The Economics of Peak Oil

Stephen P. Holland
University of North Carolina Greensboro *

July 19, 2011

Abstract

"Peak oil" refers to the future decline in world production of crude oil and to the accompanying potentially calamitous effects. The peak oil literature typically rejects economic analysis. This chapter, following Holland (2008), argues that economic analysis is indeed appropriate for analyzing oil scarcity since standard economic models can replicate the observed peaks in oil production. Moreover, the emphasis on peak oil is misplaced since peaking is not a good indicator of scarcity, peak oil techniques are overly simplistic, the catastrophes predicted by the peak oil literature are unlikely, and the literature does not contribute to correcting identified market failures. Efficiency of oil markets could be improved by instead focusing on remedying market failures such as excessive private discount rates, environmental externalities, market power, insufficient innovation incentives, incomplete futures markets, and insecure property rights.

JEL codes: Q3 and Q4.
Keywords: Depletable resources, Hotelling, peak oil.

*Send correspondence to: Stephen P. Holland, Dept. of Economics, University of North Carolina, Greensboro, NC 27402-6165, sphollan@uncg.edu.

The Economics of Peak Oil

Abstract

"Peak oil" refers to the future decline in world production of crude oil and to the accompanying potentially calamitous effects. The peak oil literature typically rejects economic analysis. This chapter, following Holland (2008), argues that economic analysis is indeed appropriate for analyzing oil scarcity since standard economic models can replicate the observed peaks in oil production. Moreover, the emphasis on peak oil is misplaced since peaking is not a good indicator of scarcity, peak oil techniques are overly simplistic, the catastrophes predicted by the peak oil literature are unlikely, and the literature does not contribute to correcting identified market failures. Efficiency of oil markets could be improved by instead focusing on remedying market failures such as excessive private discount rates, environmental externalities, market power, insufficient innovation incentives, incomplete futures markets, and insecure property rights.

JEL codes: Q3 and Q4.
Keywords: Depletable resources, Hotelling, peak oil.

1 Introduction

The concern over "peak oil" extends well beyond a mere decline in the production of crude oil and has become practically synonymous with the end of the world as we know it. (See Porter (2006) for a critical review of this literature.) To predict peak oil, a large literature attempts to estimate whether or not global oil production has peaked by using a simple logistic model first developed by Hubbert (1956). Hubbert's model attained fame by successfully predicting the peak in U.S. crude oil production in 1970, and peak oil researchers following Hubbert's techniques generally reject economic analysis. Indeed simple versions of the standard economic model of exhaustible resources as first developed in Hotelling (1931) do not predict a peak in production. This raises several important questions. Could the observed oil peak have arisen from an economic model or is it indicative of some sort of market failure or disequilibrium? Is peak oil just a series of coincidences that is not amenable to economic analysis?

Holland (2008) demonstrates that economic models can indeed predict peaking and thus that peaking is consistent with efficient use of exhaustible resources and is not indicative of some market failure. Following Holland, this chapter provides a nontechnical description of production peaking in standard models of exhaustible resources. I then argue that the focus on peak oil is misplaced.

2 How fast should we consume oil?

To begin, consider how fast society should utilize an exhaustible resource such as oil. Answering this question requires some model of oil usage. Since the question is ultimately an economic question, it is reasonable to begin with a standard economic model. To be concrete, I illustrate the model with prices and quantities, which are meant to be reasonable but not necessarily realistic.

For this modeling exercise, it is useful to make explicit several assumptions. First, assume that consumption of energy from oil has some benefits to society (*e.g.,* providing transportation services) and that these benefits are quantifiable. Moreover, assume that in each time period the social benefit has a declining marginal benefit. Declining marginal benefit implies that, for example, the benefit from consuming the 101st barrel of oil is less than the benefit from consuming the 100th barrel of oil. For concreteness assume that this marginal benefit function is constant over time and is \$500 per barrel when consumption is 30 million barrels per day and \$50 per barrel when consumption is 85 million barrels per day (approximately current world oil consumption).

Second, assume that there is some substitute resource, for example, solar energy, which is available at a much higher cost than oil. For simplicity, assume that solar is available at constant marginal cost (in this example \$500 per barrel of oil equivalent) and at this high price could satisfy the energy demand of 30 million barrels per day.

Third, assume that the supply of oil is finite. Importantly, if oil is used at some time, it cannot be used at some other time, *i.e.,* oil is depletable and has a scarcity cost. Moreover assume the oil can be extracted at some constant cost per barrel. Thus the opportunity cost to society for using a barrel of oil is the scarcity cost plus the extraction cost plus any environmental cost associated with its extraction. For concreteness, assume that the oil can be extracted at a cost of \$50 per barrel which includes both extraction and environmental costs. The scarcity cost will depend on how much oil remains. Note that these extraction and environmental costs are relatively low by assumption. This implicitly assumes that oil resources with high environmental damages have been protected from development, so only low damage resources are available. Moreover this ignores oil deposits with exceptionally high extraction costs.

Finally, assume that there is some (potentially small) social time preference for current consumption. For simplicity assume that this preference is captured by a social discount rate.

Given these assumptions, which are standard in economics, the *equimarginal principle*, a central principle in economics, describes how energy should be used to maximize benefits to society. The simplest statement of the equimarginal principle is that net benefit is maximized when marginal benefit equals marginal cost. The intuition is straightforward: if the social marginal benefit of energy consumption were greater than the social marginal cost of energy, society could be made better off by increasing energy production. Similarly, if the social marginal benefit of energy consumption were *less than* the social marginal cost of energy, society could be made better off by *decreasing* energy production. Social net benefit is only maximized when these margins are equal.

Applying similar logic to oil production, the equimarginal principle implies that social net benefits from oil are maximized when the net marginal benefit of the oil is equal across all time periods. This principle holds since if net marginal benefits were higher in some periods, society could be made better off by reallocating oil consumption to those periods in which net marginal benefits were greater.

Despite the simplicity of this principle, it has several implications in this model for socially optimal energy consumption, solar energy production, oil extraction, and the transition from oil to solar. First, at any time when solar energy is being produced, energy consumption should be 30 million barrels per day equivalent so that the marginal benefit of consumption would be exactly \$500. Note that when production is 30 million barrels per day, net marginal benefit from solar is zero, *i.e.,* marginal benefit equals marginal cost. Furthermore, this implies that the marginal benefit of energy should never exceed \$500 per barrel since additional solar production is always available at the cost of \$500 per barrel equivalent.

Now how should the oil be used? If society used 30 million barrels per day, the net marginal benefit would be \$450 per barrel of oil. At this rate the oil would last for quite some time. However, due to the social time preference for current consumption, this would violate the equimarginal principle. Since the \$450 net marginal benefit in the future is worth less than a \$450 net marginal benefit today, society would be better off by shifting some oil consumption from the future to today. How much consumption should be shifted to the present depends precisely on the social rate of time preference. The equimarginal principle is satisfied when net marginal benefits from oil are equal in present value. Thus in any two periods in which oil is used, the present value of the net marginal benefit should be equal, or equivalently, the net marginal benefit should grow at the social rate of time preference. In general, this implies that consumption is declining over time.

Now consider the transition from oil to solar. The key question is what the marginal benefit should be when oil is exhausted. Since solar production optimally prevents the marginal benefit from exceeding \$500, the marginal benefit when the oil is exhausted should not exceed \$500 per barrel. Note that the marginal benefit should also not be substantially less than \$500.

For example, if the marginal benefit were \$400 when the oil were exhausted, social benefit could be increased by delaying consumption if the present value of the \$450 net marginal benefit from oil in the next period were greater than the \$350 net marginal benefit from oil in the period of exhaustion. Thus the marginal benefit at the transition should be \$500. Note that the oil should be completely exhausted at this point.

In summary, social net benefits from energy in this model are maximized when the marginal benefit net of the \$50 extraction and environmental costs grows smoothly at the rate of social time preference to \$500 per barrel. At this time the oil is completely exhausted, and consumption transitions to solar energy.

This simple model abstracts from many important details about energy production and consumption. For example, it ignores different extraction costs for oil, different grades of oil, increasing marginal extraction costs, limited solar capacity, innovation, other renewables and nonrenewables, etc. The basic model has been extended to analyze these additional features (see Krautkraemer (1998) for a survey) but the central insight remains: the net marginal benefit should grow at the rate of social time preference until the marginal benefit reaches the cost of producing substitute resources.

3 How fast would markets use oil?

Having analyzed society's optimal energy use, now consider how markets would use the oil and transition to solar energy. One of the central insights of economics is that perfectly competitive markets lead to consumption which is optimal for society, *i.e.,* which is Pareto efficient. This result also holds for markets in exhaustible resources such as oil. However, it is useful to analyze in detail the competitive equilibrium in oil markets to illustrate the key assumptions upon which efficiency relies.

The competitive equilibrium is characterized by the demand for energy, the supply of solar energy, and the supply of oil. The demand for energy and supply of solar are straightforward. If income effects are negligible, the demand curve for energy is equivalent to the marginal benefit function. Thus for a given price the demand curve implies that the price equals the marginal benefit.

The solar supply depends on the energy price. If the price of energy were above $500 per barrel equivalent in any period, solar producers would supply as much as possible in that period. However, if the energy price were below $500, no solar would be supplied. If the price were exactly $500, then solar producers would be willing to supply any amount.

The supply of oil, an exhaustible resource, is more complex. The supply in the first period depends on the current price and all future prices. For example, if the current price were $80 per barrel and the price next year were $100, at reasonable market interest rates, the oil company would maximize profits by delaying production until next year since it would make $30 per barrel from producing today but $50 per barrel from producing next year. In fact, the only price series for which the firm would be willing to supply non-zero production in each period is a price series that yields the same present value profit in each period. In other words, a price series in which the price, net of costs, grows at the rate of interest.

In equilibrium, the price cannot exceed $500 since otherwise there would be an excess supply of solar. At prices below $500, the price net of extraction costs must grow at the rate of interest since otherwise there would be an excess supply of oil in some periods and excess demand in other periods. For any such price path, the quantities determined by the demand curve are a potential equilibrium since the supply of oil would equal demand in each period. There are many potential equilibrium price paths with this growth rate. However, only one of these price paths exactly uses all the available oil when the price reaches $500. Since all higher price paths would have excess supply of oil and all lower price paths would have excess demand, this is the unique equilibrium price path.

The equilibrium is thus characterized by the price path which grows (net of costs) at the rate of interest and exactly uses all the available oil when the price reaches $500. Thereafter, the price remains at $500 and demand for energy is satisfied by solar power. Note that since the consumer's optimization implies that the price equals the marginal benefit, the competitive equilibrium is exactly the same as the social optimum, *i.e.,* is efficient, if the market interest rate and social discount rate are equal and if the firm faces the true environmental costs of extraction.

This efficiency relies upon several important assumptions.[1] First, the market interest rate must equal the social discount rate. There are a number of reasons that market participants may be less patient than would be socially optimal. Excessive market impatience could cause the oil to be produced and consumed too quickly. Second, the costs the firm faces must be exactly the social costs. In particular, the firm must pay for the full environmental costs of the oil extraction. Thus all externalities must be internalized, for example, through environmental taxes or fees. If firms do not face the full environmental costs, then the market equilibrium will use the oil too quickly. Third, the firm must be a price taker, *i.e.,* have no market power. This holds if producers are small enough that they are unable to influence market prices. If firms have market power, the market equilibrium may use the oil too slowly. Fourth, the gains to innovation must go to whoever bears the costs of innovation. If consumers gain from innovation (through lower prices) but firms must pay for licensing or research and development costs, then the innovation incentives may be insufficient. Innovation incentives could be insufficient for both oil and solar production. Fifth, firms must be assured that if they do not produce oil today, they can produce the oil in the future. This assumption is violated, for example, if oil fields are subject to open access (rule of capture) or if future governments might revoke mineral rights, *e.g.,* after a regime change. If future extraction rights are not secure, the market equilibrium will use the oil too quickly. Finally, firms must be forward looking, *i.e.,* cannot be myopic. The transversality condition ensures that the equilibrium price path is the price path which exactly uses all the oil. However, many other price paths also have the same growth rate. If for some reason the price path had the right growth rate but were too low, there would be an arbitrage opportunity: a firm could increase its profit by holding oil until the price jumped. However, the price would not jump until the mistake became widely known, possibly after all the oil had been used. A variety of factors (such as borrowing constraints, contractual production requirements, etc.) may prevent firms from exploiting such arbitrage opportunities, which could lead to an inefficient competitive equilibrium. Complete forward markets could also help markets realize such arbitrage opportunities. However, oil futures markets, which in the near term are quite liquid and well-functioning relative to other commodities, are virtually non-existent

[1] Asymmetric information, another crucial assumption for the efficiency of markets, is not a primary concern for oil markets.

beyond the near term (*i.e.,* beyond five to ten years).

4 Modeling peak oil

The basic intuition of the economic models is that over time the oil becomes scarcer and thus the opportunity cost of the oil—the extraction cost plus the scarcity cost—increases over time. In the simplest model, this implies that production is gradually declining over time. In other words, there is no peak oil. But Brandt (2006) finds evidence of peaking in many oil markets. Does peaking suggest that there is some other important factor that the economic model is missing? Does this missing factor indicate additional inefficiency?

Holland shows that economic models of exhaustible resources can imply a peak in oil production if there is some economic factor which offsets the declining production due to the exhaustion of the oil. In fact, Holland describes four countervailing factors which can lead consumption to increase: demand growth, technological change, increasing reserves, and site development. Each of these factors can be strong enough to cause equilibrium oil production to increase for a time. However, in each case, the increasing scarcity cost eventually offsets the factor leading to increasing production and thus oil production peaks.

The first model, in which demand growth is the countervailing factor, is the simplest. If demand for energy is growing, for example, due to increases in income or population, then this increasing demand can tend to increase consumption in the early years. Initially this demand growth can offset the increasing scarcity cost, and consumption can increase. But eventually increasing scarcity dominates and the increasing energy price leads to a decrease in oil consumption. In other words, oil production peaks.

Although the demand growth model predicts peaking, it also predicts that the oil price should steadily increase. The evidence reviewed in Holland suggests that oil prices do not steadily increase and may even follow a U-shaped pattern. This suggests that additional factors may be causing the observed peak.

The second model focuses on the supply side and shows that technological change can lead to peaking. In this model, technological advances, for example, in drilling, extracting, or finding oil, can decrease the cost of producing oil. This reduction in extraction costs can offset the oil's rising scarcity cost so that the opportunity cost of the oil falls, and oil production increases. However, eventually scarcity dominates and production decreases, *i.e.,* oil production peaks.

This model relies on, but does not explain, technological change. Thus its predictive power is limited if technological change does not occur or even slows down. Thus the countervailing factor in this model is not necessarily inherent to oil production.

The third model relies on the relationship between extraction costs and reserves. Pindyck (1978) argues that there is an inverse relationship between extraction costs and reserves. This could arise since a larger resource base may make finding low cost deposits easier. Thus reserve growth can be a countervailing factor which can decrease opportunity costs and cause production to increase. As in the other models, eventually scarcity dominates and production decreases, *i.e.,* oil production peaks.

Since this model explicitly explains the source of the cost reduction, it has more predictive power than the technological change model. However, the precise mechanism by which reserve growth decreases costs is not well understood.

The fourth model, in which site development is the countervailing factor, does not have clear antecedents in the literature and is a novel contribution in Holland. In the model, there are a limited number of oil development sites. Each year, firms choose how many sites to develop and how much capacity to install at each site. Once a site is developed, oil production at the site depends on the installed capacity and continues until the oil at the site is exhausted. In equilibrium, the firms develop smaller sites and install less production capacity over time. Since production overlaps at a number of sites, aggregate production increases initially if sufficient new sites are developed to offset the production declines at all the existing sites. Since the number of development sites is limited, eventually new production cannot offset production declines at the existing sites and aggregate production decreases, *i.e.,* production peaks.

Each of these four models shows that peaking is consistent with economic models of oil markets. Thus peaking does not indicate an additional failure of economic models and should not be considered a surprise. In fact, the countervailing factors in the model are such fundamental factors inherent to so many markets that it would probably be more surprising if oil production did *not* peak than that it peaked!

5 The misplaced emphasis on peak oil

Since peaking is consistent with economic models and is predicted by reasonable analysis of economic factors, it is perhaps surprising that so much attention has been given to peak oil. Perhaps this emphasis reflects a general distrust of economic analysis. However, the emphasis on peak oil is misplaced for a number of reasons.

First, oil peaking is not a good indicator of oil scarcity. An important focus of the peak oil literature has been on predicting the peak in world oil production as an indicator of pending scarcity. However, Holland shows that in the economic models of peaking the peak can occur when none, some, or all of the oil has been consumed. Thus peaking is not a good indicator of pending scarcity.

Smith (2011) argues that scarcity indicators should be evaluated using a comparative static criterion developed by Brown and Field (1978). According to this criterion, if, for example, there is a surprise decrease in scarcity due to a new oil discovery, the indicator should go down. Smith argues that "peaking is an ambiguous indicator that provides inconsistent signals regarding resource scarcity." For instance, he shows an example in which scarcity increases due to a permanent increase in demand but with no change in the growth rate of demand. In this example, the peak occurs earlier, suggesting that an earlier peak indicates scarcity. He contrasts this with another example in which scarcity increases due to faster than expected demand growth. In this example, the peak occurs *later* suggesting that a *later* peak indicates scarcity. Thus, it is not clear whether an earlier peak is good or bad news, again suggesting that oil peaking is not a good indicator of scarcity.

Smith also compares the date of the oil peak to traditional economic indicators of resource scarcity such as unit cost, resource rent, and market price. Unfortunately, Smith notes that unit cost and resource rent can also provide inconsistent signals. For example, unit cost (extraction cost per barrel) is a good indicator of supply side scarcity, but does not respond to scarcity caused by demand side shifts or by shifts in the resource base. The second indicator, resource rent, is not directly unobservable, which limits its applicability. Moreover, Smith argues that resource rent is an inconsistent indicator. For example, increasing scarcity from a decrease in the resource stock causes the resource rent to increase, but increasing scarcity from an increase in the unit cost causes the resource rent to *decrease*. Holland notes that in each of the models of peaking, the price begins to rise before the peak occurs. Thus he argues that if we wish an early indicator of pending scarcity, price is a better early indicator of pending scarcity.[2]

A second reason that the emphasis on estimating peak oil is misplaced is that it suggests that there is an easy way to estimate oil scarcity. In fact, the simple logistic model first developed by Hubbert, estimates ultimately recoverable reserves of oil using only cumulative oil production. If we truly trusted this model, public agencies would not need to expend resources collecting accurate measures of prices, costs, and reserves. They could simply track cumulative production and use that to estimate remaining oil resources. The focus on peak oil suggests that the costly work of developing and updating consistent, high quality estimates of reserves is superfluous and/or unnecessary. I suspect that even the most ardent peak oil advocates do not follow their logic to this extreme conclusion.

A third reason for the misplaced emphasis on peak oil is that the types of catastrophes suggested by the peak oil literature are unlikely to occur. The history of oil markets is characterized by highly volatile prices with relatively small supply disruptions leading to huge price swings. This volatility likely arises from the inelastic short-run supply and demand which require large price swings to ration relatively small supply or demand shocks. The peak oil literature predicts even larger price swings as society inexorably transitions to a post carbon world.

The models described in this chapter efficiently have a smooth transition to the substitute solar energy. However, it is easy to image models in which the transition from oil to solar is not smooth and is not efficient. For example, if oil were extracted as an open access resource, any time the price exceeded $50 the supply would be huge as all producers would extract oil in order to capture the oil *e.g.,* before some other producer did or before the regime changed. In the simple

[2]Smith plans to evaluate market price using the comparative static criterion in future work.

model of open access, the equilibrium price would be exactly \$50 until all the oil is exhausted. At that point, the energy price would jump to \$500 as the economy transitioned to solar energy. This sudden transition seems to be exactly the type of chaotic, costly transition predicted by the peak oil literature. Moreover, it is not unreasonable to characterize oil markets as having less than secure property rights in oil especially given the large proportion of oil production from countries with unstable governments.

However, the chaotic, costly transition in this model hinges on one key assumption (namely a uniform extraction cost) which is not particularly reasonable. In reality, there are many different grades of crude oil, some which are quite cheap to extract and others which are quite costly to extract. Without the simplifying assumption of uniform extraction cost, the equilibrium looks quite different. No longer would the energy price jump from \$50 to \$500, but rather it would jump from \$50 to the extraction cost of the next cheapest deposits. After these deposits were all exhausted, the price would jump to the extraction costs of the next cheapest deposits. Although this equilibrium would be quite inefficient and would use the oil too quickly, it would not have a sudden, large, potentially costly transition. The existence of a variety of grades of oil with a variety of extraction costs prevents a large price jump with the kind of costly transition predicted by the peak oil literature.

The final reason the emphasis on peak oil is misplaced is that it distracts from remedying the well-known inefficiencies of oil markets. In general, markets can be good at efficiently allocating resources. However, if markets are not perfectly competitive, then they will not generally be efficient. The theoretical discussion of oil markets above described the assumptions required for oil markets to be efficient. Clearly most of these assumptions are not satisfied by current oil markets. Most prominent among the market failures is that extraction costs do not accurately incorporate environmental costs due to insufficient environmental taxes and implicit or explicit subsidies to oil production. Also important are factors which hinder the efficient delay of oil production such as excessive private discount rates, incomplete futures markets, firm myopia, minimum production lease requirements, and insecure property rights. Each of these factors leads the market equilibrium to produce oil too quickly. Inefficient innovation incentives may cause oil to be produced to slowly and may delay the transition to substitute resources. Market power may lead firms to produce oil too slowly which implies that firms capture surplus at the expense of consumers. Each of these market failures is well-studied, and economists have developed remedies for many of them, *i.e.,* Pigouvian taxes to cover environmental costs; price regulation or antitrust enforcement for market power; subsidies or public research and development to increase innovation incentives; and increasing security of property rights. Since it is no simple matter to calibrate these corrective mechanisms correctly, a misplaced emphasis on peak oil can distract from correcting these identified market failures and can prevent markets from becoming more efficient.

6 Conclusion

Peak oil has recently captured the public attention with dire predictions about pending oil scarcity. This literature generally rejects economic analysis. This chapter argues that economic analysis is indeed appropriate for analyzing the fundamentally economic question about oil scarcity since standard economic models can replicate the observed peaks in oil production. Moreover, I argue that the emphasis on peak oil is misplaced since peaking is not a good indicator of scarcity, peak oil techniques are overly simplistic, the catastrophes predicted by the peak oil literature are unlikely, and the literature does not contribute to correcting identified market failures. Efficiency of oil markets could be improved by instead focusing attention on the less exciting but well-understood market failures such as excessive private discount rates, environmental externalities, market power, insufficient innovation incentives, incomplete futures markets, and insecure property rights.

References

[1] Brandt, Adam. 2006. “Testing Hubbert,”*Energy Policy*, 35: 3074-3088.

[2] Brown, Gardner M., Jr. and Barry C. Field. 1978. “Implications of Alternative Measures of Natural Resource Scarcity,” *Journal of Political Economy*, 86(2): 229- 243.

[3] Holland, Stephen P. 2008. “Modeling Peak Oil,” *Energy Journal*, 29(2): 61-80.

[4] Hotelling, H. 1931. “The Economics of Exhaustible Resources,” *Journal of Political Economy*, 39: 137-75.

[5] Hubbert, M. King. 1956. “Nuclear Energy and the Fossil Fuels,” *American Petroleum Institute Drilling and Production Practice, Proceedings of Spring Meeting*, San Antonio, pp. 7-25.

[6] Krautkraemer, J. 1998. “Nonrenewable Resource Scarcity,” *Journal of Economic Literature*, 36: 2065-2107.

[7] Pindyck, Robert S. 1978. “The Optimal Exploration and Production of Nonrenewable Resources,” *Journal of Political Economy*, 86(5): 841-861.

[8] Porter, Richard C. 2006. “Beyond Oil: The View from Hubbert’s Peak/The End of Oil: On the Edge of a Perilous New World/The Long Emergency: Surviving the Converging Catastrophes of the Twenty-First Century,” *Journal of Economic Literature*, 44(1):186-190.

[9] Smith, James L. 2011. “On The Portents of Peak Oil (And Other Indicators of Resource Scarcity)” mimeo.