

Understanding Crude Oil Prices*

James D. Hamilton

Department of Economics
University of California, San Diego

May 22, 2008

Revised: June 4, 2008

ABSTRACT

This paper examines the factors responsible for changes in crude oil prices. The paper reviews the statistical behavior of oil prices, relates these to the predictions of theory, and looks in detail at key features of petroleum demand and supply. Topics discussed include the role of commodity speculation, OPEC, and resource depletion. The paper concludes that although scarcity rent made a negligible contribution to the price of oil in 1997, it may be an important feature of the most recent data.

*I thank Menzie Chinn for helpful comments on an earlier draft.

1 Introduction.

How would one go about explaining what oil prices have been doing and predicting where they might be headed next? This paper explores three broad ways one might approach this. The first is a statistical investigation of the basic correlations in the historical data. The second is to look at the predictions of economic theory as to how oil prices should behave over time. The third is to examine in detail the fundamental determinants and prospects for demand and supply. Reconciling the conclusions drawn from these different perspectives is an interesting intellectual challenge, and necessary if we are to claim to understand what is going on.

In terms of statistical regularities, the paper notes that changes in the real price of oil have historically tended to be (1) permanent, (2) difficult to predict, and (3) governed by very different regimes at different points in time.

From the perspective of economic theory, we review three separate restrictions on the time path of crude oil prices that should all hold in equilibrium. The first of these arises from storage arbitrage, the second from financial futures contracts, and the third from the fact that oil is a depletable resource. We also discuss whether commodity futures speculation by investors with no direct role in the supply or demand for oil itself could be regarded as a separate force influencing oil prices.

In terms of the determinants of demand, we note that the price elasticity of demand is challenging to measure but appears to be quite low and to have decreased in the most recent data. Income elasticity is easier to estimate, and is near unity for countries in an

early stage of development but substantially less than one in recent U.S. data. On the supply side, we note problems with interpreting OPEC as a traditional cartel and with cataloging intermediate-term supply prospects despite the very long development lead times in the industry. We also relate the challenge of depletion to the past and possible future geographic distribution of production.

Our overall conclusion is that the low price-elasticity of short-run demand and supply, the vulnerability of supplies to disruptions, and the peak in U.S. oil production account for the broad behavior of oil prices over 1970-1997. Although the traditional economic theory of exhaustible resources does not fit in an obvious way into this historical account, the profound change in demand coming from the newly industrialized countries and recognition of the finiteness of this resource offers a plausible explanation for more recent developments. In other words, the scarcity rent may have been negligible for previous generations but is now becoming significant.

2 Statistical predictability.

Figure 1 plots the inflation-adjusted real price of oil since 1947. One's first thought might be that someone has pasted together two or more radically different series. We'll offer in this paper a perspective on why that indeed might be a good way to think about this series. But for the moment let's just take the data from 1970:Q1 to 2008:Q1 as a single sample and ask, How predictable statistically is the change in the real price of oil over this period?

Let p_t denote 100 times the natural log of the real oil price in Figure 1 as of the third

month of quarter t and let Δp_t denote the quarterly percentage change. The average value of Δp_t over 1970-2008 is 1.12– the real price of oil has increased on average by over 1% per quarter for the last 40 years. However, the t statistic for that average growth estimate is 0.91, meaning one readily accepts the hypothesis that the expected oil price change could be zero or even negative.

One can also explore simple forecasting regressions of the form

$$\Delta p_t = \beta' \mathbf{x}_{t-1} + \varepsilon_t \tag{1}$$

where \mathbf{x}_{t-1} is a vector of variables one would have known the quarter prior to t that might have helped predict the oil price change in quarter t . Table 1 reports the results of testing for such predictability when \mathbf{x}_{t-1} is based on the observed lagged behavior of real oil prices, U.S. nominal interest rates, or U.S. GDP growth rates. Those tests for predictability are summarized by the p -value associated with the hypothesis test– if a p -value is below 0.05, we would reject the null hypothesis at the 5% level, and conclude that the indicated \mathbf{x}_{t-1} could help predict the change in oil prices. The table shows that in fact there is no basis in the historical data for claiming to be able to predict oil price changes using any of the variables listed.

How about predicting the level of p_t rather than the rate of change? One test for whether we want to be specifying forecasting regressions in levels or rates of change is the augmented Dickey-Fuller test (e.g., Hamilton, 1994, pp. 528-9), in which one looks for whether the lagged level gives you information helpful for predicting the change over and above that contained in lagged changes. This can be implemented by testing the null hypothesis that

$\eta = 0$ in the following regression:

$$\Delta p_t = \eta p_{t-1} + \zeta_1 \Delta p_{t-1} + \zeta_2 \Delta p_{t-2} + \zeta_3 \Delta p_{t-3} + \zeta_4 \Delta p_{t-4} + \varepsilon_t.$$

The t statistic for testing this hypothesis turns out to be +0.69, whereas one would need a value less than -1.95 to reject the hypothesis. Alternatively, as in Kwiatowski, et. al. (1992) one can take as the null hypothesis that the forecasting regressions should really be estimated in levels. The KPSS $\hat{\eta}_\tau$ statistic exceeds 0.32 for all lag windows ℓ between 0 and 4; for any value above 0.22 we would reject the null hypothesis at the 1% level.

All of the above test results are consistent with the claim that the real price of oil seems to follow a random walk without drift. The price increased over the sample by 172% (logarithmically), but a process like this one could just as easily have decreased by a comparable amount. While one might have forecasting success with more detailed specifications over shorter samples, the broad inference with which we come away is that the real price of oil is not easy to forecast. To predict the price of oil one quarter, one year, or one decade ahead, it is not at all naive to offer as a forecast whatever the price currently happens to be.

Although you might be fully justified in offering “no change” as your “best” short- and long-run prediction for oil prices, it’s worth emphasizing how far wrong the forecast is likely to prove to be. Let’s take for illustration the price of oil as of this writing (\$115/barrel). The standard deviation of Δp_t over the sample is $\sigma = 15.28\%$. If one took these log changes as having a Gaussian distribution, that would mean your forecast one quarter from now (2008:Q2) would have a 95% confidence interval ranging from a low of \$85 dollars a barrel

to a high of \$156.¹ As you try to forecast s quarters into the future, the standard error for a random walk becomes $\sigma\sqrt{s}$. Table 2 gives some flavor for how the forecasts deteriorate the farther you try to peer into the future. Four years from now, we may “expect” the price of oil still to be at \$115 a barrel, though we would in fact not be all that surprised if it is as low as \$34 or as high as \$391!

3 Predictions from theory.

We turn next to a discussion of what economic theory predicts for the dynamic behavior of crude oil prices. We will describe three separate conditions, all of which should hold in equilibrium.

3.1 Returns to storage.

Consider the following possible investment strategy. You borrow money today (denoted date t) in order to purchase a quantity Q barrels of oil at a price P_t dollars per barrel. Suppose you pay a fee to the owner of the storage tank of C_t dollars for each barrel you store for a year. Then you’ll need to borrow $(P_t + C_t)Q$ total dollars, and next year you’ll have to pay this back with interest. Let i_t denote the interest rate measured as a fraction of 1; for example, $i_t = 0.05$ would correspond to a 5% annual interest rate. Then next year you’ll have to pay back $(1 + i_t)(P_t + C_t)Q$ dollars. But you’ll have the Q barrels of oil that you can sell for next year’s price, P_{t+1} . If oil prices go up by so much that you can sell the

¹ Note that the confidence intervals are symmetric in logs but asymmetric in levels.

oil for more than your costs, that is, if

$$P_{t+1}Q > (1 + i_t)(P_t + C_t)Q,$$

then you'll make a profit from putting more oil into storage today.

Of course, you don't know today what next year's price of oil will be, but you have some expectation based on information currently available, which we'll denote $E_t P_{t+1}$. Working with the above expression, we see that you'd expect to make a profit from oil storage whenever

$$E_t P_{t+1} > P_t + C_t^* \tag{2}$$

where C_t^* reflects your combined interest and physical storage expenses:

$$C_t^* = i_t P_t + (1 + i_t)C_t.$$

Suppose people did expect $E_t P_{t+1}$ to be greater than $P_t + C_t^*$. Then anyone could expect to make a profit by buying the oil today, storing it, and selling it next year. If there are enough potential risk neutral investors, the result of their purchases today would be to drive today's price P_t up. Knowledge of all the oil going into inventory today for sale next year should reduce a rational expectation of next year's price $E_t P_{t+1}$. As long as the inequality (2) held, speculation would continue, leading us to conclude that (2) could not hold in equilibrium.

What about the reverse inequality,

$$E_t P_{t+1} < P_t + C_t^*?$$

Then anyone putting oil into storage is expecting to lose money, and it would not pay to do so for purposes of pure speculation. That doesn't mean that every storage tank will be

empty, because inventories of oil are essential for the business of transporting and refining oil and delivering it to the market. We could think of such factors as equivalent to a “negative” storage cost for oil in the form of a benefit to your business of having some oil in inventory, which is referred to as a “convenience yield”. We might then refine the above specification, subtracting any convenience yield from physical and interest storage costs C_t^* to get a magnitude $C_t^\#$, the net cost of carry. But it’s clear that if people expect oil prices to fall so much that

$$E_t P_{t+1} < P_t + C_t^\#,$$

then there is an incentive to sell oil out of inventories today, driving P_t down and $C_t^\#$ up.

We’re then led to the conclusion that the following condition should hold in equilibrium

$$E_t P_{t+1} = P_t + C_t^\#. \tag{3}$$

We could in principle modify our definition of the cost of carry $C_t^\#$ further to incorporate any risk premium that may induce investors to want to hold more or less inventories.

Insofar as expectations, convenience yield and risk premia are impossible to observe directly, one might think that (3) does not imply any testable restrictions on the observed relation between P_{t+1} and P_t . However, recall that the quarterly change in real oil prices has a standard deviation of 15% (see Figure 2), and increases much larger than this are observed quite often. It seems inconceivable that risk aversion or convenience yield would exhibit quarterly movements of anywhere near this magnitude. One refutable implication of (3) is that the big changes in crude oil prices should be mostly unpredictable, and given that it is the big changes that dominate this series statistically, the finding in the previous section

that oil price changes are very difficult to predict is exactly what the theory sketched here would lead us to expect.

It is sometimes argued that if economists really understand something, they should be able to predict what will happen next. But oil prices are an interesting example (stock prices are another) of an economic variable which, if we really understand it, we should be completely unable to predict.

3.2 Futures markets.

If you thought oil prices were headed higher, there is an alternative investment strategy to buying oil today and physically storing it. You could instead enter into a futures contract, which would be an agreement you reach today to buy oil one year from now at some price, F_t , to which price you and the counterparty agree today. To implement this contract through an organized exchange, you and the counterparty would both have to set aside some funds today (the margin requirement) to prove you could fulfill the contract, and add to these funds as needed if the potential obligation grows. Abstracting from margin requirements and broker's costs, if you've agreed to buy oil at the price F_t , you will make money whenever $F_t < P_{t+1}$, because you could in this event sell the oil for which you pay F_t to someone else on next year's spot market at price P_{t+1} , pocketing the difference as pure profit. If your expectations were such that $F_t < E_t P_{t+1}$, everybody would want to be on the buy side of such contracts, bidding the terms of the contract F_t up. Equilibrium requires

$$F_t = E_t P_{t+1} + H_t^\# \tag{4}$$

where $H_t^\#$ is again a term incorporating any risk premium or complications induced by margin requirements.

Note that (4) is not an alternative theory to (3)—*both* conditions have to hold in equilibrium. For example, if there were an increase in F_t without a corresponding change in P_t , that would create an opportunity for someone else to buy spot oil at time t for price P_t , store it for a year, and sell it through the futures contract at currently guaranteed price of F_t . Taking oil off the current spot market and putting it into inventory would continue until P_t increased to the point at which

$$F_t = P_t + C_t^\#.$$

If we chose to ignore cost of carry and risk premia, conditions (4) and (3) together would imply that the futures price simply follows the current spot price

$$F_t = P_t. \tag{5}$$

In practice, one finds in the data that the futures price and spot price differ, but often not by much, and when news causes the spot price to go up or down on a given day, futures prices at every horizon usually all move together in the same direction as the change in spot prices. Figure 3 plots the futures prices for a couple of representative days. On August 21, 2007, one could buy oil at any future horizon between 4 months and 8 years for between \$67.49 at \$68.70 per barrel. Over the next two months, spot and futures prices at every horizon rose substantially, though the spot and near-term contracts went up more quickly than the farther-out contracts, so that by October 4, the near-term futures prices were substantially above those for longer-term contracts.

To the extent that F_t and P_t differ, studies by Bopp and Lady (1991), Abosedraa and Baghestani (2004), and Chinn, LeBlanc and Coibion (2005) found that P_t provides as good or even a better forecast of P_{t+s} than does the futures price F_t . Interestingly, all three studies nevertheless also accepted the hypothesis that F_t embodies a rational expectation of the future spot price. The overall conclusion we might draw is that P_t offers about as good a forecast of the future spot price as one can achieve, but, recalling Table 2, even the best forecast is none too accurate.

3.3 Scarcity rent.

If we were discussing the price of a standard competitively produced good, we would have the additional equilibrium condition that the price should equal the marginal cost of production. However, oil is a depletable resource— it is mined rather than produced, and once burned, cannot be reused. Harold Hotelling pointed out back in 1931 that in the case of an exhaustible resource, price should exceed marginal cost and would exhibit particular dynamic behavior over time even if the oil market were perfectly competitive.

To understand Hotelling’s principle, suppose we take it as given that as a result of unavoidable geological limits, global production of crude oil next year could only be 90% of the amount being produced this year. If we assumed say a short-run demand price elasticity of -0.10, that would imply a price of oil next year that is twice its current value. As we noted above, under such a hypothetical scenario it would pay anyone to buy the oil today in order to store it in a tank for a year, waiting to sell into next year’s more favorable market.

It would be more efficient, however, for the owner of any oil reservoir to “store” the oil

directly by just leaving it in the ground, waiting to produce it until the price has risen. In a competitive equilibrium, the owner of the reservoir will receive a compensation for surrendering use of the nonreproducible resource that leaves them just indifferent between producing today and producing in the future.² We can think of that scarcity rent per barrel at time t , denoted λ_t , as the difference between price P_t and marginal production cost M_t :

$$\lambda_t = P_t - M_t.$$

If the owner produces the oil today and invests the profits at interest rate i_t , next year the owner would have $(1 + i_t)\lambda_t$. If that's bigger than the benefit from producing next year, that is, if

$$(1 + i_t)\lambda_t > \lambda_{t+1},$$

the owner is better off producing more today and leaving less in the ground. If the inequality were reversed, the owner is better off producing less. Thus in equilibrium Hotelling's principle postulates that $\lambda_{t+1} = (1 + i_t)\lambda_t$ or

$$P_{t+1} - M_{t+1} = (1 + i_t)(P_t - M_t); \tag{6}$$

the gap between price and marginal cost should rise over time at the rate of interest.

Although this is an elegant theory, a glance at Figure 1 gives us an idea of the challenges in using it to explain the observed data. The real price of oil declined steadily between

² Mathematically, with perfect information, λ_t would correspond to the Lagrange multiplier (sometimes referred to as the "shadow price") associated with the constraint that the sum of production over all time cannot exceed a given finite number corresponding to ultimate recoverable reserves; see for example Krautkraemer (1998, p. 2067).

1957 and 1967, and fell quite sharply between 1982 and 1986. One can try to modify the simple Hotelling framework to allow for technological progress, which could induce a downward trend in marginal production cost that for a while at least causes P_t to fall even though $P_t - M_t$ is rising. Alternatively, one can allow for unanticipated resource discoveries producing an unanticipated downward shift in an otherwise upward-trending time path for λ_t . Krautkraemer (1998) surveys some of the literature in this area, a fair summary of which might be that efforts along these lines are ultimately not altogether satisfying. As a result, many economists often think of oil prices as historically having been influenced little or none at all by the issue of exhaustibility.

There is certainly no theoretical problem with postulating that in 1997, future supply prospects were sufficiently strong, and the perceived date at which the limit of ultimately recoverable reserves would begin to affect decisions was sufficiently far into the future, that the scarcity rent λ_t at that time could have been negligible relative to costs of extraction for the marginal producer. New information about surprisingly strong demand growth prospects and limits to expanding production could in principle account for a sudden shift to a regime in which λ_t is positive and quite important.

Such an interpretation would still be inconsistent with the downward-sloping futures term structure in October 2007 noted in Figure 3, which from (4) would be difficult to square with the view that λ_t comprises a significant component of P_t and furthermore is expected, as the theory predicts, to rise over time. On the other hand, it is sovereign governments rather than private firms that control the vast majority of remaining petroleum reserves,

and although their decisions may not perfectly implement (6), one can make a case that the intertemporal calculation has started to influence current production decisions. For example, Reuters news service reported the following story on April 13, 2008:

Saudi Arabia's King Abdullah said he had ordered some new oil discoveries left untapped to preserve oil wealth in the world's top exporter for future generations, the official Saudi Press Agency (SPA) reported.

"I keep no secret from you that when there were some new finds, I told them, 'no, leave it in the ground, with grace from god, our children need it'," King Abdullah said in remarks made late on Saturday, SPA said.

The November 2006 Energy Information Administration Country Analysis Brief on Kuwait included the following:

"Project Kuwait," to be developed over 25 years, was first formulated in 1997 by the SPC, to increase the country's oil production by 500,000 (and to help compensate for declines at the mature Burgan field)....

However, the controversy over Kuwait's reserve figure could have a significant impact on the country's capacity expansion plans. Opposition MPs have called for production to be kept within 1 percent of reserves in order to ensure that oil is available for future generations, though the proposal has not yet been passed into law. Even taking the 100-billion-barrel figure, the 1 percent limit would restrict

Kuwait's production to under 3 million bbl/d, increasing difficulty of efforts to pass the Project Kuwait legislation.

3.4 Role of speculation.

Masters (2008) estimated that assets allocated to commodity index trading strategies had risen from \$13 billion at the end of 2003 to \$260 billion as of March 2008. These funds hold a portfolio of near-term futures contracts (of which about 70% represent energy prices), following a strategy of selling the expiring contract the second week of the month and using the proceeds to buy the subsequent month's contract.

If investors were risk neutral and equally informed, we would not expect the volume on the buy side to have any effect on the price. In such a world, there would be an unlimited potential volume of investors out there willing to take the other side of any bets if the purchases were to result in a price that was anything other than the market fundamentals value. But with risk-averse investors or with differing information, the answer is a little different. For example, I might read your willingness to buy a large volume of these contracts as a possible signal that you know something I don't. For this reason, standard financial market micro-structure theory (e.g., Dufour and Engle, 2000) predicts that a large volume of purchases may well cause the price to increase, at least temporarily, until I have a chance to verify what the true fundamentals value would be. DeLong, et. al. (1990) described a case in which risk-averse investors would never fully arbitrage away ill-informed speculators who are simply pouring money into any asset that has recently experienced high rates of return. In the case of a product for which the Hotelling Principle applies, Jovanovic (2007) noted

that self-fulfilling bubble paths could be indexed by the residual quantity of oil that never gets produced. Determining the current price associated with hitting complete exhaustion (that is, the price path that satisfies the intertemporal Hotelling constraint) is a daunting task given real-world uncertainties, and one could imagine that considerable time might be required for any price impact of commodity “noise investor” speculators to be undone by other market participants.

Suppose we believed that speculation as a force in and of itself could succeed in driving the futures price up. The buyer of spot crude oil would be a refiner, whose primary decision given gasoline demand is an intertemporal one. It can meet that demand with crude oil that it purchases at the current spot price, or produce out of inventory buying its crude forward at the futures price. If the futures price were to increase with the spot price fixed, there would be a big increase in the demand for spot oil. If we thought of gasoline demand as completely price-inelastic in the short run, the demand curve for spot crude would shift up by \$1 per barrel when the futures price increased by \$1. As a result, the speculators who are selling the expiring near-term contracts would find that they have indeed made a profit in an environment in which an ever-increasing volume of futures purchases drives ever-increasing futures and spot prices.

Although it might appear that we have described a self-fulfilling speculative price bubble here, in reality it is not, because the demand for gasoline is in fact not completely price inelastic at all prices. Ultimately there are physical producers of crude oil and physical consumers of gasoline, and insofar as the activities of either have any response at all to the

price, incentives for consumption would be reduced and incentives for production increased whenever the price of crude oil is driven up. For this reason, an ongoing speculative price bubble would have to result in continuous inventory accumulation, or else be ratified by cuts in production. The former is clearly unsustainable, and if it is the latter, one might make the case that the supply cuts rather than the speculation itself has been the ultimate cause of the price increase.

4 Petroleum demand.

4.1 Price elasticity.

The price elasticity of petroleum demand measures the percentage change in quantity demanded divided by the percentage change in price as we move along a given demand curve. Although this is easy enough to define as a theoretical concept, Figure 4 reminds us why it is difficult to measure in practice. Both the supply and demand in any given year t are responding to any of a number of factors besides the current price. Important among these other factors are income (a key determinant of demand) and previous years' prices. The latter is important for both demand, since it can take many years for the fleet of existing cars to reflect changes in purchasing habits, and supply, since tremendous lead times are required between initial exploration and eventual production. In any given year, both the demand curve and supply curve are shifting as a result of these factors, and one cannot simply look at how price and quantity move together to infer anything about the slope of either curve.

There are, however, some episodes in which we can be pretty confident that the most important factor was a shift in the supply curve brought about by exogenous geopolitical

events. Figure 5 plots monthly oil production figures for Iran, Iraq, and Kuwait. Iranian production fell by 5.4 million barrels per day in the immediate aftermath of the 1978 revolution, a loss representing 8.9% of total world production at the time (Hamilton, 2003, p. 390). Production from Iraq fell an additional 3.1 mb/d when that country went to war with Iraq in 1980. These supply disruptions were the primary reason that the real price of crude oil increased 81.1% (logarithmically) between January 1979 and the peak in April 1980. Between 1978 and 1981, U.S. oil consumption fell 16.0% while U.S. real GDP increased by 5.4%. If we assumed a unit income elasticity, one would have expected oil consumption to have risen by 5.4% rather than declined by 16%, for a net decrease in quantity demanded of 21.4% and an implied short-run price elasticity of

$$\frac{\Delta \ln(Q)}{\Delta \ln(P)} = \frac{0.214}{0.811} = 0.26. \quad (7)$$

Similar estimates of the short-run price elasticity were arrived at econometrically by Edelstein and Kilian (2007) and many of the studies surveyed by Dahl and Sterner (1991). The long-run elasticity is often estimated to be more than twice as large (Roya, et. al., 2006; Dahl and Sterner, 1991).

On the other hand, the relative price of oil increased 108.8% (logarithmically) between January 2002 and January 2006, despite which U.S. oil consumption actually increased 4.1% between 2002 and 2006. With U.S. real GDP growth of only 11.9% over this period, it is difficult to reach any conclusion other than that the price-elasticity of demand is even smaller now than it was in 1980. For example, Hughes, Knittel, and Sperling (2008) estimated that short-run demand elasticity was in the range of 0.21 to 0.34 over 1975-1980 but between

only 0.034 and 0.077 for the 2001-06 period.

The quite low price elasticity in combination with the dramatic supply disruptions in Figure 5 account for one key aspect of oil price behavior. With supplies of the resource dependent on geopolitically unstable regions, huge price swings are necessary to restore equilibrium after events like the embargo from the Organization of Arab Petroleum Exporting Countries (1973-74), Iranian Revolution (1978), Iran-Iraq War (1980), and Iraq's invasion of Kuwait (1990).

4.2 Income elasticity.

Figure 6 conveys the comovement between U.S. real GDP and oil consumption since World War II. The horizontal axis measures the cumulative logarithmic change in real GDP at a given date relative to where it was in 1949, so that two years separated by a distance of 0.1 on the horizontal axis correspond to a growth of real GDP of about 10% between those two years. The vertical axis measures the cumulative logarithmic change in U.S. oil consumption. If the points were all on the 45 degree line, it would mean that oil consumption always grows at exactly the same percentage rate as real GDP. In fact U.S. oil consumption grew faster than GDP over the first decade, consistent with an income elasticity of 1.2. The slope of the curve decreased slightly over the next decade, though the 1960s could still be claimed to be characterized by an income elasticity greater than unity. One then sees a significant adjustment following the 1973-74 oil shock and the much more dramatic 1979-82 adjustment already mentioned. It is interesting however that over the period from 1985-1997, oil use in percentage terms grew half as fast as real GDP, despite the fact that the real price of oil fell

43% over this period, suggesting that the income elasticity of U.S. petroleum demand has decreased significantly over time.

The combination of an income and price elasticity both well below unity accounts for the broad trends we see in the share of oil purchases in total expenditures over time. Price inelasticity means that if the price of oil goes up, total expenditures on oil go up. Income inelasticity means that as GDP goes up, the share of oil expenditures should fall. Figure 7 reveals that big price drops and growing GDP during the 1980s and 1990s together brought the dollar value of oil expenditures as a share of total GDP down to 1.1% in 1998, a small fraction of the 8.3% share reached at the peak in 1980. The price increases since 1998 have brought the share back up to 4.0% for 2007.

Note that the 2007 figure is based on the average oil price for 2007 of \$72/barrel, whereas the price in 2008 has typically been well over \$100 a barrel. It is interesting to calculate what the share would be if oil use and GDP remained at their 2007 values but we instead assumed a price of \$100 or \$150 a barrel. The latter would put us back to the all-time peak of 1980. This calculation should remind us that a price elasticity cannot be globally less than unity—price increases would then always increase expenditure shares, but the expenditure share can never exceed unity. The low expenditure share enabled consumers to shrug off the price increases of the early 2000s, but it seems extremely unlikely that such low price elasticities would continue to be observed if the price continues to increase.

The impression from U.S. data that the income elasticity has declined as GDP per person has increased is confirmed in data from a number of different countries. Figure 8 plots

average growth rates in petroleum consumption for France, Germany, and Japan over the last half century. These began at phenomenal rates in 1960 but have continued to decline over time. Figure 9 establishes that for a group of 11 important countries, the poorer the country was in 1960, the faster its growth in oil demand over the last half of the twentieth century. Gately and Huntington (2002) estimated an average income elasticity over 1971-1997 of 0.55 for 25 OECD countries but 1.17 for 11 other countries characterized by rapid income growth over the period and 1.11 for 11 oil-exporting countries..

And it is the latter countries where petroleum growth is coming from at the moment, aggravated by gasoline subsidies in many of the oil producing countries. Although the U.S. and Europe still account for almost half of all the oil used globally, these areas account for less than 1/5 of the increase in world consumption between 2003 and 2006.³ Instead the growth is coming from the rapidly growing countries and oil exporters, with the countries in the Middle East accounting for 17% of the growth and China alone accounting for 33%. China's demand grew at a phenomenal 7.2% annual logarithmic rate between 1991 and 2006. If that trend were to continue, by 2020 China would be consuming 20 million barrels per day (about as much as the U.S. is currently consuming), and by 2030 that would have doubled again to 40 mb/d (see Figure 10).

Are such extrapolated demand figures plausible? Despite its remarkable growth already, China still has a long way to go before we might expect the income elasticity of oil demand to fall significantly. During 2006, China used about 2 barrels of oil per person. For

³ World consumption numbers were taken from Energy Information Administration, "World Petroleum Consumption, Most Recent Annual Estimates, 1980-2007".

comparison, Mexico used 6.6– Chinese oil consumption could triple and they’d still be using less per person than Mexico is today. The U.S. used almost 25 barrels per person. There were 3.3 passenger vehicles per 100 Chinese residents in 2006, compared with 77 in the United States.⁴

But is the world capable of producing oil in such volumes? We turn to this question in the next section.

5 Petroleum supply.

Figure 11 plots global oil production levels over the last quarter century. Global production has stagnated over the last three years. Given the strong demand growth from China and the Middle East, that required a big increase in price to restore equilibrium. The key question is why supply failed to increase.

5.1 The role of OPEC.

Although there was once a time in which a few oil companies played a big role in world oil markets, that era is long past. ExxonMobil, the world’s largest private oil company, produced 2.6 mb/d of oil in 2007, which is only 3.1% of the world total. The combined market share of the 5 biggest private companies is less than 12%. In the modern era, it is sovereign countries rather than private companies who would be calling the shots.

The Organization of Petroleum Exporting Countries includes 12 of the important oil pro-

⁴ U.S. statistics are from the Bureau of Transportation Statistics, Chinese kindly provided me by Maximilian Auffhammer. For more details see Auffhammer and Carson (2008) and Congressional Budget Office (2006).

ducing countries, two of which (Angola and Iraq) are currently not participating in OPEC's production agreements. The OPEC-10⁵ produced 36.7% of total world liquids production in 2007, of which Saudi Arabia alone accounted for 12.1%. The modest increase in production outside of these 10 countries during 2006 and 2007 was just offset by decreases within the OPEC-10; (see Figure 12).

If OPEC were operating as an effective cartel, in the absence of a Hotelling scarcity rent it would try to set the marginal revenue for the group equal to the marginal cost. The marginal revenue for the group associated with producing one more barrel of oil would be calculated as the price of that barrel minus the revenue that OPEC would lose if to sell that marginal barrel it had to lower the price to all its previous buyers. By contrast, the marginal revenue for an individual OPEC member would be the price minus the lost revenue to the member. Because any one member is a small fraction of the entire group, the marginal revenue for an individual member is always a bigger number than the marginal revenue for the group as a whole. As a consequence, if group marginal revenue is set equal to marginal cost, individual marginal revenue is greater than marginal cost, meaning there would always be an incentive for members to try to "cheat" on the cartel's production decisions, producing a little more for themselves than the group agreed. An effective cartel requires some mechanism to deter such behavior.

Figure 13 plots the quotas and actual production levels for the 5 biggest OPEC producers.

⁵ The OPEC-10 are Algeria, Indonesia, Iran, Kuwait, Libya, Nigeria, Qatar, Saudi Arabia, United Arab Emirates, and Venezuela. One of these (Indonesia) has actually become a net oil importer in recent years. Data are from EIA, "World Production of Crude Oil, NGPL, and Other Liquids, and Refinery Processing Gain".

There is only a loose correspondence. Kuwait has always produced more than its quota and Venezuela has always produced less. Saudi Arabia was well above its quota during 2004-2005 and Iran well below its during 2006. In fact, the “quotas” and measured production levels are themselves fairly vague. The Energy Information Administration, International Energy Agency, and private organizations such as Platts all have different estimates of what the actual production numbers are. In the description of quotas that is posted on the OPEC website, the quotas for 1996-2006 are all described in terms of actual production levels for each country, whereas the new policies implemented November 2006 are described in terms of changes from previous quotas rather than new target levels, apparently reflecting a tacit acknowledgement that deviations of actual production figures from earlier quotas were quite large, and making the new guidelines— such as a 176,000 b/d cut for Iran from some unspecified previous level— having even less clarity in terms of what was required than those that had been in place earlier. For the current guidelines implemented November 2007, OPEC seems to have given up even on this, and has announced a simple aggregate target of 27.253 mb/d target for the OPEC-10 without specifying who is supposed to produce what. The only publicly available numbers I have seen on how this 27.253 figure is supposedly allocated among the OPEC members comes from an anonymous website calling itself “Saudi Oil Production,” whose numbers are used for the final values in Figure 13. It is clear that for these numbers in particular, it is the quotas that have moved to match the production rather than the other way around.

It is hard to find any clear monitoring or enforcement mechanism for implementing

OPEC's announcements, which instead seem to have more of the character of each country deciding what it wants to do anyway and the organization then making an announcement of the collection of those individual decisions. Under such a view, the announcements of the group then serve mainly political interests, giving countries like Iran and Venezuela an opportunity to appear to their domestic constituencies to be fighting for higher oil prices, and giving countries like Saudi Arabia an ability to spread the blame for its decisions over a broader group.

Since Saudi Arabia alone accounts for a third of the production from the OPEC-10, one might alternatively consider the hypothesis that the Kingdom makes a calculation based on its unilateral monopoly power, with the rest of the world producing on a more competitive basis. The condition for Saudi marginal revenue to equal its marginal cost can be written⁶

$$P \left(1 - \frac{1}{\varepsilon_S} \right) = M_S$$

where P denotes the price of oil, ε_S the price-elasticity of demand for Saudi oil, and M_S the Saudi's marginal cost of production. Note further that if the Saudis control a share κ_S of the global market and the global demand elasticity is ε_G , then

$$\varepsilon_S = \varepsilon_G / \kappa_S$$

since a 1% increase in Saudi production would only be a κ_S percent increase in global production. Hence in the absence of a scarcity rent the Saudi's objective would be to set a

⁶ Note marginal revenue can be written

$$MR = \frac{\partial(P(Q) \cdot Q)}{\partial Q} = P \left(1 + \frac{Q}{P} \frac{\partial P}{\partial Q} \right) = P \left(1 - \frac{1}{\varepsilon} \right).$$

markup of price over marginal production cost of

$$\frac{P}{M_S} = \frac{1}{1 - \frac{\kappa_S}{\varepsilon_G}}.$$

Suppose we used the price-elasticity estimate of 0.26 derived in (7) for illustration. With a Saudi global share of $\kappa_S = 0.12$, we would expect a markup of

$$\frac{1}{1 - \frac{\kappa_S}{\varepsilon_G}} = \frac{1}{1 - \frac{0.12}{0.26}} = 1.86. \quad (8)$$

If, as in Horn (2004), we assumed a marginal production cost of \$15/barrel, that would imply an oil price of \$28. Note further that the 0.26 estimate was a short-run elasticity. It is the long-run elasticity that should be used in a formula like this one, in which case the predicted price would be even lower. The above calculation also assumed zero supply elasticity from sources outside of Saudi Arabia; adding these would again give us a smaller markup than calculated in (8).

On the other hand, we noted above that oil demand has become much less price elastic over time, in which case the predicted price would increase. Indeed, as the elasticity ε_G in (8) approaches 0.12, the predicted price goes to infinity, and the Hughes, Knittel and Sperling (2008) recent estimates are even smaller than 0.12. It certainly is the case that Saudi production decreased in 2006 and 2007 (see the top panel of Figure 13), and this has undoubtedly made a contribution to the recent price increase. However, if this is indeed the explanation for the recent run-up in prices, it raises the question of why no one elsewhere in the world is able to produce oil for under \$100 a barrel to undercut the hypothesized Saudi monopoly price. We therefore turn in the next section to an investigation of global prospects for increasing oil production.

5.2 Near-term supply prospects.

There are enormous lead times between the initial discovery of a new oil reservoir and the time at which the new oil is actually being delivered to a refinery to use. These lags mean that, in the absence of significant excess production capacity, the short-run price elasticity of oil supply is also very low, another factor contributing to the potential price implications of supply disruptions. The thin line in Figure 11 plots a linear time trend fit to global oil demand over 1983-2003. Oil use actually grew much faster than this trend during 2001-2005, and in fact remains above the trend as of the time of this writing. One might then conjecture that the strength of global demand caught producers by surprise, and that some time would be required for the investments necessary to catch up.

This was the view offered by oil analysts at Cambridge Energy Research Associates as an explanation for the high prices in 2005. They also concluded that supply increases would be substantial and forthcoming very quickly. CERA Chairman Daniel Yergin wrote in the *Washington Post* on July 31, 2005:

The oil industry is governed by a “law of long lead times.” Much of the new capacity that will become available between now and 2010 is under development. Many of the projects that embody this new capacity were approved in the 2001-03 period, based on price expectations much lower than current prices....

Our new, field-by-field analysis of production capacity [concludes that there] will be a large, unprecedented buildup of oil supply in the next few years. Between 2004 and 2010, capacity to produce oil (not actual production) could grow by

16 million barrels a day— from 85 million barrels per day to 101 million barrels a day— a 20 percent increase. Such growth over the next few years would relieve the current pressure on supply and demand.

CERA's prediction, offered when oil prices were nibbling at \$60/barrel, received a lot of attention from the press at the time. But actual global production increased only 1.5 mb/d between 2004 and 2005 and did not increase at all beyond the 2005 level throughout 2006 and 2007. It is instructive to inquire how the prediction went wrong. Figure 14 provides details for the 11 areas in which production gains were anticipated. With the exception of Iraq, every country in this group did indeed experience a net gain in production thanks to new fields. However, in every case the gain was significantly less than anticipated. In some countries, such as Iraq and Nigeria, the shortfall can be attributed to political instability that disrupted the flow of oil. In others, projects developed more slowly than anticipated. A third factor is that even with the addition of new fields, production from existing mature fields can go into decline, as we discuss in the next section.

5.3 The challenge of depletion.

There are a variety of measures that can be taken to increase production from an existing field or increase the percentage of original oil in a given reservoir that is ultimately uncovered. These options include drilling additional wells at alternative locations and pumping in water or carbon dioxide to maintain pressure. New wells typically cause the production profile of a given field to increase in the initial phase of development. However, as more oil is removed, less remains in the original deposit and it becomes increasingly difficult to continue to extract

oil at the same rate. In a given field, one inevitably observes a profile of initial increasing production flow rates followed by eventual decline. To keep total production increasing, it is necessary to find new fields continuously. Historically this has been achieved by moving to new geographical areas.

The top panel of Figure 15 displays this pattern for the rich oil producing areas in Texas, from which production has been in steady decline since 1972. Production from the Prudhoe Bay supergiant field in Alaska (middle panel) has declined on average by 8.5% per year since 1988. Overall, U.S. production today is about half of what it was in 1971.

Figure 16 documents that this fall in U.S. production has not been for a lack of effort. In the 1980s, the U.S. was producing less oil using 3 times as many wells as in the 1970s. We have also made a steady transition to relying on offshore oil and deeper wells.

Lynch (2002) noted that some analysts inferred erroneously from the U.S. production decline in the early 1970s that an analogous development on a global scale was soon to follow. In practice, such forecasts proved to be inaccurate, as huge production gains outside the United States were achieved over the next 35 years. Notwithstanding, there is one sense in which those like Akins (1973) who were alarmed by the trends in Figure 15 proved to be correct. The decline in U.S. production has been irreversible, and the huge growth in U.S. oil imports and transfer of wealth that resulted from declining U.S. production rates profoundly and permanently changed the world geopolitically, and indeed I would argue is the major reason that the price path of oil in Figure 1 begins to depart dramatically from its previous behavior after the U.S. passed its peak in production in 1970.

A number of the producing areas outside the U.S. are also unambiguously now in decline. As shown in Figure 17, production from the United Kingdom and Norway has declined by 7% per year since 2002. Mexico's Cantarell complex, second only to Saudi Arabia's Ghawar in terms of its contribution to recent production levels, is dropping precipitously. China, like the U.S., was once a net petroleum exporter. Production from its three largest fields is now in decline (Kambara and Howe, 2007), though new Chinese fields have so far been sufficient to allow total Chinese production to increase modestly despite the maturity of its major producing areas. Again, it is hard to deny that declining production rates from the mature Chinese fields has been a factor influencing the recent course of world oil prices.

Saudi production, shown in the top panel of Figure 18, has historically exhibited considerable variation, as the Kingdom dropped production in times of slack demand to keep prices from falling, and raised production to moderate the price increases occasioned by historical disruptions from Iran and Iraq. This behavior on the part of Saudi Arabia helped to make the global supply curve considerably flatter than it otherwise would have been during the era when the Kingdom had lots of excess capacity. The most recent drop in Saudi production since 2005, however, appears to represent a different regime, since these began at a time of rapidly rising prices and stagnating production elsewhere. At a minimum, this is a radically different concept of "price stabilization" than seems reflected in earlier Saudi behavior, and may indicate that the Saudis' excess production capacity has been eroded. The production declines coincided with a doubling in the number of their active oil rigs, leaving some to speculate that the magnificent Ghawar oil field had begun to decline. The necessary data

to confirm or refute that conjecture are not publicly available. But it seems likely that if production from Ghawar has indeed already started to decline, the peak in global production cannot be far off.

At any given point in history, some of the world's producing fields are well into decline, some are at plateau production, and others are on the way up. It is not clear what "average" or "typical" decline rate would be appropriate to apply to aggregate global production, but a plausible ballpark number might be 4%. That means that in the absence of new projects, global production would decline by 3.4 mb/d each year. To put it another way, a new producing area equivalent to current annual production from Iran (OPEC's second biggest producer) needs to be brought on line every year just to keep global production from falling.

Despite these discouraging observations, an update of calculations like those performed in 2005 by CERA would leave one still quite optimistic about near-term oil supplies. An open-source web database⁷ tabulates a total of 6.9 mb/d in new gross production capacity from new projects that are scheduled to begin producing in 2008. Projects in Saudi Arabia, Russia, and Mexico account for about a third of this gross increase. Data currently available for the first two months of 2008 show actual production in Saudi Arabia down 350,000 b/d from its average 2005 value and Mexican production down 400,000 b/d from 2005. Russian production is down 100,000 b/d from its average level in the second half of 2007.

Although declining production from mature fields and delays in ramping the new fields up to full production will doubtless eat up a fair bit of the 6.9 mb/d new gross production

⁷ http://en.wikipedia.org/wiki/Oil_Megaprojects/2008.

capacity, it seems there is a lot left over. In the absence of significant new geopolitical disruptions to petroleum supply, some might anticipate an end to the recent plateau in global production, and significant net gains in supply for 2008.

However, it would not take too many years of 7% demand growth from China and other economies to absorb a good part of even the most optimistic projections of what is likely over the near term.

6 Conclusions.

In this paper we have reviewed a number of theories as to what has produced the current high price of oil, including commodity price speculation, strong world demand, time delays or geological limitations on increasing production, OPEC monopoly pricing, and an increasingly important contribution of the scarcity rent. Rather than think of these as competing hypotheses, one possibility is that there is an element of truth to all of them.

Unquestionably the two key features in any account are a decrease in the price elasticity of demand and the strong growth in demand from China, the Middle East, and other newly industrialized economies. These twin facts explain the initial strong pressure on prices that may have triggered commodity speculation in the first place. Speculation could have edged producers like Saudi Arabia into the discovery that small production declines could increase current revenues and may be in their long run interests as well. And the strong demand may have moved us into a regime in which scarcity rents, while negligible in 1997, are now an important permanent factor in the price of petroleum.

Notwithstanding, different emphases among these explanations would produce profoundly different predictions as to what will happen next. If speculation and short-run price inelasticity are the key driving factors, we would expect shortly to see potentially dramatic moves downward in price. The scarcity rent, by contrast, is expected to increase, not decrease, over time.

The evidence reviewed in Section 2 highlights the hazards of offering a prediction about what happens next. But the algebra of compound growth suggests that if demand continues to grow in China and other countries at its current rate, the date at which the scarcity rent will start to make an important contribution to the price, if not here already, cannot be far away.

References

Abosedraa, Salah, and Hamid Baghestani. 2004. "On the Predictive Accuracy of Crude Oil Futures Prices," *Energy Policy* 32, pp. 1389–1393.

Akins, James E. 1973. "The Oil Crisis: This Time the Wolf Is Here," *Foreign Affairs*, pp. 462-490.

Auffhammer, Maximilian and Richard T. Carson. 2008. "Forecasting the Path of China's CO2 Emissions Using Province Level Information," *Journal of Environmental Economics and Management* 55(3), pp. 229-47.

Bopp, Anthony E. and George M. Lady. 1991. "A Comparison of Petroleum Futures versus Spot Prices as Predictors of Prices in the Future," *Energy Economics* 13(4), pp. 274-282.

Chinn, Menzie D., Michael LeBlanc, and Olivier Coibion. 2005. "The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline," NBER Working Paper. No. 11033.

Congressional Budget Office. 2006. "China's Growing Demand for Oil and Its Impact on U.S. Petroleum Markets."

Dahl, Carol and Thomas Sterner. 1991. "Analysing Gasoline Demand Elasticities: A Survey," *Energy Economics* 13, pp. 203-210.

DeLong, J. Bradford, Andrei Shleifer, Lawrence H. Summers and Robert J. Waldmann. 1990. "Positive Feedback Investment Strategies and Destabilizing Rational Speculation," *Journal of Finance* 45(2), pp. 379-395

- Dufour, Alfonso and Robert F Engle. 2000. "Time and the Price Impact of a Trade," *Journal of Finance* 55(6), pp. 2467–2498
- Edelstein, Paul, and Lutz Kilian. 2007. "Retail Energy Prices and Consumer Expenditures", working paper, University of Michigan.
- Gately, Dermot, and Hillard G. Huntington. 2002. "The Asymmetric Effects of Changes in Price and Income on Energy and Oil Demand," *Energy Journal* 23(1), pp. 19-55.
- Hamilton, James D. 1994. *Time Series Analysis*. Princeton, NJ: Princeton University Press.
- Hamilton, James D. 2003. "What Is an Oil Shock?" *Journal of Econometrics* 113(2), pp. 363-398.
- Heston, Alan, Robert Summers and Bettina Aten. 2006. Penn World Table Version 6.2, Center for International Comparisons of Production, Income and Prices at the University of Pennsylvania, September 2006.
- Horn, Manfred. 2004. "OPEC's Optimal Crude Oil Price," *Energy Policy* 32(2), pp. 269-280
- Hotelling, Harold. 1931. "The Economics of Exhaustible Resources," *Journal of Political Economy* 39(2), pp. 137-75.
- Hughes, Jonathan E., Christopher R. Knittel and Daniel Sperling. 2008. "Evidence of a Shift in the Short-Run Price Elasticity of Gasoline Demand." *Energy Journal* 29(1), pp. 93-114.
- Jovanovic, Boyan. 2007. "Bubbles in Prices of Exhaustible Resources", working paper,

New York University.

Kambara, Tatsu, and Christopher Howe. 2007. *China and the Global Energy Crisis*, Edward Elgar.

Krautkraemer, Jeffrey A. 1998. "Nonrenewable Resource Scarcity," *Journal of Economic Literature* 36(4), pp. 2065-2107.

Kwiatowski, Denis, Peter C. B. Phillips, Peter Schmidt, and Yongcheol Shin. 1992. "Testing the Null Hypothesis of tationarity Against the Alternative of a Unit Root," *Journal of Econometrics* 54, pp. 159-178.

Lynch, Michael C. 2002. "Forecasting oil supply: theory and practice," *Quarterly Review of Economics and Finance* 42, pp. 373–389.

Masters, Michael W. 2008. Testimony before the Committee on Homeland Security and Governmental Affairs, United States Senate, May 20.

Roya, Joyashree, Alan H. Sanstadb, Jayant A. Sathayeb, and Raman Khaddaria. 2006. "Substitution and Price Elasticity Estimates Using Inter-country Pooled Data in a Translog Cost Model," *Energy Economics* 28(5-6), pp. 706-719.


Table 1. *P*-values for tests of null hypothesis that indicated variables are of no use in predicting quarterly real oil price change.

variable	1 lag	4 lags	8 lags
real oil price change	0.69	0.88	0.62
U.S. nominal tbill rate	0.53	0.61	0.83
U.S. real GDP growth rate	0.24	0.48	0.49

Table 2. Ninety-five percent lower and upper bounds on forecast for inflation-adjusted price of oil assuming a Gaussian random walk.

date	forecast	lower	upper
2008:Q1	115		
2008:Q2	115	85	156
2008:Q3	115	75	177
2008:Q4	115	68	195
2009:Q1	115	62	212
2010:Q1	115	48	273
2011:Q1	115	40	332
2012:Q1	115	34	391

Figure 1. Oil price in 2008 dollars per barrel.



Notes: Calculated as monthly average price (in dollars per barrel) of West Texas Intermediate times the April, 2008 value for the seasonally adjusted consumer price index and divided by the CPI as of the indicated month.

Figure 2. Quarterly percent change in real oil price.





Figure 3. Price of crude oil contract maturing December of indicated year.



Notes: solid line: contracts traded on August 21, 2007. Dashed line: contracts traded on October 4, 2007.

Figure 4. Disentangling supply and demand




Figure 5. Monthly oil production for Iran, Iraq, and Kuwait, in thousand barrels per day.






Figure 6. Changes in U.S. real GDP and oil consumption, 1949-2006.




Notes: Horizontal axis: cumulative change in natural logarithm of U.S. real GDP between 1949 and the year for which a given data point is plotted, from Bureau of Economic Analysis Table 1.1.6. Vertical axis: cumulative change in natural logarithm of total petroleum products supplied to U.S. market between 1949 and the year for which a given data point is plotted, from Energy Information Administration, "Petroleum Overview, 1949-2006", Table 5.1.

Figure 7. Share of U.S. crude oil expenditures as a fraction of GDP.



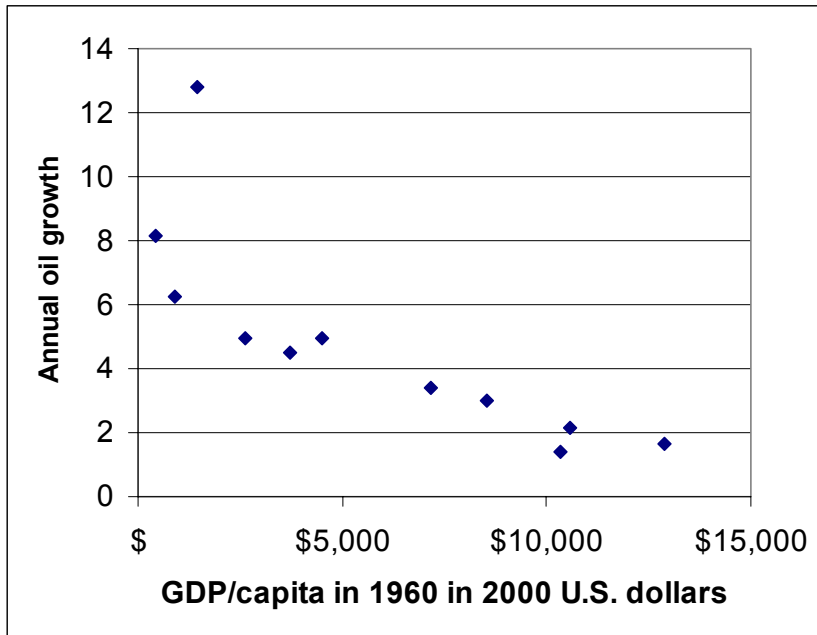
Notes: Calculated as the number of barrels of oil consumed (from EIA, World Petroleum Consumption) times the average price of West Texas Intermediate (from the FRED database of the Federal Reserve Bank of St. Louis) divided by nominal GDP.

Figure 8. Smoothed annual growth rate of petroleum consumption for France, Germany, and Japan, 1960-2002.




Notes: raw data are from Energy Information Administration, “World Petroleum Consumption, 1960-2002”, Table 11.10. Data plotted are exponentially smoothed weighted averages of 100 times the annual logarithmic growth rates with decay weight of 0.90.

Figure 9. GDP per capita and growth in petroleum demand.




Notes: Horizontal axis: GDP per person in 1960, measured in 2000 U.S. dollars, from Heston, Summers, and Aten (2006). Vertical axis: average annual logarithmic growth rate in petroleum demand between 1960 and 2002. Countries included (in order of decreasing average petroleum demand growth) are Korea, China, India, Japan, Brazil, Mexico, Italy, France, Canada, US, and UK.

Figure 10. Historical Chinese oil consumption and projection of trend.



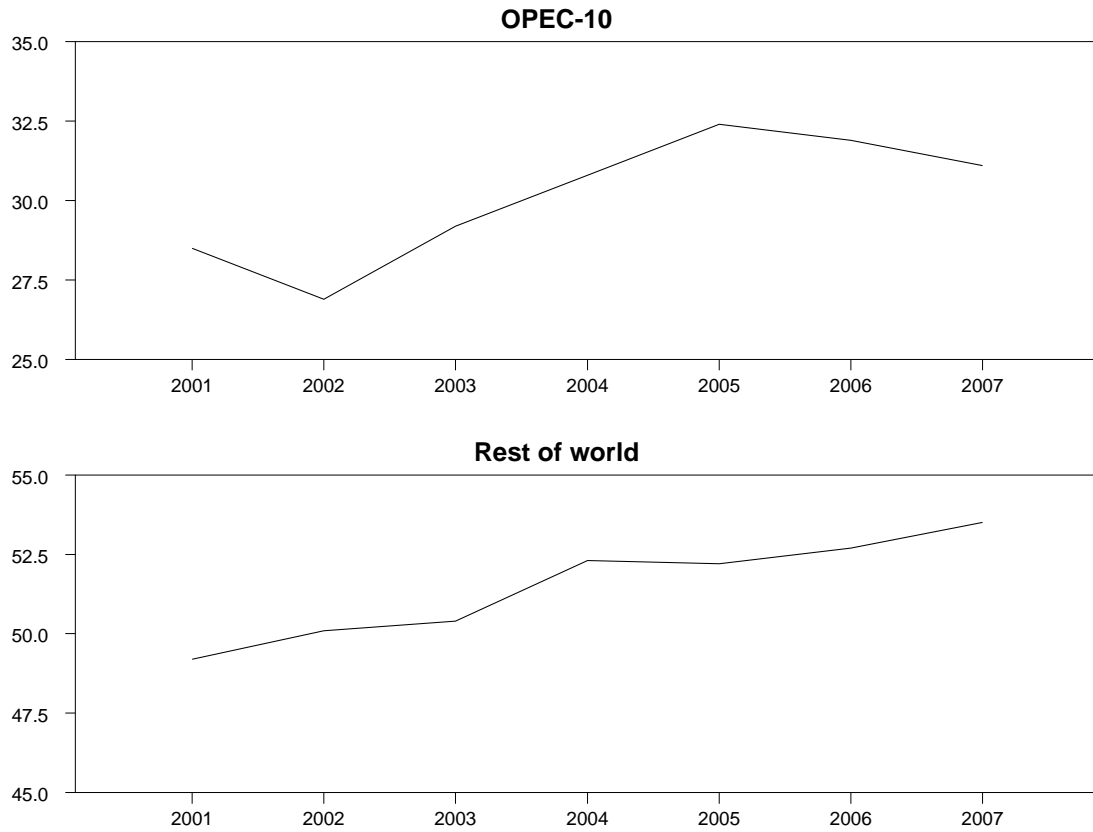
Notes: 1991-2006: Chinese oil consumption in millions of barrels per day. 2007-2030: extrapolation of 7.2% compounded growth.

Figure 11. Global production of crude petroleum.




Notes: Bold line: From EIA, “World Production of Crude Oil, NGPL, and Other Liquids, and Refinery Processing Gain”, in million barrels per day. Thin line: regression estimate of time trend fit for 1983-2003 data.

Figure 12. Production from the OPEC-10 and rest of the world.




Notes: Data from EIA, "World Production of Crude Oil, NGPL, and Other Liquids, and Refinery Processing Gain", in million barrels per day.

Figure 13. Quotas and actual production levels for 5 most important OPEC members.




Notes: Production levels from EIA Table 1.2, “OPEC Crude Oil Production (Excluding Lease Condensate)”, in thousand barrels per day. Quotas taken from OPEC website (<http://www.opec.org/home/Production/productionLevels.pdf>) with specific country allocations for quotas adopted Nov. 1, 2007 taken from <http://saudioilproduction.blogspot.com/2007/09/new-opec-quotas.html>.

Figure 14. Predicted and actual production gains, 2004-2006.



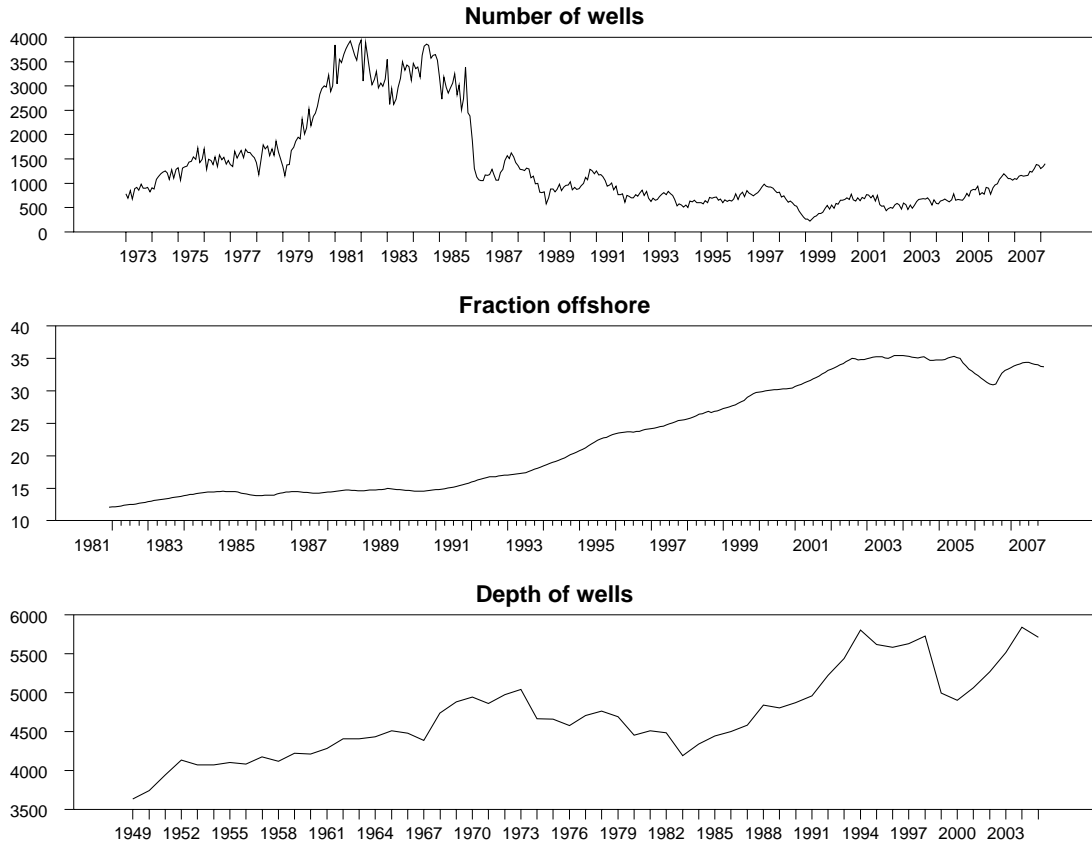
Notes: First bar: change in liquids capacity over 2004-2006 as predicted by CERA, in millions of barrels per day. Second bar: actual change according to EIA Tables 11abc and 22.

Figure 15. Production levels for state of Texas, Alaska’s Prudhoe Bay, and entire U.S.




Notes: All data reported in millions of barrels per day. Top panel: annual production from the state of Texas, 1935-2006, from Railroad Commission of Texas (<http://www.rrc.state.tx.us/divisions/og/statistics/production/ogisopwc.html>). Middle panel: annual production from Prudhoe Bay in Alaska, 1977-2005, from Alaska Department of Revenue. Bottom panel: moving average of preceding 12 months of monthly production figures for the United States, December 1920 to February 2008, from EIA, “Crude Oil Production.”

Figure 16. U.S. wells drilled, fraction of offshore production, and average well depth.




Top panel: Monthly count of the number of U.S. crude oil exploratory and developmental wells drilled, January 1973 to March 2008, from EIA, “Crude Oil and Natural Gas Exploratory and Development Wells.” Middle panel: percent of U.S. total crude oil production coming from federal and state offshore production, with both counts based on 12-month moving average of monthly production figures, December 1981 to December 2007, from EIA, “Crude Oil Production.” Bottom panel: Annual U.S. average depth of crude oil, natural gas, and dry exploratory and developmental wells drilled (feet per well), 1949 to 2005, from EIA, “Average Depth of Crude Oil and Natural Gas Wells.”

Figure 17. Oil production from the North Sea, Mexico's Cantarell, and China's Daqing.



Notes: all figures in thousand barrels per day. Top panel: sum of U.K. and Norway crude oil production, monthly moving average of preceding 12 months, December 1973 to June 2007, from EIA, Table 11.1b. Middle panel: annual production from Cantarell complex in Mexico. Data for 1996 to 2006 from Pemex 2007 Statistical Yearbook. Data for 2007 from Green Car Congress (<http://www.greencarcongress.com/2008/01/mexicos-cantare.html>). Bottom panel: annual production from Daqing field in China, 1960-2005, data from Kambara and Howe (2007), with missing observations linearly interpolated.

Figure 18. Saudi Arabian production and oil rigs.



Top panel: monthly production in thousand barrels per day, January 1973 to January 2008, from EIA, Table 11.1a. Bottom panel: monthly count of number of land and offshore oil rigs in Saudi Arabia, January 1982 to April 2008, from Baker Hughes (http://investor.shareholder.com/bhi/rig_counts/rc_index.cfm).