

How do oil producers respond to giant oil field discoveries?

by

Jochen GÜNTNER

Working Paper No. 1704

March 2017

How do oil producers respond to giant oil field discoveries?

Jochen H. F. Güntner*

Johannes Kepler University Linz

March 15, 2017

Abstract

This paper studies how petroleum producers respond to a giant oil field discovery. Using a large panel of country-level production data and a difference-in-differences identification approach, I show that domestic production levels respond *before* a newly found oil field has come on line, indicating that producers raise extraction rates from existing reservoirs. Given that domestic petroleum consumption rises by less in response to a discovery, at least part of the increase in production seems to go into (net) oil exports. I find substantial heterogeneity in the impulse responses of oil production and consumption with respect to the location and size of a giant oil field and the country's OPEC membership status.

Keywords: Giant oil field discoveries; Half life of reserves; Oil production; OPEC; Proved reserves

JEL Classification: C32, N50, Q33, Q41

*Jochen Güntner is Assistant Professor at the Department of Economics, Johannes Kepler University Linz. Address: Altenbergerstrasse 69, 4040 Linz-Auhof, Austria. E-mail: jochen.guentner@jku.at, Telephone: +43 732 2468 8360. Fax: +43 732 2468 9679.

1 Introduction

Given that the average delay between a reserves discovery and production is four to six years, Aretzki et al. (2017) use giant oil and gas field discoveries as a directly observable measure of news shocks about future output to explore the effects of such shocks in an open economy. In a large panel of countries, the authors find that — consistent with the predictions of a theoretical two-sector small open economy model with a resource sector — the current account and savings decline for the first five years before rising sharply during the following years. Investment rises shortly after the discovery, whereas GDP does not increase until after five years. Employment falls slightly and remains low for a sustained period of time.

The classical Hotelling model predicts that the optimal level of extraction increases directly after an unexpected resource discovery. In contrast, the analysis in Aretzki et al. (2017) rests on the assumption of a given extraction path from the newly found reservoir, while the possibility of adjusting production from existing reservoirs is discarded on the basis of high adjustment costs in oil and gas production (see, e.g., Hamilton, 2009; Kilian, 2009). Focusing on giant oil field discoveries from the same data set and annual panel data on petroleum production starting in 1965, I find that country-level production picks up soon after the discovery of a large oil field, while the level of statistical significance of the corresponding impulse response function depends on the exact specification of the dynamic panel distributed lag regression model.

Using information about the size, type, and location of oil fields, I show that the subsequent increase in petroleum production arises mainly from so-called “super-” and “mega-giant” fields with an estimated ultimate recovery of three billion barrels of oil equivalents (bboe) as opposed to ordinary “giant” fields, that production increases in response to offshore rather than onshore discoveries, and that non-OPEC producers tend to raise domestic production levels significantly in response to a discovery, whereas OPEC members seem to postpone production from newly discovered as well as existing reservoirs. This is consistent with prior findings in Güntner (2014) regarding the response of OPEC producers to oil demand shocks.

Using a large panel of country-level consumption data starting in 1965, I further investigate whether the observed hike in petroleum production is accompanied by a quantitatively similar increase in petroleum consumption or rather destined for raising (net) oil exports, as implicitly assumed in Aretzki et al. (2017). Based on a sample of 172 countries, I find that country-level consumption also increases in response to a giant oil field discovery — albeit by quantitatively less than production and at a lower level of statistical significance. Interestingly, consumption increases significantly for non-OPEC producers, whereas it decreases significantly or remains unchanged for OPEC member states, depending on the exact specification of the dynamic panel

distributed lag regression model.

My study is closely related to Lepore and Rastad (2011), who use discrete changes in proved reserves and oil production data from the British Petroleum (BP) Statistical Review of World Energy to investigate the response of petroleum producers to oil discoveries during 1980–2009. Relative to their study, the present paper has three important advantages. First, starting in 1965 rather than in 1980, my sample adds 15 years of data, which account for a substantial share of newly discovered oil reserves since the 1960s, as illustrated in Figure 1. Second, oil discoveries in this paper correspond to externally validated and exactly dated findings of giant oil fields rather than to “large changes” in proved oil reserves according to some arbitrary threshold chosen by the authors. This is crucial because proved reserves are defined as reservoirs for which “economic producibility is supported by actual production or conclusive formation tests (drill stem or wire line) or [...] by core analyses and/or electric or other interpretations.”¹ In particular, the amount of proved reserves is subject to the current oil price and available production techniques. As a consequence, the data set in Lepore and Rastad (2011) also contains large *reductions* in proved reserves, which cannot be explained by depletion due to production and are therefore excluded in one of their robustness check. Third, this study draws on the exact size of a newly discovered giant oil field in terms of its estimated ultimate recovery of oil in order to quantify the economic significance of each discovery, whereas Lepore and Rastad (2011) compute a dummy variable that takes on a value of one in the year of a large change in proved oil reserves and zero in all other years, effectively treating the 2009 discoveries of the 9 bboe super-giant Libra field in Brazil and the 0.75 bboe giant Tiber field in the U.S. Mexican Gulf as one and the same event. Finally, the use of a panel distributed lag regression model, as in Aretzki et al. (2017), facilitates the computation of dynamic impulse response functions beyond the analysis of selected regression coefficients, as in Lepore and Rastad (2011).

2 Related Literature

In the classic Hotelling (1931) model, the stock of an exhaustible resource is assumed to be known with certainty. Absent extraction costs, the shadow value of the resource should therefore rise over time at the market rate of interest, while all future prices and the optimal extraction path are determined by the initial price. Yet, the theoretical prediction of an upward-sloping price path seems inconsistent with the empirical observation of falling prices for many mineral commodities over much of the twentieth century. While there are many possible explanations for this discrepancy, one obvious candidate is that the resource base is *not* known with certainty

¹ Compare the “Definitions, Sources and Explanatory Notes” on https://www.eia.gov/dnav/ng/TblDefs/ng_enr_nprod_tbldef2.asp.

from the beginning. Instead, repeated downward revisions of the initial price due to unexpected resource discoveries might offset the predicted rise at the market rate of interest (compare Arrow and Chang, 1982). The present paper is related to three strands of the subsequent literature.

2.1 Exploration and extraction

First, the theoretical and empirical analysis of oil producers' optimal exploration and extraction activity. Peterson (1978) develops a model with extractive firms, which select the time paths of exploration and extraction that maximize their present value. Simulating industry behavior under competition, monopoly, and central management, the authors find that a monopolist will over- conserve the resource and hold excess reserves, whereas competitors over-explore and over-extract. Pindyck (1980) introduces demand and reserve uncertainty in a simple model of an exhaustible resource market by modeling fluctuations in the demand function and the reserve base as continuous-time stochastic processes. With constant extraction costs and risk-neutral firms, neither demand nor reserve uncertainty affects the price and the Hotelling rule continues to apply both in competitive and monopolistic markets. When extraction costs are a function of the level of reserves, however, reserve uncertainty influences the expected price dynamics. Arrow and Chang (1982) assume that the distribution of a natural resource across an unexplored territory follows a Poisson process in space. At any point in time, the socially optimal rates of consumption and exploration are chosen, assuming that reserves are depleted by consumption and increased by resource discoveries from exploration, which reduces the remaining unexplored land. The authors show that, for a large amount of unexplored land, the shadow values of land and reserves move in random cycles with only a slight upward trend, possibly explaining the failure of the so-called Hotelling rule described above.

Farrow (1985) tests whether the theoretical, privately efficient extraction path is consistent with the observed behavior of an individual mining firm. Using proprietary firm data, output price data, and an estimated trans-log cost system to compare changes in the in-situ value with the expected price path, the theoretical model is rejected even when allowing for a time-varying discount rate, an alternative expected price series, or a constraint on the rate of output. Pesaran (1990) explicitly accounts for the intertemporal nature of exploration and production decisions in the oil discovery process by deriving theoretically consistent exploration and production equations for price-taking suppliers. Applying his econometric framework to the U.K. Continental Shelf, in contrast to prior studies, significant price effects on oil production and exploration are found. In a similar vein, Farzin (2001) develops a model of additions to proved reserves that explicitly incorporates the effects of expected resource price, cumulative reserves development, and technological progress. Applying the model to U.S. data for 1950–1995, the author finds a

statistically significant but rather small price elasticity of reserve additions. Assuming steady economic growth and oil consumption, *ceteris paribus*, an annual oil price increase of 1.5–4.5% is necessary to stabilize U.S. oil import dependence in the future. Using data for three separate regions of the Norwegian Continental Shelf for 1965–2004 to estimate error-correction models that capture the longer-term relationships between drilling efforts and geological and technology variables, Mohn and Osmundsen (2008) also find robust evidence of long-term oil price effects on exploration activity in a highly regulated petroleum province.

2.2 Supply from existing reservoirs

Separating production from existing wells in a known oil field from the drilling of new wells and incorporating geological and engineering motives in a model of oil supply from known reserves, Black and LaFrance (1998) test the so-called maximum efficiency recovery (MER) hypothesis that production from established fields is invariant to the price of oil. When applying their econometric model to quarterly data from seven Montana oil fields, the MER hypothesis is strongly rejected. In contrast, Thompson (2001) refers to the “practitioner literature”, which suggests that the owners of an undeveloped resource possess compound options on information acquisition, exploration, and development drilling, whereas the daily production decisions from developed reserves resemble a corner solution. The optimal rate of production is near capacity, primarily because the marginal cost per barrel is constant and far below the market price of oil, while “backwardation” provides the incentive to drill and replace developed reserves.²

Pickering (2008) considers the relationship between extraction rates and remaining reserves. In a simple exhaustible resources model, the slope of a linear extraction rule is determined by the producers’ discount factor, whereas differences in cost and pricing behavior affect the intercept term. The reserves-production relationship is born out by panel data from the world oil industry both across countries and through time. While extraction is characterized by a robust and stable relationship over large ranges of remaining reserves, the estimated slope is significantly lower for OPEC member states. Pickering (2008) argues that this could be explained by differences in risk aversion, discount rates, and measurement error in the reserves data.

Spiro (2014) shows that, by removing any scarcity considerations of the resource owner, the assumption of a rolling planning horizon can reconcile the puzzling long-run price dynamics of exhaustible resources such as oil, gas, and metals. As a result, extraction can be non-decreasing and resource prices non-increasing for a long period of time. A calibration of the model to the oil market replicates the gradually falling real price of oil after WWII and the sharply increasing

² Backwardation means that the discounted futures price of a commodity is below the spot price. While this appears to be inconsistent with the Hotelling principle under certainty, Litzenger and Rabinowitz (1995) argue that backwardation reflects the option value of choosing the timing of production from developed reserves.

price after 1998, suggesting that long-run scarcity has recently become more important.

2.3 Effects of giant oil field discoveries

The size and arguably exogenous timing of resource discoveries make them an ideal instrument for identifying the effects of an anticipated resource boom. As a consequence, the present paper is related to a small number of recent studies that investigate the impact of giant oil (and gas) discoveries on selected economic and political variables.

In a descriptive study, Höök et al. (2009) find that the average decline rate of the world's largest oil fields, which depends on new exploration and production technologies, is increasing over time. Given that these fields represent the most important production base and that the decline rate of existing giant oil fields is already high and increasing, the authors argue that the world faces an increasing oil supply challenge in the future.

Tsui (2009) exploits exogenous variation in the timing and size of giant oil discoveries to identify the impact of changes in oil wealth on democracy and finds that a discovery of 100 billion barrels of crude oil — roughly equal to the initial endowment of Iraq — pushes a country's level of democracy almost twenty percentage points below trend after three decades. While the effect is larger for fields with higher-quality oil and lower exploration and extraction costs, it is less precisely estimated when the discovery's size is measured per capita, suggesting that politicians care about the absolute rather than the per-capita value of a country's oil wealth.

Finally, Harding et al. (2016) estimate the effects of giant oil and gas discoveries on bilateral exchange rates and find that a discovery with a value equal to the country's GDP leads to a real exchange rate appreciation by 14% within ten years following the discovery that is driven almost exclusively by non-tradable goods inflation. These results are qualitatively and quantitatively consistent with the predictions of a calibrated model with forward-looking behavior and Dutch-disease dynamics.

3 The Data

3.1 Giant oil fields of the world

In order to analyze the response of petroleum producers to a credibly exogenous change in crude oil reserves, I use a data set of giant oil and gas field discoveries compiled by Myron K. Horn, former President of the American Association of Petroleum Geologists (AAPG), which contains information about the name, deposit type (i.e. oil or gas), discovery year, hydrocarbon volumes, reservoir depth, and reservoir location of close to 1,000 oil and gas fields discovered worldwide between 1868 and 2010. In what follows, I focus on the 590 giant, super-giant, and mega-giant

oil field discoveries in Horn’s (2014) data, which are particularly suitable for our study. Figure 1 plots the hydrocarbon volumes of giant oil fields cumulated by discovery year for 1900–2010.

The research question of this paper is how oil producers react to an increase in their reserves base and whether the OPEC status of the country or the location (i.e. onshore vs. offshore) and size (i.e. giant vs. super- or mega-giant) of the reservoir matters. Given that the discovery of a giant oil field signals a substantial increase in the country’s oil reserves and thus its wealth and future production capacity and given the distribution of giant oil discoveries across countries as well as over time, each incident can be treated as a country-specific reserves shock.

Another attractive feature is that the *timing* of a giant oil discovery is arguably exogenous and unanticipated due to the uncertainty surrounding oil and gas exploration. While one may argue that discoveries are more likely in geographical regions with larger known and unknown endowments than others or a history of discoveries, the exact timing of a discovery is unlikely to be predictable, especially for giant oil and gas fields. Moreover, past discoveries have two opposing effects on the probability of current and future discoveries. On the one hand, assuming that more easily accessible locations are probed first, cumulative discoveries raise the cost of future successful drilling, as in the theoretical model of Pindyck (1978). On the other hand, knowledge about the territory’s geology and the reservoirs’ location might render future discoveries more likely.³ As a result, past discoveries do not necessarily increase the probability of or reduce the uncertainty about the timing of a giant oil field discovery (compare Aretzki et al., 2017).

Since the timing of each discovery has been independently verified and documented based on multiple sources, which are reported systematically for each incident, Horn’s (2014) data set is also immune to concerns about accidental or deliberate manipulation of the announcement date of a giant oil field discovery, for example by the government.

3.2 Oil production

For country-level petroleum production, I compile data from British Petroleum’s (BP) *Statistical Review of World Energy* and the U.S. Energy Information Administration’s (EIA) *Monthly Energy Review* in thousand barrels per day (tbpd) for a total of 80 countries starting in 1965.⁴ Each oil-producing country is treated as an independent unit of observation, except for former members of the Union of Soviet Socialist Republics (USSR), which are combined in a synthetic “Former Soviet Union” after its dissolution, both in the field discovery and oil production data.

³ In Hamilton and Atkinson (2013), for example, finding a resource today raises the cost of future discoveries, whereas extracting resources yields knowledge that reduces the cost of future discoveries.

⁴ BP production data includes crude oil, tight oil, oil sands, and NGLs (the liquid content of natural gas where this is recovered separately) and excludes liquid fuels from other sources such as biomass and derivatives of coal and natural gas. EIA production data are for crude oil and lease condensate, excluding natural gas plant liquids.

Hence, my analysis is based on an *unbalanced* panel of oil production data for 80 countries, 41 of which experienced at least one giant oil field discovery since 1945.⁵ All members of the “Former Soviet Union” experienced a total of 80 giant oil field discoveries since 1945. The panel includes 14 current OPEC members, all of which experienced at least one giant oil field discovery since 1945, ranging from exactly one for Gabon to 47 for Saudi Arabia.⁶ Table 1 reports the summary statistics for country-level petroleum production data.

3.3 Oil consumption

For country-level petroleum consumption, I compile data from BP’s *Statistical Review of World Energy* and the EIA’s *Monthly Energy Review* in thousand barrels per day (tbpd) for a total of 172 countries starting in 1965, 41 of which experienced at least one giant oil field discovery since 1945.⁷ In line with the treatment of oil production and discovery data, the consumption of former USSR member states is added together in a synthetic “Former Soviet Union”. Given that I can draw on an even broader (unbalanced) panel when analyzing the response of domestic petroleum consumption, at each point in time, there exists a large control group of countries that has not experienced a giant oil field discovery in the recent past or the entire sample period. In fact, the latter applies to most countries in the sample. Table 2 reports the summary statistics for country-level petroleum production data.

4 Econometric Methodology

Inspired by Aretzki et al. (2017), I use a dynamic panel distributed lag regression (DLR) model in order to estimate the response of petroleum-producing countries to a giant oil field discovery:

$$\Delta y_{it} = \mathbf{A}(L) \Delta y_{it} + \mathbf{B}(L) d_{it} + \alpha_i + \mu_t + \epsilon_{it}, \quad (1)$$

where Δy_{it} denotes the change in the level of oil production or consumption in tbpd in country i in year t , α_i controls for country fixed effects such as geographical location or political system, for example, and μ_t controls for time fixed effects such as fluctuations in global crude oil demand, for example. d_{it} denotes the cumulated volume of newly discovered giant oil fields in country i in year t in million barrels of oil equivalents (mmbœ), while ϵ_{it} is a country-year-specific error

⁵ Note that the relevant sample period for oil discovery data starts in 1965 – $q = 1945$ rather than in 1965.

⁶ The list of current OPEC countries includes the five founding members Iran, Iraq, Kuwait, Saudi Arabia, and Venezuela, as well as Qatar (since 1961), Indonesia (1962–2009 and since 2016), Libya (since 1962), United Arab Emirates (since 1967), Algeria (since 1969), Nigeria (since 1971), Ecuador (1973–1992 and since 2007), Gabon (1975–1995 and since 2016), and Angola (since 2007). In what follows, these countries will be treated as OPEC members throughout, as their membership status extends over the majority of the sample period, 1965–2010, with the exception of Angola and Gabon.

⁷ BP consumption data includes inland demand, international aviation and marine bunkers, refinery fuel and loss as well as consumption of biogasoline (such as ethanol), biodiesel, and derivatives of coal and natural gas.

term. $\mathbf{A}(L)$ and $\mathbf{B}(L)$ are p th and q th order lag polynomials with $p \geq 1$ and $q \geq 0$, respectively. In order to fully trace out the impulse response functions over a horizon of 20 years, I set $p = 1$ and $q = 20$ in the benchmark regression analysis.

Regarding the dynamic panel DLR model in (1), a few clarifying comments are in order. First, given that the presence of a unit root in production and consumption cannot be rejected at conventional significance levels for the vast majority of countries, I consider first differences of the dependent variable.⁸ Second, I do *not* logarithmize either the dependent variables or the volume of newly discovered petroleum in giant oil fields, because both series contain “true zeros” rather than missing values. Third, the fact that y_{it} is expressed in tbpd while d_{it} is expressed in mmbae facilitates computing the expected duration until a discrete increase in the reserves base is offset by higher production from newly discovered oil fields or existing reservoirs.⁹ Fourth, Horn’s (2014) data set reports three hydrocarbon volumes for each discovery: the estimated ultimate recovery (EUR), the estimated ultimate recovery of oil (EURO), and reserves (RSVS), all in mmbae, where $\text{EUR} \geq \text{EURO} \geq \text{RSVS}$. Given my interest in the response of country-level *oil* production and consumption, respectively, I use the EURO as a measure of the absolute size of a giant oil field discovery in what follows. Note that the EURO represents an estimate at the time of the discovery rather than a potentially revised estimate at a later date.

The panel structure of the data allows identifying the impulse responses of petroleum production and consumption to a giant oil field discovery while simultaneously controlling for country and year fixed effects. Assuming that their *timing* is exogenous to changes in country-specific production and consumption levels (after controlling for country and year fixed effects), giant oil discoveries correspond to quasi-natural experiments. Country i belongs to the *treatment group*, if an oil field has been discovered on its territory in period t , whereas all countries without a discovery on their territories in period t belong to the *control group*. At each point in time, we can therefore draw on a large group of countries that have not been treated. The dynamic panel specification in (1) also accommodates the fact that a country might have discovered a giant oil field on its territory in period $t - l$, $1 \leq l \leq q$. Accordingly, identification does not rely on a structural vector-autoregressive model with more or less arbitrary identifying restrictions.

Due to the infrequent incident of a giant oil discovery, the high adjustment costs in petroleum

⁸ At the 10% level or better, the null hypothesis of a unit root in country-level production is rejected in favor of a trend-stationary alternative in 14 out of 80 or 17.5% of all cases (see Table 3). The null hypothesis of a unit root in country-level consumption is rejected in 20 out of 172 or 11.6% of all cases (see Table 4). It is important to note that the standard caveats about the power of unit root tests in finite samples apply also in this case.

⁹ In contrast to Aretzki et al. (2017), there is no need to derive the *net present value* of a giant oil field discovery, given that I am only interested in the quantitative response of crude oil production or consumption to an exactly quantifiable discovery. Rather than applying country-specific risk-adjusted discount rates and imposing a certain production profile from the newly found reservoir, the present paper investigates the actual response of petroleum production and consumption *at the country level*.

production (see, e.g., Hamilton, 2009; Kilian, 2009; Güntner, 2014), and the gestation lag before a new reservoir can come on line, it is crucial to draw on a sufficiently long and broad panel. By including the autoregressive term $\mathbf{A}(L)$ in the panel DLR model in (1), I control for potential serial correlation of changes in country-level production or consumption and am able to construct impulse response functions beyond a forecast horizon of q years according to

$$\Phi(L) = \frac{\mathbf{B}(L)}{1 - \mathbf{A}(L)}. \quad (2)$$

4.1 The response of OPEC producers

In order to investigate whether a country's OPEC membership has any influence on its response to a giant oil field discovery, I introduce an additional regressor in (1) that interacts the dummy variable $opec_i \in \{0, 1\}$ with d_{it} , while again controlling for country and year fixed effects:

$$\Delta y_{it} = \mathbf{A}(L) \Delta y_{it} + \mathbf{B}(L) d_{it} + \mathbf{C}(L) d_{it} \cdot opec_i + \alpha_i + \mu_t + \epsilon_{it}, \quad (3)$$

where $\mathbf{C}(L)$ is a q th order lag polynomials with $q \geq 0$ and ϵ_{it} a country-year-specific error term. $opec_i = 1$ for the 14 OPEC member countries in my panel and $opec_i = 0$ else. As argued above, I do *not* account for the fact that Indonesia, Ecuador, and Gabon temporarily suspended their OPEC memberships during the sample period, while Angola joined only in 2007. If anything, this should bias any differences in results between the two groups towards zero.

From equation (3), I compute the impulse response functions of production or consumption in non-OPEC countries as in (2) and the response of OPEC countries according to

$$\Phi_{opec}(L) = \frac{\mathbf{B}(L) + \mathbf{C}(L)}{1 - \mathbf{A}(L)}, \quad (4)$$

where $\mathbf{C}(L)/[1 - \mathbf{A}(L)]$ is the *marginal* impulse response of petroleum production or consumption in OPEC member countries to a giant oil field discovery of one mmbœ in terms of EURO.

4.2 Onshore vs. offshore discoveries

Due to the likely difference in gestation lags, in an extension of (1), we also allow for differential responses to onshore and offshore oil field discoveries (compare Aretzki et al., 2017):

$$\Delta y_{it} = \mathbf{A}(L) \Delta y_{it} + \mathbf{B}(L) d_{it}^{on} + \mathbf{C}(L) d_{it}^{off} + \alpha_i + \mu_t + \epsilon_{it}, \quad (5)$$

where d_{it}^{on} and d_{it}^{off} denotes the EURO of newly discovered giant oil fields onshore and offshore, respectively, in country i and year t . Note that the specification in (5) accommodates multiple discoveries in different locations in a given country and year, weighting them by their respective

EURO in mmboe.¹⁰ In equation (3), instead, each country either is an OPEC member or not.

From equation (5), it is straightforward to compute the impulse responses of country-level production or consumption to an onshore and offshore discovery of one mmboe of EURO as

$$\Phi_{on}(L) = \frac{\mathbf{B}(L)}{1 - \mathbf{A}(L)} \quad \text{and} \quad \Phi_{off}(L) = \frac{\mathbf{C}(L)}{1 - \mathbf{A}(L)}, \quad (6)$$

respectively.

4.3 Does field size matter?

As a final extension, I investigate whether oil producers respond differently to the discovery of a giant as opposed to the discovery of a super- or mega-giant oil field, defined as an oil field with $\text{EURO} \geq 3$ billion boe (i.e. $\text{EURO} \geq 3,000$ mmboe) at the time of discovery. For this purpose, I distinguish oil discoveries according to these categories in the dynamic panel DLR model:

$$\Delta y_{it} = \mathbf{A}(L) \Delta y_{it} + \mathbf{B}(L) d_{it}^g + \mathbf{C}(L) d_{it}^{sg} + \alpha_i + \mu_t + \epsilon_{it}, \quad (7)$$

where d_{it}^g and d_{it}^{sg} denotes the EURO of newly discovered giant and super- or mega-giant fields, respectively, in country i and year t . Similar to (5), the specification in (7) accommodates the simultaneous discovery of one (or multiple) giant and super- or mega-giant oil fields in a given country and year, weighting them by their respective EURO in mmboe.

The corresponding impulse responses of country-level production or consumption to a giant and super- or mega-giant discovery of one mmboe of EURO can then be computed as

$$\Phi_g(L) = \frac{\mathbf{B}(L)}{1 - \mathbf{A}(L)} \quad \text{and} \quad \Phi_{sg}(L) = \frac{\mathbf{C}(L)}{1 - \mathbf{A}(L)}, \quad (8)$$

respectively.

5 Empirical Results

The key question of this paper is how country-level petroleum production responds to a giant oil field discovery. While newly found reservoirs cannot be tapped without a substantial gestation lag, production at the country level might nevertheless adjust in response to a reserves discovery, even in the presence of nontrivial adjustment costs. Aretzki et al. (2017) dismiss this concern by showing that their main results are robust to removing the world's ten largest oil or gas exporters from the sample. Here, I analyze whether and how fast a country's petroleum production adjusts in response to a giant oil field discovery, and whether the associated impulse response function depends on the country's OPEC member status or the field's size and location.

¹⁰ For example, Saudi Arabia discovered four giant onshore fields (Jawb, Lughfah, Samin, and Dhib) and one giant offshore field (Hamur) in 1979 but not a single field in 1980 and 1981.

In a second step, I consider the response of petroleum consumption to investigate whether any adjustment in country-level production is consumed within the country or exported to raise national income from resource extraction, in line with the narrative in Aretzki et al. (2017).

5.1 The response of oil production

Figure 2 plots the *cumulated* impulse response function of Δy_{it} to a giant oil field discovery for all countries in the panel based on the DLR model in (1). Hence, this corresponds to the average response of country-level petroleum production in tbpd to a discovery of one mmboe of EURO.¹¹ The broken and dotted lines indicate one- and two-standard-error bootstrap confidence intervals based on the heteroskedasticity- and autocorrelation- (HAC-)robust covariance matrix in Newey and West (1987).

In response to a giant oil discovery, petroleum production in the treatment group increases within the same year and continues to rise over the subsequent four years relative to production in the control group. The corresponding impulse response function is statistically significant at an approximate 5% level for the first 1–3 years, depending on the lag order p . Given a plausible gestation lag of several years between discovery and the start of production from a new oil field, this finding strongly suggests that producers raise their output from *existing* fields in response to a boost of reserves — exactly the type of behavior dismissed by Aretzki et al. (2017).

Six years after the discovery, petroleum production in the treatment group has increased by 0.04 tbpd per mmboe of EURO relative to the control group, corresponding to an increase of 94 and 32 tbpd for a giant oil field of average and median EURO, respectively. While seemingly small, the former rivals the production level of Italy in the 1990s. The difference becomes even more pronounced about seven years after the discovery, peaking at an additional 0.08 tbpd per mmboe of EURO after 13 years. At this horizon, the relative increase in petroleum production is highly statistically significant, independent of the lag order p .

To put the response of country-level production into quantitative perspective, we can also compute the hypothetical lifetime until exhaustion of the newly discovered reserves. Assuming an average increase of 0.08 tbpd or $365 \cdot 0.08 = 29.2$ tbpy per mmboe of EURO, the discovered reserves would last for about 34.25 years. Given that I abstract from any type of heterogeneity, such as in field size and location, this figure should be taken with a grain of salt.

On the one hand, my results for the full panel confirm the common wisdom of a substantial gestation lag before a newly discovered reservoir comes on line. On the other hand, I find strong evidence that country-level production increases *before* a discovered oil field has been tapped,

¹¹ Recall that the mean and median EURO of giant oil field discoveries in Horn (2014) is equal to 2,348 mmboe and 800 mmboe, respectively.

indicating that producers raise extraction rates from existing reservoirs in order to exploit any remaining spare capacity. In contrast to a demand-driven change in the real price of crude oil, which is of uncertain persistence, a giant oil discovery signals a substantial increase in reserves and production capacity in the near future. For this reason, high adjustment costs in petroleum production (see, e.g., Hamilton, 2009; Kilian, 2009) might be less of a deterrent over a horizon of several years. Accordingly, the above finding complements rather than competes with the finding of a near-zero short-run price elasticity of supply in Güntner (2014).

5.2 The response of OPEC producers

An obvious candidate for heterogeneity in the production response to a giant oil field discovery is the country’s OPEC membership. Of the 590 oil fields listed in Horn’s (2014) data set, 251 are located on current OPEC territories. In this section, I therefore investigate whether OPEC membership has any influence on the response to a discovery of one mmbobe of EURO.

The upper panels of Figure 3 plot the cumulated impulse responses of OPEC production, while the lower panels plot the cumulated impulse responses of non-OPEC production based on the extended DLR model in (3) for $p = 1$ and $p = 4$, respectively. The broken and dotted lines indicate one- and two-standard-error bootstrap confidence intervals based on the HAC-robust covariance matrix in Newey and West (1987).

Figure 3 reveals both qualitative and quantitative heterogeneity in the production responses to a giant oil field discovery. While, on average over the sample period, non-OPEC production increases significantly already during the first five years, OPEC production barely responds until about seven years after the discovery. Moreover, while production of treated OPEC members peaks at +0.08 tbpd after 13 years, that of non-OPEC producers is up by the same amount within four years and peaks at +0.11 tbpd after 18 years relative to the control group. Finally, while the response of non-OPEC production is statistically significant at an approximate 5% level throughout, the response of OPEC production is only significant (at the 5% level) between nine and 15 years after the discovery. Note that these findings are robust to the lag order p .

There are several candidate interpretations for the different response of OPEC producers to giant oil field discoveries. Adherents of the popular theory of OPEC as an effective cartel might argue that its members refrain from producing out of newly discovered reserves or — using the language of Peterson (1978) — “over- conserve the resource” in order to stabilize the price of oil. However, this seems inconsistent with the observation that the impulse response function of OPEC production is (at least partially) statistically significant after about 10 years. Similarly, members of OPEC might postpone production to conserve newly discovered oil reserves, which represent the main or only source of government revenue in these countries. Instead, the delayed

response might also be due to the fact that OPEC producers operated at or close to full capacity during large parts of the sample period. If this was the case, then OPEC production could not respond until a newly discovered oil reservoir has come on line. Very likely, the heterogeneity in Figure 3 arises from a combination of the above reasons.

5.3 Onshore vs. offshore discoveries

Aretzki et al. (2017) report qualitative differences in the responses of consumption, saving/GDP, investment/GDP, and employment to giant onshore and offshore discoveries. In this section, I therefore investigate whether the documented heterogeneity with respect to field location carries over to the impulse response functions of country-level petroleum production.

The upper panels of Figure 4 plot the cumulated impulse responses to a giant onshore oil field discovery, while the lower panels plot the cumulated impulse responses to a giant offshore oil field discovery based on the extended DLR model in (5) for $p = 1$ and $p = 4$, respectively. The broken and dotted lines indicate HAC-robust one- and two-standard-error confidence intervals.

The upper panels suggest that, in response to an *onshore* discovery, country-level production increases on impact, while the impulse response function is not statistically different from zero at an approximate 5% level for the rest of the 20-year horizon. Relative to the control group, production of the treated countries increases by less than six tbpd per mmboe of EURO.

Instead, the lower panels indicate that country-level production does not respond on impact but becomes increasingly positive and statistically significant from about four years after a giant *offshore* discovery. It is important to note that the impulse response function peaks at +0.17 tbpd per mmboe of EURO relative to the control group after 16 years, i.e. three times the amount for an onshore discovery. The differences are statistically significant at an approximate 32% level and robust to the lag order p .

Of the 590 giant oil field discoveries in Horn's (2014) data, 394 (66.8%) are located onshore. The remaining 196 offshore discoveries (33.2%) account for the significant increase in petroleum production relative to the control group. One might thus argue that field location is correlated with ownership, explaining the similarity between response patterns for OPEC and non-OPEC countries in Figure 3 with that for onshore and offshore discoveries, respectively, in Figure 4. Indeed, of the 251 giant oil fields discovered by OPEC members, 70 (27.9%) are located offshore, while 126 (37.2%) of the 339 oil fields discovered on non-OPEC territories are located offshore. Hence, the larger propensity to produce out of offshore reservoirs seems to be linked to the fact that a larger fraction of giant oil fields discovered offshore are located on non-OPEC territories and thus more likely to be used either directly or by raising production from existing fields.

5.4 Does field size matter?

While, in the previous analysis, each discovery is quantified by its EURO, Horn’s (2014) data set also classifies the respective oil field as a “giant” (89.7%), “super-giant” (9.8%), or “mega-giant” (0.5%). In what follows, I therefore investigate whether oil production responds differently to the discovery of a giant as opposed to the discovery of a super- or mega-giant, roughly defined as an oil field with an EURO ≥ 3 billion boe (i.e. EURO $\geq 3,000$ mmboe) at the time of discovery.

The upper panels of Figure 5 plot the cumulated impulse responses to a giant oil discovery, while the lower panels plot the cumulated impulse responses to a super- or mega-giant oil field discovery based on the extended DLR model in (7) for $p = 1$ and $p = 4$, respectively, where broken and dotted lines indicate HAC-robust one- and two-standard-error confidence intervals.

Surprisingly, the discovery of an “ordinary” giant oil field does not trigger a statistically significant increase in petroleum production relative to the control group, except for a partially significant increase between one and three years after the discovery. This finding is robust to the lag order of the DLR model.

Considering now the lower panels, the discovery of an exceptionally large “super-” or “mega-giant” oil field seems to raise petroleum production in the treatment group by up to 0.09 tbpd per mmboe of EURO after 13 years. This increase is robust to the DLR lag order and significant at an approximate 5% level between six and 16 respectively nine and 13 years after the discovery.

Our results suggest that the discovery of an ordinary giant oil field induces producers to extract more from existing reservoirs shortly after the event, whereas the discovery of an exceptionally large super- or mega-giant oil field, such as the Saudi Arabian Ghawar, triggers a delayed and persistent increase in production levels. It is important to note that a field’s EURO is *as of the time of its discovery*. Hence, the above findings seem to reflect effective differences in the countries’ production decision conditional on the size of a newly discovered giant oil field. Apparently, field size matters.

5.5 The response of oil consumption

In this section, I investigate whether the increase in petroleum production in Figure 2 is accompanied by an equally sized increase in domestic consumption or rather destined for raising (net) exports. Ideally, I would like to use a direct measure of net exports as the dependent variable. However, missing data on crude oil inventories above the ground for most countries implies that

$$\{net\ exports\}_t = \{production\}_t - \{consumption\}_t - \Delta \{inventories\}_t$$

contains an unknown variable on the right-hand side.

For this reason, I substitute petroleum consumption for petroleum production in the DLR models in (1) and (3). The availability of annual consumption data starting in 1965 for a total of 172 countries implies that we can draw on an even larger control group of countries without a giant oil field discovery on their territory at each point in time.

Figure 6 plots the cumulated impulse response function of Δy_{it} to a giant oil field discovery for all countries in the panel based on the DLR model in (1), corresponding to the average response of country-level petroleum *consumption* in tbpd to a discovery of one mmbob of EURO. As before, the broken and dotted lines indicate one- and two-standard-error HAC-robust bootstrap confidence intervals. In response to a giant oil discovery, petroleum consumption in the treatment group hardly responds on impact. From period one onwards, the level of consumption starts to rise relative to the control group, peaking at a cumulated +0.025 and +0.045 tbpd after ten to twelve years for $p = 1$ and $p = 4$, respectively. While the impulse response function is partially significant at an approximate 32% level for $p = 1$, it is statistically significant at the 5% level for $p = 4$ around five years after the discovery.

A comparison of Figures 2 and 6 reveals that domestic petroleum consumption increases by less than domestic production. Although the differences in the corresponding impulse response functions are not statistically significant, my findings suggest that at least part of the increase in production raises either the country's net oil exports or crude oil inventories above the ground. Given that the difference between petroleum production and consumption persist over a horizon of 20 years, a giant oil field discovery eventually translates into higher net exports.¹²

Now consider Figure 7, which plots the impulse response functions of cumulated changes in OPEC and non-OPEC consumption after a giant oil field discovery based on the DLR model in (3), where broken and dotted lines indicate one- and two-standard-error HAC-robust bootstrap confidence intervals. While the upper-right panel suggests that, for $p = 1$, an initial increase of petroleum consumption in OPEC member states is followed by a persistent decrease over the remaining forecast horizon that is statistically significant at the 5% level after twelve years, the response for $p = 4$ in the upper-left panel seems to be virtually flat. In contrast, the lower panels of Figure 7 show that petroleum consumption in non-OPEC countries increases in response to a giant oil field discovery, peaking at +0.07 tbpd after twelve years. The latter impulse response functions are statistically significant at an approximate 5% level between three and 13 years after the discovery and robust to the lag order p .

It is important to note that the difference between the impulse response functions of OPEC

¹² Note that the cumulated impulse response functions in Figures 2 and 6 correspond to flow variables, whereas crude oil inventories are a stock variable. A persistent difference between petroleum production and consumption must therefore translate into a continuous increase in crude oil inventories in response to a giant oil field discovery.

and non-OPEC consumption are statistically significant at the 5% (32%) level between ten and 13 years after the discovery for $p = 1$ ($p = 4$). Accordingly, the increase of petroleum production in the treated group relative to the control group in Figure 6 is driven by non-OPEC countries, whereas the level of consumption in OPEC member states remains constant or even decreases.

6 Conclusion

In this paper, I investigate how domestic petroleum production and consumption respond to a sizeable increase in a country’s resource base. The use of a comprehensive data set of giant oil field discoveries with arguable exogenous timing facilitates analyzing the endogenous behavior of oil producers in a quasi-natural experiment. By using an unbalanced panel of oil production and consumption data starting in 1965 for 80 and 172 countries, respectively, many of which have never experienced a giant oil field discovery, I can control for country and year fixed effects. Hence, each impulse response function corresponds to the difference between a country treated with a giant oil discovery in the recent past and all countries in the control group.

I find that country-level production starts to increase before the newly discovered field has come on line, indicating that the reserves discovery spurs production from existing reservoirs in the presence of spare capacity. When distinguishing between OPEC and non-OPEC producers, I find that the response of the former is delayed relative to the response of the latter, while both groups raise their production levels over the subsequent twenty years following a giant oil field discovery. This increase in production seems to arise mainly from offshore rather than onshore discoveries, which are relatively more concentrated on non-OPEC territories, and from “super-” and “mega-giant” rather than ordinary “giant” fields.

Finally, I show that country-level petroleum consumption also increases in response to a giant oil field discovery — albeit by quantitatively less and statistically less significantly. Depending on the specification of the panel distributed lag regression model, OPEC consumption decreases whereas non-OPEC consumption increases significantly in response to a giant oil field discovery, suggesting that the increase in reserves might be devoted to different aims in the two groups.

These findings shed light on the actual response of country-level petroleum production to a giant oil field discovery rather than imposing a certain production profile and dismissing the possibility of adjusting extraction rates from existing reservoirs, as in Aretzki et al. (2017). By comparing the impulse response functions of country-level production and consumption, I can comment on the likely use of increased domestic petroleum supply. However, a more rigorous analysis of the effects of a giant oil discovery on the current account requires panel data on (net) oil exports or inventories. Due to the current lack of such data, this is left for future research.

Acknowledgements

I would like to thank Christiane Baumeister and Yuriy Gorodnichenko for helpful comments. The research on this topic was conducted while the author visited the University of California at Berkeley and Ludwig-Maximilians-Universität Munich as a Postdoctoral Researchers International Mobility Experience (P.R.I.M.E.) Fellow of the German Academic Exchange Service (DAAD). Financial support by the DAAD and the European Union is gratefully acknowledged.

References

- Aretzki, R., V. A. Ramey, and L. Sheng (2017). News shocks in open economies: Evidence from giant oil discoveries. *Quarterly Journal of Economics* 132(1), 103–155.
- Arrow, K. J. and S. Chang (1982). Optimal pricing, use, and exploration of uncertain natural resource stocks. *Journal of Environmental Economics and Management* 9, 1–10.
- Black, G. and J. T. LaFrance (1998). Is Hotelling’s rule relevant to domestic oil production? *Journal of Environmental Economics and Management* 36, 149–169.
- Farrow, S. (1985). Testing the efficiency of extraction from a stock resource. *Journal of Political Economy* 93(3), 452–487.
- Farzin, Y. H. (2001). The impact of oil price on additions to US proven reserves. *Resource and Energy Economics* 23, 271–291.
- Güntner, J. H. F. (2014). How do oil producers respond to oil demand shocks? *Energy Economics* 44, 1–13.
- Hamilton, J. D. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings Papers on Economic Activity* 40(1), 215–283.
- Hamilton, K. and G. Atkinson (2013). Resource discoveries, learning, and national income accounting. *World Bank Policy Research Working Paper 6505* (available at SSRN: <https://ssrn.com/abstract=2285342>).
- Harding, T., R. Stefanski, and G. Toews (2016). Boom goes the price: Giant resource discoveries and real exchange rate appreciation. *School of Economics and Finance Discussion Paper 1608* (available at <https://research-repository.st-andrews.ac.uk/handle/10023/8864>).
- Höök, M., R. Hirsch, and K. Aleklett (2009). Giant oil field decline rates and their influence on world oil production. *Energy Policy* 37(6), 2262–2272.

- Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy* 39(2), 137–175.
- Kilian, L. (2009). Not all oil price shocks are alike: Disentangling demand and supply shocks in the crude oil market. *American Economic Review* 99(3), 1053–1069.
- Lepore, J. J. and M. Rastad (2011). Do oil producers extract more as they become richer? A natural experiment using unanticipated reserve discoveries. *Working Paper* (available at SSRN: <http://dx.doi.org/10.2139/ssrn.1927652>).
- Litzenberger, R. H. and N. Rabinowitz (1995). Backwardation in oil futures markets: Theory and empirical evidence. *Journal of Finance* 50(5), 1517–1545.
- Mohn, K. and P. Osmundsen (2008). Exploration economics in a regulated petroleum province: The case of the Norwegian Continental Shelf. *Energy Economics* 30, 303–320.
- Newey, W. K. and K. D. West (1987). A simple, positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica* 55(3), 703–708.
- Pesaran, H. M. (1990). An econometric analysis of exploration and extraction of oil in the U.K. Continental Shelf. *The Economic Journal* 100(401), 367–390.
- Peterson, F. M. (1978). A model of mining and exploring for exhaustible resources. *Journal of Environmental Economics and Management* 5, 236–251.
- Pickering, A. (2008). The oil reserves production relationship. *Energy Economics* 30, 352–370.
- Pindyck, R. S. (1978). The optimal exploration and production of nonrenewable resources. *Journal of Political Economy* 86(5), 841–861.
- Pindyck, R. S. (1980). Uncertainty and exhaustible resource markets. *Journal of Political Economy* 88(6), 1203–1225.
- Spiro, D. (2014). Resource prices and planning horizons. *Journal of Economic Dynamics and Control* 48, 159–175.
- Thompson, A. C. (2001). The Hotelling principle, backwardation of futures prices and the values of developed petroleum reserves – the production constraint hypothesis. *Resource and Energy Economics* 23, 133–156.
- Tsui, K. K. (2009). More oil, less democracy: Evidence from worldwide crude oil discoveries. *The Economic Journal* 121, 89–115.

Tables and Figures

Table 1: Summary statistics of country-level petroleum production data (in tbpd)

Country	Sample mean	Standard deviation	Maximum	Minimum
Albania	22.931	20.928	5.4	75.0
Algeria	1328.348	338.323	577.4	1992.3
Angola	634.793	608.857	10.9	1900.9
Argentina	593.970	187.974	275.7	910.4
Australia	490.396	193.778	7.0	818.8
Austria	26.823	1.785	23.2	30.5
Bahrain	51.148	4.180	43.9	63.9
Bangladesh	3.200	2.448	0.2	6.9
Belgium	10.598	1.850	5.6	13.1
Bolivia	40.322	13.850	23.3	67.4
Brazil	841.890	718.118	96.0	2346.3
Brunei	174.275	37.676	80.0	261.0
Bulgaria	2.536	1.450	0.0	6.0
Cameroon	106.830	38.501	58.0	185.0
Canada	2264.095	798.754	920.0	4292.3
Chad	117.020	42.148	23.6	173.4
Chile	27.446	13.543	10.9	54.0
China	2520.111	1174.292	227.1	4246.0
Colombia	426.212	261.219	129.0	1004.0
Congo (Brazzaville)	144.244	100.699	0.3	302.3
Congo (Kinshasa)	24.760	4.332	19.7	33.0
Ivory Coast	22.954	18.118	0.8	63.3
Cuba	32.907	17.703	6.0	59.2
Denmark	162.931	129.806	2.0	390.0
Ecuador	297.511	169.427	4.0	556.4
Egypt	626.590	267.580	108.6	940.7
Equatorial Guinea	187.004	135.951	0.1	358.0
Finland	7.931	3.160	0.8	15.3
Former Yugoslavia	62.358	18.370	39.0	90.0
France	93.043	16.217	61.4	119.3
Gabon	217.381	82.641	25.3	364.5
Germany	146.846	15.901	108.5	193.2
Ghana	13.918	28.332	0.0	105.9
Greece	13.125	8.341	1.0	31.0
Guatemala	12.158	6.961	3.0	23.7
Hungary	49.484	12.524	25.3	68.1
India	536.274	284.729	62.3	916.1
Indonesia	1243.859	343.275	474.0	1685.0
Iran	3646.285	1146.691	1321.0	6060.3
Iraq	1884.866	810.909	285.4	3488.6
Israel	5.128	7.925	0.1	36.0
Italy	75.137	37.305	23.4	126.8
Japan	94.900	35.467	38.0	142.6
Kuwait	2163.889	736.729	185.0	3339.0

Continued on next page

Table 1 – *Continued from previous page*

Country	Sample mean	Standard deviation	Maximum	Minimum
Libya	1592.878	558.205	478.8	3357.0
Malaysia	460.743	278.360	1.0	776.0
Mexico	2401.182	1187.857	362.0	3830.2
Mongolia	4.247	5.877	0.1	20.8
Myanmar	18.866	6.609	8.4	32.0
Netherlands	86.072	25.184	30.7	132.3
New Zealand	43.805	16.230	9.0	68.2
Nigeria	1816.621	613.883	141.0	2509.1
Norway	1733.464	1173.318	6.0	3418.0
Oman	614.394	263.510	57.0	960.7
Pakistan	55.145	20.658	11.2	98.0
Papua New Guinea	57.841	31.586	0.2	125.9
Peru	115.024	39.376	64.0	196.0
Philippines	13.400	10.832	0.4	38.6
Poland	17.856	12.871	4.4	40.2
Qatar	717.063	498.581	233.0	1997.8
Romania	194.347	79.781	83.1	312.7
Saudi Arabia	7915.486	2730.347	2219.0	11634.5
South Africa	151.076	68.878	4.0	234.0
Spain	32.288	9.900	3.9	46.8
Sudan	213.945	168.631	2.0	483.1
Suriname	8.500	4.484	1.0	15.4
Sweden	4.349	4.016	0.1	12.3
Syria	320.502	199.414	21.0	676.7
Thailand	178.532	149.715	1.3	459.0
Timor-Leste	84.407	14.039	51.9	100.9
Trinidad and Tobago	161.625	28.742	112.0	230.0
Tunisia	87.125	20.590	16.2	118.2
Turkey	54.813	12.236	39.6	85.3
United Arab Emirates	2062.978	882.492	282.0	3711.6
United Kingdom	1543.125	1024.236	1.6	2930.2
United States	9180.894	1387.321	6783.7	11644.1
Venezuela	2802.600	579.976	1744.0	3754.0
Vietnam	237.896	135.523	0.8	423.6
Yemen	289.877	128.164	10.0	457.3
Former Soviet Union	10194.760	2596.579	4858.0	13802.0

Note: The sample period is 1965–2014.

Table 2: Summary statistics of country-level petroleum consumption data (in tbpd)

Country	Sample mean	Standard deviation	Maximum	Minimum
Afghanistan	13.359	12.558	2.0	43.0
Albania	27.197	9.966	11.7	46.0
Algeria	177.595	99.249	26.8	394.9
American Samoa	3.605	0.553	2.3	4.6
Angola	43.884	29.993	19.0	112.0
Antigua and Barbuda	3.466	0.871	2.0	4.9
Argentina	474.268	62.966	388.1	669.5
Aruba	5.383	1.886	0.7	7.7
Australia	715.558	177.971	346.3	1028.6
Austria	228.181	42.708	107.1	291.7
The Bahamas	21.691	4.685	14.6	30.9
Bahrain	26.978	11.409	15.4	50.0
Bangladesh	53.659	28.812	2.2	115.4
Barbados	8.559	1.103	6.7	10.8
Belgium	547.027	105.509	312.1	730.4
Belize	3.312	1.960	1.1	7.3
Benin	11.912	10.278	2.0	34.8
Bermuda	4.093	0.731	3.2	6.3
Bolivia	39.293	15.608	21.1	71.4
Botswana	9.856	5.617	2.0	19.4
Brazil	1551.522	752.074	306.8	3228.8
Brunei	10.965	4.425	4.0	18.2
Bulgaria	145.967	67.005	71.2	271.3
Burkina Faso	6.478	3.322	2.9	13.4
Myanmar	25.891	8.364	15.2	43.0
Cambodia	14.420	11.901	1.9	34.3
Cameroon	24.584	5.992	16.0	38.7
Canada	1845.202	346.809	1108.1	2404.5
Cape Verde	1.391	0.752	0.4	2.6
Cayman Islands	2.273	1.071	0.7	4.5
Central African Republic	1.925	0.383	1.0	2.5
Chad	1.560	0.378	0.9	2.4
Chile	180.112	96.751	68.1	371.6
China	3661.067	3101.135	216.1	11056.5
Colombia	198.858	60.316	83.8	309.6
Congo (Brazzaville)	7.148	3.388	4.0	16.6
Congo (Kinshasa)	17.723	4.819	8.1	25.4
Costa Rica	31.385	13.122	13.4	50.0
Ivory Coast	28.424	3.496	21.7	37.0
Cuba	194.891	22.336	141.4	228.5
Cyprus	40.675	14.321	17.0	60.2
Czech Republic	180.380	38.256	78.1	241.9
Denmark	230.663	59.362	156.9	366.5
Djibouti	9.273	2.434	4.5	11.8
Dominican Republic	82.637	30.801	36.6	122.3
Ecuador	103.611	68.630	13.5	258.8
Egypt	412.113	210.247	91.7	813.2

Continued on next page

Table 2 – *Continued from previous page*

Country	Sample mean	Standard deviation	Maximum	Minimum
El Salvador	29.707	12.868	11.7	45.5
Equatorial Guinea	1.525	1.507	0.4	5.2
Eritrea	5.089	1.544	2.7	7.9
Ethiopia	24.942	12.458	11.6	52.4
Faroe Islands	4.180	0.592	2.9	4.9
Fiji	9.007	3.907	5.3	17.9
Finland	214.111	28.748	110.8	260.9
Former Yugoslavia	264.199	47.217	153.1	330.7
France	1907.703	287.339	1069.8	2507.9
French Guiana	5.013	1.648	2.2	7.0
French Polynesia	5.560	1.395	1.5	7.2
Gabon	14.191	2.540	9.8	19.0
Gambia	1.830	0.714	1.1	3.4
Germany	2688.320	321.967	1714.2	3341.8
Ghana	33.605	18.207	13.7	78.0
Gibraltar	20.732	14.894	2.4	54.0
Greece	282.616	102.241	85.0	442.4
Greenland	3.992	1.313	1.6	7.0
Guadeloupe	11.121	3.509	3.5	17.0
Guam	18.527	7.910	4.6	35.0
Guatemala	47.491	19.825	23.0	76.1
Guinea	7.698	1.219	5.3	9.3
Guinea-Bissau	1.899	0.727	0.6	2.7
Guyana	8.971	2.151	4.4	11.5
Haiti	9.097	4.074	3.7	16.6
Honduras	27.916	14.436	10.8	51.2
Hong Kong	176.808	99.476	40.6	368.6
Hungary	162.932	41.504	73.1	249.3
Iceland	14.838	2.908	9.0	19.0
India	1511.159	1101.585	252.6	3845.9
Indonesia	730.275	489.367	114.8	1641.0
Iran	1055.963	608.453	133.9	2038.4
Iraq	436.025	151.570	205.0	750.0
Ireland	118.846	40.541	46.7	195.0
Israel	184.653	63.878	80.8	288.1
Italy	1750.152	260.595	982.2	2035.8
Jamaica	55.407	14.282	30.8	82.4
Japan	4735.258	926.149	1705.1	5802.0
Jordan	85.347	25.954	37.0	137.2
Kenya	52.771	16.270	34.5	87.2
Kuwait	204.242	141.952	63.5	505.5
Laos	2.357	0.738	1.1	3.2
Lebanon	79.189	31.343	33.0	128.9
Lesotho	1.326	0.488	0.7	3.7
Liberia	5.223	3.428	2.2	13.0
Libya	192.119	59.758	100.0	331.2
Luxembourg	43.041	15.014	21.0	64.8
Macau	8.897	3.409	3.1	15.1

Continued on next page

Table 2 – *Continued from previous page*

Country	Sample mean	Standard deviation	Maximum	Minimum
Madagascar	10.318	3.597	5.7	17.4
Malawi	4.736	1.509	2.0	7.2
Malaysia	339.667	249.401	46.3	815.2
Maldives	2.939	2.641	0.1	7.0
Mali	3.738	0.757	1.6	4.9
Malta	18.474	11.700	6.0	47.1
Martinique	11.447	4.487	3.9	16.9
Mauritania	14.653	7.990	3.4	23.8
Mauritius	15.469	6.762	5.2	25.1
Mexico	1366.706	614.938	296.4	2067.3
Mongolia	14.258	4.050	8.0	24.6
Morocco	154.130	61.179	85.0	293.0
Mozambique	10.831	3.633	6.5	18.1
Namibia	14.914	5.457	6.9	22.0
Nepal	10.684	6.755	2.0	21.9
Netherlands	789.600	142.871	478.9	1065.1
Netherlands Antilles	81.808	25.014	62.7	156.8
New Caledonia	10.718	5.540	6.0	30.2
New Zealand	108.063	30.046	57.2	154.4
Nicaragua	35.219	62.338	12.1	280.0
Niger	4.182	1.361	0.0	5.4
Nigeria	250.319	35.044	170.0	311.6
North Korea	42.606	22.422	13.3	76.5
Norway	195.800	34.574	99.8	243.2
Oman	52.441	36.624	16.0	163.2
Pakistan	230.186	131.776	71.3	457.8
Palestinian Territories	13.838	6.275	3.1	22.1
Panama	76.423	17.106	35.0	124.9
Papua New Guinea	29.379	49.885	12.0	226.0
Paraguay	37.121	70.133	9.0	314.0
Peru	135.174	35.440	72.4	231.1
Philippines	241.916	80.809	84.9	390.0
Poland	354.897	125.858	108.7	576.4
Portugal	215.785	88.513	50.4	342.1
Puerto Rico	167.068	45.324	20.0	222.8
Qatar	68.393	59.936	16.6	200.6
Reunion	216.401	825.432	4.3	3493.0
Romania	250.502	63.165	142.5	387.1
Rwanda	4.445	1.204	1.1	5.6
Saudi Arabia	1269.161	816.450	366.4	3185.5
Senegal	24.472	9.123	6.8	39.1
Seychelles	76.224	293.361	1.1	1240.0
Sierra Leone	10.932	19.083	3.4	73.2
Singapore	486.996	364.162	71.5	1273.5
Slovakia	90.080	25.428	44.9	139.1
Somalia	111.829	350.465	3.7	1300.9
South Africa	358.155	141.330	117.7	607.0
South Korea	1249.441	950.317	25.0	2458.2

Continued on next page

Table 2 – *Continued from previous page*

Country	Sample mean	Standard deviation	Maximum	Minimum
Spain	1055.532	359.950	269.1	1613.1
Sri Lanka	54.578	24.023	17.0	100.2
Sudan	45.488	31.724	4.8	114.8
Suriname	35.753	82.672	6.7	309.1
Swaziland	23.615	67.687	2.0	250.7
Sweden	406.041	80.868	306.3	569.4
Switzerland	254.145	25.700	164.8	298.4
Syria	209.703	77.110	12.0	346.6
Taiwan	576.255	350.459	43.9	1110.3
Tanzania	18.812	7.851	1.3	34.4
Thailand	530.072	397.403	47.6	1274.0
Togo	9.468	9.493	3.5	43.1
Trinidad and Tobago	34.887	12.405	16.8	68.3
Tunisia	77.745	21.534	51.5	146.3
Turkey	451.157	207.491	89.8	724.2
Uganda	48.990	161.886	3.2	694.0
United Arab Emirates	287.971	248.389	0.2	873.0
United Kingdom	1740.495	168.088	1448.7	2226.2
United States	17394.674	2235.350	11522.2	20802.2
Uruguay	35.793	13.165	1.0	60.9
Venezuela	447.367	174.951	181.5	825.0
Vietnam	136.771	112.362	16.3	405.9
Virgin Islands, U.S.	73.612	29.715	9.3	117.0
Wake Island	16.206	29.381	8.6	134.0
Western Sahara	2.485	4.111	1.1	19.3
Yemen	85.929	36.061	14.0	161.6
Zambia	12.723	1.421	10.7	15.8
Zimbabwe	19.094	5.882	12.0	31.0
Former Soviet Union	5633.537	1934.583	3314.0	8455.3

Note: The sample period is 1965–2014.

Table 3: Augmented Dickey-Fuller test for a unit root in country-level petroleum production

Country	Adj. sample size	Lag length	ADF t -statistic	ADF p -value
Albania	30	4	-1.4475	0.825
Algeria	49	0	-1.8454	0.665
Angola	49	0	-1.5145	0.810
Argentina	48	1	-1.1862	0.900
Australia	49	0	-1.4229	0.840
Austria	34	0	-3.5499**	0.050
Bahrain	34	0	-1.7048	0.725
Bangladesh	31	0	-0.5882	0.975
Belgium	34	0	-3.8794**	0.025
Bolivia	34	0	-2.6466	0.265
Brazil	49	0	-0.9320	0.945
Brunei	49	0	-2.1876	0.485
Bulgaria	34	0	-3.4229*	0.065
Cameroon	34	0	-3.7696**	0.030
Canada	48	1	-1.1592	0.905
Chad	8	3	-4.6285**	0.025
Chile	33	1	-1.0070	0.930
China	47	2	-2.7746	0.215
Colombia	48	1	-2.5481	0.305
Congo (Brazzaville)	49	0	-2.7171	0.235
Congo (Kinshasa)	31	3	-5.6874***	0.010
Ivory Coast	34	0	-1.5249	0.800
Cuba	34	0	-0.8422	0.950
Denmark	41	1	-0.3380	0.985
Ecuador	49	0	-3.6346**	0.035
Egypt	48	1	-1.1233	0.915
Equatorial Guinea	22	1	-1.2751	0.865
Finland	33	0	-2.9025	0.175
Former Yugoslavia	34	0	-2.1927	0.480
France	34	0	-2.2750	0.435
Gabon	48	1	-2.2625	0.445
Germany	34	0	-3.5609**	0.050
Ghana	34	0	-0.0969	0.990
Greece	30	4	-0.6151	0.970
Guatemala	34	0	-1.0894	0.915
Hungary	34	0	-2.8605	0.190
India	49	0	-1.3149	0.870
Indonesia	49	0	-1.6924	0.740
Iran	48	1	-2.9393	0.160
Iraq	48	1	-3.1019	0.115
Israel	33	1	-41.0519***	0.010
Italy	49	0	-2.3081	0.420
Japan	34	0	-1.3168	0.865
Kuwait	49	0	-1.3567	0.860
Libya	49	0	-2.7861	0.210
Malaysia	49	0	0.3616	0.990
Mexico	48	1	-0.8416	0.955

Continued on next page

Table 3 – *Continued from previous page*

Country	Adj. sample size	Lag length	ADF <i>t</i> -statistic	ADF <i>p</i> -value
Mongolia	12	4	6.3312	0.990
Myanmar	34	0	−1.4543	0.825
Netherlands	34	0	−2.7990	0.210
New Zealand	34	0	−2.0721	0.540
Nigeria	46	3	−4.6211***	0.010
Norway	42	1	−0.7975	0.960
Oman	46	1	−2.0591	0.555
Pakistan	34	0	−1.5918	0.775
Papua New Guinea	20	3	−1.1534	0.890
Peru	48	1	−2.2276	0.465
Philippines	35	0	−2.6744	0.255
Poland	33	1	−2.7147	0.235
Qatar	49	0	0.4324	0.990
Romania	48	1	−2.8559	0.185
Saudi Arabia	49	0	−2.1277	0.515
South Africa	33	0	−0.7290	0.960
Spain	31	3	−1.8028	0.680
Sudan	21	0	−0.9952	0.925
Suriname	28	1	−3.5153*	0.055
Sweden	30	0	−1.1269	0.905
Syria	45	1	−0.2055	0.990
Thailand	34	0	−1.4850	0.815
Timor-Leste	7	3	−12.0093***	0.010
Trinidad and Tobago	48	1	−2.5758	0.295
Tunisia	48	0	−4.3545***	0.010
Turkey	33	1	−2.3232	0.410
United Arab Emirates	48	1	−2.6982	0.240
United Kingdom	48	1	−1.0194	0.930
United States	48	1	0.2273	0.990
Venezuela	49	0	−1.4252	0.840
Vietnam	28	0	−0.6871	0.965
Yemen	26	2	0.2969	0.990
Former Soviet Union	48	1	−3.1968*	0.100

Note: The sample period is 1965–2014. Under the alternative hypothesis, production follows a trend-stationary process. ***/**/* indicates that the null hypothesis of a unit root is rejected at the 1/5/10% significance level.

Table 4: Augmented Dickey-Fuller test for a unit root in country-level petroleum consumption

Country	Adj. sample size	Lag length	ADF t -statistic	ADF p -value
Afghanistan	32	1	-1.5839	0.775
Albania	33	0	-1.2356	0.885
Algeria	48	1	-0.9870	0.935
American Samoa	30	3	-1.5021	0.805
Angola	33	0	-0.2242	0.990
Antigua and Barbuda	29	4	-2.0040	0.575
Argentina	49	0	-0.3081	0.990
Aruba	26	1	-2.5089	0.320
Australia	49	0	-2.2812	0.435
Austria	49	0	-2.7126	0.235
The Bahamas	32	1	-3.78331**	0.030
Bahrain	33	0	-1.0888	0.915
Bangladesh	42	1	-2.3495	0.400
Barbados	33	0	-3.6359**	0.040
Belgium	49	0	-1.9321	0.620
Belize	33	0	-1.3530	0.855
Benin	33	0	-1.5949	0.770
Bermuda	33	0	-3.3415*	0.075
Bolivia	33	0	-2.0131	0.570
Botswana	33	0	-3.6617**	0.040
Brazil	48	1	-0.9321	0.945
Brunei	33	0	-2.5682	0.295
Bulgaria	48	1	-3.0745	0.125
Burkina Faso	33	0	-3.0152	0.140
Myanmar	32	1	-1.7520	0.705
Cambodia	33	0	-1.5501	0.790
Cameroon	33	0	-0.9867	0.930
Canada	48	1	-2.4726	0.340
Cape Verde	30	3	-2.7203	0.235
Cayman Islands	33	0	-2.5794	0.290
Central African Republic	33	0	-3.4486*	0.060
Chad	33	0	-3.4782*	0.055
Chile	48	1	-2.1828	0.485
China	49	0	1.2650	0.990
Colombia	47	2	-2.8201	0.195
Congo (Brazzaville)	33	0	-1.0838	0.915
Congo (Kinshasa)	33	0	-1.2739	0.875
Costa Rica	33	0	-2.8913	0.175
Ivory Coast	33	0	-3.3237*	0.080
Cuba	33	0	-2.8214	0.200
Cyprus	33	0	-0.2157	0.990
Czech Republic	48	1	-2.6409	0.265
Denmark	48	1	-2.9869	0.145
Djibouti	33	0	-1.5787	0.780
Dominican Republic	33	0	-0.8506	0.950
Ecuador	49	0	-0.5136	0.980
Egypt	49	0	-2.1461	0.505

Continued on next page

Table 4 – *Continued from previous page*

Country	Adj. sample size	Lag length	ADF <i>t</i> -statistic	ADF <i>p</i> -value
El Salvador	33	0	-1.4919	0.810
Equatorial Guinea	33	0	-2.1793	0.485
Eritrea	19	0	-2.9680	0.165
Ethiopia	33	0	-1.2446	0.885
Faroe Islands	33	0	-3.0052	0.145
Fiji	32	1	-2.3698	0.385
Finland	49	0	-3.6271**	0.040
Former Yugoslavia	33	0	-1.6119	0.765
France	48	1	-3.5948**	0.040
French Guiana	32	1	-0.0456	0.990
French Polynesia	29	4	-2.8153	0.205
Gabon	33	0	-1.9578	0.600
Gambia	33	0	-1.6609	0.745
Germany	49	0	-3.6668**	0.035
Ghana	33	0	-1.1373	0.905
Gibraltar	33	0	-3.0048	0.145
Greece	49	0	1.8889	0.990
Greenland	33	0	-2.5237	0.315
Guadeloupe	33	0	-2.1896	0.480
Guam	33	1	-2.8731	0.180
Guatemala	33	0	-1.9867	0.585
Guinea	33	0	-3.9839**	0.020
Guinea-Bissau	30	3	-2.1847	0.480
Guyana	32	1	-2.8762	0.185
Haiti	33	0	-2.7893	0.210
Honduras	33	0	-2.7099	0.240
Hong Kong	49	0	-1.7619	0.705
Hungary	49	0	-2.3560	0.395
Iceland	34	0	-0.8769	0.945
India	49	0	-0.2524	0.990
Indonesia	49	0	-2.4291	0.360
Iran	49	0	-3.1569	0.105
Iraq	33	0	-2.9118	0.170
Ireland	48	1	-1.8595	0.660
Israel	49	0	-1.2078	0.895
Italy	49	0	-2.0495	0.560
Jamaica	32	1	-1.8304	0.665
Japan	49	0	-2.6704	0.250
Jordan	33	0	-2.3839	0.380
Kenya	33	0	-2.3102	0.415
Kuwait	49	0	-1.2534	0.885
Laos	30	3	-0.2245	0.990
Lebanon	33	0	-2.2132	0.465
Lesotho	33	0	-5.9807***	0.010
Liberia	33	0	-1.3801	0.850
Libya	33	0	-3.9823**	0.020
Luxembourg	33	1	-1.4248	0.835
Macau	33	0	-1.5074	0.805

Continued on next page

Table 4 – *Continued from previous page*

Country	Adj. sample size	Lag length	ADF <i>t</i> -statistic	ADF <i>p</i> -value
Madagascar	33	0	−1.5979	0.770
Malawi	33	0	−1.8363	0.660
Malaysia	49	0	−1.8905	0.645
Maldives	33	0	−2.7696	0.215
Mali	32	1	−3.6629**	0.040
Malta	33	0	−1.7093	0.720
Martinique	33	0	−2.0828	0.535
Mauritania	33	0	−0.9561	0.935
Mauritius	33	0	−2.5584	0.300
Mexico	49	0	0.6512	0.990
Mongolia	33	0	0.0113	0.990
Morocco	33	0	−0.4795	0.980
Mozambique	33	0	−1.8210	0.670
Namibia	19	4	−2.6599	0.260
Nepal	33	0	−1.9466	0.605
Netherlands	49	0	−1.8209	0.680
Netherlands Antilles	34	0	−0.5557	0.975
New Caledonia	33	0	−0.7687	0.960
New Zealand	49	0	−1.4113	0.845
Nicaragua	33	0	−0.6729	0.965
Niger	33	0	−1.0842	0.915
Nigeria	34	0	−2.2462	0.450
North Korea	32	1	−2.3616	0.390
Norway	49	0	−3.8401**	0.025
Oman	30	3	−7.7716***	0.010
Pakistan	46	3	−2.8753	0.180
Palestinian Territories	16	0	−1.6218	0.740
Panama	30	3	−3.5672**	0.050
Papua New Guinea	33	0	−0.7451	0.960
Paraguay	33	0	−0.5885	0.970
Peru	48	1	−1.4156	0.845
Philippines	49	0	−1.3669	0.855
Poland	48	1	−2.7663	0.215
Portugal	46	3	−1.0507	0.925
Puerto Rico	33	0	−0.1843	0.990
Qatar	32	1	−1.7025	0.725
Reunion	33	0	−0.6138	0.970
Romania	49	0	−2.3297	0.410
Rwanda	33	0	−1.3933	0.845
Saudi Arabia	49	0	0.1454	0.990
Senegal	33	0	−0.8011	0.955
Seychelles	33	0	−0.6179	0.970
Sierra Leone	34	0	−0.7923	0.955
Singapore	49	0	−0.6049	0.975
Slovakia	49	0	−2.6181	0.275
Somalia	34	0	−1.2223	0.890
South Africa	49	0	−2.9739	0.150
South Korea	49	0	−1.1181	0.915

Continued on next page

Table 4 – *Continued from previous page*

Country	Adj. sample size	Lag length	ADF <i>t</i> -statistic	ADF <i>p</i> -value
Spain	46	3	−2.5675	0.300
Sri Lanka	33	0	−1.0886	0.915
Sudan	30	3	−3.0661	0.130
Suriname	34	0	−0.8060	0.955
Swaziland	34	0	−1.2300	0.885
Sweden	49	0	−3.2349*	0.090
Switzerland	49	0	−3.3080*	0.075
Syria	29	4	−1.1917	0.895
Taiwan	49	0	−0.7636	0.960
Tanzania	31	2	−2.4139	0.365
Thailand	48	1	−2.0430	0.565
Togo	33	0	−1.2563	0.880
Trinidad and Tobago	49	0	−2.1302	0.515
Tunisia	33	0	−2.4816	0.335
Turkey	49	0	−1.7619	0.705
Uganda	29	4	6.1265	0.990
United Arab Emirates	49	0	0.0588	0.990
United Kingdom	49	0	−3.2175*	0.095
United States	48	1	−3.0728	0.125
Uruguay	33	0	−1.8204	0.670
Venezuela	49	0	−2.2747	0.435
Vietnam	46	3	−1.8901	0.645
Virgin Islands, U.S.	33	0	−1.9992	0.580
Wake Island	29	4	0.8315	0.990
Western Sahara	33	0	−1.1380	0.905
Yemen	33	0	−0.8129	0.955
Zambia	31	0	−2.2099	0.470
Zimbabwe	31	0	−1.1188	0.910
Former Soviet Union	48	1	−2.4878	0.335

Note: The sample period is 1965–2014. Under the alternative hypothesis, consumption follows a trend-stationary process. ***/**/* indicates that the null hypothesis of a unit root is rejected at the 1/5/10% significance level.

Figure 1: Cumulated hydrocarbon volumes in mmboe of giant oil fields by their year of discovery



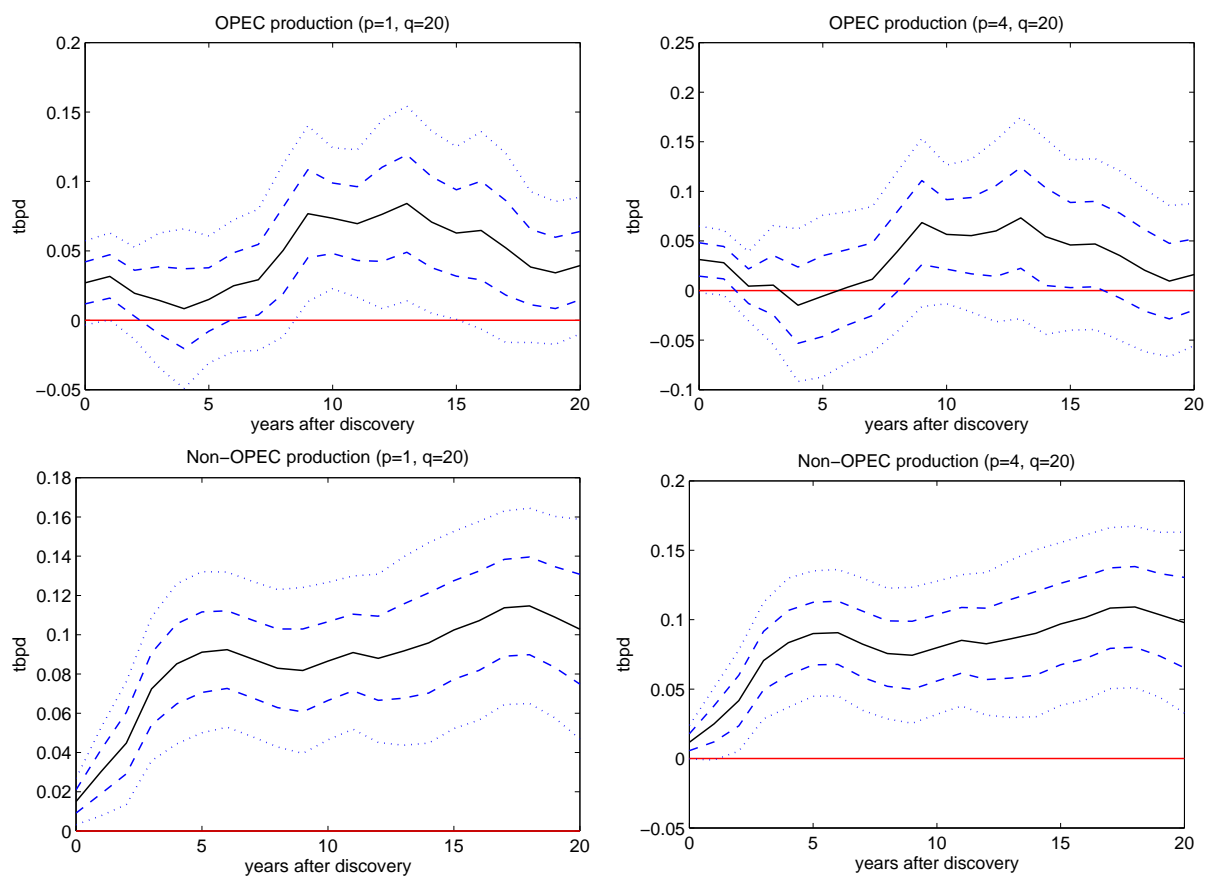


Figure 2: Impulse response of cumulated changes in country-level petroleum production to a giant oil field discovery based on the dynamic panel DLR model in (1)



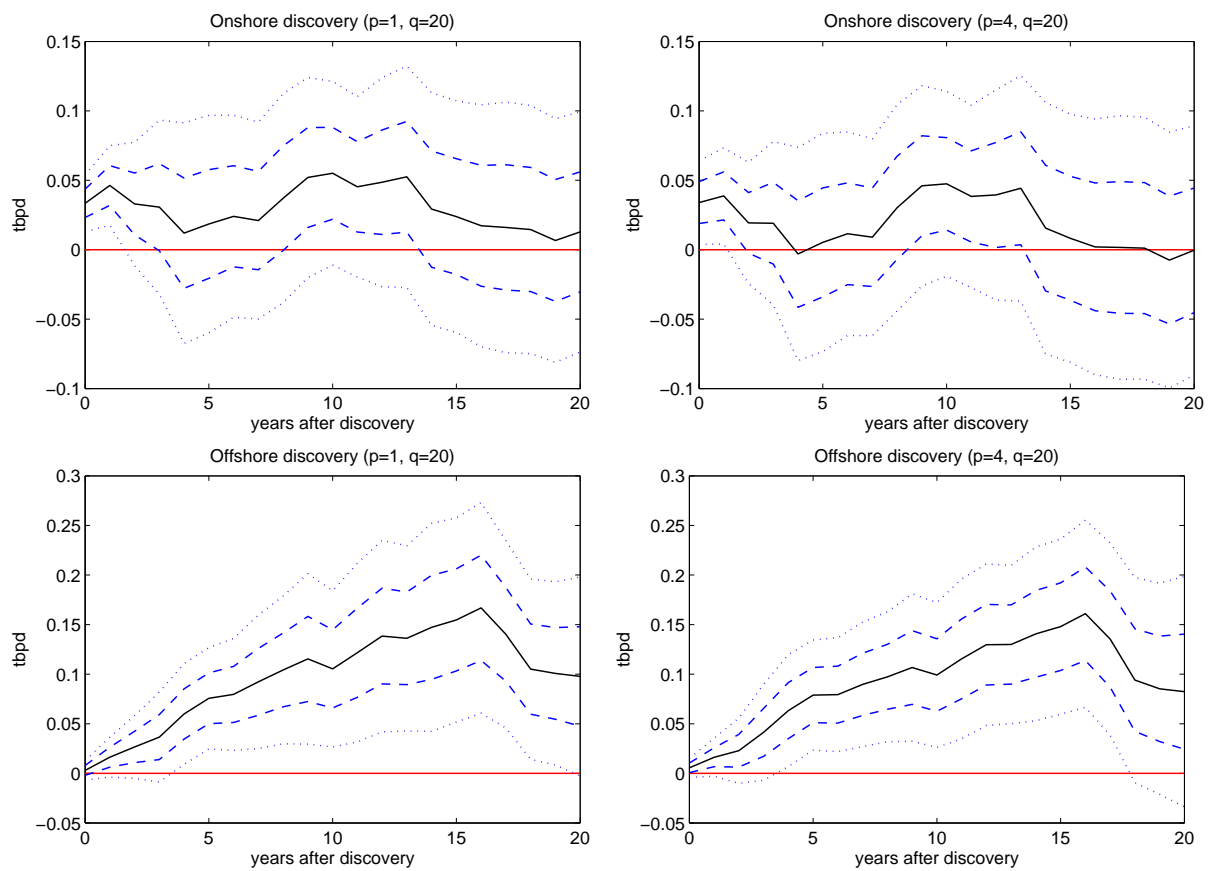
Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications

Figure 3: Impulse response of cumulated changes in OPEC and non-OPEC production to a giant oil field discovery based on the dynamic panel DLR model in (3)



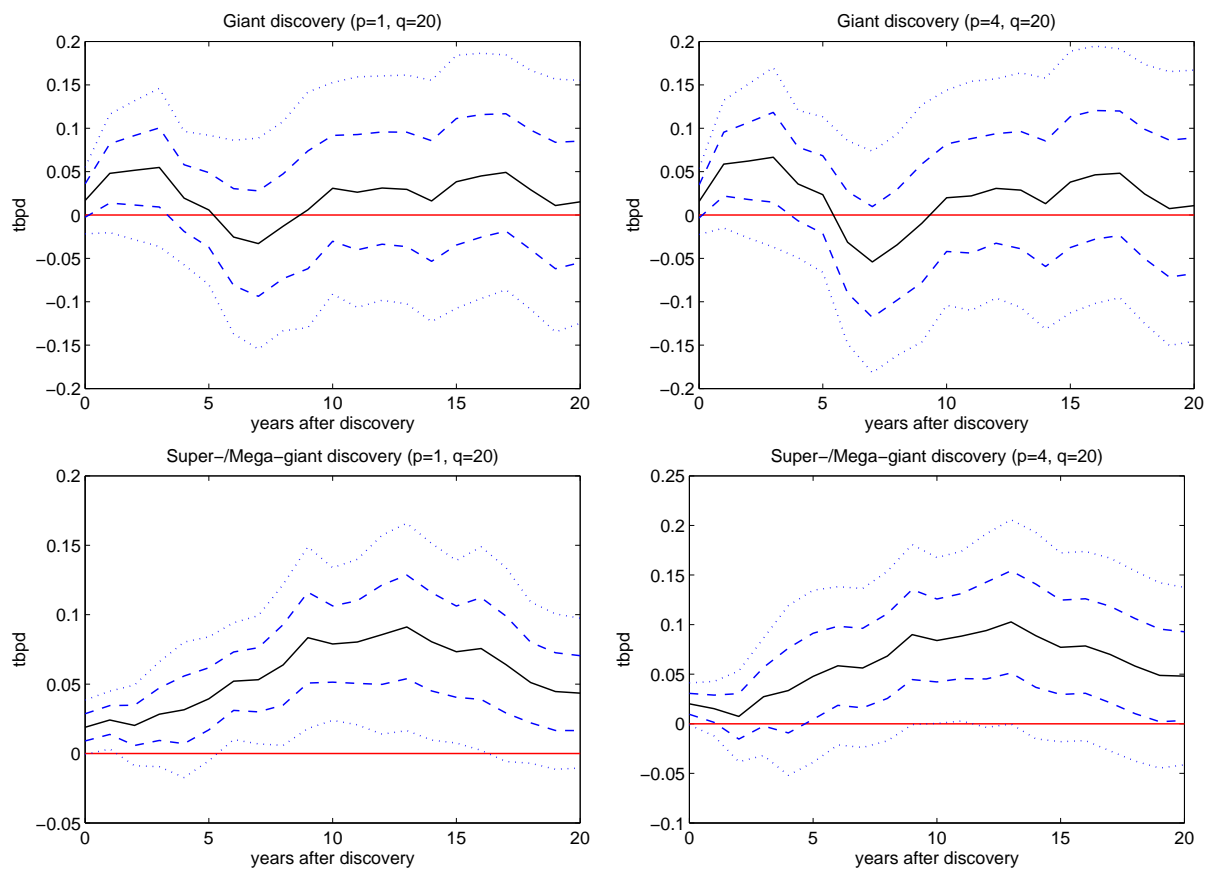
Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications

Figure 4: Impulse response of cumulated changes in production to giant onshore and offshore oil field discoveries based on the dynamic panel DLR model in (5)




Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications

Figure 5: Impulse response of cumulated changes in production to giant and super-/mega-giant oil field discoveries based on the dynamic panel DLR model in (7)



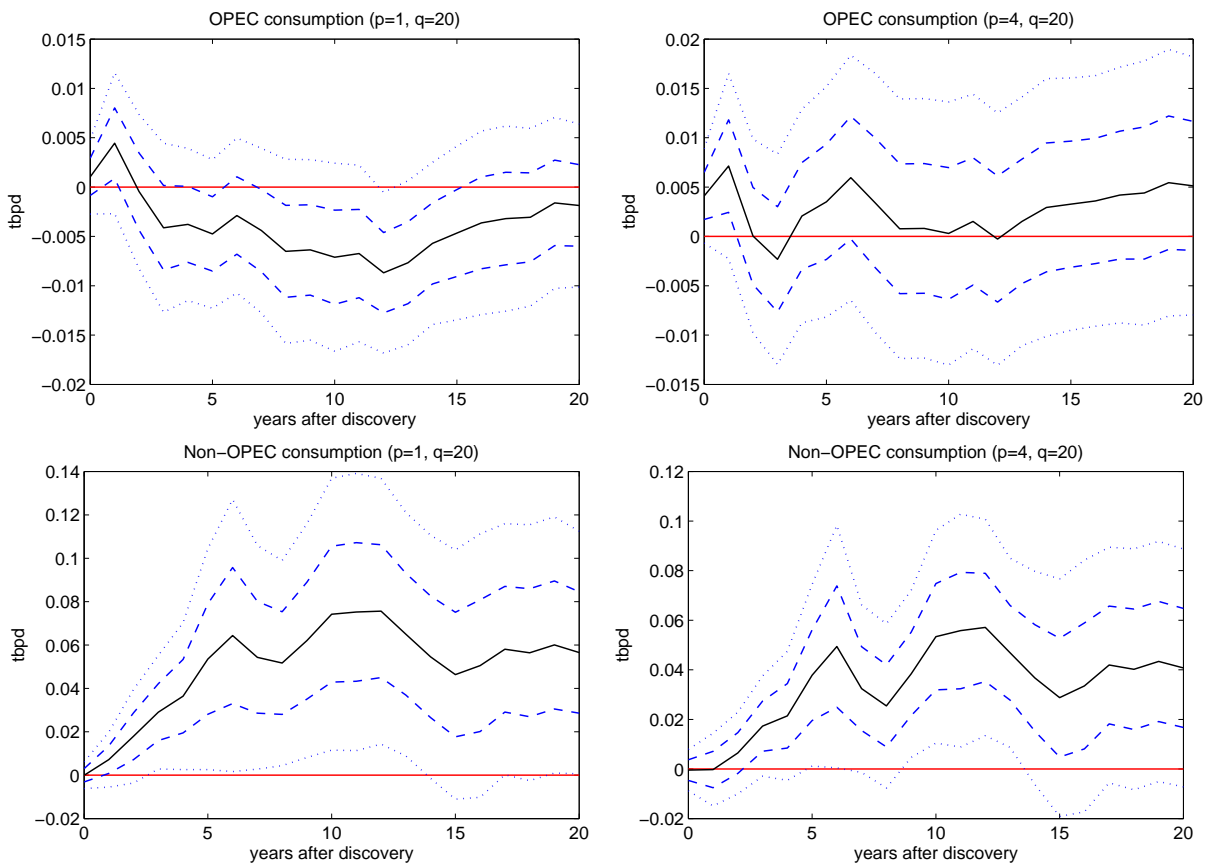
Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications

Figure 6: Impulse response of cumulated changes in country-level petroleum consumption to a giant oil field discovery based on the dynamic panel DLR model in (1)



Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications

Figure 7: Impulse response of cumulated changes in OPEC and non-OPEC consumption to a giant oil field discovery based on the dynamic panel DLR model in (3)



Note: Point estimates with one- and two-standard-error HAC-robust confidence intervals based on 1,000 bootstrap replications