International Journal of Forecasting 31 (2015) 207–221

Contents lists available at ScienceDirect

International Journal of Forecasting

journal homepage: www.elsevier.com/locate/ijforecast

The future of oil: Geology versus technology[☆]

Jaromir Benes[a], Marcelle Chauvet[b,*], Ondra Kamenik[c,1,2], Michael Kumhof[a], Douglas Laxton[a], Susanna Mursula[a], Jack Selody[d,1,3]

[a] *International Monetary Fund, United States*
[b] *University of California Riverside, United States*
[c] *OGResearch, Czech Republic*
[d] *Promontory Financial Group, Canada*

ARTICLE INFO

Keywords:
Exhaustible resources
Fossil fuels
Oil depletion
Hubbert's peak oil
Bayesian econometrics

ABSTRACT

We discuss and reconcile the geological and economic/technological views concerning the future of world oil production and prices, and 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, though the error bands are wide, reflecting sharply differing judgments on the ultimately recoverable reserves, and on future price elasticities of oil demand and supply.

1. Introduction

Future real oil prices are notoriously difficult to predict in real time, particularly over the medium and long run. Economists, government officials, and market oil specialists all experience this first hand, generally obtaining oil price forecasts that display no improvement, or only a marginal improvement, over the no-change forecast. However, the no-change forecast itself does a very poor job of predicting oil prices. This is particularly the case during sharp increases in prices, such as in the mid-1970s and the 2000s, together with the abrupt oscillations during the Great Recession in 2007–2009, which professional forecasters were slow to recognize. This result is well-known within the oil industry and the academic literature.

Several papers have shown, however, that the real price oil has some predictability in the short run. In a recent paper, Alquist, Kilian, and Vigfusson (2013) report that out-of-sample monthly forecasts from a reduced-form vector autoregressive model (VAR) of the global oil market are more reliable than forecasts from the random walk model at short horizons.[4] Nevertheless, at medium and long

☆ 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.

* Correspondence to: Department of Economics, University of California Riverside, CA 92521, United States. Tel.: +1 951 827 1587.

E-mail addresses: jbenes@imf.org (J. Benes), chauvet@ucr.edu (M. Chauvet), ondra.kamenik@gmail.com (O. Kamenik), mkumhof@imf.org (M. Kumhof), dlaxton@imf.org (D. Laxton), smursula@imf.org (S. Mursula), jselody@rogers.com (J. Selody).

1 Research Department, International Monetary Fund, 700 19th St NW, Washington, DC 20431, United States.

2 Present address: OGResearch, Sibeliova 41, Prague 6, 16200, Czech Republic.

3 Present address: Promontory Financial Group, Toronto, Canada. Tel.: +1 416 863 8500.

4 The forecasting performance is sensitive to variable selection and the lag length. In particular, Alquist et al. (2013) find, like Baumeister and Kilian (2012), that the real price of oil, defined as US refiners' acquisition cost for imported crude oil, is easier to forecast than the real price of West Texas Intermediate (WTI) crude oil. These results are based on mean square predictive errors.

http://dx.doi.org/10.1016/j.ijforecast.2014.03.012

horizons of one year and above, no-change forecasts systematically beat all models studied, and also professional forecasters. This result is also found by Baumeister and Kilian (2012, in press), who extend the analysis to include real time forecast restrictions at the monthly and quarterly frequencies, respectively. The econometric models of Alquist et al. (2013) and Baumeister and Kilian (2012, in press) use macroeconomic and financial indicators, as well as global crude oil production, as predictors of future oil prices. Many of these indicators are highly correlated with fluctuations in aggregate demand, so that the forecasts capture changes in the price of oil caused by variations in demand. In order to identify the roles of oil demand and oil supply shocks, Kilian (2009) proposes a structural VAR model of the global crude oil market. The model distinguishes between three drivers of real oil prices: global demand for industrial commodities, precautionary demand for oil, and oil supply, with the latter capturing the possibility of supply disruptions due to political events in oil producers, the dominant supply shock in historical data. The paper finds that the two demand shocks have been very important as drivers of oil prices, while supply shocks have had a negligible effect.

However, there is an alternative explanation for the recent persistent price movements that has received very little attention in the economics literature, despite there being considerable evidence to support it. This is the idea that one key driver of recent events may have been a highly persistent or even permanent shock to oil production 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.[5]

The most prominent economist who does not discount this possibility is James Hamilton. Hamilton (2009) finds that temporary disruptions in physical oil production have already played a major role in explaining the historical dynamics of oil price movements. Furthermore, he argues that stagnating world oil production, meaning a very persistent reduction in the growth rate of oil production, may have been one of the reasons for the run-up in oil prices in 2007–08.[6] According to Hamilton (2009), the main reasons why oil supply shocks affect output is their disruptive effects on key industries such as automotive manufacturing, and their impact on consumers' disposable incomes. In other words, the main effect is on the 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. It is often argued that such elasticities of substitution would tend to get larger over longer horizons, as agents find possible substitutes for oil, fueled by high prices that stimulate the technological change that can increase both the recovery of oil and the availability of substitutes for oil. Hamilton (2013), however, argues that the main reason for the historic growth in oil production has been the exploration of new geographic areas, rather than the application of better technology to existing sources, and that the end of that era could come soon. His paper goes on to explore the potentially very problematic implications of a slower future growth in oil production for future GDP growth.

Other than Hamilton, most proponents of the geological view of future oil production are found among physical scientists. They argue 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 stagnation in oil production in the face of persistent and large oil price increases is a sign that a physical scarcity of oil is already here, or at least is imminent, and that it must eventually overwhelm the stimulative effects of higher prices. Furthermore, based on extensive studies of alternative technologies and resources, they state that suitable substitutes for oil simply do not exist on the required scale, and that technologies for improving oil recovery must eventually run into limits dictated by the laws of thermodynamics, specifically entropy.

This view of oil production has its origins in the work of the geoscientist Hubbert (1956, 1962, 1967). Hubbert (1956) fitted historical production data to a symmetric bell-shaped curve and predicted correctly that US oil production would peak in 1970. Subsequently, Hubbert (1962, 1967) projected the ultimate quantity of oil to be recovered, and the rate at which it would be produced in the lower 48 US states. Hubbert (1962) adjusted logistic curves to cumulative production and discoveries, while Hubbert (1967) proposed an analysis of the quantity of oil discovered per foot of well drilled (yield per effort, YPE), fitting a negative exponential in order to form forecasts of the ultimate oil recovery (UOC). Hubbert's gloomy projections both spurred awareness and attracted criticism from the oil industry, government agencies, and academics. Some of the criticisms were related to the fact that the models were only based on physical oil production and discovery, and ignored the role of economics and technological changes. The response of Hubbert, and of subsequent studies validating his work, was that geological features were ultimately the main drivers of oil discovery, production and distribution, and that factors other than those were already embedded in the historical series used in the model.

The empirical success of Hubbert's seminal approach motivated various important academic studies that incorporated additional economic, institutional and/or technological factors into the original model, and that proposed alternative estimation methods. A partial list includes the studies by Cleveland and Kaufmann (1991), Kaufmann (1991), Kaufmann and Cleveland (2001), and Pesaran and Samiei (1995). Cleveland and Kaufmann (1991) extend Hubbert's (1962) model to account for the non-random historical drilling pattern in the oil industry in the lower

5 Kilian's (2009) analysis does not consider the possibility of shocks to the supply of oil that are driven by terminal geological limits.

6 In particular, Hamilton (2009) argues that the main dynamic was strong demand, at a low price elasticity of demand, meeting stagnating world oil production. Hamilton also finds that the flow of investment dollars into commodity futures contracts was important, but not the key factor, in explaining the late 2000s increase in real oil prices, the largest in history. By contrast, Baumeister and Peersman (2013), Kilian (2008, 2009), and Kilian and Hicks (2013) stress the role of oil demand shocks rather than oil supply shocks in causing the 2007–2008 oil price surge.

48 US states, and to incorporate the effects of the real price of oil and the annual rate of drilling effort on the YPE. The analysis is also extended to account for political/institutional factors, such as pro-rationing by the Texas Railroad Commission (TRC). The model is based on an extended version of Hubbert's production curve, and is estimated via ordinary least squares (OLS). Kaufmann (1991) proposes a two-step approach to study the impact of geological, economic, and political/institutional changes on oil production in the lower 48 states. It combines Hubbert's fitting approach with econometric methods. In the first stage, changes in physical resources are represented by a bell-shaped curve. In the second stage, the difference between actual and estimated production is modeled as a function of political/institutional and economic factors: average real oil prices, the relative price of oil to natural gas, and pro-rationing by the TRC. This model is estimated using OLS, and a grid search is used to identify the bell-shaped curves from the two-step procedure. Kaufmann and Cleveland (2001) argue that the forecast success of Hubbert's bell-shaped curve is due to the fact that it is a good approximation of the nonlinear long-run oil cost function. However, Kaufmann and Cleveland (2001) claim that Hubbert's results might be spurious because he does not take into account stochastic trends which are shared by oil production and other variables of the model. Their paper proposes a vector error correction model of oil production that includes real oil prices, average cost of production, and pro-rationing by the TRC.[7] Essentially, the analysis allows for economic factors (oil prices) and institutional factors (decisions from the TRC) having an impact on the dynamics of oil production. Pesaran and Samiei (1995) evaluate alternative estimation methods for the model of Hubbert (1956, 1962, 1967) and for the extensions of Cleveland and Kaufmann (1991) and Kaufmann (1991). The paper discusses parameter identification and estimation of the models, and proposes a forecasting formula for the ultimate oil recovery based on the rate of recovery rather than the cumulative production function.[8] The paper argues that when economic factors are taken into account, the estimates of ultimate recovery become state-dependent. The model is estimated for the lower 48 US states under this assumption.

More recently, Hubbert's work was discussed in a study for the US Department of Energy by Hirsch, Bezdek, and Wendling (2005), and in a subsequent book by Hirsch, Bezdek, and Wendling (2010).[9] The most thorough research available on this topic is by the UK Energy Research Centre (2009), and is summarized succinctly by Sorrell, Miller, Bentley, and Speirs (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 therefore, by implication, in aggregate activity, can be improved dramatically by combining the geological and economic/technological views of oil production, and by estimating their respective contributions.[10] We develop a simple macroeconomic model that combines a conventional linear specification for world oil demand with a nonlinear equation for world oil supply, with the latter integrating a mathematical formalization of the geological view with a conventional price-sensitive view of oil production. The world oil supply equation is an augmented version of the nonlinear Hubbert specification, in which oil is assumed to be more difficult to extract as the cumulative production increases (geological constraint), while production also responds positively to higher current and past real oil prices (economic/technological view). The model considers both short-run and medium-run effects of oil prices on production. For example, in the short run, increases in oil prices raise oil production to the extent that producers utilize any spare capacity from existing fields. In the medium run, high real oil prices lead to new exploration and/or better technologies, with the effects on production occurring after a few years. Finally, the model includes a decomposition of output into trend and gap components for the determination of world GDP. The model is estimated as a system of equations using nonlinear Bayesian techniques.[11]

We find that this model can predict oil prices far better out-of-sample, at all horizons, than a random walk. In addition, it can also predict oil production, at horizons greater than one year, far better than the historical track record of either official energy agencies on the one hand, or advocates of pure versions of the geological view on the other hand. We use the proposed model to identify which driving force has been most responsible for the recent run-up in oil prices, and find that the geological, price insensitive component of the oil supply equation is the key reason for the accuracy of the model's recent predictions, because it captures the underlying trend in prices. However, we also find that shocks to the demand for goods and the demand for oil have been key to explaining persistent and sizeable deviations from that trend, the latter probably due to phenomenal recent growth in China and India. 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. However, there is substantial degree of uncertainty about these future trends that is rooted in our fundamental lack of knowledge, based on current

7 In their analysis, real oil prices are decomposed to account for asymmetric effects of price increases. In addition, the inclusion of the average cost of production implies that firms do not rank and produce their fields as a function of quality.

8 These problems in identification and estimation of the models would cast doubt in Kaufmann and Cleveland's (2001) results.

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

10 This paper uses data for world real GDP from the IMF. Data for the real supply of oil come from the IEA's World Oil Statistics database. The nominal world oil price is computed as the US dollar average of UK Brent, Dubai, and West Texas Intermediate. The real oil price uses the US CPI as the deflator.

11 We have verified that, unlike those studied by Kaufmann and Cleveland (2001), the series studied in this paper do not share common stochastic trends.

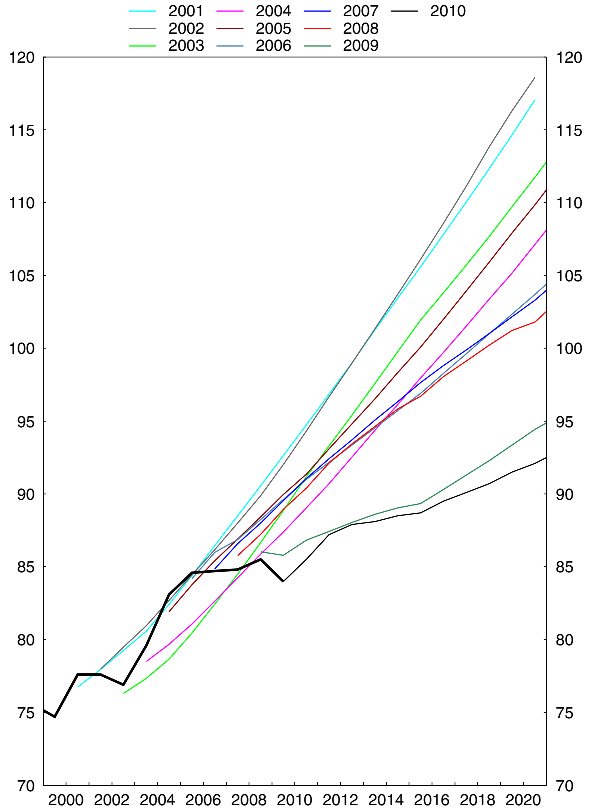

Fig. 1. EIA forecasts of oil production, 2001–2010 (EIA definition of world total oil supply, in Mbd).

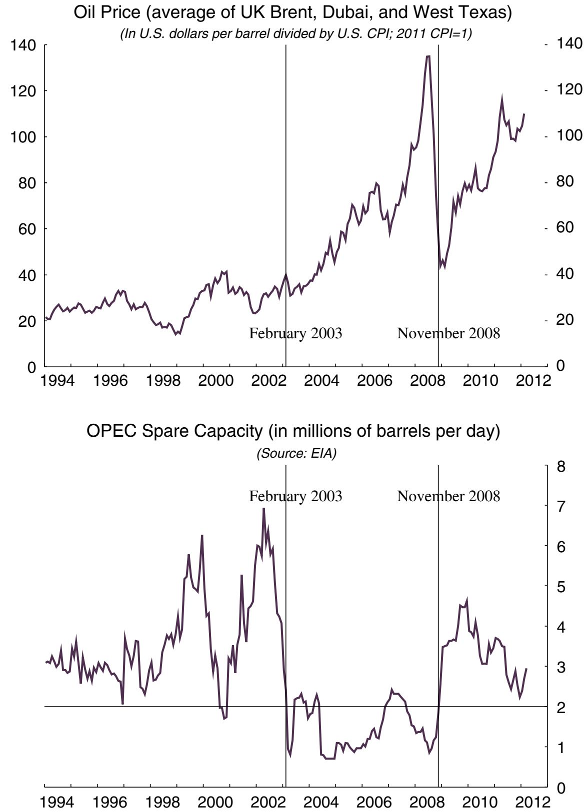

Fig. 2. World US$ oil prices and spare capacity.

data, about ultimately recoverable oil resources, and about the long-run price elasticities of oil demand and oil supply.

The rest of the paper is organized as follows. Section 2 presents various historical forecasts of oil production made by proponents of the technological and geological views. Section 3 presents and discusses the model specification and parameter estimates. Section 4 contains a detailed analysis of the estimation results. Section 5 concludes.

2. Historical forecasts of world oil production

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

The reason why this may be the case is illustrated in Fig. 2, which displays real world oil prices in 2011 US dollars,[12] alongside the OPEC spare capacity in millions of barrels per day (Mbd). Until the end of 2002, spare capacity was high in historical terms, and this was accompanied by oil prices that had not been growing significantly in real terms. This changed abruptly in early 2003, around the time of the Iraq war, when spare capacity dropped below the 2 Mbd mark, which is considered by many in the industry to be the critical point at which supply becomes a constraining factor. From that moment until the onset of the Great Recession, real oil prices started a long-term increase whereby they ultimately more than tripled, before the demand destruction of the Great Recession led to a sudden increase in spare capacity and a steep decline in oil prices. However, this only brought temporary relief to the demand-supply balance in the oil market, for two reasons. First, as we have seen in Fig. 1, oil production never regained its historical growth rate of 1.5%–2% per annum after 2005, and, in fact, actually remained on a plateau for several years. Second, partial demand recoveries restarted from 2009 onwards in many economies. Spare capacity therefore approached 2 Mbd again, and oil prices ratcheted up. With the exception of the period of deep recession, the combination of a plateau in actual oil production and renewed pressure on spare capacity indicates that the

12 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 was US$ 104. The same normalization is adopted in all subsequent charts of the real oil price.

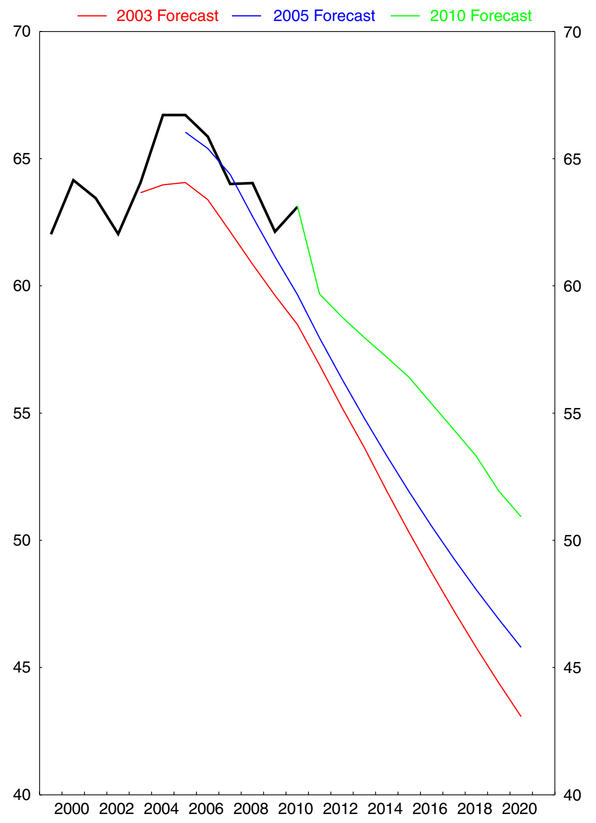

Fig. 3. Colin Campbell's forecasts for oil production, 2003–2010 (Campbell's definition of regular conventional oil, in Mbd).

physical constraints on oil production started 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. Fig. 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. For the EIA, we showed the forecasts for the world's total oil supply, which is defined as crude oil plus Natural Gas Liquids (NGL) and other liquids, plus refinery processing gains. For Campbell, we show historical forecasts for regular conventional oil. This definition covers over 75% of the world total oil production; although it is based on EIA data, it excludes heavy oil (<17.5 deg API), bitumen, oil shale, shale oil, deepwater oil and gas (>500 m), polar oil and gas, and NGL from gas plants. Furthermore, the International Energy Agency (IEA) uses yet another definition that is slightly less encompassing than the EIA's, but more encompassing than Campbell's, namely crude oil plus NGL. We will use IEA data in our empirical analysis. We have used EIA data in Fig. 1 because the EIA produces annual forecasts while the IEA does not. Fig. 3 shows that Campbell's forecasts have also erred, but on the pessimistic side this time. The differences from ex-post realized production data are somewhat smaller than those for 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 go 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.[13] As this approach yields econometrically testable equations for the production profile, we will pursue this in detail in this paper. A particularly tractable specification is known as Hubbert linearization. This is based on Deffeyes (2005), who develops a considerably simplified version of the analysis by 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. Hubbert states that annual production can be approximated usefully by the logistic curve

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

This is a bell-shaped curve, and it states that, in any given year, the 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 from exactly 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 Eq. (1) by Q_t, which produces a linear relationship between cumulative production and the ratio of current to cumulative production:

$$\frac{q_t}{Q_t} = \alpha_s - \frac{\alpha_s}{\bar{Q}} Q_t. \tag{2}$$

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

$$\frac{q_t}{Q_t} = \alpha_s - \beta Q_t. \tag{3}$$

Deffeyes (2005) finds that this relationship fits both US and world data very well until 2003, the last data point in his 2005 study, with the two 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 Fig. 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 Fig. 2,

13 See UK Energy Research Centre (2009) for a very detailed discussion.

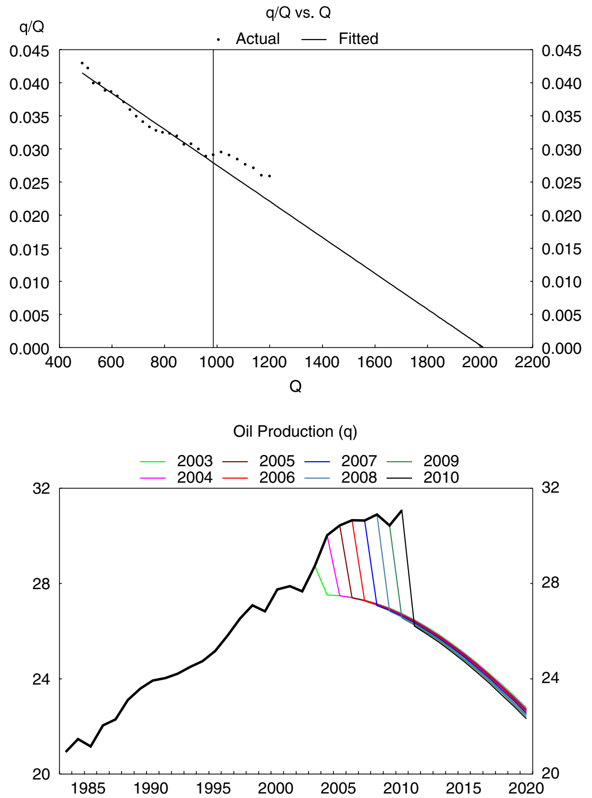

Fig. 4. Oil production forecasts in the Deffeyes (2005) model (Q in gigabarrels, q in gigabarrels per annum).

the critical feature of the 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 the period 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. Instead, it reached a plateau, where it remained, with some fluctuations, for several years. In other words, prices did not matter enough to allow production to regain its historical growth rate. The recent pick-up in the US production of shale oil has lifted world oil production above this plateau, but it is too early to know whether this will be sustainable (see the discussion by Benes et al., in press).

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

3.1. Data and estimation methodology

We estimate this system of equations using data for the world real oil price, the world real GDP, and the world real oil supply. The world nominal oil price is computed as the US dollar average of UK Brent, Dubai, and West Texas Intermediate. The real oil price uses the US CPI as the deflator. Data for world real GDP come from the IMF. Data for the world real oil supply come from the IEA's World Oil Statistics database. We use the term "oil supply" to correspond to a specific IEA definition of oil aggregates.

Specifically, on the production side of oil we have the definition

$$\text{Oil Supply} = \text{Oil Production} + \text{Refinery Processing Gains},$$

where processing gains refer to volumetric (but not energy) gains during the refining process from crude oil to the final product. Thus, "oil supply" corresponds to the production volume of the final product. In this paper we therefore generally use the terms oil supply and oil production interchangeably, with the understanding that, in terms of the data, both refer to the IEA's definition that includes processing gains.

Separately, on the demand side of oil we have the definition

$$\text{Oil Supply} = \text{Oil Demand} + \text{Change in Oil Inventories}.$$

We assume implicitly that there are no changes in oil inventories, in that the model assumes that prices equalize the demand and supply in every period.

We use annual data from 1983 to 2011, with lags that use data back to 1972 for oil prices. The model has multiple factors that drive the oil price and oil production dynamics in a fairly short sample, which can potentially lead to difficulties in obtaining sensible parameter estimates. In order to overcome this problem, we employ nonlinear Bayesian estimation techniques, using priors based on previous studies. Nonlinear techniques need to be used because the world oil supply equation is an augmented version of the nonlinear Hubbert linearization specification in Eq. (3). The model is estimated as a system. Given that 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 provided in Table 1. Posterior modes are also displayed underneath the parameter symbols in the model equations displayed below.

Table 1
Parameter estimates.

	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]

3.2. Oil supply

The oil supply equation combines the geological view embodied in the Hubbert linearization equation (Eq. (3)), whereby oil is more and more difficult to extract as the 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.[14] 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 by between four and six years. For estimation purposes, the units of cumulative production Q_t are gigabarrels, while the units of annual production q_t are hundreds of thousands of barrels. The supply equation is:

$$\frac{q_t}{Q_t} = \underset{(507.7)}{\alpha_s} - \underset{(0.243)}{\beta_1} Q_t + \underset{(0.624)}{\beta_2} p_t + \underset{(0.056)}{\beta_3} \frac{1}{3}\sum_{k=4}^{6} p_{t-k} + \epsilon_t^s, \tag{4}$$

with the auxiliary relationship

$$Q_t = Q_{t-1} + q_t/10\,000. \tag{5}$$

where $\epsilon_t^s = \phi_s \epsilon_{t-1}^s + \eta_t^s$ is the AR(1) oil supply shock and η_t^s is distributed *i.i.d.*, with mean zero and standard deviation σ_s^2. The parameter α_s 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. The coefficient β_1 is given a fairly loose uniform prior distribution. The priors for β_2 and β_3 were also given a loose uniform distribution. The reason for this is that our knowledge about the oil supply response to price increases is limited, as most estimated economic models have focused only on demand elasticities.

The estimated coefficient $\beta_1 = 0.243$ supports the 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 loose prior. The coefficients β_2 and β_3 can be converted to price elasticities of supply,[15] but, given the levels specification of Eq. (4), these elasticities depend on actual oil production, and oil prices in particular. 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 production and prices. Our forecasts show upper and lower bands, and also something of a sensitivity analysis, which indicate what is at stake. Most importantly, the fact that the main production response to prices has been contemporaneous may be a cause for concern, because this indicates that production 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, a 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. This delays and raises the peak of oil production, and perhaps also delays

[14] This involves an important technical consideration: excessively fast extraction of oil from an existing field can destroy geological structures and reduce the quantity of oil which can ultimately be recovered (see Simmons, 2005).

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

the point at which $q_t = 0$. For example, estimating 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. However, this does not remove the tendency for oil production to decline eventually, unless real oil prices were to keep rising steeply and indefinitely.

3.3. Oil demand

Oil demand is specified according to 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 increasing the incentive to find substitutes for 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.[16] 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

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

where $\epsilon_t^d = \phi_d \epsilon_{t-1}^d + \eta_t^d$ is the AR(1) demand shock and η_t^d is distributed *i.i.d.*, with mean zero and standard deviation σ_d^2. 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 a considerable degree of 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.[17] Specifically, elasticities are unaffected at the average oil prices seen prior to 2008, rise by roughly a factor of 1.75 at the average prices of 2008 and 2011, and eventually rise by a factor of at most 2.5 at the much higher prices projected by the model out to 2021.

The estimate for the income elasticity of oil demand γ_1 is consistent with other studies, which have found that, on average, industrialized countries display a lower income elasticity 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 by the 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 the available supply, or even an 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.

3.4. GDP equations

The feedback from oil prices to economic activity is captured by the reduced-form specification

$$gdp_t = pot_t * y_t, \tag{7}$$

where pot_t is potential output and y_t is the output gap, and where oil prices enter into the equations determining both of these terms, as will be discussed below. Furthermore, this specification allows us to introduce shocks to the output gap (transitory shocks to the level of output), potential output (permanent shocks to the level of output), and the potential output growth (transitory but persistent shocks to the growth rate of output) separately. The richness of this specification helps us to model the complicated interactions of oil price movements and GDP, where both the 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.

3.4.1. Level of potential GDP

Potential GDP is given by

$$\Delta \ln pot_t = \ln g_t + \epsilon_t^{pot}, \tag{8}$$

where ϵ_t^{pot} is a shock to the level of potential output and g_t is the growth rate of potential output. Oil prices do not enter this equation, since we assume that the dynamic effects of oil prices on potential output will be captured in the potential growth rate equation.

3.4.2. Growth rate of potential 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

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

[16] There are grounds for doubt as to whether long-run elasticities can continue indefinitely to be much higher than short-run elasticities. See the discussion in Section 4.5.

[17] To keep the exposition simple, this is not shown in Eq. (6).

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 is 4%. 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 7%. The results indicate that an oil price growth rate that is higher than the 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.

3.4.3. 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. Specifically, we have

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

where ϵ_t^y represents a shock to the level of aggregate demand. Similarly 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 parameter estimates that have been discussed for the model's impulse response functions, interpretation of history, and forecast accuracy, and for current forecasts of oil production, oil prices and GDP.

4.1. Impulse response functions

Fig. 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 percentage deviations from control, after removing any underlying trend.

Oil supply shocks occur separately from, and in addition to, the geological tightening effects of Hubbert's curve in Eq. (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 upward spikes observed in real oil prices when these shocks did occur have been significant. The top row of Fig. 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 return quickly 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 been less 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 Fig. 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 the 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 Fig. 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 the 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 the overall variability in the model. When they do occur, the effects on output, oil production and oil prices are of course highly persistent.

4.2. Interpretation of history

Fig. 6 shows the estimated shocks of the model. Figs. 7 and 8 show model decompositions of the post-2002 movements in oil prices and oil supply into the contributions of the three shocks that account for most of the variability in the model. In each case, the broken line represents the no-shock paths. The solid line in the top left panel corresponds to all shocks in the model, and is therefore, by construction, identical to the data. The solid lines in the remaining panels show the separate contributions of the estimated shocks to oil demand, oil supply, and the output gap.

We begin with Fig. 7, the decomposition of oil prices. The most important observation is 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.[18] The reason for this is the significant estimate

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

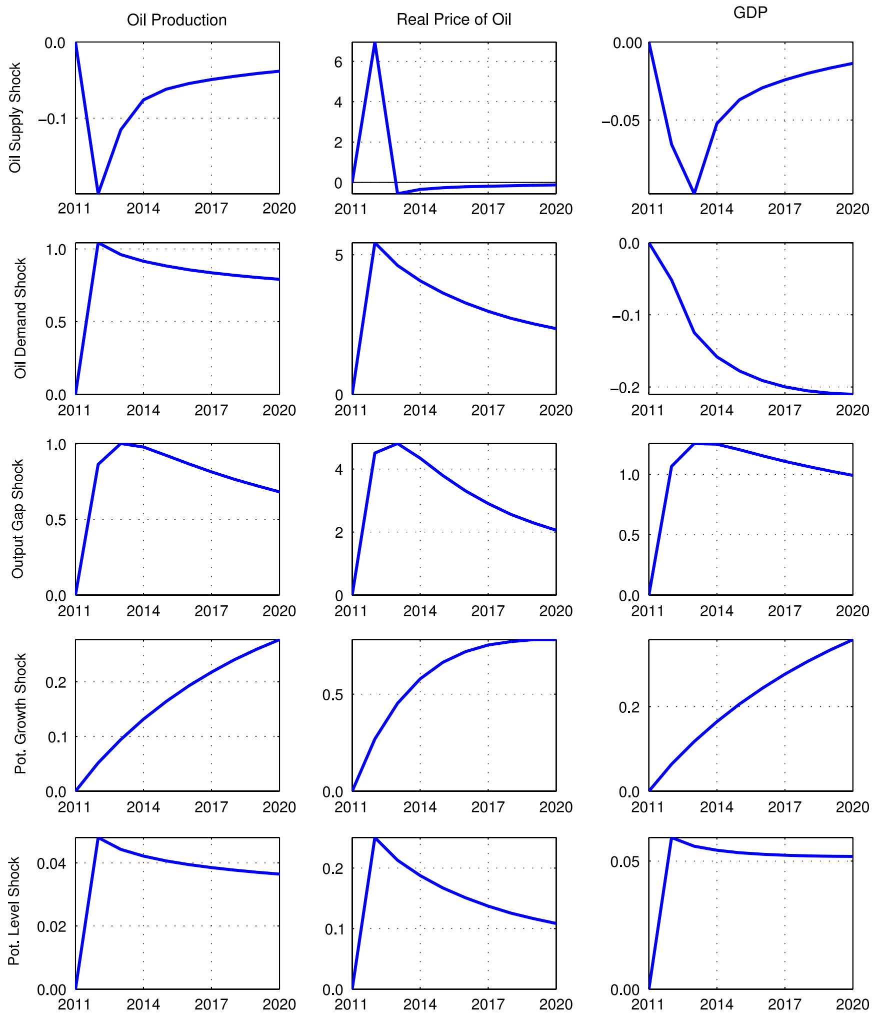

Fig. 5. Impulse responses (in percent level deviations from control; shocks in rows, variables in columns).

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 increased oil prices significantly, regardless of shocks.

As for the contribution of shocks, 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. In the earlier years, the major contributing factors were a very strong oil demand, mainly from booming emerging economies, and a positive world output gap. Oil supply shocks, at least until some time in 2005, actually helped, ceteris paribus, to keep oil prices lower than they would otherwise have been. However, as we have seen, world oil production stayed on a plateau from that time onward, and by 2008, an insufficient world oil supply had become the major factor behind high oil prices. The impact of the financial crisis in 2008 associated with the Great Recession was so severe that, in 2009, 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 returned to 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.

Fig. 8 decomposes oil production, in gigabarrels per annum.[19] We observe that, except for 2009, production was 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 to 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

19 As was explained in Section 3.1, this means that it decomposes the data series "oil supply" from the IEA database.

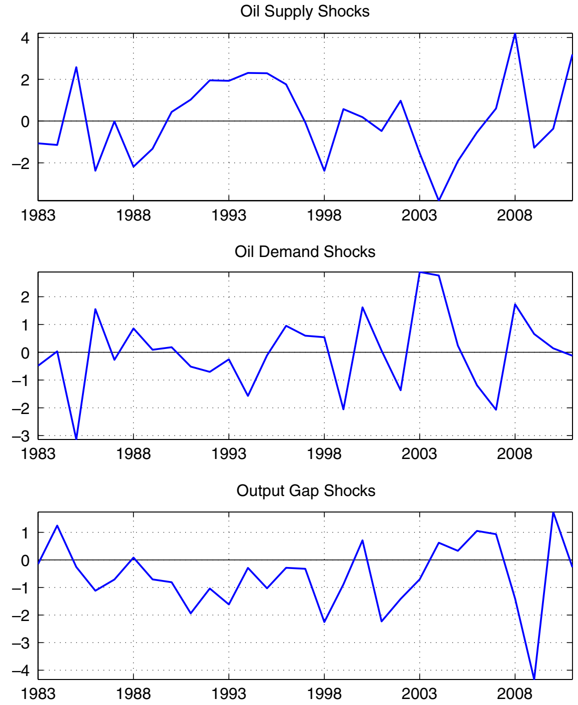

Fig. 6. Historical residuals (in percentages).

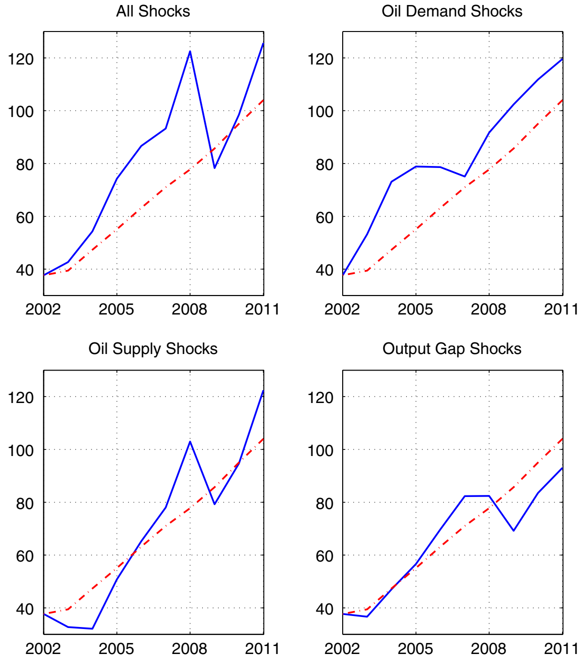

Fig. 7. Contributions of different shocks to oil prices (in real 2011 US dollars).

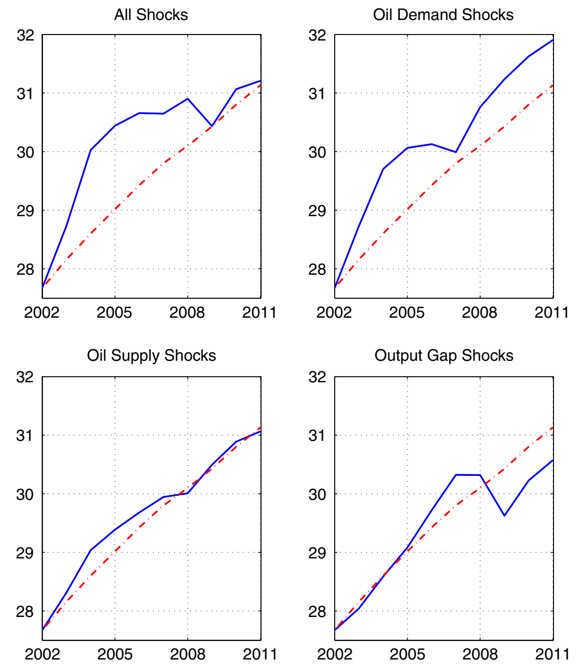

Fig. 8. Contributions of different shocks to oil production (in gigabarrels per annum).

post-2003 deviations from the pure geological explanation of oil production and prices. However, 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

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

The predicted average annual growth rates of oil production are well below the historical forecasts of the EIA, but above the forecasts made by proponents of the geological view. We therefore find that our model's accommodation of both the geological and 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 production (on a plateau until recently) and spare capacity on the one hand, and of the clear effects of prices in overturning the pure Deffeyes (2005) model.

However, the 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 in order to maintain an expansion of production 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 effect of higher oil prices on the GDP 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 which are similar to the nonlinearities in our oil demand equation. There is likely to be a critical

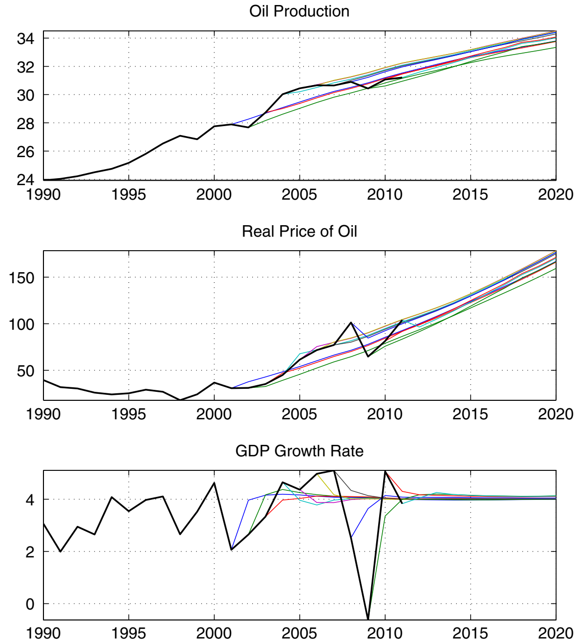

Fig. 9. Out-of-sample rolling forecasts, 2001–2011. (Oil production: gigabarrels per annum; real price of oil: real 2011 US dollars; GDP growth rate: percentage points. Each line starts in the year in which the respective forecast is made.)

range of oil prices at which 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 oil experts. For example, the chief economist of the International Energy Agency, Fatih Birol, has repeatedly warned that oil prices have reached a point that could push the world economy back into recession.[20] We will study this possibility quantitatively in future work.

Fig. 9 shows that our model predicts neither a mean-reverting oil price, like 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 predictions for oil production in the early 2000s are far more accurate than either the contemporaneous EIA forecasts or the forecasts using either Deffeyes' or Campbell's methods. In order to formalize these comparisons of forecast accuracies, Table 2 shows the root mean square errors (RMSEs) of our model's rolling forecasts over 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 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 existing literature, these results are dramatic. The gains are less dramatic for GDP, but are very substantial nevertheless.[21]

Table 2
Root mean square errors—comparisons (based on out-of-sample rolling forecasts, 2003–2011).

Horizon	Real price of oil		Oil production		GDP level	
	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

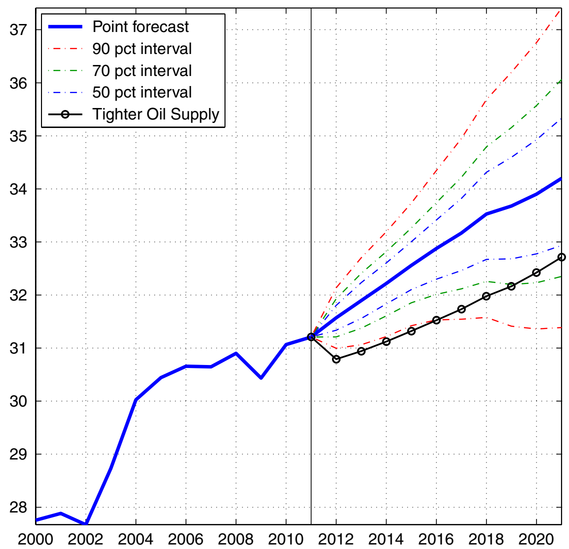

Fig. 10. Oil output forecast with error bands (in gigabarrels per annum).

4.4. Current forecasts

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

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

21 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.

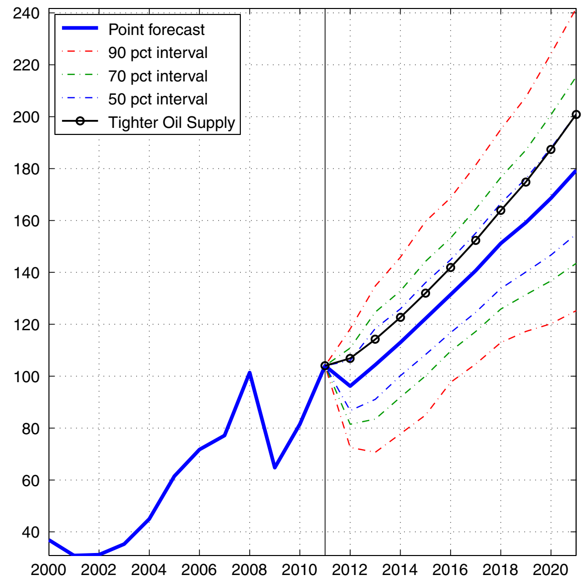

Fig. 11. Oil price forecast with error bands (in real 2011 US dollars).

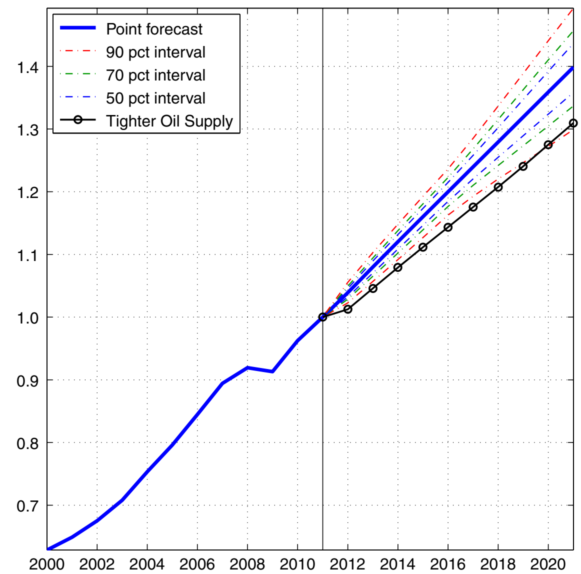

Fig. 12. World GDP (in logs) forecast with error bands (normalized index, with the log of 2011 world real GDP equal to 1).

Fig. 10 shows oil production, in gigabarrels per annum. The point forecast is for a mean annual growth rate of oil supply of around 0.9% over the coming decade, which is 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 the ultimately recoverable resources (implicit in β_1), as well as the supply and demand elasticities with respect to the oil price. The lower 90% band indicates flat oil production for the entire decade, while the upper band indicates annual production 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 which is approximately as large as the most recent EIA forecasts, the forecast for the oil price that is behind this production forecast is far higher than that anticipated by the EIA.

This is seen in Fig. 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% confidence 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.

Fig. 12 shows forecasts for GDP, with the 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 production. Nevertheless, the 90% confidence interval contains average growth rates as low as 3% per annum, and as high as 5% per annum. In other words, for more pessimistic coefficient values of ultimately recoverable resources and elasticities, average world growth would be one percentage point lower.

Finally, Figs. 10–12 also report the point forecast for an alternative scenario where β_2 takes a 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. The average growth rate of oil production 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 GDP growth by around one 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 production 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 with this is that, at this moment, our historical data contain very little information about what that relationship might look like. However, we are not entirely without information, since a number of authors in other sciences have started to ask pertinent questions, and have done some early pioneering work.[22]

There are two key questions under the maintained hypothesis of much a lower growth in oil production. 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 a stagnating or declining world oil production rate.

[22] See Kumhof and Muir (in press) for a more comprehensive overview and analysis.

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 US economy,[23] 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. These contributions include those of Ayres and Warr (2005, 2010), Hall and Klitgaard (2012), Kümmel (2011), and Kümmel, Henn, and Lindenberger (2002), who propose aggregate production functions that are based on concepts from engineering and thermodynamics. Several of these contributions estimate their own production functions. The estimations are based on technologies that use energy, rather than just oil, but, given the very limited substitutability between oil and other forms of energy, this nevertheless offers important insights.[24] These authors find output contributions of energy of up to around 50%, despite the low cost share of energy. Clearly, if this can be confirmed in further rigorous econometric studies, the implications for GDP of a lower growth in oil production 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 multi-regional Dynamic Stochastic General Equilibrium model, the Global Integrated Monetary and Fiscal model (GIMF), assuming a technology where oil's output contribution far exceeds its cost share. The simulations find that, following permanent declines in the growth rate of world oil production, the model generates much larger negative output effects than the conventional neoclassical model, because a share of the stock of technology would become obsolete.[25] This channel has never yet been of sufficient importance to explain the historical data, and our empirical model therefore does not contain it. Changing this would lead to simulation results with lower GDP growth levels.

The other key future concern concerns elasticities of substitution. Several important contributions have challenged economists' automatic assumption that the 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 the fact that this assumption is not consistent with the historical facts (Smil, 2010),[26] real-world practical limits (Ayres, 2007), or the laws of thermodynamics, specifically entropy (Reynolds, 2002, Chapter 10). Our empirical model currently makes the conventional assumption that, after some time, elasticities will 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 the world oil supply, because there is a finite limit to the extent that machines (and labor) can substitute for energy. If we were to incorporate this assumption, the model would forecast significantly higher oil prices in the event of a sufficiently large and persistent shock to world oil supply.

5. Conclusion

The main objective of this paper has been to propose a model of the world oil market that does not take an a-priori view of the relative importance of resource constraints and the price mechanism, and to evaluate it empirically. We do not want to rule out either of these mechanisms, because the recent data indicate convincingly 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 the oil production.

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 into 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 permanent near-doubling 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 the GDP is based on historical data, and indicates perceptible but small and transitory output effects. However, we suspect that there must be a pain barrier, a level of oil prices above which the effects on GDP will become 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 also have aspects of a negative technology shock. In that case the macroeconomic effects of binding resource constraints could be much larger and more persistent, and would extend well beyond the oil sector. Studying these issues in greater depth will be a priority in our future research.

Acknowledgments

The authors thank the referees and Fred Joutz for their comments and suggestions for improvements. Serhat Solmaz and Kadir Tanyeri provided excellent research assistance.

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

24 For the US economy, the historical cost share of total energy has been around 7%.

25 See e.g. Atkeson and Kehoe (1999) and Kim and Loungani (1992).

26 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 identified clearly. The latter is clearly not the case today, as renewables are not even nearly of sufficient scale.

References

Alquist, R., Kilian, L., & Vigfusson, R. (2013). Forecasting the price of oil. In G. Elliott, & A. Timmermann (Eds.), *Handbook of economic forecasting, Vol. 2* (pp. 427–507). Amsterdam: North-Holland.

Atkeson, A., & Kehoe, P. J. (1999). Models of energy use: putty-putty versus putty-clay. *American Economic Review*, *89*(4), 1028–1043.

Ayres, R. (2007). On the practical limits to substitution. *Ecological Economics*, *61*, 115–128.

Ayres, R., & Warr, B. (2005). Accounting for growth: the role of physical work. *Structural Change and Economic Dynamics*, *16*, 181–209.

Ayres, R., & Warr, B. (2010). *The economic growth engine—how energy and work drive material prosperity*. Edward Elgar Publishing.

Baumeister, C., & Kilian, L. (2012). Real-time forecasts of the real price of oil. *Journal of Business and Economic Statistics*, *30*(2), 326–336.

Baumeister, C., & Kilian, L. (in press) What central bankers need to know about forecasting oil prices. *International Economic Review*.

Baumeister, C., & Peersman, G. (2013). The role of time-varying price elasticities in accounting for volatility changes in the crude oil market. *Journal of Applied Econometrics*, *28*(7), 1087–1109.

Benes, J., Kumhof, M., Laxton, D., Maih, J., Solmaz, S., & Zhang, F. (2014). *The future of oil: mind the gap*. IMF working papers (in press).

Cleveland, C., & Kaufmann, R. K. (1991). Forecasting ultimate oil recovery and its rate of production: incorporating economic forces into the models of M. King Hubbert. *Energy Journal*, *12*(2), 17–46.

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. A. S., & Klitgaard, K. A. (2012). *Energy and the wealth of nations: understanding the biophysical economy*. Springer-Verlag.

Hamilton, J. (2009). Causes and consequences of the oil shock of 2007–08. *Brookings papers on economic activity* (pp. 215–261).

Hamilton, J. (2013). Oil prices, exhaustible resources, and economic growth. In R. Fouquet (Ed.), *Handbook of energy and climate change*. Edward Elgar Publishing.

Hirsch, R., Bezdek, R., & Wendling, R. (2005). *Peaking of world oil production: impacts, mitigation and risk management*. United States Department of Energy.

Hirsch, R., Bezdek, R., & Wendling, R. (2010). *The impending world energy mess*. Apogee Prime.

Hubbert, M. K. (1956). Nuclear energy and the fossil fuels. In: *American Petroleum Institute Drilling and Production Practice Proceedings* (pp. 5–75).

Hubbert, M. K. (1962). *Energy resources*. Washington, DC: National Research Council, National Academy of Sciences, 1000-D.

Hubbert, M. K. (1967). Degree of advancement of petroleum exploration in United States. *American Association of Petroleum Geologists Bulletin*, *51*, 2207–2227.

Hubbert, M. K. (1982). Techniques of prediction as applied to the production of oil and gas. In S. I. Gass (Ed.), *Special publication*: *Vol. 631*. *Oil and gas supply modeling* (pp. 16–141). Washington: National Bureau of Standards.

IMF (2011). Oil scarcity, growth and global imbalances. In *World economic outlook* (Chapter 3) International Monetary Fund.

Kaufmann, R. K. (1991). Oil production in the lower 48 states: reconciling curve fitting and econometric models. *Resources and Energy*, *13*, 111–127.

Kaufmann, R. K., & Cleveland, C. J. (2001). Oil production in the lower 48 states: economic, geological and institutional determinants. *Energy Journal*, *22*(1), 27–49.

Kilian, L. (2008). The economic effects of energy price shocks. *Journal of Economic Literature*, *46*(4), 871–909.

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.

Kilian, L., & Hicks, B. (2013). Did unexpectedly strong economic growth cause the oil price shock of 2003–2008? *Journal of Forecasting*, *32*(5), 385–394.

Kim, I.-M., & Loungani, P. (1992). The role of energy in real business cycle models. *Journal of Monetary Economics*, *29*(2), 173–189.

Kümmel, R. (2011). *The second law of economics—energy, entropy, and the origins of wealth*. Springer-Verlag.

Kümmel, R., Henn, J., & Lindenberger, D. (2002). Capital, labor, energy and creativity: modeling innovation diffusion. *Structural Change and Economic Dynamics*, *13*, 415–433.

Kumhof, M., & Muir, D. (2014). Oil and the world economy: some possible futures. *Philosophical Transactions of the Royal Society A*, *372* (in press).

Pesaran, M. H., & Samiei, H. (1995). Forecasting ultimate resource recovery. *International Journal of Forecasting*, *11*(4), 543–555.

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., & 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.