UC San Diego

UC San Diego Previously Published Works

Title

News Shocks in Open Economies: Evidence from Giant Oil Discoveries

Permalink

<https://escholarship.org/uc/item/9pz945h6>

Authors

Arezki, Rabah
Ramey, Valerie A
Sheng, Liugang

Publication Date

2015-09-01

Peer reviewed

NEWS SHOCKS IN OPEN ECONOMIES: EVIDENCE FROM GIANT OIL DISCOVERIES*

Rabah Arezki, Valerie A. Ramey and Liugang Sheng*

July 22, 2016

Abstract

This paper explores the effect of news shocks in open economies using worldwide giant oil and gas discoveries as a directly observable measure of news shocks about future output—the delay between a discovery and production is on average 4 to 6 years. We first analyze the effects of a discovery in a two-sector small open economy model with a resource sector. We then estimate the effects of giant oil and gas discoveries on a large panel of countries. Our empirical estimates are consistent with the predictions of the model. After an oil or gas discovery, the current account and saving rate decline for the first 5 years and then rise sharply during the ensuing years. Investment rises robustly soon after the news arrives, while GDP does not increase until after 5 years. Employment rates fall slightly and remain low for a sustained period of time.

JEL Classification: E00, F3, F4.

Word count: 17,016

* Research Department, International Monetary Fund, OxCarre and Brookings Institution (Arezki), University of California, San Diego and NBER (Ramey) and The Chinese University of Hong Kong (Sheng). Contact e-mail: rarezki@imf.org; vramey@ucsd.edu; lisheng@cuhk.edu.hk. We are grateful to Mike Horn and the Global Energy Systems research group at Uppsala University for sharing their datasets on giant oil discoveries. We thank Robert Barro, Ritwik Banerjee, Olivier Blanchard, Raouf Boucekine, Yi Chen, Giancarlo Corsetti, Domenico Fanizza, Thorvaldur Gylfason, Nir Jaimovich, Jean-Pierre Laffargue, Prakash Loungani, Akito Matsumoto, Gian Maria Milesi-Ferretti, Maury Obstfeld, Jonathan Ostry, Rick van der Ploeg, Richard Rogerson, Tony Venables, Henrik Wachtmeister, Philippe Wingender, Hongyan Zhao, and referees for detailed comments and discussions. We also thank seminar and conference participants for helpful comments. We thank Rachel Yuting Fan for excellent research assistance. Sheng thanks the Research Grants Council of Hong Kong for financial support. The views expressed in this paper are those of the authors and do not necessarily reflect those of the International Monetary Fund, its Board of Directors or the countries they represent. All remaining errors are ours.

I. INTRODUCTION

Economists have long explored how changes in expectations affect the behavior of forward-looking agents. This literature dates back at least to Pigou (1927) and Keynes (1936) who suggested that changes in expectations may be important in driving economic fluctuations. In closed-economy macroeconomics, a seminal paper by Beaudry and Portier (2006) triggered a resurgence of interest in news-driven business cycles by providing evidence that news about future productivity could explain half of business cycle fluctuations in the United States. Since then, there has been a growing number of studies using various identification methods to explore the importance of so-called “news shocks” in driving business cycles. In open-economy macroeconomics, the intertemporal approach to the current account has sought to explain fluctuations in the current account as the optimal response to changing expectations of future output growth (e.g. Sachs (1981), Obstfeld (1982), Persson and Svensson (1985), Engel and Rogers (2006)). The main challenge in both literatures has been to identify news shocks and to provide evidence of “anticipation effects” following those shocks. Unfortunately, there is little direct evidence of the empirical relevance of the effect of news shocks on macroeconomic variables.²

This paper provides empirical evidence of the effect of news shocks on the current account and other key macroeconomic variables using plausibly exogenous variation in the timing of worldwide giant oil and gas discoveries as a directly observable measure of news shocks about higher future output. The delay between a discovery and production is on average 4

² Some of the few examples are in the fiscal literature, which has employed measures of news of future fiscal actions (e.g. Ramey (2011), Barro and Redlick (2011), Mertens and Ravn (2012), Kueng (2012)). Alexopoulos (2011) measures productivity shocks with book publications, but the publications represent information about contemporaneous innovations, not news about future innovations.

to 6 years. A giant oil or gas discovery is defined as a discovery of an oil and/or gas field that contains at least a total of 500 million barrels of ultimately recoverable oil equivalent.³ Hereafter we refer to discoveries of giant oil (including condensate) and gas fields as simply “giant oil discoveries.”

We first extend the Jaimovich and Rebelo (2008) small open economy model to include two sectors, where one sector is a resource sector with oil discoveries. We use this model to develop the macroeconomic predictions for news about oil discoveries and to determine how they might differ from the standard aggregate total factor productivity (TFP) news shock. In the empirical work, we construct a net present value (NPV) of the oil discovery as a percentage of GDP at the time of the discovery. We then estimate a dynamic panel distributed lag model over a sample covering the period 1970-2012 for up to 180 countries. Our empirical estimates of the effects of oil discoveries on key macroeconomic variables are largely consistent with the predictions of our model.

A historical example of giant oil discoveries is Norway. The country borrowed extensively to build up its North Sea oil production facilities following the first several discoveries in the late 1960s and early 1970s (see Obstfeld and Rogoff, 1995 pp. 1751 and Figure 2.3). Meanwhile, Norway’s saving rate also declined due to the expectation about higher future output. The rise in investment and the decline in saving translated into a sharp current account deficit approaching minus 15 percent of GDP at its trough in the year 1977. The current account then started to improve as saving began to rise and investment demand declined following the start of massive oil exports.

³ Ultimately recoverable reserves refer to the amount that is technically recoverable given existing technology.

This example illustrates three unique features of giant oil discoveries that make them an ideal candidate for a measure of news about future production possibilities: the relatively significant size, the production lag, and the plausible exogenous timing of discoveries. First, giant oil discoveries represent a significant amount of oil revenue for a typical country of modest size. The median value of the constructed NPV as a percentage of country's GDP is about 9 percent. Giant oil discoveries provide a unique source of macro-relevant news shocks since it is difficult to find other direct measures of news shocks at the country level that have similar significance. Second, giant oil discoveries do not immediately translate into production. Instead, there is an initial burst of oil field investment for several years and production typically starts with a substantial delay of 4 to 6 years on average following the discovery. Giant oil discoveries thus constitute news about future output increases. This feature is unique in the sense that other plausibly exogenous and directly observable shocks used in other strands of literature such as natural disasters are contemporaneous. Third, the timing of giant oil discoveries is plausibly exogenous and unpredictable due to the uncertain nature of oil exploration. Thus exploiting the variation in the timing of giant oil discoveries provides a unique way to identify the news effect on macro variables.⁴

To estimate the dynamic impact of giant oil discoveries on macro variables, we adopt a dynamic panel distributed lag model. Panel techniques including year- and country- fixed effects allow us to control for global common shocks and cross-country difference in time invariant factors such as countries' geographical location, institutions, and culture. In addition, exploiting

⁴ A limited number of papers have used giant oil discoveries in the context of studies of democratization and conflicts. Tsui (2011) explores the impact of giant oil discoveries on medium run democratization. Cotet and Tsui (2013) and Lei and Michaels (2011) study the relationship between giant oil discoveries and civil conflicts. To the extent of our knowledge, we are the first to exploit giant oil discoveries as news shocks to test the predictions of a standard macro model with news.

solely within-country variations in the timing of the giant oil discoveries allays concerns about endogeneity bias that would have otherwise resulted from omitted variable problems. The impulse responses are qualitatively consistent with the predictions of the model. In the years immediately following the discoveries, the current account decreases significantly as investment rises and the saving rate declines. Five years after the discovery, the average effect of giant oil discoveries on the current account turns positive and significant, as output and saving rise and investment declines. A peak effect is reached about eight years following the discovery after which the effect of giant oil discoveries gradually declines. Interestingly, employment rates decline after the news arrives and remain below normal for over 10 years. We also explore numerous empirical extensions, such as the difference between onshore and offshore discoveries, the effects of financial market openness, and the respective roles of the private and public sectors in explaining our main results.

Our results are robust to a wide array of checks. First, our results are robust to numerous permutations of the oil discovery variables, including simple dummy variables and alternative ways of constructing the net present value of oil revenues. Second, we find that our results are not driven by a particular group of countries. Removing groups of countries including countries in Middle East and North Africa, major oil exporters or countries without any discoveries do not alter the pattern of the dynamic effects of giant oil discoveries. Third, because discoveries that immediately follow a discovery could be seen as predictable, we check whether our main results still hold if we remove them. We also selectively use discoveries that occurred when no discoveries happen in the last past three years and separately control for current and lagged values of exploration expenditures. All our results are virtually unchanged. Finally, our results

are also robust to using different model specifications, particularly including higher order lags for the dependent variable and for giant oil discoveries.

Our paper contributes to both the closed economy and open economy literatures on news-driven fluctuations. In the closed economy literature, Barro and King (1984) and Cochrane (1994) pointed out that news about future TFP could not be a driver of business cycles in a standard real business cycle (RBC) model since news about future production possibilities should lead to an initial rise in consumption and *fall* in labor because of the wealth effect. Using time series techniques to identify news shocks from stock prices and TFP, Beaudry and Portier (2006) found empirical evidence that labor increased in response to news, and that news shocks could account for 50 percent of the business cycle variation of output. Beaudry and Lucke (2009), Schmitt-Grohe and Uribe (2012), Blanchard, L’Huillier and Lorenzoni (2012), and Kurmann and Otrok (2013) used other techniques to reach similar conclusions. In response, Beaudry and Portier (2004), Jaimovich and Rebelo (2008, 2009), den Haan and Katlenrunner (2009), and others developed models that could produce an increase in labor input in response to news. Schmitt-Grohe and Uribe (2012) and Miyamoto and Nguyen (2015) estimated DSGE models allowing for news about a variety of shocks (not just TFP) and found that news shocks were a major driver of business cycles. More recently, however, Barsky and Sims (2011) and Barsky, Basu and Lee (2014) have used time series techniques to identify TFP news shocks from consumer confidence and found that news shocks did not generate business cycle fluctuations. Moreover, Fisher (2010) and Kurmann and Mertens (2014) have highlighted problems with Beaudry and Portier’s identification method. Ramey (2015) finds very low correlations between the TFP news shocks

identified using the different methods. Thus, the empirical work based on time series identification is in flux.⁵

The unique timing characteristic of oil discoveries provides a methodological contribution to the identification problem of news shocks and the associated anticipation effects. Standard approaches in this literature rely on VARs or DSGE models, both of which require many untested identification assumptions and are thus subject to debate. Exploiting the natural lags between giant oil discoveries and the subsequent increase in production provides a unique way to directly measure news shocks about future output increase. In turn, that allows us to conduct a quasi-natural experiment that does not rely on identification using VARs or parametric DSGE models. Our approach provides direct evidence on how news shocks affect macroeconomic variables. Since our approach identifies only one type of news shock, it cannot reveal what fraction of output or current account fluctuations are driven by news shocks. However, our new results can be used to shed light on other methods for identifying news. For example, one could test a time series identification method to see whether it can accurately uncover the oil discovery shocks and produce responses that match our estimated responses.

In a similar vein, this paper also provides direct evidence for the classic intertemporal approach to the current account (e.g. Obstfeld and Rogoff (1995)). That approach uses insights from the permanent income hypothesis to make predictions about the current account based on the intertemporal budget constraint of an open economy. Testing the intertemporal model is difficult, though, because there are few direct measures of expectations about future output or productivity. Typically, time series methods are used for empirical testing, but often the results

⁵ See Beaudry and Portier (2014) and Krusell and McKay (2010) for recent surveys of the literature on news shocks and business cycle fluctuations.

are sensitive to the particular assumptions used (e.g. Ghosh and Ostry (1995), Bergin and Sheffrin (2000), Corsetti and Konstantinou (2012)). We find evidence for a statistically and economically significant anticipation effect on the current account through both the saving and investment channels following the announcement of a giant oil discovery, supporting the view that expectations can be an important driving force for the current account dynamics. Thus, this empirical finding also contributes to a broader literature exploring the empirical determinants of the current account and its adjustment to shocks (e.g. Chinn and Prasad, 2003; Chinn and Wei, 2013).

The remainder of the paper is organized as follows. Section II presents a two-sector small open economy model to develop the implications of news from giant oil discoveries. Section III discusses the relevance of using giant oil discoveries. Section IV lays out the empirical strategy and Section V presents the main results. Section VI presents some extensions and Section VII discusses robustness checks. Section VIII concludes.

II. OIL DISCOVERIES IN A SMALL OPEN ECONOMY

In a simple endowment open economy, news of a future increase in output should produce an immediate rise in consumption and an immediate fall in the saving rate and current account as the country borrows abroad. Once the new resources become available, the saving rate and current account should swing from negative to positive as the country pays off its debt and also saves for the future.

Oil discoveries in a production economy add complications, however, because exploitation of the resources requires sector-specific investment. Moreover, as we discuss below,

the oil sector has a much lower labor share and higher capital share than the rest of the economy. In order to understand how these complications change the predictions for key macroeconomic variables, we analyze a stylized two-sector model that extends Jaimovich and Rebelo's (2008) (JR) one-sector model of news in a small open economy. We add a resource sector to their model in order to capture important features of news about oil discoveries. We use this model both to generalize the intuition from the endowment economy and to compare the effects of news of oil discoveries to the canonical case of news about future TFP which has been the main focus of the news literature.

We find that news about oil discoveries causes the current account/GDP ratio to swing negative initially and then to swing positive once the oil production starts. The responses of both the investment/GDP ratio and the saving rate drive the behavior of the current account. GDP does little for the first several years and then rises once the oil production starts, but consumption jumps as soon as the news arrives. Thus, the qualitative predictions of the intertemporal approach to the current account extend to the special case of oil discoveries. The behavior of labor input, however, is heavily dependent on the details of the model; it falls in some cases and rises in other cases.

II.A. Model Setup

Consider an economy populated by identical agents who maximize their lifetime utility U defined over sequences of consumption (C_t) and hours worked (N_t).

$$U = E_0 \sum_{t=0}^{\infty} \beta^t \frac{(C_t - \psi N_t^\theta)^{1-\sigma} - 1}{1-\sigma} \quad (1)$$

where $0 < \beta < 1, \theta > 1, \psi > 1$, and $\sigma > 0$. In our baseline model, we use Greenwood, Hercowitz, and Huffman (1988) (GHH) preferences, which shut down the wealth effect on labor supply and are now standard in open economy models (e.g. Correia, Neves, and Rebelo (1995), Uribe and Schmitt-Grohe (2015)).⁶ The household provides capital and labor in a competitive market.

There are two sectors in the economy: an oil sector and non-oil sector. The non-oil goods sector uses capital, K_1 , and labor, N_1 , with a constant returns to scale Cobb-Douglas production function of their inputs:

$$Y_{1,t} = A_{1,t} N_{1,t}^{\alpha_1} K_{1,t-1}^{1-\alpha_1}, \quad (2)$$

A_1 is total factor productivity (TFP) in sector 1 and $K_{1,t-1}$ is defined to be capital in sector 1 at the *end* of period $t - 1$ (or beginning of period t). Sector 2 is the oil sector, which uses capital, labor, and the stock of *producing oil reserves* with a Cobb-Douglas production:

$$Y_{2,t} = A_{2,t} N_{2,t}^{\alpha_2} K_{2,t-1}^{\alpha_k} R_{t-1}^{1-\alpha_2-\alpha_k}, \quad (3)$$

⁶ Jaimovich and Rebelo (2008, 2009) use more general preferences that nest both GHH and King, Plosser, and Rebelo (1988) preferences. However, Jaimovich and Rebelo calibrate their parameters so that the preferences are very close to GHH preferences.

where $0 < \alpha_1, \alpha_2, \alpha_k < 1$, and R_{t-1} is the stock of oil reserves available for production in period t .⁷ We will discuss more details of oil reserves below.

Following Jaimovich and Rebelo (2008, 2009), we assume that there are adjustment costs on investment, I . The adjustment costs are on sectoral investment, so that intratemporal reallocation of capital between the two sectors is impeded, which is plausible given the sectoral specificity of capital. Thus, the capital accumulation equation for each sector is:

$$K_{h,t} = I_{h,t} \left[1 - \frac{\phi}{2} \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right)^2 \right] + (1 - \delta)K_{h,t-1}, h = 1, 2 \quad (4)$$

with adjustment cost parameter $\phi > 0$ and depreciation rate δ between 0 and 1. The functional form implies that there are no adjustment costs in the steady state.

For simplicity, we assume that all goods are tradeable and that households consume only good 1, but can exchange oil for good 1 on international markets.⁸ Thus, the flow budget constraint is given as follows:

$$B_t = (1 + r_t)B_{t-1} + (Y_{1,t} + p_t Y_{2,t}) - \{C_t + I_{1,t} + I_{2,t}\} \quad (5)$$

where B_t is net foreign assets at the end of period t , which are denominated in the non-oil good, r_t is the interest rate, and p_t is the relative price of oil determined by the world market.⁹ To

⁷ Reserves appear in the value-added production function as an oil specific physical capital rather than an intermediate material to incorporate the standard assumption used in the natural resource literature that the marginal cost of extraction rises as the oil reserves are depleted.

⁸ See Pieschacón (2012) for an analysis of the effects of oil price shocks on oil exporters using a small open economy model with both tradeable and nontradeable produced goods. She assumes that oil is a non-produced endowment in order to simplify the analysis.

⁹ We are implicitly assuming that the country does not immediately sell the oil field, since this is rare in practice for giant oil discoveries. If the country sells the oil field upfront, then the responses of the current account and saving rate may be different from those we feature, depending on how the transaction is recorded in the Balance of Payments and whether the payment is contingent on successful production.

induce stationarity of foreign bond holdings, we follow the external debt-elastic interest rate proposed by Schmitt-Grohe and Uribe (2003),

$$r_t = r_w + \chi[\exp(\bar{B} - B_{t-1}) - 1] \quad (6)$$

where r_w is the world interest rate, and $\chi > 0$ is the interest rate debt elasticity. The second term on the right-hand side is the risk premium which is decreasing in the country's aggregate net foreign assets. We assume these effects are not internalized by the representative agent.

The current account is defined as

$$CA_t = B_t - B_{t-1} = SA_t - I_{1t} - I_{2t}, \quad (7)$$

where SA_t is saving.

Aggregate output, capital, investment, and labor are defined as:

$$Y_t = Y_{1,t} + p_t Y_{2,t}, \quad K_t = K_{1,t} + K_{2,t}, \quad I_t = I_{1,t} + I_{2,t}, \quad N_t = N_{1,t} + N_{2,t} \quad (8)$$

Even if a country starts with an oil sector, there is typically no capital in place at the site of a newly-discovered oil field. Moreover, most of the investment in capital in the new oil field must be completed before the first barrel of oil is extracted. Figure I shows oilfield investment and production for two oilfields in Norway. Note how investment displays a dramatic hump after discovery, but that oil production starts only after investment falls toward zero. At the Jotun oilfield, production rises rapidly before gradually declining; at the Draugen oilfield, production rises more gradually and declines more gradually.

To capture these features, we would ideally analyze the effect of a discovery when there is no initial capital or labor at the site of the oil discovery. Unfortunately, this approach is not computationally feasible since standard perturbation methods cannot be used in models with

values of zero in steady states. We are then left with the problem that, even with time to build on capital, a social planner would reallocate labor immediately to combine with the positive pre-existing stock of capital to exploit the newly discovered oil.¹⁰ We circumvent this problem by making a distinction between *known reserves* and *producing reserves* and by introducing a *time to connect*. *Known reserves* appear as soon as the oil is discovered but become *producing reserves* only when the pipelines have been connected to the capital and labor, which takes time. This *time to connect* feature captures the time delay between oil discovery and the first oil production. The stock of *producing* reserves evolves as follows:

$$R_t = \bar{R} + R_{t-1} - Y_{2,t} + \varepsilon_{t-j} \quad (9)$$

Producing reserves at the end of year $t-1$, R_{t-1} , are augmented with an exogenous stream of new inflows, $\bar{R}$, and are endogenously depleted by the production of oil, $Y_{2,t}$.¹¹ ε_{t-j} captures the interaction of news of an oil discovery and the time-to-connect feature; in period $t-j$, news of an oil discovery arrives. *Known* oil reserves rise immediately at $t-j$, but *producing* reserves R do not rise until period t because it takes time to connect them to the capital and labor. Thus, the lag on ε_{t-j} captures the key feature that the reserves are not immediately available for production when the news about the discovery is revealed.

¹⁰ Even very high labor adjustment costs do not slow down the reallocation much because the returns to exploiting the oil immediately are very high. Our assumption of adjustment costs on investment mimics many aspects of time to build for investment dynamics (Lucca (2007)), but does not overcome the problem of the initial positive stock of capital being used.

¹¹ We assume the constant stream of exogenous reserve inflow only to avoid the computational problems caused by steady-states with zero reserves. One could endogenize the exploration and discovery process, as in Pindyck (1978), Bohn and Deacon (2000), and Gross and Hansen (2013), but doing so would add nothing to the intuition about the effect of news.

The first-order conditions for the representative agent are presented in the theoretical appendix. Our baseline calibration is summarized in Table I. Many of the parameters are similar to those in Jaimovich and Rebelo (2008), with relevant ones converted to an annual basis to match our data. The new parameters for the resource sector are set to match some key facts. Following Gross and Hansen (2013), we set the labor share to 13 percent and the capital share to 49 percent, leaving a resource share of 38 percent. These numbers are also broadly consistent with U.S. data.¹²

II.B. Model Simulation Results

The typical lag between discovery and initial oil production is five years, as discussed in more detail in the next section. Thus, we explore the effects of a news shock ε_{t-5} in equation (9). The shock is normalized so that the present value of the rise in oil revenue is equal to one percent of initial GDP in the baseline model.

Figure II shows the predictions of our stylized model for the effects of oil discovery news that arrives in year 0. The solid lines show the results for the baseline simulation with GHH preferences. The upper left graph shows that news of an oil discovery leads the current account to turn negative for five years before becoming sharply positive. The two lower left graphs show that the initial decline in the current account comes from both a decline in the saving rate and an increase in the investment rate. The saving rate declines initially because of the wealth effect from the anticipation of the new resources. The aggregate investment rate rises because of the

¹² For example, in the U.S. labor share is 13 percent of value added in the oil and gas extraction sectors. A comparison of the estimates of the value of resources in *The Survey of Current Business*, April 1994, pp. 50-72 with the BEA estimates of fixed capital by industry suggest that our capital and resource shares are roughly consistent.

boom in investment in the oil sector.¹³ After the oil sector capital stock is built up, the aggregate investment rate falls below normal for a number of years.

The upper right graph of Figure II shows that GDP does not respond much for the first five years, but then rises significantly at year 5 when the reserves become available for production. It then gradually falls as the extra reserves are depleted. Consumption rises on the arrival of the news and remains permanently higher. Hours fall slightly for the first four years (note the scale of the graph) before beginning to rise. With GHH preferences, the response of hours depends solely on the current wage. Wages fall by a very small amount initially because investment in Sector 1 falls temporarily, reducing the capital-labor ratio.¹⁴ The shift in capital to Sector 2 does not compensate because Sector 2 has much lower labor share.

To determine how many of these effects are due to GHH preferences, the graphs in Figure II also show results from a model with standard King, Plosser, Rebelo (1988) (KPR) preferences, displayed as the dashed lines. The qualitative differences across simulations for the current account, saving, investment, and GDP are small. Thus, the results for those four variables are robust to the differences in preferences. In contrast, the responses of consumption and hours are somewhat different across the two experiments. With KPR preferences, hours decline more as a result of the wealth effect on labor supply, in addition to the reallocation effects from Sector 1 to Sector 2.

It is noteworthy that even with GHH preferences, the macroeconomic effects of oil discovery news do not look anything like a business cycle. Jaimovich and Rebelo (2008, 2009)

¹³ The responses for each sector are shown in the supplemental online appendix.

¹⁴ How much Sector 1 investment and hours fall depends on how much interest rates rise in the short-run. Our calibration follows Jaimovich and Rebelo and sets the debt elasticity of interest rates to be very low, so interest rates barely move. If the elasticity is higher, so that interest rate rises more, investment and hours in Sector 1 fall more.

specifically introduced their preferences and calibrated them to be very close to GHH so that news about future TFP could induce business cycles. However, in our case, the news causes investment rates to move in the opposite direction of both hours and output during key times. For example, investment collapses just when output is rising.

Another difference is the response in hours, which decrease in the short-run even for GHH preferences. The key difference between the effects of oil discoveries and the canonical TFP news shock is the differential labor and capital shares in the oil discovery case. As discussed above, the oil sector has a lower labor share and a higher capital share than the rest of the economy. It is for this reason that the oil news shock does not induce a rise in hours for the first several years.

To illustrate this point, we compare the results of our baseline model with endogenous depletion and different factor shares to an exogenous TFP model in which both sectors have identical factor shares: a labor share of 0.58 and a capital share of 0.32.¹⁵ News is about future TFP in the small sector; to match our baseline experiment, we use the reserves process from the baseline depletion model as an exogenous process for TFP. Figure III shows the results for both the baseline oil discovery model (the solid line) and the identical factor shares model (the dashed line). The results for the current account, the saving rate, the investment rate, and GDP are qualitatively similar across the two simulations, though the responses of output and the investment rate are substantially greater in the identical factor share model. The consumption and the hours responses are very different from the oil news simulation. In this alternative experiment, both consumption and hours rise slowly when the news arrives and then spike up

¹⁵ These parameter values imply slight decreasing returns in each sector. Decreasing returns are necessary for an interior solution since both goods are traded on world markets and only good 1 is consumed domestically.

when the TFP increase is realized. These results show that an oil news discovery shock has different effects on hours and consumption relative to a news shock to a sector that has similar factor shares to the rest of the economy.¹⁶

To summarize, our theoretical analysis shows that the current account, saving rate, investment rate, and output responses to news are robust qualitatively to a variety of specifications. The current account and the saving rate become significantly negative for the five years between the arrival of the news and the increase in resources. Investment booms for several years after the news arrives, and then falls. Output rises only after the investment is made. In contrast, the behavior of some of the other variables, such as labor input and consumption, depends significantly on the details of the types of preferences assumed and whether the sector has factor shares that differ from the rest of the economy.

The theoretical analysis also highlights the importance of expanding the study of the usual “aggregate TFP” news shock to more realistic shocks. The existence of true news shocks that are expected to affect aggregate TFP is not self-evident. In fact, Atalay (2015) presents evidence that sector-specific shocks contribute over half of aggregate volatility in the U.S. Even general purpose technologies are typically not recognized initially for their ability to transform most sectors of the economy. Rather, the recognition of the potential of those technologies to transform most sectors of the economy develops very slowly. The more plausible type of “sudden” news shock is one that affects a few key sectors. The results of our model show that the effects of those news shocks on some variables can depend very much on the specifics of the sectors they hit.

¹⁶ Further explorations showed that the key difference across the experiments shown in Figure III is the difference in factor shares, not whether reserves are subject to endogenous depletion or are modeled as an exogenous TFP process. We also found that an aggregate TFP shock has similar effects to the sectoral TFP shock shown in the graph. See Figure C.II in the Supplemental Appendix.

III. WHY USE GIANT OIL DISCOVERIES?

Evaluating the empirical relevance of news shocks is quite challenging. Difficulties arise on two main fronts. First, theory suggests that the main driving force is agents' perception of future availability of output, but it is empirically difficult to measure agents' expectation as is well-known from the literature on news shocks. The literature generally relies on subtle identification assumptions in the context of VARs, which extract news shocks from stock prices or surveys of expectations about the future (e.g. Leduc and Sill (2013)), or estimation of DSGE models, which are subject to controversies (see for instance, Beaudry and Portier, 2014). This approach is even less promising if we want to test the effect of news shocks on the current account because, as pointed out by Glick and Rogoff (1995), the current account responds to country-specific shocks, rather than global shocks.

We adopt a quasi-natural experiment approach to test the dynamic impact of news shocks on output, the current account, saving, investment, consumption and employment by using giant oil discoveries for a sample covering the period going from 1970 to 2012 and up to 180 countries. The giant oil discovery dataset is from Horn (2014).¹⁷ Three unique features of giant oil discoveries make them ideal candidates for measures of news about future output increase. In turn, exploiting variation in the timing of giant oil discoveries allow us to adopt a quasi-natural experiment approach that does not rely on a VAR structure and on subtle identification assumptions.

¹⁷ We are heavily indebted to Mike Horn, former President of the American Association of Petroleum Geologists, for his guidance through some of the technical considerations discussed in this section.

The first attractive feature of giant oil discoveries is that they signal significant increases in production possibilities in the future. To be able to test the effect of news shocks on the dynamics of macroeconomic aggregates, particularly to isolate a significant anticipation effect, those shocks must be significant for the whole economy. It might be difficult to find other output shocks at the country level that have the macro-relevance of giant oil discoveries. Moreover, giant oil discoveries are relatively rare events within a country-specific location, so we can treat them as country-specific shocks.

Second, there is a significant delay between the discovery and the start of production. Discoveries involve years of delay for platform fabrication, environmental approvals, pipeline construction, refinery and budgetary considerations. Figure I showed the delay for two Norwegian oilfields. Experts' empirical estimates suggest that for a giant oil discovery, it takes between 4 and 6 years to go from drilling to production.¹⁸ Based on our own calculation using an alternative data source that is less comprehensive but contains more detailed information at the field level, we find that the average delay between discovery and production start is 5.4 years.¹⁹ Obviously, there is some heterogeneity between oil and gas fields. One potential source of heterogeneity is the difference between onshore and offshore discoveries. Using the aforementioned alternative dataset, we find that the average delay is 6.7 years for offshore discoveries and 4.6 for onshore discoveries. All in all, the lag between the announcement of oil discoveries and production can be substantial and thus allows us to treat giant oil discoveries as news shocks about future output.

¹⁸ See for instance, <http://www.ellipticalresearch.com/drillingandoilproduction.html>. Mike Horn relies on a 7 year time lag between discovery and production.

¹⁹ The data are from Global Energy Systems, Uppsala University. The dataset includes 358 discoveries of giant oil fields and covers 47 countries. The number of discoveries however shrinks to 157 when considering the period 1970 onwards.

The last attractive feature of giant oil discoveries is that their timing is arguably exogenous and unexpected due to the uncertainty surrounding oil and gas exploration, after controlling for country- and year- fixed effects.²⁰ This feature is crucial for our identification of the anticipation effect on macroeconomic aggregates including the current account because the latter adjusts only after the agents receive the news about giant oil discoveries. Resource exploration is an uncertain activity because it is affected by technological innovation in exploration and drilling, and by the relative knowledge of geological features for a particular location including knowledge about the detailed structure of the oil field, its depth or whether the oil is located in deep water. Some may argue that oil discoveries are somewhat predictable because some countries appear to have larger oil endowments, or because they have had discoveries in the past.²¹ The exact timing of giant oil discoveries is however less likely to be predictable. Moreover, ex ante no one has information about the potential size of discoveries which we will also exploit in our empirical strategy.

Thus, the timing of giant oil discoveries constitutes a unique source of within-country variation that can be used to test both directly and precisely whether news shocks about future output shocks may affect macroeconomic aggregates. Our data covers giant oil discoveries for

²⁰ One might also argue that the precise timing of the announcement of a giant oil discovery could be manipulated by governments or other entities. Based on conversations with Mike Horn, we understand that these concerns about a possible manipulation have little ground. In addition, Mike Horn's dataset is immune from such concerns, as each discovery date included in his dataset has been independently verified and documented using multiple sources which are reported systematically for each discovery date.

²¹ Past discoveries may have two opposite effects on the likelihood of current and future discoveries. On the one hand, cumulative discoveries may drive up discovery costs so that future discoveries become less likely (see Pindyck, 1978). On the other hand, past discoveries foster learning about the geology and render future discovery more likely (see Hamilton and Atkinson, 2013). Thus, past discoveries do not necessarily increase the likelihood of new discoveries, nor reduce the uncertainty about the timing of new discoveries. In order to control for possible serial correlations in oil discoveries, we do include previous discoveries and country and year fixed effect in our empirical regression presented in the next section.

the period 1970-2012 and for a wide range of countries in the world.²² This allows us to adopt panel data estimation techniques which control for country and year fixed effects.

Table II shows the spatial and temporal distribution of giant oil discoveries recorded in Horn (2014)'s data during 1970-2012. In total, 64 countries have had at least one giant oil discovery during the sample period. While the Middle East and North Africa region experienced a total of 97 discovery events out of a total of 371 in the world, other regions such as Asia (74), the Western Hemisphere (74) and the Commonwealth of Independent States and Mongolia (51) also experienced significant numbers of discovery events during the same period.²³ The 1970s is the peak period for giant oil discoveries, but the number of discoveries has been growing since the 1980s. This contradicts the commonly held view that it became more and more difficult to discover new oil fields. Figure IV presents the distribution of the logarithm of the size of giant oil discoveries measured in million barrels of ultimately recoverable oil equivalent. It shows that there is significant heterogeneity in the size of oil discoveries.

IV. EMPIRICAL STRATEGY AND DATA

IV.A. Empirical strategy

To test the theoretical predictions and in particular the existence of an anticipation effect, we use a dynamic panel model with a distributed lag of giant oil discoveries, as follows:

²² The dataset excludes shale-oil formations because these do not constitute discovery news shocks. Most if not all large reserves of synthetic oil and gas in shale rocks in the United States have been known for a (very) long time, as early as the 1920s. Until the mid 2000s, oil extraction from shale rock formations was thought to be too costly and technologically impossible. The breakthrough in technological innovation allowed oil to be extracted from shale formation, but there is very little lag (less than a year) between the first investment and shale production. Thus, fracking is not "news."

²³ A discovery event is a dummy variable that takes a value of 1 if during a given year at least one discovery of either a giant oil or gas field is made in any given country, and zero otherwise. The country grouping is from the International Monetary Fund.

$$y_{it} = A(L)y_{it} + B(L)Disc_{it} + \alpha_i + \mu_t + \gamma_1'Z_{it} + \epsilon_{it}, \quad (10)$$

where y_{it} is the dependent macroeconomic variables including log real GDP in local currency, current account-GDP ratio, saving-GDP ratio, investment-GDP ratio, log real consumption in local currency, and the employment-population ratio; α_i controls for country fixed effects which capture unobserved time invariant characteristics such as geographical location; μ_t are year effects controlling for common shocks, such as global business cycles and international crude oil and gas prices; Z_{it} are other control variables used in the robustness exercises, such as exploration expenditures; and ϵ_{it} is the disturbance. $Disc_{it}$ is the net present value of giant oil discoveries which we describe in greater detail below. $A(L)$ and $B(L)$ are p th and q th order lag operators with $p \geq 1$ and $q \geq 0$. In the benchmark regression, we use $p = 1$ and $q = 10$. In regressions using log levels of variables (rather than as a percent of GDP) and employment rate, we also include country-specific quadratic trends.

The panel structure allows us to identify the dynamic effect of oil discoveries on macroeconomic aggregates, while controlling for country-specific and year fixed effects. Controlling for country fixed effects is important because it allows us to estimate the within-country variation in giant oil discoveries on within-country variation in macroeconomic aggregates and thus to control for any unobservable and time invariant characteristics which may affect giant oil discoveries and macroeconomic aggregates.²⁴ The extensive panel data (both in terms of the number of cross-sectional units, N and time span, T) allows us to utilize fully within

²⁴ It is worth noting that the estimates of the dynamic panel with fixed effect are inconsistent if the time span of the panel, T , is small. In our case, our sample period covers at least thirty years, thus the Nickell bias of order $(1/T)$ is seemingly negligible. However, the Nickell bias relies on asymptotic assumptions. Indeed, Barro (2012) shows that there could be substantial bias in relatively small samples. Relying on the plausible exogenous nature of giant oil discoveries, we also tried excluding country fixed effects and verified that our main results were qualitatively and quantitatively similar. We include the country fixed effects in our benchmark model because they are jointly significant.

country variation in giant oil discoveries. Because of the infrequent nature of giant oil discoveries, and because of the long gestation period surrounding the production process, it is crucial to use a large panel dataset to capture the dynamic effect of those discoveries. The dynamic feature of the panel regression in the form of an autoregressive model with distributed lags allows us to use impulse response function to capture the dynamic effect of giant oil discoveries, which is given by $IRF(L) = B(L)/(1 - A(L))$.

IV.B. Data Construction

Our data set consists of an oil discovery measure combined with macroeconomic data for many countries. We begin by discussing the oil discovery measure. Mike Horn's dataset contains information on the country and year of the discovery, in addition to other key information such as whether the field contains oil and/or gas and the estimated total ultimately recoverable amount in oil equivalent. The ultimately recoverable size for each discovery is based on the estimation of the value at the time of the discovery, rather than potentially revised estimates in subsequent years. It contains the timing of announcements of giant oil discoveries independently of whether discoveries do eventually pan out or not. Because agents should respond to the net present value of the output shock revealed by the discovery news, we construct a measure of the net present value of a giant oil discovery as a percent of GDP, NPV , as follows:

$$NPV_{i,t} = \frac{\sum_{j=5}^J \frac{q_{i,t+j} * oilprice_t}{(1 + r_i)^j}}{GDP_{i,t}} \times 100 \quad (11)$$

NPV for a given country, i , at the time the discovery is made, t , is the discounted sum of gross revenue derived from an approximated oil production profile, $q_{i,t+j}$, from the fifth year

following the discovery to the exhaustion year, J , valued at the oil price prevailing at the time of the discovery. The approximated production profile follows a piece-wise process in the form of reserve specific plateau production followed by an exponential decline (see Höök et al., 2014 and Robelius, 2007).²⁵ Appendix B.I. describes in detail the approximation method relying on estimates using an alternative oil field database. Gross revenues are valued at current international prices. The rationale behind using current international prices to value the production is that oil price series typically follow a random walk process so that current price is the best price forecast.²⁶

To account for the fact that giant discoveries may happen in countries where the perceived political risk is high, we allow for country-specific risk adjusted discount rates. Indeed, exploiting oil and gas fields can be rendered difficult if not impossible in countries where political risk is high. Discoveries in countries where political risk is elevated should thus be discounted more than places where risk is lower. We thus compute the adjusted discount rate as the sum of the risk free rate set to 5 percent and a country specific risk premium.²⁷ The risk free rate is assumed to be the rate prevailing in the United States. Considering that measures of risk premia based on that sovereign bond spreads are not readily available for all countries and they are not necessarily comparable, we use predicted values for risk premia based on the historical relationship between observed (and consistent) measures of sovereign bond spreads and political

²⁵ We choose not to use the so-called Hubbert curve to approximate oil production profiles since it is regarded by petroleum engineers as a good fit for aggregated field production profiles for a whole region or at the global level. For single fields, a reserve specific piecewise process consisting in a plateau production and then an exponential decline is commonly used. We use the engineering-determined depletion schedule rather than the endogenous depletion rule assumed in our model. As Anderson, Kellogg, and Salant (2014) show, rates of depletion are constrained by reservoir pressure and do not appear to respond to price.

²⁶ See Hamilton (2009) and references therein for a discussion on forecasting oil prices.

²⁷ Some researchers have however argued that an annual interest rate as high as 14 percent is needed to be consistent with United States' consumption-income relationships in a closed economy setting (see Bernanke, 1985). Using alternative values for the risk free rates does not significantly affect our main results.

risk ratings. The data on spreads on sovereign bonds are from the Emerging Markets Bond Index Global (EMBI Global) that is available for 41 emerging market economies for the period 1997-2007.²⁸ Emerging markets are a set of countries for which risk ratings can vary substantially and thus provide significant statistical variation for estimating a relationship between risk ratings and sovereign bond spreads. Bond spreads are measured against a comparable US government bond and are period averages for the whole year. The political risk rating is available for 138 countries in International Country Risk Guide (2015), which covers most of countries with at least one giant oil discovery. To examine the effects that political risk has on sovereign bond spreads, we estimate the following econometric model:

$$\ln(\text{Spread}_{i,t}) = \theta_0 + \theta_1 \ln(\text{PoliticalRisk}_{i,t}) + \alpha_i + \mu_t + u_{i,t} \quad (12)$$

where β_i are country fixed effects, μ_t are year effects and $u_{i,t}$ is an error term.²⁹ We estimate the elasticity of the sovereign spreads to political risk ratings using our sample. We then predict the $\widehat{\text{Spread}}_t$ given country's political risk rating and compute the NPV of giant oil discoveries accounting for country specific discount rates.

Figure V presents the histogram of the logarithm of NPV, and it shows the significant heterogeneity in the NPV of oil discoveries. The median NPV is 9 percent of GDP (2.2 in logarithms), and the largest one is estimated to be 63 times of the country's GDP. It should be

²⁸ The availability of the sovereign bond spread data limits the sample size to the following countries: Argentina, Bulgaria, Brazil, Chile, China, Colombia, Cuba, Dominican Republic, Algeria, Ecuador, Egypt, Gabon, Ghana, Greece, Hungary, Indonesia, Iraq, Jamaica, Kazakhstan, Republic of Korea, Lebanon, Morocco, Mexico, Malaysia, Nigeria, Pakistan, Panama, Peru, Philippines, Poland, Russia, El Salvador, Seychelles, Spain, Thailand, Tunisia, Turkey, Ukraine, Uruguay, Venezuela, Vietnam.

²⁹ The estimated coefficients used in the prediction are as follows:

$$\ln(\text{Spread}_{i,t}) = 14.10 - 1.93 \times \ln(\text{PoliticalRisk}_{i,t}),$$

(3.22) (-1.85)

The t-statistics in parenthesis indicates that political risk is a significant determinant for the sovereign bond spreads for emerging markets. R-squared is 0.34.

noted however that the results presented below are robust to using alternative measures for the giant oil discoveries such as NPV with common discount rates and uniform production profile, and a dummy variable for a discovery event.

Our macroeconomic variables are from the IMF (2013), the World Bank (2013), and the International Labor Organization. The data appendix Table B.II gives a more detailed description of the data definition and sources. Because our benchmark measure NPV of giant oil discoveries starts in 1970, and we include 10 lags of oil discovery sizes, our baseline regression uses macro variables from 1980-2012. The data appendix Table B.III provides the summary statistics for our key macro variables.

V. BENCHMARK RESULTS

We now present our benchmark results for the dynamic impact of the risk-adjusted NPV of giant oil discovery on relevant macroeconomic aggregates. Figures VI and VII show the dynamic responses to an oil discovery news shock based on the estimates of the panel autoregressive distributed lag model with country and year fixed effects.³⁰ The shaded areas are 90% and 68% (darker grey) confidence bands based on Driscoll-Kraay (1998) standard errors and the delta method.³¹

Figure VI displays the responses of the current account, saving, and investment. The top panel shows that giant oil discoveries have a negative effect on the current account-GDP ratio in the years immediately following the announcement. Five years after the discovery, the average

³⁰ The supplementary appendix provides the coefficient estimates. Both the country and year fixed effects are jointly significant with a p-value of 0.000. For variables for which we do not include country-specific quadratic trends, we also adopt formal panel unit root tests which rejects unit root hypothesis at standard significance levels.

³¹ We use *xtscc* and *nlcom* commands in Stata. The bands are similar if we instead use a non-parametric bootstrapping method.

effect of giant oil discoveries turns positive. A peak effect is reached eight years following the discovery after which the effect starts declining. Those results are consistent with the theoretical predictions of the two-sector production economy presented earlier. The negative effect of giant oil discoveries on the current account immediately following the announcement strongly supports the existence of an anticipation effect. The timing of the anticipation effect is also consistent with the fact that oil production occurs with a delay of 4-6 years on average. The effect starts to be positive five years after the discovery which is consistent with the timing at which oil production starts and output increases. The second and third panels of Figure VI show that the anticipation effect plays out through both the saving and investment channels. The saving-GDP ratio becomes negative for about five years following the announcement of the discovery, and then becomes robustly positive. On the other hand, the investment-GDP ratio starts to rise one year after the giant oil discovery. It hits a peak around 5 years once oil production starts and then returns to normal quite quickly.

The top panel of Figure VII shows the effects of a giant oil discovery on GDP. On average a giant oil discovery has a slightly negative impact on aggregate GDP initially, but then has a robust positive effect after the start of oil and gas production, about 5 years after the discovery announcement is made. Output peaks about 8 years after discovery and then slowly returns to normal over the following years. The pattern of the response of aggregate output to the news of a giant oil discovery is thus qualitatively consistent with the theoretical predictions from the two-sector model presented earlier.

The second panel of Figure VII shows the effect on log consumption (not as a percent of GDP). The estimates indicate that consumption does not respond much at first, but then does start to rise three years after the discovery. The estimates, however, are imprecise. This

imprecision could reflect the fact that there is substantial measurement error in our consumption variable. As we will discuss below, another issue is that the consumption variable includes both private and public consumption.

The bottom panel of Figure VII shows the response of the employment rate.³² The graph shows that the employment rate begins to fall immediately after a giant oil discovery and continues to fall even after the start of production. The employment rate remains depressed for quite a few years before returning to normal. The estimates are small in magnitude, but precisely estimated.

A comparison of the empirically estimated responses in Figures VI and VII with the theoretical responses in Figure II reveals similar qualitative effects. In both the data and the model, the current account and saving rate fall when the news arrives and then swing positive after several years. Investment spikes up in the short-run and GDP rises only after several years. The persistent decline in employment in the data is more consistent with the model with King, Plosser, Rebelo (1988) preferences, indicating both the wealth effect and the reallocation effect are important for the decline in employment.

Not all of the point estimates of the estimated impulse responses are different from zero at conventional levels of significance. The hypotheses we really want to test, though, are about the general patterns, not whether a response at one particular horizon is statistically different from zero. In particular, we want to test whether the *integral* of the response between the discovery and the start of oil production is different from zero and whether the *integral* of the response after production is different from zero. Table III shows the hypothesis tests for the relevant integrals. We develop the alternative hypotheses to be consistent with our theory. For

³² We use the employment rate rather than total hours because the latter were not available for all of the countries in our sample. Even with the employment rate, the data are available only starting in 1990.

example, we test the null hypothesis that the response of the current account-GDP ratio is greater than or equal to zero against our theoretical prediction that it is negative during the first five years (horizons 0 to 4). Similarly we then test the null hypothesis that the response is less than or equal to zero against the theoretical prediction that it is positive for horizons 5 to 11.

The results show that in most cases we can reject the null hypothesis in favor of the theoretical prediction at standard levels of statistical significance. For example, the response of the current account-GDP ratio is significantly negative between discovery and production, indicating a significant anticipation effect, and is significantly positive after oil production starts. The results are similar for the saving-GDP ratio. The investment-GDP ratio is significantly positive (with a p-value of 0.02) for the first five years, but not for the following years. The GDP response is not different from zero during the first few years, but is significantly positive for the years after the oil production starts up. The consumption response is also significant at conventional level with a p-value of 0.096. Moreover, the employment response is (statistically) significantly negative.

Quantitatively, the empirical estimates suggest that a typical giant oil discovery with the NPV equal to the median value, 9 percent of initial GDP, leads to a peak in GDP of 0.28 percent 8 years after the discovery and a cumulative (undiscounted) change of 1.7 percent (in log points). The investment rate reaches the peak in the sixth year by increasing 0.15 percent of GDP and the cumulative change is about 0.42 percent of GDP. The same size of shock leads the current account to fall in the short-run by 0.21 percent of GDP and to rise in the intermediate run to a peak of 0.4 percent of GDP. The quantitative effect of a typical discovery on saving is roughly

about a half of the effect on the current account. The same size of shock causes consumption to increase by 1.8 percent and the employment rate to decrease by -0.17 percent accumulatively.³³

VI. Extensions

We now explore several extensions of the empirical model. In the first extension, we examine whether the timing of the effects differs for onshore versus offshore discoveries. Second, we study whether the degree of (external) borrowing constraints affects the responses to giant discoveries. Finally, we explore the responses of government spending, real exchange rates, and stock markets to giant oil discoveries.

As mentioned earlier, offshore discoveries typically have longer delays between oil discoveries and first oil production than onshore discoveries. Thus, it is interesting to check whether the current account and saving rate switch from negative to positive later in the case of offshore relative to onshore discoveries, and whether investment-GDP ratio, output and employment respond earlier in the case of onshore discoveries. To do so, we estimate the following extended version of equation (10):

$$y_{it} = A(L)y_{it} + B(L)OnshoreDisc_{it} + C(L)OffshoreDisc_{it} + \alpha_i + \mu_t + \gamma'Z_{it} + \epsilon_{it}, \quad (13)$$

³³ These empirical quantitative results are smaller than the ones implied by the baseline stylized theoretical model. The theoretical analysis is not intended as a quantitative matching exercise, so we did not include additional frictions that would dampen the response, such as KPR preferences, imperfect information, uncertainty, and differences between actual and expected oil production.

where $OnshoreDisc_{it}$ and $OffshoreDisc_{it}$ denote the NPV of onshore and offshore discoveries respectively, and $B(L)$ and $C(L)$ are the q th order lag operators.³⁴ The IRFs for onshore and offshore discoveries are given by $B(L)/(1 - A(L))$ and $C(L)/(1 - A(L))$ respectively.³⁵

Figure VIII compares the responses of the six key variables for the two types of oil discoveries separately. Overall, results are qualitatively similar to our baseline results pooling both types of discoveries. However, the current account and saving rate turn from negative to positive earlier for onshore than for offshore discoveries, as one would expect from the differential lag length. Moreover, the results suggest that offshore discoveries necessitate bigger and longer-lived investments than onshore discoveries, and output also increases later for offshore than for onshore discoveries. The trough of the IRF for the employment rate is earlier for onshore than for offshore discoveries. These results are consistent with the fact that the delay between the announcement of the discovery and the start of production is longer for offshore than onshore discoveries. Most of the responses are not statistically different, though. The responses of consumption are puzzling. Consumption rises earlier for offshore discoveries than for onshore discoveries.

One important implicit assumption embedded in our model is the absence of external borrowing constraints. If a country cannot borrow from the world market, however, then following a giant oil discovery saving would increase rather than decrease along with investment,

³⁴ Our measures of NPV for onshore and offshore discoveries also account for the differences in the time delay and production profiles for the two types of discoveries. More specifically, the time delays are specified as 4 and 6 years for onshore and offshore discoveries respectively, and the parameters for production profiles are listed in Table B.I. Our results still hold if we construct the NPV for two types of discoveries by using the same parameters in the baseline.

³⁵ This specification imposes the assumption that the autoregressive coefficients (and other control variables) are the same for the two types of discoveries. However, we achieve similar results if we run regression for two types of discoveries separately. This alternative specification allows different coefficients in all independent variables and it is valid in our case because the correlation between two types of discoveries is close to zero.

and the current account would equal zero since saving equals investment. To investigate these potential effects, we test whether the macroeconomic effects of giant oil discoveries are different across countries depending on their degree of financial openness. To capture the level of financial openness, we use a *de facto* measure—the ratio of total asset and liability to GDP, constructed from updated data on external wealth of nations from Lane and Milesi-Ferretti (2006).³⁶ We calculate the average of this index for each country, and take the median as the threshold to determine whether a country is financially open (above the threshold) or financially closed (below the threshold). We then re-estimate our main regressions separately for the two groups of countries. Figure IX shows that the responses of the current account, saving, investment and output are roughly similar for the two groups of countries. The employment rates decline less in financially open countries. The consumption responses have opposite signs for the financially closed and open cases after oil production starts, but the estimates are not precise.

Thus, the effects of a giant oil discovery on the current account and saving rate are not sensitive to the degree of financial openness. One possible explanation is that even those countries that typically face borrowing constraints may see their constraints relaxed after a giant oil discovery since they can use the giant oil fields as collateral when borrowing. Indeed, we do not find evidence that the macroeconomic responses to giant oil discoveries for Sub-Saharan African countries is any different than the overall response for other countries as shown in the Supplementary Appendix D.

Next we seek to shed some light on two elements we did not incorporate in our simple theoretical model – the potential government response and real exchange rates. We begin by

³⁶ An alternative measure of financial openness is the Chinn-Ito (2006) index, which is widely used in the international finance literature for capital account openness. Notice this index is a *de jure* measure for a country's degree of capital account openness as it is based on the binary dummy variables that codify the tabulation of restrictions on cross-border financial transactions reported in the IMF's *Annual Report on Exchange Arrangements and Exchange Restrictions*. Our results reported here is robust to this alternative index.

considering the behavior of government. Our data on investment and consumption aggregates private and public. Our theoretical model did not model public investment and public consumption separately, and instead (implicitly) considered government consumption and investment to be perfect substitutes for private consumption and investment. Pieschacón (2012) models the government spending and tax responses to oil price increases and shows that a differential response is important for understanding the differences in the effects of oil price increases on Norway versus Mexico. To determine the extent to which government is playing a role after giant oil discoveries, we investigate the responses of public vs. private investment and public vs. private consumption using data from the IMF (2013), as well as the response of the government spending-GDP ratio. As shown in Figure X, the private investment-GDP ratio increases, but the public investment-GDP ratio decreases. Thus, all of the increase in the aggregate investment-GDP ratio we saw in the earlier graph was due to the response of private agents. Private consumption (log levels) increases somewhat while public consumption does not respond to the initial discovery but then jumps up once the oil production begins. Thus, the government is an important part of the response of total consumption. The final graph of Figure X shows that while government spending on consumption and investment does rise, it rises less than the increase in GDP. Thus, it is unlikely that the government behavior is significantly changing the outcomes relative to our simple model.³⁷

In order to keep our theoretical model simple, we did not include a non-tradeable sector, and thus our model does not yield any prediction for the real exchange rate. Standard models

³⁷ Lei and Michaels (2014) also find that government spending rises after an oil discovery. In contrast to our results they find no effect on private investment. There are multiple differences in our implementation that could explain the different results. Our time period covers 1970 – 2012, whereas theirs covers 1946 – 2008; we measure our macroeconomic variables in country-specific currency units (which is appropriate for treating each country as an experiment) whereas they convert everything to U.S. dollars and are thus subject to exchange rate fluctuations; and they use a simple dummy variable for a discovery event whereas we use a richer method that takes into account both the size of the discovery and the size of the country's economy.

would suggest that the announcement of a giant oil discovery would lead to an immediate appreciation of the real exchange rate, operating through an increase in the relative price of non-tradeables (e.g. Eastwood and Venables, 1982).³⁸ Empirically, we found that the real exchange rate appreciated during the first five years following oil discoveries, either using a CPI-based or a GDP deflator-based measure, as shown in the Supplementary Appendix D. However, in both specifications, the point estimates were very imprecise so no response was significantly different from zero. Our results are somewhat consistent with the empirical literature on the so-called Dutch disease, which finds mixed evidence.³⁹ It should be noted however that one key difference between our estimates and results from the empirical literature on Dutch Disease is that the latter has so far been focused on the effect of contemporaneous windfall shocks as opposed to news shocks. Besides obvious measurement issues associated with real exchange rates, there are several potential explanations as to why we do not find evidence of significant real exchange rate appreciation following the announcement of a giant discovery. Perhaps, most importantly, is the fast moving nature of (real) exchange rates combined with the fact that we rely on annual frequency data for oil discovery announcements. Indeed, as argued by Ramey (2011) getting the timing right is essential to avoid bias in empirically investigating the effect of news shocks, and perhaps even more so when considering the exchange rate responses.

Finally, we also explore the effect of oil discoveries on country stock markets. Using data on financial structure updated from Beck et al. (2000), we tested whether giant oil discoveries

³⁸ Pieschacón (2012) analyzes the effects of oil price increases on Norway and Mexico, and finds mixed effects on the relative price of non-tradeables. In Mexico, the relative price of nontradeables increases, whereas in Norway it increases for a couple of quarters but then decreases. Wills (2013) explores the optimal response of monetary policy to a giant oil discovery in a standard small open economy model. He finds that the real exchange rate appreciate twice: first when forward-looking households and then the government increases their consumption.

³⁹ See Arezki and Ismail (2013) and references therein for a discussion of the empirical literature on the Dutch Disease. Arezki and al. (2013) discuss the difficulty of assessing the effect of windfalls on real exchange rates for oil exporters and find some evidence of asymmetrical Dutch Disease along boom-bust cycles.

affect stock market capitalization (as a percentage of GDP). The results, shown in in the Supplementary Appendix D, indicate that stock market indices respond earlier than the GDP response to a discovery, though the estimates are not very precise.

VII. ROBUSTNESS CHECKS

In this section, we discuss the results of extensive robustness checks for the benchmark specification. We explore different measures of giant oil discoveries, the effects of removing groups of countries, differences across oil and gas, and whether oil discoveries are predictable.

We first determine whether our main results are robust to alternative measures of NPV of giant oil discoveries. We first consider a simple dummy variable for an oil discovery event. This variable relies only on the timing of the oil discovery, not on the method used to construct a NPV. However, the disadvantage of this variable is that it ignores the size of the oil discovery and how it compares to the size of the economy, so it omits potentially important information. Figure XI presents the estimated responses of six key variables to an oil discovery event. Their patterns are similar to our baseline results except that the estimates of output are not precise enough to be significant different from zero.

We also explore alternative versions of our NPV measure of oil discoveries. Instead of using a more realistic projected oil production profile, we assume a constant production rate for 20 years after production starts. Moreover, alternative to the country-specific risk adjusted discount rate, we also use common discount factors of 10 percent. The impulse responses of main macro variables are virtually the same, implying that our results are not subject to the way we construct our measure of NPV of discoveries. (Figure D.I in the Supplementary Appendix.)

Second, we check whether removing groups of countries affects the estimates of the effect of giant oil discoveries on the macro variables. Recognizing the relative concentration of giant oil discoveries in certain Middle East and North African countries, we removed all countries belonging to this region from the sample and found that our main results still hold. Moreover, we explore the robustness of our main results to using countries which have experienced at least one discovery. Results are again robust, suggesting that the lack of comparison from countries where discoveries are absent is not an issue. We also tried using a broader income categorization based on whether countries are considered “high income” or “low income” countries according to a classification used by the World Bank. Our main results are robust to such sample split (Figure D.II in the Supplementary Appendix D).

Third, we check whether the impact of giant discoveries on the macroeconomy is different for oil (including condensate) and natural gas discoveries. We adopt the approach similar to the specification in Equation (13), this time including oil and gas discoveries. The responses of six key variables show similar patterns to our baseline results except the employment response is insignificant for natural gas discoveries (Figure D.III in the Supplementary Appendix D).

Fourth, we investigate formally the predictability of giant oil discoveries. It is possible that agents might have other hidden information that is unobserved by the econometrician and that could help predict discoveries. In this case, agents may adjust their behaviors including their saving and investment (and thus the current account) as a response to anticipated discoveries. For example, agents could borrow through current account deficits before oil discoveries, and then the Granger causality between discoveries and the current account would be reversed. To test this possibility, we used both a linear probability model and a logit model to test whether lagged

values of the current account, saving and investment have predictive power on the incidence of oil discoveries. Table IV presents the results of the test of the hypothesis that three annual lags of current accounts and investment do not have significant predictive power on the incidence of oil discoveries.⁴⁰ The conclusion that emerges from these tests is that there is no evidence of predictive power of lagged values of the current account and investment for giant oil discoveries.

We go further in exploring whether our results are robust to removing any potentially predictable giant oil discoveries. Because discoveries that followed others are more likely to be subject to the view that they are predictable, we test whether our main results still hold if we remove discoveries in the year immediately following a pre-existing discovery. The impulse responses for our key macro variables are virtually unchanged compared to our benchmark results, except the response of employment becomes less informative. We also tried retaining as news shocks only discoveries that happened without a prior history of discoveries in the last past three years, and the impulse responses are again qualitatively similar to our benchmark results (Figure D.IV in the Supplementary Appendix D). Finally, we also controlled for current and lagged exploration efforts using data from Global Energy Systems and our still find significant anticipation effects in the current account and saving rate (Figure D.V in the Supplementary Appendix D).⁴¹

One important feature for giant oil discoveries is the lag between the announcement of discoveries and the start of production. One could argue that the discovery of an oil field might induce a substitution effect between the newly found oil field and existing ones so much so that

⁴⁰ Saving is excluded due to obvious redundancy.

⁴¹ The responses of output, investment and employment are insignificant when we include exploration efforts. The decline in significance owes to the fact that the measure of exploration expenditures limits our sample to about one-fifth of the original sample. Conditional on this limited sample, the responses of the six key variables are not sensitive to the inclusion of the exploration effort and its lagged values.

future oil production and future output may remain unchanged, and the current output may increase. This substitution effect could potentially reduce the anticipation effect. However, due to the nature of oil extraction, adjustment costs tend to be high and further investment is necessary to increase the oil output in existing fields. Thus, in reality it might be difficult to speed up the oil pumping in existing fields in a short period. Moreover, conceptually the substitution effect might be more relevant for large oil exporters because they would tend to internalize the effect of their production on international oil and gas prices. We thus tested the robustness of our main results to removing the top ten largest oil or gas exporters in the world.⁴² Our main results are virtually unchanged (Figure D.VI in the Supplementary Appendix D).

Our results are also robust to using different dynamic specifications. In particular, our results are robust to using alternative econometric specifications that consist in higher order lags for the dependent variable such as $p = 2$, and different orders in lag independent variables, say $q = 8$, or 12. Results are indeed virtually unchanged to the case when we only use the first lag of the dependent variable. We also removed the controls for country-specific quadratic trend for output, employment rate and other level variables. The patterns of impulse responses are also similar (Figure D.VII in the Supplementary Appendix D).

Finally, we estimated alternative impulse responses using the Chang and Sakata (2007) “long autoregression” method, which is equivalent to the Jordà (2005) local projections method. This method has the advantage of imposing fewer dynamic restrictions. Unfortunately, it results in the loss of many years of data. Nevertheless, the results for the most part give the same

⁴² According to the U.S. Energy Information Administration, the five largest oil exporters in 2012 are Saudi Arabia, Russia, United Arab Emirates, Kuwait, and Nigeria, and the five largest gas exporters are Russia, Norway, Qatar, Canada, and Netherlands.

patterns as the baseline dynamic model for the relevant horizons. (Figure D.VIII in the Supplementary Appendix D).

VIII. Conclusion

In this paper, we have examined the effect of news about giant oil and gas discoveries on macroeconomic aggregates. We first presented a stylized two-sector open economy model to highlight the predicted effect of a sector-specific news shock. To identify the news shock in the data, we exploited the plausibly exogenous within-country variation in the timing and size of giant oil discoveries. We then estimated a dynamic panel distributed lag model over a sample covering 1970 to 2012 for about 180 countries. Results from the estimation provided evidence for a significant anticipation effect following the announcement of a giant oil or gas discovery. In particular, we found that immediately following the news shock, the current account and saving decreased, while investment increased. Only after the beginning of energy production did we find that the current account and saving rate increased along with GDP. In contrast, we found that employment fell when the news arrived and remained low for a number of years.

The canonical shock analyzed in the news-driven business cycle literature is a shock to aggregate TFP. Our oil news shock has special characteristics, though, related to the fact that resource sectors have lower labor shares and higher capital shares than the rest of the economy. Our stylized model showed, however, that these unique features did not alter the qualitative predictions for the responses of the current account, saving, investment, or GDP, relative to a more standard TFP news shock, though the magnitudes were markedly different in several cases. The response of consumption and hours did vary across types of shocks for some

parameterizations of preferences. Thus, the model sheds light on conditions under which we would expect our empirical results to hold more generally.

Our analysis of a specific, measureable news shock also highlights the need to think more seriously about the nature of the news shocks identified using time series and estimated DSGE model methods. While aggregate TFP shocks are useful for theoretical exercises, we would argue that it would be difficult to find actual instances of shocks that affect the aggregate production function and are fully anticipated in advance. We suspect that the potential of general purpose technologies that ultimately raise aggregate TFP is initially perceived by firms in only a few sectors.

Our oil news shock cannot answer the question of what fraction of output or current account fluctuations are driven by all news shocks. Nevertheless, our measure may be a useful tool for others who seek to answer this question. For example, our direct news measure may be employed as a specification test for times series identification methods. One could use our data set to test whether a given times series identification technique can properly extract the news shocks and produce estimated impulse responses that match those that we have estimated. If so, then those methods may be more credible when they are used to uncover news shocks in other data sets.

More broadly, our finding that news about future economic events leads to immediate responses is consistent with two other literatures. It is consistent with the fiscal literature results of Ramey (2011) and Barro and Redlick (2011), who study the effects of news about future government spending changes, and of Mertens and Ravn (2012), who study the effects of news about future tax changes. Like Mertens and Ravn (2012), we find that some key macroeconomic

aggregates oscillate from negative to positive after the news arrives. Our results are also consistent with the notion that news about future output can be an important source of fluctuations in the current account. Thus, our findings are supportive of the work of Engel and Rogers (2006) and Corsetti and Konstantinou (2012).

RESEARCH DEPARTMENT, INTERNATIONAL MONETARY FUND AND OXCARRE

UNIVERSITY OF CALIFORNIA, SAN DIEGO AND NBER

THE CHINESE UNIVERSITY OF HONG KONG

A. Theoretical Appendix

This appendix provides the first-order conditions for the two-sector model. In the equations below, λ is the Lagrange multiplier on the net foreign asset accumulation equation (5), the η 's are the multipliers on the capital accumulation equations in equation (4), and ζ is the multiplier on the resource accumulation equation (9). The first-order conditions for this economy are:

$$C_t : \quad (C_t - \psi N_t^\theta)^{-\sigma} = \lambda_t, \quad (\text{A.1})$$

$$N_{1t} : \quad (C_t - \psi N_t^\theta)^{-\sigma} \psi \theta N_t^{\theta-1} = \lambda_t \alpha_1 \frac{Y_{1,t}}{N_{1,t}} \quad (\text{A.2})$$

$$N_{2t} : \quad (C_t - \psi N_t^\theta)^{-\sigma} \psi \theta N_t^{\theta-1} = \alpha_2 \frac{Y_{2,t}}{N_{2,t}} (p_t \lambda_t - \zeta_t) \quad (\text{A.3})$$

$$B_t : \quad \lambda_t = \beta E_t (\lambda_{t+1} (1 + r_{t+1})) \quad (\text{A.4})$$

$$I_{ht} : \quad \lambda_t = \eta_{h,t} \left[1 - \frac{\phi}{2} \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right)^2 - \phi \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right) \frac{I_{h,t}}{I_{h,t-1}} \right] \quad (\text{A.5}),$$

$$h = 1, 2 \quad + \beta E_t \left[\eta_{h,t+1} \phi \left(\frac{I_{h,t+1}}{I_{h,t}} - 1 \right) \left(\frac{I_{h,t+1}}{I_{h,t}} \right)^2 \right] \quad (\text{A.6})$$

$$K_{1t}:\quad \eta_{1,t} = \beta E_t \left[\lambda_{t+1}(1 - \alpha_1) \frac{Y_{1,t+1}}{K_{1,t}} + \eta_{1,t+1}(1 - \delta) \right] \quad (\text{A.7})$$

$$K_{2t}:\quad \eta_{2,t} = \beta E_t \left[(\lambda_{t+1} p_{t+1} - \zeta_{t+1}) \alpha_k \frac{Y_{2,t+1}}{K_{2,t}} + \eta_{2,t+1}(1 - \delta) \right] \quad (\text{A.8})$$

$$R_t:\quad \zeta_t = \beta E_t \left[(1 - \alpha_2 - \alpha_k) \frac{Y_{2,t+1}}{R_t} (p_{t+1} \lambda_{t+1} - \zeta_{t+1}) + \zeta_{t+1} \right] \quad (\text{A.9})$$

There are 22 endogenous variables: $Y, Y_l, Y_2, C, I, I_l, I_2, K, K_l, K_2, N, N_l, N_2, B, r, R, CA, SA, \lambda, \eta_l, \eta_2, \zeta$.

The 22 equations are the 9 FOCs plus (2), (3), two equations in (4), (5), (6), two equations in (7), the four equations in (8), and (9).

B. Data Appendix

B.I. Approximation of production profile of giant oil discovery

Following the typical 4-6 year delay after the discovery announcement of a giant oil field, production is set to start.⁴³ To compute the net present value of giant oil field discoveries used in

⁴³ The net present value computation is based on 5-year delay between the discovery announcement and the start of oil production.

our econometric analysis, an approximation of the production profile following each discovery is needed considering the unavailability of comprehensive actual data on field level production.

The approximation relies on empirical production rates from limited available production data and the size of the reserve associated with the discovery. Given the assessment of the initial reserve, a production profile is derived using a plateau production level that is field size dependent, q_p , and a maximum depletion rate of remaining reserves, d_m . Production equals the plateau level as long as the depletion rate $d(t)$ is lower than the maximum rate. When exceeded, the depletion rate is held constant, causing exponential decline in production.

The depletion rate of remaining reserves at time t is defined as:

$$d(t) = \frac{q(t)}{RRR(t)},$$

where $q(t)$ is annual production rate and $RRR(t)$ is remaining reserves defined as:

$$RRR(t) = URR - Q(t),$$

where URR is ultimately recoverable reserves (equivalent to field size or initial reserve) and $Q(t)$ is cumulative production. Accordingly, the field production profile can be described as follows:

$$q(t) = \begin{cases} 0 & t \leq 5, \\ q_p & t > 5 \text{ \& } d(t) < d_m, \\ d_m RRR(t) & \text{otherwise.} \end{cases}$$

Approximations of q_p and d_m are derived using an oil field database maintained by the Global Energy Systems research group at Uppsala University (see e.g. Höök et al., 2014 and Robelius, 2007). Based on an empirical power law relationship between field size and plateau production, q_p , is specified as:

$$q_p = \alpha URR^\beta.$$

Furthermore, based on a general tendency of higher depletion and decline rates of smaller fields compared to larger fields, d_m is specified as

$$d_m = \gamma URR^\delta.$$

In Table A1 estimates of parameters α , β , γ and δ are presented for three categories of giant oil fields: all, OPEC and Non-OPEC fields.

Table B. I: Parameter estimates and R2-values for plateau production and maximum depletion rate approximation functions

Plateau production rate			
Fields	α	β	R ²
All giants	0.57	0.65	0.65
Onshore giants	0.19	0.78	0.78
Offshore giants	1.53	0.53	0.48
Maximum depletion rate			
Fields	γ	δ	R ²
All giants	0.64	-0.31	0.31
Onshore giants	0.39	-0.25	0.27
Offshore giants	0.7	-0.30	0.23

For simplicity, in the baseline setting a single set of parameter estimates corresponding to the one obtained from pooling all giant oil fields are used for the computation of production profiles. However, our baseline result is also robust if we construct the NPV by using different sets of parameters for onshore and offshore discoveries.

To illustrate the approximation, Figure A1 shows the production and depletion rate for a typical giant oil field of 500 million barrels of ultimately recoverable reserves. After the 5 years following the discovery announcement, production starts to rise and reaches a production plateau

that last 5 years before exponentially declining at constant depletion rate. It should be noted that the production profile would display a different duration for the plateau production depending on the size of the ultimately recoverable reserves. Formally, the duration of the production plateau, N , is given by the solution to the following equation:

$$d_m = \frac{q_p}{URR - q_p N}.$$

Figure B. I: Production and depletion rate for a field of 500 million barrels of ultimately recoverable reserves.

B.I. Data Description

Table B.II presents the description of the macroeconomic variables used and Table B.III presents summary statistics.

Table B. II: Data Definition and Sources

Variable	Definition and transformations	Source
Real GDP	Logarithm of GDP in constant prices, local current unit.	IMF (2013)
Current account as a percentage of GDP		IMF (2013)
Investment % of GDP as gross fixed capital formation a percentage of GDP	Gross fixed capital formation, both public and private	World Bank (2013)
Saving as a percentage of GDP	Constructed as the sum of current account and investment, to ensure consistency. The estimated dynamic effect of giant oil discoveries on saving is virtually unchanged if we instead use the saving data also provided by the World Bank (2013).	
Real consumption	Logarithm of final consumption expenditures in constant local current unit.	IMF (2013)
Employment	Defined as employment rate, defined as the employment to population ratio (in percentage), both male and female, age 15+. Available from 1991.	"emploare" from International Labor Organization website at (www.ilo.org/kilm).
Real private investment	Logarithm of (constant price, local currency unit) private gross fixed capital formation	IMF (2013)
Real public investment	Logarithm of (constant price, local currency) public gross fixed capital formation.	IMF (2013)
Real private consumption	Logarithm of (constant price, local currency unit) private consumption expenditures.	IMF (2013)
Real public consumption	Logarithm of (constant price, local currency) public consumption expenditures.	IMF (2013)
Real exchange rate	Real exchange rate based on CPI and GDP deflator	IMF (2013)
Stock market capitalization as a percentage of GDP		Beck et al. (2000)

Table B. III: Summary statistics of macro variables (1980-2012)

Variable	Years	Maximum number of countries	Obs	Min	Median	Max
Ln(GDP)	1980-2012	183	5504	-356.1	578.0	1477.8
CA/GDP	1980-2012	180	5408	-242.2	-3.2	106.8
Saving/GDP	1980-2011	171	4711	-202.9	18.1	107.2
Investment/GDP	1980-2011	178	4956	-2.4	21.1	113.6
Ln(final consumption)	1980-2012	162	4567	-82.3	590.7	1435.6
Employment rate	1991-2012	161	3541	28.9	57.6	88.1
Private investment/GDP	1980-2012	105	1959	-13.7	15.4	95.9
Public investment/GDP	1980-2012	105	1959	-4.5	4.3	47.1
Ln(final private consumption)	1980-2012	160	4511	-113.9	568.4	1426.5
Ln(final public consumption)	1980-2012	161	4523	-268.6	421.7	1223.3
Government spending/GDP	1980-2012	104	1926	1.8	21.5	87.6
Ln(real exchange rate based on CPI)	1980-2010	179	4745	285.1	463.4	1528.4
Ln(real exchange rate based on GDP deflator)	1980-2011	179	5075	225.6	460.7	2608.5
Stock market capitalization/GDP	1989-2011	112	1958	0.0	28.5	569.5

References

- Alexopoulos, Michelle.** 2011. "Read All About It!! What Happens Following a Technology Shock?" *American Economic Review*, 101(4), 1144-1179.
- Anderson, Soren T., Ryan Kellogg, and Stephen W. Salant.** 2014. "Hotelling Under Pressure" NBER Working Paper, July 2014.
- Arezki, Rabah and Ismail, Kareem.** 2013. "Boom–bust cycle, asymmetrical fiscal response and the Dutch disease," *Journal of Development Economics*, vol. 101(C), pages 256-267.
- Atalay, Enghin.** 2015. "How Important are Sectoral Shocks?" University of Wisconsin working paper, December.
- Barro, Robert J.** 2012. "Convergence and Modernization Revisited," NBER Working Papers 18295, National Bureau of Economic Research, Inc.
- Barro, Robert J., and Robert G. King.** 1984. "Time-Separable Preferences and Intertemporal Substitution Models of Business Cycles," *Quarterly Journal of Economics*, 99 (4), 817-839.
- Barro, Robert J., and Charles J. Redlick.** 2011. "Macroeconomic Effects from Government Purchases and Taxes," *Quarterly Journal of Economics*, 126 (1), 51-102.
- Barsky, Robert B., and Eric R. Sims.** 2011. "News Shocks and Business Cycles." *Journal of Monetary Economics*, 58(3): 273-289.
- Barsky, Robert B., and Eric R. Sims.** 2012. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence." *American Economic Review*, 102(4): 1343-77.
- Barsky, Robert B., Susanto Basu, and Keyoung Lee.** 2014. "Whither News Shocks?" NBER Working Paper 20666. November.
- Beaudry, Paul, and Franck Portier.** 2004. "An Exploration into Pigou's Theory of Cycles." *Journal of Monetary Economics* 51(6): 1183-1216.

- Beaudry Paul and Franck Portier.** 2006. "Stock Prices, News, and Economic Fluctuations." *American Economic Review*, 96(4): 1293-1307.
- Beaudry, Paul, and Franck Portier.** 2007. "When Can Changes in Expectations Cause Business Cycle Fluctuations in Neo-classical Settings?" *Journal of Economic Theory* 135(1): 458-477.
- Beaudry, Paul, and Bernd Lucke.** 2010. "Letting Different Views about Business Cycles Compete." *NBER Macroeconomics Annual 2009*, 24: 413-455. University of Chicago Press.
- Beaudry Paul and Franck Portier.** 2014. "News Driven Business Cycles: Insights and Challenges." *Journal of Economic Literature* 52 (4): 993-1074.
- Beck, Thorsten, Asli Demirgüç-Kunt and Ross Levine.** 2000. "A New Database on Financial Development and Structure," *World Bank Economic Review*, 14, 597-605.
- Bergin Paul R. and Steven M. Sheffrin.** 2000. "Interest Rate, Exchange Rates and Present Value Models of the Current Account." *The Economic Journal*, 110: 535-558.
- Bernanke, Ben.** 1985. "Adjustment Costs, Durables, and Aggregate consumption." *Journal of Monetary Economics*, Elsevier, 15(1), 41-68.
- Blanchard, Olivier J., Jean-Paul L’Huillier, and Guido Lorenzoni.** 2013. "News, Noise, and Fluctuations: An Empirical Exploration." *American Economic Review*, 103(7), 345-70.
- Blanchard, Olivier J., and Gian Maria Milesi-Ferretti.** 2011. "(Why) Should Current Account Imbalances Be Reduced?" International Monetary Fund (IMF) Staff Discussion Note SDN/11/03.
- Bohn, Henning and Robert T. Deacon.** 2000. "Ownership Risk, Investment, and the Use of Natural Resources." *American Economic Review*, 90(3), 526-49.
- Caselli, Francesco and James Feyrer.** 2007. "The Marginal Product of Capital." *Quarterly Journal of Economics*, 122(2), 535-568.

Chang, Pao-Li and Shinichi Sakata. 2007. "Estimation of Impulse Response Functions Using Long Autoregression," *Econometrics Journal*, 10, 453-469.

Chinn, Menzie D., and Eswar S. Prasad. 2003. "Medium-term Determinants of Current Accounts in Industrial and Developing Countries: An Empirical Exploration." *Journal of International Economics*, 59(1): 47-76.

Chinn, Menzie D., and Shang-Jin Wei. 2013. "A Faith-Based Initiative Meets the Evidence: Does a Flexible Exchange Rate Regime Really Facilitate Current Account Adjustment?" *Review of Economics and Statistics*, 95(1): 168-184.

Cochrane, John. 1994. "Shocks." *Carnegie-Rochester Conference Series on Public Policy*. 41(2), 295-364.

Correia, Isabel, João Neves, and Sergio Rebelo. 1995. "Business Cycles in a Small Open Economy." *European Economic Review*, 39 (6): 1089-1113.

Corsetti, Giancarlo and Panagiotis Konstantinou. 2012. "What Drives US Foreign Borrowing? Evidence on the External Adjustment to Transitory and Permanent Shocks." *American Economic Review*, 102(2), 1062-1092.

Cotet, Anca M. and Kevin K. Tsui. 2013. "Oil and Conflict: What Does the Cross Country Evidence Really Show?" *American Economic Journal: Macroeconomics*, 5(1), 49-80.

Den Haan, Wouter J., and Georg Kaltenbrunner. 2009. "Anticipated Growth and Business Cycles in Matching Models." *Journal of Monetary Economics*, 56(3): 309-327.

Den Haan, Wouter J., and Matija Lozej. 2011. "Pigou Cycles in Closed and Open Economies with Matching Frictions." *NBER International Seminar on Macroeconomics 2010*. Ed. Richard Clarida and Francesco Giavazzi, 193-233.

Driscoll, John C., and Aart C. Kraay. 1998. "Consistent covariance matrix estimation with spatially dependent panel data." *Review of economics and statistics*, 80.4: 549-560.

Eastwood, Robert K and Anthony J. Venables. 1982. "The Macroeconomic Implications of a Resource Discovery in an Open Economy," *Economic Journal*, Royal Economic Society, Vol. 92(366), pages 285-99, June.

Engel, Charles and John H. Rogers. 2006. "The U.S. Current Account Deficit and the Expected Share of World Output." *Journal of Monetary Economics*, 53(5), 1063-1093.

Fisher, Jonas D.M. 2010. "Comment on 'Letting Different Views of the Business Cycle Compete'" *NBER Macroeconomics Annual* 2009, Vol. 4, 457-474.

Ghosh, Atish R and Jonathan D. Ostry. 1995. "The Current Account in Developing Countries: A Perspective from the Consumption-Smoothing Approach," *World Bank Economic Review*, World Bank Group, vol. 9(2), pages 305-33, May.

Glick, Reuven and Kenneth Rogoff. 1995. "Global Versus Country-specific Productivity Shocks and the Current Account." *Journal of Monetary Economics*, Elsevier, 35(1), 159-192.

Greenwood, Jeremy, Zvi Hercowitz, and Gregory H. Huffman. 1988. "Investment, Capital Utilization, and the Real Business Cycle." *American Economic Review*, 78(3), 402-417.

Gross, Isaac and James Hansen. 2013. "Reserves of Natural Resources in a Small Open Economy." Reserve Bank of Australia Research Discussion Paper, RDP 2013-14.

Hall, Robert Ernest. 1978. "Stochastic Implications of the Life Cycle-Permanent Income Hypothesis: Theory and Evidence." *Journal of Political Economy*, 86(6): 971-87.

Hamilton, James D. 2009. "Understanding Crude Oil Prices," *The Energy Journal*, International Association for Energy Economics, International Association for Energy Economics, vol. 0(Number 2), pages 179-206.

Hamilton, Kirk and Giles Atkinson. 2013. "Resource discoveries, learning, and national income accounting." *The World Bank Policy Research Working Paper Series* 6505.

Hoffmann, Mathias, Michael U. Krause, and Thomas Laubach. 2013. “The Expectations-driven US Current Account”. No. 10/2013. Discussion Paper, Deutsche Bundesbank, 2013.

Höök, M., Davidsson, S., Johansson, S., Tang, X.. 2014. “Decline and depletion rates of oil production: a comprehensive investigation.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372, 20120448.

Horn, (Mike) Myron K. 2014. “Giant Oil and Gas Fields of the World”.

<http://www.datapages.com/AssociatedWebsites/GISOpenFiles/HornGiantFields.aspx>.

International Country Risk Guide. 2015. *Political Risk Services*. East Syracuse, New York.

International Monetary Fund. 2013a, *World Economic Outlook (WEO)*.

International Monetary Fund. 2013b, *United States: Article IV Consultation*, IMF Country Report No. 13/237 (Washington).

Jaimovich, Nir and Sergio Rebelo. 2008. “News and Business Cycles in Open Economies,” *Journal of Money, Credit and Banking*, 40 (8), 1699–1711.

Jaimovich, Nir and Sergio Rebelo. 2009. “Can News About the Future Drive the Business Cycle?” *The American Economic Review*, 99 (4), 1097–1118.

Jordà, Òscar. 2005. “Estimation and Inference of Impulse Responses by Local Projections.” *American Economic Review* 95 (1): 161-82

Keynes, John Maynard. 1936. *The General Theory of Employment, Interest and Money*. London: Macmillan.

King, Robert G., Charles I. Plosser, and Sergio T. Rebelo. 1988. “Production, Growth, and Business Cycles: I. The Basic Neoclassical Model,” *Journal of Monetary Economics*, 21(2-3), 195-232.

Krusell, Per and Alisdair McKay. 2010. "News Shocks and Business Cycles." *FRB Richmond Economic Quarterly*, 96(4): 373-397.

Kueng, Lorenz. 2014. "Tax News: Identifying Tax Expectations from Municipal Bonds with an Application to Household Consumption." NBER Working Paper 20437, August.

Kurmann, André and Elmar Mertens. 2014. "Stock Prices, News, and Economic Fluctuations: Comment," *American Economic Review*, 104(4): 1439-1450.

Leduc, Sylvain and Keith Sill, 2013. "Expectations and Economic Fluctuations: An Analysis Using Survey Data". *Review of Economics and Statistics* 95(4), October 2013, 1352-1367".

Leeper, Eric, Alexander Richter and Shu-Chun Susan Yang. 2012. "Quantitative Effects of Fiscal Foresight," *American Economic Journal: Economic Policy* 4 (2): 1-27.

Leeper, Eric, Todd B. Walker, and Shu-Chun Yang. 2013. "Fiscal Foresight and Information Flows," *Econometrica*, 81(3): 1115-1145.

Lei, Yu-Hsiang and Guy Michaels. 2011. "Do Giant Oilfield Discoveries Fuel Internal Armed Conflicts?" C.E.P.R. Discussion Papers 8620.

Lane, Philip R. and Gian Maria Milesi-Ferretti, 2007. "The external wealth of nations mark II: Revised and extended estimates of foreign assets and liabilities, 1970–2004", *Journal of International Economics* 73, November, 223-250.

Lucca, David. 2007. "Resuscitating Time-To-Build." Manuscript.

<http://ideas.repec.org/p/red/sed007/909.html> .

Mertens, Karel and Morten O. Ravn. 2012. "Empirical Evidence on the Aggregate Effects of Anticipated and Unanticipated US Tax Policy Shocks," *American Economic Journal: Economic Policy*, 4(2): 145-81.

Obstfeld, Maurice. 1982. "Aggregate Spending and the Terms of Trade: Is There a Laursen-Metzler Effect?" *Quarterly Journal of Economics*, 97 (2): 251-270.

- Obstfeld, Maurice, and Kenneth Rogoff.** 1995. "The Intertemporal Approach to the Current Account." In *Handbook of International Economics*, Vol. 3, edited by Grossman, G., Rogoff, K., 1731-1799. Amsterdam: Elsevier.
- Obstfeld, Maurice.** 2012. "Does the Current Account Still Matter?" *American Economic Review*, 102(3): 1-23.
- Persson, Torsten and Lars E.O. Svensson.** 1985. "Current Account Dynamics and the Terms of Trade: Harberger-Laursen-Metzler Two Generations Later." *Journal of Political Economy*, 93(1): 43-65.
- Pieschacón, Anamaría.** 2012. "The Value of Fiscal Discipline for Oil-Exporting Countries." *Journal of Monetary Economics*, 59 (3): 250-268.
- Pigou, Arthur C.** 1927. *Industrial fluctuations*. London: Macmillan.
- Pindyck, Robert S.** 1978. "The Optimal Exploration and Production of Nonrenewable Resources." *Journal of Political Economy*, 86(5): 841-862
- Ramey, Valerie A.** 2011. "Identifying Government Spending Shocks: It's all in the Timing." *Quarterly Journal of Economics* 126.1: 1-50.
- Ramey, Valerie A.** 2015. "Macroeconomic Shocks and Their Propagation." Prepared for the *Handbook of Macroeconomics*, ed. John B. Taylor and Harald Uhlig.
- Robelius, Fredrik.** 2007. "Giant Oil Fields - The Highway to Oil : Giant Oil Fields and Their Importance for Future Oil Production." *Doctoral dissertation*. Uppsala University, Uppsala.
- Sachs, Jeffrey.** 1981. "The Current Account and Macroeconomic Adjustment in the 1970s." *Brookings Papers on Economic Activity*: 201-258.
- Schmitt-Grohé, Stephanie, and Martin Uribe.** 2003. "Closing Small Open Economy Models." *Journal of international Economics*, 61(1): 163-185.

Schmitt-Grohé, Stephanie, and Martin Uribe. 2012. "What's News in Business Cycles." *Econometrica*, 80(6): 2733-2764.

Sheffrin, S. and Woo, W. T. 1990. "Present Value Tests of an Intertemporal Model of the Current Account." *Journal of International Economics*, 29(3-4): 237-53.

Tsui, Kevin K. 2011. "More Oil, Less Democracy: Evidence from Worldwide Crude Oil Discoveries." *Economic Journal*, Royal Economic Society, 121(551): 89-115.

Uribe, Martín and Stephanie Schmitt-Grohé. 2014. *Open Economy Macroeconomics*. Manuscript, September 21, 2014.

Wills, Samuel. 2013. "Optimal Monetary Responses to Oil Discoveries," Economics Series Working Papers OxCarre Research Paper 12, University of Oxford.

World Bank Group. 2013, *World Development Indicators* (WDI).

Table I: Baseline Calibrated Parameters For Two-Sector Model

Parameter	Name	Value
β	Discount factor	0.943
ψ	Governs disutility of labor, set so steady-state labor is 20%.	0.4623
θ	Exponent on labor in the utility function, governing intertemporal substitution.	1.2
σ	Governs intertemporal substitution of the consumption-hours bundle	1
φ	Investment adjustment cost parameter	0.1
δ	Capital depreciation	0.1
α_l	Labor share in non-oil sector	0.64
α_2	Labor share in oil sector	0.13
α_k	Capital share in oil sector	0.49
χ	Elasticity of interest rate with respect to net foreign assets	0.0001
$\bar{B}$	Parameter in interest rate function; set so that the steady-state tb/y = 0.04	-22.285
p	Relative price of oil	1
$\bar{R}$	Steady-state flow of oil reserve inflow, set so that steady-state oil sector is around 6% of GDP	2
A_i	TFP in Sector i, i = 1,2	1

Table II: The Spatial And Temporal Distribution Of Giant Oil Discoveries (1970-2012)

Region	1970s	1980s	1990s	2000s	2010s	Total
Sub-Saharan Africa	5	6	9	9	9	38
Asia	17	14	20	23	0	74
Commonwealth of Independent States and Mongolia	22	12	4	10	3	51
Europe (include Central and Eastern Europe)	17	5	7	3	5	37
Middle East and North Africa	36	15	23	18	5	97
Western Hemisphere	20	15	16	21	2	74
World total	117	67	79	84	24	371

Note: the figures in the table reflect the total number of “discovery events” for a given decade and a given region. A discovery event is a dummy variable takes a value of 1 if during a given year at least one discovery of either a giant oil or gas field was made in any given country, and zero otherwise. The data are from Mike Horn and the country grouping is from the International Monetary Fund.

Table III: Hypothesis Tests On Responses To An Oil News Shock

Variable	Theoretical Prediction for Alternative Hypothesis H_1	Hypothesis Test	p-value
Current Account/GDP	Negative response at horizons 0-4	$H_0: \sum_{h=0}^4 b_h \geq 0$ vs. $H_1: \sum_{h=0}^4 b_h < 0$	0.00
	Positive response starting at horizon 5	$H_0: \sum_{h=5}^{11} b_h \leq 0$ vs. $H_1: \sum_{h=5}^{11} b_h > 0$	0.00
Saving/GDP	Negative response at horizons 0-4	$H_0: \sum_{h=0}^4 b_h \geq 0$ vs. $H_1: \sum_{h=0}^4 b_h < 0$	0.05
	Positive response starting at horizon 5	$H_0: \sum_{h=5}^{11} b_h \leq 0$ vs. $H_1: \sum_{h=5}^{11} b_h > 0$	0.08
Investment/GDP	Positive response at horizons 0-4	$H_0: \sum_{h=0}^4 b_h \leq 0$ vs. $H_1: \sum_{h=0}^4 b_h > 0$	0.02
	Negative or zero response starting at horizon 5	$H_0: \sum_{h=5}^{11} b_h > 0$ vs. $H_1: \sum_{h=5}^{11} b_h \leq 0$	0.53
GDP	Positive or zero response at horizons 0-4	$H_0: \sum_{h=0}^4 b_h < 0$ vs. $H_1: \sum_{h=0}^4 b_h \geq 0$	0.71
	Positive response starting at horizon 5	$H_0: \sum_{h=5}^{11} b_h \leq 0$ vs. $H_1: \sum_{h=5}^{11} b_h > 0$	0.00
Consumption	Positive response at horizons 0-11	$H_0: \sum_{h=0}^{11} b_h \leq 0$ vs. $H_1: \sum_{h=0}^{11} b_h > 0$	0.10
Employment- Population ratio	Negative response at horizons 0-11	$H_0: \sum_{h=0}^{11} b_h \geq 0$ vs. $H_1: \sum_{h=0}^{11} b_h < 0$	0.03

Notes: b_h denotes the estimated impulse response at horizon h . P-values were obtained from Delta method. The hypotheses are constructed based on the theory presented in Section II.

Table IV: Test For Predictability Of Giant Oil Discoveries

	1980-2012	1970-2012	1980-2012	1970-2012
Linear	0.62	0.47	0.68	0.60
Logit	0.74	0.54	0.72	0.63
Previous discoveries in 10 years	No	No	Yes	Yes
Country and year fixed effects	Yes	Yes	Yes	Yes

Note: The table reports the outcome of tests of non-predictability of the giant oil discovery event dated by their announcements. The row denoted "linear" contain the p-value of F test of the hypothesis that three lags of current account and investment have no predictive power for the oil discovery event on the basis of panel linear probability model with fixed effects. The row denoted "logit" report the p-value for the likelihood ratio test based on panel logit model. All tests are specified as $H_0: \alpha_x = 0$ against $H_1: \alpha_x \neq 0$.

**Figure I: Typical Oil field Investment And Production Patterns:
Examples From Two Norwegian Oil Fields**

Note: The investment data is based on nominal data divided by the GDP deflator. The oil production data is in one thousand barrels per day. The data is from the Norwegian Petroleum Directorate (NPD), <http://www.npd.no/en/>.

Figure II: Effect Of Oil Discovery News
Baseline Model, 5 Year Lag Of News

Note: The vertical axis shows percentage changes. The solid blue line is the baseline model with GHH preferences. The green dashed line is the model with KPR preferences. The shock is normalized so that the present value of the rise in oil revenue is equal to 1% of initial GDP in the baseline model.

Figure III: Effect of Sectoral TFP News
Two-Sector Model, 5 Year Lag of News

Note: The vertical axis shows percentage changes. The solid blue line is the baseline oil discovery model with GHH preferences. The red dashed line is the two-sector model with identical factor shares and exogenous TFP. Each sector has a labor share of 58% and a capital share of 32%. Exogenous TFP takes the values of the baseline endogenous reserves process with an exponent of 0.38.

Figure IV: The Size Distribution Of Giant Oil Discoveries: 1970-2012

Note: The figure presents the logarithm of million barrels of ultimately recoverable oil equivalent for giant discoveries in our sample.

Figure V: The Distribution Of Net Present Value Of Giant Oil Discoveries: 1970-2012

Figure VI: The Impact of Giant Oil Discoveries on the Current Account, Saving and Investment

Note: The figure presents the impulse response of an oil discovery with NPV equal to 1% of GDP. The line with circles indicates point estimates, and grey areas are 90 percent and 68 confidence intervals. The vertical axis shows percentage changes.

Figure VII: The Impact of Giant Oil Discoveries on GDP, Consumption and Employment

Note: The figure presents the impulse response of an oil discovery with NPV equal to 1% of GDP. The line with circles indicates point estimates, and grey areas are 90 percent and 68 confidence intervals. The vertical axis shows percentage changes.

Figure VIII: Onshore Discoveries v.s. Offshore Discoveries

Note: The figure presents the impulse responses of an oil discovery with NPV equal to 1% of GDP for onshore and offshore discoveries respectively. The vertical axis shows percentage changes.

Figure IX: Financial Openness

Note: The figure presents the impulse responses of an oil discovery with NPV equal to 1% of GDP for countries with high and low financial openness respectively. The vertical axis shows percentage changes.

Figure X: Private Versus Public Responses

Note: The figure presents the impulse response of discovery on logarithm of real private and government consumption (in local currency units) with control for country specific quadratic trend, as well as the response of the government spending as a percentage of GDP. The line with circles indicates point estimates, and grey areas are 90 percent and 68 confidence intervals respectively.

Figure XI: The Impact Of Giant Oil Discovery Event On Macro Variables

Note: The figure presents the impulse response of an oil discovery event. The line with circles indicates point estimates, and grey areas are 90 percent and 68 confidence intervals. The vertical axis shows percentage changes.