



HARVARD Kennedy School

MOSSAVAR-RAHMANI CENTER
for Business and Government

Oil Intensity: The curious relationship between oil and GDP

**Christof Rühl
Tit Erker**

May 2021

M-RCBG Associate Working Paper Series | No. 164

The views expressed in the M-RCBG Associate Working Paper Series are those of the author(s) and do not necessarily reflect those of the Mossavar-Rahmani Center for Business & Government or of Harvard University. The papers in this series have not undergone formal review and approval; they are presented to elicit feedback and to encourage debate on important public policy challenges. Copyright belongs to the author(s). Papers may be downloaded for personal use only.

Oil Intensity:

The curious relationship between oil and GDP

By

Christof Rühl¹ and Tit Erker²

Abstract

Since the mid-1980s, volumetric improvements in oil intensity have followed a linear time trend. We verify the existence and explore the causes and consequences of this trend. It is the result of a shift in the composition of oil demand from intermediate to final consumption. The typically single-fuel appliances associated with final consumption restrict inter-fuel substitutability and price elasticity while, in a globalising world, efficiency gains in end-user appliances spread fast and uniformly, mirroring global turnover patterns. We argue that on a deeper level, the resulting time trend reflects a gradual regime change from a supply to a demand constrained global oil market. In percentage terms, oil intensity first improves at rates below global GDP growth, allowing global oil demand to rise. Over time, given its linear functional form, the rate of intensity improvements will accelerate to outpace the rate of global GDP growth, and global oil demand will start to decline. It is conceivable for relative price changes to break the trend, by reconfiguring the degree of substitutability embodied in the capital stock. However, so far neither shale oil (accelerating oil demand) nor the energy transition (decelerating it) have had this effect. For now, the linear trend soldiers on.

JEL Classification: Q41, Q43

¹ Senior Research Scholar, Center on Global Energy Policy, Columbia University and Senior Fellow, Mossavar-Rahmani Center for Business and Government, Harvard Kennedy School, Harvard University, USA.

² Strategy and Planning Manager, Abu Dhabi Investment Authority, UAE.

I. Introduction

Oil markets are notoriously difficult to model.³ They are complex and volatile and at the same time, of course, of vital importance to the economy. Despite this importance, it is hard to find a simplified, common narrative connecting different phases of their historic development, much less a generally accepted model of lasting empirical relevance.

This void has good reasons. Oil prices and production are subject to idiosyncratic and hard-to-model parameters, often interlocking with what they are supposed to explain (e.g., increasing returns in production or, on the demand side, feedback loops with the economy). These interdependencies make it hard to identify even basic ingredients of aggregate demand and supply schedules such as, notoriously, price or income elasticities. Finally, it is hard to model the structure of a market which sometimes appears to be highly cartelized, and at other times populated by a large flock of peaceful price takers.

'Too many moving parts' is a most fitting description.

All the more surprising then how in the midst of all these moving parts, a stable trend stands out. Oil intensity, i.e. oil consumed per unit of economic output, has improved (declined) steadily over many decades, with remarkable regularity, and in almost perfectly linear fashion.

The following note investigates this time trend. We revisit the historic context in which it first emerged (Part II), examine the reasons for its stability (Part III), and explore whether this trend can be exploited to improve our understanding of longer-term oil market developments (Part IV).

II. Oil intensity: A very stable relationship

Oil remains the largest primary fuel, moving into pole position in 1964, almost 60 years ago. Like the consumption of all other commercially traded fuels, global oil consumption continues to rise in absolute terms⁴ – by 6% per year since records exist (1870); 1.2% p.a. since 1973, when its share in global primary energy peaked; and 1.6% over the last ten years. Its market share in the global primary energy mix, however, has fallen from 50% in 1973 to 33% in 2019 and continues to decline (Fig. 1).⁵


Analytically, oil consumption growth is a function of price and income changes. The literature has long emphasized the relative price inelasticity of demand, especially in the short term. The extent to which oil price fluctuations may impact GDP growth adds a curious identification problem to income elasticities. However, the attempts at estimating with great accuracy functional elasticities which are recognized as unstable over time and space, may not be the only strategy to gain insights into the future path of production and consumption. A bird's-eye, macro view of changes in the ratio of aggregate oil consumption and global income generation provides a surprisingly constructive addition.

³'Oil' in this paper is defined as per the primary data source (BP 2020), i.e. including in-country demand plus international aviation and marine bunkers plus refinery fuel and loss. Biofuel consumption is excluded; derivatives of coal and natural gas are included. All energy data in this paper is BC ('Before Covid'), up until and including 2019.

⁴ With the possible exception of biomass (wood) and coal. There is no data on commercially traded biomass for energy consumption, and it is too early to tell whether coal consumption has peaked.

⁵ Oil consumption growth was very volatile until after WW II and steadied after the 1960s (1870-1965 = 8.3% p.a. vs. 1966-2019 = 2.1% p.a., cf. BP (2011, 2020)).

Figure 1: Primary fuel shares and Global oil demand



Oil intensity is defined as oil consumed per unit of economic output (or income, or GDP), and measured best in barrels (bbl) per dollar. Just as GDP, it can be studied from the income or the expenditure side. Oil intensity is often viewed as the broadest efficiency measure we have – a gauge of the economic importance of oil (it is the inverse of oil productivity, after all). Effectively, it is the result of multiple influences, including changes in end-user preferences, technical efficiency improvements, and changes in the structural composition of the economy, such as industrialisation or urbanisation.

(a) The linear time trend of oil intensity

The key observation motivating this paper is that oil intensity on a global scale has for many decades exhibited linear improvements (declines) year after year. The unceasingly regular decline of oil consumption in global GDP has been accompanied by convergence of oil intensity across major consuming countries.

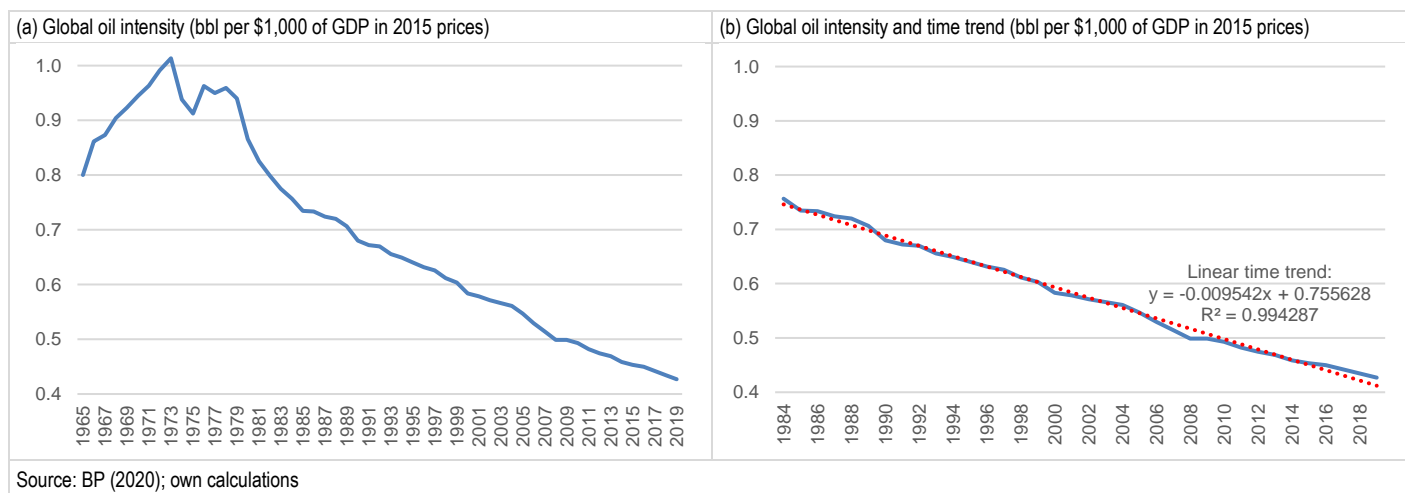
The regularity of this time trend is startling. At a global level, oil intensity starts to decline in roughly constant (annual) increments from 1984 onward, after a decade of high price and consumption volatility. In retrospect, the global oil economy at this point appears to have settled into a new modus operandi, of which the time trend we are discussing is a key indicator. The trend has stayed intact ever since, running silently and rather neglected in the background.⁶ The linear improvements continued through the only other period of comparable price volatility, after the turn of the millennium when, in contrast to the 1970s, oil consumption remained relatively insensitive to price volatility. At a first glance, prices do not seem to be the main driver of the oil intensity trend, an issue we will come back to.

Panel (b) in Fig. 2 illustrates the linear decline of oil intensity from 1984 until 2019. The data for this period transform into an almost perfect straight line with an impressive statistical fit around the linear time trend (R^2 of 0.99). The time trend itself depicts constant volumetric declines per unit of GDP over the entire period. In numbers, the regression line tells us that it took on average 0.0095 barrels (1.5 litres or 0.4 US gallons) less to produce \$1,000 worth of global GDP (in 2015 prices) – year after year, for the last three and a half decades, with very high regularity.⁷

⁶ The only explicit treatment of reasonably recent vintage we are aware of are Fagan and Duan (2019), Fagan (2020) and Mirzoev T., et al. (2020). We were unable to replicate the basic oil intensity data cited in Fagan and Dunn (2019) and in Fagan (2020), and therefore their results. The authors did not follow through with a promised explanation despite repeated attempts. Mirzoev T., et al. (2020) confirm the basic results and observations around the integrity, longevity and stability of the oil intensity time trend presented here, while the purpose of their publication is not to explain it. The authors infer that the long-term robustness of this time trend may lead to the possibility of estimating future oil demand growth even without recourse to prices – which is a tad further than we would be prepared to go.

⁷ Over the period 1984-2019 oil intensity declined by 1.6% p.a., on average.

Figure 2: Global oil intensity and time trend since 1984



One of the first questions when encountering this linear rate of improvement will be whether it would withstand disaggregation to the region or country level. However, this time trend does not depend on data manipulation or source, is not a statistical fluke, and it is robust.⁸ It can be replicated for regions and countries. Individual characteristics or deviations displayed on that level are consistent with the explanation for the trend's existence we will advance below (see Appendix). The same is true if it is broken down by sector. The trend shows the highest explanatory power in retail-based, final consumption sectors such as transport or residential use, and less in lumpy intermediate segments such as industry and petrochemicals.⁹

There are fluctuations around the trend over time, of course, but they tend to be limited in scope and duration. For instance, we have seen a significant deceleration in the incremental annual decline in oil intensity (i.e., a flattening of the curve) during the years 1991-97, and an even more pronounced deceleration from 2001-04.¹⁰ However, in both instances this was reversed by a subsequent acceleration in the oil intensity decline (i.e., a steepening of the curve), with a record improvement after the second episode.


This is an important point because, had the data shown the increments declining in step with the decline in energy intensity over time, it would have been an indication of oil intensity improvements following an exponential, rather than a linear trend. So far this has not been the case; we have seen oscillations instead. In other words, the trend was stable as periods of relatively slower decline were alternating with and reversed by periods of faster intensity decline. The duration of deviations in either direction has been 4 to 8 years in the past, and the amplitude was between 1-2% of the oil intensity level (or 0.005 to 0.010 barrels per \$1,000 of GDP) before the financial crisis in 2008/09 (Fig. 3).

⁸ On the matter of data sources, we have tested alternative GDP and oil data sources and combinations, including GDP data from The World Bank at constant market exchange rates and at purchasing power parity, and oil consumption from the IEA and BP. In all four combinations the linear trend performs marginally better than the exponential trend, compared to the data series ultimately reported here (Oxford Economics for GDP and BP for oil demand), a combination chosen for depth, consistency and timeliness of coverage.

⁹ The residential and transport sectors have an R^2 of 0.98 and 0.96, respectively; compared to industry and petrochemicals, which have an R^2 of 0.93 and 0.82, respectively (all over the 1984–2018 period as per IEA 2020 data availability).

¹⁰ In 1991-97, oil intensity declines flattened to 0.008 barrels (bbl) per \$1,000 p.a., down from 0.011 bbl per \$1,000 p.a. during the preceding 7 years when the long term oil intensity trend first emerged (1984-90); In 2001-04 incremental oil intensity improvements fell below 0.006 bbl per \$1,000 p.a., compared to 0.010 bbl per \$1,000 during 1984-2000. The subsequent acceleration in 2005-08 saw oil intensity improving by a record 0.016 bbl per \$1,000 p.a..

Figure 3: Oil intensity deviations from the linear trend



The most pronounced slowdown in the rate of oil intensity improvements has happened more recently, namely since the financial crisis in 2008/09, and throughout the high and low price periods since then. Although this moderation stands out in terms of depth and duration, it seems already obvious that it will be reversed by the impact of the COVID-19 pandemic, which hit global oil demand much harder than global GDP. Cheeky historical resolutions aside, the current deviation is of interest because it appears to be linked to the US shale 'revolution', which could bring it much closer to the possibility of a structural break in the system than previous episodes (discussed in Part III (e) below).

Like most economists would when catching sight of a linear time trend, we harboured lingering suspicions about Stein's law and hence performed due diligence by applying an exponential trend to the data as well.¹¹ The statistical fit turns out to be literally almost identical, with the flattening of the curve since 2009 also following a significant acceleration in intensity declines in the period 2005-08.¹² However, for regions, like the EU or the OECD, in which the regular oil intensity decline is most advanced and where decline rates have already surpassed GDP growth and have translated into declining oil demand, the linear trend is the better fit.¹³ That statistical fit, the observed episodes of reversed deviations and oscillation around the trend; and not least our understanding of the reasons for a linear time trend led to the conclusion that oil intensity improvements have been linear, and to adopt the linear trend as the most appropriate specification.

The predictive power of this linear approximation confirms the exceptional data fit. Out-of-sample forecasts, 1 and up to 14 years ahead, based on at least 20 years of data (from 1984 to 2003, and up to 2017) have an average error of 1.7% to 5.1% (increasing with the length of the forecast period). In absolute terms, this is equivalent to an average root mean squared deviation of 0.008 barrels per \$1,000 for all 1-year ahead within sample oil intensity forecasts for the years 2004 to 2018, and up to 0.022 barrels per \$1,000 for the three 14 years ahead within sample forecasts.¹⁴ It is a truly extraordinary performance, especially over the longer horizons – even more so in the case of the EU and the OECD, two regions where oil demand has been declining for some time, and where forecasting accuracy starts to improve over longer horizons (see Appendix (3) for detail).

¹¹ The dictum 'If something cannot go on forever, it will stop' is referred to as Stein's Law (Samuelson 2013). One could argue, of course, that oil demand will eventually disappear, like firewood did – to stay in the genre, this would be an example from the broader category of 'What goes up must come down but the reverse is seldom true' (Hogan 1993, p.).

¹² R^2 is 0.996 for the exponential trend (i.e. logarithmic specification), and 0.994 for the linear trend specification.

¹³ R^2 in the linear specification is 0.991 in both cases, whereas with the exponential trend in the logarithmic specification R^2 is 0.982 for the OECD and 0.983 for the EU.

¹⁴ To put the absolute deviations of 0.008 and 0.022 bbl per \$1,000 into perspective, note the average actual oil intensity level between 2004 and 2018 was 0.487 bbl per \$1,000, and the average actual oil intensity level between 2017 and 2019 (which are the three 14-year forecasts) was 0.435 bbl per \$1,000.

It bears repeating just how remarkable the consistency and the ex-post predictive power of this linear trend are. 35 years is a long time, with price booms and busts, OPEC successes and failures, wars, revolutions, the rise of China and the dissolution of the Soviet Union all exerting their influence on the global market for oil. Without much deviation, the constant improvements in oil consumption per unit of GDP have been sustained through all of it.

It was not always like that, however.

(b) How did this trend emerge?

Historically, oil intensity has been rising during the years after World War II, with the share of oil in the global energy mix increasing fast. Oil became the largest fuel in 1964 and continued to globalize, on the back of large finds in the Middle East. Based on economic competitiveness, i.e. a competitive relative price for superior energy density and convenience properties, it substituted for other fuels, specifically coal, in power generation, space heating and industrial applications. The rise of oil was of course propelled forward by the global roll-out of a transport sector based on the internal combustion engine; but it was not confined to new applications.

Global oil intensity peaked in 1973, the same year oil's share in global primary energy reached its zenith (at 50%). At that time, it took just over 1 barrel (159 litres or 42 US gallons) of oil to support the generation of \$1,000 worth of global GDP (in 2015 dollars) – a number that helps to appreciate the prominence of oil in contemporary debates of economic performance. (Fig. 2 panel (a))

The time when oil rapidly permeated the nooks and crannies of the global economy came to an end with the two oil crises in the 1970s.

After the Iranian revolution, as real crude prices reached levels unrivalled for the next three decades, oil intensity went into a tailspin for a few years. Oil retreated from applications where it faced competition from other fuels, which now had become relatively cheaper and less prone to the threat of supply disruptions and associated price volatility. Oil use declined 'under the boiler' and especially in power generation; its concentration in transport increased significantly. The steep decline of oil intensity coincides with a steep decline in fuel market share, both lasting until 1984 when the new regime of a secular trend decline in oil intensity sets in.


By 2019, oil consumption per \$1,000 of global GDP was down to 0.43 barrel (68 litres or 18 US gallons), a decline of 58% since oil intensity had peaked 46 years earlier (and still a number giving pause for thought). More than half (56%) of this achievement has been attained after the earlier period of high prices and rapid oil demand decline, i.e. after 1984, and as a result of the steady oil intensity improvements embodied in the new time trend,

What matters here, of course, is less the observation that oil intensity has declined over time but the regular pattern which has been marking this decline. A new regime with income and demand growth moving in lockstep has asserted itself. In its wake, the relationship between prices and demand also had to change.

The 1970s and early 1980s have been a period of great volatility of price and quantity adjustments in global oil markets, with (political) disruptions triggering huge price movements and price movements triggering swift demand adjustments (and a lagged supply response). Until the mid-1980s, global demand and therefore oil intensity mirrored price fluctuations during the years of turbulence as one would expect – dropping after price spikes and rising after prices fell or stabilized. After that, oil intensity continued to decline, despite a drop in oil prices: Between 1980 and 1985 oil intensity declined by 15%, while real oil price dropped by more than 40%.

An important component of this transition is how the interaction between the price of oil and oil intensity changed. With low and declining real oil prices until the early 1970s, oil intensity had increased; during and after the oil price spike of the first oil shock, oil intensity dropped precipitously. From the mid-1980s on, when oil intensity embarked on its newly found steady course, the umbilical link to prices seems to have become undone (Fig. 4).

Figure 4: Global oil intensity and real oil price, before and after 1985



When the price volatility of the 1970s returned on a comparable scale decades later, with the run-up of prices in the first decade of the new millennium and the shale oil induced price collapse from 2014, oil intensity improvements showed little perturbation. The standard explanation for the difference is that the price volatility was caused by supply disruptions in the first and by increasing demand in the second instance, and that price and income elasticities therefore have been affected in different ways.¹⁵

Viewed from the macro perspective of oil intensity, however, there is more to this asymmetric reaction than different price triggers alone. In both instances, not only oil intensity but the share of oil in global primary energy behaved very differently – volatile in the 1970s and gradually continuing to decline in the 2000s. A bigger transition was afoot. Had oil markets become less flexible, i.e., less prone to inter fuel substitution than it used to be? If so, why would that be?

(c) Oil demand

The new configuration after the turbulent 1970s was not lost on contemporary observers.

The beginning of the secular trend in regular oil intensity improvements coincides with years of weak demand and falling oil prices. It puzzled observers why the demand response to falling prices was so sluggish.¹⁶ In standard economic theory, the demand reaction to falling or rising energy (or oil) prices would have been symmetric, with capital, labour and energy close substitutes and their configuration (or 'technology') flexible and driven by relative prices. The lack of a more enthusiastic demand response to falling and stagnating oil prices in the 1980s and 90s, in contrast to the decline in demand following the price spikes in the 1970s, did not fit into this picture.

One popular attempt at explaining the data was the hypothesis of an asymmetric response of oil demand to price movements: Once high prices had triggered efficiency gains, either baked into the capital stock ('isolating the attic') or directly incentivising technology (a more efficient car engine), these gains would to an extent persist and not go away in response to a fall in the relative price of oil.¹⁷ The result would be a ratchet effect for oil demand which, with half way steady GDP growth, should leave oil intensity improving in steps: High price periods would meet elastic oil demand and lead to irreversible improvements in oil intensity, whereas for low price periods the predicted oil intensity adjustments would be disproportionately small.¹⁸

¹⁵ There is extensive and not always uncontroversial literature on the differences in causes, symptoms and consequences of oil demand and supply 'shocks', including the two episodes referred to in the text. See Blanchard O.J. and Gali, J. (2007), Baumeister C. et al (2009), Chain P. et al (2014); specifically for the 2000s in the US Hamilton J.D. (2009a,b), Kilian L. (2009).

¹⁶ It took 9 years for global oil consumption to regain the level of 1979. For real prices (which had reached their apex in 1980), this didn't happen until 2008.

¹⁷ See Gately and Huntington (2002), and Dargay (1992), Gately (1992, 1993), Dargay and Gately (1994) for the early stages of this discussion.

¹⁸ In theory, price oscillations would thus drive oil intensity toward zero, which is why slightly more sophisticated versions had to be added.

Not everyone was convinced. There is a discord between the practical appeal of an idea such as efficiency gains being irreversible (unlearning, we like to think, is hard to do) and the tenets of standard economic theory. To the extent that efficiency gains are the result of extensive investments in response to a shift in the relative price of oil (the famous isolation of the attic), they will depreciate and wither away eventually, i.e. not be replaced over time, once the initial price trigger has been removed. To the extent that they reflect endogenous technological change triggered by price movements (e.g., a more efficient car engine following an oil price increase), other parts of the economy will catch up to neutralize the comparative advantage of this innovations over time, once the underlying structure of relative prices has been restored, i.e. the original incentive removed. Hence, in the long term, there should be no ratchet effect in the data.

Yet, there was no denying that oil intensity fell continuously, with GDP growth in the 1980s outpacing oil demand growth, despite a declining relative price for oil. Prices low and falling, demand growth stalling despite reasonable GDP growth, oil intensity steadily improving: it must have seemed as if market demand was doomed *and* as if there was a missing ingredient for economic modelling to fit the data. It may be obvious in retrospect (as with so many good ideas) to identify oil-saving technological change as the most trustworthy guide out of the analytical impasse. To incorporate a time trend of oil conserving technology which was independent of price would allow the co-existence of positive oil demand growth with declining oil intensity into the future, without compromising long term price effects.

Hogan (1993) estimates a number of aggregate oil demand specifications, symmetric and asymmetric. He encapsulates the general conclusion, with the winner to emerge from his exercise a reduced form model which combines symmetric price reactions and an oil-conserving technology trend which is not price dependent.¹⁹ “Apparently,” the author notes with slight disbelief when inspecting modelled long term intensity estimates (ibid, p, 149), “there should be a significant further reduction [of] oil intensity over the 1990s, independent of price, even though there will be continuing growth in total oil demand”. Which, of course, is exactly what has happened.

This was perhaps the high point in a debate which subsequently petered out gradually, just as the oil markets themselves lost much of their lustre and excitement throughout the 1990s. Oil markets remained subdued until catching analysts by surprise again, first by the lack of economic impact of sky-high prices in the 2000s, and then by shale oil. As to the economic debate, variables such as oil intensity lost their place in the sun, as attention turned inward, toward exercises driven by data availability and at any rate, increasingly divorced from the macroeconomic narrative of the day.

However, there was a clear understanding by the early 1990s of how a growing economy over the long term would reconcile demand growth and oil intensity decline (or why oil intensity was declining less than GDP was growing), based on an exogenous series of oil-saving technological changes.²⁰ The analytical insights of the debate would allow to forecast the direction of the future path of oil intensity. But to tame quantity adjustments by adding irreversible exogenous efficiency improvements is not the same as to rule out that symmetric price responses might increase oil intensity again. The second half of our puzzle, the fact that energy intensity improvements turned out to be not only relentless but regular, was not something anyone had on the radar, forecast or not.

There is no way around finding an explanation for that linear time trend, if the goal is to improve our understanding of the path of oil consumption.

¹⁹ Based on the Harvard Oil Market Simulation Model, the estimates were carried out for OECD economies and product prices (as distinct from crude prices). Technological change was exogenous.

²⁰ These models would typically use putty clay versions of technological change, with changes in the capital stock irreversible.

III. What explains the linear trend?

It is not one but three questions that need an answer. First, what explains the linearity of the observed path of oil intensity over the last few decades. Second, why the two most volatile oil price periods in living memory had such a profoundly distinct impact on real variables, including on oil consumption. And third what, if anything, this regime change can teach us about a future in which people may actively desire to switch away from oil consumption.

(a) Derived demand

Oil demand is derived demand. There is no direct utility from consuming oil. To benefit from its use requires a parallel investment in appliances. For final consumption demand, these appliances will enable consumers to derive the services oil provides – a car for transport or a boiler for space heating. Alternatively, for production purposes, capital goods are required to convert oil from an intermediary input into a final good or into another intermediary, such as plastic or power.²¹

Conceptually it pays to distinguish end-user appliances which enable final consumption demand (the household car or boiler) from capital goods which absorb oil as an intermediary input in producing something else (a power plant or a petrochemical factory). A second distinction will be useful for our purposes, namely that between single and multi-fuel appliances – that is, between appliances (or systems) which allow for fuel switching and substitution across fuels, and those which do not.

Appliances of all sorts imply sunk costs. Single-fuel appliances diminish the price elasticity of oil demand directly: the higher the ratio of fixed to variable cost, with oil the variable component, the less will variations in fuel prices affect the cost of the service provided.²² This same principle should apply for final consumption and for intermediate use. Conversely, the option of fuel substitution increases the price elasticity of oil demand. Note that fuel switching may be due to multi-fuel appliances (dual fired power plants or, if batteries lived up to expectations, hybrid cars); or to the possibility of fuel switching by switching between appliances with different input fuels in systems with built-in redundancy (for example, if a power system for security reasons offers the option to switch between oil fired and other power stations).²³

Traditionally, final demand and the consumer appliances it is associated with have been more likely to consist of single fuel devices; whereas capital goods, which use oil as an intermediary input, are more likely to allow for fuel substitution, especially if policy issues (energy security, environment, industrial planning) are involved.²⁴ New and most efficient technologies tend to favour single fuel specialisation as well.

Our first proposition, on the impact of derived demand on global oil intensity, would thus be the following: over time, as income levels rise, the share of oil used in final consumption and therefore in conjunction with single fuel appliances, increases. As a result, oil price elasticity diminishes. This has been the case on a global scale since the 1980s.


²¹ Installations to refine and process crude oil belong into the second category as well. For empirical purposes, we will use the IEA sectoral classification, with transport and residential use representing final consumption, and the remainder representing intermediate demand (cf. IEA 2020).

²² 'Oil' here refers to products not crude, but empirical studies seem to be doing fine by employing a fixed conversion schedule and taking crude prices as a proxy.

²³ Close substitutes (e.g., public vs individual transport) would also fall into this category.

²⁴ Also for technical reasons, i.e. because they tend to use the heavy end of the barrel.

Figure 5: Share of oil in primary energy and share of oil used in final consumption




In this reasoning, as the lion’s share of oil consumption shifts toward final consumption, the increased efficiency of end-user devices will eventually slow the growth of oil in the economy; the dwindling options for fuel substitution will stabilize the share of oil in the fuel mix; and both lower the price elasticity of oil. Oil consumption growing faster than GDP (oil intensity rising), then GDP growing faster than oil intensity falls (intensity improvements combined with positive oil demand growth) and finally, the point from which the intensity decline starts to outpace GDP growth and oil demand starts to decline – these have been the thresholds oil intensity had to cross in regions where the full process has already played itself out.

(b) Globalization

Globalisation itself is another factor encouraging both regularity and direction of travel for the trend toward lower global oil intensity. Accelerating international competition and the growing ‘flow of everything’ (trade, investment, people and capital) across borders have bolstered the uniform spread of technology over the last four decades, including efficiency improvements embodied in machinery. Fig. 6, representing 47% of global consumption in 2019, illustrates the convergence of oil intensities across the most important oil consumers which mirrors deeper economic integration as well as intensifying global competition.²⁵

²⁵ The diffusion of oil conserving efficiency improvements across developing countries relative to high income economies cannot always be depicted as smooth convergence of oil intensities. The later industrialisation has started, for example, the less likely that the initial position has large scope for substituting oil in intermediary uses, such as power or industry. The more elastic the demand for end user appliances is with respect to per capita income growth, the higher the likelihood that, as GDP grows, the volume demand caused by the sheer number of new appliances outweighs their efficiency improvements in aggregate. Under these conditions, even leapfrogging (resulting, for example, in a new and highly efficient vehicle fleet) may, at least during the initial period of rapid oil demand growth, not show up as the smooth convergence of oil intensity depicted in the text.

Figure 6: Oil intensity convergence across the three largest consumers



It may appear that speed and reach of the diffusion of technological change embodied in mechanical devices should be the same for end-user appliances and for capital goods; which would imply that globalisation would not affect the scope for fuel substitution determined by the sum total of these devices. However, the data suggest otherwise. It stands to reason that the diffusion of appliances for retail consumption (transport or residential oil consumption) has indeed been faster and more uniformly distributed than the (even on a global scale) lumpier infrastructure required to use oil as an intermediary input (industry and power generation). Shorter life spans for equipment and regular turnover cycles will both accelerate diffusion and smoothen into a more uniform distribution over time, especially in the thick markets for common devices such as cars.

A number of subordinate reasons support the argument: The global retail market for applications is competitive also because it is for the most part private, with consumers meeting profit maximizing suppliers; whereas the use of oil as an intermediary, from power generation to oil intensive industries in particular producer countries, is more likely to encounter governmental interference. Second, 'strategic' industries, of which the latter group is often considered a part, often come with a degree of redundancy reflecting security concerns, not economics; third, the shorter lifespan and turnover period of end-user appliances should encourage a uniform spread of efficiency improvements; finally, it is easier to pass on the cost of slow innovation to customers in oligopolistic and capital intensive (and sometimes government controlled) industries; investments in these industries may be 'lumpy' even on a global scale, and measured oil intensity improvements therefore more discretionary. All of this indicates a disproportionately faster and smoother impact of globalisation on the diffusion of oil conserving end-user appliances, reinforcing the trend toward a smooth, regular oil intensity decline over time.

A second proposition would be that globalisation encourages a uniform spread of efficiency improvements across countries. It should therefore promote convergence of oil intensity across countries, with improvements fastest and smoothest in end-user appliances. Convergence is visible on a global scale since the 1980s.

(c) Structural change

The path of economic development displays a common pattern of structural change, which is connected to energy intensity in an unambiguous fashion: as economies industrialize, i.e. as their dominant economic activity shifts from agriculture to industry, primary energy intensity rises. And as they switch from industrial to the dominance of service sector activities, energy intensity falls. There are only a handful of easily explainable exceptions to this general rule.²⁶

²⁶ The biggest exception are large oil producers in the Middle East, where energy intensity continues to rise, regardless of their sectoral composition. Cf. Rühl and Giljum (2011) and Rühl et al. (2012).

In a first approximation, i.e. keeping the fuel mix constant, this would imply that oil intensity first rises and then falls as economies industrialize and mature. Unfortunately, the fuel mix is unlikely to stay constant during transformations of this scope and magnitude. But there are two other data points which can serve as handrails.

First, as economies mature there is scope for the composition of oil consumption itself to change. For example, the share of oil for intermediate use (power and industry) may be rising during the period of industrialisation. To the extent this intra-fuel adjustment does take place, it will enlarge the playing field for inter-fuel competition later.

These changes operate through changing the stock of mechanized consumer devices and capital goods. Under the common, sensible assumption of irreversible putty-clay technology, investment in highly specific appliances cannot easily be converted back to other uses. To what an extent the capital stock changes (and how fast it can turn over) depends on actual as well as expected relative fuel prices over the lifetime of new equipment.

In principle, the scope for relative prices to change the capital stock is vast: we have seen the surge of oil in industry and in power generation when oil, able to compete on price, penetrated the global economy for the first time in the 1950s and 60s. China's industrialisation proceeded with a greatly decreased demand for oil in key sectors (though not entirely without it); and the dissolution of the Soviet Union has demonstrated how deep modern energy systems will have to cut (to the bone) if the set of relative prices suddenly changes. In all cases, though, the price elasticity of oil has fallen once income levels allowed for a rise in the share of oil destined for final consumption (and with lower substitutability). The reason why the impact of relative price changes on the composition of the capital stock is sparingly documented and often difficult to trace in the data is the obvious one: changes in relative prices exert their influence only over the long term, and they need to be accompanied by durable changes in expectations for the capital stock to become malleable.

Second, the link between changes in a country's sectoral composition and the rise and decline of energy intensity may be unambiguous, but no single fuel needs to mimic this pattern with accuracy. Oil intensity in particular is likely to 'overshoot' as final consumption in transport and residential use, i.e. the price insensitive and non-substitutable part of total oil consumption, continues to expand with income growth, which rises even after the economy has embarked on its transition toward the service sector.

The rise of the service sector decreases primary energy intensity. But the share of oil in the fuel mix, and in principle even oil intensity, may continue to rise for a while longer as the share of non-substitutable consumer appliances in total oil consumption increases. During this transition, oil demand becomes less sensitive to price changes. At the end of it will be the point at which GDP growth outpaces oil demand growth, and oil intensity will start to decline. And eventually, if the trend continues, the oil intensity of GDP will decline faster than GDP, and translate into falling oil demand.

The exact interactions between intra fuel shifts in the substitutability of oil demand, changes in the economy's fuel mix, and the ultimate translation into oil and energy intensities, are a complicated web. But disentangling it (to the extent possible on a conceptual level) leads to a number of conclusions. First, global oil intensity has declined as a larger share of global GDP moves from industry into services. Second, the scope for inter-fuel substitution baked into the capital stock diminishes as the capital stock responds to the demands of higher income levels. Third, oil demand growth will stabilize as a consequence. Through the microcosm of many developing economies the linear trend in oil intensity has been sustained.

(d) Regime change

It now becomes clearer why the impact of high and volatile prices on global economic activity as well as on oil demand has been so substantially different after the turn of the millennium, from the last comparable high price

episode in the 1970s.²⁷ The two oil price shocks in the 1970s ushered in a global recession; but neither the oil price spikes in the run-up to the financial crisis 2008, nor the high price period 2011-14 have been blamed for an adverse impact on global economic performance.

The standard explanation of course is that both events were triggered by completely different circumstances, and this is correct.²⁸ The price spike in the 2000s and its long run-up was the result of rapidly growing demand from booming economies, especially in the developing world; whereas the price shocks in the 1970s were sudden, disruptive and caused by politically motivated supply side events.

Behind the obvious loom the vast structural changes that have taken place in-between these episodes, roughly three decades apart. They can now be condensed into three data points. These are of wider interest because they do not only explain what did not happen in the 2000s, but what may be in store when the energy transition starts to interfere with oil prices and consumption.

First, the role of oil in the global economy has changed in ways which are likely to become irreversible, especially if they are explored and exploited by future energy policy: oil has simply become a less important component of the economy, globally as well as in most economic jurisdictions. This is of course the very essence of arguing that oil intensity has declined by almost two thirds and that systematic and endogenous reasons for this decline can be identified, which are likely to shape the future of oil in the economy even more than they have shaped the past. With oil needs per unit of economic output so much lower, price changes in the precious commodity will have a lower pass through – not only on inflation, but on real variables such as income, employment and productivity.

Second, oil remains the largest fuel globally, but its share is declining. Given the rather inexorable way in which this decline is likely to proceed if the current regime does not change (which is possible, caused for example by a long period of low relative oil prices and sustained expectations of more thereof), and assuming continued economic growth, peak oil demand is a serious prospect. Driven by the adjustments discussed above, this would change oil markets further, and in particular in areas not addressed here, namely on the supply side. A shrinking global market is unlikely to leave space for the role OPEC has played over the entire period from the end of the old regime in the 1970s to the dawn of the present one.²⁹

Third are changes in the composition of oil demand, namely the shift to end-user demand and the associated loss of substitutability. Just as a lower share of oil in the economy will lessen the impact of higher oil prices on the economy, so a higher share of final consumption in total oil consumption will lessen the price impact on total oil demand. This may make price spikes more likely but limits their impact for any given magnitude. If oil demand has become less price elastic, and even more so in countries with higher per capita incomes, some of the more interesting implications concern the intent of making oil more expensive, by carbon pricing or taxation, to curtail its use. Without loosening the current tight link between final consumption services and non-substitutability, i.e., without widening the scope for substitution for final consumption services (to use the hybrid car example again), this may be a tall order, with oil demand much less affected than other fuels by a uniform carbon price across the entire economy.³⁰

²⁷ After the first oil price shock in 1973, GDP slowed down to 0.6% (1976), from a five year average of 5.3% (pre-1973); similarly for the price shock in 1979/80: From a previous average of 4.3% (four years), GDP growth slowed down to 0.4% in 1982. After the run-up to higher prices in 2010-11, global growth actually accelerated (to 2.8%, from a 5-year average of 2.6% for 2007-11).

²⁸ Cf the references listed in footnote 15.

²⁹ The likelihood and likely timing of peak oil demand is explored in Rühl and Erker (2021), the consequences for market structure and OPEC, in Rühl, Mityakov, Portnykh (2021).

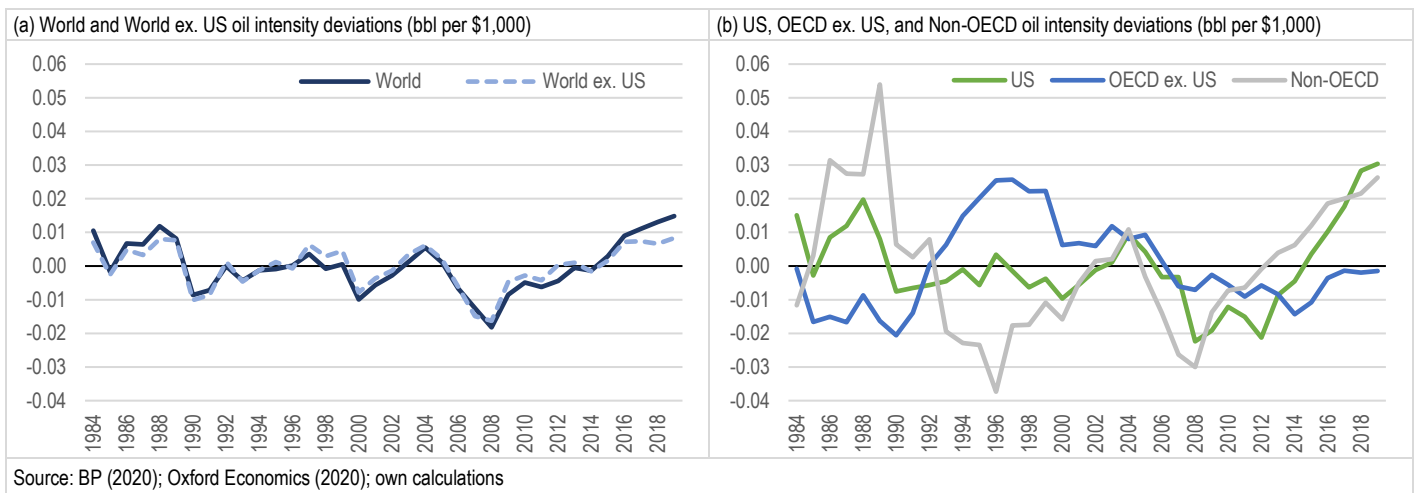
³⁰ Assuming no quantity targets, i.e. that the principle is maintained that, to minimize the social cost of decarbonisation, the price of carbon has to be uniform across all segments of the economy.

(e) Resource abundance

As noted earlier (in Part II (a)), global oil intensity improvements have slowed down since 2009. The relationship of this deceleration to oil prices is highly ambiguous; the link to the slowdown in income growth after the global financial crisis less so. However, although the intensity slowdown has been noted in most countries or regions, it appears to have been driven to a considerable extent by the US, and in the wake of the shale ‘revolution’ – which will make for an interesting test case.

The pace of oil intensity improvements in the rest of the world has fallen by 20% since 2009, to an annual average of 0.008 barrels per \$1,000. In the US, the pace of the decline over the same period halved, to 0.006 barrels per \$1,000 p.a.. The US moderation (and global difference) is even more striking for the period after 2012, after US shale oil production was in full swing (Fig. 7).³¹ Is this heralding a trend break on a global scale, or will it become part of the oscillations discussed earlier?

Figure 7: Global and regional oil intensity deviations from its linear trend line



The jury is out. However, the ongoing episode highlights two important insights: first, the crucial role that not only relative prices but expectations – of prices and resource availability – play in determining the configuration of the capital stock (between final and intermediate uses of oil) which ultimately determines the path of oil intensity; second, the ability of expectations and relative prices to reconfigure the capital stock effectively and in unexpected ways. The data indicate that the deceleration in US oil intensity improvements was not based on a simple, price induced reversal of oil consumption back into areas where it can compete with other fuels (such as substitution in power generation). Instead, it was driven by an increase in final oil consumption – mostly by an immense increase in the share of less fuel efficient CUVs (cross-over utility vehicles), pick-up trucks and SUVs in the US car fleet, i.e. by a simple switch of appliances within final consumption.³² The effect was massive enough to allow for an uptick in the share of oil in US primary energy consumption.

The rise of tight oil production (and the lowering of its production cost) has been a monumental moment in global oil markets, and it has been well publicised. The price impact, however, was global rather than local. Moreover, the price decline caused by shale oil was late (second half of 2014) and muted, courtesy initially of numerous supply disruptions and later, of the formation and activities of OPEC+.

What did change dramatically was the US position in global oil markets as the country moved from the world’s largest oil importer to an exporter (since 2020) in just a few years. It is fair to say that the rapid rise of shale production and

³¹ During 2012-19, the US trend flattened to 0.003 bbl per \$1,000 p.a., down from 0.012 bbl per \$1,000 p.a. in the period 1984-2011; compared to 0.008 bbl per \$1,000 from 0.010 bbl per \$1,000 over the same two periods for the rest of the world – a 70% reduction in the US vs. 20% in the rest of the World.

³² EIA (2018)

its bright future has completely changed perceptions of energy security and resource scarcity among US consumers. Given the absence of a more tangible price impact and the estimated low price elasticity of consumer demand, it will have been that change in long term expectations which caused the switch in the car fleet enabling the rapid adjustment of consumption patterns.³³

No matter whether shale oil will be sufficient to move the dial on the long term trend in global energy intensity, the episode already should provide food for thought for the energy transition – in many ways shale’s antidote, as an attempt to reverse current consumption patterns and to accelerate the oil intensity trend decline.

IV. Conclusion

Oil markets are hard to model, it is true. Like the global economy and global energy markets of which they are a part, they do exhibit patterns of systematic structural adjustment and reconfiguration. Such patterns are traceable and may form a narrative which can be checked for coherence; it is from these narratives that we hope to find a window into what the future may hold.

Oil markets, we have argued, have experienced a regime change over the course of a long and protracted transition. This process is far from finished. Over the last four decades, markets have reconfigured from intermediating a supply constrained, highly versatile and competitive commodity into allocating a commodity which is demand constrained, highly specialized and concentrated in its applications. We have traced this reconfiguration by putting a spotlight on oil intensity: first growing and volatile, as global oil consumption growth outpaced GDP growth; next, declining and regular, but with oil demand still rising, as increasing specialisation ensures the rate of oil intensity improvements remains below GDP growth; and in future, if the narrative holds, with the decline in oil intensity continuing to accelerate (in relative terms) until it outpaces GDP growth, which means oil demand will have to shrink. We are in stage two of this development. Note the obvious, however: oil intensity is just one parameter in a much bigger universe of change. In particular, we have not investigated the supply side at all.³⁴

Nevertheless, the focus allows for a number of conclusions. Two of them are of particular importance in linking this analysis to broader oil market developments.

There is first the intriguing question of how far the claim of a linear time trend for oil intensity is to be trusted: should it be applied as a forecasting tool? Elsewhere we have answered this question with ‘Yes’.³⁵ Given the track record (although replicated ex-post) as well as the consistency of the narrative, the presumption of a continuation of the linear trend should indeed put it on a par with other, more traditional concepts. That means it can be used, for example, to calculate how much oil demand will be lost if global GDP returns to pre-COVID (i.e., 2019) levels only in 2021 or 2022; or to estimate global peak oil demand.

The second relates to energy policy and the energy transition. In many discussions oil is being singled out, typically in its most protected habitat, transport, as a fuel which needs to be eliminated from the global fuel mix and replaced because of its greenhouse gas emissions. The attempt to break oil’s monopoly by deploying technological alternatives which allow for substitution, especially electric vehicles, has started. Here is not the place to comment on this idea’s technical and environmental merits. However, it is widely undisputed, that goodwill, subsidies and regulation alone will not be enough to decarbonize the global transport sector, and that a carbon price would at least be helpful in this regard.

³³ The forward curve appears to support this proposition. The one and the five year strip both indicate lower price expectations in reaction to the advent of sizeable tight oil production growth. However, forward curves also have a tendency to move with prompt prices, a tendency that seems to have been left intact by the rise of shale oil.

³⁴ The impact of a ‘demand peak’ on today’s oil market structure would be dramatic as it would undermine the ability of low-cost producers to raise prices by cutting production – and with it, the OPEC edifice which has been with us for so long (Rühl, Mityakov, Portnykh 2021).

³⁵ Rühl and Erker (2021).

To tease oil out of its most privileged application with the help of a carbon price may become a tall order. To be precise, the analysis gives a dual perspective on transport and the energy transition. Price sensitivity has diminished, and so any carbon price would have less impact on oil in transport than on other fuels in applications with more scope for substitution. To open up options to substitute away from oil (such as the electric car) will become crucial for any decarbonisation strategy targeting final oil consumption, creating a little bit of a 'chicken and egg' problem if alternatives are unknown, or have to be subsidized. Assuming that, to minimize the social cost of decarbonisation, the principle of a uniform carbon price across the economy is maintained (and quantity targets eschewed), the less scope for fuel substitution there is, the higher the carbon price necessary to dent oil consumption. However, this is just a variant of a more general conclusion of this analysis, namely that the scope for price volatility has become larger and the scope for volume volatility has become smaller over time.

So far, the energy transition has had no discernible effect on global oil markets, disruptive or otherwise. The gains in oil conservation from the shift toward final demand have been endogenous, as was the retreat of oil from (now sometimes more carbon intensive) intermediary sectors. We have seen how relative price changes can have an impact on the configuration of the capital stock. In transport, it will take a very large adjustment – for prices and for substitutes.

Appendix

(1) Global, OECD and Non-OECD oil intensity and share of oil in primary energy

Figure A1: Oil intensity with corresponding linear time trend and its specification




Figure A2: Share of oil in total primary energy consumption


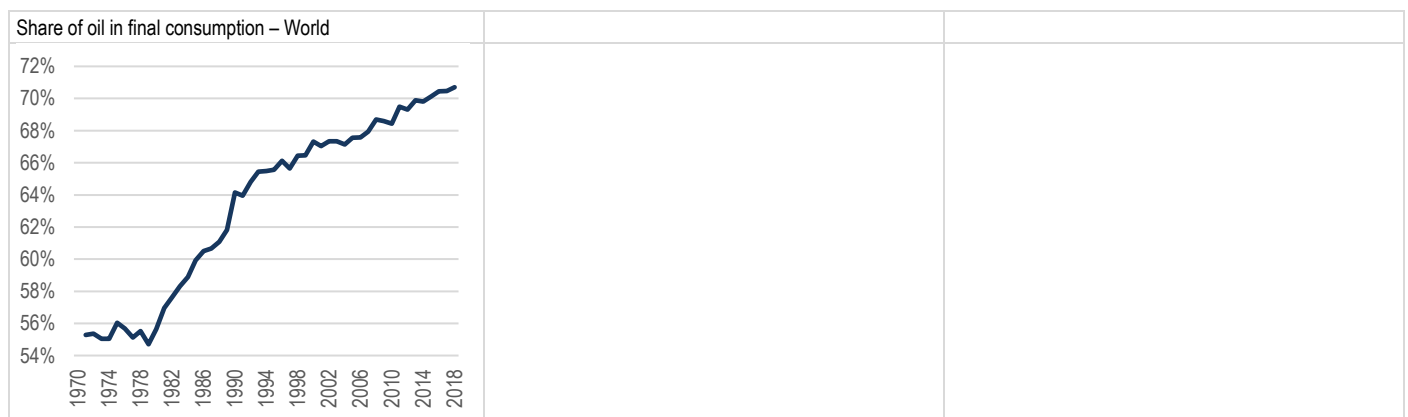


Figure A3: Share of oil used in final consumption (transport and residential use)



Data sources: BP (2020); Oxford Economics (2020); IEA (2020); own calculations.

(2) Model specification

The econometric analysis of the time trend of oil intensity is based on time series data from 1984–2019 (BP Statistical Review of World Energy 2020 and Oxford Economics Global Database June 2020).

Global oil intensity is calculated as:

$$Oilint_t = \frac{Oil_t}{GDP_t}$$

Where:

$Oilint_t$ global oil intensity in year t in barrels of oil equivalent (bbl) per \$1,000 of global GDP in 2015 constant prices and exchange rates;

Oil_t global oil demand in year t in barrels of oil equivalent (bbl) per year; and

GDP_t global GDP in year t in 2015 constant prices and exchange rates.

The time trend of oil intensity is estimated by the following regression:

$$Oilint_t = c + \beta Trend_t + \epsilon_t$$

Where:

$Oilint_t$ global oil intensity in year t in barrels of oil equivalent (bbl) per \$1,000 of global GDP in 2015 constant prices and exchange rates;

$Trend_t$ time trend where year 1984 = 0, up to year 2019 = 35

In addition to the linear trend model, we estimated logarithmic trend model with oil intensity expressed in natural logarithms. Table A1 reports the results.

Table A1: Global oil intensity trend in linear and logarithmic terms – regression results

	Linear trend specification: $Oilint_t = c + \beta Trend_t + \epsilon_t$	Logarithmic trend specification: $Ln(Oilint)_t = c + \beta Trend_t + \epsilon_t$
Sample	1984–2019	1984–2019
Constant	0.746 t = 295.5 (P = 0.000)	-0.270 t = -72.1 (P = 0.000)
Trend	-0.00954 t = -76.9 (P = 0.000)	-0.01664 t = -90.5 (P = 0.000)
Observations	36	36
R ²	0.994	0.996

Data sources: BP (2020); Oxford Economics (2020); own calculations in EViews.

Table A2 reports the comparison of both trend specification for the OECD and the EU – two economic blocks in which oil demand has peaked. For these more mature economies, the linear trend specification exhibited the better fit.

Table A2: OECD and EU oil intensity trend in linear and logarithmic terms – regression results

	Linear trend specification: $Oilint_t = c + \beta Trend_t + \varepsilon_t$		Logarithmic trend specification: $Ln(Oilint)_t = c + \beta Trend_t + \varepsilon_t$	
	OECD	EU	OECD	EU
Sample	1984–2019	1984–2019	1984–2019	1984–2019
Constant	0.654 t = 212.8 (P = 0.000)	0.525 t = 207.7 (P = 0.000)	-0.392 t = -44.0 (P = 0.000)	-0.607 t = -63.2 (P = 0.000)
Trend	-0.00917 t = -60.7 (P = 0.000)	-0.00780 t = -62.7 (P = 0.000)	-0.01906 t = -43.5 (P = 0.000)	-0.0206 t = -43.7 (P = 0.000)
Observations	36	36	36	36
R ²	0.991	0.991	0.982	0.983

Data sources: BP (2020); Oxford Economics (2020); own calculations in EViews.


(3) Cross-validation of the linear trend specification

In a first step we evaluate the predictive power of our linear specification for the global, the OECD and the EU oil intensity model. The three regressions were run on 15 overlapping samples (so-called training periods) starting with period 1984–2003 and up to 1984–2017, in annual increments. Based on the estimated coefficients we made out-of-sample forecasts from 1 year ahead up to 14 years ahead, which we compared to the actual oil intensities, to calculate root mean squared deviations (RMSDs).

For example, based on the 1984–2003 training period (of 30 years) we estimated the linear time trend coefficient, which was then used to generate forecasts from 1 and up to 14 years ahead (i.e. for years 2004, 2005, and up to 2017). We repeated the same procedure for the 1984–2004 training period, and up to the 1984–2017 training period. We generated 15 one and two year ahead forecasts, 14 three year ahead forecasts, up to three 14 year ahead forecasts.

We calculated the corresponding RMSDs and expressed them in relative terms, compared to the corresponding actual average oil intensity.

Figure A3: Cross-validation results



Data sources: BP (2020); Oxford Economics (2020); own calculations in EViews.

(4) Cross-validation comparison of the linear and logarithmic model specification


We compared the predictive power of the linear and the logarithmic specification for the global, the OECD and the EU oil intensity model. The six regressions were run on 15 overlapping samples starting with the period 1984–2003, up to the period 1984–2017, in annual increments.

We compared the predictive power of all 15 regressions for 1 to 14 forecast years ahead (or lags), for the global and each regional pair.

The results are consistent with the more general R^2 results: For global oil intensity the logarithmic specification exhibits slightly better predictive power at all forecast legs; for the OECD and the EU the linear trend specification shows a significant forecasting advantage.

Figure A3 shows the difference between the root mean squared deviations (RMSDs), normalised by the mean intensity, between the two regression specifications. A positive difference indicates that the logarithmic specification has better predictive power.

Figure A4: Cross-validation results comparing linear and logarithmic specifications



Data sources: BP (2020); Oxford Economics (2020); own calculations in EViews.

References

- Baumeister, C., Peersman, G. Robays, I (2009), The Economic Consequences of Oil Shocks: Differences across Countries and Time, Reserve Bank of Australia, Conference 2009.
- Blanchard, O.J., Galí, J. (2007), The Macroeconomic Effects of Oil Price Shocks: Why Are the 2000s so Different from the 1970s?, NBER Working Paper, No 13368.
- BP (2011), Energy Outlook 2030, London. <https://www.bp.com/content/dam/bp/business-sites/en/global/corporate/pdfs/energy-economics/energy-outlook/bp-energy-outlook-2011.pdf>
- BP (2020), The Statistical Review of World Energy, London. <https://www.bp.com/en/global/corporate/energy-economics/statistical-review-of-world-energy.html>
- Cashin, P., Mohaddes, K., Raissi, M., Raissi, M. (2014), The differential effects of oil demand and supply shocks on the global economy, Energy Economics, 44 (July).
- Dargay, J. (1992). "The Irreversible Effects of High Oil Prices: Empirical Evidence for the Demand for Motor Fuels in France, Germany, and the UK," in: Howdon D (ed.), Energy Demand: Evidence and Expectations, ed. D. Hawdon, London: Academic Press.
- Dargay, J., Gately D. (1994), Oil Demand in the Industrialized Countries, The Energy Journal Special Issue, Frank H. (ed.) The Changing World Petroleum Market.
- EIA (2018), Today in Energy, Crossover utility vehicles overtake cars as the most popular light-duty vehicle type, July 2018, <https://www.eia.gov/todayinenergy/detail.php?id=36674#>
- Fagan M., Duan J. (2019), Up the Down Staircase: What History Teaches Us about Oil Demand after a Crisis, Boston: London Economics International LLC, <https://ssrn.com/abstract=3592443>
- Fagan, M. (2020), Taking a look ahead: the long term impact of a crisis on oil demand, Presentation at the United States Energy Association (USEA), May 27, <https://usea.org/event/taking-look-ahead-long-term-impacts-oil-demand-after-crisis>
- Gately, D., Huntington, H. (2002) The Asymmetric Effects of Changes in Price and Income on Energy and Oil Demand, The Energy Journal, Vol 23, No. 1
- Gately, Dermot (1992). "Imperfect Price-Reversibility of U.S. Gasoline Demand: Asymmetric Responses to Price Increases and Declines," The Energy Journal 13(4): 179-207.
- Gately, Dermot (1993). "The Imperfect Price Reversibility of World Oil Demand," The Energy Journal 14(4): 163-8
- Hamilton, J.D. (2009), Causes and Consequences of the Oil Shock of 2007-08, Brookings Papers on Economic Activity, 40.
- Hamilton, J.D. (2009b), Understanding Crude Oil Prices, Energy Journal, 30 (2).
- Hogan, W. (1993), OECD Oil demand dynamics: Trends and Asymmetries, The Energy Journal, Vol 14, No 1.
- IEA (2020); *Oil total final consumption by sector, 1971-2018*, IEA, Paris <https://www.iea.org/data-and-statistics/charts/oil-total-final-consumption-by-sector-1971-2018>
- Kilian, L. (2009), "Not All Oil Price Shocks Are Alike: Disentangling Demand and Supply Shocks in the 26 Crude Oil Market", American Economic Review, 99, 1053-1069. <https://doi.org/10.1257/aer.99.3.1053>
- Mirzoev T., Zhu L., Yang Y., Zhang T., Roos E. (2020), The Future of Oil and Fiscal Sustainability in the GCC Region, IMF Working Paper No.20/01, <https://www.imf.org/en/Publications/Departmental-Papers-Policy-Papers/Issues/2020/01/31/The-Future-of-Oil-and-Fiscal-Sustainability-in-the-GCC-Region-48934>

Oxford Economics (2020), Global Economics Databank, June 2020 www.oxfordeconomics.com

Rühl, C., Giljum, J. (2011), BP Energy Outlook 2030, International Association of Energy Economists: Energy Forum, 3rd Quarter,

http://www.google.co.uk/url?sa=t&rct=j&q=joe%20giljum%20iaee%202030&source=web&cd=1&ved=0CC0QFjAA&url=http%3A%2F%2Fwww.iaee.org%2Fen%2Fpublications%2Fnewsletterdl.aspx%3Fid%3D136&ei=I7VtT7jJLsrZ4QSlIJS_Ag&usg=AFQjCNHFVAtSDO4Dxxd_oVyG7HAMMRvrKg

Rühl, C., P. Appleby, J. Fennema, A. Naumov, M. Schaffer (2012), Economic Development and the Demand for Energy: A Historical Perspective on the Next 20 Years, Energy Policy, 50 (November).

Rühl, C., Erker, T. (2021), Causes and Consequences of Peak Oil Demand, manuscript, CGEP, Columbia University

Rühl, C., Mityakov, S., Portnykh, M (2021), Oil Market Structure and the Future of OPEC, manuscript, Florida State University.

Samuelson, R. (2013), The Washington Post (30 May), https://www.washingtonpost.com/opinions/robert-j-samuelson-is-steins-law-real/2013/05/30/716942f2-c942-11e2-8da7-d274bc611a47_story.html