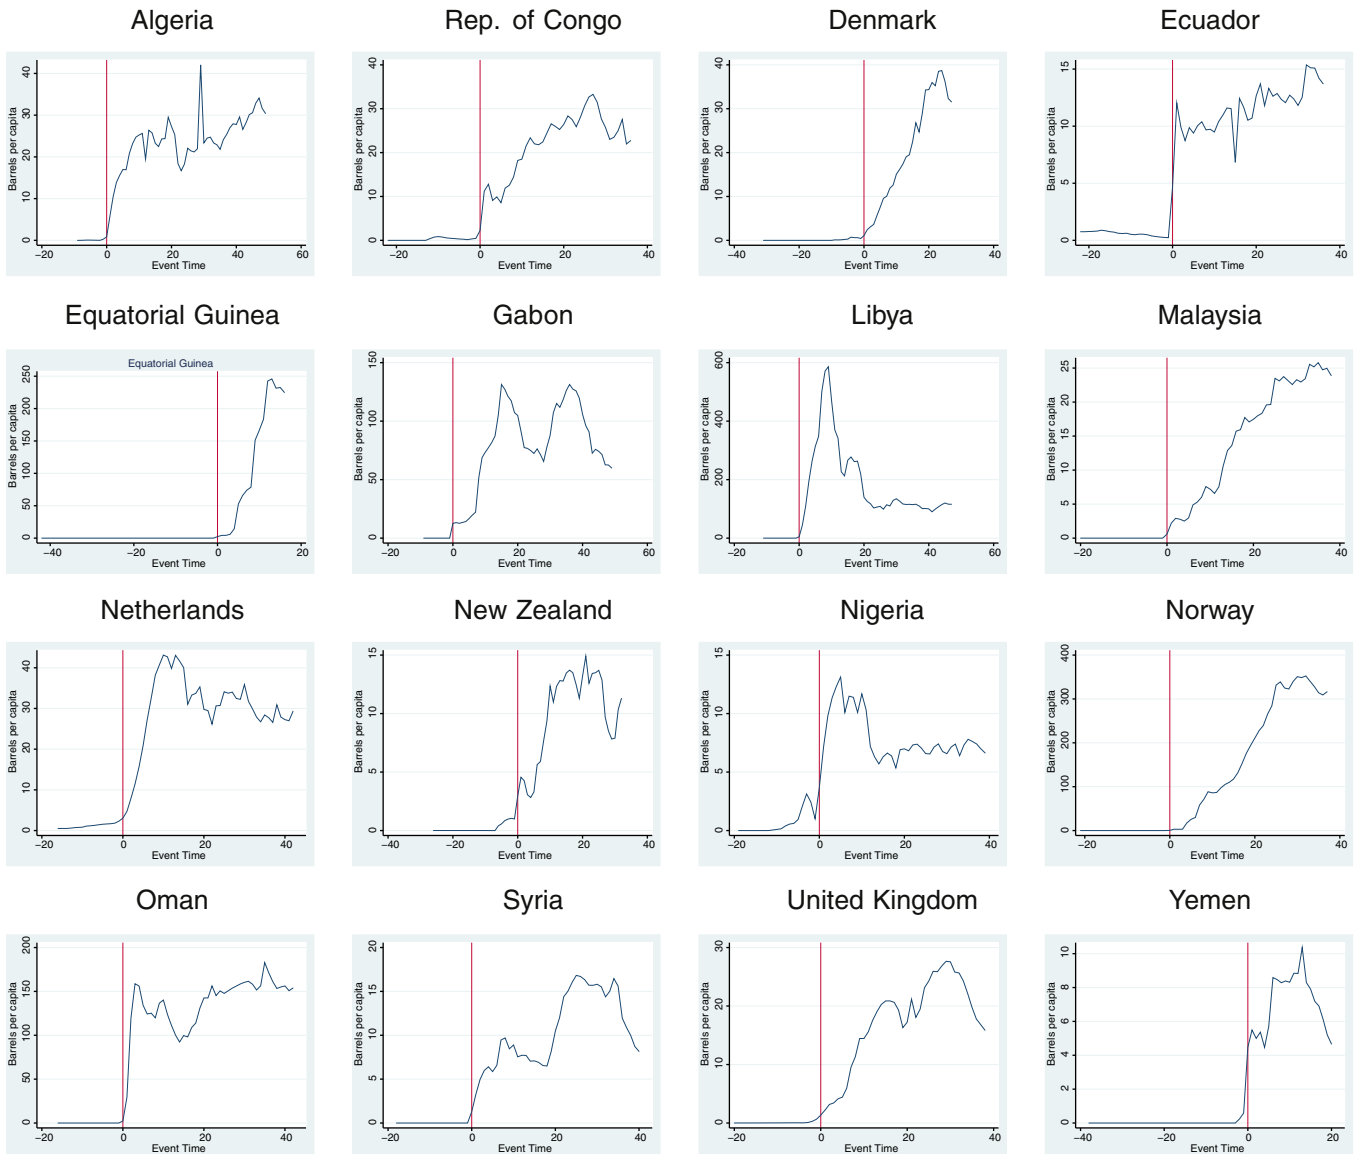


Fig. A1. Hydrocarbon production and event years.

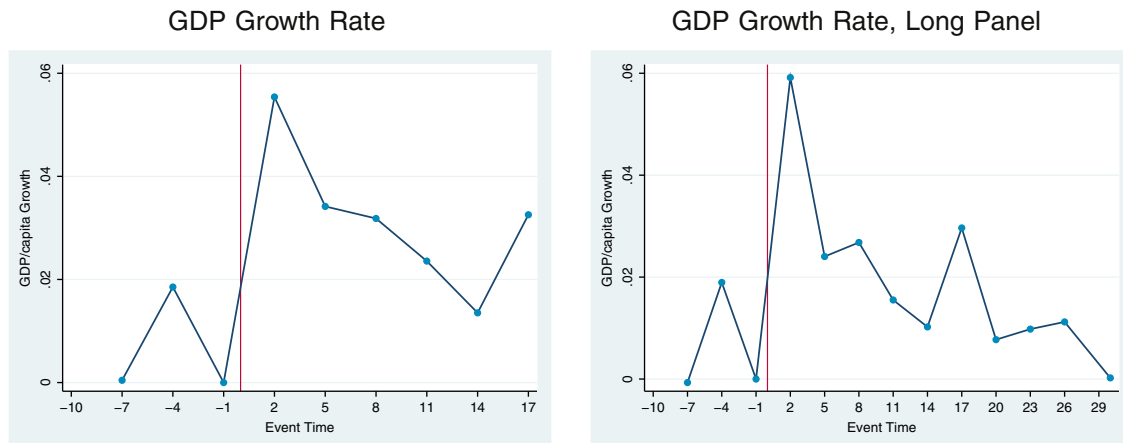


Fig. A2. Event studies, GDP growth rates.

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
R ²	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.

+ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

Table A2
Robustness checks.

Robustness, panel A					
	(1) Reduced T group	(2) Increased T group	(3) Production lag omit	(4) Event lag omit	(5) Year FEs
Post	0.506* (0.206)	0.308* (0.126)	0.390* (0.177)	0.445* (0.203)	0.354** (0.131)
N	5900	6077	6112	6065	6765
R ²	0.690	0.672	0.685	0.687	0.517
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.

+ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

Table A3
Event study: GDP proximate causes.

	(1) Capital	(2) TFP	(3) Labor Force
Exploitation year – 7–9	– 0.141 (0.142)	0.070 (0.055)	0.010 (0.012)
Exploitation year – 4–6	– 0.082 (0.079)	0.042 (0.052)	0.011 (0.011)
Exploitation year + 0–2	0.225 ⁺ (0.127)	0.030 (0.056)	0.018 (0.017)
Exploitation year + 3–5	0.406* (0.179)	0.095 (0.070)	0.031 ⁺ (0.017)
Exploitation year + 6–8	0.502* (0.239)	0.210* (0.082)	0.033** (0.011)
Exploitation year + 9–11	0.433* (0.214)	0.270* (0.114)	0.041** (0.013)
Exploitation year + 12–14	0.476* (0.227)	0.340* (0.141)	0.051** (0.016)
Exploitation year + 15–17	0.539* (0.232)	0.335* (0.146)	0.075*** (0.019)
Exploitation year + 18–20	0.581* (0.268)	0.361* (0.156)	0.108** (0.028)
Exploitation year + 21–23	0.554 ⁺ (0.301)	0.358* (0.148)	0.108** (0.036)
Exploitation year + 24–26	0.565 ⁺ (0.318)	0.351* (0.158)	0.117** (0.040)
<i>N</i>	2766	2786	3261
<i>R</i> ²	0.829	0.210	0.927

Notes: The dependent variable is the natural log of the outcome given in the column header. 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.

⁺ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

Table A4
Event study: education.

	(1) All countries
Exploitation year – 6–10	0.14 (0.16)
Exploitation year – 1–5	0.13 (0.21)
Exploitation year + 0–4	0.27 (0.26)
Exploitation year + 5–9	0.40 (0.33)
Exploitation year + 10–14	0.71 ⁺ (0.42)
Exploitation year + 15–19	1.02* (0.50)
Exploitation year + 20–24	1.12* (0.56)
<i>N</i>	993
<i>R</i> ²	0.888

Notes: The dependent variable is the natural log of average years schooling. 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.

⁺ Significant at 10%.

* Significant at 5%.

** Significant at 1%.

*** Significant at .1%.

Appendix B. Additional data sources

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.³⁰

Data on TFP and capital stock are drawn from the UN World Productivity Database, which provides data for a global sample of countries going back to 1960. Labor Force data comes from Penn World Tables 6.1 (which is also used to construct the World Productivity Database).

Education data is drawn from the Barro and Lee (2013) 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) with varying coverage for all countries from 1800–2009.

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

Infant Mortality data comes from the United Nations World Population Prospects, 2010 Revision.

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.


³⁰ In 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 Appendix Fig. A1 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.

Appendix D. Synthetic control GDP/capita results


Algeria




Botswana



Rep. of Congo




Denmark




Appendix D (continued).


Nigeria




Norway




Oman




Syria



United Kingdom



Yemen



Appendix D (continued).

Appendix E. Synthetic control weights by country

The following tables show the weights given to each control country in the synthetic control analysis. 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

References

- Abadie, A., Gardeazabal, J., 2003. The economic costs of conflict: a case study of the Basque Country. *Am. Econ. Rev.* 113–132.
- Abadie, A., Diamond, A., Hainmueller, J., 2010. Synthetic control methods for comparative case studies: estimating the effect of California's Tobacco Control Program. *J. Am. Stat. Assoc.* 105 (490).
- Alesina, A., Devleeschauwer, A., Easterly, W., Kurlat, S., Wacziarg, R., 2003. Fractionalization. *J. Econ. Growth* 8 (2), 155–194.
- Alexeev, M., Conrad, R., 2009. The elusive curse of oil. *Rev. Econ. Stat.* 91 (3), 586–598.
- Aslaksen, S. (2007). Corruption and oil: evidence from panel data. Unpublished Manuscript.
- Atkinson, G., Hamilton, K., 2003. Savings, growth and the resource curse hypothesis. *World Dev.* 31 (11), 1793–1807.
- Barro, R.J., 1989. Economic growth in a cross section of countries. Technical Report. National Bureau of Economic Research.
- Barro, R.J., Lee, J.W., 2013. A new data set of educational attainment in the world, 1950–2010. *J. Dev. Econ.* 104, 184–198.
- Bohn, H., Deacon, R.T., 2000. Ownership risk, investment, and the use of natural resources. *Am. Econ. Rev.* 526–549.
- Bravo-Ortega, C., De Gregorio, J., 2005. The relative richness of the poor? Natural resources, human capital, and economic growth. *Natural Resources, Human Capital, and Economic Growth* (January 2005). World Bank Policy Research Working Paper, (3484).
- Brückner, M., Ciccone, A., Tesei, A., 2012. Oil price shocks, income, and democracy. *Rev. Econ. Stat.* 94 (2), 389–399.
- Brunnschweiler, C., Bulte, E., 2008. The resource curse revisited and revised: a tale of paradoxes and red herrings. *J. Environ. Econ. Manag.* 55 (3), 248–264.
- Caselli, F., Michaels, G., 2013. Do oil windfalls improve living standards? Evidence from Brazil. *Am. Econ. J. Appl. Econ.* 5 (1), 1–31.
- Caselli, F., Tesei, A., 2011. Resource windfalls, political regimes, and political stability. Technical Report. National Bureau of Economic Research.
- Collier, P., Goderis, B., 2012. Commodity prices and growth: an empirical investigation. *Eur. Econ. Rev.* 56 (6), 1241–1260.
- Collier, P., Hoeffler, A., 1998. On economic causes of civil war. *Oxf. Econ. Pap.* 50 (4), 563.
- Cotet, A.M., Tsui, K.K., 2013a. Oil and conflict: what does the cross country evidence really show? *Am. Econ. J. Macroecon.* 5 (1), 49–80.
- Cotet, A.M., Tsui, K.K., 2013b. Oil, growth, and health: what does the cross-country evidence really show? *Scand. J. Econ.* 115 (4), 1107–1137.
- David, A.P., Wright, G., 1997. Increasing returns and the genesis of American resource abundance. *Ind. Corp. Chang.* 6 (2), 203–245.
- Davis, G.A., 2013. Replicating Sachs and Warner's working papers on the resource curse. *J. Dev. Stud.* 49 (12), 1615–1630.
- Gylfason, T., 2001. Natural resources, education, and economic development. *Eur. Econ. Rev.* 45 (4–6), 847–859.
- Gylfason, T., Herbertsson, T., Zoega, G., 1999. A mixed blessing. *Macroecon. Dyn.* 3 (02), 204–225.
- Haber, S., Menaldo, V., 2011. Do natural resources fuel authoritarianism? A reappraisal of the resource curse. *Am. Polit. Sci. Rev.* 105 (1), 1–26.
- Hodler, R., 2006. The curse of natural resources in fractionalized countries. *Eur. Econ. Rev.* 50 (6), 1367–1386.
- Jones, C.I., 1995. R & D-based models of economic growth. *J. Polit. Econ.* 759–784.
- Lederman, D., Maloney, W., 2003. Trade structure and growth. World Bank Policy Research Working Paper, (3025).
- Lei, Y.-H., Michaels, G., 2014. Do giant oilfield discoveries fuel internal armed conflicts? *J. Dev. Econ.* 110, 139–157.
- Leite, C., Weidmann, J., 2002. Does mother nature corrupt? Natural resources, corruption, and economic growth. *Governance, Corruption, and Economic Performance* pp. 159–196.
- Mehlum, H., Moene, K., Torvik, R., 2006. Institutions and the resource curse. *Econ. J.* 116 (508), 1–20.
- Michaels, G., 2011. The long term consequences of resource-based specialisation. *Econ. J.* 121 (551), 31–57.
- Mideksa, T.K., 2013. The economic impact of natural resources. *J. Environ. Econ. Manag.* 65 (2), 277–289.
- Papayrakis, E., Gerlagh, R., 2004. The resource curse hypothesis and its transmission channels. *J. Comp. Econ.* 32 (1), 181–193.
- Robinson, J., Torvik, R., Verdier, T., 2006. Political foundations of the resource curse. *J. Dev. Econ.* 79 (2), 447–468.
- Ross, M., 2001. Does oil hinder democracy? *World Polit.* 53 (3), 325–361.
- Sachs, J., Warner, A., 1995. Natural resource abundance and economic growth. NBER Working Paper.
- Sala-i Martin, X., Subramanian, A., 2003. Addressing the natural resource curse: an illustration from Nigeria. NBER Working Paper.
- Sarr, M., Bulte, E., Meissner, C., Swanson, T., 2011. On the looting of nations. *Public Choice* 148 (3), 353–380.
- Torvik, R., 2002. Natural resources, rent seeking and welfare. *J. Dev. Econ.* 67 (2), 455–470.
- Tsui, K.K., 2011. More oil, less democracy: evidence from worldwide crude oil discoveries. *Econ. J.* 121 (551), 89–115.
- Van der Ploeg, F., 2011. Natural resources: curse or blessing? *J. Econ. Lit.* 366–420.
- Van der Ploeg, F., Poelhekke, S., 2010. The pungent smell of red herrings: subsoil assets, rents, volatility and the resource curse. *J. Environ. Econ. Manag.* 60 (1), 44–55.
- Yergin, D., 2011. *The prize: The epic quest for oil, money & power.* Simon and Schuster.