

The Future of Oil: Geology versus Technology

Jaromir Benes, International Monetary Fund

Marcelle Chauvet, University of California, Riverside

Ondra Kamenik, Consultant

Michael Kumhof, International Monetary Fund

Douglas Laxton, International Monetary Fund

Susanna Mursula, International Monetary Fund

Jack Selody, Consultant

June 4, 2012

Abstract

We discuss and reconcile the geological and the economic/technological views concerning the future of world oil production and prices. We present a nonlinear econometric model of the world oil market that encompasses both views. The model performs far better than existing empirical models in forecasting oil prices and oil output out of sample. Its point forecast is for a near doubling of the real price of oil over the coming decade. The error bands are wide, and reflect sharply differing judgments on ultimately recoverable reserves, and on future price elasticities of oil demand and supply.

JEL Classification: C11, C53, Q31, Q32.

Keywords: Exhaustible resources; fossil fuels; oil depletion; Hubbert's Peak; Bayesian econometrics.

The views expressed here are those of the authors, and do not necessarily reflect the position of the International Monetary Fund or any other institution with which the authors are affiliated. The authors thank Serhat Solmaz and Kadir Tanyeri for excellent research assistance.

1 Introduction

Future oil prices have been notoriously difficult to predict. In a recent paper, Alquist, Kilian, and Vigfusson (2011) conclude that forecasts based on monthly futures prices, monthly surveys of forecasts, simple econometric models, or other commonly employed forecasting techniques cannot consistently beat a random-walk forecast out of sample. This result is well known within the oil industry.

The simple econometric models used by Alquist, Kilian, and Vigfusson (2011) emphasize macroeconomic indicators as predictors of future oil prices. These indicators are highly correlated with fluctuations in aggregate demand, and will therefore mainly capture changes in the price of oil caused by variations in demand. But they are unlikely to be effective in capturing temporary oil supply disruptions. Moreover, since aggregate demand tends to revert to a trend, these variables are not likely to be successful in predicting a long-lasting increase in the price of oil such as the one we have recently observed. However, there is an alternative explanation for the recent persistent price movements that, despite considerable evidence in its support, has received very little attention in the economics literature. This is that one key driver of recent events may have been a highly persistent or even permanent shock to oil supply that is due to geological limits on the oil industry's ability to maintain the historical growth rate of production. The extent to which the literature discounts or embraces this possibility is critical for its interpretation of recent events in the oil markets.

Kilian (2009), in analyzing the U.S. economy, distinguishes between three drivers of oil prices, aggregate demand for goods, precautionary demand for oil, and supply of oil, where the latter captures only the possibility of temporary supply disruptions due to political events in oil producers, the dominant supply shock in historical data. He finds that the two demand shocks have been far more important as drivers of oil prices, while supply shocks have had a negligible effect. Kilian's (2009) analysis does not allow for the possibility of highly persistent shocks to the supply of oil that are driven by terminal geological limits.

Hamilton (2009), on the other hand, finds that temporary disruptions in physical oil supply have already had a major role in explaining historical dynamics of oil price movements. Furthermore, he argues that stagnating world oil production, meaning a very persistent reduction in oil supply growth, may have been one of the reasons for the run-up in oil prices in 2007-08. The main reasons why oil supply shocks affect output according to Hamilton is their disruptive effect on key industries such as automotive manufacturing, and their effect on consumers' disposable incomes. In other words, the main effect is on aggregate demand. As for aggregate supply effects, there may be large short-run impacts due to very low short-run elasticities of substitution between oil and other factors of production. Such elasticities of substitution would tend to get larger over longer horizons, as agents find possibilities to substitute away from oil, because high prices start to stimulate technological change that can both increase the recovery of oil, and the availability of substitutes for oil. But Hamilton (2012) argues that the main reason for the historic growth in oil output has been exploration of new geographic areas rather than application of better technology to existing sources, and that the end of that era could come soon. The paper goes on to explore the potentially very problematic implications of slower future oil supply growth for future GDP growth.

Proponents of the geological view of future oil production suggest that oil reserves are ultimately finite, easy-to-access oil is produced first, and therefore oil must become harder and more expensive to produce as the cumulated amount of oil already produced grows. According to many scientists in this group, the recently observed stagnant oil production in the face of persistent and large oil price increases is a sign that physical scarcity of oil is already here, or at least imminent, and that it must eventually overwhelm the stimulative effects of higher prices. Furthermore they state, on the basis of extensive studies of alternative technologies and resources, that suitable substitutes for oil simply do not exist on the required scale, and that technologies to improve oil recovery must eventually run into limits dictated by the laws of thermodynamics, specifically entropy. This view of oil supply

traces its origins back to the work of M. King Hubbert (1956), a geoscientist who in 1956 correctly predicted that U.S. oil output would peak in 1970. It is discussed in a study for the U.S. Department of Energy¹, Hirsch et al. (2005), and in a subsequent book, Hirsch et al. (2010). The most thorough research available on this topic is UK Energy Research Centre (2009), which is succinctly summarized in Sorrell et al. (2010). Based on a wealth of geological and engineering evidence, these authors conclude that there is a significant risk of a peak in conventional oil production before 2020, with an inexorable decline thereafter.

In this paper we find that our ability to forecast future developments in the oil market, and by implication in aggregate activity, can be dramatically improved by combining the geological and economic/technological views of oil supply, and by estimating their respective contributions. We develop a simple macroeconomic model that combines a conventional linear demand specification with a nonlinear supply equation, the latter combining a mathematical formalization of the geological view with a conventional price sensitive oil production. We find that this model can predict oil prices far better out of sample than a random walk, and that it can predict oil production far better than the historical track record of official energy agencies on the one hand, and of advocates of pure versions of the geological view on the other hand. We also use the model to identify which driving force has been most responsible for the recent run-up in oil prices. We find that the geological, price-insensitive component of supply is the key reason for the recent accuracy of the model's predictions because it captures the underlying trend in prices. But we also find that shocks to excess demand for goods and to demand for oil, the latter probably due to phenomenal recent growth in China and India, have been key to explaining persistent and sizeable deviations from that trend. These deviations work through the price channel. Looking into the future, both of these factors continue to be important, and point to a near doubling of real oil prices over the coming decade. But there is substantial uncertainty about these future trends that

¹ Other studies by official U.S. agencies that have warned about this issue include Government Accountability Office (2007) and United States Joint Forces Command (2010).

are rooted in our fundamental lack of knowledge, based on current data, about ultimately recoverable oil reserves, and about long-run price elasticities of oil demand and supply.

The rest of the paper is organized as follows. Section II presents data on historical oil supply forecasts by proponents of the technological and geological views. Section III presents and discusses the model specification and parameter estimates. Section IV presents a detailed analysis of the estimation results. Section V concludes.

2 Historical Forecasts of World Oil Production

The complicated dynamics of world oil supply and oil demand makes oil production forecasting very difficult. Figure 1 shows the track record of the U.S. Energy Information Administration (EIA). Strikingly, between 2001 and 2010 their forecasts have exhibited an almost continuous decline, with the forecast for 2020 declining by over 20%, or by 25 million barrels per day. Earlier EIA forecasts were based on the simple notion that supply would be available to satisfy demand, so that these forecasts essentially only considered the drivers of demand. This turned out to be far too optimistic, and more recent forecasts may be starting to reflect the recognition that constraints on oil supply are starting to influence production and prices.

The reason why this may be the case is illustrated in Figure 2, which displays real world oil prices in 2011 U.S. dollars² alongside OPEC spare capacity in millions of barrels per day (Mbd). Until the end of 2002 spare capacity had been high in historical terms, and this was accompanied by oil prices that had not been growing significantly in real terms. But this changed abruptly in early 2003, around the time of the Iraq war, when spare capacity dropped below the 2 Mbd mark, which by many in the industry is considered the critical mark where supply becomes a constraining factor. From that moment until the onset of the

² The figure is normalized so that the real oil price in 2011 equals 104. This makes the units intuitive, given that the average 2011 nominal oil price equalled US\$ 104. The same normalization is adopted in all subsequent charts of the real oil price.

Great Recession real oil prices started a long-term increase that ultimately saw them more than triple, before the demand destruction of the Great Recession led to a sudden increase in spare capacity and a steep decline in oil prices. This however only brought temporary relief to the demand-supply balance in the oil market, for two reasons. First, as we have seen in Figure 1, oil production never regained its historical growth rate of 1.5%-2% per annum after 2005, and has in fact been on what looks like a plateau ever since that time. And second, partial recoveries in many economies restarted demand from 2009 onwards. Spare capacity is therefore again approaching 2 Mbd, and oil prices are ratcheting up again. The combination of a plateau in actual oil production, and of repeated pressure on spare capacity except at a time of deep recession, indicate that physical constraints on oil production are starting to have an increasing impact on prices.

Proponents of the geological view of oil production have a track record that can be compared to that of the EIA. Figure 3 shows the track record of Colin Campbell, a former oil geologist who has become one of the most influential proponents of the geological view. The one caveat in such a comparison is that different agencies and individuals produce forecasts for different aggregates of oil production. While for the EIA we showed the forecasts for world total oil supply, which is defined as crude oil, plus NGL and other liquids, plus refinery processing gains, for Campbell we have historical forecasts for regular conventional oil. This definition covers over 75% of world total oil production. It is based on EIA data but excludes heavy oil (<17.5 deg API), bitumen, oil shale, shale oil, deepwater oil and gas (> 500m), polar oil and gas, and NGL from gas plants. Furthermore, the International Energy Agency (IEA) uses yet another definition that is slightly less encompassing and therefore smaller than the EIA's, but more encompassing than Campbell's, namely crude oil plus NGL. We will use IEA data in our empirical analysis, but have used EIA data for Figure 1, because the EIA produces annual forecasts while the IEA does not.

Figure 2 shows that Campbell's forecasts have also erred, but this time on the pessimistic side. The differences to ex-post realized production data are somewhat smaller than those

of the EIA, whose 2001 estimate for 2010 overestimates actual production by 8.7 Mbd, compared to a 2003 underestimate by Campbell of 4.5 Mbd.

Campbell's methodology is based on an extremely detailed knowledge, country by country, of production and exploration data that goes back to his participation in the construction of an industry database in the early 1990s. Another methodology that is used by proponents of the geological view is curve fitting for world oil production.³ As this yields econometrically testable equations for the production profile, we will pursue this in some detail in this paper. A particularly tractable specification is known as Hubbert linearization. This is based on Deffeyes (2005), who develops a much simplified version of the analysis in Hubbert (1982). We adopt the notation that q_t represents annual oil production at time t , Q_t represents cumulative production until time t , and \bar{Q} represents ultimately recoverable reserves, or cumulative production by the time the last oil well in the world runs dry. Then Hubbert states that annual production can be usefully approximated by the logistic curve

$$q_t = \alpha_s Q_t \left(\frac{\bar{Q} - Q_t}{\bar{Q}} \right) . \quad (1)$$

This is a bell-shaped curve, and it states that in any given year actual production is determined by the cumulative production that has already taken place, and by the fraction of oil that remains to be produced. The latter dominates exactly from the point where half of all oil has been produced, $Q_t = \bar{Q}/2$. At that point annual oil production peaks, and subsequent production starts to decline. This logistic function can be transformed by dividing (1) by Q_t , which produces a linear relationship between cumulative production and the ratio of current versus cumulative production:

$$\frac{q_t}{Q_t} = \alpha_s - \frac{\alpha_s}{\bar{Q}} Q_t . \quad (2)$$

Given that for econometric purposes both α_s and \bar{Q} are unknowns, this can be written as

$$\frac{q_t}{Q_t} = \alpha_s - \beta Q_t . \quad (3)$$

³ See UK Energy Research Council (2009) for a very detailed discussion.

Deffeyes, a Princeton professor of geology, finds that this relationship fits both U.S. and world data very well until 2003, the last datapoint in his 2005 study, with both series being very close to a straight line relationship for the period 1983-2003. His fit of the data indicates a logistic curve with a peak in late 2005, and a decline in world oil production thereafter. Deffeyes responds to the economic/technological view, that higher prices should spur additional technological development and hence production that might delay the peak, by stating that “improved technologies and incentives have been appearing all along, and there seems to be no dramatic improvement that will put an immediate bend in the straight line”. As we show in the top half of Figure 4, this prediction was not borne out by subsequent events, as significant positive deviations from Deffeyes’ straight line started to appear immediately after 2003. As we have seen in Figure 2, the critical feature of post-2003 data that can account for this development is that oil prices started to increase to much higher levels than at any point during 1983-2003. This appears to have significantly spurred production relative to what it might otherwise have been, so that production did not peak in late 2005. In other words, prices did matter. However, and this is critical, production did not increase either from that point onwards, it rather reached a plateau, where it has, with some fluctuations, remained until the present day. In other words, prices did not matter enough to allow production to regain its historical growth rate.

In summary, we observe that both the advocates of the economic/technological view and the advocates of the geological view have had to significantly revise their projections over the last decade, the former downwards and the latter upwards. There does seem to be a tendency for both sets of views to eventually converge, but the differences in forecasts are at this moment still large, and improvements in forecast accuracy would greatly assist an informed debate. We believe the foregoing illustrates very clearly that what is needed is an analytical and empirical approach that allows for both views in an integrated framework. This is what the rest of this paper is designed to do.

3 The Model

In this section we present our econometric model of the world oil market, and comment on parameter estimates for the key coefficients. The model is kept as simple as possible, and consists only of a conventional equation for world oil demand, an equation for world oil supply that combines the geological and economic/technological views, and a set of conventional trend and gap equations for the determination of world GDP.

We estimate this system of equations using data for world real GDP (IMF data), the real quantity of oil produced (IEA definition and data), and the real oil price (U.S. CPI based). We use annual data from 1983 through 2011, with lags that use data back to 1972 for oil prices. The model has multiple factors that drive oil price and output dynamics in a fairly short sample, which can potentially lead to difficulties in obtaining sensible parameter estimates. To overcome this problem we employ nonlinear Bayesian estimation techniques, using priors based on other studies. Nonlinear techniques need to be used because the world oil supply equation is an augmented version of the nonlinear Hubbert linearization specification in (3). The model is estimated as a system. Because the model contains an endogenous breakdown of output into trend and gap components, we impose a diagonal covariance matrix as an identifying assumption. A summary of the model's key parameters, including their distributions, prior and posterior modes, and 90% confidence intervals, is shown in Table 1. Posterior modes are also displayed underneath the parameter symbols in the displayed model equations below.

3.1 Oil Supply

The oil supply equation combines the geological view embodied in the Hubbert linearization equation (3), whereby oil is more and more difficult to extract as cumulative production increases, with the economic/technological view of a standard supply curve, whereby production responds positively to current and past oil prices p_t . The short-run

effects of oil prices on production arise to the extent that producers can and want to speed up production from existing fields.⁴ In other words, they utilize existing spare capacity. Over the medium run additional price effects can arise as high prices lead to new exploration and/or better technologies, but these projects tend to have lead times of at least four years. We therefore introduce an additional response of production to real oil prices lagged between four and six years. For estimation purposes, the units of cumulative production Q_t are in gigabarrels, while the units of annual production q_t are in hundreds of thousands of barrels. The supply equation is

$$\frac{q_t}{Q_t} = \alpha_s - \beta_1 Q_t + \beta_2 p_t + \beta_3 \frac{1}{3} \sum_{k=4}^6 p_{t-k}, \quad (4)$$

(507.7) (0.243) (0.624) (0.056)

with the auxiliary relationship

$$Q_t = Q_{t-1} + q_t/10000. \quad (5)$$

The parameter $\alpha_s < 1$ indicates the speed at which oil production increases in the early years, before depleted reserves constrain growth, and the parameter $\beta_1 > 0$ indicates the effect of depleted reserves on production. The parameters $\beta_2 > 0$ and $\beta_3 > 0$ indicate that the production of oil increases with the current and lagged prices of oil.

Our prior for the coefficient β_1 was taken from the Deffeyes (2005) study of peak oil. It is given a fairly loose uniform distribution. The priors for β_2 and β_3 were also given a uniform distribution, and not set tightly. The reason is that our knowledge about the oil supply response to price increases is limited, as most estimated economic models focus only on demand elasticities.

The estimated coefficient $\beta_1 = 0.243$, which is slightly lower than the prior, supports a role for the geological channel advocated by Deffeyes (2005), as values much closer to zero, which would have minimized the importance of that channel, were not ruled out by our

⁴ This involves an important technical consideration: Excessively fast extraction of oil from an existing field can destroy geological structures and reduce the ultimately recoverable quantity of oil. See Simmons (2005).

loose prior. The coefficients β_2 and β_3 can be converted to price elasticities of supply⁵, but given the levels specification of (4) these elasticities depend on actual oil production and, especially, oil prices. We find that, during the pre-2003 period of relatively low oil prices, the elasticity with respect to current prices, computed from β_2 , was around 0.05, while the elasticity with respect to lagged prices, computed from β_3 , was well below 0.01. During the most recent period these values increased to around 0.15 and 0.02, respectively. Whether price elasticities of this magnitude can be maintained for the foreseeable future is a critical question that determines the outlook for future output and prices. Our forecasts will show upper and lower bands, and also some sensitivity analysis, that indicate what is at stake. Most importantly, the fact that the main output response to prices has been contemporaneous may be a reason for concern, because this indicates that output has mainly been able to respond to high prices by producers immediately dipping into spare capacity, rather than by increasing exploration or improving technology to increase longer-run capacity. To the extent that the future may be characterized by much tighter supply constraints and therefore much lower spare capacity, this option may no longer be available to the same extent as in the past.

The effect of $\beta_2 > 0$ and $\beta_3 > 0$ is to flatten the line of the Hubbert linearization, and to shift it upward, as oil prices embark on their upward trend. The effect is to delay and raise the peak of oil production, and perhaps also to delay the point at which $q_t = 0$. For example, estimation of the curve, with β_2 and β_3 set to zero, over the period 1983-2003, when oil prices were relatively low and steady on average, produces estimates that generate a steeply downward sloping line. Extending the sample period to 1983-2010 and allowing for $\beta_2 > 0$ and $\beta_3 > 0$, to include data points with higher oil prices that raise the average price of oil over the sample, raises and flattens the curve. But this does not remove the tendency for oil production to eventually decline, unless real oil prices were to keep rising steeply and indefinitely.

⁵ The units of the coefficients are of course affected by the fact that in our data q_t and Q_t are expressed in different units.

3.2 Oil Demand

Oil demand is determined by the standard view that a combination of economic activity (GDP) and oil prices drives world oil demand. Higher economic activity increases the demand for oil since production requires oil as an input, and higher oil prices reduce the demand for oil by raising the incentive to substitute away from oil. The price elasticity is expected to be small in the short run, but it may rise in the long run as substitution takes place. For example, the stock of cars turns over very slowly, over more than a decade.⁶ We therefore include both current oil prices and a 10-year moving average of oil prices in our explanatory variables. The demand equation is estimated in differences. We have

$$\Delta \ln q_t = \underset{(-0.018)}{\alpha_d} + \underset{(0.910)}{\gamma_1} \Delta \ln gdp_t - \underset{(0.021)}{\gamma_2} \ln \frac{p_t}{p_{t-1}} - \underset{(0.06)}{\gamma_3} \left(\ln \frac{p_{t-1}}{p_{t-10}} / 9 \right). \quad (6)$$

The prior for γ_1 was set to reflect the tight relationship between GDP and oil demand that has been found in numerous previous studies, including a recent analysis in the April 2011 IMF World Economic Outlook (IMF (2011)). The distribution is also set tightly to reflect the robustness of this link in the literature. The prior distributions for γ_2 and γ_3 are also set tightly, reflecting considerable consensus about these values in the literature. The prior modes are set so that the short-run elasticity of demand is less than the long-run elasticity. We also allow for the possibility that γ_2 and γ_3 may be up to 2.5 times larger at very high oil prices, because such prices would dramatically increase the incentives to substitute away from oil.⁷ Specifically, at the average oil prices seen prior to 2008 elasticities are unaffected, at the average prices of 2008 and 2011 elasticities rise by roughly a factor of 1.75, and at the much higher prices projected by the model out to 2021 elasticities eventually rise by a factor of maximally 2.5.

The estimate for the income elasticity of oil demand γ_1 is consistent with other studies, which have found that industrialized countries on average display a lower income elasticity

⁶ There are grounds for doubt as to whether long-run elasticities can continue to be much higher than short-run elasticities. See the discussion in Section IV.E.

⁷ To keep the exposition simple this is not shown in (6).

around 0.5, reflecting a less oil-intensive and more service-intensive production structure, while many key emerging markets, which have been the main drivers of recent world economic growth, display income elasticities of around 1. The estimated price elasticities of demand are in line with the estimates reported in IMF (2011), with a very low short-run elasticity of 0.02 and a long-run elasticity (after 10 years) of 0.08. The combination of low price elasticities of supply and demand implies that any reduction in available supply, or even inadequate growth of supply relative to past trends, must lead to either much higher oil prices or an economic contraction, or a combination of the two.

3.3 GDP Equations

The feedback from oil prices to economic activity is captured by a reduced-form production function that allows us to separately specify shocks to the output gap (transitory shocks to output), shocks to potential output (permanent shocks to the level of output), and shocks to potential output growth (permanent shocks to the growth rate of output). The richness of this specification helps us to model the complicated interactions of oil price movements and GDP, where both trend and gap decline if oil prices increase. However, there is not enough variation in the historical data to provide well-determined estimates of these separate effects based on a single observed variable. One advantage of adopting Bayesian estimation techniques is that we can adopt reasonable and tightly set priors that help with the identification of these three different shocks to output. GDP is given by

$$gdp_t = pot_t * y_t , \quad (7)$$

where pot_t is potential output and y_t is the output gap.

3.3.1 Potential Level of GDP

Potential GDP is given by

$$\Delta \ln pot_t = \ln g_t + \epsilon_t^{pot} , \quad (8)$$

where ϵ_t^{pot} is a shock to the level of potential output and g_t is the growth rate of potential

output. This states that the level of potential fluctuates around its trend path. Oil prices do not enter this equation, since we assume that the dynamic effects of oil prices on potential output are captured in the potential growth rate equation.

3.3.2 Potential Growth Rate of GDP

The growth rate of potential world GDP is specified as fluctuating around an exogenous long-run trend, with oil price changes making the fluctuations more severe. Oil prices are allowed to have persistent but not permanent effects on the growth rate of GDP. We have

$$\ln g_t = \underset{(0.899)}{\lambda_1} \ln g_{t-1} + (1 - \lambda_1) \underset{(0.04)}{g} - \underset{(0.005)}{\lambda_2} \left(\Delta \ln p_t - \underset{(0.07)}{\rho} \right) - \underset{(0.005)}{\lambda_3} \left(\Delta \ln p_{t-1} - \underset{(0.07)}{\rho} \right) + \epsilon_t^g, \quad (9)$$

where ϵ_t^g is a shock to the growth rate of potential output, g is the average or steady state growth rate of GDP, and ρ is the average growth rate of real oil prices. The estimated steady state world annual growth rate of potential GDP equals four percent. The average annual growth rate of real oil prices, which is the growth in oil prices at which the model assumes zero effects of oil prices on output growth, is seven percent. The results indicate that an oil price growth rate that is higher than that historical average has a small but significant negative effect on the growth rate of potential. Both exogenous shocks ϵ_t^g and oil price fluctuations cause the growth rate to deviate quite persistently from its long-run value, given that the estimated coefficient on the lagged growth rate equals 0.9.

3.3.3 Output Gap

Apart from allowing for an effect of higher oil prices on the growth rate of potential output, the model also allows for the possibility that higher oil prices can cause fluctuations in the amount of excess demand in the economy. As is standard with equations of this type, fluctuations in the output gap are modeled as persistent. Specifically, we specify the process

$$\Delta \ln y_t = \left(\underset{(0.956)}{\phi_1} - 1 \right) \ln y_{t-1} + \underset{(0.257)}{\phi_2} \Delta \ln y_{t-1} - \underset{(0.005)}{\phi_3} \left(\Delta \ln p_t - \underset{(0.07)}{\rho} \right) - \underset{(0.005)}{\phi_4} \left(\Delta \ln p_{t-1} - \underset{(0.07)}{\rho} \right) + \epsilon_t^y, \quad (10)$$

where ϵ_t^y represents a shock to the level of aggregate demand. Similar to the equation for potential, the coefficient estimates show that higher oil prices have a small but significant negative effect on excess demand, and that this effect is highly persistent.

4 Analysis

We now study the estimation results in more detail, by analyzing the implications of the already discussed parameter estimates for the model's impulse response functions, for its interpretation of history, for forecast accuracy, and for current forecasts of oil output, oil prices and GDP.

4.1 Impulse Response Functions

Figure 5 shows the impulse response functions of the model, with three columns for the responses of oil production, the real price of oil and GDP, and five rows for the five shocks, oil supply shocks, oil demand shocks, output gap shocks, potential growth shocks, and potential level shocks. All impulse responses are shown in percent deviations from control, after removing an underlying trend.

Oil supply shocks occur separately from, and in addition to, the geological tightening effects of Hubbert's curve in equation (4). We find that, relative to oil demand shocks and output gap shocks, such shocks have been comparatively small and transitory in the recent data, and consequently their effects on real oil prices have been transitory as well, although the implied upward spikes in real oil prices when these shocks did occur have been significant. The top row of Figure 5 shows that a negative oil supply shock creates a five year cycle in which output is below potential, and where the contraction in GDP is about half as large as the contraction in oil supply. Due to very low short-run demand and supply elasticities, oil prices increase dramatically in the short run, by more than 30 times the magnitude of the supply contraction, but they subsequently quickly return to trend.

Oil demand shocks have been significantly larger in size, and have been a major

contributor to high oil prices especially in the period prior to the Great Recession, and in the recent partial recovery from that recession. Oil demand shocks have also had much more persistent effects on oil production and GDP than oil supply shocks. Their effect on the real price of oil has not been as sharp, but again more persistent.

The main shocks that explain the behavior of oil prices during the crisis are output gap shocks, which are illustrated in the third row of Figure 5. Estimated output gap shocks have very large and persistent effects on GDP that lead to similarly large and persistent effects on oil demand. Of course the dominant output gap shock during the crisis has been a negative shock that reduced economic activity and oil demand. The resulting large effect on the oil price is a major part of the model's explanation for the steep drop in oil prices following the onset of the Great Recession.

The impulse responses for potential growth rate shocks are illustrated in the fourth row of Figure 5. These shocks are smaller in size than output gap shocks, but they have much more persistent effects on output and oil production. Their effects on the real price of oil are less dramatic, because these shocks only lead to a gradual increase in oil demand, so that low short-run price elasticities of demand and supply do not play a significant role.

Finally, potential level shocks do not contribute much to overall variability in the model. When they do occur, the effects on output, oil demand and oil prices are of course highly persistent.

4.2 Interpretation of History

Figure 6 shows the estimated shocks of the model. Figures 7 and 8 show model simulations that decompose the post-2002 movements in oil prices and oil output into the contributions of the three shocks that account for most of the variability in the model. The model simulation without further shocks is in each case represented by the broken line. The top left simulation compares this to the model simulation with all shocks (solid line), where the latter is by construction identical with the data. The remaining simulations show the

separate contributions of the estimated shocks to oil demand, oil supply and the output gap (solid lines).

We begin with Figure 7, the decomposition of oil prices. By 2008 oil prices had reached a level that was 60% higher than what the model would have predicted on the basis of 2002 information. The major contributing factors in the earlier years were very strong oil demand, principally from booming emerging economies, and a positive world output gap. Oil supply, at least until some time in 2005, actually helped to, *ceteris paribus*, keep oil prices lower than what they would otherwise have been. From that time onward however, as we have seen, world oil production stayed on a plateau, and by 2008 insufficient world oil supply had become the major factor behind high oil prices. The Great Recession, from 2009, was so severe that oil prices dropped below the original 2002 forecast. The model attributes roughly half of this drop to a negative output gap shock, and the other half to a positive oil supply shock. The latter is the model's interpretation of the increase in oil excess capacity in 2009. By 2011 real oil prices had regained their 2008 average (not peak) levels. The model attributes almost all of this to negative oil supply shocks, with oil demand and output gap shocks showing no major trend reversal after 2008. In other words, the insufficient growth of world oil supply that had begun to assert itself between 2005 and 2008 returned to center stage, as production remained on the same approximate plateau that it had reached in late 2005. It is very important, and evident from Figure 7, that it is not the shocks that are the major driving force behind the trend increase in oil prices in our model. Rather, the no-shocks scenario predicts an increase in oil prices that is not far from the actual trend.⁸ The reason is the significant estimate of the Hubbert linearization coefficient β_1 in the oil supply curve. This confirms that the problem of oil becoming harder and harder to produce in sufficient quantities was an important factor that would have significantly increased oil prices regardless of shocks.

⁸ The actual trend does show a positive deviation from the no-shocks scenario. The main reason is unexpectedly strong demand from emerging economies post-2002.

Figure 8 decomposes oil production. We observe that production was, except for 2009, consistently and sometimes significantly above the trend predicted by the model in 2002. However, oil supply shocks only made a minor contribution to this development, with the major driving forces coming from booming oil demand and, from 2006 through 2008, positive output gaps. Because both of these shocks lead to higher oil prices, the price mechanism that we added to Deffeyes' (2005) Hubbert linearization specification is key to being able to account for the post-2003 deviations from the pure geological explanation of oil production and prices. But it is of course this geological explanation that is able to account for the strong underlying trends in the model, especially the upward trend in oil prices.

4.3 Relative Forecast Performance

Figure 9 shows our model's out-of-sample rolling forecasts, from 2001 through 2011, for oil production, oil prices and the growth rate of real GDP. The figure shows only the point forecasts, we will discuss error bands in the next subsection.

The predicted average annual growth rates of oil output are well below the historical forecasts of the EIA, but above the forecasts by proponents of the geological view. We therefore find that our model's accommodation of both the geological and the economic/technological views leads to estimation results that provide partial support for both, while rejecting pure versions of either. This is not unexpected, given our discussion of recent trends in oil output (plateau since 2005) and in spare capacity on the one hand, and of the clear effects of prices in overturning the pure Deffeyes (2005) model.

However, this projected positive trend in oil production comes at a steep cost, because the model finds that it requires a large increase in the real price of oil, which would have to nearly double over the coming decade to maintain an output expansion that is modest in historical terms. Such prices would far exceed even the highest prices seen in 2008, which according to Hamilton (2009) may have played an important role in driving the world economy into a deep recession.

This negative GDP effect of higher oil prices is present in the model's forecasts for GDP growth, but as we will see it is modest. This raises the question of whether future versions of the model should include nonlinearities in the output response similar to the nonlinearities in our oil demand equation. There is likely to be a critical range of oil prices where the GDP effects of any further increases become much larger than at lower levels, if only because they start to threaten the viability of entire industries such as airlines and long-distance tourism. If this is correct, the effect of real oil prices on GDP should be modeled as convex. There is support for this conjecture among the experts. For example, the chief economist of the International Energy Agency, Fatih Birol, has repeatedly warned in recent months that current high oil prices, which are nearly back to their levels in 2008, are at a point that could push the world economy back into recession.⁹ We will study this possibility quantitatively in future work.

Figure 9 shows that our model predicts neither a mean-reverting oil price, as do most empirical models of the oil market, nor even a random walk, which has been shown to outperform such models in many studies. Rather it predicts a clear upward trend, which is exactly what we have been observing in the data, with the exception of the demand destruction of the Great Recession. Furthermore, our model's out-of-sample oil output predictions in the early 2000s turned out to be far more accurate than either the contemporaneous EIA forecasts or the forecasts using the Deffeyes or Campbell methods. To formalize these comparisons of forecast accuracy, Table 2 shows the root mean square errors (RMSEs) of our model for the period 2003-2011, and compares the forecasts for the level of oil production to the EIA's forecasts, the forecasts for the level of oil prices to a random walk, and the forecasts for the level of world GDP to those of contemporaneous editions of the IMF's World Economic Outlook (WEO). For production, our RMSEs are lower than those of the EIA's historical forecasts at all but the one-year horizon, and less than half as large

⁹ See the IEA website at <http://www.worldenergyoutlook.org/quotes.asp> for a collection of Birol's recent quotes on this subject.

at longer horizons. For prices the gains from using our model are even larger, especially at longer horizons. For example, at the five-year horizon our model's RMSE is about a quarter of the RMSE of a random walk. Against the background of the discussion in Alquist, Kilian, and Vigfusson (2011), these results are dramatic. For GDP the gains are less dramatic but nevertheless very substantial.¹⁰

4.4 Current Forecasts

Figures 10, 11 and 12 show the model's current projections, for the decade from 2012 through 2021, for oil production, oil prices and GDP. The figures contain a point forecast and error bands around this forecast. They also show an alternative scenario that assumes a tighter future oil supply due to a lower future elasticity of oil supply with respect to contemporaneous oil prices. We will comment on this scenario at the end of this subsection.

Figure 10 shows oil production. The point forecast is for a mean annual growth rate of oil output of around 0.9% over the coming decade, positive but well below its historical growth rate of around 1.5%-2.0%. The 90% confidence interval is very wide, and reflects high levels of uncertainty concerning ultimately recoverable reserves (implicit in β_1) as well as supply and demand elasticities with respect to the oil price. The lower 90% band indicates flat oil output for the entire decade, while the upper band indicates annual output growth rates that are almost as large as historical ones. It is important to observe that, while the point forecast is for an annual growth rate approximately as large as the most recent EIA forecasts, the forecast for the oil price that is behind this output forecast is far higher than anticipated by the EIA.

This is shown in Figure 11, which shows a point forecast that implies a near doubling of real oil prices over the coming decade, and an increase in prices over and above the very high recent levels even under a very optimistic scenario, at the lower 90 percent confidence

¹⁰ We will not emphasize the RMSE differences for GDP further in this paper, partly because this result may have less to do with our modeling of the oil sector and more with our modeling of the different component processes of output.

interval. The world economy has never experienced oil prices this high for anything but short transitory periods, and we reiterate our previous statement that this might take us into uncharted territory, where a nonlinear, convex effect of oil prices on output might be a more prudent assumption.

Figure 12 shows forecasts for GDP, with 2011 world real GDP normalized to one. The point forecast is for a roughly 4% per annum real GDP growth rate. The error bands may appear narrow relative to those for oil prices and oil output, but the 90% confidence interval nevertheless contains average growth rates as low as 3% per annum, and as high as 5% per annum. In other words, at more pessimistic coefficient values for ultimately recoverable reserves and elasticities, average world growth would be one percentage point lower.

Finally, Figures 10-12 also report the point forecast for an alternative scenario where β_2 takes the lower value corresponding to its lower 90% confidence band. The baseline value for β_2 was estimated over a period when, at most times, it was possible for producers to respond to high prices by immediately utilizing ample spare capacity, an option that may not be available to the same extent in a future of tighter supply constraints. We find that the lower value for β_2 has very large effects on the results, even though β_2 only drops fairly modestly, from 0.624 to 0.505. Average oil output growth drops from 0.9% to 0.5% per annum, the oil price now fully doubles by 2021, and the path for GDP is approximately equal to the lower 90% confidence band. This last result implies that this one change alone reduces the point forecast for average world output growth by approximately 1 percentage point.

4.5 Oil and Output - Open Questions

Our data and analysis suggest that there is at least a possibility that we may be at a turning point for world oil output and prices. A key concern going forward is that the relationship between higher oil prices and GDP may become nonlinear if oil prices become sufficiently high. The problem is that, at this moment, historical data contain very little information about what that relationship might look like. But we are not entirely without information,

because a number of authors in other sciences have started to ask pertinent questions, and have done some early pioneering work.

There are two key questions, under the maintained hypothesis of much lower oil output growth. First, what is the importance of the availability of oil inputs for continued overall GDP growth? Second, what is the substitutability between oil and other factors of production? We emphasize that these concerns focus not on the demand side but rather on the supply side effects that could result from stagnating or declining world oil production.

As for the contribution of oil to GDP, the main problem is that conventional production functions imply an equality of cost shares and output contributions of oil, which for a long time has led economists to conclude that, given its historically low cost share of around 3.5% for the U.S. economy¹¹, oil can never account for a massive output contraction, even with low elasticities of substitution between oil and other factors of production. This view has been challenged in several recent articles and books by natural scientists, who state that it need not hold with a more appropriate modeling of the aggregate technology. The contributions include Ayres and Warr (2005, 2010), Hall and Klitgaard (2011), Kümmel (2011), and Kümmel et al. (2002), who propose aggregate production functions that are based on concepts from engineering and thermodynamics. Several of these contributions estimate their production functions. The estimations are based on technologies that use energy, rather than more narrowly oil, but given the very limited substitutability between oil and other forms of energy this nevertheless offers important insights.¹² These authors find output contributions of energy of up to around 50%, despite the low cost share of energy. It is clear that if this can be confirmed in further rigorous econometric studies, the implications of lower oil output growth for GDP could be very large. This view is explored in oil shock simulations in the IMF's April 2011 World Economic Outlook (IMF (2011)), using the IMF's global DSGE model GIMF and a technology where oil's output contribution far

¹¹ See http://www.eia.gov/oiaf/economy/energy_price.html.

¹² For the U.S. economy the historical cost share of total energy has been around 7%.

exceeds its cost share. The simulations find that following permanent declines in the growth rate of world oil output, the model generates much larger negative output effects than the conventional neoclassical model, because a share of the stock of technology would become obsolete. This channel has never yet been of sufficient importance to explain the historical data, and our empirical model does not contain it. Changing this would lead to simulation results with lower GDP growth.

The other key concern going forward concerns elasticities of substitution. Several important contributions challenge economists' automatic assumption that elasticities of substitution between oil and other factors of production must be much higher in the long run than in the short run. The objections include that this assumption is not consistent with the historical facts (Smil (2010))¹³, with real-world practical limits (Ayres (2007)), or with the laws of thermodynamics, specifically entropy (Reynolds (2002), Ch. 10). Our empirical model presently makes the conventional assumption that elasticities will after some time be higher at higher prices. A plausible alternative that could reconcile the economists' view with the above objections is to assume that elasticities are very low in the short run (due to rigidities, adjustment costs, etc.), significantly higher in the medium run (as the rigidities are overcome), but much lower again in the long run if there is a sufficiently large shock to the growth rate of world oil supply, because there is a finite limit to the extent that machines (and labor) can substitute for energy. If this assumption was incorporated, the model would forecast significantly higher oil prices in the event of a sufficiently large and persistent shock to world oil supply.

¹³ This book describes the major energy transitions in world history, from biomass to coal, oil and nuclear energy. The critical observation is that all these transitions took many decades to complete, were enormously expensive and, crucially, happened at times when a new major energy resource of sufficient scale had already been clearly identified. The latter is clearly not the case today, as renewables are not even nearly of sufficient scale.

5 Conclusion

The main objective of this paper has been to propose and to empirically evaluate a model of the world oil market that does not take an a-priori view of the relative importance of binding resource constraints versus the price mechanism for world oil supply. We do not want to rule out either of these mechanisms, because the recent data tell a convincing story that both must have been important. Our empirical representation of this view models oil supply as a combination of the Hubbert linearization specification of Deffeyes (2005) and a price mechanism whereby higher oil prices increase oil output.

Our empirical results vindicate this choice. Our model performs far better than competing models in predicting either oil production or oil prices out of sample, in a field where predictability has historically been low. Our empirical results also indicate that, if the model's predictions continue to be as accurate as they have been over the last decade, the future will not be easy. While our model is not as pessimistic as the pure geological view, which typically holds that binding resource constraints will lead world oil production onto an inexorable downward trend in the very near future, our prediction of small further increases in world oil production comes at the expense of a near doubling, permanently, of real oil prices over the coming decade. This is uncharted territory for the world economy, which has never experienced such prices for more than a few months. Our current model of the effect of such prices on GDP is based on historical data, and indicates perceptible but small and transitory output effects. But we suspect that there must be a pain barrier, a level of oil prices above which the effects on GDP becomes nonlinear, convex. We also suspect that the assumption that technology is independent of the availability of fossil fuels may be inappropriate, so that a lack of availability of oil may have aspects of a negative technology shock. In that case the macroeconomic effects of binding resource constraints could be much larger, more persistent, and they would extend well beyond the oil sector. Studying these issues further will be a priority of our future research.

References

- Alquist, R., Kilian, L. and Vigfusson, R. (2011), "Forecasting the Price of Oil", Working Paper.
- Ayres, R. (2007), "On the Practical Limits to Substitution", *Ecological Economics*, **61**, 115-128.
- Ayres, R. and Warr, B. (2005), "Accounting for Growth: The Role of Physical Work", *Structural Change and Economic Dynamics*, **16**, 181-209.
- Ayres, R. and Warr, B. (2010), *The Economic Growth Engine - How Energy and Work Drive Material Prosperity*, Edward Elgar Publishing.
- Deffeyes, K. (2005), *Beyond Oil: The View from Hubbert's Peak*, Hill and Wang.
- Government Accountability Office (2007), "Crude Oil: Uncertainty about Future Oil Supply Makes It Important to Develop a Strategy for Addressing a Peak and Decline in World Oil Production", Report to Congressional Requesters.
- Hall, C. and Klitgaard, K. (2011), *Energy and the Wealth of Nations: Understanding the Biophysical Economy*, Springer Verlag (forthcoming, June 2011).
- Hamilton, J. (2009), "Causes and Consequences of the Oil Shock of 2007-08", *Brookings Papers on Economic Activity*, 215-261.
- Hamilton, J. (2012), "Oil Prices, Exhaustible Resources, and Economic Growth", in: *Handbook of Energy and Climate Change* (forthcoming).
- Hirsch, R., Bezdek, R. and Wendling, R. (2005), "Peaking of World Oil Production: Impacts, Mitigation and Risk Management", United States Department of Energy.
- Hirsch, R., Bezdek, R. and Wendling, R. (2010), *The Impending World Energy Mess*, Apogee Prime.
- Hubbert, M.K. (1956), "Nuclear Energy and the Fossil Fuels", *American Petroleum Institute Drilling and Production Practice Proceedings*, pp. 5-75.
- Hubbert, M.K. (1982), "Techniques of Prediction as Applied to the Production of Oil and Gas", in: S.I. Gass, ed., *Oil and Gas Supply Modeling*, Special Publication 631, Washington, National Bureau of Standards, pp. 16-141.
- IMF (2011), "Oil Scarcity, Growth and Global Imbalances", World Economic Outlook, April 2011, Chapter 3, International Monetary Fund.
- Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the Crude Oil Market", *American Economic Review*, **99**(3), 1053-1069.
- Kümmel, R. (2011), *The Second Law of Economics - Energy, Entropy, and the Origins of*

Wealth, Springer Verlag.

Kümmel, R., Henn, J. and Lindenberger, D. (2002), "Capital, Labor, Energy and Creativity: Modeling Innovation Diffusion", *Structural Change and Economic Dynamics*, **13**, 415-433.

Reynolds, D. (2002), *Scarcity and Growth Considering Oil and Energy: An Alternative Neo-Classical View*, Edwin Mellen Press.

Simmons, M. (2005), *Twilight in the Desert: The Coming Saudi Oil Shock and the World Economy*, Hoboken, New Jersey: John Wiley & Sons.

Smil, V. (2010), *Energy Transitions - History, Requirements, Prospects*, Praeger.

Sorrell, S., Miller, R., Bentley, R. and Speirs, J. (2010), "Oil Futures: A Comparison of Global Supply Forecasts", *Energy Policy*, **38**, 4990-5003.

UK Energy Research Centre (2009), "Global Oil Depletion - An Assessment of the Evidence for a Near-Term Peak in Global Oil Production".

United States Joint Forces Command (2010), "The Joint Operating Environment 2010".

	Parameter	Distribution	Prior Mode	Prior St. Dev. (or Bounds)	Posterior Mode	90% Confidence Interval
Oil Supply	α_s	uniform		[0 1000]	507.6483	[501.9955 514.8299]
	β_1	uniform		[0 100]	0.2427	[0.2353 0.2538]
	β_2	uniform		[0 100]	0.6238	[0.5053 0.7422]
	β_3	uniform		[0 100]	0.0546	[0.0043 0.1322]
Oil Demand	α_d	uniform	0	[-0.1 0.1]	-0.0177	[-0.0237 -0.0119]
	γ_1	lognormal	0.9	0.09	0.9098	[0.7844 1.0352]
	γ_2	invgamma	0.02	0.002	0.0213	[0.0181 0.0252]
	γ_3	invgamma	0.06	0.006	0.06	[0.0507 0.0707]
Output Growth	λ_1	beta	0.9	0.009	0.8987	[0.8833 0.9128]
	λ_2	normal	0.005	0.0005	0.0048	[0.0039 0.0056]
	λ_3	normal	0.005	0.0005	0.0048	[0.0040 0.0056]
Output Gap	ϕ_1	normal	0.85	0.085	0.9556	[0.9058 0.9873]
	ϕ_2	normal	0.25	0.025	0.2565	[0.2156 0.2967]
	ϕ_3	normal	0.005	0.0005	0.005	[0.0042 0.0058]
	ϕ_4	normal	0.005	0.0005	0.005	[0.0042 0.0058]

Table 1: Parameter Estimates

Horizon	<u>Real Price of Oil</u>		<u>Oil Production</u>		<u>GDP Level</u>	
	Model	Random Walk	Model	EIA	Model	WEO
1 year	14.7	27.7	1.69	1.59	1.82	1.83
2 years	17.6	47.4	1.97	2.57	3.03	3.41
3 years	19.9	57.9	2.31	3.51	3.62	4.69
4 years	22.4	79.0	2.41	4.66	3.74	5.55
5 years	25.1	100.0	2.69	5.72	3.05	5.00

Table 2: Root Mean Square Errors - Comparisons




Figure 1: EIA Forecasts 2001-2010
 (EIA Definition of World Total Oil Supply, in Mbd)





Figure 2: World Real Oil Prices and Spare Capacity




Figure 3: Colin Campbell Forecasts 2003-2010
 (Campbell Definition of Regular Conventional Oil, in Mbd)





Figure 4: Oil Production Forecasts in the Deffeyes (2005) Model
(Q in gigabarrels, q in gigabarrels per annum)

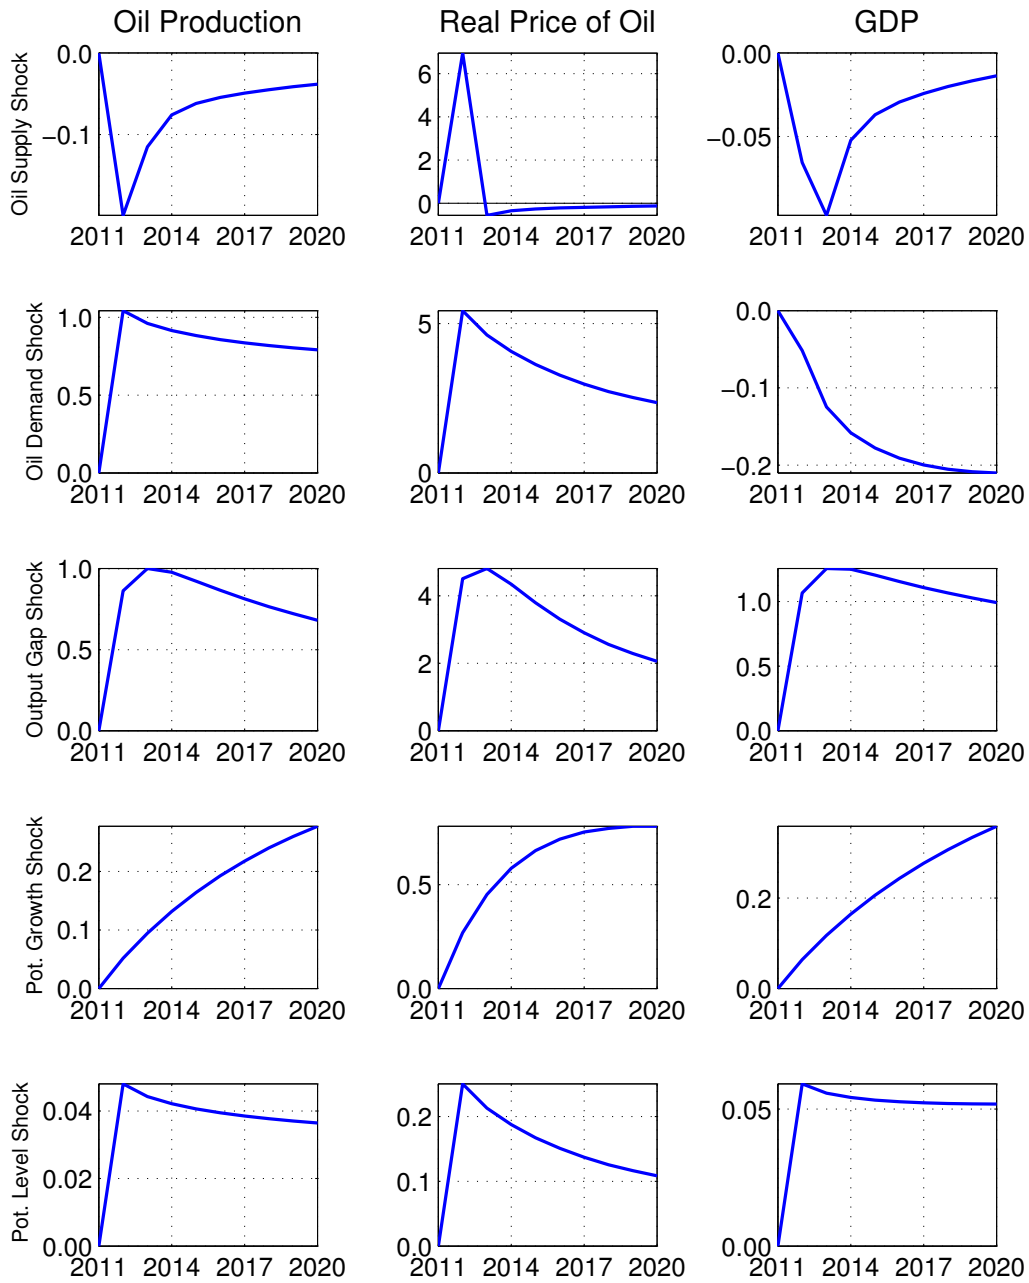


Figure 5: Impulse Responses
(in percent level deviation from control)




Figure 6: Historical Residuals
(in percent)

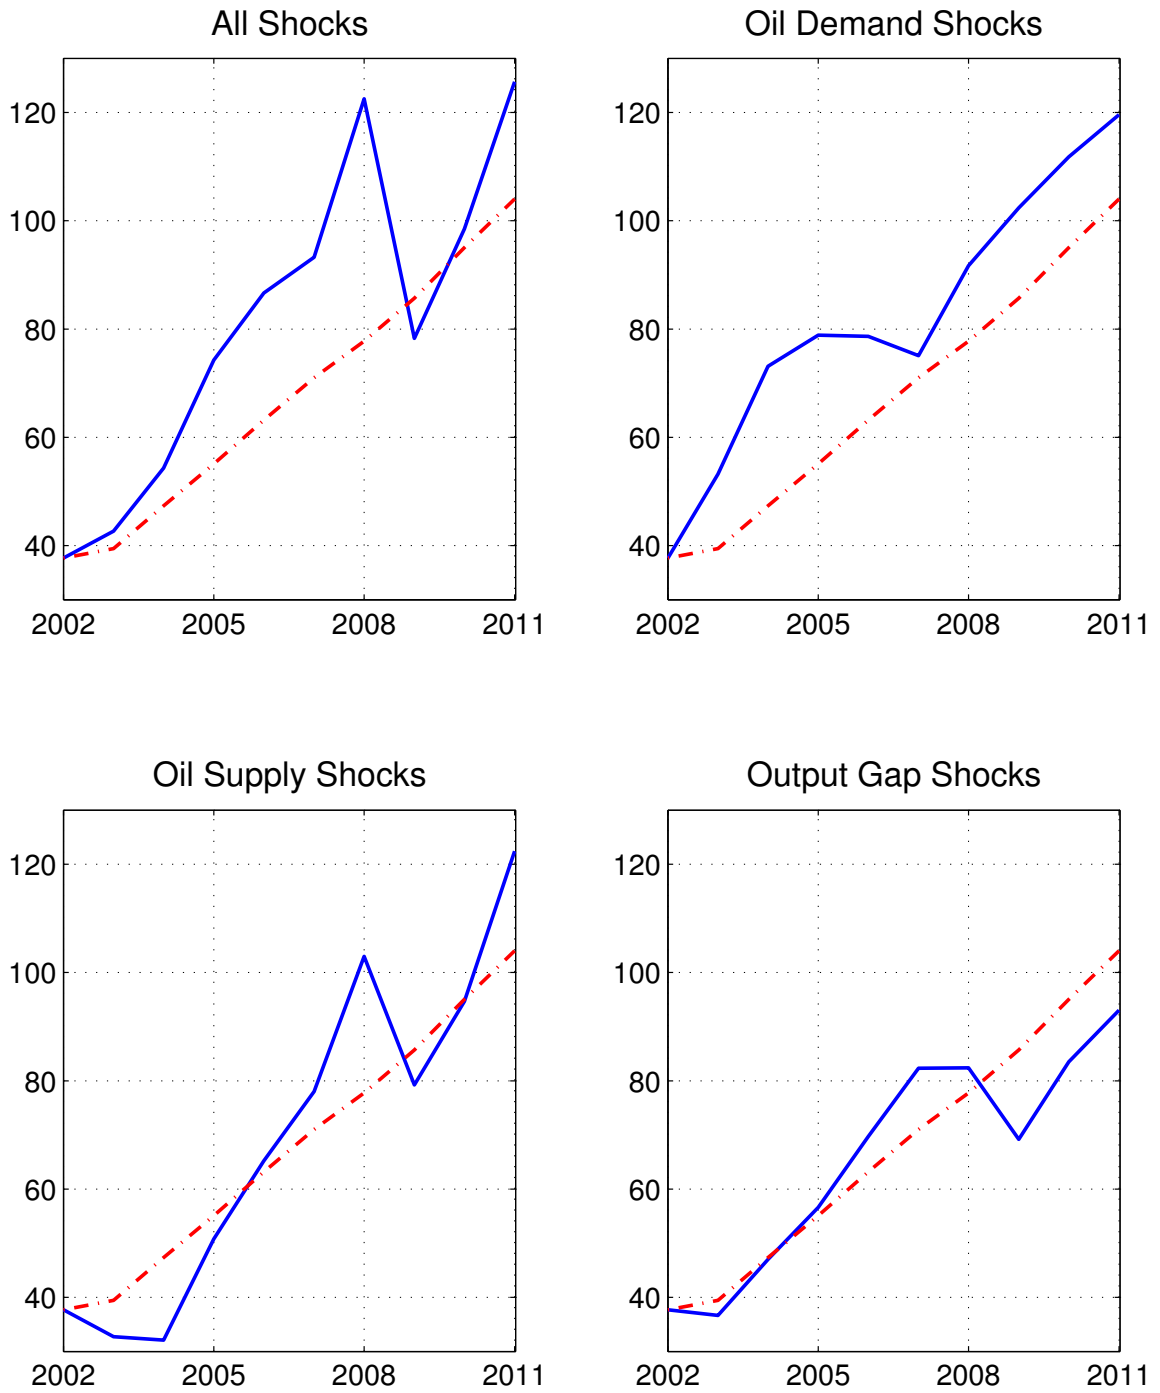


Figure 7: Contributions of Different Shocks to Oil Prices
(in real 2011 U.S. dollars)

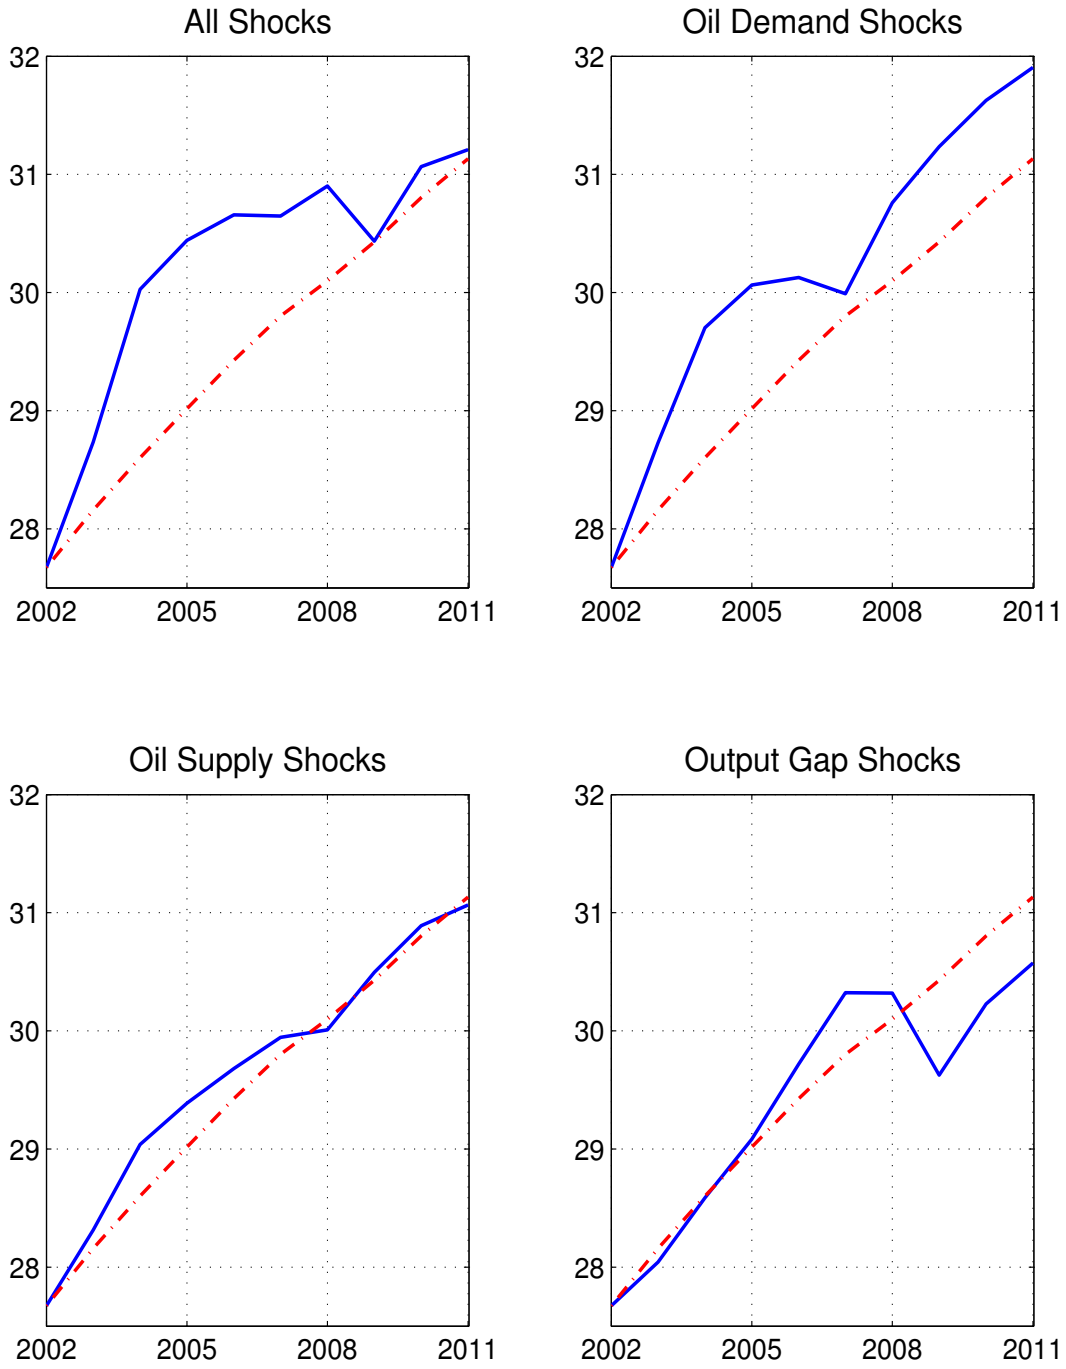


Figure 8: Contributions of Different Shocks to Oil Production (in gigabarrels per annum)

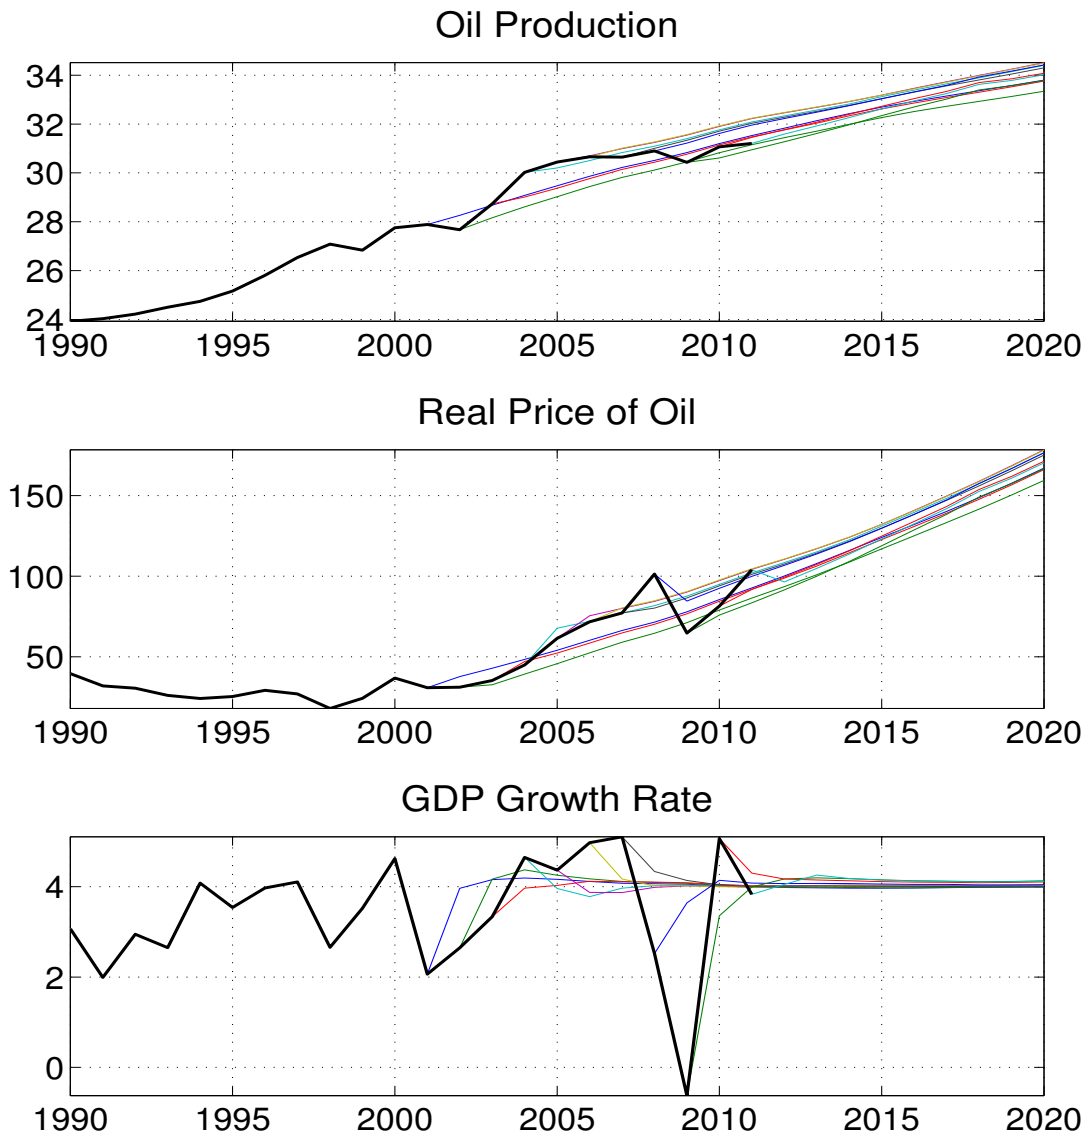


Figure 9: Rolling Forecasts




Figure 10: Oil Output Forecast with Error Bands
(in gigabarrels per annum)




Figure 11: Oil Price Forecast with Error Bands
(in real 2011 U.S. dollars)

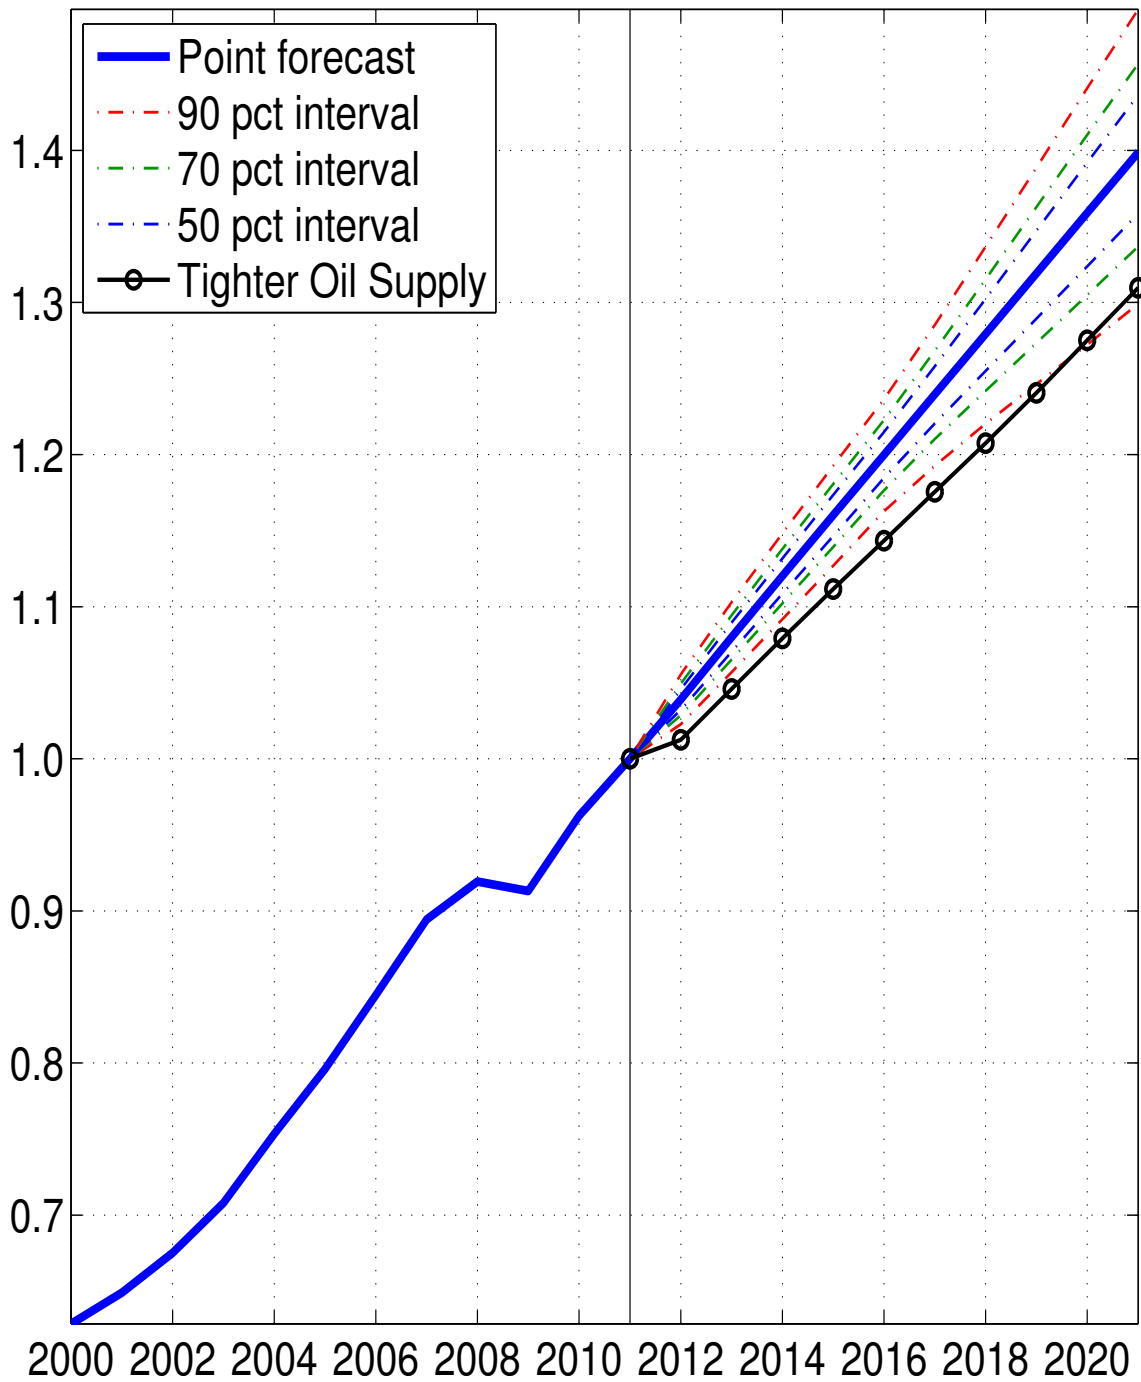


Figure 12: GDP (in logs) Forecast with Error Bands