

DEPARTMENT OF ECONOMICS
OxCarre (Oxford Centre for the Analysis of
Resource Rich Economies)

Manor Road Building, Manor Road, Oxford OX1 3UQ
Tel: +44(0)1865 281281 Fax: +44(0)1865 281163
reception@economics.ox.ac.uk www.economics.ox.ac.uk



OxCarre Research Paper 193

Evidence for a Presource Curse? Oil Discoveries, Elevated
Expectations, and Growth Disappointments

James Cust
World Bank & OxCarre
&

David Mihalyi
Natural Resource Governance

Evidence for a Presource Curse? Oil Discoveries, Elevated Expectations, and Growth Disappointments*

James Cust[†]
David Mihalyi[‡]

July 12, 2017

Abstract

Oil discoveries can constitute a major positive and exogenous shock to economic activity, but the *resource curse* hypothesis would suggest they might also be detrimental to growth over the long run. This paper utilizes a new methodology for estimating growth underperformance to examine the extent to which discoveries depress the growth path of a country following a discovery and prior to production starting. The study finds causal evidence of a significant negative effect on short-run growth and growth relative to counter-factual forecast growth in countries with weak institutions; creating growth disappointments prior to private and public resource windfalls. This effect is termed the **presource curse**.

For a giant oil or gas discovery, between 1988 to 2010, the study estimates an average *growth disappointment* effect of 0.83 percentage points, measured as the average annual gap between forecast and actual growth over the five years following a discovery. Further, the estimated effect varies by the size of the discovery, increasing to a 1.77 percentage points gap in the case of super giant discoveries. The estimated effect is inversely related to the quality of political institutions, and driven by countries with lower institutional quality at the time of the discovery, consistent with the similar long-run results documented in the resource curse literature. For countries with below-threshold institutional quality, the growth disappointment effect is larger, measured as 1.35 percentage points in annual terms. There is no measured growth disappointment effect for countries with strong institutions. Using the synthetic control method we confirm our findings for a selection of countries above and below the institutional quality threshold.

The findings suggest that studies of the resource curse that focus only on the effects of resource exploitation or examine only long-run growth effects may overlook important short-run growth disappointments following discoveries, and the way countries respond to news shocks.

Keywords: resource curse, economic growth, forecasting, forecast errors, news shocks, institutions.

JEL classification: O40, O43, Q33, Q35

*This work was supported by the Natural Resource Governance Institute and the Office of the Chief Economist, Africa Region at the World Bank. We would like to thank Luis-Diego Barrot, Cesar Calderon, Punam Chuhan-Pole, Rick van der Ploeg and Michael Ross as well as participants at the CSAE Conference 2017, University of Oxford for valuable comments. The views expressed in this paper are entirely those of the authors. They do not necessarily represent the views of the International Bank for Reconstruction and Development/World Bank, the Natural Resource Governance Institute or any affiliated organizations.

[†]World Bank. Office of the Chief Economist Africa Region. Email: jcust@worldbank.org

[‡]Natural Resource Governance Institute. Email: dmihalyi@resourcegovernance.org

1 Introduction

Oil discoveries can constitute a major positive exogenous shock to national wealth and hence to economic activity. However, the *resource curse* hypothesis suggests that such discoveries may ultimately prove detrimental to long run growth if they put countries on a path to resource dependence. There exists no short-run corollary to the resource curse hypothesis, however recent years have seen a spate of countries facing severe economic challenges in the years following major natural resource discoveries, including Ghana, Mozambique and Brazil. In this paper we seek to examine whether these anecdotal cases generalise to a widespread empirical phenomenon. We investigate the short run growth response to major oil and gas discoveries around the globe.

We utilise novel data and a newly developed methodology for estimating growth under-performance, to examine the extent to which discoveries depress the growth path of a country following a discovery - and prior to resource wealth being converted into private and public revenues. By combining information on the timing and location of giant oil and gas discoveries with IMF growth forecasts we can estimate the extent to which discoveries lead to *growth disappointment* effects whereby the actual growth path under-performs relative to the counter-factual expected growth. Using synthetic control method we verify these effects for a selection of countries. We find causal evidence of a significant negative effect on short-run growth relative to counter-factual forecast growth; creating growth disappointments even before production begins. We term this the **presource curse**.

Under the resource curse hypothesis, the curse is normally defined and measured in terms of lower growth in resource abundant economies compared to non-resource rich economies. These may or may not constitute a meaningful control group. An ideal test would be to compare resource-rich countries' growth compared with the counter-factual growth trajectory had they successfully harnessed their resource wealth. In other words, an ideal measure of the resource curse would evaluate what growth *could have been* compared to what it actually was. Unfortunately, we do not observe this counter-factual. Our paper takes advantage of the closest we get to this growth counter-factual over the short-run - expert growth projections from the IMF's World Economic Outlook. While limited to five year forecasts only, these nonetheless provide a setting for us to test the growth performance against some benchmark counter-factual of what the growth effects of a discovery are expected to be.

We find that while growth forecasts across the global sample are consistently higher on average for the period following a major oil or gas discovery, actual growth paths are, on average, lower or unchanged, relative to periods without such discovery. This divergence between the expected effect of a discovery on output and actual economic performance can be thought of as a growth disappointment - in the cases where

projections exceed performance. For an oil or gas discovery exceeding 500 million barrels of oil equivalent, we estimate an average growth disappointment effect of 0.83 percentage points, measured as the average annual gap between forecast and actual growth over the 6 years following a discovery. Our estimated effect varies by the size of discovery, increasing to 1.77 in the case of super-giant discoveries.

Evidence suggests that the growth disappointment effects are mediated by the quality of political institutions. Our estimated effect is inversely related to various measures of political institutions, consistent with the similar long run results documented in the resource curse literature. For countries with below-threshold institutional quality prior to discovery, measured in terms of constraints on executive, the growth disappointment effect is larger - 1.35 percentage points. We find no evidence for such a result in countries with strong institutions.

We identify discoveries as the source of these growth disappointments and document a phenomenon we term the presource curse,¹ a corollary to the long run resource curse. We follow Blanchard and Leigh (2013) in utilising the difference between forecast growth and actual growth performance to evaluate the outcomes of government policy in response to a shock.

Oil discoveries can constitute a major exogenous news shock to the economy. While a discovery typically will not result in oil production for over half a decade, it can be considered a step-change in known national wealth, implying increased economic output in the future and a permanent increase in consumption potential for the country. A discovery can therefore generate expectations of increased economic activity and consumption in the future (Arezki et al., 2016), however any subsequent impact on output is also affected by policy choices made by government, including borrowing in the short run to finance increases in investment and consumption (Eastwood and Venables (1982), Pieschacon (2012) & Wills (2014)). Furthermore, it has been proposed that oil discoveries can trigger a country's accumulation of unsustainable levels of debt (Mansoorian, 1991), while government decisions can be distorted by citizen expectations leading to excessive consumption spending (Collier, 2017).

Under rational expectations, a giant oil discovery should lead to a jump in both forecasted and actual growth as the country moves to a new steady state output equilibrium consistent with the shock to national wealth and the associated increase in investment and economic activity. Meanwhile, economic forecasts generated subsequent to a discovery should adjust to reflect the higher growth path anticipated as the economy adjusts to this shock.

However, if a discovery is associated with a presource curse effect, whereby the actual growth path is

¹The term was originally used by energy reporter Leigh Elston describing the case of Mozambique: <http://interfaxenergy.com/gasdaily/article/20380/mozambique-and-the-presource-curse>

lower than the expected growth path, we would expect to observe a systematically negative growth differential (the amount forecast growth exceeds actual growth). This negative growth differential reflects the growth under-performance relative to some forecasted counter-factual. Therefore, if countries are systematically underachieving on growth compared to what is expected, we would predict the positive news shock to be positively correlated with a growth differential; while being negatively or uncorrelated with actual growth. This is what we find.

Our contribution complements the existing literature on the resource curse. Here many studies have sought to document the relationship between resource abundance and disappointing growth rates, however these are typically associated with, and blamed on, responses to resource extraction, resource exports and the rents government derives from this production. See for example Van der Ploeg (2011) and Ross (2015) who discusses the evidence for various economic and political mechanisms that may underlie this curse including Dutch disease, the impacts of commodity price volatility, rent-seeking, rapacious depletion and conflict.

In our case, we find evidence for growth disappointments setting-in prior to production starting and before governments see windfall revenues. Our findings suggest the challenges posed by resource wealth are present and can have detrimental consequences to the economy long before a single barrel of oil is extracted. Further this effect is associated with the way in which government responds to resource discoveries, mediated by the quality of government policies and institutions.

Our contribution is threefold: we present the first empirical evidence - to the best of our knowledge - of a ‘presource curse’ whereby growth disappointments are observed following resource discoveries rather than subsequent to production and revenue windfalls. Second, we present new evidence on the primary importance of institutions and governance in mediating any such presource curse effects. Third, we build on earlier work such as Blanchard and Leigh (2013) to demonstrate how growth forecast errors can be useful to evaluate the divergence between economic performance and official forecasts, and its implications for policy making. Here our findings suggest that this divergence is associated with weaker governance as well as elevated expectations. Elevated expectations themselves may further exacerbate growth disappointments, for example if they drive overly favourable investor perceptions or sovereign credit scores.

1.1 Context and literature

The relationship between economic growth and natural resource wealth

The relationship between economic growth and resource wealth has been subject to extensive study and debate (for recent surveys see Van der Ploeg (2011) & Ross (2015)). Some argue that on average

resource wealth creates a curse, causing reduced or even negative growth, with notable examples including Nigeria and Sierra Leone. In contrast, there are countries whose overall economic development was spurred or unhindered by natural resources - ranging across countries such as the UK, USA, Norway, Malaysia, Botswana and Chile.

An emerging consensus agrees that any overall resource curse effect is best understood as mediated by the quality of institutions (Mehlum et al. (2006) & Robinson et al. (2006)). Here it is argued that countries with strong political institutions at the time of discovery are better placed to reap the benefits of resource wealth, in contrast, countries with weak institutions at the time of discovery are more susceptible to the various resource curse mechanisms and are likely to see a reduction in growth compared to various counter-factuals.

What these studies share in common is an examination of the relationship between resource wealth's contribution to the economy, typically measured via production value, export dependence or government revenue windfalls, and economic performance. Additionally, these studies typically examine the long run growth effects of resource abundance- examining this relationship across decades.

Our approach differs in three key ways: first, we examine the post-discovery effects of *anticipated* resource wealth on countries- testing for a potential effect prior to production commencing; second, we focus on estimating the short run growth impacts of resources. And third, we examine the role for government policy and institutions in mediating the growth effects of discovery.

The economic effects of a resource discovery

A number of countries have recently faced severe macroeconomic problems only a few years after major oil discoveries. Brazil's economy plunged into a recession, while Ghana's and Mozambique's government turned to the IMF for financial assistance shortly following a series of seemingly transformative oil and gas finds.

Major gas discoveries were made in Mozambique in 2009 and 2010. At the time growth rates were averaging 6.5%. In the years immediately following the discovery the IMF forecast growth would reach an average of 7.8% for the 2012-2016 period. Actual growth by the end of the period had slumped to 3.3%. Similarly, Ghana saw major oil and gas discoveries in 2007 and 2010. The IMF continued to forecast strong economic performance above 6% when in fact it reached historical lows below 4% between 2014-2016. Brazil too saw growth slump from a pre-discovery average of 4%, to a post-discovery recession within 7 years of the first of the major offshore 'pre-salt layer' oil discoveries. Each of these cases raises the question of whether the discoveries themselves, rather than windfall revenues or production triggered actions that ultimately

endangered macroeconomic performance and generated growth disappointments. New research suggests several African countries may have experienced *resource curse* type symptoms even without discoveries (Frynas et al., 2016). Here the authors argue that in the cases of São Tome & Príncipe and Madagascar elevated expectations that discoveries would be made prompted spending pressures and deterioration in governance.²

The theoretical foundations for the economic consequences of a resource discovery were developed in the 1980s, including Eastwood and Venables (1982). The authors show how a significant oil discovery can induce a recession under certain conditions, in contrast to the New Classical macroeconomic view at the time that an oil discovery should create no special macroeconomic problems. More recent papers have also considered the short-run consequences of a resource discovery. Arezki et al. (2016) consider the impacts on macroeconomic variables such as employment, savings, investment and the current account, Harding et al. (2016) examine the impact on relative prices, Toews et al. (2016) find a FDI bonanza following giant discoveries, while Pieschacon (2012) and Wills (2014) consider the ideal government policy responses to these shocks.

A resource discovery can be thought of as a one-time shock to national wealth, constituting a news shock about higher levels of output in the future. The delay between discovery and production determines how far into the future increased revenues may be, however both private and public consumption may rise prior to production, financed via borrowing or by running a current account deficit.

Arezki et al. find that in anticipation of oil production following a discovery, employment falls and the current account deteriorates. Meanwhile in the immediate years following a resource discovery output remains flat, until production begins. In contrast, the IMF consistently projects positive output growth in the period following a resource discovery, as the country converges to a new steady state with its increased natural wealth endowment.

This is supported by the theory: in a simple endowment open economy, news of a future increase in output should produce an immediate rise in consumption and therefore output, and an immediate fall in the savings rate and current account as the country borrows from abroad. Once the new resource begins being extracted, the savings rate and current account should swing from negative to positive as the country pays off its debt and also saves for the future.

Wills identifies potential sources for a short-run growth disappointment, relative to the theoretical boon to output implied by the discovery. First, under nominal rigidities, a resource discovery can induce a

²These cases are not included in our analysis since we rely on data on actual giant discoveries for the timing of our news shock. São Tome & Príncipe and Madagascar are yet to make major discoveries of oil or gas.

recession without appropriate policy responses by government. Countries who operate an exchange rate peg are particularly vulnerable to this risk, constituting 75 percent of resource rich countries in Wills' analysis. Second, even under a flexible price regime output may dip in the pre-production years as households work less and consume more, in anticipation of increased government spending and therefore increase household income in the future. Notably, Wills argues that government policy, and specifically monetary policy responses can mitigate the risk of a recession by allowing the interest rate to track the natural rate of inflation. If the monetary policy rule properly responds, the predicted inflation and recession can be overcome.

Countries may also choose to consume the anticipated flow of revenues while simultaneously postponing extraction, an anticipation of better times effect hypothesized by Van der Ploeg (2011). Under these circumstances economic output gains might be reduced by countries saving less in anticipation of higher prices in the future, which in turn may be delayed or never realized.

Furthermore, it has been suggested that countries that lack strong legal and institutional safeguards can find themselves succumbing to a 'voracity effect', whereby powerful elites compete for fiscal spoils (Tornell and Lane, 1999). Such forces can generate outcomes where the loss in output following a discovery may exceed the size of the discovery itself (Collier, 2017).

These studies variously argue that the economic effects of discoveries may depend on institutional quality and adequate government policy response; where failure can trigger growth disappointments. This is what we find reflected in the data.

Forecasting and Economic Growth

We follow Blanchard and Leigh (2013) in utilizing the divergence between forecast and actual growth rates to evaluate short run policy performance in response to a shock. This approach builds on work by Timmermann and Granger (2004) who argue that growth forecasts internalize relevant available information at the time – therefore forecast errors indicate a divergence of performance from expectation, which in turn reflects an incorrectly calibrated growth multiplier. Blanchard and Leigh apply this approach to the fiscal consolidation episodes in Europe following the financial crisis. Here they find evidence of systematic under-forecasting of growth rates by the IMF - providing evidence, they argue, for a realization by European countries of a true growth multiplier that is larger than that used in the forecasts. Fatás and Summers (2016) have since extended this analysis using the same methodology with extrapolated growth forecasts to examine longer run effects.

We take this same technique and apply it to the impact of a resource discovery. Here we find the opposite- namely evidence that the short run growth effect of a resource discovery is systematically smaller

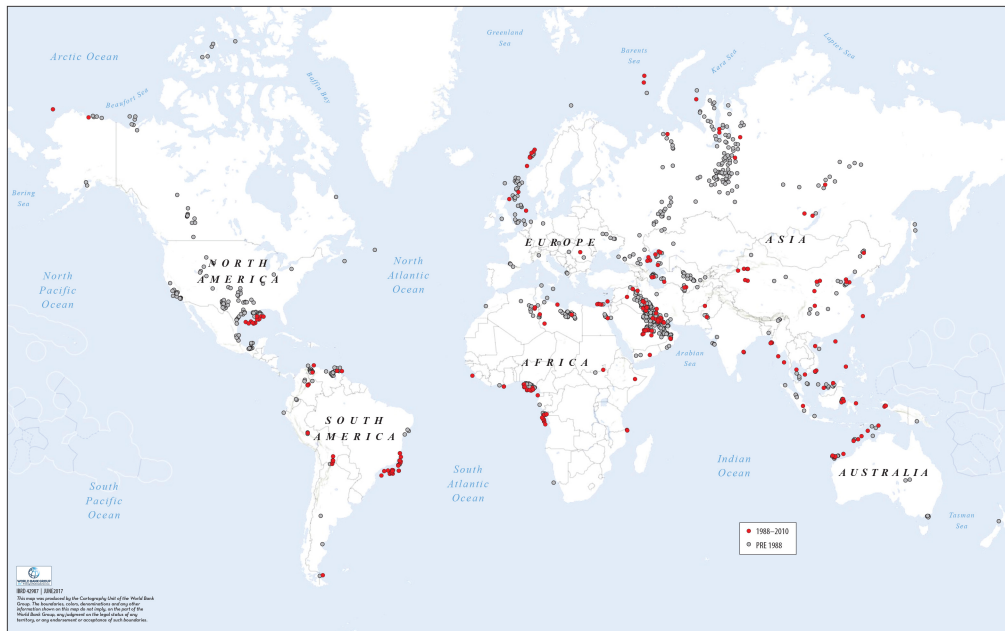


Figure 1: Giant oil discoveries in the full sample and between 1988-2010

than that modelled in forecasts, and in some cases may actually be negative. While experts agree that resource discoveries should be good for growth, in certain policy settings the opposite appears to be true.

2 Data

Our analysis is based on a global cross country panel from 1990 to 2010 and builds on three core datasets. The dataset of giant oil and natural gas field, compiled by Horn (2012), reports on fields of over 500 million barrel equivalent of ultimate recoverable reserves discovered between 1868 and 2010. The dataset provides information on the location and size of the field.³ In terms of measures of oil field size, we construct a dummy variable on reported discoveries as our main variable. We also create an additional dummy variable for supergiant discoveries.⁴ Our third variable of interest on discovery size is the Estimated Ultimate Recovery figure as measured in million of barrel equivalent (*EURMMBOE*). The dataset reports on 996 discoveries since 1868 and 236 discoveries since 1988. The discoveries in our sample are widely distributed geographically, covering 46 countries, and are of significant magnitude for the country's development: the average discovery equals 1.4 percent GDP in NPV terms.

³We refer to oil and gas discoveries as oil hereafter.

⁴We also include the single, even larger mega-giant discovery reported during the period of analysis as part of the supergiant category.

Forecasts are from the IMF’s database of global historical World Economic Outlook (WEO) reports which covers growth, balance of payment and inflation forecasts in every WEO published from 1990 until 2016. Reports are published both in spring and fall, but for simplicity, we use the fall edition each year.⁵ For each year (t), we compile the geometric mean of growth forecasts (averaged growth rate) starting from year $t+1$ until $t+5$, which is period which the forecasts are reported for. Growth is measured in real GDP terms (at constant price).⁶ Actual growth over our sample period is also taken from the Fall 2016 IMF WEO report. We use the figures reported to compile similar 5 year ($t+1-t+5$) average growth performance variables.

We supplement the dataset with additional data on WTI oil price history from the World Bank (to construct discovery value variable) and GDP per capita (IMF) to construct our developing country dummy. Institutional quality measures are taken from a variety of sources including the Augmented Freedom House index of Political Rights and the various components of the Polity IV scores. Our sovereign credit scores are taken from the database of Institutional Investor Ratings, compiled by Reinhart and Rogoff (2011).

Summary statistics are presented below. We have 3,781 country-year observations, comprising 185 countries over the period 1990 to 2012. We observe 285 country-years post discovery, for all countries with a giant oil discovery in the previous 2 years. The remaining 3,496 country-year observations include all country-years for which there is no recorded giant discovery over the previous two years.

Our overall sample records an average forecast growth rate of 4.5 while actual growth is measured at an average of 4.0- implying an average over-forecasting of growth across the country-year sample of around 0.5 percentage points. The absolute forecast error- i.e. the measure of overall mis-forecasting- positive or negative is 2.2 percentage points.

In the case of resource rich countries, (as classified by the IMF in 2012), the average forecast growth rate was of 5.2 while the actual growth was 5.8 implying an average under-forecasting of growth of about 0.50. This is shown in Table 2.

The summary statistics describe a forecast error of notable size. For the full sample it shows is an over-forecasting of growth compared to actual, while for resource rich countries we observe the exact opposite: an under-forecasting of growth of similar magnitude. The absolute forecast error is larger for resource rich countries on average.

IMF WEO Forecasts

⁵We show that our results are robust to the choice of WEO between spring and fall editions.

⁶We also estimate results using the compounded growth in output of the period i.e. the total growth rate.

Variable	Mean	Std. Dev.	N
Actual minus forecast growth annual over five years	-0.472	3.616	3781
Absolute difference annual over five years	2.15	2.945	3781
Averaged actual growth rate annual over five years	3.985	3.913	3781
Averaged forecast growth rate annual over five years	4.457	2.054	3781

Table 1: Summary statistics, full sample

Variable	Mean	Std. Dev.	N
Actual minus forecast growth annual over five years	0.504	6.53	583
Absolute difference annual over five years	3.194	5.717	583
Averaged actual growth rate annual over five years	5.776	6.818	583
Averaged forecast growth rate annual over five years	5.272	2.547	583

Table 2: Summary statistics, resource rich countries

The IMF World Economic Outlook provides forecasts on key macroeconomic aggregates on a 5-year forward looking horizon twice a year for 191 economies. Country forecasts are carried out independently by each IMF country team. They are based on the information and data provided by the national authorities such as Statistics Office, Ministry of Finance, Central Bank as well as common assumptions and guidance provided by a central IMF team for consistency. As the IMF website reports “because forecasts are made by the individual country teams, the methodology can vary from country to country and series to series depending on many factors”.⁷ Despite the differences an important assumption underlying all forecasts is that the present policies of national authorities will be maintained.⁸

2.1 Identification strategy

We test the extent to which growth disappointments arise following a resource discovery. In particular, we investigate the period immediately following discoveries and prior to production taking place, and how performance is associated with government policies and institutions. Our ideal test would be to randomly assign major resource discoveries to different countries and then measure their performance against some calibrated trajectory of growth. We could then observe deviations by countries with weak institutions or poor policy making in the years following a resource discovery and examine how this performance compares to

⁷www.imf.org/external/pubs/ft/weo/faq.htmq1g

⁸www.imf.org/external/pubs/ft/weo/data/assump.htm

their counter-factual trajectory. To mimic this set-up we exploit the giant oilfields data as having plausibly exogenous timing once we account for country and year fixed effects. This dataset has been utilized by various studies in the literature such as Harding et al. (2016), Lei and Michaels (2014) & Arezki et al. (2016).

The empirical test for a resource curse is typically formulated as the difference in long run growth rates between resource-rich and non-resource-rich countries. Here non-resource-rich countries serve as a counter-factual to the growth path of resource-rich countries. However, this overlooks two challenges. The first is that we do not observe the counter-factual growth rate for the resource-rich countries themselves. Second, we do not consider the under-performance of the country in terms of harnessing its resource wealth. Since such wealth should be a blessing, the curse is not merely a divergence from a non-resource-rich average, but the divergence from the true growth potential from resource-wealth (Ross, 2012). Our methodology of comparing growth to counter-factual forecast growth addresses both these challenges, albeit in the short-run only.

By utilizing IMF WEO growth forecasts we are able to construct a growth trajectory counter-factual based on expert analysis and assuming a stable policy path.⁹ This allows us to examine how much weak institutions and policy actions might lead to systematic deviations from the growth path; here any systematic deviation from forecast growth that is associated with a resource discovery can be thought of as evidence of a divergence between the modelled growth and actual growth.

We follow others in the literature (Harding et al., 2016) to argue that the within-country timing of individual discoveries is plausibly exogenous due to the uncertainty surrounding explorations and the limited ability countries and companies have in triggering giant discoveries. While the country’s characteristics may be associated with whether exploration activity takes place or not (see for example the role of institutions versus geology (Cust and Harding, 2014)), and whether a giant discovery is feasible or not, the timing of such a discovery is hard to predict and from a country’s point of view can be considered an unanticipated news shock.

We seek to test two hypotheses. First, whether growth rates are systematically lower in countries following a resource discovery, even before production begins, compared to their predicted counter-factual. And two, whether such under-performance is systematically associated with weak institutions and policy making.

We estimate the effect of giant oil discoveries on the subsequent differential between forecast economic performance and actual economic performance in the period immediately following the discovery. We

⁹www.imf.org/external/pubs/ft/weo/data/assump.htm

then distinguish between countries with varying quality of prior institutions and policy making. Finally, we test against long run growth performance to estimate where short-run growth disappointments are balanced by accelerated growth over the medium or long run.

We discuss a variety of potential channels, including the role for IMF growth forecasts influencing sovereign debt rating and the cost of borrowing. We present several examples and country cases that illustrate cases of growth disappointments following major resource discoveries.

For robustness, we test whether our estimates of an overall growth disappointment can be explained by increased imprecision of forecasting i.e. an increase in the absolute forecasting error. We find no evidence for the average imprecision of forecasting increasing following a resource discovery in the global sample and judge this unlikely to be driving our results of a growth disappointment effect.

Further, our evidence that IMF forecasts overshoot performance more in weakly governed countries than compared to the estimated effect from discoveries in general supports our interpretation that our results are driven by countries growth under-performance rather than analyst exuberance alone. Where we examine the estimated effect of discoveries in strong institution contexts we find no significant effect on any of our main outcomes of interest, namely forecast growth, actual growth or the growth forecast error, further supporting this interpretation.

2.2 Empirical specification

We are interested in both how discoveries impact growth rates relative to their counter-factual, and how these effects might differ across countries with different quality of political institutions. We compare the growth forecast errors in the period immediately following a resource discovery to those prior to the discovery and further away in time. To identify the effect we measure the within-country variation in growth forecast errors, controlling for time invariant country characteristics via country fixed effects, as well as common time-varying effects via year fixed effects. For our main specification we regress our oil discovery on the subsequent growth performance: defined in terms of forecast growth, actual growth and the resulting growth differential.

Under the null hypothesis, the estimated coefficient on the difference between forecast growth and actual growth should equal zero. In other words, the forecast growth effect of a giant oil discovery should not differ from the growth path the country subsequently experiences. Where this estimated coefficient is less than zero, average forecasts exceed the average observed growth experience, implying oil discoveries have proved worse for growth, over the period of 6 years, than we expected.




Figure 2: Timeline and specification

We analyze a cross country panel between 1990 and 2010. In every year (t) and country (n) we compare cumulative average $t+1$ - $t+5$ year growth forecasts as reported in year (t) with actual growth trajectory for the same $t+1$ - $t+5$ period. We label the difference of growth forecast and actual growth “diff”. For example, in 1990 the IMF WEO reported an average expected growth rate of 5.1% for country Z for the 1991 to 1995 period. The actual average growth rate turned out to be 4.7% in the same country and period. Hence the difference (shortfall) in growth rates for country Z in 1990 is - 0.4%. $Forecast(Z, 1990) = 5.1\%$, $Actual(Z, 1990) = 4.7\%$. $Diff(Z, 1990) = - 0.4\%$.

$$GrowthDifferential_{i,t}^{(t+1)-(t+5)} = ForecastGrowth_{i,t}^{(t+1)-(t+5)} - ActualGrowth_{i,t}^{(t+1)-(t+5)} \quad (1)$$

We utilize the forecast growth rate, the actual and the differential across various specifications.

Our analysis looks at whether the lagged variable of oil discoveries affect growth and growth forecasts, and by extension the forecasting gap or error.

Our preferred treatment variable is two year lags of oil discoveries, based on both evidence from other papers and confirmed by our own analysis. This would mean we look at whether 1990 forecasts versus actual growth patterns are affected by oil discoveries in both 1988 and 1989. Discoveries in 1990 are unlikely to be announced in time to be reflected in forecasts published in April or September of that given year. In our robustness check section, we also review a variety of different treatment variables and lags between discovery and the date when forecasts are being prepared. However, using additional lags might contaminate our identification strategy with early signs of the presource curse becoming incorporated into the forecast as more time passes between discovery and forecast. Where we expect a negative growth response to discovery- both forecast and actual- this effect, if present, should bias our estimates towards zero.¹⁰

We use a country-year panel regression with country-level fixed effects, year fixed effects and cluster-robust standard errors clustered by country. We run one regression on effects of lagged discoveries on

¹⁰We run our main specification with a variety of different lags for our discovery shock. Our results do not differ qualitatively from our main findings.

forecasts, one on actual growth performance and one on growth performance against forecast.

$$Y_{i,5} = \beta X_{it} + \alpha_i + year_t + \epsilon_{it} \quad (2)$$

where Y_t is the five-year averaged values of growth, forecast growth and the growth differential. X_t is a dummy taking a value of 1 if there was a major oil discovery in either of the two preceding years. α_i gives country time-invariant characteristics which we control for using country fixed effects. Common time varying effects $year_t$ are captured in our year fixed effects term.

$$X_{it} = Discoverydummy_{i,t-1} + Discoverydummy_{i,t-2} \quad (3)$$

Our alternative difference in difference specification takes the form:

$$Y_{it} = \beta post_{it} + \alpha_i + year_t + \epsilon_{it} \quad (4)$$

where $post_{it}$ takes the value of 1 in all countries with oil discoveries for all years after the discovery.

3 Results

Our main results support the hypothesis that resource discoveries have a causal impact on the growth performance of countries relative to expectations. We find that following giant oil discoveries, the growth rate forecasts as reported in the IMF WEO increase significantly, while changes in actual growth remains, on average, zero. Hence, we find that growth forecasts have a statistically significant upward bias relative to the actual growth path.

3.1 Testing for a growth disappointment effect following a discovery

Our first test examines the direct relationship between a discovery and growth rates- as they deviate from forecast growth. Here we are primarily interested in the effect on three variables; the measured rate of growth, the forecast rate of growth over the same period, and the difference between these two - the growth forecast error which provides our estimate of the *growth disappointment* effect.

Table 3 presents results for our primary specification where our dependent variable represents the changes in our three outcomes of interest. Columns 3,4 and 5 provide additional tests when we include a measure of institutional quality - defined by the Polity IV measure of constraints on the executive (XConst), and measured prior to our first discovery, in 1980.

Our estimates suggest a growth forecast error of -0.83 on average, while the estimated impact on actual growth is -0.95 and significant when we control for initial institutional quality. Our measure of forecast growth alone cannot be distinguished from zero for the global sample. The interpretation on the growth differential coefficient states that the WEO forecast growth for the five year period is on average higher than actual growth by roughly 0.833 percentage points, and this is driven by both negative growth and growth being lower than anticipated. In other words, a negative estimated coefficient on ‘*Diff gr*’ reflects a growth disappointment effect.

The result is stronger in the case of a ‘super-giant’ discovery, where forecast growth rates rise by an average of 1.14 percentage points, adding up to an estimated differential of -1.77 percentage points. This indicates a larger growth disappointment effect.




Figure 3: Actual and forecast growth

3.2 The role of differences in institutional quality

We examine how our main results vary according to differences institutional quality, holding other things equal. Our main results in Table 3 includes a continuous measure of institutions interacted with our discovery

	(1)	(2)	(3)	(4)	(5)	(6)
	Diff gr	Actual gr	Forecast gr	Diff gr	Actual gr	Forecast gr
Lag disc dum,two years	-0.833*** (0.308)	-0.095 (0.325)	0.458 (0.299)	-1.556** (0.645)	-0.952* (0.572)	1.018* (0.603)
Inst*lag_disco				0.138 (0.140)	0.162 (0.123)	-0.159 (0.099)
Fixed effects	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Sample	Full	Full	Full	Full	Full	Full
Obs- Country Year)	3781	5413	4773	2807	3972	3536
R-sq	0.06	0.05	0.03	0.06	0.06	0.04
F-test	13.78	8.95	7.34	13.11	8.82	6.63
Time trend, Standard errors clustered at country level						
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$						
	(1)	(2)	(3)	(4)	(5)	(6)
	Diff gr	Actual gr	Forecast gr	Diff gr	Actual gr	Forecast gr
Lag supergiant dum	-1.766* (1.062)	-0.336 (0.890)	1.138*** (0.351)	1.677 (4.801)	-4.615*** (0.816)	1.800 (1.272)
Inst*lag_Super_disco				-2.690 (3.584)	1.332*** (0.136)	-0.590 (0.714)
Fixed effects	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Sample	Full	Full	Full	Full	Full	Full
Obs- Country Year)	3781	5413	4773	2807	3972	3536
R-sq	0.06	0.05	0.03	0.05	0.06	0.03
F-test	14.18	8.80	7.62	16.10	12.81	7.05
Time trend, Standard errors clustered at country level						
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$						

Table 3: Main results, all sample, five year cumulative

dummy. Table 4 presents a wider range of possible governance scores, including our preferred measure, XConst, which captures constraints on the executive. Additionally we present results using Freedom House scores, measures of democracy, autocracy and the overall Polity IV scores.

We want to distinguish our effects across different levels of institutional quality. However, institutional quality is measured on an ordinal rather than cardinal scale- as such interpretation of the effect on a continuous variable is problematic. We therefore select a threshold to split the sample between strong and weak institutional quality. Our preferred threshold for institutional quality is set at 4 out of a score from 1 to 7 on the XConst Constraints on the Executive measure from Policy IV, measured for countries in 1980 or the nearest available year. 1980 is chosen as preceding the earliest discovery year we include- 1988. Our calibration of the threshold value is discussed in section 5. The timing allows us to be sure we are measuring institutional quality prior to any discovery in our sample, thus reducing potential simultaneity between oil discoveries and our measure of institutional quality; for example, if oil wealth had an effect on subsequent




Figure 4: Change in growth forecast error, by institutions

political institutions.

Table 5 and Table 6 provide separate estimates for the sample above and below an institutional quality threshold, respectively. We present estimates decomposed by our three growth related variables: the growth differential, the actual growth rate and the forecast growth rate.

The results presented in Table 5 show the increased impact of discoveries in countries with observed weaker institutions. For a given oil discovery with estimate an increase in growth differential of almost 1.4 percentage points, while forecast growth rates measured alone are not distinguishable from zero. The effect on actual growth performance is estimated to be negative, at -0.8 percentage points per year.

Table 6 presents the same estimates but now for our sample of countries above the institutional quality threshold. We do not observe any significant differential between forecast and actual growth, and the increase in actual growth and forecast growth are both indistinguishable from zero.

These results suggest our global full sample results are being driven by those countries with weaker institutions seeing a bigger growth disappointment effect, rather than mis-forecasting of the general effect of discoveries.

	(1)	(2)	(3)	(4)	(5)
	XConst	FH	Democ	Autoc	Polity
Lag disc dum,two years	-1.556** (0.645)	-2.317 (1.809)	-1.468*** (0.451)	-0.769 (0.531)	-0.979*** (0.312)
Inst*lag_disco	0.138 (0.140)				
Inst*lag_disco		2.162 (2.331)			
Inst*lag_disco			0.133* (0.077)		
Inst*lag_disco				-0.066 (0.097)	
Inst*lag_disco					0.053 (0.044)
Fixed effects	Country, Year	Country, Year	Country, Year	Country, Year	Country Year
Sample	All sample	All sample	All sample	All sample	All sample
Obs- Country Year)	2807	1890	2807	2807	2807
R-sq	0.06	0.03	0.06	0.06	0.06

Time trend, Standard errors clustered at country level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 4: Main results, diff, weak institutions

	(1)	(2)	(3)
	Diff gr	Actual gr	Forecast gr
lag_disco	-1.354*** (0.419)	-0.831* (0.427)	0.631 (0.445)
Fixed effects	Country, Year	Country, Year	Country, Year
Sample	Weak inst	Weak inst	Weak inst
Obs	1950	2765	2451
R-sq	0.05	0.08	0.04

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 5: Main results, weak institutions

	(1)	(2)	(3)
	Diff gr	Actual gr	Forecast gr
sinst.lag_disco	-0.365 (0.373)	0.296 (0.453)	0.158 (0.141)
Fixed effects	Country Year	Country Year	Country Year
Sample	Strong inst	Strong inst	Strong inst
Obs- Country Year)	2807	3972	3536
R-sq	0.05	0.06	0.03

Standard errors in parentheses

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 6: Main results, strong institutions




Figure 5: Plots above institutions threshold: $xconst1980=4$
From left: Malaysia, Israel, Trinidad & Tobago




Figure 6: Plots below institutions threshold: $xconst1980=4$
From left: Algeria, Colombia, Mozambique

3.3 Synthetic controls

In this section we present synthetic control estimates for a variety of different countries. The use of synthetic controls allows us to demonstrate the variation in the effect of discoveries on different countries with different prior characteristics. Since our identification strategy uses the exogenous timing of a major discovery, we can use this to show effects at the country specific level.

We follow the literature in generating our synthetic controls Abadie et al. (2010). We use five predictor variables to construct each synthetic control based on conditions in the country prior to discovery: ethnic fragmentation, population in year prior to discovery, GDP per capita one, three, and five years before the discovery event.

Figures 5 and 6 gives illustrative examples of the short run GDP per capita effects of a giant oil or gas discovery. Here the examples show the differential effects above and below our chose threshold of institutional quality, constraints on the executive measured in 1980. Here we see countries such as Malaysia, Israel and Trinidad & Tobago exceeding their counter-factual synthetic GDP per capita path. In contrast for countries such as Algeria, Colombia and Mozambique, we see the discovery coinciding with countries undershooting their counter-factual growth.

4 Threats to identification

In this section we examine a range of potential channels and estimate heterogeneous effects of a resource discovery on growth. We seek to determine what may drive our headline result and thus determine the policy implications that may derive from this finding.

There are a range of competing interpretations for our result, including that they may be driven not by growth under-performance, but instead by forecasting error, persistent over-forecasting, or exuberance following discoveries. We address each of these in turn.

Forecasting error A number of articles have scrutinized the accuracy of cross-country macroeconomic forecasts. Given its long history, transparency and widespread country coverage, the majority of this research has examined IMF WEO forecasts, which has a coverage that dwarfs other similar datasets (e.g. OECD, Consensus Forecasts).

The accuracy of forecasts is generally evaluated against the performance of a naïve forecast such as random walk or recursive mean. The ratio of the root-mean squared forecast errors (RMSFE ratio) of actual and naïve forecast provides a measure of accuracy. Global reviews on the accuracy of the WEO forecasts find that it generally perform better than naïve forecasts (Timmermann and Granger (2004) & Genberg and Martinez (2014)). However, country and period selection matters; for a subset of developing countries over certain periods the IMF WEO did not outperform a random walk forecast (Beach et al. (1999) & Arora and Smyth (1990)).

Our results suggest that there is a large and consistent divergence in growth forecasts and measured growth following major resource discoveries. It could be the case that the size and complexity of such an external shock may create forecasting error simply due to the wide range of possible growth responses and difficulties in modelling and predicting how different economies may respond. For example a resource discovery may trigger capital flows which may have differential effects based on the monetary regime or the depth of the financial system.

We can disentangle any potential forecasting error effect from our growth disappointment effect via a test on the **absolute** growth differential as compared to our main variable of interest - the simple growth differential. For details on construction of the variables please refer to Appendix A.

Under conditions of increased but symmetric forecasting error we would expect the estimated coefficient on the discovery effect on absolute growth differentials to be positive. In other words, we would expect the absolute spread between growth forecasts and actual growth to rise following a discovery since

it may simply be harder to forecast growth after such a large external shock. In contrast, we would not expect a result significantly different from zero for our measure of the simple growth differential- here any over-estimate should be cancelled by under-estimates of the true rate of growth across our global sample.

	(1)	(2)	(3)	(4)	(5)
	Abs diff	Abs diff	Abs diff	Abs diff	Abs diff
lag_disco	0.200 (0.197)	0.399 (0.277)	0.221 (0.201)	0.132 (0.234)	0.464* (0.267)
Fixed effects	Country, Year	Country Year	Country, Year	Country Year	Country, Year
Sample	All sample	Weak institutions	Strong institutions	Low and Middle income	High income
Obs- Country Year)	3781	1950	857	2647	1134
R-sq	0.02	0.03	0.05	0.03	0.13

Standard errors clustered at country level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 7: Absolute forecast error, different samples, five year cumulative.

Table 7 presents our results. We find that for the global sample there is no evidence for an increase in the imprecision of the forecasts (absolute error) after discovery. The imprecision error increases significantly in countries with strong institutions, on which we found no evidence of over-forecasting relative to actual growth (simple growth differential - shown in Table 6). In countries with weak institutions, there is no consistent increase in imprecision, but forecasts are more optimistic after discovery relative to actual growth (Table 5).

Persistent over-forecasting As shown in the descriptive statistics (Table 1), WEO forecasts are on average larger than actual growth rates on the global sample. The Independent Evaluation Office’s review of WEO forecasts estimated there is a statistically significant persistent over-forecasting of 0.3 percentage point across all countries and year (Genberg and Martinez, 2014). Previous research has also demonstrated that that persistent over-forecasting is not a unique to the IMF. Its widespread occurrence has been documented across forecasts prepared by national governments (Frankel (2011) and Frankel and Schreger (2012)), central banks (Alessi et al. (2014) and Lansing and Pyle (2015)), multilateral institutions (Genberg and Martinez, 2014) and those of the private sector such as the Consensus Forecast (Loungani, 2001). Studies have also found similar levels of efficiency forecasts and a near-perfect collinearity between these sources (Loungani (2001) and Timmermann and Granger (2004)).

Frankel (2011) postulates that over-optimism in official government forecasts explains why governments may have adopted overly pro-cyclical fiscal policy and run excessive budget deficits in boom times. Cimadomo (2012) finds that pro-cyclicality can also be fueled by biased GDP data available in real time, which is then subject to subsequent revisions. In response to these papers, Avellan and Vuletin (2015) argue that pro-cyclicality cannot be explained by biases in forecasts and data revisions alone, it is being enabled and aggravated by weak institutions.

Our estimation technique is not contaminated by any such global over-forecasting effect since we measure variation of forecasting errors within countries across time. Any ‘average’ IMF bias would be captured in our constant term. If persistent over-forecasting is worse in certain contexts, this is handled via our use of country fixed effects. Here, we are able to be sure our variation comes from the within-country variation from before and after a resource discovery, rather than an average difference between types of country. Finally, if we are concerned about an evolution in forecasting bias across time, our use of year fixed effects should capture any such common time effects over our sample period.

For our results to be driven by such bias the over-forecasting bias would have to time-vary at the country level and be associated with the timing of major resource discoveries. We turn our attention to such a possibility.

Analyst exuberance following a discovery A related but separate concern would arise if mis-forecasting occurs, but is systemically biased in an upward direction following a discovery - an effect we term “analyst exuberance”. Here we might consider that instead of growth differentials arising from below-par growth performance, instead analysts may be systematically and “incorrectly” over-estimating the growth path resulting from a giant oil discovery. There are similar types of issues discussed around the optimism bias of megaprojects (Kahnemann, 2011). The passage from discovery to production is analogous to the period from announcement and planning to eventual execution of megaprojects, where many studies have established the underestimation of cost and complexity that subsequently emerges.

Oil discoveries of this magnitude generate headlines, their news spreads to policy makers and analysts. On the other hand, details of the true commercial potential of these oil finds are generally not well known at the outset, and sometimes kept confidential or exaggerated by unscrupulous investors seeking to generate buzz. Some discoveries are not developed at all or may take several years before any investment decision is made. While the typical oil field takes 6-8 years to develop and produce (according to Smith (2015)), they can also take much longer, depending on geology, market conditions and relations between the investor and the government. All of this can contribute to a challenge for any forecaster to evaluate the true potential and likely economic contribution of a given discovery.

If growth forecasters systematically mis-perceive the size of the discovery, the speed at which it may be developed or the resulting growth dividend the government and country can generate, they may in turn systematically over-forecast the true growth potential of the discovery. This would bias our global estimated results, and suggest growth-under-performance is occurring when in fact it is not relative to the true size of the discovery.

In order to test for an “analyst exuberance” effect we consider a range of heterogeneity estimates

that help us unpack the source of our growth disappointment effect.

The first test is to consider different types of countries receiving the oil discovery shock. Under conditions of “analyst exuberance” we might consider over-forecasting to be by resource discoveries. Here the effect would be common to discoveries across different settings. In other words, if the effect is deriving from a forecaster effect rather than real economic under-performance the estimated forecast error should not differ by countries’ institutional quality or other country features.

In contrast, if the effect stems from economic performance which is mediated by government policy and institutional quality, we may expect our results to vary by institutional quality. Further, we may hypothesize that the growth disappointment effect would be stronger in countries with poorer institutions, consistent with the resource curse literature. In both cases this is what we find, presented above in section 3.2.

Table 6 presents our estimates for a sub-sample of only countries with strong institutions prior to discovery. Consistent with our priors there is no evidence of a growth disappointment effect in these countries. This suggests countries are able to realize the positive growth potential predicted by a discovery shock, and reflected in growth forecasts.

Our second test is to examine the role of a country’s first discovery. Here we might consider this another form of heterogeneity whereby the effect on the growth forecast of a first discovery should be no different to a second, or fifth discovery, in terms of their forecast accuracy, unless it is conditioned on the country’s response to the discovery. In contrast, for a country experiencing its first major oil and gas discovery there may be wide ranging political and economic consequences, including a higher likelihood of policy missteps, and an absence of defensive institutions to protect against voracity effects or citizen pressures for consumption or preferential redistribution.

In this case we find divergent effects depending on the sequential number of discovery. Section 5.1 examines this in more detail. Here, we see for a country’s first major discovery a larger growth disappointment effect, including a negative and significant effect on GDP growth rates relative to the period prior to discovery. Further, our results suggest that the estimated growth disappointment effect is driven by country conditions rather than the forecaster.

Our findings are consistent with the literature examining where and why systematic forecasting errors occur. Studies have found that any growth bias effect across the global sample disappears once all recession years are excluded from the calculations. This implies that any such bias stems from an inability to foresee crises and forecast periods with negative growth (Genberg and Martinez, 2014). While countries

receiving IMF financial support and under an IMF program have higher forecast errors and overly optimistic growth projections (Beach et al. (1999) & Faust (2013)), any such effect in our data would be handled via country fixed effects and otherwise bias our forecasts to zero for those with IMF programs. Similarly, it has been found that forecasts of inflation are systematically optimistic for countries voting in-line with the US in the General Assembly (Dreher et al., 2007). Finally, as previously referenced, Blanchard and Leigh (2013) find that there were larger growth forecast errors for countries that were planning to undertake major fiscal consolidation during the crises in Europe. Here they interpret this forecast error as evidence for an over-performance by European countries relative to the expected size of fiscal multipliers.

Salience of forecast errors: estimating the impact on sovereign credit scoring We find evidence for a divergence between growth forecasts and actual growth performance. One question is what consequences may arise from this divergence, beyond the measured growth disappointment.

One feature of interest is whether other agents are influenced by growth forecasts above and beyond historical growth performance. One test for the salience of the forecasts for public policy is to see how other variables might respond to IMF WEO growth forecasts. Sovereign credit scores given by investors provide a means to test the degree to which private agents internalize growth forecasts.

Table 8 presents estimates for the effect of WEO forecasts and actual growth performance on credit scores. We find that lagged increases in growth forecasts improve subsequent credit scores, while lagged increases in actual growth performance also improve credit scores but with a smaller magnitude. When we control for both forecasts (future) and actual growth performance (past), we find both to be significant, but forecasts carry a larger magnitude. When we include a lagged dummy for resource discoveries we find no change in the direction or significance of the result; indicating that our effect is driven by forecasts and growth performance in general, rather than via discoveries specifically.

Our estimates suggest that investor perceptions of sovereign risk are positively influenced by WEO growth forecasts, even when controlling for the actual growth performance of the country. In other words, when forecasts improve we would expect to see a subsequent improvement in a country’s sovereign investor perception as measured by the Institutional Investor Rating dataset. This holds even when we control for actual growth performance historically. This implies forecasts may factor into investor decisions and other factors such as the cost of borrowing for the country.

If our identification and interpretation holds, it suggests IMF forecasts may be giving an unduly optimistic assessment of the true growth potential of a major oil and gas discovery in countries with weak institutions. In other words, in the presence of a presource curse, growth forecasts may need to be carefully calibrated to take into account the political institutions and governance environment, or else risk distorting

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	CScore	CScore	CScore	CScore	CScore	CScore	CScore	CScore	CScore
lag_forecast	2.377*** (0.484)				1.346*** (0.337)		2.257*** (0.494)		1.252*** (0.338)
L.lag_forecast		2.154*** (0.466)				1.009** (0.324)			
lag_actual			0.931*** (0.124)		0.722*** (0.120)			0.795*** (0.101)	0.611*** (0.110)
L.lag_actual				0.997*** (0.123)		0.861*** (0.130)			
Lag disc dum (disc in past 2 years (excl. t=0))							-2.457 (2.513)	-3.674 (2.085)	-1.894 (2.491)
Lag disco*Lag forecast							0.787 (0.612)		-0.262 (0.660)
Lag discovery*Lag actual								0.953 (0.506)	1.113 (0.674)
Controls	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Standard errors	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust	Robust
Obs- country year	1095	1034	1601	1540	1080	1019	1095	1419	1072
R-sq	0.54	0.52	0.50	0.54	0.61	0.62	0.54	0.55	0.61
F Test	23.71	19.70	31.88	19.67	37.74	24.45	21.14	34.74	40.42

Standard errors in parentheses
* $p < 0.05$, ** $p < 0.01$, *** $p < 0.001$

Table 8: Effects of growth measures on CR

expectations and potentially behavior.

Systemically over-forecasting growth following a resource discovery compared to what a country might reasonably achieve given its institutional context is not just evidence for a presource curse. It can have knock-on effects which may drive or exacerbate any growth disappointments. For example, a mis-perception of the feasible growth path following a resource discovery may also be held by policy-makers, citizens and financial markets, all of whom may be influenced by growth forecasts.

5 Robustness testing and heterogeneity

In this section we provide a variety of tests and alternative specifications to examine the robustness of our main results.

5.1 Discovery sequencing: First, second and multiple prior discoveries

Next we examine, whether there is any difference between the economic effects of a first discovery within a country versus subsequent discoveries. Here, we are interested in whether the impact of a resource discovery is stronger for new producing countries or whether it persists, regardless how mature the petroleum sector

may already be.¹¹

First, and first or second discoveries lead to much larger and much bigger growth differential than subsequent ones. As a result, first, and first or second discoveries result in a very large (over 1.5 and 2.5 percentage points respectively) forecasting error, which drops magnitude for subsequent (multiple) discoveries.

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	First Diff	First Act	First Fore	Second diff	Second Act	Second fore	Multi diff	Multi Act	Multi fore
lag_Fdisco	-1.695*	0.894	0.654						
	(0.957)	(0.621)	(0.638)						
lag_Secdisco				-2.579*	-0.453	0.458			
				(1.310)	(0.929)	(0.781)			
lag_Multidisco							-0.597*	-0.201	0.427
							(0.325)	(0.353)	(0.324)
Obs	3781	5413	4773	3781	5413	4773	3781	5413	4773
R-sq	0.06	0.05	0.03	0.06	0.05	0.03	0.06	0.05	0.03
Standard errors in parentheses									
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$									
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
	First Diff	First Act	First Fore	Second diff	Second Act	Second fore	Multi diff	Multi Actual	Multi fore
inst_lag_Fdisco	-2.586**	0.730	1.004						
	(1.157)	(0.640)	(0.860)						
inst_lag_Secdisco				-2.557**	-1.464*	0.449			
				(1.267)	(0.827)	(0.766)			
inst_lag_Multidisco							-1.001**	-0.887*	0.558
							(0.449)	(0.493)	(0.499)
Obs	2807	3972	3536	2807	3972	3536	2807	3972	3536
R-sq	0.05	0.06	0.03	0.05	0.06	0.03	0.05	0.06	0.04
Standard errors in parentheses									
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$									

Table 9: First, first or second and multiple discoveries, top panel all sample, lower panel weak institutions

5.2 The short run effect on the measured rate of growth

Our main results find evidence for a growth disappointment effect. However, for our global sample there does not appear to be a clear negative effect on actual growth rates from the lagged discovery, compared to years prior to and later in the sample period. This may help explain why post-discovery growth disappointments have gone unnoticed in the literature until now.

We next examine the link between discoveries and measured actual growth rates in different settings. Table 10 presents the results on actual growth rates for different event study samples. Here we take a period symmetric around the year of discovery t .

Our results indicate some evidence for a short-run negative response in growth rates compared

¹¹The ordering of discoveries is determined across the full giant oil discovery dataset, so discoveries are only recorded as first in the analysed period (1988 onwards), where there was no preceding discovery recorded in the full sample (going back to 1868).




Figure 7: Time series plot of pre/post treatment, annual growth rates

to the period immediately preceding discovery. For the five year post-discovery growth is on average 0.14 percentage points lower than the same period pre-discovery. This effect is smaller for the 10 year comparison- only 0.06 percentage points lower on average.

However, when we extend our difference-in-means test to the 15 year period and to the full sample of years prior and post discovery, this effect is no longer measured as significantly different from zero.

	(1) Actual gr	(2) Actual gr	(3) Actual gr	(4) Actual gr	(5) Actual gr	(6) est6
post	-5.166*** (1.876)	-7.508** (2.891)	-0.911 (1.267)	-1.760* (0.949)	-2.941** (1.378)	0.559 (1.055)
Fixed effects	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Period	Five years	Five years	Five years	Five years	Five years	Five years
Sample	All sample	Weak institutions	Strong institutions	All sample	Weak institutions	Strong institutions
SE	Robust	Robust	Robust	Robust	Robust	Robust
Obs- Country Year)	432.00	242.00	179.00	779.00	440.00	313.00
R-sq	0.11	0.19	0.35	0.08	0.12	0.18

inst_pol, Standard errors clustered at country level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 10: Actual annual growth rates, Polity, different event samples

5.3 Testing for persistence in growth disappointments (long run growth)

Our study finds diverging short-term growth effects as a result of discoveries across countries depending on their institutional quality. A question of interest is whether these results persist or even accentuate in the longer term, or whether these effects are only temporary, with long-run convergence to a similar steady state between countries.

Table 10 shows that for the five and ten year pre vs. post discovery periods, average GDP growth is lower following a discovery. However, when we examine longer comparison groups we find no evidence for an average divergence from zero, showing in Table 11. Similarly, when we divide the sample by institutional quality we find no evidence for an average long run effect.

Our findings on the growth effects of resource discoveries suggest that the conventional estimation approaches used in the literature may overlook a short-run presource curse effect. For example, Smith (2015) notes: “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”. Instead Smith focuses on estimating the resource curse effect from the year of exploitation, which can occur 8 or more years beyond the year of discovery. We find, to the contrary of Smith’s hypothesis, a systematic short-run growth under-performance following discoveries across our sample. Indeed, where comparisons are made for growth in the period prior and post the year of production starting, such estimates risk contamination by the presence of short-run presource curse effect.

It is worth noting, when extended to the long-run which captures the effects of oil production as well as discovery, we find no evidence for a long run (unconditional) resource curse effect. As found by Warner (2015), oil production can mechanically raise the aggregate level of growth in an economy, even where the rest of the economy fails to benefit fully from the boom (for example when measured in terms of growth in non-resource GDP). Our findings would be consistent with this observation.

	(1)	(2)	(3)
	AGrowth	AGrowth	AGrowth
post	-0.029 (0.921)	0.004 (0.966)	0.009 (0.870)
Controls	CFE YFE	CFE, YFE	CFE YFE
Standard errors	Robust	Robust	Robust
Event sample	Fifteen years	Twenty years	Full sample
Obs- country year	1097	1382	1946
Actual growth rate			
* p 0.10, ** p 0.05, *** p 0.01			

Table 11: Actual annual growth, long run, different event samples

5.4 Controlling for prior drilling and production

We may be concerned that our result is being driven by the effects of prior oil sector activities. For example if discovery today is correlated with prior exploration drilling, we may worry that our estimated treatment effect is actually capturing some measure of the effect from earlier drilling or oil production, rather than a presource curse effect.

	(1)	(2)	(3)	(4)	(5)	(6)
	Actual gr	Actual gr	Actual gr	Actual gr	Actual gr	Actual gr
post	0.591 (0.680)	0.723 (0.937)	1.604* (0.829)	1.096 (0.809)	0.924 (1.236)	1.416 (0.826)
Fixed effects	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Period	Fifteen years	Fifteen years	Fifteen years	All years	All years	All years
Sample	All sample	Weak institutions	Strong institutions	All sample	Weak institutions	Strong institutions
SE	Robust	Robust	Robust	Robust	Robust	Robust
Obs- Country Year)	1097.00	668.00	280.00	1946.00	1187.00	472.00
R-sq	0.06	0.08	0.17	0.06	0.07	0.17

inst, Standard errors clustered at country level
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 12: Actual 5yr growth rates, different event samples.

We therefore control for prior drilling, in terms of country well count at the time of discovery. We also control for whether the country is already an oil producer at the time of discovery.

Our forecast error results remain unchanged with the inclusion of these controls, suggesting such factors are not driving our main results.

	(1)	(2)	(3)	(4)
	Diff gr	Diff gr	Diff gr	Diff gr
Lag disc dum,two years	-0.833*** (0.308)	-0.827*** (0.311)	-0.824*** (0.308)	-0.819*** (0.310)
Exploration drilling prior sum		-0.000 (0.000)		-0.000 (0.000)
Production dummy			-0.157 (0.470)	-0.156 (0.470)
Fixed effects	Country, Year	Country, Year	Country, Year	Country, Year
Sample	All sample	All sample	All sample	All sample
Obs- Country Year)	3781	3781	3485	3485
R-sq	0.06	0.06	0.06	0.06

Time trend, Standard errors clustered at country level, proddum control
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 13: Controlling for prior exploration well count and whether country is producing

5.5 Testing for a threshold in institutional quality

Our main estimates show that when we control for, and interact our discovery variable with institutional quality, the magnitude and significance of our results increases. Indeed, we find that for countries with lower institutional quality, the growth disappointment effects appear stronger. In contrast, countries with strong institutions prior to discovery appear to suffer no growth disappointment effect following a resource discovery.

	(1)	(2)	(3)	(4)	(5)	(6)
	Actual gr	Actual gr	Actual gr	Actual gr	Actual gr	Actual gr
lag_disco	-0.095 (0.325)	-0.446 (0.678)	-0.480 (0.380)	-0.662 (0.502)	-0.831* (0.427)	-0.670 (0.423)
Fixed effects	Country Year	Country Year	Country Year	Country Year	Country Year	Country Year
Sample	Full sample	XConst 1	XConst 2	XConst 3	XConst 4	XConst 5
Obs	5413	313	1132	2063	2765	3112
R-sq	0.05	0.16	0.06	0.09	0.08	0.07

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 14: Main results, actual growth, weak institutions

	(1)	(2)	(3)	(4)	(5)	(6)
	Diff gr	Diff gr	Diff gr	Diff gr	Diff gr	Diff gr
lag_disco	-0.833*** (0.308)	1.142* (0.556)	-0.527 (0.473)	-1.157*** (0.398)	-1.354*** (0.419)	-1.337*** (0.402)
Fixed effects	Country Year	Country Year	Country Year	Country Year	Country Year	Country Year
Sample	Full sample	XConst 1	XConst 2	XConst 3	XConst 4	XConst 5
Obs	3781	220	795	1438	1950	2206
R-sq	0.06	0.16	0.03	0.05	0.05	0.06

Standard errors in parentheses
* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 15: Main results, forecast diff, weak institutions

Below we test for whether our growth disappointment effect exhibits a threshold- whereby countries below the threshold show qualitatively different results compared to countries above the threshold. Table 14 and Table 15 provide estimates by actual growth and growth forecast error respectively. Each column gives a different cut-off of institutional quality, limiting the sample. In both cases the magnitude of our growth disappointment effect is maximized for countries below the threshold XConst=4. This is the threshold we use in our main split-sample estimates.

5.6 Testing for different size of forecasting windows

The IMF produces 5-year forecasts in each WEO edition. We use this five yearly forecast as our default testing horizon, since it captures the maximum information regarding a country's growth trajectory in any given year. In order to check the robustness of our results we can narrow this forecast window by excluding outer years from our variables and re-run our main specification.

Table 16 below presents our results using forecast windows of 1, 2, 3 and 4 years in addition to our

baseline estimates using the 5 year window. For windows shorter than 4-years we fail to find a significant effect. This suggests that shorter windows we are unable to measure with precision our negative effect on the forecast error. The magnitude and significance is maximized at the longest window available in the WEO forecasts.

	(1) Diff 1y	(2) Diff 2y	(3) Diff 3y	(4) Diff 4y	(5) Diff 5y
Lag disc dum (disc in past 2 years (excl. t=0))	-0.599 (0.509)	-0.527 (0.381)	-0.350 (0.305)	-0.597** (0.295)	-0.833*** (0.308)
Dep var	1 yr window	2 yr window	3 yr window	4 yr window	5 yr window
Fixed effects	CFE YFE	CFE YFE	CFE YFE	CFE YFE	CFE YFE
Sample	Full	Full	Full	Full	Full
Obs- Country Year)	4522	4335	4148	3964	3781
R-sq	0.06	0.07	0.07	0.06	0.06
F-test	15.20	12.41	11.89	12.82	13.78

Time trend, Standard errors clustered at country level

* $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$

Table 16: Forecast error results, all sample, different time windows

6 Conclusions

We find that in the years following a giant oil or gas discovery, growth forecasts increase significantly while actual growth performance does not. Following a major discovery in countries with weaker prior institutional quality, actual growth performance is negative relative to the same years prior to the discovery. These results add up to a systematic growth differential reflecting growth under-performance relative to a counter-factual anticipated growth path. Our results increase in magnitude in the case of supergiant discoveries. Furthermore when we limit our analysis to only oil discoveries (excluding gas) the magnitude also increases.

We interpret our growth disappointment estimates as causal evidence that discoveries may trigger growth under-performance relative to expectations in countries with weaker governance. We find these effects set in following major discoveries but prior to production starting; a phenomenon we term the *presource curse*.

Meanwhile, we find no evidence that discoveries lead to increased forecasting error on average in our global sample, and instead find evidence that our results are driven by growth under-performance in countries with weaker institutional quality. Furthermore, our approach controls for any persistent over-forecasting that may be present across time or space that is unrelated to the timing of resource discoveries.

Our contribution is threefold: we present the first empirical evidence of a ‘presource curse’ whereby growth disappointments are observed following resource discoveries. Second, we present new evidence on the primary importance of institutions and governance in mediating any such presource curse effects. Third, we build on earlier work such as Blanchard and Leigh (2013) to demonstrate how growth forecast errors can be useful to evaluate the divergence between economic performance and official forecasts, and its implications for policy making. Here, our findings suggest that this divergence is associated with weaker governance and political institutions.

We leave it to future research to examine the underlying mechanisms that may generate the growth effects we estimate.

A Definition of variables

$GDP_{i,t}$	Real GDP for country I in period t
$FGDP_{i,t}^{t+5}$	Real GDP forecast for country I as per WEO published in year t for period t+5
$ActualGrowth_{i,t}^{t+1-t+5}$	$= (GDP_{i,t+5}/GDP_{i,t})^{1/5} - 1$ The 5 year average growth rate for for country i as per WEO published in year t
$ForecastGrowth_{i,t}^{t+1-t+5}$	$= (FGDP_{i,t+5}/GDP_{i,t})^{1/5} - 1$ The 5 year average forecast growth rate for country i as per WEO published in year t
$GrowthDiff_{i,t}^{t+1-t+5}$	$= ForecastGrowth_{i,t}^{t+1-t+5} - ActualGrowth_{i,t}^{t+1-t+5}$ The 5 year average growth differential for country i as per WEO published in year t
$AbsGrowthDiff_{i,t}^{t+1-t+5}$	$= ForecastGrowth_{i,t}^{t+1-t+5} - ActualGrowth_{i,t}^{t+1-t+5} $ Absolute value of 5 year average growth differential for country i as per WEO published in year t
$LagDisco_{i,t}$	A dummy variable which equals one if one or more giant oil discovery was made in country i in the two year prior to the WEO being published in year t
$LagSuperDisco_{i,t}$	A dummy variable which equals one if one or more supergiant oil discovery was made in country i in the two year prior to the WEO being published in year t
$LagSignDisco_{i,t}$	A dummy variable which equals one if the estimated value of discoveries made are larger than 1 percent GDP in country i in the two year prior to the WEO being published in year t
$LagFirstDisco_{i,t}$	A dummy variable which equals one if country i made its first giant oil discovery in the two year prior to the WEO being published in year t
$Inst_i$	A dummy variable which divides country into two groups based on whether they scored above or below mean on the Augmented Freedom House index of Political Rights in 1985, the last year we have data for before our sample starts

References

- Abadie, A., A. Diamond, and J. Hainmueller (2010). Synthetic control methods for comparative case studies: Estimating the effect of california’s tobacco control program. *Journal of the American statistical Association* 105(490), 493–505.
- Alessi, L., E. Ghysels, L. Onorante, R. Peach, and S. Potter (2014). Central bank macroeconomic forecasting during the global financial crisis: The european central bank and federal reserve bank of new york experiences. *Journal of Business & Economic Statistics* 32(4), 483–500.
- Arezki, R., V. A. Ramey, and L. Sheng (2016). News shocks in open economies: Evidence from giant oil discoveries. *The Quarterly Journal of Economics*, qjw030.
- Arora, H. K. and D. J. Smyth (1990). Forecasting the developing world: An accuracy analysis of the imf’s forecasts. *International Journal of Forecasting* 6(3), 393–400.
- Avellan, L. and G. Vuletin (2015). Fiscal procyclicality and output forecast errors. *Journal of International Money and Finance* 55, 193 – 204. Macroeconomic and financial challenges facing Latin America and the Caribbean after the crisis.
- Beach, W. W., A. B. Schavey, and I. M. Isidro (1999). *How reliable are IMF economic forecasts?*
- Blanchard, O. J. and D. Leigh (2013). Growth forecast errors and fiscal multipliers. *The American Economic Review* 103(3), 117–120.
- Cimadomo, J. (2012). Fiscal policy in real time. *Scandinavian Journal of Economics* 114(2), 440–465.
- Collier, P. (2017). The institutional and psychological foundations of natural resource policies. *The Journal of Development Studies* 53(2), 217–228.
- Cust, J. and T. Harding (2014). Institutions and the location of oil exploration. *OxCarre Research Paper* 127.
- Dreher, A., S. Marchesi, and J. R. Vreeland (2007). The politics of IMF forecasts.
- Eastwood, R. K. and A. J. Venables (1982). The macroeconomic implications of a resource discovery in an open economy. *The Economic Journal* 92(366), 285–299.
- Fatás, A. and L. H. Summers (2016). The permanent effects of fiscal consolidations. Technical report, National Bureau of Economic Research.
- Faust, J. (2013). A report of the predictive accuracy of the IMF’s WEO forecast. *Version as of February 5*.
- Frankel, J. A. (2011, July). Over-optimism in forecasts by official budget agencies and its implications. Working Paper 17239, National Bureau of Economic Research.
- Frankel, J. A. and J. Schreger (2012, August). Over-optimistic official forecasts in the eurozone and fiscal rules. Working Paper 18283, National Bureau of Economic Research.
- Frynas, J. G., G. Wood, and T. Hinks (2016). The resource curse without natural resources: expectations of resource booms and their impact. *African Affairs*.
- Genberg, H. and A. Martinez (2014). On the accuracy and efficiency of IMF forecasts: A survey and some extensions. *IEO Background Paper No. BP/14/04 (Washington: Independent Evaluation Office of the IMF)*.
- Harding, T., R. R. Stefanski, and G. Toews (2016). Boom goes the price: Giant resource discoveries and real exchange rate appreciation.

- Lansing, K. J. and B. Pyle (2015). Persistent overoptimism about economic growth. *FRBSF Economic Letter*.
- Lei, Y.-H. and G. Michaels (2014). Do giant oilfield discoveries fuel internal armed conflicts? *Journal of Development Economics* 110, 139–157.
- Loungani, P. (2001). How accurate are private sector forecasts? cross-country evidence from consensus forecasts of output growth. *International Journal of Forecasting* 17(3), 419 – 432. Reassessing Modern Business Cycles.
- Mansoorian, A. (1991). Resource discoveries and excessive external borrowing. *The Economic Journal* 101(409), 1497–1509.
- Mehlum, H., K. Moene, and R. Torvik (2006). Institutions and the resource curse. *The economic journal* 116(508), 1–20.
- Pieschacon, A. (2012). The value of fiscal discipline for oil-exporting countries. *Journal of Monetary Economics* 59(3), 250–268.
- Reinhart, C. M. and K. S. Rogoff (2011). From financial crash to debt crisis. *The American Economic Review* 101(5), 1676–1706.
- Robinson, J. A., R. Torvik, and T. Verdier (2006). Political foundations of the resource curse. *Journal of development Economics* 79(2), 447–468.
- Ross, M. (2012). *The oil curse: How petroleum wealth shapes the development of nations*. Princeton University Press.
- Ross, M. L. (2015). What have we learned about the resource curse? *Annual Review of Political Science* 18, 239–259.
- Smith, B. (2015). The resource curse exorcised: Evidence from a panel of countries. *Journal of Development Economics* 116, 57–73.
- Timmermann, A. and C. W. Granger (2004). Efficient market hypothesis and forecasting. *International Journal of forecasting* 20(1), 15–27.
- Toews, G., P.-L. Vezina, et al. (2016). Resource discoveries and fdi bonanzas. Technical report, Oxford Centre for the Analysis of Resource Rich Economies, University of Oxford.
- Tornell, A. and P. R. Lane (1999). The voracity effect. *american economic review*, 22–46.
- Van der Ploeg, F. (2011). Natural resources: Curse or blessing? *Journal of Economic Literature* 49(2), 366–420.
- Warner, A. M. (2015). *Natural Resource Booms in the Modern Era*. International Monetary Fund.
- Wills, S. (2014). Optimal monetary responses to oil discoveries.