

News Shocks in Open Economies: Evidence from Giant Oil Discoveries

Rabah Arezki, Valerie A. Ramey and Liugang Sheng*

January 5, 2015

Abstract

This paper explores the effect of news shocks on the current account and other macroeconomic variables using worldwide giant oil 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 present a two-sector small open economy model in order to predict the responses of macroeconomic aggregates to news of an oil discovery. We then estimate the effects of giant oil discoveries on a large panel of countries. Our empirical estimates are consistent with the predictions of the model. After an oil 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 for a sustained period of time.

JEL Classification: E00, F3, F4.

Keywords: news shocks, current account, saving, investment, employment, oil, discovery

* Research Department, International Monetary Fund 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; lsheng@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, Kirk Hamilton, Thomas Helbling, Nir Jaimovich, Ayhan Kose, Jean-Pierre Laffargue, Prakash Loungani, Akito Matsumoto, Gian Maria Milesi-Ferretti, 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 participants at the Aix Marseille School of Economics, CERDI, the International Monetary Fund, The Hong Kong University of Science & Technology, Norges Bank, Peking University, Rice University, Tsinghua University, the Martyfest Conference, Cambridge University, UCL, University of Oxford, University of Siegen, University of Virginia, University of Warwick, CESifo, and the ECB for helpful comments. We thank Hites Ahir and Daniel Greenwood for excellent research assistance. Liugang thanks for the financial support from Research Grants Council of Hong Kong. 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. Most of the existing studies rely on structural vector autoregressive models (VAR) or on structural estimation of standard dynamic stochastic general equilibrium models. 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. Obstfeld and Rogoff (1995), 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.¹

This paper provides empirical evidence of the effect of news shocks on the current account, saving, investment, GDP, consumption, and employment using plausibly exogenous variation in the timing of worldwide giant oil discoveries as a directly observable measure of news shocks about higher future output. The delay between a discovery and production is on average 4 to 6 years.² 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

¹ 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.

² Hereafter we refer to discoveries of giant oil (including condensate) and gas fields as simply giant oil discoveries. A giant oil 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. Ultimately recoverable reserve refers to the amount that is technically recoverable given existing technology.

they might differ from the standard aggregate 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. The NPV is calculated as the sum of gross oil revenue derived from an approximated oil production profile discounted by country-specific discounting factors, and valued at the oil price prevailing 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 130 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 Obsfeld 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.

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 shocks. This feature is unique in the sense that other plausibly exogenous shocks used in other strands of literature and based on directly observable measures 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.³

The unique timing characteristic of oil discoveries provides a methodological contribution to the identification problem of news shocks and the associated anticipation effects in the recent literature on “expectation driven” business cycles and current account movements. 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. While our approach cannot answer the question of what fraction of output or current account fluctuations are driven by news shocks, our results can help answer this question indirectly because our response estimates can be used to test which of the competing VAR or DSGE identification schemes yield responses that are consistent with our estimates.

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 solely within-country variations in the timing of the giant oil discoveries allay 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

³ 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 Michael (2011) study the relationship between giant oil discoveries and civil conflicts. To the extent of our knowledge, we are the first to exploit the timing of giant oil discoveries to test the predictions of a macro model with news.

decline after the news arrives and remain below normal for over 10 years. We also explore empirically the respective roles of the private and public sectors in explaining our main results. We find that the private investment-GDP ratio rises robustly, but that the public investment-GDP ratio falls. In contrast, the total consumption response is driven by both an increase in private and public consumption, though the estimates are very imprecise.

Our results are robust to a wide array of checks. First, we use two alternative measures of oil discoveries: the NPV with constant production profile and the NPV with common discount factors across countries. The impulse responses are virtually the same, which implies that our results are not subject to the way we construct the NPV. 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 has implications for the news-driven business cycles hypothesis. Barro and King (1984) and Cochrane (1994) pointed out that in a standard RBC model news about future TFP could not be a driver of business cycles 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. More recently, however, Barsky and Sims (2011) and Barsky, Basu and Lee (2014) have used

time series techniques to identify 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. Thus, the empirical work based on time series identification is in flux.⁴ To our knowledge, we are the first to employ direct measures of news about future output. As we will demonstrate later, the results are not consistent with news as a driver of business cycle fluctuations.

This paper also provides direct evidence for the classic intertemporal trade model of the current account (e.g. Obstfeld and Rogoff (1995)). That model 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 present value model is difficult, though, because there are few direct measures of expectations about future output or productivity. Typically, time series methods are used to identify effects, but often the results 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. THEORETICAL EFFECTS OF AN OIL DISCOVERY

In this section, we study the theoretical predictions for the effects of oil discoveries on macroeconomic variables in a small open economy. Before examining the effects in a full two-

⁴ See Beaudry and Portier (2014) and Krusell and McKay (2010) for recent surveys of the literature on news shocks and business cycle fluctuations.

sector dynamic production economy, it is useful to review the intuition for the effects of a news shock versus a contemporaneous shock on consumption, saving and the current account in an endowment economy.

II.A News Shocks in an Endowment Economy

Consider a small open economy populated by a large number of infinitely lived households. There is one tradeable good and in each period households receive an exogenous endowment. Households discount future utility by the factor β and have the ability to borrow and lend at the exogenous world interest rate r . Thus, the representative household's maximization problem is as follows:

$$\max E_0 \sum_{t=0}^{\infty} \beta^t U(C_t) ,$$

$$s. t. CA_t \equiv B_t - B_{t-1} = rB_{t-1} + Y_t - C_t,$$

where C_t is representative household's consumption, CA_t is the current account, B_t is household's holding of the riskless bond at the end of period t , and Y_t denotes the exogenous and stochastic endowment of goods. Assuming for simplicity that the utility function, U , is linear-quadratic as in Hall (1978), that the discount factor $\beta = \frac{1}{1+r}$, and imposing a No-Ponzi scheme, the Euler condition yields the classical condition on consumption under the permanent income hypothesis, i.e., $C_t = E_t(C_{t+1})$. To introduce news, we assume that output Y_t , follows an exogenous stochastic process that is an auto-regressive process of order one ($AR(1)$) with coefficient, ρ , and that households receive news about future output. More specifically, we define $y_t = Y_t - \bar{Y}$, where \bar{Y} is the long run steady state, so that y_t rewrites as follows:

$$y_t = \rho y_{t-1} + \epsilon_t , \quad 0 \leq \rho \leq 1$$

where ϵ_t is the error term.

Following Schmitt-Grohe and Uribe (2012), suppose that ϵ_t consists of two components:

$$\epsilon_t = \eta_t^0 + \eta_{t-1}^1 ,$$

where η_t^0 is an unanticipated contemporaneous shock, and η_{t-1}^1 is the one-period ahead news shock, which materializes in period t , but that households learn about in period $t - 1$. For $j = 0,1$, both η_t^j , are of mean zero and are uncorrelated across time and across anticipation horizons.

Consider now the differentiated effects of news shocks and contemporaneous shocks on consumption and the current account. The optimal change in consumption is given as follows:

$$\Delta C_t = \frac{r}{1+r-\rho} \left(\eta_t^0 + \frac{\eta_t^1}{1+r} \right)$$

Both the contemporaneous shock and the news shock in t change agents' expectations of future output, inducing consumers to update their optimal consumption path. However, the materialized news shock, η_{t-1}^1 , disappears in the equation because households have learnt about the shock in $t - 1$, and there is no information updating although it materializes in t . In other words, consumers adjust their consumption upon the time they receive the news, rather than at the time the output materializes. However, the observational equivalence between contemporaneous shocks and news shocks in t presents a challenge for empirical testing as one needs to identify news shocks that are orthogonal to contemporaneous shocks prior to testing their effects on consumption.

We now turn to exploring the effect of news shocks on the current account. The current account is given by $CA_t = B_t - B_{t-1}$ and is equal to saving in this endowment economy with no investment. One can show that the change in the current account is given by:

$$\Delta CA_t = \frac{1}{1+r-\rho} \left[-(1-\rho)^2 y_{t-1} + (1-\rho) \eta_t^0 - \frac{r}{1+r} \eta_t^1 + \left(1-\rho + \frac{r}{1+r} \right) \eta_{t-1}^1 \right]$$

It clearly shows that contemporaneous shocks (η_t^0) and news shocks in period t (η_t^1) have opposite effects on the current account (and saving). A positive and temporary contemporaneous shock causes an increase in the current account, while the anticipation effect of a positive news shock causes the current account to decline. In contrast, last period's news shock that is realized in this period, η_{t-1}^1 , causes the current account to rise. In sum, as households receive the news

shock, the current account should decrease first due to the anticipation effect and then increase as the output materializes.

This simple endowment economy model is useful to review the intuition for the effects of a news shock versus an unanticipated change on the current account through the saving channel. However, the model does not consider endogenous responses of investment and labor, as well as some special features of oil sector. Next we present a standard two-sector open economy model with oil discovery news shocks, and analyze the effects of news on relevant macroeconomic variables.

II.B A Two-Sector Small Open Economy Model with a News Shock

Our two-sector model is an extension of the one-sector model of Jaimovich and Rebelo (2008), who study the effect 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.

II.B.1 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).

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

where

$$X_t = C_t^\gamma X_{t-1}^{1-\gamma}, \quad (2)$$

and $0 < \beta < 1, \theta > 1, \psi > 1$, and $\sigma > 0$. These preferences were introduced by Jaimovich and Rebelo (2008, 2009) and are convenient because they nest both King, Plosser, Rebelo (1988) (KPR) preferences ($\gamma = 1$) and Greenwood, Hercowitz, and Huffman (1988) (GHH)

preferences ($\gamma = 0$). GHH preferences shut down the wealth effect on labor supply and are often used in open economy models (e.g. Correia, Neves, and Rebelo (1995), Uribe and Schmitt-Grohe (2014)). The presence of X_t makes preferences non-time-separable in consumption and hours worked. Agents internalize the dynamics of X_t in their maximization problem.

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:

$$Y_{1,t} = A_{1,t} N_{1,t}^{\alpha_1} K_{1,t-1}^{1-\alpha_1}, \quad (3)$$

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 a resource factor (R) with a Cobb-Douglas production:

$$Y_{2,t} = A_{2,t} N_{2,t}^{\alpha_2} K_{2,t-1}^{\alpha_k} R_{t-1}^{1-\alpha_2-\alpha_k}, \quad (4)$$

where $0 < \alpha_1, \alpha_2, \alpha_k < 1$, and R_{t-1} is the resource factor at the *end* of period $t - 1$ (or beginning of period t).

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:

$$K_{h,t} = I_{h,t} \left[1 - \frac{\phi_h}{2} \left(\frac{I_{h,t}}{I_{h,t-1}} - 1 \right)^2 \right] + (1 - \delta_h) K_{h,t-1}, \quad (5)$$

$$h = 1, 2$$

We assume that the parameter $\phi_h > 0$. The functional form implies that there are no adjustment costs in the steady state.

As in Jaimovich and Rebelo (2008, 2009), we also introduce sector-specific adjustment costs on labor, which are subtracted from the economy's flow budget constraint. 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.⁵ Thus, the flow budget constraint is given as follows:

$$B_t = (1 + r_t)B_{t-1} + (Y_{1,t} + p_t Y_{2,t}) - \left\{ C_t + I_{1,t} + I_{2,t} + \sum_{h=1}^2 N_{h,t} \left[\frac{\Psi_h}{2} \left(\frac{N_{h,t}}{N_{h,t-1}} - 1 \right)^2 \right] \right\} \quad (6)$$

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, p_t is the relative price of oil, and the final terms capture the labor adjustment costs. We assume that the relative price of oil is determined by the world market. To induce stationarity of foreign bond holdings, we follow the external debt-elastic interest rate proposed by Schmitt-Grohe and Uribe (2003),

$$r_t = r_w + \chi[\exp(\bar{B} - B_{t-1}) - 1] \quad (7)$$

where r_w is the world interest rate, and $\chi > 0$ is the interest rate debt elasticity. The second part 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 trade balance is defined as

$$TB_t = (Y_{1,t} + p_t Y_{2,t}) - \left(C_t + I_{1,t} + I_{2,t} + \sum_{h=1}^2 N_{h,t} \left[\frac{\Psi_h}{2} \left(\frac{N_{h,t}}{N_{h,t-1}} - 1 \right)^2 \right] \right), \quad (8)$$

and the current account is defined as

$$CA_t = B_t - B_{t-1}, \quad (9)$$

and thus $CA_t = TB_t + r_t B_{t-1}$.

Saving is defined as

$$SA_t = CA_t + I_{1t} + I_{2t}, \quad (10)$$

⁵ 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.

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

$$Y_t = Y_{1,t} + p_t Y_{2,t}, \quad K_t = K_{1,t} + K_{2,t}, \quad I_t = I_{1,t} + I_{2,t}, \quad N_t = N_{1,t} + N_{2,t} \quad (11)$$

In the typical model, the news shock is represented as news that aggregate TFP will increase at some later date. Since the economy starts in a steady-state with positive capital and labor inputs, output will rise when aggregate TFP rises at the future date even if there is no new investment or change in labor supply. This is not the case with giant oil discoveries. Although the reserves are known to exist as soon as the discovery is made, no oil can be extracted until capital and labor have been reallocated. Moreover, most of the investment in capital in the new oil field must occur 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 Heidrun oilfield, production rises rapidly before gradually declining. At the Draugen oilfield, production rises more gradually and declines at a faster rate after the peak.

We capture these essential features with the following AR(2) specification for the resource factor:

$$\ln R_t = (1 - \rho_1 - \rho_2) \ln R + \rho_1 \ln R_{t-1} + \rho_2 \ln R_{t-2} + \varepsilon_{t-j}^j \quad (12)$$

where $\rho_1 + \rho_2 < 1$ and $j \geq 0$. We interpret R_t as a “resource factor” rather than actual reserves. The lag on ε_{t-j}^j incorporates the key feature that the resources are not immediately available when the news arrives. Combined with the sector-specific adjustment costs on investment, this model mimics the key features of time-to-build.⁶ A more realistic process would take into account the fact that investment in exploration increases the probability of discovery and that production leads to resource depletion. Because optimal resource extraction is not central to the current application, we model this feature with an exogenous depletion rate governed by the ps.

⁶ Lucca (2007) shows that adjustment costs in the change in investment mimic the effects of time-to-build.

The first-order conditions and steady-state equations for the representative agent are presented in the theoretical appendix. We calibrate the model to match the annual data used in our empirical analysis. Our baseline calibration is summarized in Table I. Most of the parameters are similar to those in Jaimovich and Rebelo (2008), with relevant ones converted to an annual basis to match our data. For example, the value of γ is set so that preferences are very close to GHH preferences and χ is set so that the elasticity of interest rates is very low.⁷ We set the world interest rate to 10 percent to match the average of the risk adjusted interest rate in 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.⁸

II.B.2 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}^5 in equation (12). We first compare the effects of a discovery shock with delay in the available resource to a discovery that results in immediately available additional oil resources (i.e. we use ε_t^0 in equation (12)).⁹ Both shocks arrive during period 0, and are normalized so that the peak response of GDP is the same as we will see later in the empirical estimates.

Figure II shows the differential effects of news shocks about future resources versus contemporaneous shocks about current resources. The responses for the case of discovery with a five-year delay in availability of oil resources for production are shown by the blue solid lines and the case of discovery with immediate availability is shown by the green dashed line. The upper left graph of Figure II shows that news about future resources leads the current account to turn sharply negative for five years before becoming positive. The ratio of saving to GDP is also negative for several years before becoming very positive. The declines in the current account

⁷ We choose a value for χ half way between Jaimovich-Rebelo's and Schmitt-Grohe-Urbe's values.

⁸ 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.

⁹ Note that this latter case is similar to an unanticipated rise in the relative price of oil that is expected to be persistent but not permanent.

and the ratio of saving to GDP are due to the anticipation effect of future resource and output increases. In contrast, when the resources are immediately available the current account-GDP ratio turns only slightly negative for couple of years and the saving-GDP ratio is always positive. The saving-GDP ratio rises because households smooth their consumption over time since the shock is not permanent. The slight decline in the current account-GDP ratio is due to the rising investment-GDP ratio, as shown in the lower left graph, which makes the response of the current account-GDP ratio for the immediate availability case different from the pure endowment economy. As discussed by Uribe and Schmitt-Grohé (2014), unanticipated persistent shocks (e.g. the contemporaneous oil shock) can have negative impact effects on the current account in an economy with modest investment adjustment cost. However, when the shock is contemporaneous, rising investment is the only channel leading to the counter-cyclical movement of the current account. In contrast, the news shocks leads to counter-cyclical of the current account through two channels: rising investment and declining savings.

The upper right graph shows the results for GDP. In the case of news about future resources, GDP rises very little during the first five years, and then shoots up rapidly before gradually declining. In contrast, shocks to current resources cause output to jump immediately. It stays high for several years during which the reallocation of capital and labor to the oil sector overcomes the exogenous depletion, and then falls.

Consumption jumps immediately in both experiments, as one would expect from a positive wealth shock. It jumps more for the contemporaneous shock because the wealth effect is greater since the present discounted value of the increase in resources is greater.

The biggest qualitative difference is in the response of hours. The contemporaneous shock to the oil sector leads to a rise in hours during the first five years, followed by oscillations around zero thereafter. In contrast, hours fall in both the short-run and the medium run in the case of the news shock.¹⁰ Thus, even when we use preferences that are close to GHH preferences, thus shutting down most of the wealth effect on labor supply, we find that hours *fall* in response to a news shock. This is in contrast to Jaimovich and Rebelo's (2008) findings.

¹⁰ It should be noted, though, that in both cases the response of hours is small, only one-tenth of the response of aggregate GDP (note the scale of the graph).

In order to understand whether the special features of the oil sector – low labor share and high capital share – account for the differences in our model’s predictions relative to Jaimovich and Rebelo’s model, we conduct two additional experiments. The first experiment compares the effect of the oil news shock to the effect of a news shock on TFP in a standard one-sector economy, similar to the experiment analyzed in Jaimovich and Rebelo (2008). In both cases, the news arrives five years before the increase in resources or TFP. Figure III shows these results. The paths of GDP look very similar across the experiments, in part because we normalized their peaks to be the same. The qualitative responses of the current account-GDP ratio, the saving-GDP ratio, and the investment-GDP ratio are also similar across experiments, but the swings are much more pronounced for the oil shock than for the aggregate TFP. The magnification is due to the required capital reallocation in the case of the oil shock, particularly since the capital share in the oil sector is higher than in sector 1. The responses of consumption and hours are hump-shaped and both are positive for the TFP shock. Because the Jaimovich-Rebelo preferences induce nonseparability in consumption and hours, the consumption responses tend to mirror the hours responses. Thus, even with our slightly different calibration of several (non-oil sector related) parameters, our use of an annual frequency, and longer news lags, the aggregate TFP news shock still has a positive effect on hours.

Our second experiment is designed to determine whether the difference in results is due to the special nature of the oil sector, with its low labor share and high capital share, or to the assumption that the shock hits only one sector. To answer this question, we compare our baseline model, in which the Sector 2 labor share is 0.13, the capital share is 0.49, and the resource share is 0.38, to a model in which Sector 2 has a labor share equal to 0.4, capital share equal to 0.22, and a resource share of 0.38. Thus, the latter parameterization sets the ratio of labor share to capital share equal to their ratio in Sector 1, which are standard values used for one-sector economies.¹¹

Figure IV shows the results of the second experiment. The qualitative responses of GDP, the current account-GDP ratio and the saving-GDP ratio are similar across experiments. The investment-GDP response is substantially muted in the experiment with the lower capital share.

¹¹ Alternatively, if we raise the labor share in the oil sector to 0.6 so that it is close to the one in Sector 1, and set the capital share and resource share to 0.3 and 0.1 respectively, the response of investment and hours is very similar.

This implies that a large portion of the spike in investment in our baseline model is due to the high capital share in the resource sector. The differences in the responses of hours are more complex. We can see the responses of hours look very similar whether the oil sector is high or low labor share in the interval between the arrival of the news and the start of oil production, but differ significantly after oil production starts. The decline in hours during the initial interval is due to the anticipated wealth effect, which is similar across experiments. However, after the oil production begins, the impact of reallocation from Sector 1 to Sector 2 on hours crucially depends on whether Sector 2 is labor intensive or capital intensive. Thus, when oil production commences, the hours in the baseline case continue to decline as the reallocation effect leads to a decline in labor demand, while the hours become significantly positive in the fifth year for the case of high labor share in the oil sector.

Thus, our theoretical finding of a negative news effect on hours is not due to the special calibration required for a realistic oil sector. However, our finding of a sustained decline in hours after the oil is produced *is* dependent on the special features of the oil sector. Perhaps these results on hours should not be a surprise, given Jaimovich and Rebelo's (2009) two-sector model results in a closed economy setting. They model consumption and investment goods sectors, and assume the same labor and capital shares. They show that the parameter space that is consistent with a positive effect of sector-specific news on aggregate hours is much smaller than for an aggregate shock; indeed, the lower panel of their Table 1 shows that the space is exceedingly small.

To summarize, our theoretical analysis shows that the current account response to news is robust to sector-specific news shocks (whether labor share is high or low) and to aggregate TFP shocks. In every case, the current account becomes significantly negative for the five years between the arrival of the news and the increase in resources. In contrast, the behavior of some of the other variables, such as labor input and consumption, depends significantly on the details of the shock. In particular, we find that the positive response of hours to the news shock is a feature only of the one-sector model. In the two-sector model, hours decline when the news arrives but before it materializes, no matter whether the labor share in the oil sector is low or high. After oil production starts, hours fall more if the labor share is low, while hours rise robustly if the labor share is high.

Testing the theory empirically presents serious challenges because it is difficult to find a source of macro-relevant and country-specific news shocks. In the following sections, we use a unique panel data of announcement of giant oil discoveries as news shocks to test the theory.

III. WHY USE GIANT OIL DISCOVERIES?

Evaluating the empirical relevance of news shocks is quite challenging. Difficulties arise on two main fronts. First, our theory suggests that the main driving force is agents' perception of future availability of resource input, but it is empirically difficult to measure agents' expectation as is well-known from the literature on news shocks. As discussed earlier, 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, 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 130 countries.¹² 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.

The first attractive feature of giant oil discoveries is that they signal significant increases in production possibilities in the future.¹³ To be able to test the effect of news shocks on the dynamics of macroeconomic aggregates, particularly to isolate a significant anticipation effect,

¹² 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.

¹³ For example, the IMF (2013a) released its latest estimation suggesting that the recent energy booming in the U.S. could increase the real GDP in the U.S. by about 1.2 percent and employment by 0.5 percent over the next 12 years, if the energy production is assumed to increase during this period due to the so called shale revolution.

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 with 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. Anecdotally, the average delay in the United Kingdom during the period 1954-2011 was 4.5 years.¹⁴ Experts' empirical estimates suggest that for a giant oil discovery, it takes between 4 and 6 years to go from drilling to production.¹⁵ Based on our own calculation using an alternative data source to Mike Horn's that is less comprehensive but for which we have more detailed information at the field level, we find that the average delay between discovery and production start is 5.4 years.¹⁶ 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 6.4 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.

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 and previous discoveries.¹⁷ This feature is crucial for our identification of the anticipation effect on macroeconomic aggregates including the current

¹⁴ See for instance report from United Kingdom, Department of Energy and Climate Change, 2013: https://www.gov.uk/.../130718_decc-fossil-fuel-price-projections.pdf

¹⁵ See for instance, <http://www.ellipticalresearch.com/drillingandoilproduction.html>. Mike Horn relies on a 7 year time lag between discovery and production.

¹⁶ 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.

¹⁷ 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.

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.¹⁸ 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 the period 1970-2012 for up to 130 countries. This allows us to adopt panel data estimation techniques which control for country and year fixed effects.

The giant oil discovery dataset is from Horn (2004). Some summary statistics of giant oil discoveries around the world are now discussed. Table II shows the spatial and temporal distribution of giant oil discoveries 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 Common Wealth of Independent States and Mongolia (51) also experienced significant numbers of discovery events during the same period.¹⁹ The 1970s is the peak period for giant oil discoveries, but the number of discoveries has been growing since 1980s. This contradicts the commonly held view that it is more and more

¹⁸ 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.

¹⁹ 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.

difficult to discover new oil fields.²⁰ Figure V 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:

$$y_{it} = A(L)y_{it} + B(L)Disc_{it} + \alpha_i + \gamma_0' dt + \gamma_1' Z_{it} + \epsilon_{it}, \quad (13)$$

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 employment-population ratio. α_i controls for country fixed effects which capture unobserved time invariant characteristics such as geographical location, dt 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 in which we describe in greater details 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 richness of the data has three advantages. First, the panel structure allows us to identify the dynamic effect of oil discoveries on macroeconomic aggregates, while controlling for country-

²⁰ Technological innovations for exploration and drilling render future discoveries more likely. One notorious example of the role of technology in oil exploration and drilling is George Mitchell's innovative use of horizontal wells and hydraulic fracturing in the 1990s to release gas from a previously-impermeable rock formation near Fort Worth, Texas. Those drilling breakthroughs have paved the way for tapping into previously inaccessible and vast oil and gas reserves including in the United States, Poland and Argentina.

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.²¹ Second, including lagged values of oil discoveries allows us to control for the possible correlations in oil discoveries. Thus, we can identify the conditional effect of oil discoveries at a given point in time on macroeconomic variables. Third, 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 $IRF(L) = B(L)/(1 - A(L))$. Moreover, we constructed an extensive panel data (both in terms of the number of cross-sectional units, N and time span, T) to fully utilize within 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 large panel dataset to capture the dynamic effect of those discoveries.

IV.B. Data

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.²² 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 giant oil

²¹ 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.

²² 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.

discoveries. An approach using solely variation in the timing would not take into account the heterogeneity in the size of discoveries.²³

We construct a measure of the net present value of a giant oil discovery as a percent of GDP, NPV , as follows:

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

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).²⁴ 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.²⁵

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.²⁶ The risk free rate is assumed to be the rate prevailing in the United States. Considering that measures of risk

²³ In practice, though, our qualitative results are robust to using discoveries events in lieu of our benchmark NPV measures.

²⁴ 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.

²⁵ See Hamilton (2009) and references therein for a discussion on forecasting oil prices.

²⁶ 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.

premia based on sovereign bond spreads are not readily available for all countries and when 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 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 1997-2007.²⁷ Emerging markets are a set of countries for which risk ratings can vary substantially and thus providing with significant statistical variation to estimate 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 from International Country Risk Guide (2012) covering 133 countries in the world. To examine the effects that political risk has on sovereign bond spreads, we estimate the following econometric model:

$$\ln(\text{Spread}_{i,t}) = \alpha_0 + \alpha_1 \ln(\text{PoliticalRisk}_{i,t}) + \beta_i + \mu_t + u_{i,t}$$

where β_i are country fixed effects, μ_t are year effects and $u_{i,t}$ is an error term.²⁸ We estimate the elasticity of the sovereign spreads to political risk ratings using our sample. We then predict the $\widehat{\text{Spread}}_i$ given country's political risk rating and compute the NPV of giant oil discoveries accounting for country specific discount rates.

Figure VI 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

²⁷ 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.

²⁸ The estimated coefficients used in the prediction are as follows:

$$\ln(\text{Spread}_{i,t}) = 13.27 - 1.76 \times \ln(\text{PoliticalRisk}_{i,t}),$$

(3.71) (-2.03)

The t-statistics in parenthesis indicates that political risk is a significant determinant for the sovereign bond spreads for emerging markets. R-squared is 0.41. We further adjust the estimated equation to obtain a predicted spread of zero for countries with the level of risk equal to the one of the United States or lower.

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.

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 covers 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 VII and VIII 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.²⁹ The shaded areas are 90% and 68% (darker grey) confidence bands based on Driscoll-Kraay (1998) standard errors and the delta method.³⁰

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

²⁹ 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 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.

³⁰ We use *xtscc* and *nlcom* commands in Stata. The bands are similar if we instead use a non-parametric bootstrapping method.

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 VII 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.³¹

The top panel of Figure VIII 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 7 to 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. We will discuss the quantitative match below.

The second panel of Figure VIII 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 very 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 VIII 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.

³¹ These swings in variables from negative to positive in response to news about future opportunities are reminiscent of Mertens and Ravn's (2012) analysis of the effects of news about future tax changes.

³² We use the employment rate rather than total hours because the latter were not available for all of the countries in our sample.

Not all of the point estimates of the 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 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).³³ 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 0.06) 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.³⁴ The consumption response, however, has a p-value of 0.13, so even the integral response is not statistically significant at conventional levels. On the other hand, the employment response is (statistically) significantly negative.

Figure IX compares the empirical estimates of the responses of all six variables against the responses derived from our theoretical model. The latter are normalized so that the peak response of GDP is the same as the one from the empirical model; we will return to the issue of magnitude in the next paragraph. The graphs show that the empirical and theoretical responses are broadly

³³ The null hypothesis is consistent with the response of the current account to a contemporaneous positive shock in the oil sector, thus we can reject that an oil discovery is a contemporaneous shock if we reject the null.

³⁴ The two high p-values occur only because the estimates are not precise enough to rule out some values opposite in sign to those predicted by the theory.

consistent in shape, magnitude, and timing. There are a few notable differences. For example, the (positive) current account and savings responses are larger in the empirical responses than in theoretical model ones, following the start of oil production. Furthermore, the saving response reaches a peak later in the empirical model than in the theoretical one. The response of consumption is larger than predicted by the theoretical model, but the empirical estimates are imprecise. The negative employment response in the theoretical model is somewhat more pronounced than in the empirical one. That said, the responses from empirical and theoretical models are overall remarkably close, particularly since we used calibrations from other papers.

There is one notable discrepancy between the model and the data that is not evident from the graphs. As discussed earlier, we normalized the shock in the model so that the GDP peak would match the empirical responses. The shock used in the empirical estimates is an oil discovery that we estimate to lead to a potential stream of oil revenues equal to one percent of initial GDP. The shock in the model represents a much smaller shock. If we had used a shock that represented an increase in the value of oil production that was one percent of GDP in NPV, we would have found that all of the theoretical model responses would have been substantially greater. For example, the peak of GDP would have been 0.19 rather than 0.03. This discrepancy implies that either our estimate of the oil discovery NPV is higher than the value perceived by the agents at the time or that there are important elements that our model does not incorporate. For example, the discounting factor we used in the NPV calculation might be not high enough as the risk premium does not fully capture all the uncertainty associated with oil extraction including unknown features of oil fields and the related social tensions. However, those factors are either difficult to measure precisely or to incorporate into the theoretical model.

VI. Extensions

We now seek to shed some light on two elements we did not incorporate in our theory – the potential government response and real exchange rates. We begin by considering the behavior of government. Our data on investment and consumption is for total amounts, aggregating 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 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 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.

In order to keep our 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 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 (See the early contribution by Eastwood and Venables, 1982).³⁵ Empirically, we found that the real exchange rate appreciated during the first five years following oil discoveries, either using CPI-based or GDP deflator-based measures. However, in both specifications, the point estimates were very imprecise so no response was significantly different from zero. Real exchange rate responses are shown in the Supplementary Appendix. Our results are somewhat consistent with the empirical

³⁵ 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: when forward-looking households and then the government increases their consumption.

literature on the so-called Dutch disease, finding mixed evidence for the latter.³⁶ 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.

VII. ROBUSTNESS CHECKS

In the following, we discuss the results of extensive robustness checks for the benchmark specification. First, our main results are robust to alternative measures of NPV of giant 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.³⁷ 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. (Figures C.I in the Supplementary Appendix.)³⁸

Also, we check whether removing groups of countries affects the robustness of the effect of giant oil discoveries on the current account. Given the relative concentration of giant oil discoveries in certain Middle East and North African countries, we tried removing all countries belonging to

³⁶ See for instance 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 rate for oil exporters and find some evidence of asymmetrical Dutch Disease along boom-bust cycles.

³⁷ One advantage of this measure is that we can expand our sample to cover about 170 countries for our key macro variables, although most of countries in the extended sample do not have giant oil discoveries.

³⁸ As we noted earlier, most of our results are also qualitatively robust to using a simple dummy variable for oil discoveries. In fact, we used the dummy variable for our baseline in an earlier version of the paper.

this region, and our main results still hold. Also, we explore the robustness of our main results to using solely countries which have experienced at least one discovery (Figure C.II in the Supplementary Appendix). Results are again robust thus suggesting that the lack of comparison from countries where discoveries are absent is not an issue.

Next, we investigate formally the predictability of giant oil discoveries. Indeed, agents might learn about giant oil discoveries from information observable to the econometrician such as previous discoveries which has been controlled for in the empirical specification. However, 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 in turn 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 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 lagged values of current accounts and investment do not have significant predictive power on the incidence of oil discoveries.³⁹ The conclusion that emerges from these tests is that there are no evidence of predicting power of giant oil discoveries derived from lagged values of the current account and investment. However, 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 C.III in the Supplementary Appendix). Finally, we also controlled for current and lagged exploration expenditures using data from Global Energy Systems and our main results are robust to doing so (Figure C.IV in the Supplementary Appendix).

³⁹ Saving is excluded due to obvious redundancy.

One may also wonder whether the anticipation effect of news shocks on the current account may depend on countries' borrowing constraints. This is true in our theory where borrowing constraints are modeled as a debt-elastic interest-rate premium over the world interest rate for small open economies which is a function of countries' foreign asset/debt ratio (see for instance, Schmitt-Grohé and Uribe, 2003). However, borrowing constraints are possibly endogenous to giant oil discoveries. Because poor countries initially may have limited access to international financial capital markets due to collateral constraints, they however could use the giant oil discoveries as collateral to borrow from abroad once they discovered such new oil fields. Thus, it is still possible to identify the anticipation effect of giant oil discoveries on the current account for countries with ex ante borrowing constraints. Indeed, we found evidence that anticipated effect to giant oil discoveries on the current account for African countries only is present and similar to our benchmark results (Figure C.V in the Supplementary Appendix).

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 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.⁴⁰ Our main results are virtually unchanged (Figure C.V in the Supplementary Appendix).

Our results are also robust to using different 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

⁴⁰ 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.

dependent variable. We also try to remove the controls for country-specific quadratic trend for output, employment rate and other level variables. The patterns of impulse responses are also similar. (Figure C.VI and C.VII in the Supplementary Appendix).

VIII. Conclusion

In this paper, we have examined the effect of news about giant oil discoveries on macroeconomic aggregates. We first presented a two-sector open economy model to highlight the predicted effect of sector-specific news on variables such as the current account, saving, investment, GDP, consumption, and employment. 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. To test this prediction, we estimated a dynamic panel distributed lag model over a sample covering 1970 to 2012 for about 130 countries. Results from the estimation provided evidence for a statistically and economically significant anticipation effect following the announcement of a giant oil 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 oil production did we find that the current account and saving increased along with GDP. In contrast, we found that employment fell when the news arrived and continued to decline for a number of years.

One might wonder whether the special nature of our oil news shock limits the degree to which it can inform the news-driven business cycle literature. Our theoretical analysis shows that the initial negative news effect on hours is not driven by the special feature of low labor share in the oil sector. Rather, the negative effect occurs more generally when the news is about a sector-specific, rather than an aggregate, TFP shock. While aggregate TFP shocks are useful for theoretical exercises, it would be difficult to find one in reality. Even general purpose technologies that ultimately raise aggregate TFP will initially be anticipated in only a few sectors.

Our finding that news about future economic events leads to immediate responses 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 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.

Finally, our results also provide strong support for 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).

References

- Alexopoulos, Michelle.** 2011. "Read All About It!! What Happens Following a Technology Shock?" *American Economic Review*, 101(4), 1144-1179.
- 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.
- 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.
- 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.
- Caselli, Francesco and James Feyrer.** 2007. "The Marginal Product of Capital." *Quarterly Journal of Economics*, 122(2), 535-568.
- 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.** 2011. "Giant Oil and Gas Fields of the World".
<http://www.datapages.com/AssociatedWebsites/GISOpenFiles/HornGiantFields.aspx>.
- International Country Risk Guide.** 2012. *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.

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.

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.

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.

Obsfeld, 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.

Obsfeld, Maurice. 2012. "Does the Current Account Still Matter?" *American Economic Review*, 102(3): 1-23.

- 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." *The Quarterly Journal of Economics* 126.1: 1-50.
- 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.
- Schmitt-Grohé, Stephanie, and Martín Uribe.** 2003. "Closing Small Open Economy Models." *Journal of international Economics*, 61(1): 163-185.
- Schmitt-Grohé, Stephanie, and Martín 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.
- Samuel Wills.** 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).

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 (6), μ is the multiplier on the utility function state variable definition in equation (2), and the η 's are the multipliers on the capital accumulation equations in equation (5).

The first-order conditions for this economy are:

$$(C_t - \psi N_t^\theta X_t)^{-\sigma} + \mu_t \gamma C_t^{\gamma-1} X_{t-1}^{1-\gamma} = \lambda_t, \quad (\text{A.1})$$

$$(C_t - \psi N_t^\theta X_t)^{-\sigma} \psi N_t^\theta + \mu_t = \beta E_t [\mu_{t+1} (1 - \gamma) C_{t+1}^\gamma X_t^{-\gamma}], \quad (\text{A.2})$$

$$\begin{aligned} & (C_t - \psi N_t^\theta X_t)^{-\sigma} \psi \theta N_t^{\theta-1} X_t \\ &= \lambda_t \left[\frac{\alpha_1 Y_{1,t}}{N_{1,t}} - \frac{\Psi_1}{2} \left(\frac{N_{1,t}}{N_{1,t-1}} - 1 \right)^2 - \Psi_1 \left(\frac{N_{1,t}}{N_{1,t-1}} - 1 \right) \frac{N_{1,t}}{N_{1,t-1}} \right] \\ &+ \beta E_t \left[\lambda_{t+1} \Psi_1 \left(\frac{N_{1,t+1}}{N_{1,t}} - 1 \right) \left(\frac{N_{1,t+1}}{N_{1,t}} \right)^2 \right] \end{aligned} \quad (\text{A.3})$$

$$\begin{aligned} & (C_t - \psi N_t^\theta X_t)^{-\sigma} \psi \theta N_t^{\theta-1} X_t \\ &= \lambda_t \left[\alpha_2 \frac{p_t Y_{2,t}}{N_{2,t}} - \frac{\Psi_2}{2} \left(\frac{N_{2,t}}{N_{2,t-1}} - 1 \right)^2 - \Psi_2 \left(\frac{N_{2,t}}{N_{2,t-1}} - 1 \right) \frac{N_{2,t}}{N_{2,t-1}} \right] \\ &+ \beta E_t \left[\lambda_{t+1} \Psi_2 \left(\frac{N_{2,t+1}}{N_{2,t}} - 1 \right) \left(\frac{N_{2,t+1}}{N_{2,t}} \right)^2 \right] \end{aligned} \quad (\text{A.4})$$

$$\lambda_t = \beta E_t (\lambda_{t+1} (1 + r_{t+1})) \quad (\text{A.5})$$

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

$$\eta_{1,t} = \beta E_t \left[(\lambda_{t+1} (1 - \alpha_1) Y_{1,t+1} / K_{1,t} + \eta_{1,t+1} (1 - \delta_1)) \right] \quad (\text{A.7})$$

$$\eta_{2,t} = \beta E_t \left[(\lambda_{t+1} \alpha_k p_{t+1} Y_{2,t+1} / K_{2,t} + \eta_{2,t+1} (1 - \delta_2)) \right] \quad (\text{A.8})$$

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.⁴¹ To compute the net present value of giant oil field discoveries used in 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:

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

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

⁴¹ The net present value computation is based on 5-year delay between the discovery announcement and the start of oil production.

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

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 |
| OPEC giants | 0.36 | 0.70 | 0.75 |
| Non-OPEC giants | 0.67 | 0.64 | 0.58 |
| Maximum depletion rate | | | |
| Fields | γ | δ | R ² |
| All giants | 0.64 | -0.31 | 0.31 |
| OPEC giants | 0.40 | -0.26 | 0.36 |
| Non-OPEC giants | 0.66 | -0.30 | 0.26 |

The difference between parameter estimates for giant fields in OPEC versus non-OPEC is relatively small. For simplicity, 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.

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:

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

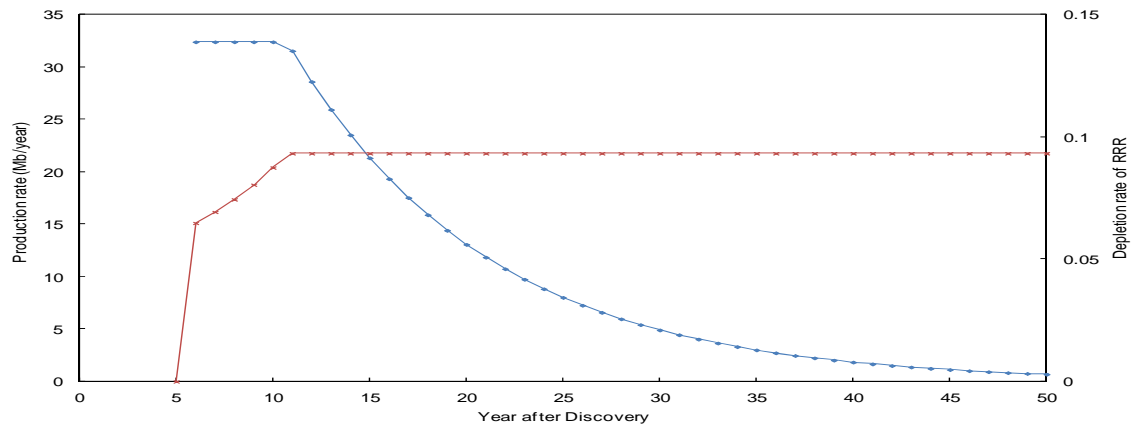


Figure B.I. Production and depletion rate for a field of 500 million barrels of ultimately recoverable reserves.

B.II. 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 % 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) |

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

| Variable | Years | Maximum number of countries | Obs | Min | Median | Max |
|--|-----------|-----------------------------|------|---------|--------|---------|
| Ln(GDP) | 1980-2012 | 184 | 5537 | -356.15 | 579.65 | 1477.83 |
| CA/GDP | 1980-2012 | 180 | 5408 | -242.19 | -3.2 | 106.84 |
| Saving/GDP | 1980-2011 | 171 | 4711 | -202.89 | 18.19 | 107.18 |
| Investment/GDP | 1980-2011 | 178 | 4956 | -2.42 | 21.13 | 113.58 |
| Ln(final consumption) | 1980-2012 | 163 | 4600 | -82.29 | 593.84 | 1435.58 |
| Employment rate | 1991-2012 | 165 | 3630 | 28.9 | 57.4 | 88.1 |
| Public investment/GDP | 1980-2012 | 106 | 1992 | -13.67 | 15.43 | 95.87 |
| Private investment/GDP | 1980-2012 | 106 | 1992 | -4.5 | 4.31 | 47.08 |
| Ln(final private consumption) | 1980-2012 | 161 | 4544 | -113.86 | 569.92 | 1426.54 |
| Ln(final public consumption) | 1980-2012 | 162 | 4556 | -268.61 | 425.7 | 1223.26 |
| Government spending/GDP | 1980-2012 | 105 | 1959 | 1.82 | 21.53 | 87.63 |
| Ln(real exchange rate based on CPI) | 1980-2010 | 180 | 4776 | 285.07 | 463.42 | 1528.4 |
| Ln(real exchange rate based on GDP deflator) | 1980-2011 | 180 | 5107 | 225.58 | 460.69 | 2608.46 |

Table I. Baseline Calibrated Parameters for Two-Sector Model

| Parameter | Name | Value |
|------------------|--|---------------------|
| β | Discount factor | 0.9 |
| ψ | Governs disutility of labor, set so steady-state labor is 20. | 0.0197 |
| θ | Exponent on labor in the utility function, governing intertemporal substitution. | 1.2 |
| σ | Governs intertemporal substitution of the consumption-hours bundle | 1 |
| γ | Governs the wealth effect; GHH preferences = 0 | 0.01 |
| ϕ_h | Investment adjustment cost parameter | 0.1 in both sectors |
| Ψ_h | Labor adjustment cost parameter | 0.5 in both sectors |
| δ_h | Capital depreciation | 0.1 in both sectors |
| α_1 | Labor share in non-resource sector | 0.64 |
| α_2 | Labor share in resource sector | 0.13 |
| α_k | Capital share in resource sector | 0.49 |
| χ | Elasticity of interest rate with respect to net foreign assets | 0.00038 |
| \bar{B} | Parameter in interest rate function; set so that the steady-state tb/y = 0.02 | -13 |
| ρ_1, ρ_2 | Parameters governing exogenous depletion rate of reserves | 1.4, -0.5 |
| p | Relative price of oil, set so that steady-state oil share of GDP is 7% | 1.1 |
| 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 |
|---|-------|-------|-------|-------|-------|-------|
| 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 | 45 |
| 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. $H_1: \sum_{h=5}^{11} b_h > 0$ | 0.05 |
| Investment/GDP | Positive response at horizons 0-4 | $H_0: \sum_{h=0}^4 b_h \leq 0$ vs. $H_1: \sum_{h=0}^4 b_h > 0$ | 0.06 |
| | Negative or zero response starting at horizon 5 | $H_0: \sum_{h=5}^{11} b_h > 0$ vs. $H_1: \sum_{h=5}^{11} b_h \leq 0$ | 0.58 |
| GDP | Positive or zero response at horizons 0-4 | $H_0: \sum_{h=0}^4 b_h < 0$ vs. $H_1: \sum_{h=0}^4 b_h \geq 0$ | 0.83 |
| | 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. $H_1: \sum_{h=0}^{11} b_h > 0$ | 0.13 |
| 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 | 0.61 | 0.45 | 0.69 | 0.57 |
| Logit | 0.75 | 0.53 | 0.72 | 0.62 |
| 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 $H_1: \alpha_X \neq 0$.

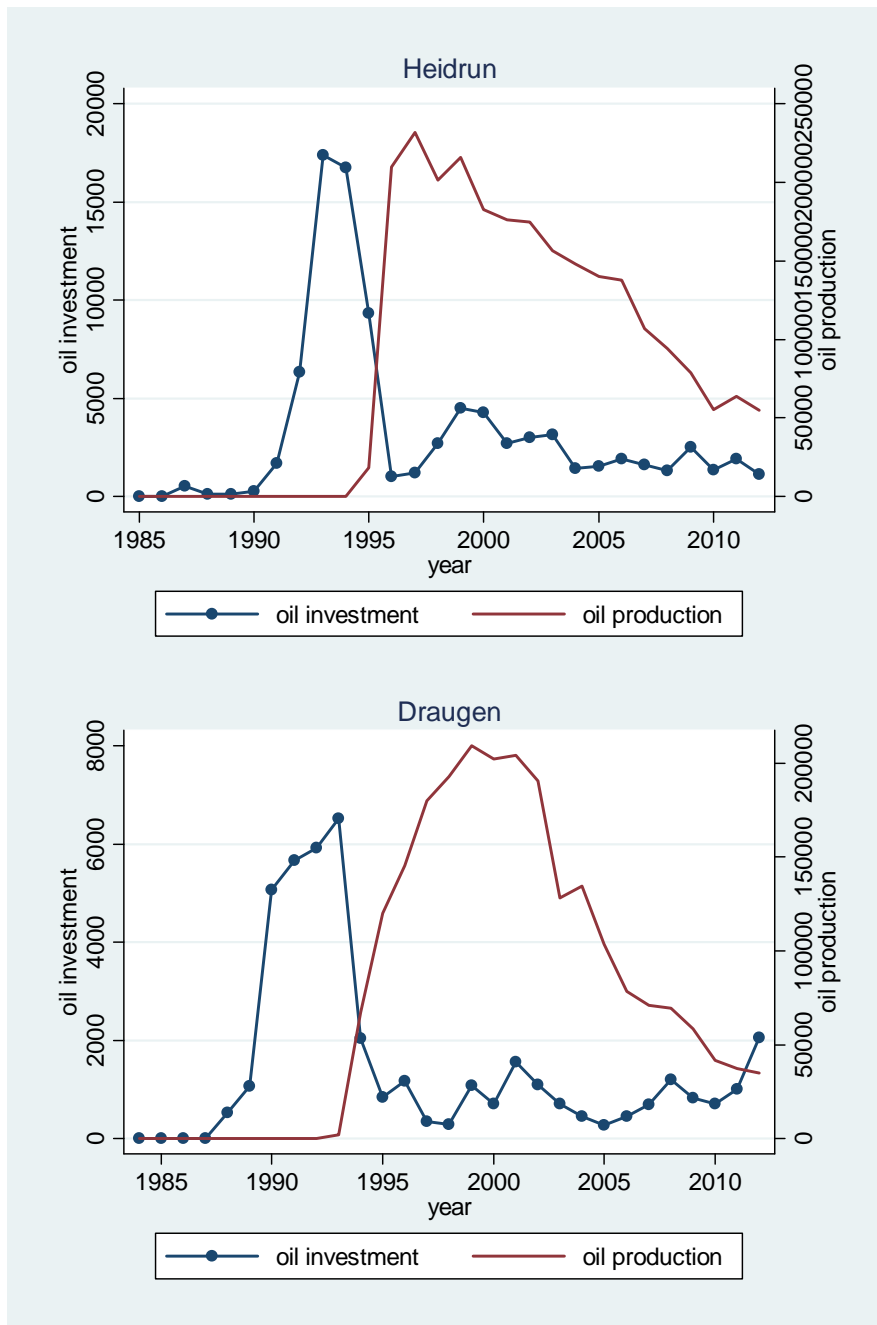


Figure I. Typical Oilfield Investment and Production Patterns: Examples from Two Norwegian Oil Fields

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

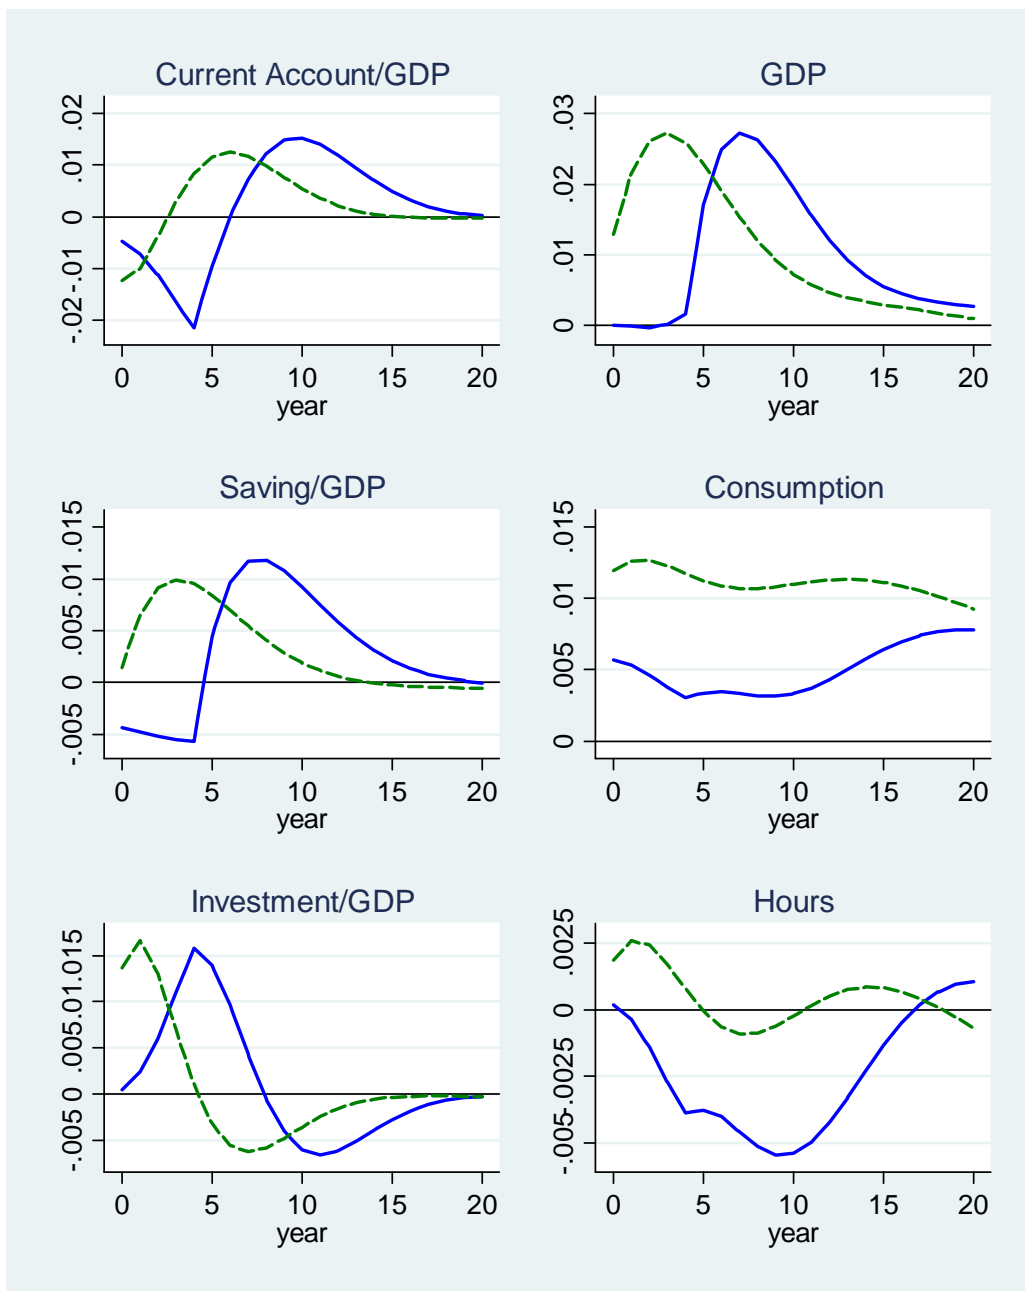


Figure II. Effect of Oil News versus Unanticipated Oil Shock

Baseline Model, 5 year lag of news

Note: The vertical axis shows percentage changes. The solid blue line is the oil news. The green dashed line is the unanticipated oil shock. Shocks are constructed so that the peak of GDP is the same across experiments.

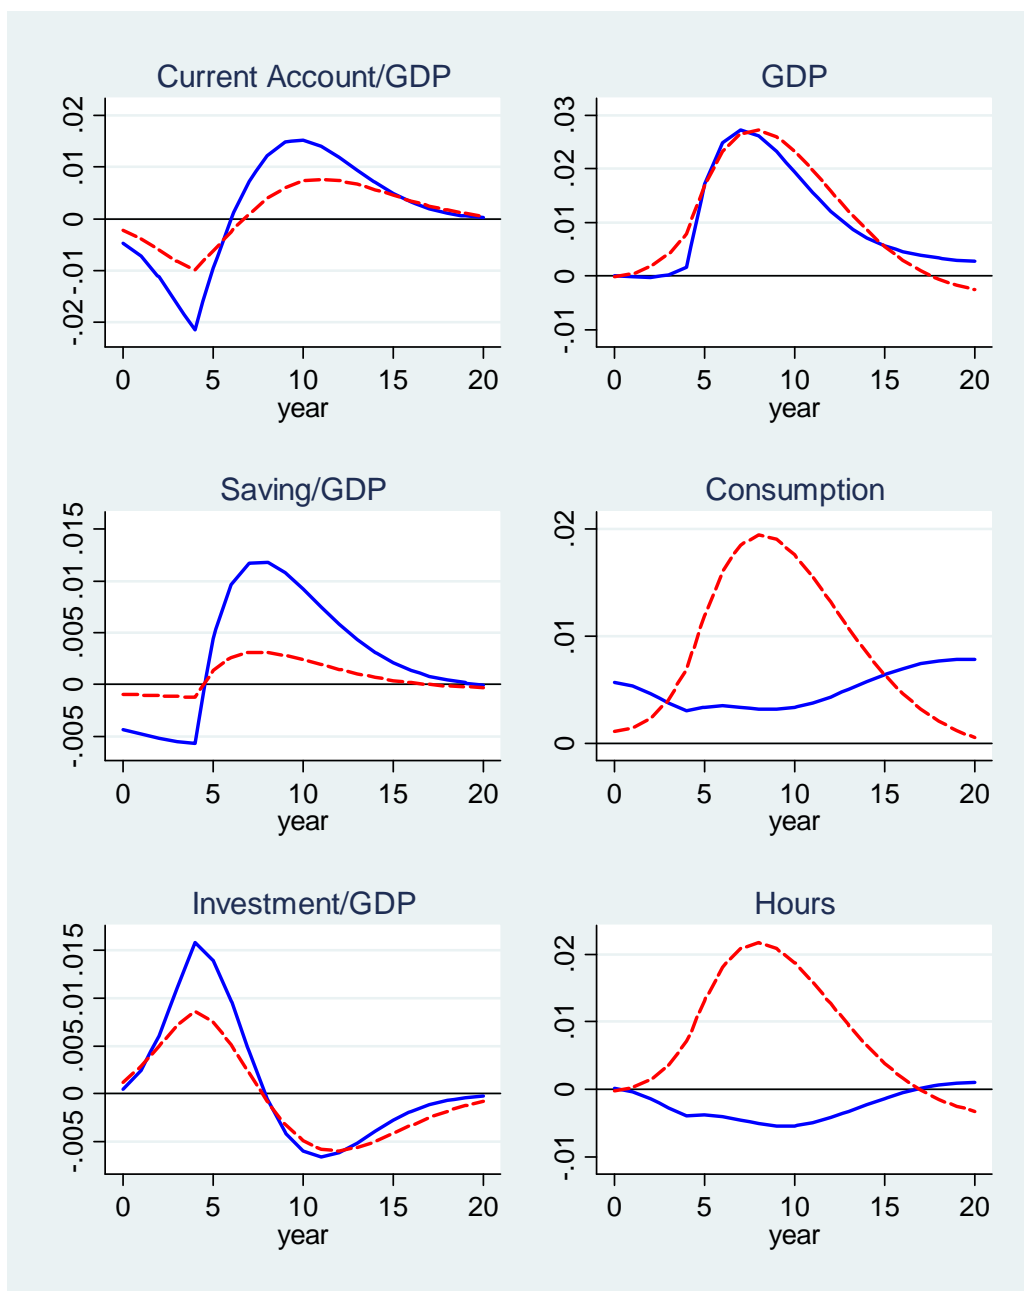


Figure III. Effect of Oil News vs. Aggregate TFP News in One Sector Model

Baseline Model with near GHH preferences, 5 year lag of news

Note: The vertical axis shows percentage changes. The solid blue line is the oil news. The red dashed line is the aggregate TFP shock. Shocks are constructed so that the peak of GDP is the same across experiments.

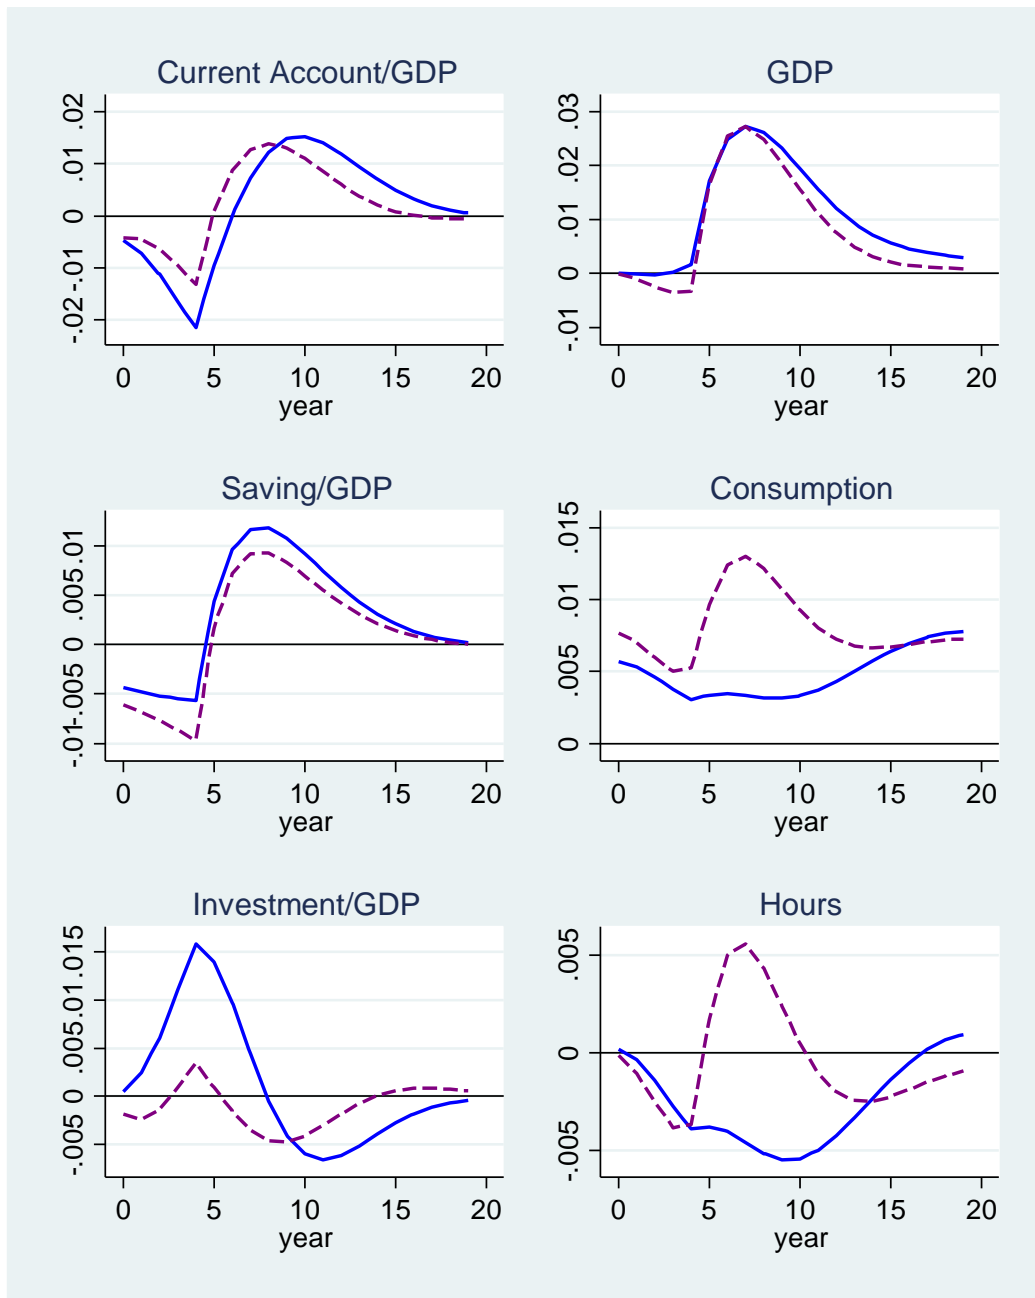


Figure IV. News in 2 Sector Models: The Impact of Low Labor Share,
5 year lag of news

Note: The vertical axis shows percentage changes. The solid blue line is the baseline model, where sector 2 has low labor share. The purple dashed line shows the effect of news when sector 2 has high labor share as details are explained in text. Shocks are constructed so that the peak of GDP is the same across experiments.




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

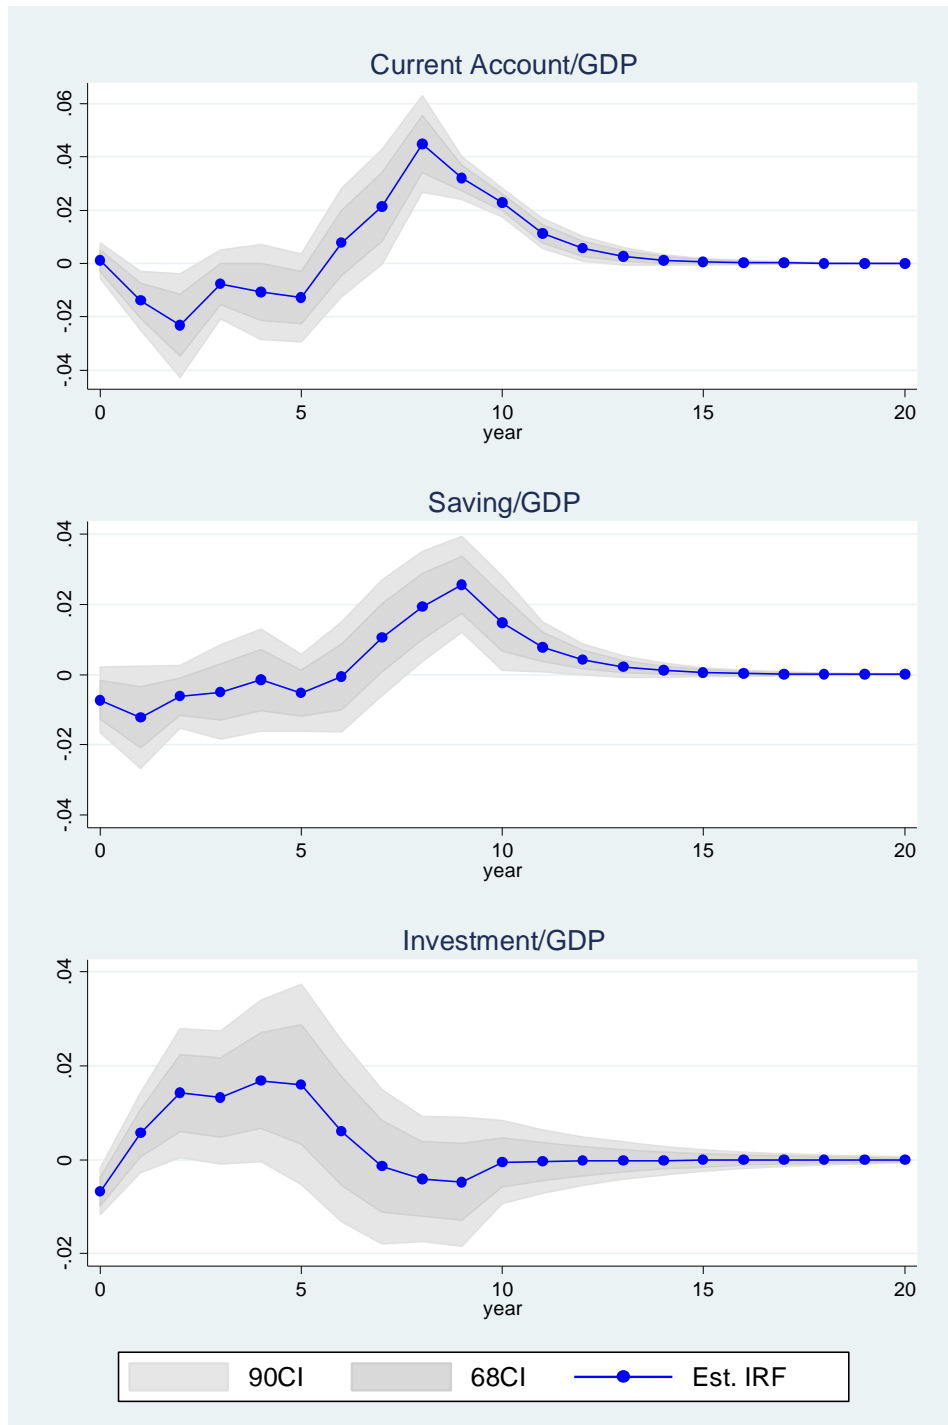


Figure VII: The Impact of Giant Oil Discoveries on the Current Account, Saving and Investment

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.

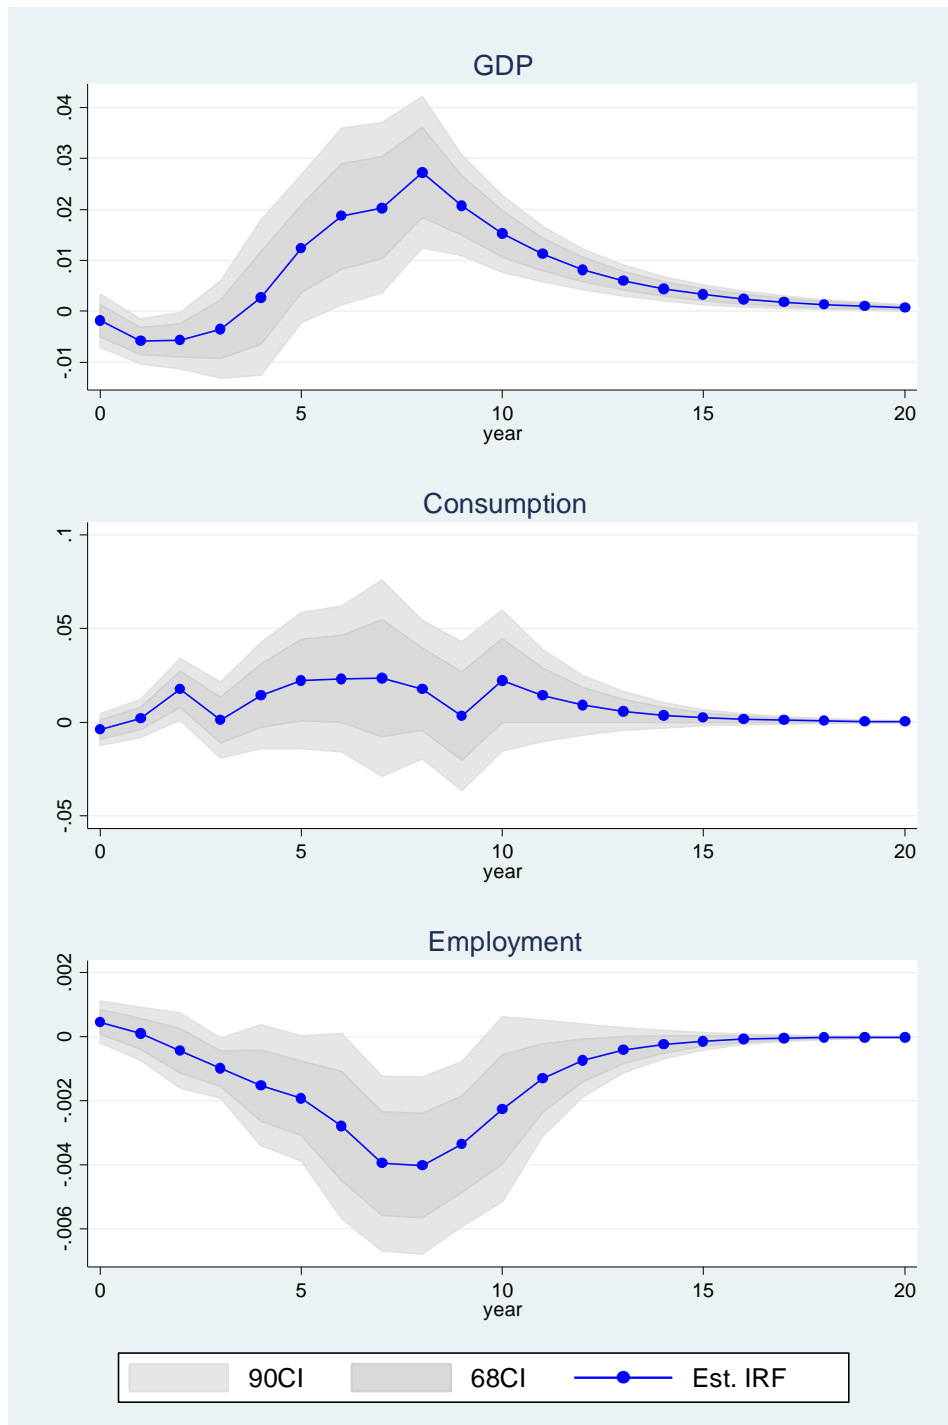


Figure VIII: The Impact of Giant Oil Discoveries on GDP, Consumption and Employment

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.

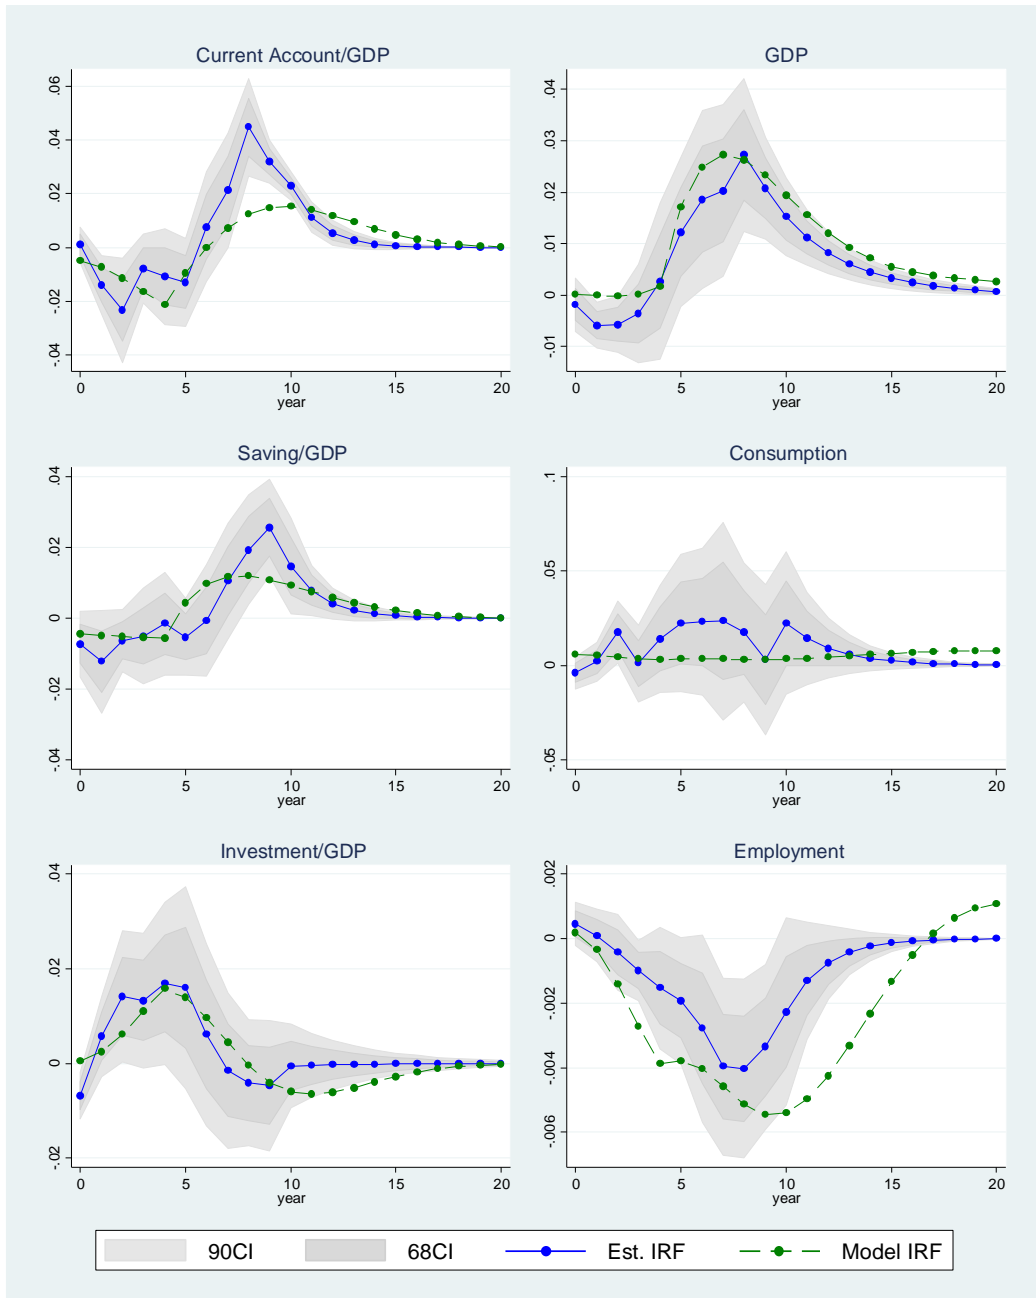


Figure IX: Comparison of the Estimates to the Theoretical Predictions

Note: The green dash lines are the theoretical impulse response. The solid blue lines are the point estimates, and grey areas are 90 percent and 68 confidence intervals. The responses are normalized so that the peak of GDP is the same. The vertical axis shows percentage changes.

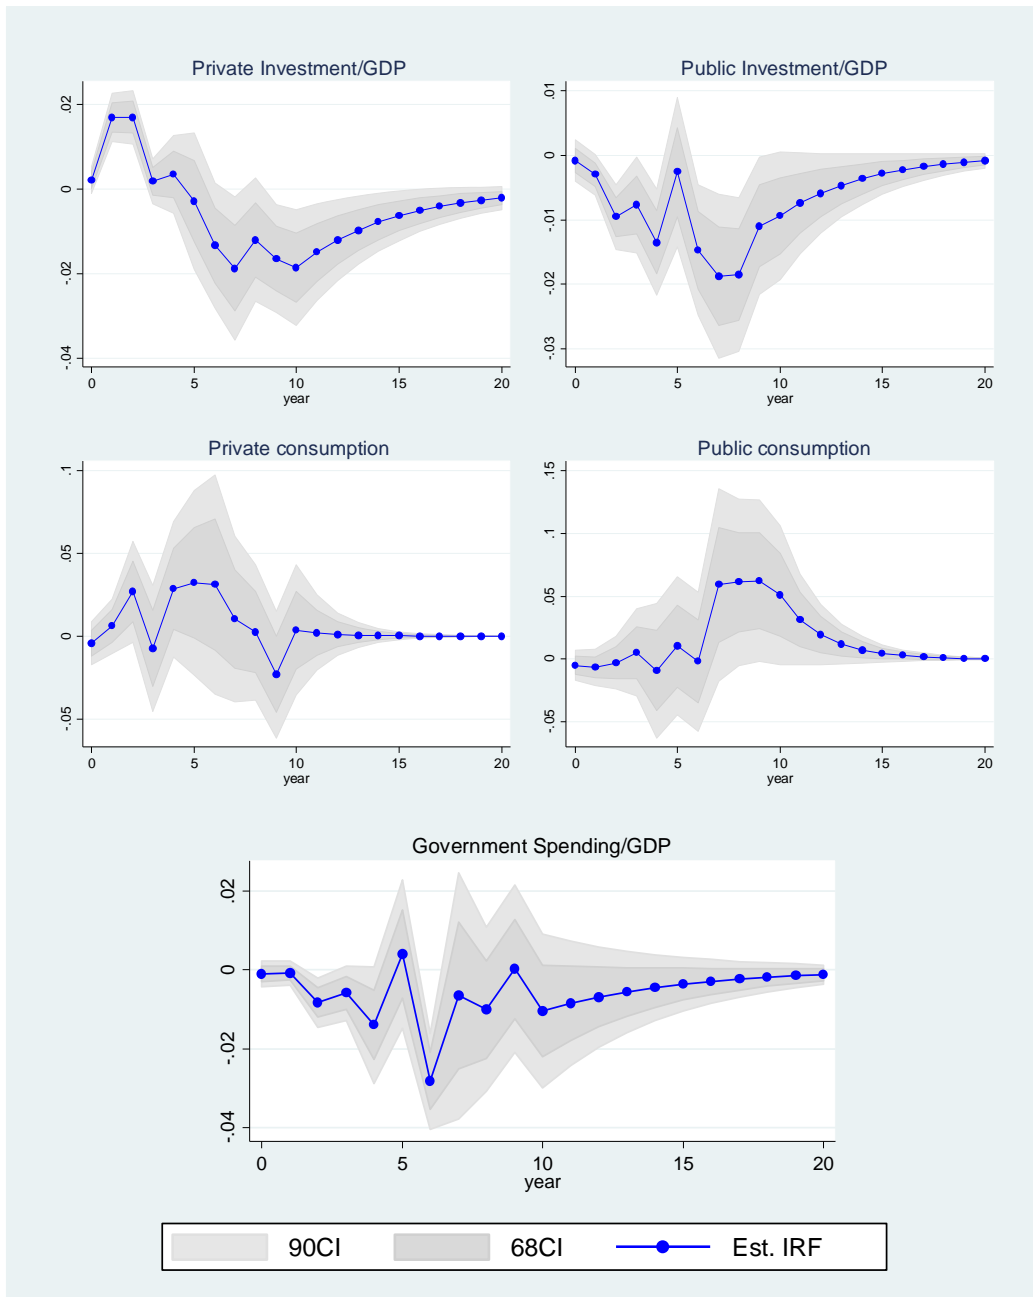


Figure X: Private versus Public Responses

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.