Research

CrossMark
click for updates

Author for correspondence:
Michael Kumhof
e-mail: mkumhof@imf.org

Oil and the world economy: some possible futures

Michael Kumhof[1] and Dirk Muir[2]

[1]International Monetary Fund, Research Department, Suite 11B-338, 1900 Pennsylvania Avenue NW, Washington, DC 20431, USA
[2]International Monetary Fund, Research Department, Suite 9-603G, 700 19th Street NW, Washington, DC 20431, USA

This paper, using a six-region dynamic stochastic general equilibrium model of the world economy, assesses the output and current account implications of permanent oil supply shocks hitting the world economy. For modest-sized shocks and conventional production technologies, the effects are modest. But for larger shocks, for elasticities of substitution that decline as oil usage is reduced to a minimum, and for production functions in which oil acts as a critical enabler of technologies, output growth could drop significantly. Also, oil prices could become so high that smooth adjustment, as assumed in the model, may become very difficult.

1. Introduction

Over the past decade, the world economy has experienced a persistent increase in oil prices. While part of this has been owing to continued rapid demand growth in emerging markets, stagnant production has also played a major role. Figure 1 shows the sequence of downward shifts in the historical trend growth rate of world production of crude oil since the late 1960s. The latest trend break occurred in late 2005, when the average growth rate of 1.8% per annum of the 1981–2005 period could no longer be sustained, and production entered a fluctuating plateau that it has maintained ever since.

This paper attempts to analyse the implications of downward shifts in the growth rate of world oil production for the world economy.[1] The focus is on gross domestic product (GDP), current accounts and oil prices. We use simulation analysis based on the International Monetary Fund (IMF)'s Global Integrated Monetary and Fiscal Model (GIMF), a six-region dynamic general equilibrium model of the world economy that is

---

[1]The paper represents a further development of the analysis contained in ch. 3 of the April 2011 IMF *World Economic Outlook* [1].
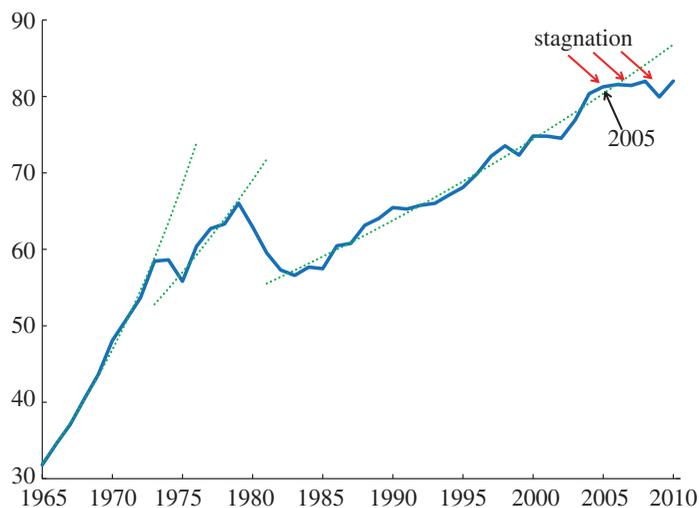
Figure 1. World crude oil production (in million barrels per day). (Online version in colour.)

frequently used at the IMF for policy and scenario analysis. In GIMF, oil is a separate input into production (in addition to capital and labour) and consumption (in addition to other goods and services), with price elasticities of demand and of production that are empirically based and very low.[2]

The paper, following many papers in the economics literature, does not separately introduce non-oil energy sources as inputs into production. The reason, apart from limiting model complexity, is that we think of the economic role of oil broadly, as the services provided by oil's most important applications. To the extent that non-oil substitutes are available in these applications, this affects our judgement concerning the appropriate model calibration of the scarcity of oil, and of its substitutability with capital and labour.

The analysis begins with a simulated baseline scenario in which the economy experiences a negative shock to oil production.[3] This scenario makes two assumptions. First, the reduction in the trend annual growth rate of world oil production, from 1.8% per annum to 0.8% per annum, is relatively modest but highly persistent. Second, a conventional macroeconomic model, with oil entering the economy's production and consumption technologies as part of simple constant elasticity of substitution (CES) aggregators,[4] is adequate under conditions of increasing oil scarcity. We find that under those two assumptions oil scarcity may not be a major constraint on global growth, nor would it dramatically worsen current account imbalances. Under an alternative upside scenario, we find that if long-run price elasticities of oil demand are increasing functions of the oil price, the effects on growth and current account imbalances are even smaller. This is followed by an analysis of three downside scenarios. The first downside scenario features an explicit lower limit, or boundary, on the amount of oil use per unit of production and consumption, with the price elasticity of oil demand approaching zero as the economy approaches the boundary owing to reduced availability of oil. The second downside scenario features a contribution of oil to economic output that is much higher than suggested by its small historic share of total input costs, because the availability of oil is a critical precondition for the continued

---

[2]The elasticity of variable a_t with respect to variable b_t is given by $\partial \ln(a_t)/\partial \ln(b_t)$, holding all other variables in the system under consideration constant. It measures the change in a_t, in percent, in response to a 1% change in b_t.

[3]In the economics literature, a 'shock' is a sudden and unanticipated change in one of the economy's exogenous driving forces, in this case the growth rate of world oil supply. The term exogenous refers to the fact that a variable is fixed outside the model, that it is not a function of economic decisions within the model.

[4]We use the word aggregators rather than production functions because oil also enters the consumption bundle of households.

viability of many key technologies that contain materials or use fuels derived from oil.[5] The third downside scenario features a reduction in the growth rate of world oil production that is assumed to be much larger than in the baseline scenario, at 3.8 percentage points rather than 1 percentage point. The resulting 2% per annum absolute decline in world oil output is in line with several recent forecasts in the scientific literature. These downside simulations, alone but especially in combination, can lead to a reduction in the growth rate of world GDP of several percentage points. But more ominously, they predict oil prices of such magnitude that a smooth adjustment, as assumed in the model, can no longer be taken for granted.

Our paper, based on empirical evidence, pays serious attention to the view that geology may at some point in the not-too-distant future start to constrain world oil production,[6] but without taking a stand on the precise year in which this will occur (or may already have occurred). According to this geological view, 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 cumulative amount of oil already produced grows. According to many scientists who advocate this view, 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 oil prices on oil production. Furthermore, they argue, 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 over the required time horizon,[7] and that technologies to improve oil recovery from existing fields, and to economize on oil use, must eventually run into limits dictated by the laws of thermodynamics. This view of oil production traces its origins back to the work of Hubbert [7], a geoscientist who in 1956 correctly predicted that US oil production would peak around 1970. It is discussed in a study produced for the US Department of Energy[8] [10] and in a subsequent book [11]. The most thorough scientific research available on this topic is with UK Energy Research Centre [12], which is succinctly summarized in Sorrell *et al.* [13]. 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, and a near certainty of a peak by 2030, with an inexorable decline thereafter. Waisman *et al.* [6] come to very similar conclusions. Given that this still allows for a wide range of possible dates for the next major trend break, we will not be specific concerning the year at which the shock hits the world economy in our simulations.

The rest of the paper is organized as follows. Section 2 presents the relevant details of the model and its calibration. Section 3 discusses detailed rationales for the main alternative scenarios. Section 4 presents simulation results. Section 5 concludes.

2. The model

(a) Brief introduction to neoclassical production theory

The basic building blocks of a neoclassical dynamic general equilibrium model are firms and households that maximize profits and utility, subject to budget constraints. Utility and production functions are assumed to satisfy standard conditions, most importantly concavity,[9] that guarantee unique optima for the respective optimization problems. Models are closed by way of a set of

[5]See Ayres [2,3] and Wrigley [4] for excellent discussions of the historically critical role of energy and oil in facilitating technological advances.

[6]For a detailed exposition of this view, see Aleklett [5]. Other authors have stressed the importance of technical uncertainty and adjustment costs, e.g. Waisman *et al.* [6].

[7]As we will discuss, existing technologies may permit significant substitution away from oil towards gas and coal once oil prices reach very high levels. This may well delay the moment at which oil supply problems start to have serious effects, but probably by years rather than decades.

[8]Other studies by official US agencies that have warned about this issue include United States Government Accountability Office [8] and United States Joint Forces Command [9].

[9]A function $f(.)$ is strictly concave if, for any positive weights ω_1 and ω_2 such that $\omega_1 + \omega_2 = 1$, then $\omega_1 f(a_1) + \omega_2 f(a_2) < f(\omega_1 a_1 + \omega_2 a_2)$.

market clearing conditions that equate the supplies of labour, capital, raw materials and produced goods to their demands.

The model used in this paper differs from other dynamic general equilibrium models primarily in its specification of production functions, specifically in the way in which oil enters these functions. To clarify this, the following provides a brief introduction to neoclassical aggregate production functions and related concepts.

In neoclassical economics, a production function relates the physical inputs of a production process to the physical output. The primary purpose of this production function is to address *allocative* efficiency in the use of inputs and the resulting distribution of income to those inputs, while making no attempt to represent the *physical* realities of real production processes, including most importantly the laws of thermodynamics.[10]

The neoclassical model of economic growth goes back to Solow [14]. It consists of a production function with only capital and labour inputs, where the productivity of those inputs is assumed to increase over time as a consequence of 'technology'. To model long-run economic growth consistently with the observed national accounts shares of capital and labour in total input costs, and with the empirical features identified by Kaldor [15], Solow [14] and the subsequent literature have assumed production functions that exhibit constant returns to scale, which means that a 1% increase in both capital and labour increases output by exactly 1%.

Economic models use different functional forms to relate inputs to economic output while satisfying the requirement of constant returns to scale. The choice of functional form is generally based on a combination of analytical tractability and empirical realism for the particular production process at hand, including most importantly existing knowledge concerning the ease of substitution between different inputs. The two most widely used functional forms are the Cobb–Douglas (CD) function and the CES function.[11] We denote inputs by x_t^j, $j = 1, 2, \ldots$, input prices divided by the price of output by p_t^j, $j = 1, 2, \ldots$, and output by y_t^i, $i \in \{CD, CES, \ldots\}$, where the subscripts t denote time.

Then the two-input CD function is given by

$$y_t^{CD} = (x_t^1)^{\alpha_1}(x_t^2)^{\alpha_2}, \tag{2.1}$$

where $\alpha_1 + \alpha_2 = 1$. The latter condition ensures constant returns to scale. The two-input CES function, which also exhibits constant returns to scale, is given by

$$y_t^{CES} = ((\gamma_1)^{1/\epsilon}(x_t^1)^{(\epsilon-1)/\epsilon} + (\gamma_2)^{1/\epsilon}(x_t^2)^{(\epsilon-1)/\epsilon})^{\epsilon/(\epsilon-1)}, \tag{2.2}$$

where $\gamma_1 + \gamma_2 = 1$. Since the 1970s, it has become more common to use production functions with multiple inputs, notably energy. We will encounter one such form below, which nests a CD function within a CES function. In general terms, this three-input function can be written as

$$y_t^{NESTED} = ((\gamma_1)^{1/\epsilon}((x_t^1)^{\alpha_1}(x_t^2)^{\alpha_2})^{(\epsilon-1)/\epsilon} + (\gamma_2)^{1/\epsilon}(x_t^3)^{(\epsilon-1)/\epsilon})^{\epsilon/(\epsilon-1)}. \tag{2.3}$$

The marginal product mp_t^j of an input j is the extra output that can be produced by using one more unit of the input j, assuming that the quantities of other inputs to production are held constant. The marginal product is the first derivative of the production function with respect to that input. For example, for (2.1) we have

$$mp_t^1 = \frac{\partial y_t^{CD}}{\partial x_t^1} = \frac{\alpha_1 y_t^{CD}}{x_t^1}. \tag{2.4}$$

---

[10]With this emphasis on allocation rather than on actual physical processes, it is therefore not inappropriate to think of neoclassical production theory as the science of the shopkeeper, as opposed to the science of the factory owner.

[11]Kümmel *et al.* [16] discuss the much more general LINEX production function, which also satisfies constant returns to scale.

As long as individual firms take prices as given, because their production is small relative to the overall level of production, which in economic jargon is referred to as the assumption of perfect competition, it will furthermore be true that $p_t^1 = mp_t^1$. This in turn implies that

$$\alpha_1 = \frac{p_t^1 x_t^1}{y_t^{CD}} \tag{2.5}$$

and that

$$y_t^{CD} = p_t^1 x_t^1 + p_t^2 x_t^2. \tag{2.6}$$

In other words, the coefficients α_1 and α_2 correspond exactly to the input cost shares of inputs x_t^1 and x_t^2 (equation (2.5)),[12] and the sum of input costs equals the sales revenue of each firm (equation (2.6)). Under perfect competition, the latter result holds for any production function that exhibits constant returns to scale, including the CES function.

Furthermore, it can be shown that, under constant returns to scale, the elasticity of output with respect to an input $\partial \ln(y_t^i)/\partial \ln(x_t^j)$ is exactly equal to that input's cost share. This is obvious for CD functions from (2.1), but it also holds for CES functions, and other functions exhibiting constant returns to scale. It follows that reductions in the availability of inputs that account for a small share of total costs, which include oil, should not have a significant effect on economic output.

The elasticity of substitution eos_t measures how a producer's relative choices over inputs change as the ratio of their marginal products change, holding output constant. Under the assumption of perfect competition, the ratio of marginal products of two inputs can furthermore be replaced by the relative prices of the same inputs. Intuitively, the elasticity of substitution therefore measures how easy it is to substitute one input for another if the latter becomes relatively more expensive. Formally, the formula for eos_t for two-input production functions (2.1) and (2.2), holding output constant at $\bar{y}^i$, and assuming perfect competition, is

$$eos_t^i = \frac{d \ln(x_t^2/x_t^1)}{d \ln(p_t^1/p_t^2)}\bigg|_{y_t^i = \bar{y}^i}. \tag{2.7}$$

This measures the change in the input ratio, in per cent, in response to a 1% change in the relative price ratio. It can be shown that for the CD case we have $eos^{CD} = 1$, while for the CES case we have $eos^{CES} = \epsilon$. Note that, for both functional forms, the substitutability among inputs is independent of the current input mix, and of time. Expression (2.7) also shows that the price elasticity of demand for good 1, which holds output y_t^i and the quantity and price of good 2 constant (see footnote 2), is equal to the elasticity of substitution and is independent of the current input mix for both CD and CES functions.

Two-input production functions can be represented graphically as in figure 2. Here, the horizontal and vertical axes, which span what we refer to as the input space, represent the physical quantities of the two inputs into production, and the curve, which is referred to as an isoquant, represents all combinations of inputs that can be used to produce a given level of output. There is of course an infinite family of such curves, which represent a three-dimensional topography in which the third dimension is output. Crucially for our later discussion, isoquants are asymptotic to the horizontal and vertical axes, which implies that any quantity of output can in principle be produced using an infinitesimally small amount of one or more of (but of course not all of) the inputs.

(b) Critiques of neoclassical production theory

Critiques of neoclassical production theory have a long history that goes back at least to the 1950s. The target of most if not all of these critiques is the above-mentioned lack of real physical content in production functions.

---

[12]In the case of CES functions, higher coefficients γ_1 or γ_2 also increase the cost shares of the respective inputs, but coefficients and cost shares are not equal.
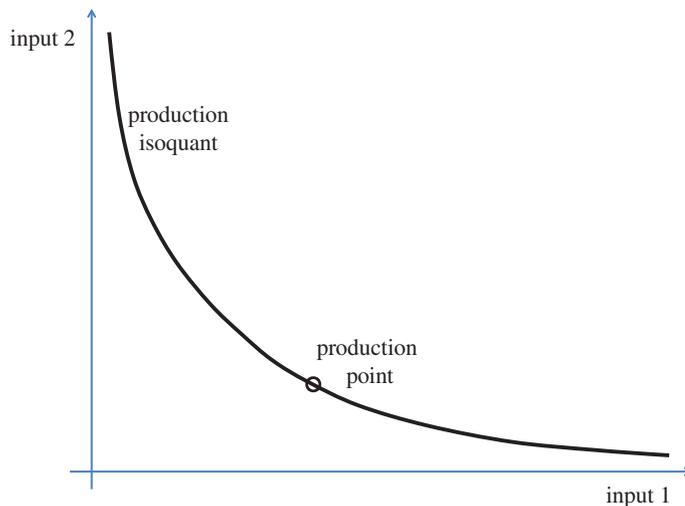
Figure 2. Input space of the neoclassical production function. (Online version in colour.)

The first example is the Cambridge (USA)–Cambridge (UK) controversies of the 1950s–1970s (see [17–19]), which centred on issues concerning the meaning and measurement of capital in neoclassical production functions. In the 1970s, various authors extended the model and methodology of Solow by adding oil or energy. Such production functions have continued to represent the dominant approach to introducing oil or energy into aggregate technologies, at least in mainstream economics journals.[13] This includes the baseline version of the IMF's GIMF model used in this paper, but also similar work at other policy institutions such as the Federal Reserve [20].

The economist Georgescu-Roegen [21] criticized this approach as a 'conjuring trick' that ignores the laws of thermodynamics, as this variant allows capital and labour to be infinitely substituted for natural resources, including energy or more narrowly oil. Daly [22] argues that any production function that obeys the laws of thermodynamics cannot avoid strict complementarity, or zero substitutability, between resources on the one hand and capital and labour on the other hand.

The critical question of the substitutability between energy and other inputs therefore started to receive attention in the empirical literature.[14] However, Frondel & Schmidt [25], in an extensive survey of such studies, which were all based on some variant of constant returns to scale neoclassical production functions, arrived at a pessimistic conclusion: '...inferences obtained from previous empirical analyses appear largely to be an artifact of cost shares,[15] and have little to do with statistical inference about technology relationships. Thus, one cannot expect to elicit the true "substitution relationships"...' [25, p. 72].

This suggests that what is needed is a return to economic theory in order to improve the specification of aggregate production functions. The literature has followed two different approaches in attempting this.

One approach started with the pathbreaking work of Georgescu-Roegen [26], which in its emphasis on the laws of thermodynamics has even older antecedents in the work of Soddy [27].

---

[13]This should be distinguished from specialized journals in energy economics, which have provided a forum for some alternative approaches.

[14]See [23,24]. To ensure that this discussion is not misunderstood by non-economists, we emphasize that while capital and labour can never substitute for energy in the sense of doing the same work as energy, they can substitute in the sense that an input combination with less energy and more capital or labour may be able to produce the same amount of economic output.

[15]Specifically, they find that the data have no chance of displaying complementarity for two factors if the cost shares of these factors are sufficiently high.

But unfortunately, this framework cannot be readily incorporated into modern dynamic general equilibrium frameworks and has since been largely neglected by both the mainstream economics literature and the energy economics literature.

The alternative approach is to continue working with static aggregate production functions, but to augment them in order to achieve a more realistic representation of real physical processes, based on an understanding that conventional constant returns to scale production functions, for sufficiently large departures from business-as-usual scenarios, provide no certainty that the laws of thermodynamics are not violated (see [28,29]). This approach is dominant in energy economics today, and uses 'hybrid models' that combine 'top-down' elements for macroeconomic consistency with 'bottom-up' elements that incorporate very detailed engineering knowledge. However, unlike in the modern macroeconomics literature, decision-making at the 'bottom-up' level is typically not modelled as an intertemporal problem.

Our approach is similar to the hybrid approach in that we aim for a more realistic modelling of real physical processes. But we do so in a highly simplified form that allows us to stay within the modern tradition of macroeconomic dynamic general equilibrium models. Specifically, one key alternative specification of our model exhibits strict complementarity, or zero substitutability, between oil and other inputs into production on the boundary of the space of feasible input combinations, and greater than zero but nevertheless very low substitutability inside that space.[16] This is therefore a compromise between the strict neoclassical view of infinite substitutability, and the alternative view of zero substitutability. Stern [31] and Reynolds [32] have presented similar arguments, including useful graphical interpretations.

Under strict complementarity, which is known as the Leontief production function, the proportions in which different inputs need to be used are completely fixed. The elasticity of output with respect to every single input is therefore exactly equal to 1, which implies that a 1% decline in the availability of oil reduces output by exactly 1%. This is the maximum long-run output decline that a decline in oil availability can cause in any model. For infinitesimal deviations from this case, in other words for elasticities of substitution between oil and other inputs extremely close to 0, any reduction in the availability of oil leads to such a dramatic increase in the oil price that oil's cost share, even if it is initially modest, quickly approaches one as oil availability keeps declining, meaning that the elasticity of output with respect to oil still approaches one. But this is no longer true for larger elasticities of substitution, because in that case the elasticity of output with respect to oil equals oil's modest cost share at the margin, and more importantly it remains modest even for sizeable increases in oil prices, given that the cost share remains small. Given the very low initial cost shares of oil, this implies a low importance of oil for aggregate output. This feature has however recently also been criticized by a number of important papers and books in the natural sciences, including Ayres & Warr [33,34], Hall & Klitgaard [35], Kümmel [36] and Kümmel et al. [37], who propose specifications of production functions with positive elasticities of substitution, but where elasticities of output with respect to energy can be much larger than cost shares of energy. An alternative specification of our model will pursue a similar idea.

(c) Overview of the Global Integrated Monetary and Fiscal Model

The theoretical structure and calibration of the IMF's GIMF is fully documented in Kumhof et al. [38]. GIMF includes several features found to be important for replicating real-world behaviour, including finite planning horizons for households and firms, gradual adjustment of prices, nominal wages, consumption, investment and imports to unexpected changes in exogenous driving forces, a financial system where losses constrain borrowers' future activity through higher financing costs, and a fully specified fiscal sector. The version used here has six economic

---

[16]Hassler et al. [30], for the case of conventional CES production functions in oil and capital/labour, make a convincing case that very low substitutability inside the space of feasible input combinations is critical to explain the historical data.

regions—oil exporters, the USA, the euro area, Japan, emerging Asia and remaining countries. For reporting purposes these will be aggregated as 'oil exporters', 'USA and the euro area' and 'rest of the world'. All regions are assumed to have flexible exchange rates.

In GIMF, oil is the third input into production, in addition to capital and labour, and the second component of final consumption, in addition to goods and services. The price and availability of oil therefore influence production as well as consumption possibilities. The updated version of GIMF used in this paper extends the theoretical framework of Kumhof *et al.* [38] in a number of ways that are discussed in detail in §3. The extensions are: first, allowing for elasticities of substitution between oil and other inputs that can rise with the oil price; second, putting 'entropy boundary' (the term is owing to [32]) limits on the extent to which other inputs can substitute for oil in both production and consumption; and third, allowing for 'technology externalities' in production (not in consumption) that act to increase the contribution of oil to economic output.

In the model, prior to the shock to oil production, all real variables, except for labour supply, grow at the exogenous and constant rate of technological progress $g = T_t/T_{t-1}$. We will choose the calibration $g = 1.018$ so that the economy grows at a real rate of 1.8% per annum. Conventional techniques for model solution and simulation require that all variables be stationary—that they have a finite long-run or steady-state value. The computational representation of the model is made to satisfy this assumption by detrending by T_t, in other words by dividing all growing variables at time t by T_t. In all equations presented below, real variables are therefore shown in this detrended form. Steady-state values of this detrended model are denoted by a bar above the respective variable.

(d) Oil production

Each region's production of oil is an exogenously given quantity (in economics jargon, an endowment), except for a quantitatively small elasticity of oil production with respect to the oil price. It is given by

$$\log(O_t^{\text{prod}}) = \log(O_t^{\text{prod}*}) + \epsilon_s \log\left(\frac{p_{t-4}^{O,\text{avg}}}{\bar{p}^O}\right), \tag{2.8}$$

where O_t^{prod} is an individual region's oil production, p_t^O is the oil price in local currency and $p_t^{O,\text{avg}} = (p_t^O(p_{t-1}^{O,\text{avg}})^3)^{1/4}$. In the initial steady state of the economy, the exogenous part of oil production $\bar{O}^{\text{prod}*}$ is assumed to be a constant when normalized by T_t. In other words, oil production prior to the simulated shock grows at the same rate as output. The simulations will subject $O_t^{\text{prod}*}$ to exogenous shocks whereby oil production grows, for three decades, at a significantly lower rate than its historic trend growth rate of 1.8% per annum.

The long-run price elasticity of oil production (or oil supply) is given by ϵ_s. Specifically, equation (2.8) states that the output of oil rises by ϵ_s percent for a 1% deviation of the lagged moving average of oil prices $p_{t-4}^{O,\text{avg}}$ from the initial steady-state oil price $\bar{p}^O$. The particular parametrization of the moving average expression adopted here implies that the moving average oil price $p_t^{O,\text{avg}}$ starts to fully reflect permanent changes in the actual oil price p_t^O after a period of around 5 years.[17] Also, the moving average term itself enters with a 4-year lag, because we are interested in capturing the effects of higher oil prices on exploration activity and new field development, and it is well known in the industry that the lead time for bringing new capacity online equals 4 years or more. It is possible to also introduce a responsiveness of oil production

---

[17]We use a moving-average rather than a single-period oil price to ensure that only long-lasting price changes affect the growth of oil production capacity, while short-run price fluctuations have only muted effects.

to the current oil price, which would correspond to oil producers using existing spare capacity[18] when prices are favourable. However, we decided not to pursue this, because persistently high spare capacity would be very unlikely to occur under the scenarios we study here.[19]

(e) Oil demand

(i) Baseline scenario

The economy has three sectors that require oil, the non-tradables (superscript N) and tradables (superscript T) manufacturing sectors and the consumption (superscript C) sector. We use the general notation J for sectors $J \in \{N, T, C\}$. The baseline production function is a CES aggregate over oil O_t^J and a CD composite M_t^J consisting of capital K_t^J and labour L_t^J,[20] and with an elasticity of substitution of ϵ_d that for simplicity is assumed to be equal across sectors.[21] Finally, an adjustment cost $G_{O,t}^J$ makes it costly to rapidly vary the use of oil in response to shocks. This has the effect of making the short-run price elasticity of oil demand lower than the long-run price elasticity ϵ_d, the more so the larger is the coefficient ϕ_O. The detrended production function for the baseline scenario is given by

$$
\left.
\begin{aligned}
Z_t^J &= ((1 - \eta^J)^{1/\epsilon_d}(M_t^J)^{(\epsilon_d-1)/\epsilon_d} + (\eta^J)^{1/\epsilon_d}(O_t^J(1 - G_{O,t}^J))^{(\epsilon_d-1)/\epsilon_d})^{\epsilon_d/(\epsilon_d-1)}, \\
M_t^J &= (K_t^J)^{1-\alpha}(L_t^J)^\alpha
\end{aligned}
\right\}
$$

and
$$
G_{O,t}^J = \frac{\phi_O}{2}\left(\frac{O_t^J - O_{t-1}^J}{O_{t-1}^J}\right)^2.
$$
(2.9)

(ii) Growing elasticity scenario

Here, we assume the same specification as in (2.9), but with the long-run elasticity ϵ_d replaced by a time varying elasticity $\epsilon_{d,t}$ that is given by

$$
\log(\epsilon_{d,t}) = \log(\epsilon_d) + \epsilon_\epsilon^p \log\left(\frac{p_t^{O,\text{avg}}}{\bar{p}^O}\right).
$$
(2.10)

The parameter ϵ_ϵ^p, which equals 0 in the baseline scenario, and 1 or 2 in the growing elasticity scenario, is the price elasticity of the price elasticity of oil demand. In other words, it measures the percentage change of the price elasticity of oil demand in response to a 1% increase in the average oil price $p_t^{O,\text{avg}}$. For $\epsilon_\epsilon^p = 1$, a doubling of the oil price relative to its initial steady state leads to a doubling of the long-run price elasticity of oil demand.

(iii) Entropy boundary scenario

In this alternative specification, oil inputs per unit of production (and consumption) are constrained to be above a minimum level given by an entropy boundary. Specifically, the

---

[18]Spare capacity is the amount of oil that producers could technically bring to market immediately if prices (and politics) justified it. The evidence shows that when officially reported Organization of the Petroleum Exporting Countries (OPEC) spare capacity drops below 2 million barrels per day, oil prices start to increase sharply. In other words at that level, which was reached in 2008 and approached more recently, the price elasticity of oil production drops sharply.

[19]Furthermore, globally significant spare capacity in Saudi Arabia, for example, is more theoretical than practical, because a large portion of it is 'Arab Heavy' oil, a grade not easily refined except by dedicated refineries, and therefore nearly useless in a shortage of light sweet crude oil.

[20]Similarly, the consumption aggregator is a CES aggregate over oil and non-oil consumption goods.

[21]The aggregate price elasticity of oil demand are available. We are not aware of estimates of its sectoral counterparts for the sectoral decomposition chosen in our model.
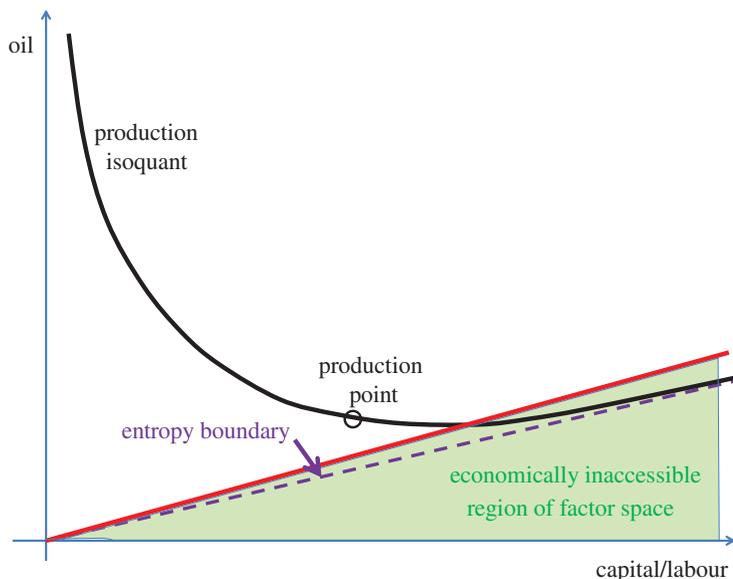
10

rsta.royalsocietypublishing.org Phil. Trans. R. Soc. A **372**: 2012032

Figure 3. The entropy boundary in input space. (Online version in colour.)

production function of an individual firm takes the form

$$Z_t^J = \left((1 - \eta^J)^{1/\epsilon_d^J} (M_t^J)^{(\epsilon_d^J - 1)/\epsilon_d^J} + (\eta^J)^{1/\epsilon_d^J} \left(O_t^J (1 - G_{O,t}^J) - \beta \frac{\bar{O}^J}{\bar{M}^J} \tilde{M}_t^J \right)^{(\epsilon_d^J - 1)/\epsilon_d^J} \right)^{\epsilon_d^J/(\epsilon_d^J - 1)}, \qquad (2.11)$$

where M_t^J is the firm-specific capital–labour bundle and is defined as in (2.9), $\bar{M}^J$ is the original steady-state value of M_t^J, and $\tilde{M}_t^J$ is the aggregate capital–labour bundle, which is taken as exogenous by the firm when deciding on its optimal input use.[22] This means that an individual firm acts as though it cannot affect the overall scale of production, while in equilibrium we have that $M_t^J = \tilde{M}_t^J$, because all firms have the same objective function and constraints, and therefore behave identically.[23]

Production function (2.11) differs from the previous specification in two ways. First, the coefficients ϵ_d^J vary across sectors. Second, oil inputs are modified by the term $-\beta(\bar{O}^J/\bar{M}^J)\tilde{M}_t^J$, which represents the entropy boundary, with $(\bar{O}^J/\bar{M}^J)$ denoting the initial steady-state ratio of oil inputs to other inputs, β denoting the proximity of the entropy boundary input ratio to that ratio and $\tilde{M}_t^J$ representing the macroeconomic scale of production as measured by the aggregate capital–labour bundle. With $\beta = 0$, this model reverts to standard specification (2.9), but as β approaches one the initial steady state moves closer and closer to the boundary. Figure 3 illustrates the input combinations, the entropy boundary and the inaccessible region of the input space for this specification. We note that the production isoquants in this model are still asymptotic to the boundaries of the input space, as in figure 2, but because that boundary is not horizontal for oil, the isoquants become upward-sloping before they approach it. Because the upward-sloping portions of isoquants are not compatible with cost minimization,[24] the economically feasible region of the input space is smaller than that demarcated by the entropy boundary itself.

---

[22]The aggregate production function is identical to (2.11), but with firm-specific magnitudes M_t^J and O_t^J replaced by aggregate quantities.

[23]Households' consumption aggregator takes the same form as (2.11), but with M_t^C representing non-oil consumption goods rather than capital and labour.

[24]In these regions, the producer could produce the same amount of output with less of each input.

For the baseline scenario, the long-run elasticity of substitution between O_t^J and M_t^J is simply given by ϵ_d. For the alternative specification, the formula for the long-run elasticity of substitution is now more complex and time-dependent, namely $\varepsilon_{d,t}^J = \epsilon_d^J (1 - \beta(\bar{O}^J/\bar{M}^J)(M_t^J/O_t^J)) < \epsilon_d^J$, which is endogenously determined by the intensity of the respective sector's oil use. Specifically, with a tight entropy boundary (high β) the elasticity of substitution quickly declines towards zero after a shock to oil production drives the economy even closer to the boundary, rather than remaining constant as in the baseline specification. We will calibrate the initial steady state of our model by adjusting each sector's ϵ_d^J such that the elasticity of substitution equals $\bar{\varepsilon}_d^J = \epsilon_d$, where ϵ_d is the constant long-run elasticity of substitution from the baseline scenario. Finally, the economy-wide average elasticity of substitution across the three oil-using sectors will be one of our reporting variables for this scenario. It is given by the weighted formula $\varepsilon_{d,t} = (\varepsilon_{d,t}^N O_t^N + \varepsilon_{d,t}^T O_t^T + \varepsilon_{d,t}^C O_t^C)/(O_t^N + O_t^T + O_t^C)$.

In essence, the entropy boundary scenario asserts that available quantities of oil could drive elasticities of substitution into the opposite direction from that into which the corresponding prices drive them under the growing elasticity scenario. Functional form (2.11) corresponds exactly to the graphical intuition for such a phenomenon discussed in Stern [31] and Reynolds [32].

(iv) Technology externality scenario

In this alternative specification, the production function of an individual firm is

$$
Z_t^J = \left((1-\eta^J)^{1/\epsilon_d^J} (M_t^J)^{(\epsilon_d^J - 1)/\epsilon_d^J} + (\eta^J)^{1/\epsilon_d^J} \left(\left(\frac{\tilde{O}_t^J}{\bar{O}^J} \right)^{\xi^J} O_t^J (1 - G_{O,t}^J) \right)^{(\epsilon_d^J - 1)/\epsilon_d^J} \right)^{\epsilon_d^J/(\epsilon_d^J - 1)}, \qquad (2.12)
$$

where O_t^J is firm-specific oil use, $\bar{O}^J$ is the original steady-state value of O_t^J and $\tilde{O}_t^J$ is aggregate oil use, which is taken as exogenous or external by the firm when it decides on its optimal input use.[25] This means that an individual firm acts as though it cannot affect the overall scale of oil use, while in equilibrium we have $O_t^J = \tilde{O}_t^J$, because all firms have the same objective function and constraints, and therefore behave identically. The term $(\tilde{O}_t^J/\bar{O}^J)^{\xi^J}$ represents oil-augmenting technology. This specification implies that the cost share of oil remains below its output contribution when $\xi^J > 0$. The beneficial effects of the availability of oil on the economy's technological possibilities are therefore not captured exclusively by the suppliers of oil, but rather by all inputs into production.

(f) World oil market equilibrium

Letting i index the six regions of the world economy, the market clearing condition for the world oil market is given by

$$
\sum_{i=1}^{6} O_t^{\mathrm{prod}}(i) = \sum_{i=1}^{6} (O_t^N(i) + O_t^T(i) + O_t^C(i)), \qquad (2.13)
$$

where the world oil price adjusts to equilibrate oil production and oil demand.

(g) Calibration

The long-run price elasticity of oil demand in both production and consumption is assumed to equal 0.08, while the short-run elasticity, which reflects the interaction of the long-run elasticity

---

[25]In economic jargon, an externality is a cost or benefit that is incurred by a party who was not involved as a significant buyer or seller of the goods or services causing the cost or benefit. Prices in a competitive market therefore do not reflect the full costs or benefits of producing or consuming such a product or service. In terms of formulating agents' utility or profit maximization problems, an externality is characterized by the fact that some inputs or outputs are treated by the agent as not being affected by its own decisions, even though the collective decisions of all economic agents do affect this variable.

and the size of adjustment costs, is around 0.02. This is consistent with estimates for 1990–2009 contained in Helbling *et al.* [1] and also with the Bayesian estimation results in Benes *et al.* [39].

In the baseline, the contribution of oil to output is equal to the oil cost share. Based on a careful evaluation of recent historical data averages for the six regions of the model economy, this cost share has been calibrated at between 1.5% and 4.5%, depending on the sector and region. The definition of oil used for quantifying this was taken from *BP Statistical Review of World Energy* (various issues), and includes crude oil, shale oil, oil sands and natural gas liquids.

The literature contains far fewer studies that estimate the price elasticity of oil production. But Benes *et al.* [39], who use a nonlinear Bayesian estimation methodology,[26] are able to separately identify demand and production elasticities. They find short-run price elasticities of oil production, which represent the ability of oil producers to use existing spare capacity, of between 0.05 and 0.15. They also estimate a long-run price elasticity, and find it to be only very slightly higher, at between 0.05 and 0.17. The difference between long- and short-run price elasticities, of between 0.005 and 0.02, corresponds almost exactly to the coefficient ϵ_s in our specification. Our specification omits short-run price elasticities, based on the reasonable assumption that the decline in the growth rate of world oil production in all our scenarios will eliminate spare capacity going forward. At the same time, our calibration of $\epsilon_s = 0.03$ is slightly higher than the estimate of Benes *et al.* [39].

Because we model oil production as an exogenous endowment, we need to specify how the revenue from oil sales is divided into extraction costs and payments to owners. We assume that, prior to the decline in the growth rate of oil production, 40% of oil revenue must be used to pay for intermediate goods inputs. The remaining 60%, which is referred to as the oil rent, is distributed between the private sector and the government. We assume that the real extraction cost of oil increases at a constant annual rate of 2%, which at unchanged oil prices would slowly reduce the oil rent over time. In the net oil-importing regions of our model economy, the government is assumed to receive only a very small portion of the oil rent. However, in oil exporters it is assumed to receive 90%, reflecting the fact that in many of these countries the oil sector is dominated by state-owned oil companies. Critically, we assume that governments do not immediately spend the additional funds, but that they accumulate them in a US dollar-based fund that is spent gradually over time, at a rate of 3% per annum.[27] One of the key effects of an increase in the oil price is therefore a dramatic increase in world savings owing to the low propensity to consume out of oil revenues of oil exporters' governments.

3. Further discussion of the alternative specifications

(a) Growing elasticity scenario

Section 3*b*,*c* will discuss scenarios in which shocks to oil production can have more serious effects on output and current accounts than in the baseline scenario. But in the near future, substitution away from oil and towards gas or electricity may be sufficient to delay the moment at which lower oil availability could cause more serious economic problems. Given the proven technical feasibility of some substitution technologies, such as gas-powered vehicles, coal-to-liquids and gas-to-liquids, it may therefore be reasonable to assume that price elasticities of demand have scope to rise for a number of years. It is for this reason that we include the growing elasticity scenario in this study.

(b) Entropy boundary scenario

Economists generally assume that elasticities of substitution between oil and other inputs into production must be higher in the long run than in the short run. There are two possible reasons for

[26]For an exposition of this methodology, see ch. 10 in Durbin & Koopman [40].

[27]This matches the modus operandi of the Norwegian oil fund very closely. See Norwegian Ministry of Finance [41].

this view. First, as in our baseline scenario, adjustment costs drive the short-run elasticity below the constant long-run elasticity ϵ_d. Second, as in our growing elasticity scenario, for persistently very high oil prices, ϵ_d may no longer be constant, but may instead grow with oil prices.

On the other hand, several contributions from the natural sciences have objected that the assumption of a constant or growing long-run elasticity is not consistent with the historical facts [42], with real-world practical limits to substitution [43] or with the laws of thermodynamics [22,31,32]. In this paper, we mathematically formalize Reynolds' entropy boundary by way of production function (2.11). This implies that after a shock to oil production elasticities are very low in the short run (owing to adjustment costs), significantly higher in the medium run (as adjustment costs are overcome), but potentially much lower again in the long run if the shock is sufficiently large, because there is a finite limit to the extent that machines (and labour) can substitute for oil.

Entropy relates to energy rather than specifically to oil. It therefore affects the scope for substitution between energy and capital/labour. Entropy, or the second law of thermodynamics, states that any ordered system of matter naturally tends towards disorder through energy dissipation. Maintaining the system therefore requires the constant addition of a flow of energy. Any capital or labour used to substitute away from energy is not independent of energy, but rather needs a continuing input of energy in order to maintain it. Critically, this implies that only energy in excess of a certain minimum amount can contribute to the production of goods and services. This puts a boundary on the input combinations that can be realized, with near-zero energy input not an option. When actually used input combinations approach that boundary, the elasticity of substitution between energy and those inputs must also approach zero. We do not take a stand on how close the minimum required energy input may be to the actually used energy input— we only argue that the implications of different assumptions concerning this relationship should be explored.

To relate this concept to feasible input combinations between oil and capital/labour, which is what interests us in this paper, we need to consider the relationship between oil and energy. The point is that, if oil and other forms of energy have to be used in proportions that cannot change sufficiently, or sufficiently rapidly, an entropy boundary for energy implies an entropy boundary for oil. The two critical questions are therefore whether other energy sources can technically substitute for oil over realistic time horizons and at the required scale, and whether such energy sources have their own supply limitations.

(i) Supply limitations

The renewables solar and wind exhibit limited substitutability with oil, primarily because they can only be used to generate electricity, and electricity exhibits limited substitutability with oil— see §3*b*(ii). But solar and wind also exhibit limited substitutability with other forms of electricity generation, because of their high intermittency. Kunz [44] therefore argues that solar and wind cannot supply more than 20–30% of overall electricity without causing serious problems for grid stability, while WWF [45] suggests that much higher contributions can be accommodated. Storage or smart grids will play a key role here, but at significant cost, and probably not to the point where intermittency ceases to be an issue altogether. Furthermore, solar and wind cannot be operated without a large and costly backup capacity of fossil-fuel-operated power plants to match demand when there is no wind or sun. Another problem with solar and wind is that their energy return on energy invested (EROEI) is generally much lower than that of conventional oil.[28]

Biofuels are of course much more substitutable for oil than other renewables. But their problems include technical limitations on use as a commercial fuel, a low EROEI, the demands placed upon scarce water resources, the very large quantities of land required, and the potential negative impacts on global food production (see [47]).

As for coal, several authors have recently claimed that coal reserves have been greatly overstated in a number of countries (see [48–50]), and put the global peak of coal production about

---

[28]The subject of EROEI, or EROI, is covered elsewhere in this Theme Issue (see [46]).

20 years away. The peak of coal production would of course come much faster if an attempted substitution towards coal was to be of a large scale.[29]

The peak of global gas production may be further away, especially with the hydraulic fracturing (fracking) of shale gas formations that has recently been much publicized. But fracking is claimed in some quarters to have serious problems with environmental pollution and has been stopped in a number of countries and jurisdictions for that reason. Shale gas development is also costly because the rate of production from individual wells declines rapidly, creating the need for continuous drilling of new wells to maintain aggregate production. Berman [52] contains an excellent overview of the problems with natural gas for the US case. They include an acceleration of sector-wide average decline rates of gas wells,[30] and therefore of extraction costs, and a large overstatement of field reserves for financial reporting reasons (see also Rogers [54]). If Berman's prognoses turn out to be correct, natural gas prices could soon begin an increase that would significantly reduce its cost advantage over oil, and therefore the incentive to switch from oil to gas. The expectation that cost advantage is not just large today but also highly persistent is critical, because a switch to electricity produced by natural gas would require the building of costly power plants, which have long amortization periods, while a direct switch to natural gas as a transportation fuel, which is technically among the most realistic substitution possibilities, would require the installation of a costly network of pipelines and filling stations.[31]

As for nuclear power, Dittmar [55] recently estimated that if large-scale substitution from fossil fuels towards nuclear power was attempted, the benefits would be very short lived because the world would hit a peak in uranium production in short order. But this finding is highly controversial, with others claiming that uranium resources are underestimated, and that different reactor types could eliminate fuel supply problems for the foreseeable future. In any event, the worldwide trend seems to be going away from rather than towards nuclear power, owing to the fear of low-probability but extremely high-cost events such as Fukushima.

To summarize, while the evidence is at this point far from complete, other non-renewable energy sources may also have limits on the extent to which production can be expanded, and there is at least a possibility that they may not be able to provide energy security over the long run.

(ii) Technical substitutability

Critical considerations for technical substitutability include, relative to a specific application, storability, transportability, volumetric and gravimetric density, cleanliness and, most importantly, the ability to deliver *useful* energy (see Smil [42]). An example is the impossibility of running an aeroplane with coal, which is because of coal's physical properties rather than simply its energy content.

The most important use of oil today is in transportation. The properties of oil, or more generally liquid fuels, in this application are far superior to the alternatives. In some cases, a transition towards a transportation infrastructure based on electricity, natural gas and/or liquid fuels derived from non-oil sources may be possible by retrofitting existing equipment or building differently configured new equipment, together with the associated infrastructure. But a large-scale transition would be enormously expensive in terms of dollars,[32] in terms of energy and most

---

[29]Oil, coal and natural gas currently account for about 33%, 28% and 23%, respectively, of primary energy consumption in the world [51].

[30]See also Hughes [53], who shows that both shale gas and shale oil suffer from this problem.

[31]Natural gas filling stations require significantly greater safety precautions than conventional gasoline filling stations. They also have significant practical limitations—for example, the time required to fill up a personal vehicle currently equals around 20 min.

[32]This is why coal-to-liquids or gas-to-liquids industries have not yet been established on a significant scale, except in countries that have faced economic sanctions or economic isolation.

importantly in terms of time—the transition would require several decades.[33] This, according to the best available scientific evidence [10,12], is time that we may not have.

The second main use of oil today is as a feedstock for the chemical industry. Oil products are embedded in virtually every single product we consume, either directly, or in the mining of the respective raw materials, in transportation, or in processing. In most cases, there simply is no easy substitute for oil derivatives in these applications, although in some specific cases there are possibilities to derive the chemicals from coal or gas instead.

To summarize, while the evidence here is again far from complete, sufficient substitutability away from oil in transportation and in the chemical industry can also not be taken for granted. It is certainly not sufficient to assert that substitution will save the day without conducting very extensive, and detailed, technical feasibility studies.

(c) Technology externality scenario

As discussed in §2a, conventional production functions, such as (2.9), imply that the contribution of oil to the growth in economic output is equal to the share of oil in total input costs. This has led most economists to conclude that, given its historically low cost share of around 3.5% for the US economy (see http://www.eia.gov/oiaf/economy/energy_price.html), oil can never account for a large output contraction.[34] There are two counterarguments to this. First, if oil prices were to permanently rise sharply owing to large and persistent shocks to oil production, then cost shares could become high enough to worry about further shocks even if output contributions equalled cost shares. Second, as argued in §2b, output contributions of energy/oil need not equal cost shares when the aggregate production function in equation (2.9) is replaced by a production function that does a better job at representing real physical processes.

Based on such production functions, Ayres & Warr [33,34], Kümmel [36] and Kümmel *et al.* [37] estimate that during the twentieth century the growth of energy inputs[35] has, depending on the country under consideration, accounted for up to around half of the growth in economic output—a far larger contribution than what is implied by neoclassical production functions. But the specification of technology in at least some of these papers may be problematic.[36] We have therefore adopted the alternative specification (2.12), whereby oil enters in a similar fashion to technology. This feature can yield output contributions of oil that are higher than cost shares as long as $\xi^J > 0$. This specification is theoretically sound and intuitive.

4. Simulation results

Figures 4–8 present the results of our simulations. These are nonlinear simulations that use a Newton solution algorithm.[37] Variables are shown, where applicable, relative to trend. It is important to remember this point when interpreting the simulations: when our figures show a negative value for, say, GDP, this represents a decline relative to the trend level of GDP, which grows at 1.8% per annum.

The first column of each figure shows the evolution of world oil production, the real oil price and the world average real interest rate. The second column shows real GDP, real absorption[38]

---

[33]Crassous *et al.* [56], based on detailed modelling of energy demand in a hybrid macroeconomic-cum-engineering model, find that one major constraint on decoupling oil and economic growth is the rise of the transport intensity of growth. The reason is that oil-based transportation fuels will retain a competitive advantage over other alternatives for a prolonged period.

[34]This is of course only true as long as elasticities of substitution between oil and other inputs into production do not approach zero, as under the entropy boundary scenario.

[35]Strictly speaking, Ayres and Warr focused on inputs of useful work rather than energy.

[36]Our main concern has to do with the mathematical treatment of inequality constraints.

[37]See Armstrong *et al.* [57] and Juillard *et al.* [58]. In applying the Newton algorithm we rely, to facilitate convergence, on a numerical linearization technique. Specifically, we first divide the size of the shock by a small integer, then we perform the nonlinear model simulation for this smaller shock, and finally we multiply the simulation result by the same integer.

[38]Absorption equals the sum of consumption, investment and government spending. GDP equals absorption plus net exports.
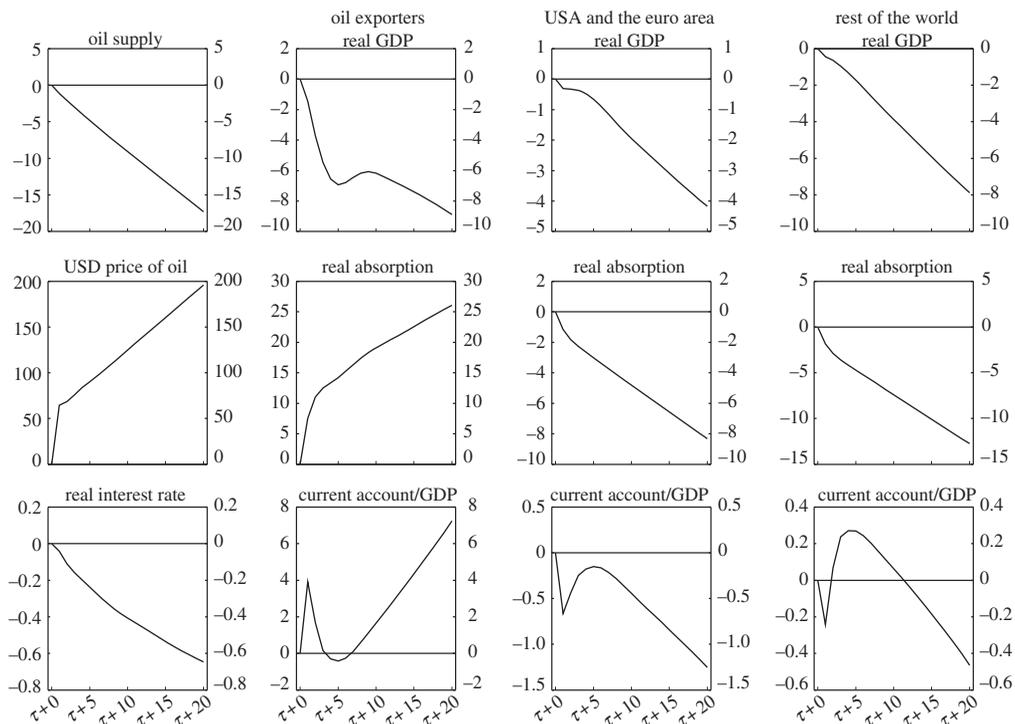
Figure 4. Baseline scenario.

and the current account-to-GDP ratio for oil exporters.[39] The third and fourth columns show the same three variables for the USA plus euro area region, and for the rest of the world group of countries. Aggregations across regions of the world economy use purchasing power parity weights to fix relative country sizes. The year in which the shock to oil production hits the economy is denoted by τ, and a range of different values for τ are justifiable.

(a) Baseline scenario

The baseline scenario, shown in figure 4, analyses the impact of a decline in the trend growth rate of world oil production by 1 percentage point below its historical trend growth rate of 1.8% per annum, starting in year τ, with an eventual return to the initial growth rate in year $\tau + 30$.[40] As in all simulations that follow, agents are assumed to be surprised by the shock, but thereafter they perfectly foresee the future evolution of world oil production.

The shock generates an immediate oil price spike of over 60%. This reflects the very low short-run price elasticity of oil demand. Because the decline in the growth rate of oil production is persistent, the real oil price continues to increase thereafter, as market equilibrium requires ongoing substitution away from oil. Over 10- and 20-year horizons, the cumulative increases in the oil price amount to just over 100 and 200%. The 10-year result is very close to the point forecast in Benes *et al.* [39].

The reduced availability of oil and the resulting higher oil prices lead to a reduction in GDP levels and to larger current account deficits in oil importers. In the short run, the global adjustment is also shaped by the wealth transfer from oil importers to oil exporters, which has effects on trade and capital flows.

---

[39] Oil exporters include the following countries: Algeria, Angola, Azerbaijan, Bahrain, Canada, Republic of Congo, Equatorial Guinea, Iraq, Kuwait, Libya, Mexico, Nigeria, Norway, Oman, Qatar, Russia, Saudi Arabia, United Arab Emirates and Venezuela.

[40] For technical reasons, it is not possible to simulate a completely permanent shock to this growth rate.

Rising oil prices imply that oil exporters experience sustained increases in income and wealth, but not higher real GDP.[41] As a result, domestic absorption increases ahead of GDP, at an initial rate of over 2% annually. Higher domestic spending leads to upward domestic price pressures and a large real appreciation.[42] This 'Dutch disease'[43] effect reduces output in the tradables sector (other than oil), thereby reducing GDP by up to 7% below trend over the first 5 years and by almost 10% after 20 years. The current account improvement in this group of countries, which equals up to 4% of GDP in the very short run and almost 8% after 20 years, is entirely because of the higher value of oil exports, with the initial spike in oil prices explaining the large current account surplus at that time. Goods exports fall substantially relative to GDP, and the non-oil current account deteriorates. But the government's very low propensity to consume out of the oil fund limits the size of that deterioration.

Oil importers' absorption contracts relative to trend as a result of the negative wealth effect[44] of higher oil prices, at an average rate of around 0.5% per annum. Their GDP also declines relative to trend, but only moderately, by 0.2–0.4% per annum.[45] World real interest rates gradually drop, and after 20 years are 0.6 percentage points below their initial value. The reason is that oil exporters' additional oil revenue, which accrues primarily to governments, leads to higher saving. This effect is reminiscent of the international lending boom in the 1970s and early 1980s following large oil price increases.

These dynamics will look somewhat unfamiliar to those accustomed to studying historical oil price shocks. For GDP, as Hamilton [60] shows, such episodes were typically characterized by a sizable but transitory contraction in oil production accompanied by a spike in oil prices, with GDP declining sharply but temporarily. Under such circumstances, real oil prices quickly return to earlier levels after the unwinding of the recession. The key feature accounting for the different predictions under our scenario is that the shock to oil production, and therefore the loss in output, is assumed to be much smaller initially, but also far more persistent than anything observed to date.

For interest rates, the earlier episodes were characterized by high inflation, not exclusively caused by higher oil prices, which led to a tightening in monetary policy. In our simulation, inflation remains subdued, partly because the oil price shock is moderate but persistent rather than large and transitory, but mainly because monetary policy is assumed to start from a situation where inflationary expectations are under control.

Global current account imbalances worsen in this scenario over the short to medium run. The USA and euro area current accounts deteriorate immediately as a result of costlier oil imports, while during a lengthy transition period the current accounts of the rest of the world remain stronger, as they export more goods to oil exporters. But they eventually also deteriorate. The long-run effects are not particularly large, however, with oil importers' current account-to-GDP ratios on average deteriorating by at most a little over 1 percentage point by year $\tau + 20$. This is explained by the relatively low aggregate cost share of oil.

[41]The reason is that real GDP is an index number that measures the real physical volume of economy-wide production across different components of output. It does so by removing the effects of price changes for these different components, including oil price changes. See Chevalier [59] for an explanation of the methodology for computing GDP. By contrast, domestic absorption, because it measures nominal expenditures divided by the aggregate price index, does reflect increases in oil prices.

[42]A real appreciation is an increase in the relative price of a representative basket of domestic versus foreign goods. This is to be distinguished from a nominal appreciation, which relates to currencies (monies).

[43]'Dutch disease' is the apparent relationship between an increase in a country's revenues from natural resources and a decline in its manufacturing sector. The mechanism is that the additional resource revenues make a nation's currency stronger in nominal and more importantly in real terms, thereby making the country's exports and thus its manufacturing sector less competitive. The term was originally coined to describe the decline in Dutch manufacturing following the discovery of a large natural gas field in 1959.

[44]In the model, consumption increases in response to increases in wealth. Wealth is equal to the present discounted value of future incomes net of expenditures, including expenditures for oil. Higher oil prices therefore reduce oil importers' wealth, and therefore their consumption.

[45]Regional differences in the size of the long-run GDP effects reflect differences in the shares of oil in production and consumption.
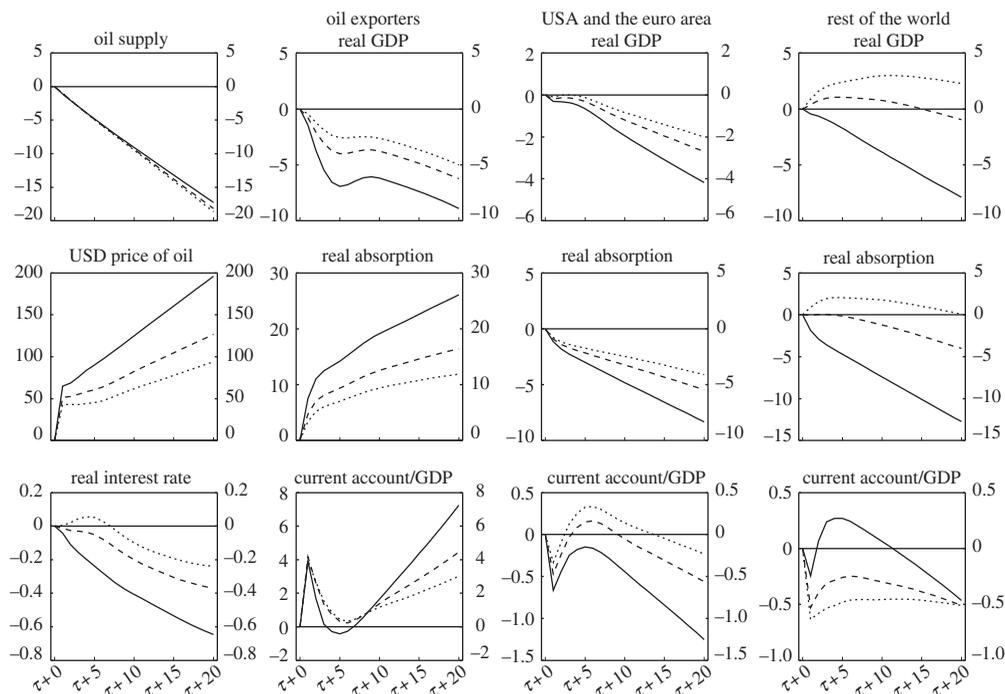
Figure 5. Growing elasticity scenario. Solid lines denote baseline scenario; dashed lines, growing elasticity scenario (parameter = 1); dotted lines, growing elasticity scenario (parameter = 2).

In the following subsections, we explore the sensitivity of our results to a number of assumptions. In most cases, the baseline scenario results are shown as a solid black line, and then compared to alternative scenarios.

(b) Growing elasticity scenario

Figure 5 illustrates how the results change, for the same shock as in the baseline scenario, if the price elasticity of demand for oil $\epsilon_{d,t}$ rises with the oil price as per equation (2.10), with the dashed line showing the case of $\epsilon_\epsilon^p = 1$ and the dotted line the case of $\epsilon_\epsilon^p = 2$. Our comments focus on the case $\epsilon_\epsilon^p = 2$. We observe that the oil price needs only to rise by half as much as in the baseline scenario to bring about the necessary substitution. Relative to the baseline the effects on GDP are positive everywhere, with roughly 50% smaller contractions in oil exporters, the USA and the euro area, and an even larger turnaround in the rest of the world countries, where an 8% output loss over 20 years turns into a 2–3% output gain.

The latter is driven mainly by emerging Asia and to a lesser extent Japan. Under the baseline scenario, these two regions, whose production is heavily manufacturing based and therefore oil-intensive, suffer very large contractionary effects of lower oil availability and higher oil prices. Under the growing elasticity scenario, where oil is less critical because of its higher substitutability, these contractionary effects are much smaller, and two countervailing effects are now strong enough to raise rather than contract output. The first of these is a surge in goods exports to oil exporters to satisfy their increasing domestic demand. Emerging Asia and Japan have particularly strong export linkages to oil exporters and therefore benefit disproportionately. The second countervailing effect is a surge in investment demand in response to lower world real interest rates, which is particularly strong in emerging Asia because of its higher steady state investment-to-GDP ratio.
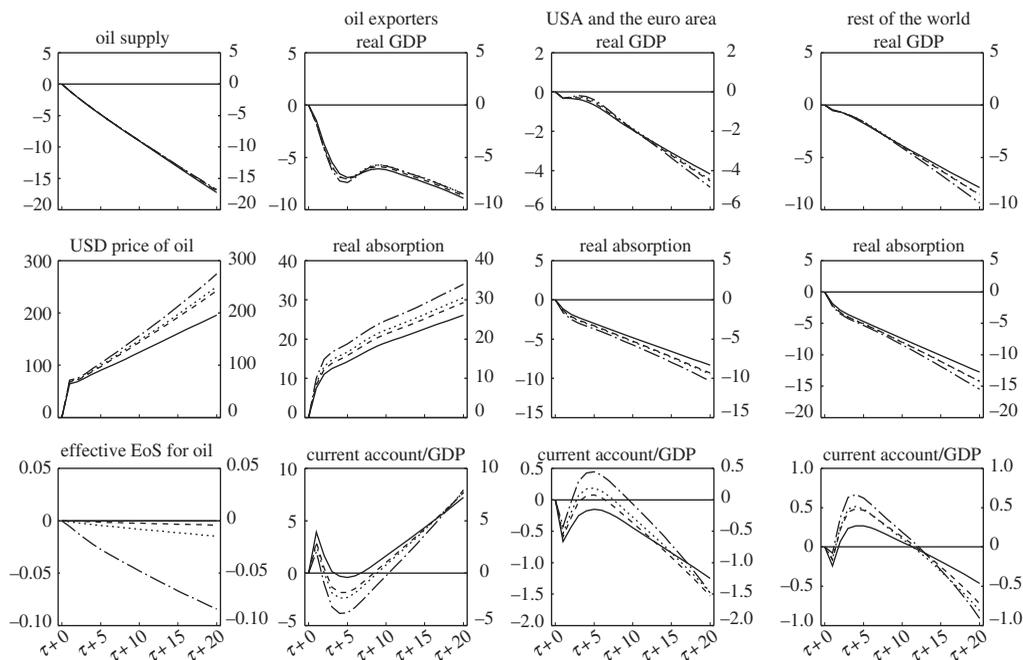
Figure 6. Entropy boundary scenario. Solid lines denote baseline scenario; dashed lines, entropy boundary scenario ($\beta = 0.3$); dotted lines, entropy boundary scenario ($\beta = 0.6$); dash-dotted lines, entropy boundary scenario ($\beta = 0.9$).

As for global current account imbalances, these are also less severe under this scenario, with smaller surpluses for oil exporters and smaller deficits for oil importers.

(c) Entropy boundary scenario

Figure 6 shows four simulations of a shock that is again identical to that of our baseline scenario, but with the technology now given by (2.11), and for the cases of $\beta \in \{0, 0.3, 0.6, 0.9\}$. The case of $\beta = 0$ is identical to the baseline scenario. For larger β the entropy boundary is tighter, in other words closer to the initial steady-state input mix, which means that further substitution away from oil becomes progressively more difficult. The bottom left panel of figure 6 shows the average elasticity of substitution $\varepsilon_{d,t}$.

We observe that for a tighter entropy boundary the increase in the oil price is significantly larger. The reason is that the elasticity of substitution moves closer to 0 as the reduced availability of oil makes further substitution away from oil harder. For the tightest boundary ($\beta = 0.9$), the oil price increases by almost 300% rather than by 200% after 20 years.

Oil exporters experience a larger positive wealth effect if β is larger. The resulting larger absorption, real exchange rate appreciation and Dutch disease effect cause a larger initial output contraction in that region, but eventually output in that region also experiences a larger rebound, because over time its governments start to spend more and more of the accumulated oil surpluses, which are now larger owing to higher oil prices. Oil importers experience worse contractions in absorption and output when the economy starts closer to the entropy boundary, and current account imbalances get larger, but the size of these effects is fairly small.

(d) Technology externality scenario

Figure 7 compares the baseline scenario to three further alternatives. The first of these assumes that the output contribution of oil is larger than its cost share because $\xi^J > 0$ in equation (2.12).
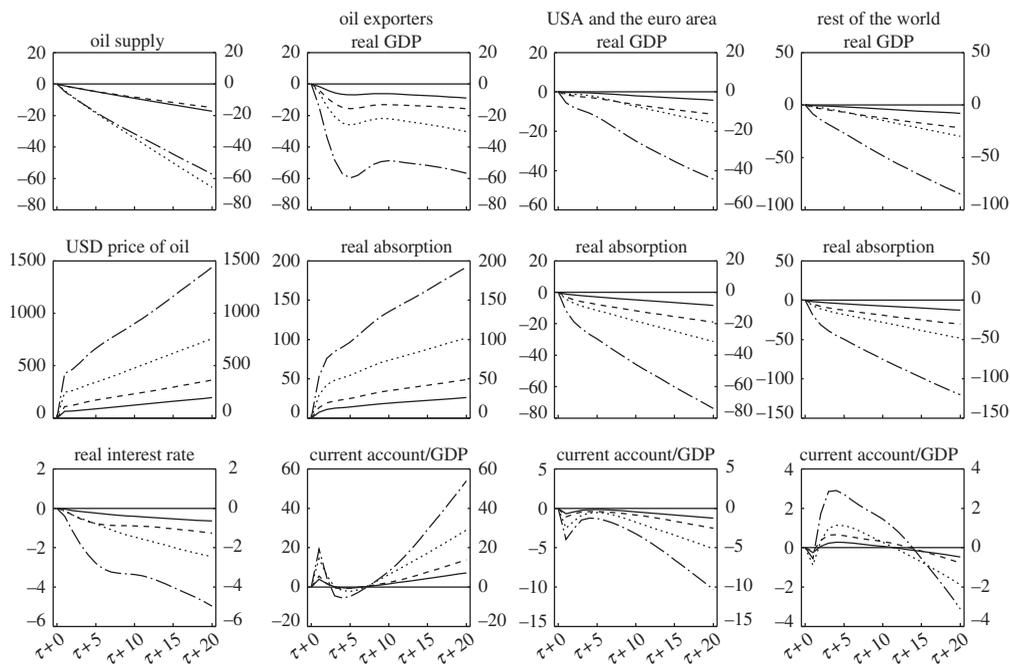
20

rsta.royalsocietypublishing.org   Phil. Trans. R. Soc. A **372**: 2012032

Figure 7. Technology externality and larger shock scenarios. Solid lines denote baseline scenario; dashed lines, technology externality scenario; dotted lines, larger shock scenario; dash-dotted lines, combined scenario: technology externality + larger shock.

Specifically, it assumes that the initial steady-state contribution of oil to output, either directly or as an enabler of technology, amounts to 25% in the tradables sector and 20% in the non-tradables sector, rather than around 4.5% and 1.5%, respectively, as in the baseline scenario. The simulation, shown as the dashed line in figure 7, shows that oil prices now increase by almost 400% after 20 years, rather than 200% as in the baseline. There are sizeable effects on growth, with the deterioration in all oil-importing regions' GDP larger by around a factor of three than in the baseline.[46] Part of the reason for the size of this effect is that with our specification both oil and the technologies dependent on oil have a very low elasticity of substitution with other inputs into production, so that a sizeable contraction in oil production cannot easily be compensated for by replacing oil with other inputs, and oil-based technologies with other technologies. A higher output contribution of oil also has significant effects on oil importers' current accounts, which in the longer run deteriorate by almost twice as much as in the baseline. Savings imbalances cause world real interest rates to drop by approximately twice as much as that in the baseline, or around 1.2 percentage points after 20 years.

(e) Larger shock scenario

Sorrell et al. [13] discuss several studies which conclude that world oil production is currently on a temporary and fluctuating plateau, and which forecast future declines in global production of around 2% per annum. The second alternative scenario in figure 7, shown as the dotted line, therefore considers the implications of a 3.8 rather than 1 percentage points annual decline in the growth rate of world oil production. Given an initial growth rate of 1.8% per annum, this implies that, except for the small production response to higher oil prices, oil production declines by 2% annually. We also assume that this outright oil output contraction is accompanied by

---

[46]Note however that these output declines are still below 1% per annum, the maximum possible decline that would obtain under a zero elasticity of substitution between oil and other factors.

an annual increase in real extraction costs per barrel of 4% rather than 2%. This is a highly conservative assumption, because recent increases in production costs have averaged around 10% per annum.[47]

In this scenario, the longer run output and current account effects are approximately four times as large as in the baseline, in other words they increase roughly in proportion to the size of the shock. Declines in absorption in oil importers now average around 2% annually over the period shown, rather than 0.5% as in the baseline. Annual GDP growth rates in the USA and the euro area drop by around one percentage point, rather than 0.25 percentage points as in the baseline. Current account deteriorations in oil importers are also much larger, averaging 5 percentage points of GDP on average in the long run in the USA and the euro area.

The most striking aspect of this scenario is however that reductions in oil production of this magnitude would require a more than 200% increase of the oil price on impact, and an 800% increase over 20 years. Relative price changes of this magnitude would be unprecedented, and would almost certainly have nonlinear effects on GDP that the model is not able to capture adequately. Furthermore, the increase in world savings implied by this scenario is so large that several regions could, after the first few years, experience nominal interest rates that approach zero, which could create difficulties for the conduct of monetary policy.

(f) Combined downside scenarios

So far we have analysed each potential aggravating factor in isolation. If they were to occur in combination, the effects could become more severe. In this paper, we only discuss one such combination, namely the technology externality scenario combined with the larger shock scenario. This is illustrated by the dash-dotted line in figure 7.

Technically, the model can still be simulated for this case, but the effects now become so large that some aspects are no longer plausible. Most importantly, real oil prices under this scenario would increase by over 400% on impact, and by around 1500% after 20 years. Despite this, there is no sharp crisis in the short run, and the subsequent reduction in annual GDP growth rates in oil importers equals a steady, crisis-free 3 percentage points.

Real-world response mechanisms to such extreme increases in oil prices could in principle take one of two possible forms. One is a much more urgent search for alternatives to oil, reflected in much higher elasticities of substitution. We study this in §4g. The other is a much sharper contraction in aggregate demand. The model in its current form is unable to deliver this. But it is clear that, if output were to contract far more sharply at the simulated oil prices, the resulting demand destruction would in turn limit the required increase in oil prices. We return to this question in §4i.

(g) Combined downside and growing elasticity scenario

In figure 8, we combine the worst case of figure 7 with the growing elasticity scenario, to simulate the much more urgent search for alternatives to oil in the presence of extremely high oil prices. For the case of $\epsilon_\epsilon^p = 2$, we observe that the increase in the oil price is only half as large, and so are, approximately, the output effects. But this still leaves an oil price increase of 800% after 20 years. More broadly, unless the increase in the price elasticity of demand in response to higher oil prices is extreme, and we have discussed above why this is not likely, then the worst of our downside scenarios would still force the economy to cope with entirely unprecedented increases in oil prices. In the real world, if such a scenario came to pass, the more likely outcome would therefore include not only higher elasticities of substitution but also a larger output contraction.

---

[47]As discussed in Mackenzie [61], Bernstein Research recently found that over the 2002–2011 period the marginal cost of oil production for non-OPEC, non-Former Soviet Union oil producers has increased at 12 and 7% compound annual real growth rates at the 90th and 50th percentiles of the marginal cost of production. Bernstein Research argues that this group of countries determines the oil price in the long run, owing to its high marginal cost levels compared with OPEC.
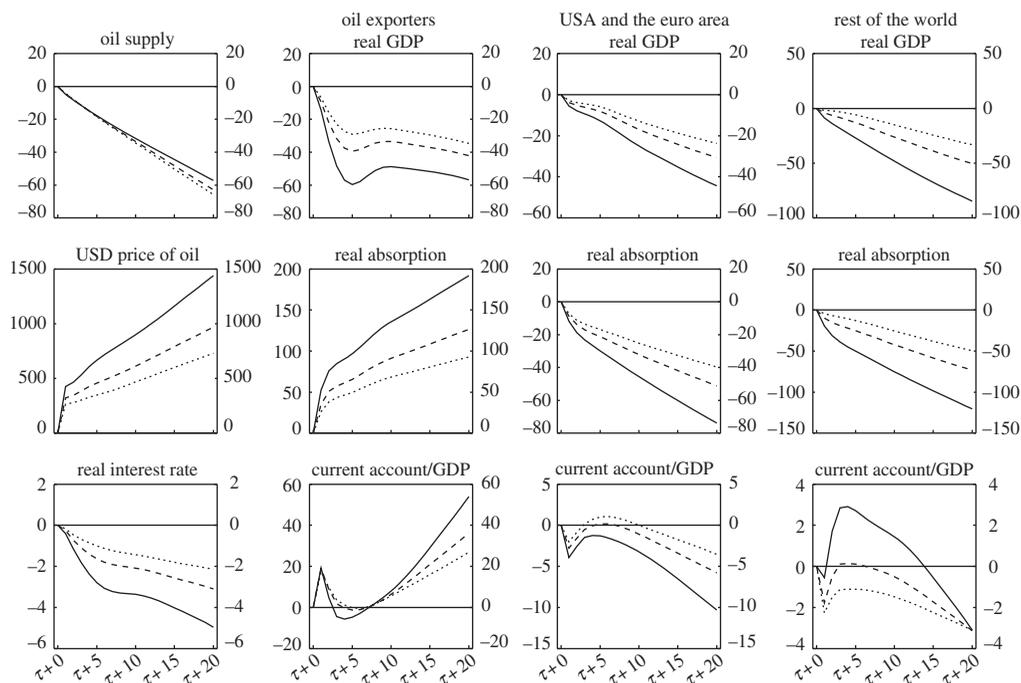
Figure 8. Combined downside and growing elasticity scenario. Solid lines denote combined scenario: technology externality + larger shock; dashed lines, combined scenario: technology externality + larger shock + growing elasticity (parameter = 1); dotted lines, combined scenario: technology externality + larger shock + growing elasticity (parameter = 2).

(h) The assumption of unitary income elasticity

Our balanced growth model, with constant returns to scale in labour, capital and oil, imposes a unitary income elasticity of oil demand, while the data suggest an income elasticity of well below one, typically around 0.7. There are two possible ways to reconcile model and data.

The first is to assume that the economy is indeed characterized by our constant returns to scale technology, but that it benefits from energy-saving technological progress that reduces the oil intensity of production over time. We have not simulated this possibility in our scenarios, because for the realistic case where energy-saving technological progress grows at a constant rate, it does not change our conclusions—only deviations from constant trends matter.

But there is a second possible explanation that is not captured by our current model at all. This is that the economy is characterized by increasing returns to scale.[48] One reason might be that most economies progress through stages that begin with agriculture and manufacturing, both of which are oil-intensive, and then proceed to a higher services component, which is less oil-intensive. If such an economy in addition remains characterized by very low elasticities of substitution between oil and other inputs, a reduction in oil availability can cause more serious problems than indicated by our simulations. Namely, if it really takes a less than one percentage point increase in oil production per annum to support additional GDP growth of one percentage point, then it must also be true that it would take a less than one percentage point decrease in oil production growth to reduce GDP growth by a full percentage point. This problem is mitigated, but not eliminated, if there is some moderate substitutability between oil and other inputs.

---

[48]Given oil's low historical cost share, the deviation from constant returns to scale could be small.

(i) The assumption of smooth reallocation

In each of the scenarios in §4*a*–*g*, the transition to a new equilibrium is assumed to be smooth. Consumers in oil exporters easily absorb large surpluses in goods exports from oil importers, financial markets efficiently absorb and intermediate a flood of savings from oil exporters, businesses respond flexibly to higher oil prices by reallocating resources and workers readily accept lower real wages. Some of these assumptions may be too optimistic.

A smooth reallocation of resources among inputs and across sectors as the economy adjusts to less oil is a very strong assumption. Unlike in the model, real economies have many and highly interdependent industries, and some of them would be affected by an oil shock much earlier and much more seriously than others. The adverse effects of large-scale bankruptcies in such industries could spread to the rest of the economy, through corporate or bank balance sheets.

Historical experience also suggests caution when it comes to the efficient intermediation by domestic banking sectors of large net capital flows from oil exporters' governments.[49]

In recent years, labour market flexibility has improved the ability of economies to adapt to oil shocks [62]. In the case of larger and more persistent oil price increases, however, workers may resist a series of real wage cuts. This, while perhaps mitigating the distributional consequences of the oil shock, could significantly increase the impact of the shock on GDP.

Finally, the simulations do not consider the possibility that a fast erosion in net oil exports is highly likely to occur regardless of what happens to the growth rate of world oil production, owing to rapid population and GDP growth in many oil exporters. Some oil exporters might withhold an even larger share of their stagnating or decreasing oil production for domestic use, for example through fuel subsidies, in order to support energy-intensive industries and to forestall domestic unrest. The amount of oil available to oil importers will therefore with certainty shrink much faster than world oil production [63], with obvious negative consequences for growth in those regions.

5. Conclusion

This paper has shown that the extent to which persistent oil scarcity could affect global economic growth and current account imbalances depends critically on a small number of key factors. If, as in our baseline scenario, the trend growth rate of oil production declined only modestly, and if the economy was adequately represented by a conventional CES production function in capital, labour and oil, world output would eventually suffer, but the effect would not be dramatic. If the substitutability between oil and other inputs into production tended to increase as the oil price increases, the effect would be even smaller. But if the reductions in oil production were more in line with the more pessimistic studies in the scientific literature, the effects could be extremely large. The same could be true if, as claimed by several authors in the scientific literature, standard production functions underestimate the economic importance of oil for output. We discussed three possibilities. First, if the economy attempted to substitute away from oil on a large scale, it might encounter a lower limit of oil use dictated by entropy. Second, the contribution of oil to output could be much larger than its cost share, because oil is an essential precondition for the continued viability of many modern technologies. Third, the income elasticity of oil demand could be well below one as in many empirical studies, rather than one as in our model.

[49]In such a situation, domestic banks suddenly find that they can obtain foreign loans at very attractive interest rates. But they may not be able to identify a sufficient number of suitable borrowers, and may therefore end up making imprudent lending decisions.

References

1. Helbling T, Kang JS, Kumhof M, Muir D, Pescatori A, Roache S. 2011 Oil scarcity, growth, and global imbalances. In *World economic outlook, April 2011*, ch. 3. Washington, DC: International Monetary Fund.

2. Ayres R. 1990 Technological transformations and long waves, part I. *J. Tech. Forecast. Soc. Change* **37**, 1–37. (doi:10.1016/0040-1625(90)90057-3)

3. Ayres R. 1990 Technological transformations and long waves, part II. *J. Tech. Forecast. Soc. Change* **37**, 111–137. (doi:10.1016/0040-1625(90)90065-4)

4. Wrigley E. 2010 *Energy and the English industrial revolution*. Cambridge, UK: Cambridge University Press.

5. Aleklett K. 2012 *Peeking at peak oil*. New York, NY: Springer.

6. Waisman H, Rozenberg J, Sassi O, Hourcade J-C. 2012 Peak oil profiles through the lens of a general equilibrium assessment. *Energy Policy* **48**, 744–753. (doi:10.1016/j.enpol.2012.06.005)

7. Hubbert MK. 1956 Nuclear energy and the fossil fuels. *American Petroleum Institute Drilling and Production Practice Proceedings*, pp. 5–75.

8. United States 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. Washington, DC: US Government Accountability Office.

9. United States Joint Forces Command. 2010 *The joint operating environment 2010*. Suffolk, VA: US Joint Forces Command Joint Futures Group.

10. Hirsch R, Bezdek R, Wendling R. 2005 *Peaking of world oil production: impacts, mitigation and risk management*. Washington, DC: US Department of Energy.

11. Hirsch R, Bezdek R, Wendling R. 2010 *The impending world energy mess*. Burlington, Canada: Apogee Prime.

12. UK Energy Research Centre. 2009 *Global oil depletion: an assessment of the evidence for a near-term peak in global oil production*. London, UK: UK Energy Research Centre.

13. Sorrell S, Miller R, Bentley R, Speirs J. 2010 Oil futures: a comparison of global supply forecasts. *Energy Policy* **38**, 4990–5003. (doi:10.1016/j.enpol.2010.04.020)

14. Solow R. 1956 A contribution to the theory of economic growth. *Q. J. Econ.* **70**, 65–94. (doi:10.2307/1884513)

15. Kaldor N. 1957 A model of economic growth. *Econ. J.* **67**, 591–624. (doi:10.2307/2227704)

16. Kümmel R, Strassi W, Gossner A, Eichhorn W. 1985 Technical progress and energy-dependent production functions. *J. Econ.* **45**, 285–311. (doi:10.1007/BF01282565)

17. Stiglitz J. 1974 The Cambridge–Cambridge controversy in the theory of capital: a view from New Haven. *J. Political Econ.* **82**, 893–903. (doi:10.1086/260245)

18. Samuelson P. 1987 Sraffian economics. In *The new Palgrave: a dictionary of economics*, vol. 3, pp. 452–460. New York, NY: Macmillan.

19. Cohen A, Harcourt G. 2003 Whatever happened to the Cambridge capital theory controversies? *J. Econ. Perspect.* **17**, 199–214. (doi:10.1257/089533003321165010)

20. Bodenstein M, Erceg C, Guerrieri L. 2011 Oil shocks and external adjustment. *J. Int. Econ.* **83**, 168–184. (doi:10.1016/j.jinteco.2010.10.006)

21. Georgescu-Roegen N. 1975 Energy and economic myths. *South. Econ. J.* **41**, 347–381. (doi:10.2307/1056148)

22. Daly H. 1997 Georgescu-Roegen versus Solow/Stiglitz. *Ecol. Econ.* **22**, 261–266. (doi:10.1016/S0921-8009(97)00080-3)

23. Berndt E, Wood D. 1975 Technology, prices and the derived demand for energy. *Rev. Econ. Stat.* **57**, 259–268. (doi:10.2307/1923910)

24. Jorgenson D, Fraumeni B. 1981 Relative prices and technical change. In *Modeling and measuring natural resource substitution* (eds E Berndt, B Field), pp. 17–47. Cambridge, MA: MIT Press.

25. Frondel M, Schmidt C. 2002 The capital-energy controversy: an artifact of cost shares? *Energy J.* **23**, 53–79. (doi:10.5547/ISSN0195-6574-EJ-Vol23-No3-3)

26. Georgescu-Roegen N. 1971 *The entropy law and the economic process*. Cambridge, MA: Harvard University Press.

27. Soddy F. 1922 *Cartesian economics: the bearing of physical science upon state stewardship*. London, UK: Students' Unions of Birkbeck College and London School of Economics.

28. Babiker M, Reilly J, Mayer M, Eckaus R, Wing I, Hyman R. 2001 *The MIT emissions prediction and policy analysis (EPPA) model: emissions, sensitivities and comparison of results*. Report no. 71.

25

rsta.royalsocietypublishing.org *Phil. Trans. R. Soc. A* **372**: 20120327

Cambridge, MA: Massachusetts Institute of Technology, Joint Program on the Science and Policy of Global Change.

29. Ghersi F, Hourcade J-C. 2006 Macroeconomic consistency issues in E3 modeling: the continued fable of the elephant and the rabbit. *Energy J.* (Hybrid Modeling Special Issue, no. 2), 39–62. (doi:10.5547/ISSN0195-6574-EJ-VolSI2006-NoSI2-3)

30. Hassler J, Krusell P, Olovsson C. 2012 *Energy-saving technical change*. Working Paper. Stockholm, Sweden: Institute for International Economic Studies.

31. Stern D. 1997 Limits to substitution and irreversibility in production and consumption: a neoclassical interpretation of ecological economics. *Ecol. Econ.* **21**, 197–215. (doi:10.1016/S0921-8009(96)00103-6)

32. Reynolds D. 2002 *Scarcity and growth considering oil and energy: an alternative Neo-classical view*. Lampeter, UK: Edwin Mellen Press.

33. Ayres R, Warr B. 2005 Accounting for growth: the role of physical work. *Struct. Change Econ. Dyn.* **16**, 181–209. (doi:10.1016/j.strueco.2003.10.003)

34. Ayres R, Warr B. 2010 *The economic growth engine: how energy and work drive material prosperity*. Cheltenham, UK: Edward Elgar Publishing.

35. Hall C, Klitgaard K. 2012 *Energy and the wealth of nations: understanding the biophysical economy*. New York, NY: Springer.

36. Kümmel R. 2011 *The second law of economics—energy, entropy, and the origins of wealth*. Berlin, Germany: Springer.

37. Kümmel R, Henn J, Lindenberger D. 2002 Capital, labor, energy and creativity: modeling innovation diffusion. *Struct. Change Econ. Dyn.* **13**, 415–433. (doi:10.1016/S0954-349X(02)00008-5)

38. Kumhof M, Laxton D, Muir D, Mursula S. 2010 The global integrated monetary and fiscal model—theoretical structure. IMF Working Paper WP/10/34.

39. Benes J, Chauvet M, Kamenik O, Kumhof M, Laxton D, Mursula S, Selody J. 2012 The future of oil: geology versus technology, IMF Working Paper WP/12/109.

40. Durbin J, Koopman S. 2012 *Time series analysis by state space methods*, 2nd edn. Oxford, UK: Oxford University Press.

41. Norwegian Ministry of Finance. 2001 *Storting: guidelines for economic policy*. Report no. 29. See http://www.regjeringen.no/en/dep/fin/Selected-topics/economic-policy/economic-policy.html?id=418083.

42. Smil V. 2010 *Energy transitions: history, requirements, prospects*. New York, NY: Praeger.

43. Ayres R. 2007 On the practical limits to substitution. *Ecol. Econ.* **61**, 115–128. (doi:10.1016/j.ecolecon.2006.02.011)

44. Kunz H. 2010 *The fake fire brigade—how we cheat ourselves about our energy future*. See http://www.theoildrum.com/node/6641.

45. WWF. 2011 *The energy report—100% renewable energy by 2050*. See http://www.ecofys.com/en/publication/11/.

46. Murphy DJ. 2014 The implications of the declining energy return on investment of oil production. *Phil. Trans. R. Soc. A* **372**, 20130126. (doi:10.1098/rsta.2013.0126)

47. Timilsina GR. 2014 Biofuels in the long-run global energy supply mix for transportation. *Phil. Trans. R. Soc. A* **372**, 20120323. (doi:10.1098/rsta.2012.0323)

48. Heinberg R. 2009 *Blackout: coal, climate and the last energy crisis*. Gabriola Island, Canada: New Society Publishers.

49. Höök M, Davidsson S, Johansson S, Tang X. 2014 Decline and depletion rates of oil production: a comprehensive investigation. *Phil. Trans. R. Soc. A* **372**, 20120448. (doi:10.1098/rsta.2012.0448)

50. Höök M, Fantazzini D, Angelantoni A, Snowden S. 2014 Hydrocarbon liquefaction: viability as a peak oil mitigation strategy. *Phil. Trans. R. Soc. A* **372**, 20120319. (doi:10.1098/rsta.2012.0319)

51. United States Energy Information Administration. 2009 *International energy outlook*. Washington, DC: US Energy Information Administration.

52. Berman A. 2012 After the gold rush: a perspective on future US natural gas supply and price. *Resilience*. See http://www.resilience.org/stories/2012-02-08/after-gold-rush-perspective-future-us-natural-gas-supply-and-price.

53. Hughes JD. 2013 *Drill baby drill—can unconventional fuels usher in a new era of energy abundance?* Santa Rosa, CA: Post Carbon Institute.

54. Rogers D. 2013 *Shale and Wall Street: was the decline in natural gas prices orchestrated?* Energy Policy Forum.

55. Dittmar M. 2011 The end of cheap uranium. Working Paper. Institute of Particle Physics, ETH, Switzerland.

56. Crassous R, Hourcade J-C, Sassi O. 2006 Endogenous structural change and climate targets: modeling experiments with IMACLIM-R. *Energy J.* (Endogenous Technological Change Special Issue, no. 1), 259–276. (doi:10.5547/ISSN0195-6574-EJ-VolSI2006-NoS11-13)

57. Armstrong J, Black R, Laxton D, Rose D. 1998 A robust method for simulating forward-looking models. *J. Econ. Dyn. Control* **22**, 489–501. (doi:10.1016/S0165-1889(97)00081-X)

58. Juillard M, Laxton D, Pioro H, McAdam P. 1998 An algorithm competition: first-order techniques versus Newton-based techniques. *J. Econ. Dyn. Control* **22**, 1291–1318. (doi:10.1016/S0165-1889(98)00013-X)

59. Chevalier M. 2003 Chain fisher volume index methodology. Research paper no. 42, Statistics Canada.

60. Hamilton J. 2009 Causes and consequences of the oil shock of 2007–08. Brookings Papers on Economic Activity, pp. 215–261.

61. Mackenzie K. 2012 *Marginal oil production costs are heading towards $100/barrel*. See http://ftalphaville.ft.com/2012/05/02/983171/marginal-oil-production-costs-are-heading-towards-100barrel/.

62. Blanchard O, Galí J. 2007 The macroeconomic effects of oil shocks: why are the 2000s so different from the 1970s? NBER Working Paper no. 13368. Boston, MA: National Bureau of Economic Research.

63. Brown J, Foucher J. 2010 Peak oil versus peak exports. *Resilience*. See http://www.energybulletin.net/stories/2010-10-18/peak-oil-versus-peak-exports.

26

rsta.royalsocietypublishing.org *Phil. Trans. R. Soc. A* **372**: 20120327