

Smith, Brock

Working Paper

The resource curse exorcised: Evidence from a panel of countries

Working Paper, No. 13-3

Provided in Cooperation with:

University of California Davis, Department of Economics

Suggested Citation: Smith, Brock (2013) : The resource curse exorcised: Evidence from a panel of countries, Working Paper, No. 13-3, University of California, Department of Economics, Davis, CA

This Version is available at:

<https://hdl.handle.net/10419/79683>

Standard-Nutzungsbedingungen:

Die Dokumente auf EconStor dürfen zu eigenen wissenschaftlichen Zwecken und zum Privatgebrauch gespeichert und kopiert werden.

Sie dürfen die Dokumente nicht für öffentliche oder kommerzielle Zwecke vervielfältigen, öffentlich ausstellen, öffentlich zugänglich machen, vertreiben oder anderweitig nutzen.

Sofern die Verfasser die Dokumente unter Open-Content-Lizenzen (insbesondere CC-Lizenzen) zur Verfügung gestellt haben sollten, gelten abweichend von diesen Nutzungsbedingungen die in der dort genannten Lizenz gewährten Nutzungsrechte.

Terms of use:

Documents in EconStor may be saved and copied for your personal and scholarly purposes.

You are not to copy documents for public or commercial purposes, to exhibit the documents publicly, to make them publicly available on the internet, or to distribute or otherwise use the documents in public.

If the documents have been made available under an Open Content Licence (especially Creative Commons Licences), you may exercise further usage rights as specified in the indicated licence.

UC DAVIS

DEPARTMENT OF ECONOMICS Working Paper Series

The Resource Curse Exorcised: Evidence from a Panel of Countries

Brock Smith
UC Davis

March 18, 2013

Paper # 13-3

This paper evaluates the impact of major natural resource discoveries since 1950 on GDP per capita and other economic and social indicators. Using panel fixed-effects estimation and resource discoveries in countries that were not previously resource-rich, I find a positive effect on GDP per capita following extraction that persists in the long term, in contrast with much of the resource curse literature that uses cross-sectional designs. I also find positive effects on education levels, reductions in infant mortality, and negative effects on democratic institutions. I further test these outcomes with synthetic control analysis, yielding results consistent with the fixed-effects model.

Department of Economics
One Shields Avenue
Davis, CA 95616
(530)752-0741

http://www.econ.ucdavis.edu/working_search.cfm

The Resource Curse Exorcised: Evidence from a Panel of Countries

Brock Smith

University of California, Davis

brosmith@ucdavis.edu

October 2012

ABSTRACT

This paper evaluates the impact of major natural resource discoveries since 1950 on GDP per capita and other economic and social indicators. Using panel fixed-effects estimation and resource discoveries in countries that were not previously resource-rich, I find a positive effect on GDP per capita following extraction that persists in the long term, in contrast with much of the resource curse literature that uses cross-sectional designs. I also find positive effects on education levels, reductions in infant mortality, and negative effects on democratic institutions. I further test these outcomes with synthetic control analysis, yielding results consistent the fixed-effects model.

I would like to thank Giovanni Peri, Ann Stevens, Hilary Hoynes, Christopher Meissner, Douglas Miller, Alan Taylor, and participants at the UC Davis Applied Micro brownbag series and Macroeconomics brownbag series for helpful comments and advice.

1 Introduction

Since the seminal work of Sachs & Warner (1995), a near-consensus has 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.

Most of the subsequent literature on natural resources and growth has more or less taken this result as a given, and extended the Sachs & Warner (1995) 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, Sarr et al 2011). 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, Bulte et al 2011). Other papers have argued that low levels of human capital (Gylfason 2001, Papyrakis 2004, Ortega & Gregorio 2005), lack of investment (Atkinson & Hamilton 2004), corruption (Torvik 2002, Papyrakis 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 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 as a percentage 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 high dependence will likely impact subsequent performance. A more appropriate measure is what Brunnschweiler & Bulte 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.

However, use of resource abundance alone would not overcome the second flaw in the literature: the use of cross-sectional data. One of the more intuitive a priori explanations of the resource curse is the historical path-dependence story 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 (the Belgian Congo and Sierra Leone for example). 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 perform a panel fixed-effects analysis, in which newly resource-rich countries are compared to countries that remain resource-poor. This method provides plausibly exogenous resource shocks and a clearly defined treatment group, allowing for

estimates free of the endogeneity concerns discussed above. It is advantageous that nearly all recent discoveries have been in oil, a commodity noted for being susceptible to the curse. Oil's high spatial concentration makes it susceptible to control by one party, promoting inequality and civil strife (Wicke & Bulte, 2009).

This paper estimates the effect of the transition to resource abundance from non-resource abundance on subsequent growth rates, investment activity, social indicators, and political institutions in post-extraction periods. In contrast to the prior literature, I find that newly-resource rich countries on average experience a large short-term boost in growth and non-negative long-run effects on growth. Discoveries are also found to be associated with higher subsequent levels of education and lower levels of infant mortality. More consistent with past literature, I find that resource shocks have harmed political freedoms, as democratization has stagnated in treatment countries compared to the rest of the world. For all of these outcomes, the effects are concentrated in developing countries, with small and insignificant effects for developed countries.

Moving beyond simple difference in difference 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. The results on GDP per capita, education and infant mortality are consistent with the average positive effects found with the fixed effects model.

To my knowledge, few papers have used panel data to examine the relationship between growth and natural resources. Collier & Goderis (2007) use a panel cointegration approach to estimate a specified long-run equilibrium relationship between growth and resource-export prices, finding a negative long-run effect of price increases. 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 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 only use 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 is rich in oil. Second, Michaels & Lei (2011) is primarily focused on the effects on civil conflict, while this paper is focused on economic and social indicators.

The following section gives a brief historical overview of the oil industry and exploration. In section 3 I outline my empirical design. In section 4 I present and discuss results. Section 5 concludes.

2 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.¹ 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

¹This section borrows heavily from the canonical book about the history of oil *The Prize* by Daniel Yergin.

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 One, 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 Two the global market was completely dominated by just a handful of countries. As of 1938, just 8 countries² 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 (which by then included Saudi Arabia) 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

²USA, Mexico, Russia, Indonesia, Romania, Iraq, Iran, Venezuela.

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

A second factor was that the decades following the war were 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 of 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,³ 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 technology 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

³Drilling 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.

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, especially given the country and region-year fixed effects used in all regressions.

Can the data tell us anything about the likelihood of oil discovery? In Table 2 I 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.⁴ Each regression includes regional fixed effects (see appendix B for list of countries by region).

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

⁴There are 39 discovery countries by this definition, compared to 78 non-discovery countries.

among more people.

When I combine all predictors into one joint regression, I lose all but 42 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 country fixed effects in the main regression specifications (which use panel data, rather than a cross section as in Table 2) should largely control for any population and fragmentation effects, since relative population and fragmentation levels are fairly stable over time.

In these regressions 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 Empirical design

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

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

Where Y_{ct} 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 extraction 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. 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. However, as a robustness check I run a specification omitting New Zealand.

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 extraction 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 equation (1). For non-discovery countries, each indicator is equal to zero for the entire sample period.

Although each event-time coefficient is estimated with a small number of observations relative to the baseline difference-in-difference design, this method has two significant advantages. First, it checks for the existence of pre-existing trends that could lead to spurious difference-in-difference 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. In particular, I can identify differences in short-run effects 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 equivalent⁵ (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 B).⁶ 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 is a very low output that generates trivial wealth for the country, whereas 10 barrels generates 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 extraction 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

⁵Natural 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.

⁶Former 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.

only includes sample countries that achieved some non-zero level of production. 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.

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 represents a reasonably representative geographic spread, and a variety of economic and political backgrounds.

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 extraction. 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. 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 Figure 2, 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 3 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 endo-

geneity. 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 difference-in-difference specification with the years between discovery and the event year omitted, so that I am only comparing pre-discovery periods with post-extraction periods.

3.2 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 a “synthetic control”. I apply this method individually for each treatment country, essentially performing 16 different case studies.⁷ 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

⁷Gabon is excluded for reasons discussed in section 4.1.

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.⁸ 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.⁹ 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.

⁸For 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.

⁹the V matrix is found through a nested optimization procedure such that the mean squared prediction error of the pre-treatment outcome variable is minimized.

4 Empirical Results

4.1 GDP Growth

Column 1 of Table 3 shows the regression results for the specification in equation (1). Treatment countries saw an average post-extraction statistically significant increase in GDP per capita of approximately 36%. In Column 2 I run the same specification using Penn World Table 7.0 GDP data. As discussed in Appendix A, 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 are not able to 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 3 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. The estimate is actually slightly smaller than the PWT estimate, suggesting that, if anything, Maddison data underestimate the treatment effect.

In Columns 4 and 5¹⁰ I perform robustness checks against the endogeneity of production lag. Column 4 excludes the years between the first recorded oil field discovery and the first year of non-zero production. Column 5 excludes the years between the first recorded discovery and the actual event year. In both cases the results are similar, but the estimates are slightly larger. In Column 6, I check that the results are not due to differing long-term trends between treatment and control countries by adding country-specific trends to the specification. The estimate is slightly smaller but more precise.

Table 4 contains further robustness checks. In Column 1 I exclude Botswana, since

¹⁰Columns 4-6 and all remaining GDP specifications use Maddison data

it is the only non-hydrocarbon producing treatment country, which could raise questions of comparability. As Botswana was one of the fastest-growing countries in the world over the period studied thanks to its diamond industry, the average effect when excluding it is smaller but still economically significant, and statistically significant at a 10% level. In Column 2 I exclude New Zealand since, as discussed earlier, it does not have a natural regional classification. In this case the effect is slightly larger.

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 3 I increase the cutoff to 20 barrels, which eliminates five treatment countries.¹¹ The treatment effect with this reduced treatment group increases considerably, as would be expected given the higher intensity of treatment. In Column 4 I decrease the cutoff to five barrels, which adds six countries.¹² The effect is slightly smaller, but still significant at 5% level.

Column 5 of Table 4 shows the results of the main specification if only non-OECD treatment countries are included, and Column 6 if only the five OECD treatment countries (Denmark, Netherlands, New Zealand, Norway, United Kingdom) are included. Northern Europe is prominent in the treatment group, owing to the discoveries of North Sea oil and gas starting in the 1960s. Since these countries have long had what is widely considered to be “good” institutions, they may have been better equipped to convert resource wealth into high growth rates while avoiding the various traps discussed in the resource curse literature. On the other hand, since their economies were already highly developed, natural resource finds may not have as big an impact.

Columns 5 and 6 reveal a striking difference in effects between OECD and non-OECD countries. The effect for non-OECD treatments is considerably larger than the overall av-

¹¹Ecuador, New Zealand, Nigeria, Syria, and Yemen.

¹²Albania, Cameroon, Egypt, Hungary, Indonesia, and Tunisia.

erage effect, while the effect on OECD countries is actually negative, although 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 large average 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 curse.

Table 5 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 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. The same pattern, to a greater degree, is followed for the sample with OECD treatments excluded. With only OECD treatments included, there is a slight negative downward trend before the exploitation year and no long-term effect. The graphical representation of this table is shown in Figure 4.¹³

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 kick in. While we do see conditional growth flatten out after about 10 years, we do not see a deterioration in the time frame used in the base 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, I only need to drop

¹³in 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 4 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.

two treatment countries from the analysis (Equatorial Guinea and Yemen). The graphical result of this specification is shown in Figure 5. 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 the start of oil production). Although there is a slight downward trend at the end of the period, there is no evidence of a long-term curse.

Finally, I run the event study specification using the first recorded oil field discovery as the event date. In this specification, all countries that made a discovery on or after 1959 are counted as treatment countries (28 in all). I cut it off at 1959 so there is enough pre-event GDP data (which starts in 1950) for a balanced panel from 9 years before to 23 years after discovery. The purpose is to see if there is a pre-trend in GDP per capita that predicts discovery. The graphical result is shown in Figure 6. There is a very slight upward trend prior to discovery, after which growth flattens out before beginning an accelerating upward trend, presumably as more countries begin extraction further into event-time.

The results above show an average positive effect of resource exploitation on growth, but as will be shown in the synthetic control results, 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 resource effect on growth.¹⁴

Column 1 of Table 6 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

¹⁴Initial 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.

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 6, 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.

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.¹⁵ 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

¹⁵There 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 extraction, 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.

and seven years before the event. The weights making up each country's synthetic control for the GDP per capita analysis are shown in Appendix D.

The graphical results for each individual treatment country is shown in Appendix C. 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-difference results, in that we see a positive average effect in the short and long term. However, there is an interesting mix of outcomes. There are five countries (Botswana, Republic of Congo, Equatorial Guinea, Nigeria, and Oman) that perform significantly better than their synthetic counterpart. There are three countries (Algeria, New Zealand, Yemen) that do noticeably worse (although the long-term 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), at least in the long term. There is generally little to no effect on OECD countries, as robust 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 7 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 7 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.

4.2 Non-Hydrocarbon sector and Investment

While the results thus far have established a positive average effect of resource discovery on GDP per capita, they have been silent on the mechanism 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* online¹⁶ 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.¹⁷ Because of less extensive GDP data coverage in Penn World Tables, there are five treatment countries that do not have pre-extraction 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

¹⁶This is an average of equally weighted Dubai, Brent, and Texas crude oil prices.

¹⁷Resource value is converted to a real value by first using the PPP conversion factor given in PWT to obtain the value in current PPP adjusted dollars, and then multiplying by the ratio of real GDP to current PPP adjusted GDP. PWT 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.

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 7 shows the results for all regions. Conditional non-hydrocarbon GDP per capita was approximately 17% lower in the post-extraction period. To check whether developed countries were more effective in diversification, I run the same regression for the subsample that includes non-OECD treatments only, and likewise for OECD treatments. The effect is nearly the same for both groups-in fact the negative impact is slightly larger for OECD countries.

Due to the diminished treatment group, the event study specification yields especially noisy results, and are suggestive of a downward trend possibly causing spurious difference-in-difference estimates. For these reasons I examine time-variant effects on non-hydrocarbon GDP per capita using synthetic control analysis on the 10 available treatment countries. Figures 8 and 9 present all results for non-OECD and OECD treatments, respectively. Among the non-OECD countries, Malaysia's non-hydrocarbon GDP is substantially better than its synthetic control. For Ecuador there is little to no difference, but for the other three countries there is a negative, though somewhat erratic effect. For the OECD countries, there is a mostly positive effect in Norway, but the difference is close to zero by the end of the period. New Zealand has a large negative effect, but the pre-event difference indicates that New Zealand's pre-event trend is not well replicated by the synthetic control, so the result may still be spurious. The remaining three countries have moderately negative effects. Overall, the synthetic control results are consistent with the average negative effect estimated in Table 7. The negative outcome suggests that the estimated GDP per capita gains come exclusively from the direct contribution of oil and gas revenues, and that other sectors of the

economy may get crowded out.

Analyzing national investment activity can provide further context to the results discussed above. Penn World Tables provides estimates of what percentage of GDP in a given year goes to consumption, investment, or government spending. We can test whether it appears that oil revenues are reinvested or consumed by inserting the investment/GDP ratio as the dependent variable of the difference-in-difference model. Since this variable comes from PWT, it has the same coverage limitations as the non-hydrocarbon GDP measure, so the sample is the same as that in Table 7, except with Equatorial Guinea and Botswana added back in, bringing the treatment group to 12 countries.

As shown in Table 8, in the simple difference-in-difference specification, there is a positive effect on the investment ratio, implying that it was 3.3% higher in the post-extraction period, but it is not statistically significant. In the event study specification, with the graphical results shown in Figure 10, there is an upward trend in the investment ratio that begins well before extraction. This could reflect the investment in drilling infrastructure in the period between discovery and extraction-if so then the event as defined thus far may not be appropriate for analyzing investment. Hence I redefine the event year as the year of discovery. Because of the earlier event dates, I must further restrict the sample and specification to obtain a balanced panel. If I restrict the event-time window to six years before discovery to 17 years after, there are 10 treatment countries with complete data. The results for this regression are shown in Figure 11. There is a very slight downward trend in investment/GDP in the six years before discovery (which runs counter to what one would expect if discovery were endogenous), but after discovery it rises steadily to a peak of roughly eight percentage points higher than in the pre-period.¹⁸ While again noting the data limitations of these regressions, it appears that more investment is indeed taking place in treatment countries due to discovery, but the results on non-hydrocarbon GDP suggest that most or all of this

¹⁸in the difference-in-difference specification, not shown, the effect is positive but not significant.

extra investment is being put back into the oil and gas sectors.

4.3 Education and Infant Mortality

The positive results for growth still leave questions on distributional effects. To what degree is the windfall shared among the population? Is there a reduction in poverty? Unfortunately, the lack of data coverage for inequality and poverty rates during the period studied in this paper makes it impossible to address these questions directly with my chosen design. However, two outcomes that are worth analyzing do have sufficient data: education and infant mortality. A strong education system is widely seen as a bulwark against inequality. Educational outcomes can also give a sense as to whether governments have reinvested resource revenues into public goods. Infant mortality is another useful measure, as it is commonly used as a proxy for poverty.

Table 9 presents the regression results for average years of schooling. Three treatment countries (Equatorial Guinea, Nigeria, and Oman) are not covered in the Barro-Lee education data set, and so are not included in this specification. In the full sample there is no significant effect. For OECD treatments, which started the sample at high education levels, the effect is negative but insignificant. When we trim the treatment group to include only non-OECD countries, which mostly come from regions where educational metrics are typically exceptionally low at the beginning of the sample period, there is a statistically significant increase of 0.89 average years of schooling. The results of the event study specification for non-OECD countries are shown in Figure 12.¹⁹ There is no trend in conditional schooling prior to exploitation, but it then begins a steady rise that reaches about 1.2 years of additional schooling by the end of the period. For additional context, Figure 13 shows unconditional educational averages for treatment and non-treatment countries in real time over the sample period, with all OECD countries excluded from each group. In 1950, before

¹⁹Since 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.

any exploitation events, treatment countries have lower average education levels than non-treatments, but over time slightly surpass the non-treatments. Thus it appears that at least some government resource revenues are reinvested to public services.

I also analyze the effect of extraction on education using synthetic control analysis for each country. With education measured every five years, the event year used is the first observation after the true event year. The control variables used to match with the treatment country are average years schooling 5, 10, and 15 years prior to the event (or only 5 and 10 if the event is too early), and population, fragmentation, and GDP per capita 5 years before the event. The results are shown in Figures 14 and 15, again split up by OECD and non-OECD. In the non-OECD group, there are two countries with a near-zero effect, and one with negative effect of a little less than two years by the end of the period. The remaining five²⁰ all experience positive effects, with Gabon and Botswana in particular strongly outperforming their synthetic controls. The non-OECD treatments are split between positive and negative results, with the effects being relatively small. The synthetic control results are again consistent with the average treatment effects estimated in Table 9.

For infant mortality, I exclude OECD countries altogether because by 2000, all OECD countries in the world had reduced infant mortality to less than 1% (compared to a world average of 5%), so natural resources could not be a factor for this particular variable and subset of countries. Table 10 shows the infant mortality results for the 12 non-OECD treatments. Conditional post-exploitation infant mortality was 1.1 percentage points lower than the pre-period rate, significant at a 10% level. Figure 16 shows the graphical event study results. Since infant mortality data does not begin until 1955, I must exclude two countries (Gabon and Algeria) to obtain a balanced panel going back to ten years before exploitation, leaving ten treatment countries. There is a steady downward trend in conditional infant

²⁰Yemen is not included in this analysis because it had lower pre-extraction education levels than all control countries in its region, so a suitable synthetic control cannot be constructed. Similarly, New Zealand has substantially higher levels of pre-extraction education than all controls, so it is likewise excluded.

mortality throughout the period, but in this case the trend begins before the exploitation event, so the effect may be spurious.

Due to the apparent pre-trends, I further test the impact on infant mortality using synthetic controls, which has the advantage of individually controlling for pre-trends. The control variables are defined similarly to those used for education. The results, shown in Figure 17,²¹ are mixed, with three countries experiencing positive effects, two experiencing little to no effect, and six experiencing negative effects, with the negative impacts being generally larger in magnitude than the positive ones. Hence resource extraction does appear on average to be associated with reductions in infant mortality.

The results for education and infant mortality run somewhat counter to the conventional wisdom. While it is still possible, perhaps likely, that most of the resource rents in treatment countries have been captured by the elite, the results suggest that the poor and the wider population, on average, are at least slightly better off than they would otherwise be. This is not to paint too rosy a picture-Nigeria and Equatorial Guinea still rank among the world's worst in infant mortality, but these countries are in a troubled region and were likewise among the worst prior to discovering oil. For these examples a "curse" might be better described as a failed opportunity to develop.

4.4 Democracy

The last outcome I examine is whether the degree of democracy is affected by resource discovery and extraction. For similar reasons as for infant mortality, OECD treatment countries are excluded from this regression. As it happens, each OECD treatment had a highest-possible 10 democracy rating throughout the sample period, so the assumption that democracy in these countries would change in a similar way to control countries (which increase in democracy during the period) absent the treatment is violated. Further, three treatment countries (Algeria, Gabon, and Yemen) do not have democracy data prior to the

²¹Yemen is again excluded as its pre-extraction infant mortality rate is higher than all regional controls.

event data, and are thus dropped. This leaves only nine treatment countries to identify the effect on democracy, but the results are still suggestive.

Column 1 of Table 11 shows the democracy score results for the full non-OECD sample. Consistent with past literature (Barro 1999, Ross 2001), there is a negative effect on democratic institutions following exploitation, significant at a 10% level. The coefficient implies that resource discovery and exploitation causes the democracy score to be 1.4 points lower on average in the post-exploitation period. Since the democracy score is not a continuous measure and arguably ordinal, in Column 2 I use a democracy dummy as the dependent variable. The dummy is equal to one if the Polity IV score is greater than or equal to 6, following Collier & Rohner (2008). The coefficient suggests that a treatment country is 19% less likely to be a democracy post-exploitation than if no discovery had occurred, however the estimate is not statistically significant. In columns 3 and 4 I further restrict the sample to the 10 African and Middle Eastern treatments only, as these are notoriously non-democratic regions. In this case the estimates are similar but more precise and both estimates are significant at a 5% level.

Data coverage limitations preclude an event study specification, but Figure 18 sheds some light on the results of Table 11. This figure shows the unconditional average democracy scores for treatment and control countries through real time, with both treatment and control countries limited to Africa and the Middle East. Rather than democracy deteriorating after extraction as might be expected, there is a failure to democratize along with much of the region. In 1970 (the first year shown since that is the first year all treatment countries, save Yemen, had democracy scores available) the average democracy score was roughly equal for both treatment and control countries. But in the late 1980s/early 1990s (past the event year of nearly all treatment countries), the world saw a dramatic increase in democratization following the collapse of the Soviet Union, while the treatment group remained at its low average democracy level.

Another way to see this is with the democracy indicator variable. In 1970, there were 32 control countries in Africa and the Middle East (for which democracy data exists) that were not democracies. Of these, 16 were democracies as of the last measurement in 2008. Out of the eight treatment countries in the same regions that were not democracies,²² none of them were in 2008.

5 Conclusion

This paper takes a novel approach to estimating the impact of natural resources, using modern discoveries, longitudinal data, and more sophisticated empirical methods to provide a more rigorous test of the existence of the resource curse than has been heretofore performed. Contrary to the past literature, I find positive growth effects for modern resource discoveries in lesser developed countries, and no effect for highly developed countries. I also find positive effects on education levels and reductions in infant mortality, indicating that the state of public investment and poverty, while still deplorable in some of the treated countries, are at least on average better off than they otherwise would have been. However, other findings provide a more nuanced picture of the benefits of resource wealth. As a whole, treated countries failed to invest the windfall into diversification of the economy, and the non-hydrocarbon sector suffered. This raises particular worries for countries where the resources are close to being exhausted, such as Gabon. In addition, the results on democracy indicate that resource windfalls have curtailed the spread of political freedoms.

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. As-

²²Botswana was a democracy in 1970, and Yemen was not measured.

sessing the impact of resource extraction 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.
- 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*.
- Bravo-Ortega, C. and De Gregorio, J. (2005). The Relative Richness of the Poor? Natural Resources, Human Capital, and Economic Growth. *World*.
- 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.
- Collier, P. and Goderis, B. (2007). Commodity prices, growth, and the natural resource curse: Reconciling a conundrum. Technical report, CSAE Working Paper.
- 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.
- 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.
- Hodler, R. (2006). The curse of natural resources in fractionalized countries. *European Economic Review*, 50(6):1367–1386.
- Lei, Y. and Michaels, G. (2011). Do giant oilfield discoveries fuel internal armed conflicts?
- 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.
- 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*.
- 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.
- Torvik, R. (2002). Natural resources, rent seeking and welfare. *Journal of Development Economics*, 67(2):455–470.
- Yergin, D. (2008). *The prize: The epic quest for oil, money, & power*. Free Pr.

Table 1: Treatment Countries

Country	Event Year	Initial Discovery	First 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: Initial Characteristics as Predictors of Discovery

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Initial Pop.	0.111*** (0.021)	0.094** (0.031)						0.254* (0.096)
Area		0.023 (0.030)						0.016 (0.063)
Initial GDP/capita			0.115 (0.107)					0.116 (0.263)
Initial Democ. Score				-0.014 (0.019)				0.006 (0.031)
Initial Avg. Schooling					-0.018 (0.029)			0.036 (0.070)
Initial investment/GDP						-0.002 (0.004)		0.006 (0.006)
Fragmentation							0.100 (0.211)	1.044+ (0.536)
N	113	111	91	61	91	71	110	40
r2	0.27	0.28	0.16	0.14	0.19	0.10	0.13	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 3: Difference in Difference: GDP/capita

	(1)	(2)	(3)	(4)	(5)	(6)
	Main Spec.	PWT data	Maddison with PWT sample	Production Lag Omitted	Event Lag Omitted	Country Trends
Post	0.350*	0.297	0.245	0.390*	0.445*	0.311**
	(0.157)	(0.196)	(0.178)	(0.177)	(0.203)	(0.111)
<i>N</i>	6195	5353	5353	6112	6065	6195
<i>R</i> ²	0.684	0.682	0.719	0.685	0.687	0.911

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. + indicates significance at a 10% level, * at a 5% level, and ** at a 1% level.

Table 4: Difference in Difference: GDP/capita - Robustness Checks

	(1)	(2)	(3)	(4)	(5)	(6)
	Botswana exc.	NZ exc.	reduced T group	increased T group	non-OECD Treatments	OECD Treatments
Post	0.252 ⁺	0.386*	0.535**	0.308*	0.518*	-0.102
	(0.138)	(0.161)	(0.201)	(0.126)	(0.200)	(0.094)
<i>N</i>	6136	6136	6077	6077	5900	5487
<i>R</i> ²	0.696	0.688	0.685	0.672	0.685	0.728

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. + indicates significance at a 10% level, * at a 5% level, and ** at a 1% level.

Table 5: Event Study: GDP/capita

	(1)	(2)	(3)
	Full Sample	non-OECD Treatments	OECD Treatments
Exploitation year - 7-9	-0.027 (0.033)	-0.049 (0.044)	0.044** (0.017)
Exploitation year - 4-6	-0.002 (0.027)	-0.008 (0.037)	0.018 (0.011)
Exploitation year + 0-2	0.108** (0.037)	0.160*** (0.046)	-0.031+ (0.017)
Exploitation year + 3-5	0.245** (0.090)	0.346** (0.116)	-0.019 (0.034)
Exploitation year + 6-8	0.339** (0.121)	0.470** (0.155)	-0.010 (0.038)
Exploitation year + 9-11	0.419** (0.145)	0.579** (0.184)	-0.013 (0.050)
Exploitation year + 12-14	0.433** (0.163)	0.599** (0.206)	-0.019 (0.067)
Exploitation year + 15-17	0.434** (0.162)	0.602** (0.205)	-0.026 (0.075)
<i>N</i>	5650	5515	5327
<i>R</i> ²	0.701	0.706	0.718

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

Table 6: 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)	6.85*** (0.98)	7.08*** (0.70)
Post*(log pop.)	-0.11+ (0.062)	-0.17*** (0.046)	-0.23*** (0.065)	-0.22*** (0.027)
Post*(log GDP/cap)	-0.62*** (0.11)	-0.67*** (0.075)	-0.60*** (0.067)	-0.67*** (0.067)
Post*(log fragmentation)	-0.15+ (0.081)	-0.018 (0.056)	-0.22* (0.11)	0.040 (0.038)
Post*(log inf. mortality)	-0.058 (0.21)	-0.41*** (0.11)	0.33 (0.43)	0.096 (0.26)
Post*(log avg. yrs school)	0.18+ (0.096)		0.26* (0.12)	
<i>N</i>	6018	6195	5723	5900
<i>R</i> ²	0.730	0.734	0.722	0.729

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. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table 7: Difference in Difference: Non-hydrocarbon GDP/capita

	(1) All countries	(2) non-OECD Treatments	(3) OECD Treatments
Post	-0.17** (0.060)	-0.19* (0.091)	-0.15+ (0.077)
<i>N</i>	5993	5698	5724
<i>R</i> ²	0.718	0.708	0.722

Notes: The dependent variable is the log of real non-resource generated GDP, except for Column 2 in which it is the log of total GDP, using Penn World Tables data. Robust standard errors clustered at the country level are reported in parenthesis. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table 8: Difference in Difference: Investment Ratio

	(1) All countries
Post	4.625 (4.551)
<i>N</i>	6234
<i>R</i> ²	0.136

Notes: The dependent variable is the investment-to-GDP ratio. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table 9: Difference in Difference: Average Years Schooling

	(1) Full Sample	(2) non-OECD Treatments	(3) OECD Treatments
post	0.41 (0.36)	0.89* (0.43)	-0.47 (0.38)
<i>N</i>	1365	1300	1248
<i>R</i> ²	0.881	0.885	0.891

Notes: The dependent variable is average years of schooling. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table 10: Difference in Difference: Infant Mortality

	(1) Non-OECD treatments
post	-0.011 ⁺ (0.0062)
<i>N</i>	1464
<i>R</i> ²	0.846

Notes: The dependent variable is infant mortality rate in percentage points. All regressions include country and region-year fixed effects. OECD countries are not included. Robust standard errors clustered at the country level are reported in parenthesis. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Table 11: Difference in Difference: Democracy Score

	(1) Democ. Score	(2) Democ. Dummy	(3) Democ. Score, Afr/ME only	(4) Democ. dummy, Afr/ME only
Post	-1.438 ⁺ (0.810)	-0.190 (0.145)	-1.287* (0.544)	-0.178* (0.084)
<i>N</i>	5184	5184	2243	2243
<i>R</i> ²	0.433	0.364	0.213	0.131

Notes: The dependent variable in Columns 1 and 3 is the Polity IV Democracy rating, and in Columns 2 and 4 a democracy indicator. Regressions do not include OECD treatment countries. All regressions include country and region-year fixed effects. Robust standard errors clustered at the country level are reported in parenthesis. + indicates significance at a 10% level, * at a 5% level, ** at a 1% level, and *** at a .1% level.

Figure 1: Maximum Barrel Production Histogram

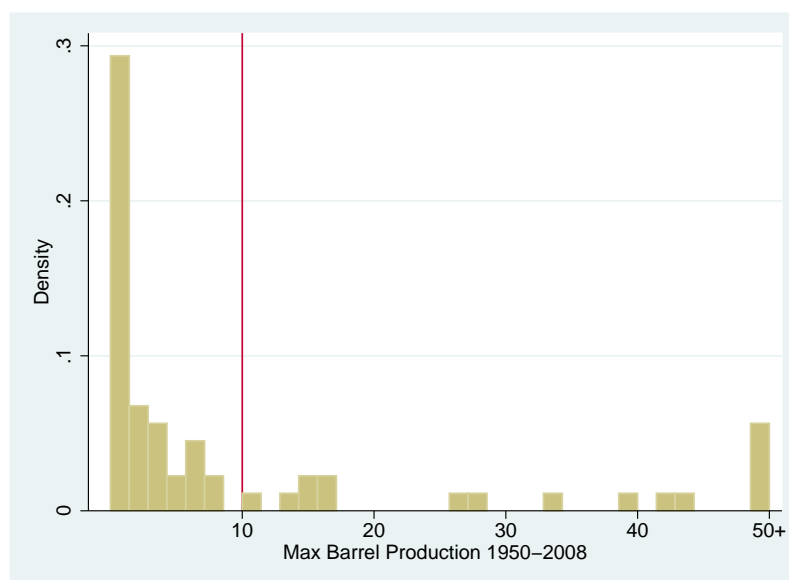


Figure 2: Hydrocarbon Production and Event Years




Figure 3: Oil Price and Event Years




Figure 4: Event Study, GDP per capita

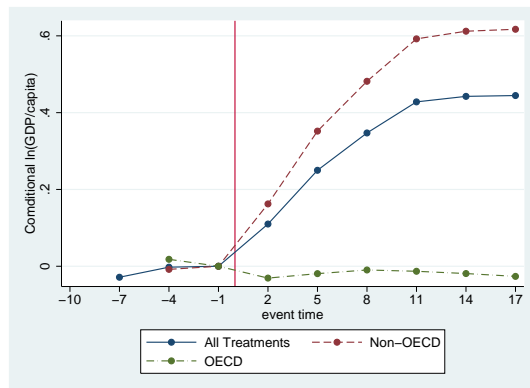


Figure 5: Event Study, GDP per capita, Long Panel




Figure 6: Event Study, GDP per capita, Discovery Event Year




Figure 7: Synthetic Control Results, Selected Countries

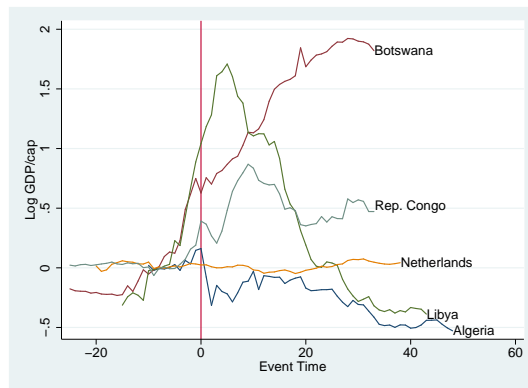


Figure 8: Synthetic Control Results, Non-hydrocarbon GDP per capita, non-OECD

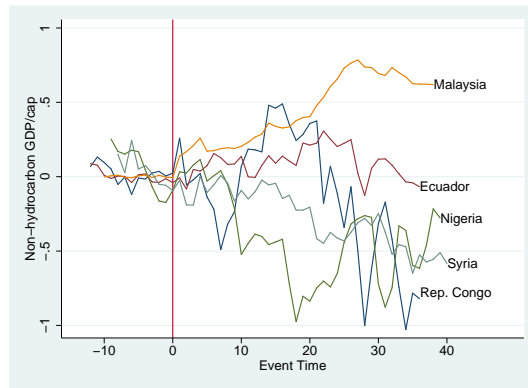


Figure 9: Synthetic Control Results, Non-hydrocarbon GDP per capita, OECD

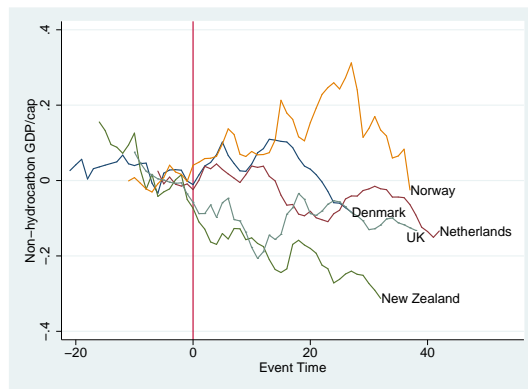


Figure 10: Event Study, Investment/GDP




Figure 11: Event Study, Investment/GDP, Discovery Event Year

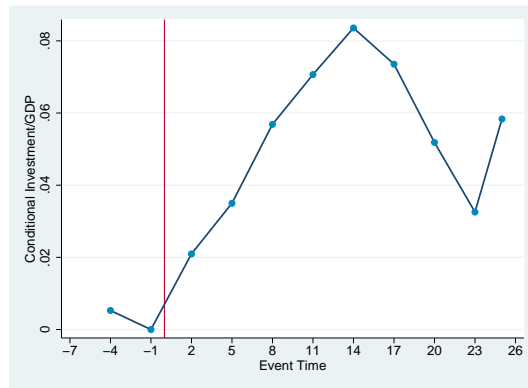


Figure 12: Event Study, Average Years Schooling, Non-OECD Treatments




Figure 13: Unconditional Average Years Schooling, Non-OECD countries

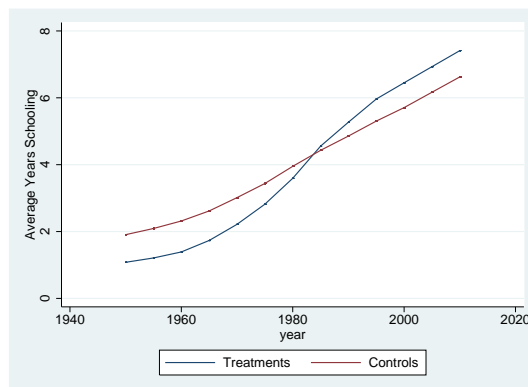


Figure 14: Synthetic Control Results, Average Years Schooling, non-OECD

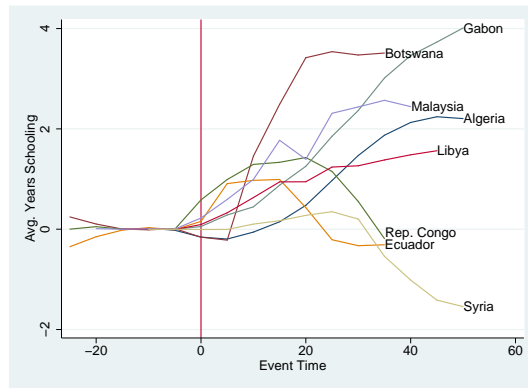


Figure 15: Synthetic Control Results, Average Years Schooling, OECD




Figure 16: Event Study, Infant Mortality rate, Non-OECD Treatments




Figure 17: Synthetic Control Results, Infant Mortality, non-OECD

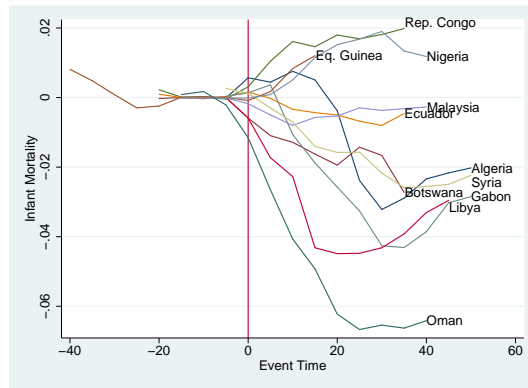



Figure 18: Unconditional Democracy Score, Middle East and Africa Only



6 Appendix A: Data Sources

For Online Publication

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 investment data, and also for purchasing power parity conversion factors (needed to find non-hydrocarbon GDP used in section 4.2), which are not found in the Maddison data. Because of the diminished data coverage, some treatment countries are cut from these parts of the analysis, which will be discussed further in the results section.

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.²³

²³In 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,

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.

Infant mortality data comes from the United Nations 2010 revision of World Population Prospects. Like education, measurements are made every five years, but start in 1955.

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 ethnic fragmentation I use the data set compiled by Alesina, et al (2003). Their formula for fragmentation is the Herfindahl index, which ranges from zero (completely homogeneous) to 1 (every citizen is a different ethnic group). This is only measured once per country, and in a range of years from 1979-2001.

7 Appendix B: List of Sample Countries by Region

For Online Publication

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.

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.

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, The, Ghana, Guinea, Kenya, Lesotho, Liberia, Madagascar, Malawi, Mali, Mauritania, Mauritius, Mozambique, Namibia, Niger, **Nigeria**, Rwanda, Senegal, Somalia, Sudan, Swaziland, Tanzania, Togo, Uganda, Zambia, Zimbabwe.

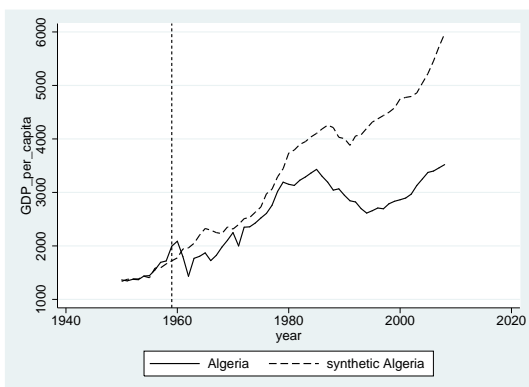
Countries dropped for being resource rich before sample period

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.

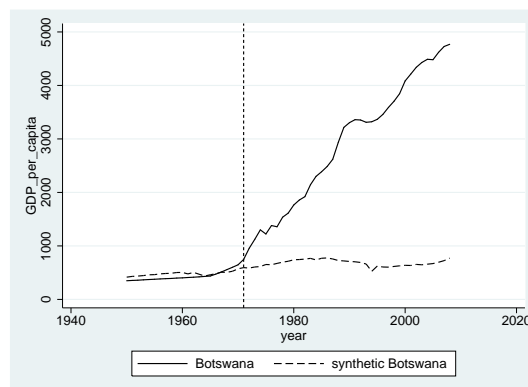
8 Appendix C: Synthetic Control GDP/capita Results

For Online Publication

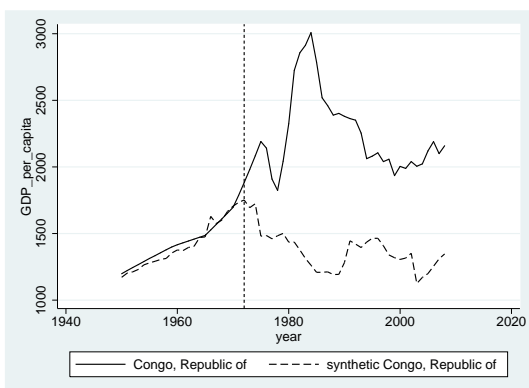
Algeria



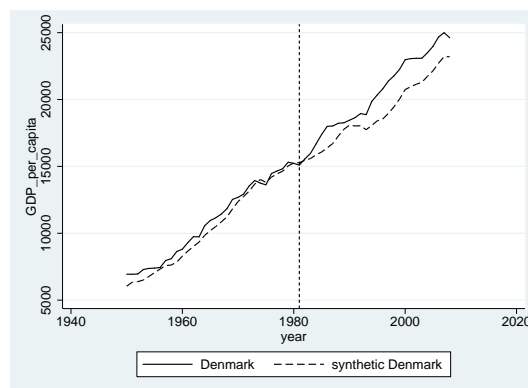
Botswana



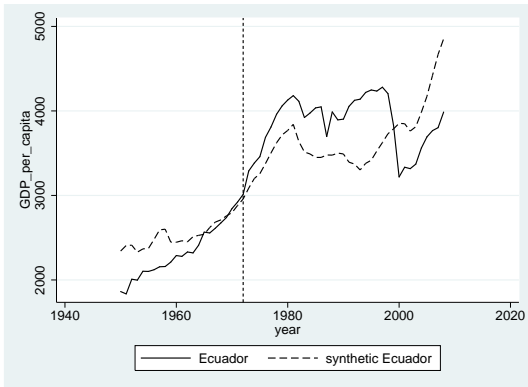
Rep. of Congo



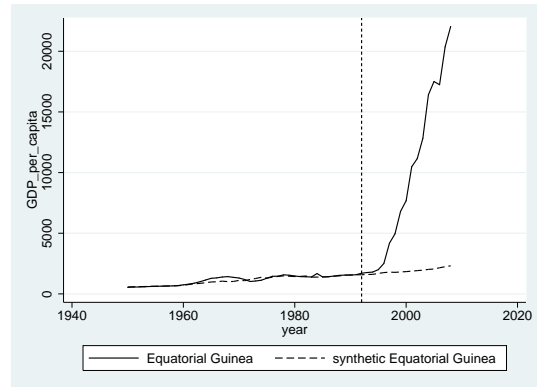
Denmark



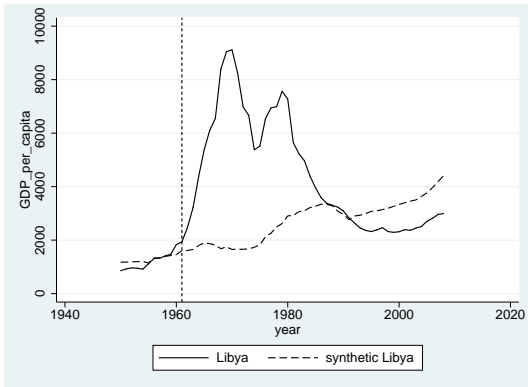
Ecuador



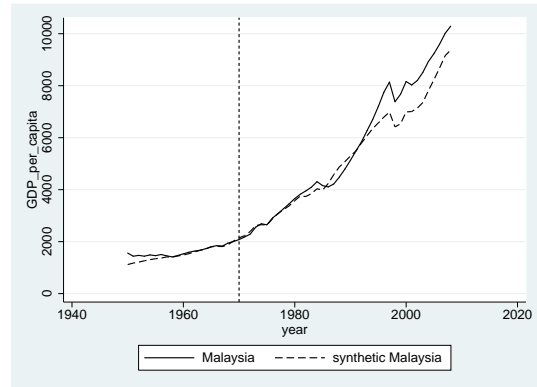
Equatorial Guinea



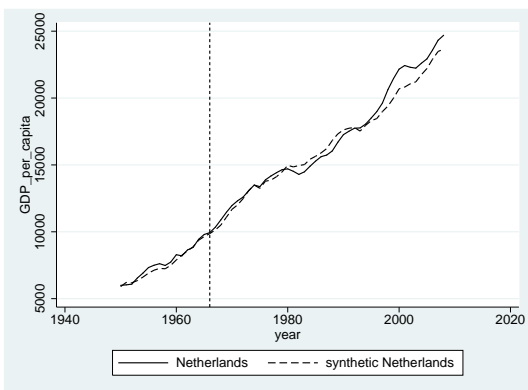
Libya




Malaysia



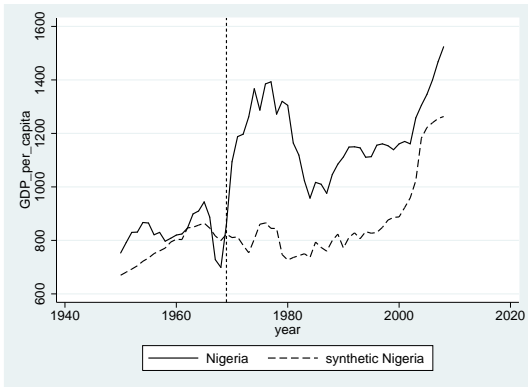
Netherlands



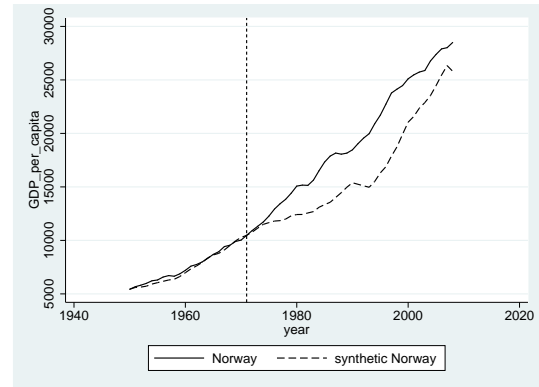
New Zealand



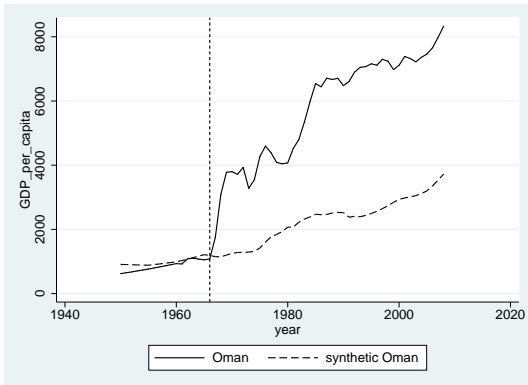
Nigeria



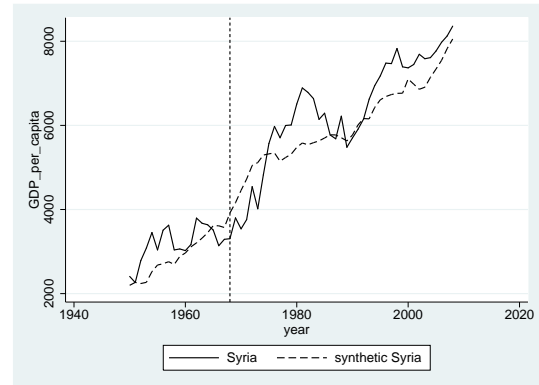
Norway




Oman



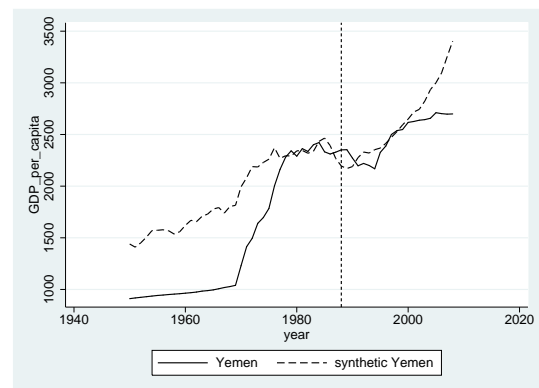
Syria



United Kingdom



Yemen



9 Appendix D: Synthetic Control Weights by Country

For Online Publication

The following tables show the weights given to each control country in the synthetic control analysis for GDP per capita in Figure 7. 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