The QUARTERLY REVIEW Of ECONOMICS And FINANCE

NORTH-HOLLAND The Quarterly Review of Economics and Finance 42 (2002) 373–389

Forecasting oil supply: theory and practice

Michael C. Lynch*[,1]

DRI-WEFA, 24 Hartwell Avenue, Lexington, MA 02421, USA

Received 21 August 2001; accepted 15 January 2002

1. Introduction

The history of oil market forecasting has been a sorry one, all the more so for the inability of many forecasters to perceive—let alone correct—their errors. Partly this reflects the role of self-interest: some forecasts are wishful thinking rather than serious analysis. But also, the behavior to be modeled is much more complex than many realize, and most researchers lack the ability or resources to analyze it except superficially. Obviously, for an outside observer to judge various forecasts is thus quite challenging.

One low-cost approach is to review the track record. This can be valuable by identifying whether there are, indeed, bias and recurring errors, although a more precise examination of methodologies is necessary to reveal the nature of those errors. Further, an understanding of the nature of petroleum supply at the microeconomic level is necessary to explain the difficulty in forecasting it at the macroeconomic (in the sense of national or regional in scope) level. This article will describe the performance of supply forecasting over the past two decades, discuss the methodological errors in the geophysical models and the difficulties in creating a valid microeconomic model. The inconsistency of results with theory, and the reasons for it, will make up the final section. First, a discussion of oil price forecasting will be provided to demonstrate the value of reviewing forecasts.

2. Price forecasting: the 3% rule

For the forecasting of oil prices, recurring errors are relatively easy to identify. Forecasts have been persistently too high since the late 1970s, such that even a moderate forecast such

* *E-mail address:* michael.lynch@dri-wefa.com (M.C. Lynch).

PII: S1062-9769(02)00126-6

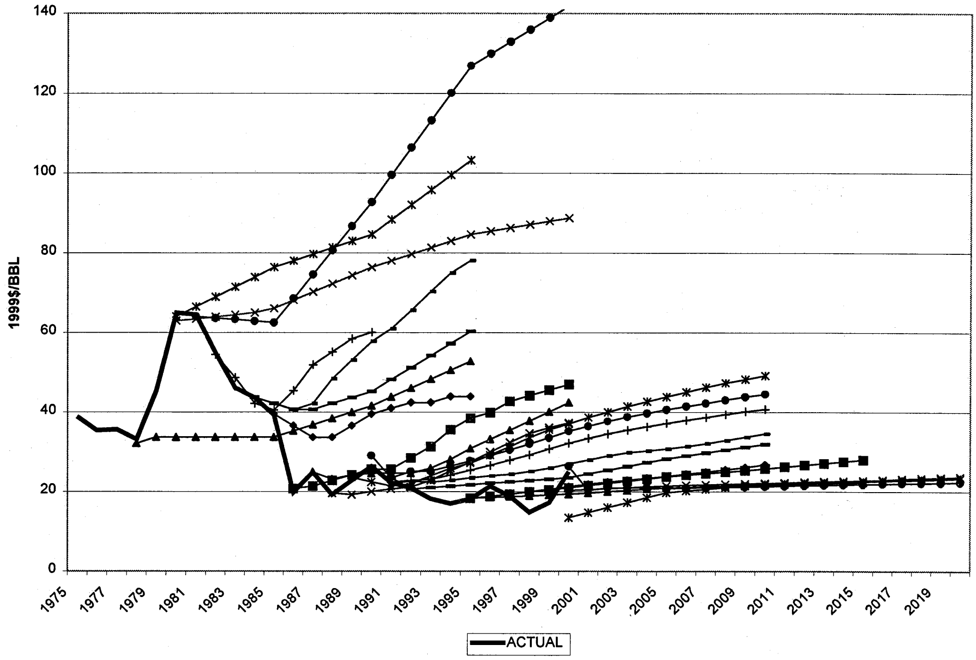

Fig. 1. The evolution of DOE oil price forecasts.

as Lynch (1989), predicting Arab Light would decline to the $13–18.5/bbl range by 2000 (all figures in 1998 dollars), was considered “heretical” (*Petroleum Economist* 9/89, p. 270). But although the mistakenly high price forecasts of the 1980s are often thought to be extrapolation of the 1970s price increases, in fact, such was not the case. Throughout most of the history of the oil industry, prices were mean-reverting, around approximately $14/bbl, under the influence of various price-stabilizing efforts, led first by Standard Oil, then the Texas Railroad Commission and the Seven Sisters and now OPEC. The increases in the 1970s were disruption-related spikes, as opposed to the gradually rising prices predicted by most forecasters—a trend never seen in mineral prices.

The flaw was partly theoretical, with many misinterpreting Hotelling (1931) to demonstrate that net mineral prices (gross price minus cost) should rise at the rate of interest. (See Adelman (1990) and Watkins (1992) for refutation of this interpretation.) For oil, whose cost was effectively zero for marginal Middle Eastern supplies, this meant that the gross (or market) price would rise at the rate of interest. The 1979 decision by OPEC’s Long-term Strategy Committee to propose a long-term price path rising at the rate of OECD growth bolstered expectations of 2–4%/year real price growth. As a result, for many years, nearly every oil price forecast called for such a trend; as the forecasts proved erroneous, the trend was retained, but applied to the new, lower initial point. Fig. 1 shows the US Department of Energy’s price forecast as it evolved over time, which clearly demonstrates this behavior.

The combination of these theoretical arguments with the oil price shock of 1979 gave great credence to these rising price forecasts, and it has proven difficult to convince casual observers that although prices might rise, this is neither inevitable nor preordained by either economic law or geology. Indeed, even computer modelers were slow to learn this lesson. The difference in the long-term price trends seen in the computer models of the world oil market reviewed in EMF6 (1982) and EMF11 (1991) is marginal; the former averages an increase of 3.6%/year, the latter, 4.6%/year. (Actual real prices declined by 2–3% per year until the spike in 2000.) Recently, an assumption of flat prices has become common in long-term forecasts, although the price spike of 2000 might change that.

3. Historical supply forecasting

Although many who write about oil include (now humorous) historical quotes about the scarcity of oil resources (regional or global), attitudes about oil supply actually appear to have varied over time. There was a clear paradigm shift in the late 1970s, when expert opinion shifted as part of a trend towards pessimism about resource availability. Before 1979, despite the publication of Malthusian tracts like *Limits to Growth* (Meadows, Meadows, Randers & Behrens, 1972) it was not uncommon to find optimistic forecasts. OECD (1974) predicted that the 1974 oil price increases would lead to US oil production of 11.5–18.8 mb/day by 1980 (vs. actual of 8.5), for example. However, by 1977, new studies began to argue that resources were limited, discoveries were not keeping pace with production, and that oil production outside the Middle East had peaked. (See Lynch (1996) for a review.) The Iranian Oil Crisis only served to confirm the presence of scarcity for many. As the CIA put it:

> "The gas lines and rapid increases in oil prices during the first half of 1979 are but symptoms of the underlying oil supply problem—that is, the world can no longer count on increases in oil production to meet its energy needs." CIA (1979), p. iii.

Not only did rapidly rising non-OPEC supply in the early 1980s fail to curb pessimism about potential production, but even the 1986 oil price collapse had little impact, partly because US production dropped sharply from 1986, as did UK production in 1988, appearing to confirm the pessimists' arguments. Fig. 2, showing a group of 1989 forecasts, is just one of many that demonstrates past oil supply forecasting failures. Note that all are pessimistic (evidence of bias), and the only forecast that doesn't foresee a decline is Lynch (1989).

4. The nature of the errors

Examining one particular set of forecasts over time can yield an indication of the nature of the recurring error. In Fig. 3, the US Department of Energy's forecast for non-OPEC Third World crude oil production is shown as it has changed over the past twenty years. Given that this region (actually many regions) is one which has not been heavily exploited to date (drilling density in sedimentary basins is about 2% that of the US), it is hard to find a rationale for the repeated pessimism other than Malthusian bias, as Lynch (1987) argued was true of most such forecasts for those countries. And it can be seen that not only have these

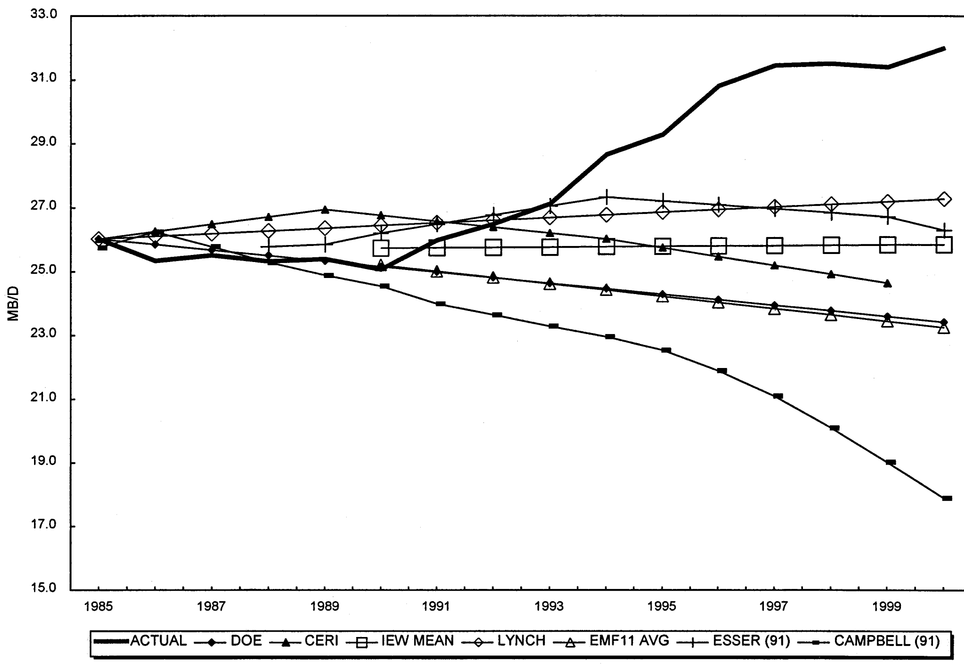

Fig. 2. Major non-OPEC supply forecasts in 1989/1991.

forecasts been too low, but there was a consistent tendency to predict a peak within a few years, and a decline. As this peak was surpassed, the correction was not to remove the peak, but to raise it and move it out, always a few years into the future. This resembles the price forecasting record, where throughout the 1980s, prices were predicted to rise from whatever the current level was, even though they were continually falling. (The change in the forecast in recent years reflects a change in the personnel involved.)

In earlier works,[2] I have reviewed a wide variety of non-OPEC oil supply forecasts and found most displayed similar errors, with a near-term peak in production—which subsequently proved to be incorrect—regardless of expected price or region being forecast. In fact, no major region has yet shown signs of such behavior, outside of the US. It is only now, after two decades of such forecasts, that a few areas like the North Sea, Egypt and Argentina, appear to be reaching their peaks.

This does not however prove that the models are inherently mis-specified. The specific nature of the errors that generated these curves will be discussed below.

5. Geophysical models

The Hubbert approach deserves special attention because it has received enormous attention lately (see Campbell & Laherrere, 1998; Edwards, 1997; Hatfield, 1997). The arguments made are fairly common: you can't produce oil that doesn't exist, regardless of

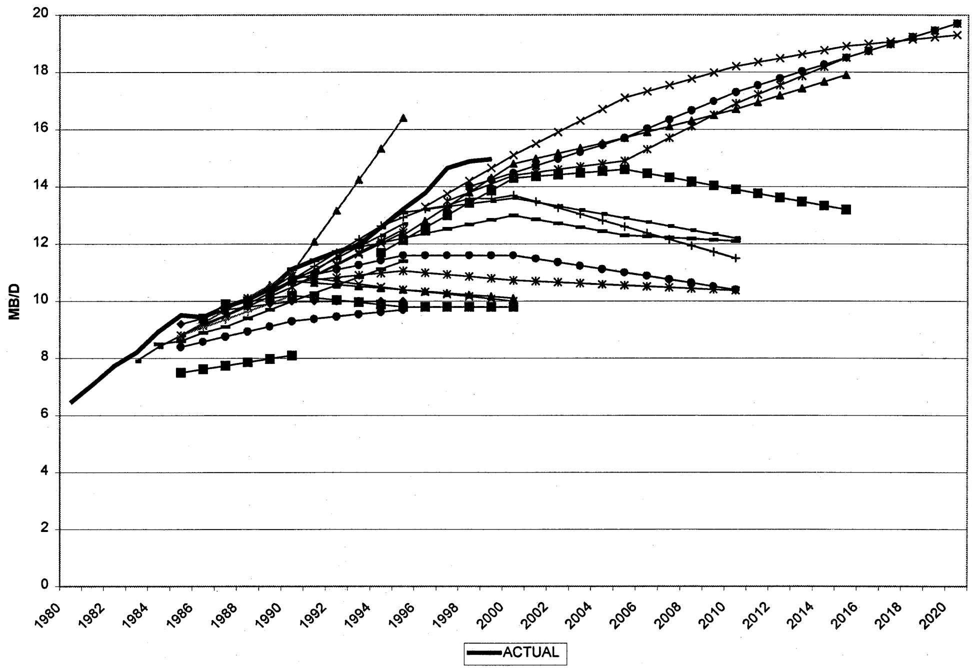

Fig. 3. DOE forecasts of non-OPEC third world oil supply over time.

price; and depletion—which is determined only by geology and chemistry—drives production curves. There is no room for economics in this view, which, given the poor record of oil price forecasting, might seem gratifying. Instead, these models are driven entirely by geophysical factors, and most are based on the work of M. King Hubbert (Hubbert, 1956). They rely on estimates of ultimately recoverable resources (URR) and historical production, fitted to a bell or parabolic curve, to generate long-term production forecasts. If production is past its peak, then the curve can be estimated in its entirety, and the area under it represents URR. Alternatively, given an estimate of URR and a production history, a production peak can be estimated, since it should appear at approximately the midpoint of depletion.

The popularity of the approach stems partly from the fact that Hubbert's 1956 prediction of lower-48 oil production was extremely accurate, even to this day. However, this is misleading, as his other three predictions were highly inaccurate. His forecast of US gas production in 2000 was 65% too low and his world oil production forecast for 2000 was 50% too low. Even production in Texas is now about twice the amount he forecast. Indeed, given his estimate of URR in Texas, production should have ceased recently as the resources were exhausted. Since Texas was a mature, heavily studied province even in 1956, this error speaks to the fallibility of the method.

Other Hubbert models exhibit the same flaw. One group at the US Department of Energy produced a series of Hubbert-style production profiles in the early 1980s (for example, US Department of Energy, 1983). For non-OPEC countries, they produced a prediction for Southeast Asia that has proven very accurate to date, but for non-OPEC South America and

Egypt, their forecasts were much too low. Similarly, Root (1991) and Masters et al. (1990) also produced forecasts with near-term peaks for regions, most if not all of which are clearly too pessimistic.

The most egregious errors have come from C.J. Campbell, who has repeatedly predicted a near-term peak for the *world*, not just non-OPEC or non-Middle East (Campbell, 1989; Campbell, 1991; Campbell, 1997) even though most in the industry have difficulty finding signs of near-term scarcity. His 1989 prediction that world production had already peaked and prices would rise to the $30–50 range in the early 1990s was clearly wrong, and his 1991 book produced forecasts for non-OPEC countries that were 10 mb/day too low (net) by 1999.

6. Methodological errors

There appear to be two primary errors in the design of these models. First, Hubbert-style forecasts take URR as a static variable when it is dynamic. This is a serious error. URR refers not to total resources, which is arguably a fixed amount, but to the proportion of the total which is recoverable. It is logical that this should increase over time, as technological advances raise the proportion of a field which can be recovered economically and as other changes (additions of pipelines, for example) lower costs and thus make it economical to produce smaller and/or deeper fields and less productive wells.

Some like Mackenzie (1996) argue that the estimates of URR have stabilized at a level of 2 trillion barrels, suggesting that technology has reached an asymptote and thus URR can be treated as a fixed value. However, reviews of historical URR estimates are usually dominated by those made in the 1970s and early 1980s, producing a false appearance of a historical asymptote. Since the amount of oil which is recoverable depends on the technology, and nearly all URR estimates rely on existing technology (rather than attempt to predict technological developments and their impacts), estimates produced in one period should be similar (see McCabe, 1998; 2001).

Indeed, URR estimates do seem to expand by time, with the average estimate from the 1950s and 1960s being 1.0 trillion barrels or less, while recently, numbers are 2.5–3.0 trillion. Examining individual authors—to correct for methodological differences—also finds the same pattern. The USGS has increased their estimates of URR from 1.7 trillion in 1984 to 3 trillion last year. Even Campbell has raised his estimates by 150 billion barrels from 1991 to 1997, an amount greater than consumption during that period.

And the impact of technology in expanding resources can be seen in a number of cases. As an example, in the British North Sea, many small fields that were discovered in the early days of exploration were not produced because they were not economic. Later, however, as other discoveries and other infrastructure were put in place nearby and as new technologies such as subsea templates were developed, it became much cheaper to develop small fields and connect them to nearby pipelines or platforms. Thus, in 1995, approximately 16% of production in the UK came from fields that had been discovered before 1980, but not put into production until after 1990.

Some modelers have argued that URR can be estimated using so-called “creaming curves” which show discovery size by companies or in a nation, such as the UK, to demonstrate the

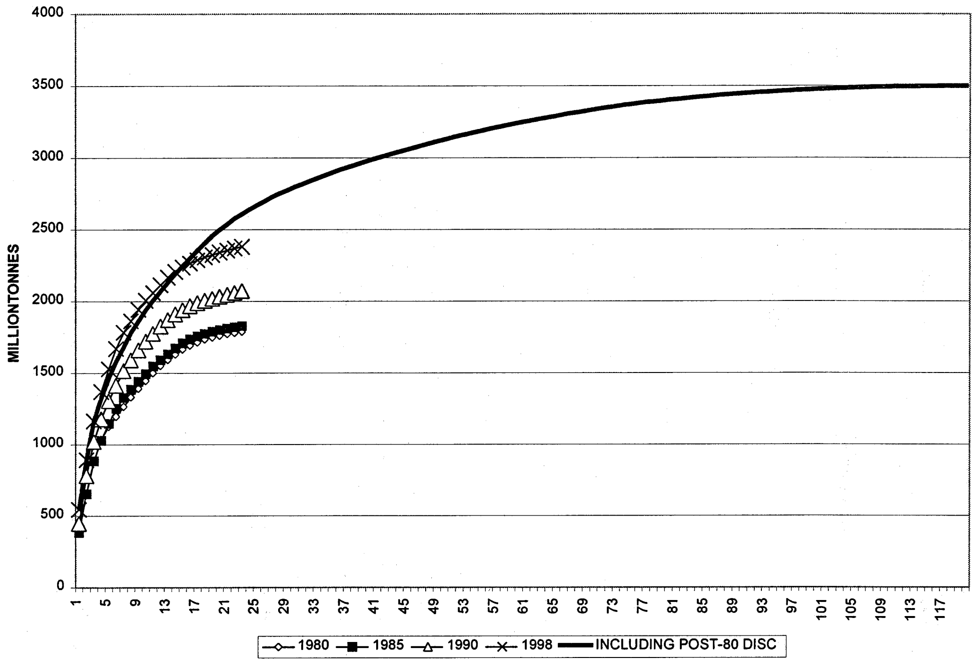

Fig. 4. Creaming curves for the U.K.

asymptote[3] (Laherrere, 1999). However, this is misleading because they compare current estimates of field size for discoveries of many different periods. This is like comparing acorns and oak trees; naturally the latter are bigger, but that doesn't prove that the former are destined to always be smaller. Using a given data base of field sizes seems to always yield an asymptote, but the asymptote moves over time. Fig. 4 shows the movement of the UK creaming curve, for example, with the four lowest curves showing the amount of oil discovered in the 23 fields that were found and producing by 1980, and the manner in which they grew over time, so that the asymptote went from 1780 million tonnes (13 billion barrels) in 1980 to 2400 million tonnes (17 bb) in 1998, a growth rate of 1.6% per year. The top curve includes all fields that had been discovered and producing by 1998, and there the asymptote reaches 3500 million tonnes (25 billion barrels).

Because of the many different definitions, interpretations and methods used to estimate URR, it is impossible to show any particular trend in URR, although I would argue 2% p.a. growth as reasonable. But this makes defining a bell curve where URR is the area under the curve problematic, since the midpoint becomes harder to determine and revisions to a URR estimate cannot be added to the historical production curve. Instead, they must be added to future production. For example, when Campbell (1997) raised his estimate of URR by 150 billion barrels, he revised his world production peak from 62 mb/day to 66 mb/day, and moved it out by 6 years, exactly the kind of correction that Lynch (1996) said would be necessary as a result of this methodological error.

7. Depletion effect

However, there is another major mis-specification. Production depends not just on discovery, but the amount of capacity lost due to depletion effects. As a field is produced, its productive capacity declines—all else being equal (that is, if no additional drilling is done, or enhanced recovery put in place). Many of the pessimistic models appear to be showing a very high rate of production decline for existing fields, and some such as Pursell (1999) explicitly argue that depletion is now so great that offsetting it—not raising capacity—is the major challenge to the industry.

But careful study suggests that the impact is being overstated, as either the rate of depletion is overstated or the ability to offset it is understated. In reviewing oil production forecasts for various nations, Lynch (1990) was able to derive the forecasters' estimates of production decline rates for existing fields and found nearly all showed an annual drop of 10–20%, close to the absolute maximum (that is, with no further investment). Since all but one of the production forecasts proved to be much too low, the implication is that either additional investment slowed the decline of production in existing fields, newer fields were offsetting more than expected, or more probably a combination of the two.

Careful study of recent forecasts demonstrates that this error has not been corrected. Laherrere (1999) estimates total oil recovery for various fields by plotting annual production levels against cumulative production, generating a figure which usually resembles a bell curve. While Laherrere doesn't explicitly extrapolate this behavior to either national or global levels, Campbell is clearly doing so, as Fig. 5 shows (using his 1991 book). The shape of the historical production curve for the Forties field in the UK precisely matches his 1991 prediction for aggregate UK production, naturally with different scaling. Obviously, he is not including new discoveries at all, presumably relying on the old adage about the vast majority of oil existing in a few large fields. Yet, as the figure shows, this is not true at all. (The production slow-down in 1988 from the Piper Alpha disaster and the resulting need to install more safety equipment at other fields also seem to have been integrated into Campbell's long-term production profile for the UK.)

8. Economic models

Economics-based supply models have also fared particularly poorly, although there are relatively few economic models, per se, and it is sometimes hard to distinguish them from either a geophysical model or an exogenous forecast. The primary flaw in this case appears to be omitted variables, particularly the inclusion of a resource depletion effect and exclusion of offsetting variables such as improvements in infrastructure, knowledge and technology. The result is that nearly all economic models have also proved to be too pessimistic (Porter, 1990; EMF11, 1991). Thus, at EMF11 in 1991, the average forecast for non-OPEC supply over the subsequent ten years called for a 0.8% per year decline if prices remained constant in real terms, whereas the actual behavior was a 2.2% per year increase while prices fell by 2–3% per year (ignoring the collapse of 1998 and the spike of 2000).

It might be thought that a simple top-down approach would suffice—estimate a price

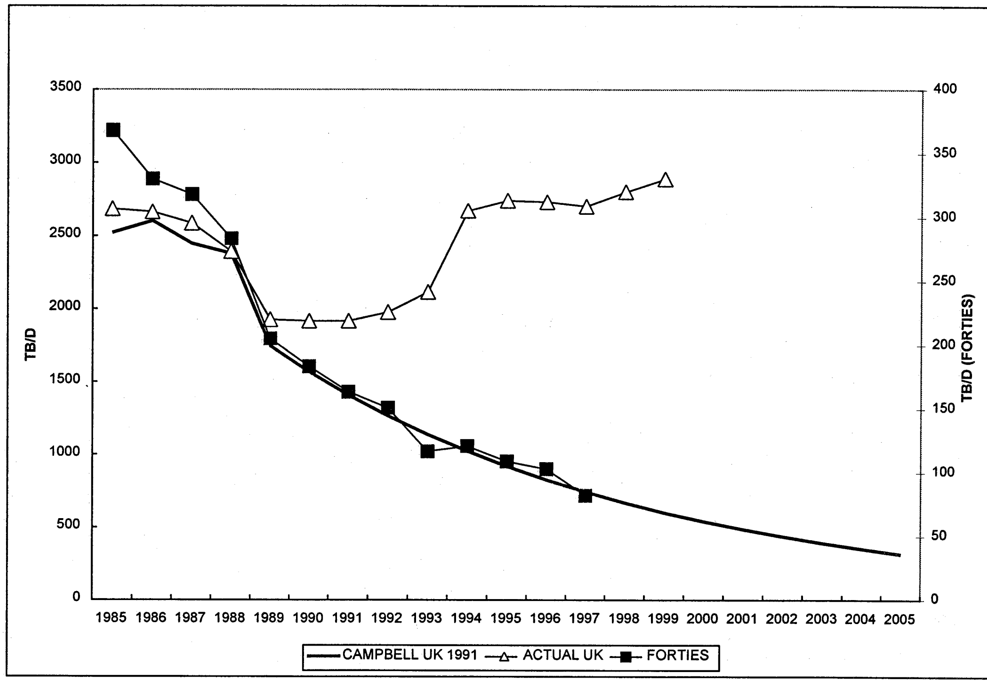

Fig. 5. The micro failing the macro: UK production by field and aggregate.

elasticity of supply and apply a price forecast. And early models (in the 1970s) did just that, assuming a price elasticity of supply such as 0.3, which naturally generated forecasts of sharply rising production—even for the US!

However, this method has proved to be underspecified. Prices soared in the 1970s and supply didn't, for one thing. But subsequently, both the non-OPEC Third World and the North Sea experienced rising production simultaneously with falling prices, implying a price elasticity of supply which is actually negative!

And in neither case can a simple correction be made. The early fix—adding a depletion effect that constrained supply by raising costs—did not improve matters much, as the inaccuracy of the EMF11 supply forecasts showed. And while productivity gains might explain rising production amid falling prices, few would credit recent technological improvements alone with such enormous impact. The true explanation is much more complex and subtle, and will be discussed below.

9. A simple supply model (and its shortcomings)

In theory, a bottom-up, microeconomic supply model can be developed with clear-cut dependent and explanatory variables: Price (or revenue) leads to exploration expenditures and thus drilling, which cause discoveries, discoveries are developed into capacity which is produced. Various factors such as cost and depletion modify the outcomes, but *in theory*, all

Table 1
A simple supply model

Dependent variables	Explanatory variables		
Exploration investment	Oil price	Costs	
Drilling	Investment	Price/rig year	
Discoveries	Drilling	Returns to drilling	
Capacity additions	Discoveries	Development expenditures	
Capacity	Capacity (t-1)	Decline/depreciation	Capacity additions
Production	Capacity	Capacity utilization	

of these would be amenable to simple econometric modeling (see Table 1). However, such a model has proved impossible to construct successfully for a variety of reasons.

Although it is easy to find research which generates results for these different variables, the reality is that there is very little data that are consistent, reliable and useful. In this section, the problems with the data for the major variables above will be discussed, including aggregation, availability, reliability, consistency and interpretation.

Investment data, for example, are reported in a number of places, often broken down by region, but are not very useful. First, publicly available data series represent private oil companies, not all oil companies. The numbers sometimes include purchases of reserves as well as drilling (depending on a company's accounting practice), but also it is often unclear what is included. Are pipelines considered an upstream expenditure? What about a dormitory for the workers? Or a helicopter pad at the drill site? Seismic analysis is often done in the head office, but that makes it difficult to allocate regionally. Further, the inherent problem of accounting for dry holes confuses cost estimation: should dry holes be averaged across all discoveries, or only those in the same region, or not at all? Governments may find it hard to understand why a dry hole drilled overseas should be charged against operations in their territory, but to a company, it is logical accounting. An exploration and production (E&P) department often wants to avoid showing regional results that will fluctuate highly with each discovery, and instead prefers to show all expenditures and drilling results as part of their overall operations.

And a semantic question complicates the issue of the results of upstream expenditure. Exploration investment adds discoveries, but discovery size is always a fraction of the ultimate size of the reserve as estimates of field size tend to grow over time—the precise amount being a subject of much debate. But this means that unit costs are hard to estimate. Too many reports of "finding costs" try to correct for this problem by including exploration and development investment and dividing it by discoveries and reserves added through development drilling to yield a per-barrel cost. However, the economics of the two processes—exploration and development—are so different that this is like mixing apples and oranges and comparing it to asparagus and cake.

Similarly, discovery data are rarely reported outside of the OECD. The primary source for non-OECD oil reserves data for most researchers is the *Oil & Gas Journal*'s yearly survey, and while historically it is generally reliable, many countries' reserves data are reported by government agencies. (There are commercial data sources but they are prohibitively expensive for most researchers and not apparently designed for trend analysis.) These data sometimes do not change for years, and other times fluctuate wildly for no apparent reason. Reserves in OPEC countries in particular have often been considered exaggerated as a result of cartel politics. But even Canadian reserves were unchanged between 1999 and 2000, which is almost certainly due to reporting error, not to the coincidence of production equaling reserve additions. In fact, Ecuador's reserves haven't changed since 1995, and this past year saw only one African country's reserves change.

Drilling data are much more available, but hardly perfect either. For instance, the shift towards more complex drilling rigs raises the average rig rate (higher costs) but improves the returns (lower costs per unit) which means that simple data on rig activity is not accurate. Even wells drilled can be misleading, not just because of the differences between onshore and offshore wells but the very significant difference between the productivity of vertical and horizontal wells. Unfortunately, published drilling data in most of the world does not separate unsophisticated land rigs from those equipped for deep drilling, or horizontal drilling, or other methods such as slim-hole drilling. Recently, the situation has improved somewhat, but the lack of historical data makes it difficult to perform statistical analyses especially outside of North America.

Further, data tend to be aggregated by region, which can be wildly misleading. Returns to drilling in Saudi Arabia are enormously different from those in Oman although both are in the Middle East, so that a decline in drilling activity in one country and increase in the other can cause an enormous change in apparent returns to investment. With the recent opening of some OPEC countries to foreign investment, analyzing historical data to predict future discoveries, capacity additions, and costs in these regions will not yield valid results.

Unfortunately, most economic models rely on aggregate investment or drilling data. A well-specified model should separate exploration and development spending, drilling for gas and oil, onshore and offshore, geological basins (a Saudi rig-year is worth more than a Syrian rig-year), rig types, and many other factors. Given data limitations, there are few cases outside North America where this can be done, and the data are not always very good in North America. And even in North America, the forecasting record has been abysmally bad.

10. Dummy variables

Although the persistence of bias and recurring errors certainly suggests model flaws, the truth is that too many variables (dummy and other) affect supply trends to allow modeling to be done with precision at a geographically disaggregate level. Tax rates are one of the most confounding elements as they not only vary country by country, and over time, but often field by field, negotiated as part of the production-sharing contracts. And since producing countries absorb the bulk of the rents from oil production—and thus most of the price risk—oil companies' profits and cash flow do not correlate precisely with prices. This is why prices do not predict investment as closely as they might.

There are also numerous policy variables that can make statistical analysis difficult if not impossible (or meaningless) and can feed modeler bias. For instance, many countries that are supposed to have reached peak production are really showing the impact of a variety of policy-related effects. The opening up of a number of OPEC countries to upstream investment by foreign oil companies has meant a relative shift in investment from non-OPEC to OPEC. Even a well-specified model of non-OPEC oil supply would probably misinterpret this as the impact of depletion.

Other factors can have similar effects (in either direction). High US or Canadian gas prices can divert drilling from oil, *ceteris paribus*. Wellhead oil price controls in Argentina discouraged foreign investment, so that supply forecasts proved too pessimistic after price deregulation occurred. Similarly, political unrest in countries like Colombia and Nigeria can reduce investment and cause a production downturn, which a Hubbert model interprets as depletion. The ultimate instance is Wattenberger's application of a bell curve to the FSU after production (and the central government) collapsed there, arguing that it was declining abnormally because the rate was greater than a Hubbert curve would provide (see Wattenberger, 1994). If ever a sector cried out for dummy variables, it was the Soviet's.

The fact is that no nation has perfectly unrestricted oil production which responds solely to economic factors. Rather a variety of regulations affect and restrict investment, including poorly designed taxes, lack of access to prospective territories, price controls, et cetera. The US is perhaps the closest case to a free market, yet it has suffered all of the above at various times (including the present). In fact, Hubbert's 1956 forecast of lower-48 oil production was as accurate as it was because of factors he did not consider. Oil import quotas raised domestic prices in the 1960s, and the termination of those quotas caused production to drop in 1970. If they had never been implemented, Hubbert's forecast would have proven too optimistic, as cheap foreign oil would have pressured domestic drillers' margins rather badly in the early 1960s, causing production to decline earlier than actually occurred.

Attempting to model the role of technology is even more difficult. In demand forecasting, many modelers simply assume an "autonomous energy efficiency improvement" factor, without attempting to show the determinants of technological advance or any specific impacts of it. On the supply side, not only is it impossible to forecast technological advances, but forecasting their impact on drilling productivity is nearly impossible. It takes at least five years for significant innovations to become widely adopted even in the US, and the impact on factors such as discovery rates, drilling productivity, and costs might not show up in the data for another five years. Advances which increase field recovery rates and growth, for example, might not be reflected in reserves data for some time. (Arguments by the pessimists that current discovery rates are abnormally low seem to be driven by this factor, but the record is not clear.)

As a result of all of these various problems, simplistic analysis of oil supply means that the modeler is confronted with the counterintuitive situation of falling prices and rising production, implying a negative price elasticity of supply, but more importantly, declining drilling and rising production, in both Europe and the non-OPEC Third World. Rigs active in the non-OPEC Third World declined from 876 in 1982 to 416 in 1997 (the nadir), while production rose from 7 to 13 mb/day in the same period.[4] If drilling drops by 5% per year and production goes up by 4% per year, it does not mean that technological improvement is adding 9% per year to drilling productivity—over and above the effects of depletion. Aggregate analysis would lead to a mistaken estimate of excessive improvements in drilling productivity, just as many have misinterpreted production declines as stemming solely from geological factors.

Both geophysical models, which ignore price, costs and investment, and the econometric models, which lack these many dummy variables, are almost automatically doomed to failure. There are so many other factors that influence oil production at both the disaggregate and the aggregate level that any econometric model of drilling, reserves additions, capacity additions and production is unlikely to be successful.

11. Returns to drilling: unfortunate convergence

There is in fact one substantive area where economists and geologists have a similar interpretation—the effect of resource depletion. Unfortunately, they appear to misinterpret this effect or at least misapply it, as they have not been able to model the effect accurately. The theory is very straightforward and not in dispute: the best deposits should be located first, since even random exploration will find the largest before the smallest and size is highly (and negatively) associated with per-unit costs. Naturally, the industry prefers to look for the cheapest available deposits, which means that over time, the quality of the remaining deposits will decline—smaller, deeper, more expensive, et cetera. The question is the degree to which this moves the supply curve, that is raises costs and/or lowers production.

In practice, it has been found that oil plays can be analyzed and a discovery curve can be developed which shows either the drop in discovery size or the change in returns to drilling (barrels found per well or foot drilled) as a function of cumulative drilling. This has been used in many geophysical models, including Kaufmann and Cleveland (1991), Reinsch et al. (1988) and a number of efforts have been made to calculate returns to drilling globally (Ivanhoe, 1983) or in regional basins Smith (1980).

However, this has not generated results which are useful in predicting oil production. Partly, some of the databases are not public (for global discoveries) and so can neither be tested nor duplicated. But also, the results used have not proven to generate valid supply forecasts. The CERI model—the only one covering numerous countries which has a lengthy track record—has been exceptionally pessimistic in most areas, even the US and Canada, where the large numbers of wells drilled and fields found should yield reliable statistical results (Reinsch et al., 1988). Indeed, Kaufmann and Cleveland (1991) are able to generate an R^2 of 0.97 for lower-48 US production by analyzing actual behavior from 1945–1989, but by 2000, their forecast is too low by about 40%, despite an assumed price that was too high.

The problems appear to be two-fold. First, the models apparently do not effectively incorporate the effects of technological improvement, which seem to have increased in recent years. In the US, the success rate for exploration wells has increased by 50% in the past ten years, for example, improving the returns to drilling.

But more important is probably the fact that nearly all models are looking at political regions rather than petroleum plays. Statistical methods of estimating discovery size appear to be valid only when analyzing a play—an area that is geologically homogenous. Since a region (country or province) can have many different plays, and the discovery of a new play cannot be predicted nor the discovery curve modeled until the play has a track record, predicting production from a geographical region defined by a political boundary is unlikely to be successful.

The aforementioned policy changes can also have a similar effect and be confused with geological constraints and/or technological improvements. For example, the privatization of the national oil company (YPF) in Argentina resulted in a significant decline in drilling, from 70 to 80 rigs active in the early 1980s, to 50–60 rigs active subsequently, but a doubling of production. Obviously the government had interfered heavily with the investment decisions of the state-owned company, and reduced its productivity accordingly. Modeling Argentina's oil supply using data predating oil price decontrol would yield very pessimistic results, while afterwards, the trend would imply enormous technological improvements.

And even if returns to drilling could be modeled, there is still the problem of predicting drilling levels. Investment is a dependent variable and subject to a number of dummy variables that are hard to forecast. Forecasting investment suffers from its own problems, especially when analyzing production at a national level and the desire to drill is influenced by other policy-related factors, including political risk but also taxes and royalties, access to desired acreage, and so forth.

12. Aggregation

John Mitchell of the Royal Institute for International Affairs has noted that non-OPEC oil supply (excluding the US and the FSU) has increased linearly for many years now, and that simple extrapolation has proved the best predictor. Although this may sound theoretically inelegant, it hints at a major solution, namely, aggregation, as a way to deal with both policy and economic uncertainties (if the model and data are unbiased) as well as the role of inertia.

Inertia refers to the stabilizing effect of the capital stock. Drilling rigs are mostly fungible, which means that changes in rigs active are much more volatile at the national level than the regional level in most of the Third World. Essentially, a reduction in drilling/investment activity in one place—whether due to political unrest, changes in the tax environment, or other factors—frees up the equipment for work in another, and given the many opportunities globally, it is rare that this equipment will go idle. The enormous drop in rigs active in the US might seem to contradict this, but it reflects primarily scrapping of older, inefficient rigs that were called into service during the boom years. Outside the US, the overall fluctuation in rigs active at the aggregate level is much smaller.

Additionally, production levels have a certain momentum, since some of the associated capital becomes "free" as existing production declines, meaning that other costs for incremental production drop. Once a pipeline is in place to serve a large field, as it declines, smaller fields should pay smaller transportation charges and thus become viable, offsetting the large field's decline. As a result, production in individual countries and regions seems to level off rather than drop at the 8–10% rates that Campbell and others expect.

However, while modeling at the aggregate level can result in relatively accurate results, it requires that the model be unbiased. Given the past two decades of Malthusian bias in oil supply forecasting, it would seem this is particularly difficult.

13. Conclusions

Most oil supply forecasting has been done very badly in the past, with many models severely underspecified. Although geology is an important determinant of discovery rate, the tendency of some modelers to interpret all supply behavior as being geologically determined is impossible to justify. And even where the models appear to be correctly specified, the results still prove to be too pessimistic, suggesting there is some remaining bias at work.

This article is of necessity much too brief to describe all of the many difficulties in analyzing and forecasting oil supply, let alone cover all aspects of the debate on the issue—the breadth of the errors, the misleading semantics, and so forth. Indeed, one of the best pieces of evidence of pessimistic bias is simply the overwhelming number of forecasts that have been produced—many from models whose design appears accurate—but which proved not only wrong, but embarrassingly too low. That many of those modelers have generated new forecasts which are nearly identical to their old ones without explaining the cause of the previous errors should make even the most casual observer skeptical.

Notes

1. Chief Energy Economist, DRI-WEFA, and Research Affiliate, Center for International Studies, M.I.T. I am indebted to M.A. Adelman for comments on a draft of this paper, but all responsibility for the contents remain my own.
2. Lynch (1990), Lynch (1996), and Lynch (1998).
3. Creaming curves are the inverse of an older method whereby field size is modeled against date of discovery to predict both future discoveries and total reserves (the area under the curve). Laherrere (2001) typically uses cumulative discoveries to show a rising curve, with the asymptote being total expected reserves (a "creaming curve"). He applies this to both individual basins and companies, to show that nearly all are rapidly approaching an asymptote. However, since he mixes current estimates of field size for older fields with initial field size estimates for new discoveries, the asymptote is likely to be incorrect, specifically too low. He argues that more recent estimates are much more accurate than earlier ones, but the evidence to date suggests that growth in field size estimates continues, so that he is comparing old orchards with newly planted saplings and extrapolating to demonstrate declining tree size.
4. Rig counts are from Baker-Hughes and production data are from BP's *Statistical Review of World Energy* (annual).

References

Adelman, M. A. (1990). Mineral depletion, with special reference to petroleum. *Review of Economics and Statistics, 72,* 1–10.

Campbell, C. J., & Laherrere, J. H. (March 1998). The end of cheap oil. *Scientific American*.

Campbell, C. J. (1997). *The coming oil crisis*. Multi-Science Publishing Company & Petroconsultants S.A.

Campbell, C. J. (1991). *The golden century of oil 1950–2050: the depletion of a resource*. Kluwer Academic Publishers.

Campbell, C. J. (December 1989). Oil price leap in the nineties. *Noroil, 12,* 35–38.

Edwards, J. D. (August 1997). Crude oil and alternate energy production forecasts for the twenty-first century: the end of the hydrocarbon era. *AAPG Bulletin, 81,* 1292–1305.

Energy Modeling Forum. (December 1982). *World Oil,* 2 vols., Stanford University, Stanford, CA.

Energy Modeling Forum. (September 1991). *International Oil Supplies and Demands,* EMF Report 11, vol. I.

Hatfield, C. B. (October 22, 1997). The oil we won't have. *Washington Post,* p. a21.

Hotelling, H. (1931). The economics of exhaustible resources. *Journal of Political Economy*.

Hubbert, M. K. (1956). Nuclear energy and the fossil fuels. *Drilling and Production Practice*.

Ivanhoe, L. F. (November 21, 1983). Free world oil discovery indexes—1945–81. *Oil & Gas Journal*.

Kaufmann, R. K., & Cleveland, C. J. (1991). Policies to increase US oil production: likely to fail, damage the economy, and damage the environment. In J. Hollander, R. H. Socolow & D Sternlight, *Annual Review of Energy and the Environment*, vol. 16.

Laherrere, J. (June 19, 2001). Estimates of oil reserves. Unpublished paper presented to the EMF/IEA/IEW meeting IIASA.

Laherrere, J. H. (April 1999). Reserve growth: technological progress, or bad reporting and bad arithmetic? In *Geopolitics of Energy*.

Lynch, M. C. (March 1998). Crying wolf: warnings about oil supply. Unpublished memo.

Lynch, M. C. (1996). The analysis and forecasting of petroleum supply: sources of error and bias. In D. H. El Mallakh, *Energy Watchers VII.* Boulder, CO: International Research Center for Energy and Economic Development.

Lynch, M. C. (1994). Bias and theoretical error in long-term oil market forecasting. In J. R. Moroney, *Advances in the Economics of Energy and Natural Resources*. Greenwich, CT: JAI Press.

Lynch, M. C. (October 1990). An omitted variable in OECD oil supply forecasting. In *Energy Supply/Demand Balances: Options and Costs*. Proceedings, International Association for Energy Economics Twelfth Annual North American Conference.

Lynch, M. C. (May 1989). *Oil prices to 2000: the economics of the oil market.* London: Economist Intelligence Unit.

Lynch, M. C. (July 1987). The Underprojection of Non-OPEC Third World Oil Production. Delivered at the Ninth International Conference of the International Association of Energy Economists, Calgary, Alberta, in Proceedings.

MacKenzie, J. (Summer 1996). Heading off the permanent oil crisis. *Issues in Science and Technology*.

Masters, C. D., Root, D. H., & Attanasi, E. D. (1990). World oil and gas resources—future production realities. *Annual Review of Energy*.

McCabe, P. J. (June 2001). Energy resources—cornucopia or empty barrel?: reply. *AAPG Bulletin*.

McCabe, P. J. (November 1998). Energy resources—cornucopia or empty barrel? *AAPG Bulletin, 11,* 2110–2134.

Meadows, D. H., Meadows, D. L., Randers, J., & Behrens III, W. R. (1972). *The Limits to Growth*, Signet.

Organization for Economic Co-Operation and Development. (1974). *Energy Prospects to 1985*, 2 vols. Paris.

Porter, E. (August 1990). Non-OPEC Supply and World Petroleum Markets: Past Forecasts, Recent Experience, and Future Prospecs. American Petroleum Institute, Research Study #054, Washington, D.C.

Pursell, D. A. (August 23, 1999). Depletion: No Longer the Forgotten Factor!, Energy Service Industry Research. Houston: Simmons & Company International.

Reinsch, A. E., Brown, K. J., & Stanford, J. O. (September 1988). *Stability within Uncertainty: Evolution of the World Oil Market.* Calgary: Canadian Energy Research Institute.

Root, D. H. (1991). Projections of Crude Oil Production Based upon U.S. Geological Survey Data on Potential Resources. In *International Oil Supplies and Demand*. Stanford: Energy Modeling Forum 11.

Smith, J. L. (1980). A probabilistic model of oil discovery. *Review of Economics and Statistics, 62,* 587–594.

U.S. Central Intelligence Agency. (August 1979). The world oil market in the years ahead. Washington: Government Printing Office.

Watkins, G. C. (1992). The Hotelling principle: autobahn or cul de sac? *The Energy Journal*, *1*, 1–24.

U.S. Department of Energy. (1983). *The petroleum resources of Libya, Algeria, and Egypt*. Foreign Assessment Series. Washington: Government Printing Office.

Wattenberger, R. A. (1994). Oil production trends in the former Soviet Union. In J. R. Moroney, *Advances in the Economics of Energy and Natural Resources*. Greenwich, CT: JAI Press.