Hotelling Meets Darcy: A New Model of Oil Extraction

Charles F. Mason* and Klaas van 't Veld*

March 4, 2013

Abstract

For decades resource economists have relied on the seminal Hotelling paper to model extraction and price paths, despite overwhelming evidence of the empirical limitations of the approach. Particularly in oil markets, the shortcomings of the Hotelling approach are evident: for over 100 years real oil prices were basically constant, in marked contrast to the prediction that rents should rise at the rate of interest. A number of explanations have been offered, including the steady flow of newly found deposits and technological advances that might have continually shifted the price path out, thereby obscuring what otherwise would have been the trend towards higher prices. While these aspects are no doubt important the conventional explanation misses what we believe to be a crucial ingredient, namely production constraints that inhibit agents' ability to shift production forward in time. These constraints are rooted in the physical constraints associated with subterranean fluid flows, as articulated by Darcy. In particular, Darcy's approach predicts that output from any given well decline exponentially, at an exogenous rate that is dictated by geological features. As such, the only variable under the extractor's control is the number of deposits to open up for extraction. We articulate a model that incorporates these features, use it to generate a predicted market equilibrium price path, and discus the impact of breakthrough innovations upon the price path.

* Department of Economics & Finance, University of Wyoming, 1000 E. University Ave., Laramie, WY 82071.
E-mail: klaas@uwyo.edu, bambuzlr@uwyo.edu

1. Introduction

Arguably the most heavily cited paper in the environmental and resource literature, and perhaps one of the most heavily cited in the economics profession as a whole, is the seminal Hotelling (1931) piece. The heavy reliance on this paper is understandable: it contains an elegant model that generates clean predictions, namely that resource extracting firms will operate in such a manner as to induce rents to rise at the rate of interest. But is is generally well-known that the Hotelling model does not fare well empirically; indeed, a number of authors have generated negative results, particularly for what is arguably the most economically important resources for modern economies, namely crude oil. The problem, as Figure 1 clearly demonstrates, is that for long periods of time the path of oil prices bears no resemblance to the hypothetical path under which rents would rise exponentially.

A range of explanations have been offered for this conundrum, mainly focusing on tendencies for potential supplies to expand over time because of the discovery of new deposits or the arrival of important technological advances. There is likely some merit in these explanations: proven oil reserves steadily increased for decades as new fields were located, or existing fields were expanded. But there is a difficulty with this explanation: if rents do not rise at the interest rate, agents should move production forward in time; in particular, they should speed up their exploratory and development drilling to bring as many fields on line as soon as possible. In practice, firms tend to add to reserves more or less to offset the reductions that occur because of extraction levels. Regarding technological advances, these innovations would indeed shift out oil supply profiles, but sporadically; between these shifts the Hotelling logic should apply, so if a stream of innovations were the major explanation for the empirical failure of the Hotelling rule one should see a sawtooth pattern to prices: they should rise between those moments when technological advances occur, dropping sharply immediately after the arrival of the new technology.[1]

[1] Of course, this presumes the advances are not foreseeable. To the extent that new advances are foreseeable agents could arbitrage across time, smoothing out the teeth in the pattern we just described. But surely most important innovations are not widely adopted until after they are proven; prior to that demonstration they are largely ignored. This was true, for example, with the relatively new techniques of hydraulic fracturing and horizontal drilling, technologies that were known but rarely if ever utilized for decades.

An alternative explanation is that prices do not appear to rise steadily because the interest rate is shifting around. But if that explanation were valid, one would expect to see a clear negative link between interest rates and oil prices: unanticipated drops in interest rates would induce agents to shift production backwards in time, pushing up current prices; likewise, unanticipated increases in interest rates should push production forward in time, leading to price drops. Either way, one would expect to see a negative relation between interest rates and prices. But consider Figure 2, which plots combinations of the daily price of crude oil, as measured by the West Texas Intermediate (WTI) spot price, and the daily Treasury bill (T-bill) rate between 1986 and 2012. There appear to be four regimes, one of which does exhibit the alleged negative relation, but two of which suggest no relation (the band of observations where WTI prices vary sharply and T-bill rates do not vary at all, and the band of observations where T-bill rates exhibit a marked variation but WTI prices are roughly constant), and a regime in which oil prices and T-bill rates are positively correlated. Digging a bit deeper into the data, Figure 3 plots WTI prices and T-bill rates over time; the four regimes are easily identified here: T-bill rates fell between 1986 and 2005, but oil prices changed very little; oil prices rose sharply after 20009 but T-bill rates did not change; both oil prices and T-bill rates rose between 2005 and 2007; and T-bill rates fell, while oil prices rose, during the first part of 2008. Certainly there is no indication of a persistent negative relation between oil prices and interest rates.

A question that is not asked has to do with the empirical relevance of the inter-temporal arbitrage that would lead to the inter-temporal equation of discounted rents. Suppose that rents today exceeded discounted rents nest period; under the Hotelling logic, firms should increase current efforts, shifting production forward in time. But if such increases were difficult to impossible to produce there would be no reason to see discounted rents inter-temporally equated. And there is good reason to expect there are constraints that preclude the moving forward of production in time: there are physical constraints that limit the rate at which oil can be extracted from any particular well. One could increase production by drilling additional wells, but that sort of action is a lumpy decision, not the continuous arbitrage implicitly assumed in the conventional dynamic optimization scenario; it is easy to see that there can be scenarios where the incremental gain in present discounted value associated with the shifting of production across time can not cover the up-front costs of drilling the additional well, in which case the putative arbitrage will not obtain, and so the downward pressure on current price will fail to materialize.

In this paper we propose a model of resource extraction patterned after these physical realities, where current actions are limited by physical constraints, and where expansion decisions are discrete events. We then extrapolate to the market as a whole, and describe the inter-temporal pattern of prices that would be predicted. We then describe a scenario where a new technology arrives, shifting out the production possibilities in the immediate future, and sketch out the impact upon the path of prices that is likely to obtain.

2. Modeling Oil Production

2.1. *Darcy's Law*

Our point of departure is the observation that oil production requires moving a fluid (crude oil) through rock substrata to a well bore. The mathematical characterization of the rule governing such flows is due to Darcy (1856), which is summarized in Darcy's law; this rule essentially guarantees that the flow of output from an oil well will be proportional to the size of the deposit from which the extraction is taken. Letting $q_t =$ the rate of extraction, R_T the remaining reserves and δ the factor of proportionality, one has the natural constraint

$$q_t \leq \delta R_t. \tag{1}$$

As is usual in resource extraction models, the size of the remaining deposit declines at the rate of production:

$$\dot{R}_t = -q_t, \tag{2}$$

which then implies

$$\dot{R}_t = -\delta R_t, \tag{3}$$

at any moment where the upper bound on the extraction rate binds. Accordingly, we refer to the proportionality factor δ as the "decline rate" of production. The particular value of the coefficient of proportionality depends on a variety of geological features, which are exogenous from the firm's perspective.[2]

This description holds at the individual well level; adding wells adds production capacity. There is good reason to believe these extra wells will increase productive capacity in a linear fashion when the number of wells is sufficiently small, but that diminishing returns will set in at some point (Hannesson, 1998). In addition to drilling incremental wells, one might imagine the firm undertaking efforts aimed at increasing the difference between surface and subsurface pressures, so as to increase the rate of production; examples include injecting water or $CO_2\dot{S}$uch an approach runs the risk of destroying the structure in the host rock, if the induced pressure becomes too large.[3]

To the extent that Darcy's law holds, one would expect to see a relatively stable linear relation between oil production and proven reserves. Figure 4 provides evidence supporting this conjecture. Each observation in this figure represents a combination of US oil production and proven reserves, for a given year, for the period from 1920 to 2012. While there are some deviations, it is apparent that the pattern closely resembles a linear pattern. This over-arching linear also relationship provides a reasonable fit at a less aggregated level, as depicted in Figure 5. Here, we plot annual combined production for the UK and Norway against proven reserves from 1993 to 2012. Again, the pattern is relatively close to linear.

[2] Examples include the porosity and permeability of the host rock, which determine how easily fluids move through the rock, and the depth of the deposit, which determines the pressure differential between reservoir and surface. Also relevant is the capacity of the pipe inserted into the well, which will be dictated by the specifications made available by firms engaging in the production of this form of capital. While a very large firm might have sufficient bargaining power to influence the characteristics of pipes available in the market, it seems unlikely that a generic oil extracting firm could exert this sort of influence. Accordingly, it seems plausible that this determinant is also mainly exogenous.

[3] Petroleum engineers speak of *fracture pressure*; if the firm raises pressure above this level, additional fractures are created. These new fractures create channels in the rock that allow injected fluids bypass oil, lowering productivity. There is also the additional concern that surrounding rock will start to break apart, allowing sand and grit to get into the well bore and accelerating depreciation of the capital.

2.2. *An Adapted Hotelling Problem*

If Darcy's Law does apply to oil extraction, then the typical firm's dynamic optimization problem must be adjusted to reflect this constraint. Accordingly, let us write the firm's extrication cost function as $c(q_t, R_t)$, and the market price of oil in period t as p_t. The firm's optimization problem is

$$\mathbf{max}_{q_t} \int [p_t q_t - c(q_t, R_t)] e^{-rt} \mathrm{dt}$$

$$\text{subject to: } \dot{R} = -q_t; 0 \leq q_t \leq \delta R_t; R_t \geq 0;$$

and the initial level of reserves given as R_0. Most of this description is standard; the upper bound on the extraction rate is new, reflecting the role played by Darcy's Law. The current value Hamiltonian for this problem is

$$\mathcal{H} = p_t q_t - c(y_t, R_t) - m_t y_t + \lambda_t (q_t - \delta R_t).$$

The solution to the firm's dynamic optimization problem satisfies Pontryagin's maximum principle:

$$p_t - \frac{\partial c}{\partial y} - m_t = -\lambda_t, \tag{4}$$

$$\dot{m}_t = r m_t + \frac{\partial c}{\partial R} + \lambda \delta, \tag{5}$$

together with the state equation (2). The interpretation of these equations is that rents rise at r — *i.e.*, Hotelling's rule applies — *iff* production constraint does not bind. But if, on the other hand, the constraint does bind, then production will decline exponentially; the firm is, in essence, relieved of any ability to control the pace of activity.

While the firm exerts no control over the pace of extraction within a regime where Darcy's Law applies, it can influence the scope of extraction by undertaking to add to its reserve base. Such actions do facilitate production increases, via the increase in reserves. One can then think of the firm's supply curve as containing two bits: a part which is perfectly inelastic, reflecting extraction from its initial reserve base, and a part which is upward-sloping, associated with the reserve additions.

To formalize this notion of expanding production via reserve additions, we introduce a new state variable, A_t, which measures cumulative additions, and a new control variable a_t, which measures the rate of reserve additions. By construction, the initial value of this variable is 0; its evolution over time is given by

$$\dot{A}_t = a_t. \tag{6}$$

Finding and developing new additions is costly, and this cost is increasing in both the level of effort a_t and the stock of cumulative additions A_t. We model these features by assuming the cost of additions can be expressed as[4]

$$C(a, A) = c(A)a, \tag{7}$$

where $c\,' > 0$, reflecting the notion that it is more costly to add to reserves the larger is the stock of past additions. This feature is intuitive: one expects those deposits that are better or cheaper to produce from to be included into the firm's portfolio earlier.

We make two further refinements to the model above. The first assumption reflects the likely feature that oil will eventually be replaced by some other resource.[5] The second assumption is closely related to the physical constraints we articulated above. In general, the rate of extraction of fluids from an oil deposit is limited by the volume associated with the installed well capacity. The extracted fluids include both oil and water, the latter increasing in importance over time in conjunction with the exponential decline in oil production. But to the extent that the cost of operation is tied to the cost of lifting fluids, this cost is fixed at a point in time. Accordingly, we assume production costs at any moment where the firm is actively extracting are equal to F. Altogether, then, the firm's flow payoff is

$$\omega(a, A, R) = p_t q_t - F - c(A_t)a_t. \tag{8}$$

---

4 Our model is similar to that of Venables (2011), whose focus is on the price path, and who does not embed Darcy's Law into his analysis.

5 This is often termed the *backstop technology*, which seems highly likely to apply to oil markets: Speaking of the market's future, the former Saudi oil minister Sheik Yamani famously said "the Stone Age came to an end, not because we had a lack of stones, and the oil age will come to an end not because we have a lack of oil" (Fagan, 2000).

The current-value Hamiltonian-Lagrangean associated with the expanded dynamic optimization is then

$$\mathcal{L}_t = p_t q_t - F - c(A_t)a_t + \mu_t(a_t - q_t) + \nu_t a_t + \lambda(\delta R_t - q_t) + \gamma q_t + \phi a_t.$$

We now have two control variables, q_t and a_t and two state variables, R_t and A_t. Accordingly, the maximum principle consists of the four equations

$$\frac{\partial \mathcal{L}}{\partial q} = p_t - \mu_t + \lambda + \gamma = 0; \tag{9}$$

$$\frac{\partial \mathcal{L}}{\partial a} = -c(A_t) + \mu_t + \nu_t + \phi = 0; \tag{10}$$

$$\dot{\mu}_t = r\mu_t + \delta\mu_t; \tag{11}$$

$$\dot{\nu}_t = r\nu_t + c\,'(A_t)a_t; \tag{12}$$

together with the two state equations (2) and (6). The solution also must satisfy the complementary slackness conditions:

$$\frac{\partial \mathcal{L}}{\partial \lambda} = \delta(R_t + a_t) - q_t \geq 0, \quad \lambda \geq 0, \quad \lambda\frac{\partial \mathcal{L}}{\partial \lambda} = \lambda[\delta(R_t + a_t) - q_t] = 0; \tag{13}$$

$$\frac{\partial \mathcal{L}}{\partial \gamma} = q_t \geq 0, \quad \gamma \geq 0, \quad \gamma\frac{\partial \mathcal{L}}{\partial \gamma} = \gamma q_t = 0; \tag{14}$$

$$\frac{\partial \mathcal{L}}{\partial \phi} = a \geq 0, \quad \phi \geq 0, \quad \phi\frac{\partial \mathcal{L}}{\partial \phi} = \phi a_t = 0. \tag{15}$$

When the firm chooses to add to reserves, a singular solution applies ($\phi - 0$):

$$-c\,'(A_t)a_t + \dot{\mu}_t + \dot{\nu}_t = 0.$$

Substituting in from the co-state equations, we have

$$0 = -c\,'(A_t)a_t + r(\mu_t + \nu_t) - \delta\lambda_t + c\,'(A_t)a_t = r(\mu_t + \nu_t) - \delta\lambda_t.$$

Substituting in for $\lambda_t = p_t - \mu_t$ yields

$$0 = r(\mu_t + \nu_t) - \delta(p_t - \mu_t).$$

Time-differentiating this expression and substituting for $\dot{\mu} + \dot{\nu}$ then gives:

$$
\begin{aligned}
rc'(A_t)a_t &= \delta\big(\dot{p}_t - \dot{\mu}_t\big), \text{or} \\
a_t &= \Big(\frac{\delta}{rc\,'(A_t)}\Big)\Big[\dot{p}_t - rp_t + (r+\delta)\lambda_t\Big]. \qquad (16)
\end{aligned}
$$

There are also transversality conditions, which are influenced in part by the backstop price. The backstop becomes viable when market price has risen to $\overline{p}$; write the time this happens as $\overline{T}$. At this point in time, the firm stops adding to reserves,[6] and so its extraction path from that moment forward is dictated by the decline curve. The firm ceases active production altogether at time T^*, when flow profits are zero; accordingly, it leaves an amount $R_{T^*} = F/(\delta\overline{p})$ of its reserves unextracted.[7]

Let the values of R and A at time $\overline{T}$ be $\overline{R}$ and $\overline{A}$, and write the associated PDV of profits be $V(\overline{R}, \overline{A})$:

$$
\begin{aligned}
V(\overline{R}, \overline{A}) &= \int_{\overline{T}}^{T^*} [\overline{p}\delta\overline{R}e^{-\delta(t-\overline{T})} - F]e^{-rt}dt \\
&= e^{-r\overline{T}} \int_0^{T^*-\overline{T}} [\overline{p}\delta\overline{R}e^{-(r+\delta)s} - Fe^{-rs}]ds \\
&= e^{-r\overline{T}} \Bigg\{ \delta\overline{p}\overline{R}\Big[\frac{1-e^{-(r+\delta)(T^*-\overline{T})}}{r+\delta}\Big] - F\Big[\frac{1-e^{-r(T^*-\overline{T})}}{r}\Big]\Bigg\}, \qquad (17)
\end{aligned}
$$

where we have changed variables in the integral from t to $s = t - \overline{T}$.

Now, the transversality conditionfor the terminal time T^* implies

$$
F = \delta\overline{p}\overline{R}e^{-\delta(T^*-\overline{T})}.
$$

6 Notice that the right side of eq. (16) become negative when $\dot{p} = 0$, which implies that $a > 0$ is no longer optimal.

7 Formally, the transversality condition at time T^* is

$$
H_{T^*} = 0 \Rightarrow \overline{p}q_{T^*} - F - c(A_{T^*})a^*_{T^*} = 0.
$$

Since $q_{T^*} = \delta R_{T^*}$ and $a_{T^*} = 0$, it follows that $R_{T^*} = F/(\delta\overline{p})$.

Substituting this into eq. (17), we get

$$
\begin{aligned}
V(\overline{R},\overline{A}) &= e^{-r\overline{T}}\delta\overline{p}\overline{R}\left\{\frac{1-e^{-(r+\delta)(T^*-\overline{T})}}{r+\delta} - e^{-\delta(T-\overline{T})}\left[\frac{1-e^{-r(T-\overline{T})}}{r}\right]\right\} \\
&= e^{-r\overline{T}}\delta\overline{p}\overline{R}\left\{\frac{1-e^{-(r+\delta)(T^*-\overline{T})}}{r+\delta} - \frac{1-e^{-(r+\delta)(T^*-\overline{T})}}{r} + \frac{1-e^{-\delta(T^*-\overline{T})}}{r}\right\} \\
&= e^{-r\overline{T}}\delta\overline{p}\overline{R}\left\{\delta\left[1-e^{-(r+\delta)(T^*-\overline{T})}\right] + \frac{1-e^{-\delta(T^*-\overline{T})}}{r}\right\}. \qquad (18)
\end{aligned}
$$

The firm's optimal program prior to time $\overline{T}$ must obey transversality conditions for the two shadow values, namely

$$\mu_{\overline{T}} = \partial V(\overline{R},\overline{A})/\partial\overline{T}; \qquad (19)$$

$$\nu_{\overline{T}} = \partial V(\overline{R},\overline{A})/\partial\overline{A} = 0. \qquad (20)$$

3. Market Equilibria

In this section we describe the market equilibrium for oil, based on the model previously described. We start by focusing on the first two periods of a multi-period story. In this setting, we refer to the current time frame as "period 0" and the future as "period 1." At the start of a given period t, firms collectively have remaining reserves r_t^{con}, where the superscript indicates that these reserves have been developed using conventional technology. Of these reserves, at most a given proportion δ can be extracted. Under a wide range of conditions, this constraint will bind, allowing us to focus on the role played by current 'developed' reserves. To produce a greater level of output, new reserves must be added to the portfolio, a process that involves exploration and development. Here we do not describe these steps in any detail; rather, we focus on the impact these 'additions' a^{con} will have on the problem.

Based on the discussion above, we assume the cost of adding to reserves is given by $c(A^{\text{con}})a^{\text{con}}$, where A^{con} is cumulative additions at the current time, with $c\,'(A^{\text{con}}) > 0$. In deciding what level of additions to bring forward at time 0, the (discounted) stream of operating profits is compared against the (up front) development costs; this implies a cutoff price, $\hat{p}_0$, that would just generate the requisite stream of profits. Viewed through this lens, one can think of the incremental cost of bringing forward greater levels of additions as comprising an increasing supply schedule. We denote the available reserves in period t (after the new additions are brought on line) as $R_t^{\text{con}} \equiv r_t^{\text{con}} + a_t^{\text{con}}$; based on these available reserves, output is $Q_t^{\text{con}} = \delta R_t^{\text{con}}$. Within the period, the market equilibrium price equates the incremental cost associated with this supply schedule to the willingness to pay for that last barrel of oil, as reflected by the market demand curve. Writing the optimal level of new additions in period 0 be a_0^{con}, total output delivered to market is then $Q_0 = \delta(r_0^{\text{con}} + a_0^{\text{con}})$. The market-equilibrium price corresponding to this level of output, illustrated in panel (a) of Figure 6, is p_0.

Now consider the next period. The remaining reserves at the start of the period are $r_1^{\text{con}} = (1 - \delta)R_0^{\text{con}}$, the fraction of period-0 available reserves that was not extracted. Again, firms may add to these reserves, up to the point where the last barrel added brings a profit stream that just covers the current development cost. Because the depletion effect associated with period 0 production reduces the amount that can be produced in period 1, the supply curve shifts in (leftward) between periods 0 and 1.[8] As a result, the market equilibrium price increases from p_0 to p_1. This point is illustrated in panel (b) of Figure 6.

In this framework, imagine firms discover the potential to add to reserves via a new technology, for expositional concreteness we refer to this new technique as "enhanced oil recovery," or EOR. This technique is similar to additions, in that it raises the available stock of reserves, but it is cheaper to incorporate. The increased production facilitated by adoption of EOR induces an outward tilting of the supply curve, above some new threshold price; panel (c) of Figure 6 illustrates. With this new supply curve, the market equilibrium price falls and equilibrium quantity rises.[9]

---

[8] In addition, the time horizon is shorter (the backstop is more imminent) in period 1, so there is a shorter time frame to enjoy these profits, which raises the price required to motivate development must be greater than it was in period 0. This effect is of second-order importance in our scenario, and so we abstract from it in the pursuant discussion.

[9] This reduction in price obtains so long as demand is not perfectly elastic.

Importantly, the increase in market quantity implies that the new output associated with EOR must more than offset the reduction in output from conventional sources that are not brought to market in period 0. That is, the oil production 'displaced' by EOR is smaller than the increase due to EOR, contrary to arguments that have been made. Note too that this 'displaced' production is associated with new additions that are no longer economic as a result of the lower price that follows naturally from the increase in supply arising from EOR. While these additions are not brought on line in period 0, they are nevertheless still available in the future.

Because of the upward-sloping nature of the incremental cost curve for additions, some additions to conventional developed reserves remain economic even with the lower price; call the corresponding optimal level of additions ${a'}_0^{\text{con}}$. The net effect on output in period 0 is an increase, from

$$Q_0^{\text{con}} = \delta R_0^{\text{con}} = \delta(r_0^{\text{con}} + a_0^{\text{con}})$$

to

$${Q'}_0^{\text{con+EOR}} = \delta {R'}_0^{\text{con+EOR}} = \delta(r_0^{\text{con}} + {a'}_0^{\text{con}} + a_0^{\text{EOR}}).$$

At the start of period 1, remaining reserves are now

$${r'}_1^{\text{con+EOR}} = (1-\delta){R'}_0^{\text{con+EOR}} = (1-\delta)(r_0^{\text{con}} + {a'}_0^{\text{con}} + a_0^{\text{EOR}}),$$

and so if without any further additions, supply would equal

$$\delta {r'}_1^{\text{con+EOR}} = \delta(1-\delta)(r_0^{\text{con}} + {a'}_0^{\text{con}} + a_0^{\text{EOR}}).$$

This supply level corresponds to the intercept of the supply curve labeled $S_1^{\text{con+EOR}}$ in panel (d) of Figure 6. The supply curve labeled S_1^{con} represents the component of overall supply that is provided by conventional sources, and has intercept

$$\delta {r'}_1^{\text{con}} = \delta(1-\delta)(r_0^{\text{con}} + {a'}_0^{\text{con}}).$$

Because at p'_0, demand exceeds $\delta r'^{\text{con+EOR}}_1$, the new market equilibrium entails a higher price p'_1. To avoid clutter, we have set this price to correspond with the original period-0 price in the absence of EOR, p_0.[10]

At this price $p'_1 = p_0$, further additions to both conventional and EOR reserves become economic and are brought into production. Importantly, by our convenient assumption that $p'_1 = p_0$, the period-1 additions to conventional reserves that become economic, a'^{con}_1, are precisely those additions that would have been economic in period 0 in the absence of EOR, but became uneconomic as a result of the price drop from p_0 to p'_0. That is, $a'^{\text{con}}_1 = a^{\text{con}}_0 - a'^{\text{con}}_0$. It follows that the conventional production out of these additions that would have been supplied in period 0, but was displaced, becomes part of overall production in period 1. The displacement is therefore only temporary: in this case it merely involves a one-period delay.

More generally, the ultimate impact of EOR upon cumulative oil production depends on the impacts in multiple periods. Panel (a) of Figure 7 compares the oil-price path with only conventional production, p^{con}_t, to that with added production from EOR, $p'^{\text{con+EOR}}_t$. Let T^{con} and $T^{\text{con+EOR}}$ denote the times at which the oil price reaches $\overline{p}$ in the scenario without and with EOR, respectively. Under both scenarios, the oil price increases over time. However, the addition of EOR production raises production in each period, as illustrated in panel (b) of Figure 7. In every such period, the increased production implies a lower price, and so the oil-price path $p'^{\text{con+EOR}}_t$ lies below the oil-price path p^{con}_t in panel (a). Eventually, all projects that are economic without the introduction of EOR are also economic when EOR is introduced, so that the accumulated amount of oil produced from conventional sources is unaffected (*i.e.*, the area under the oil-quantity paths Q^{con}_t and Q'^{con}_t are identical).

At a certain point, each path hits a price ceiling $\overline{p}$, above which oil demand drops to zero. This is because $\overline{p}$ represents the price of a renewable, "backstop" technology that is a perfect substitute for oil. At this price, oil production may still continue for some time, but no additions to reserves are made. Since the path $p'^{\text{con+EOR}}_t$ lies below the path p^{con}_t, it follows

[10] Of course, this scenario is coincidental. But our central point, that the net effect of introducing EOR will only be neutral in terms of the impact on ultimate oil production under very special circumstances, does not depend on this graphical assumption.

that $T^{\text{con+EOR}} > T^{\text{con}}$. In the example illustrated in panels (a) and (b) of Figure 7, where the backstop price does not change over time, $T^{\text{con}} = 4$, whereas $T^{\text{con+EOR}} = 8$.

If, on the other hand, the backstop price falls over time, matters are more complex.[11] In this scenario, the outward shift in the oil price means there is more time for the backstop price to decline, the implication of which is that the backstop is reached at a lower price. Panel (c) of Figure 7 illustrates this point. This lower 'terminal' oil price in turn induces a reduction in cumulative additions (*i.e.*, the area under the oil-quantity path $Q_t'^{\text{con}}$ is smaller than the area under the oil-quantity path Q_t^{con}). Ultimately, the net effect on total production turns on a comparison of this induced reduction on the one hand with the increased output associated with EOR on the other.

4. CONCLUSION

In this paper, we have proposed a new model of natural resource extraction that pays attention to important physical constraints on production. In this model, firms are commonly constrained as to the amount they can extract; to add to production they must open up new fields, or expand production capabilities from existing fields (for example by drilling new wells). In this setting, prices rise not because firms choose to reduce their extraction levels, but because the cost of bringing new fields on line naturally rises with cumulative development.

We also investigated the impact of the introduction of a new technology, which we term "EOR," upon oil markets, paying specific attention to the interplay of supply and demand across time. A key result is that the introduction of EOR displaces oil production early on, as the lower prices which result from a supply expansion render some potential additions uneconomic. But over time, as depletion occurs and prices rise, the erstwhile displaced projects become economic. Accordingly, the introduction of EOR may not displace any production, though it necessarily will delay in the development of some new sources of production. In

[11] There is, we believe, good reason to expect the backstop price to fall: historically, R&D into technologies associated with renewable resources have produced striking cost reductions over time. If this pattern is indicative of likely trends going forward, then there seems to be good reason to anticipate the backstop price will decline over time.

the end, then, introduction of EOR will necessarily increase the total accumulated amount of oil produced.

If the critical price at which the economy switches from oil to another energy source, which we have referred to as the backstop price, is constant across time then it follows that all production that would have obtained if EOR were never introduced will also (eventually) be brought on line after EOR is introduced. The end result is that the introduction of EOR must necessarily raise total accumulated oil production by the time the backstop is adopted, with the amount of extra production corresponding to the level of oil produced by EOR. By contrast, if the backstop price falls over time, as one might expect if learning occurs over time, and this learning makes the new technology more efficient, then some of the projects that are pushed out of the market after the introduction of EOR will never be developed. The implication is that the increase in total accumulated amount of oil produced is smaller than that described in the previous paragraph.

References

Darcy, H. (1856). *Les Fontaines Publiques de la Ville de Dijon*, Dalmont, Paris.

Fagan, M. (2000). Sheikh Yamani predicts price crash as age of oil ends, *The Telegraph* . Available online at http://www.telegraph.co.uk/news/uknews/1344832/Sheikh-Yamani-predicts-price-crash-as-age-of-oil-ends.html.

Hannesson, R. (1998). *Petroleum Economics: Issues and Strategies of Oil and Natural Gas Production*, Praeger, Ann Arbor, MI.

Hotelling, H. (1931). The economics of exhaustible resources, *Journal of Political Economy* **39**(2): 137–175.

Venables, A. J. (2011). Depletion and development: Natural resource supply with endogenous field opening. OxCarre Research Paper 62, Department of Economics, University of Oxford.

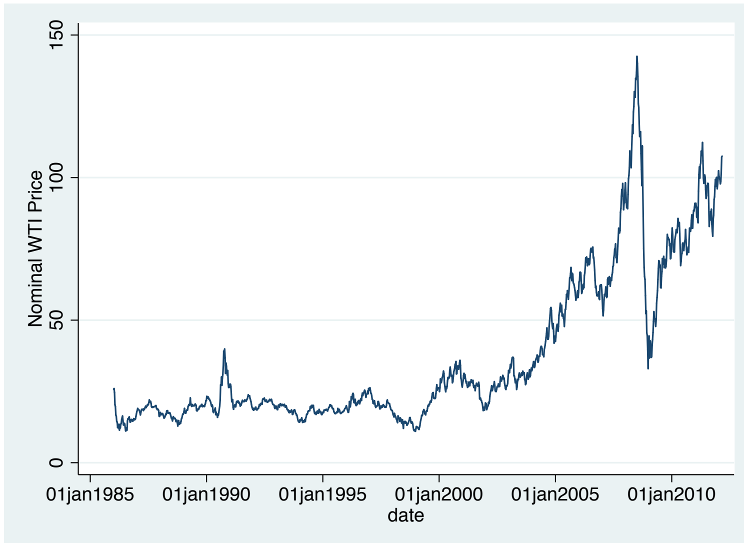

Figure 1. Weekly average crude oil SPot Prices.

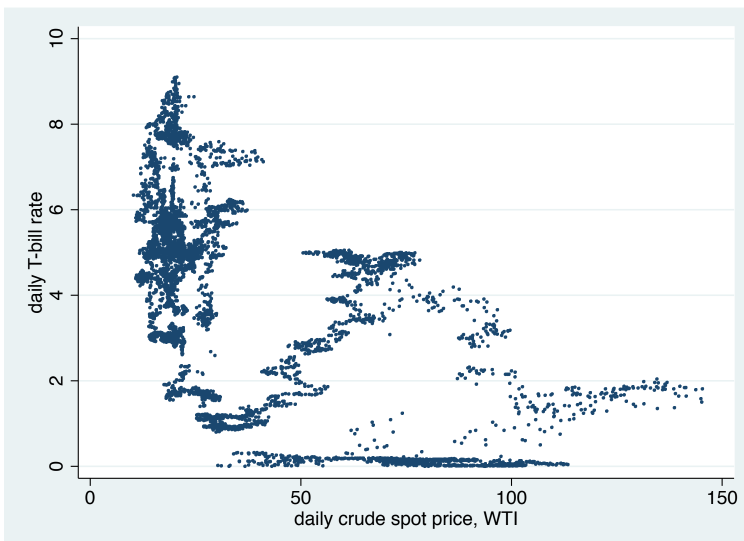

Figure 2. The relation between oil prices and interest rates.

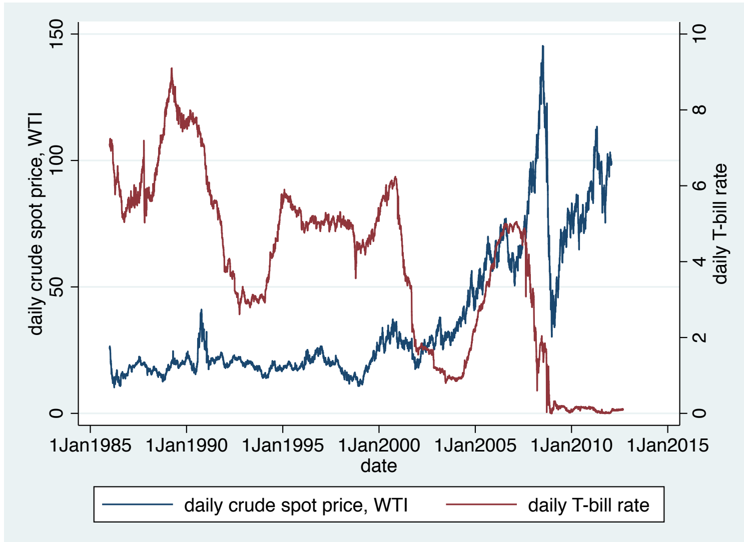

FIGURE 3. Time series of oil prices and interest rates.

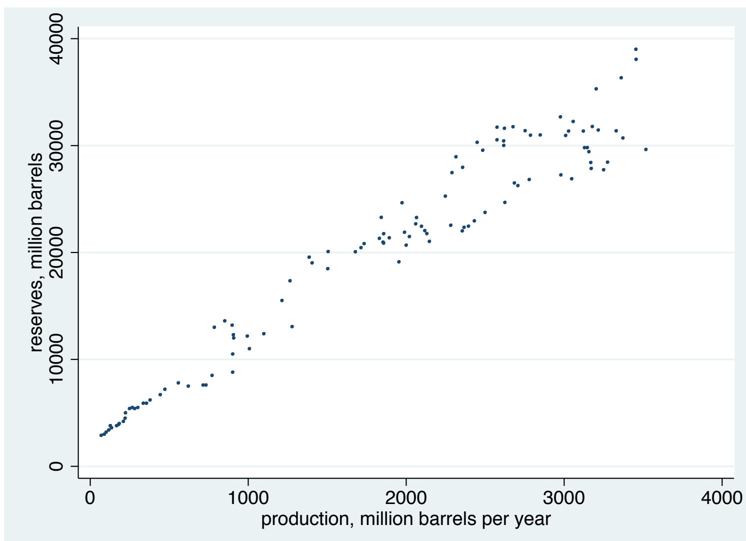

FIGURE 4. US oil production vs. proven reserves.

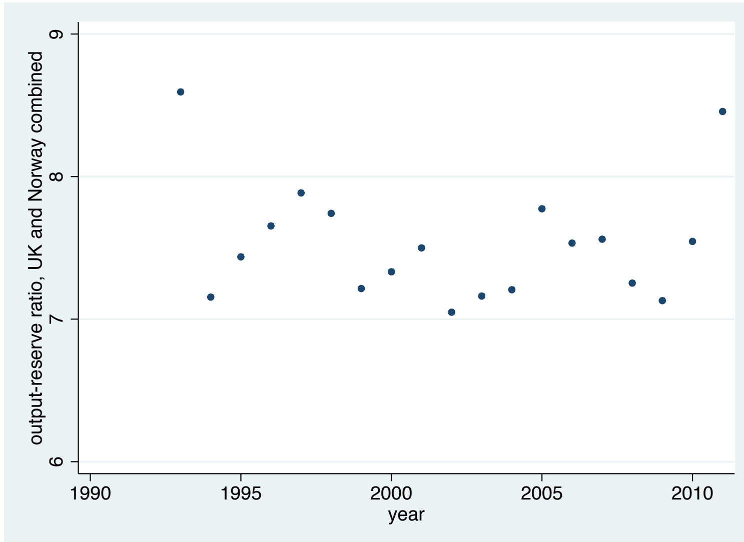

FIGURE 5. Output-reserve ratios over time.

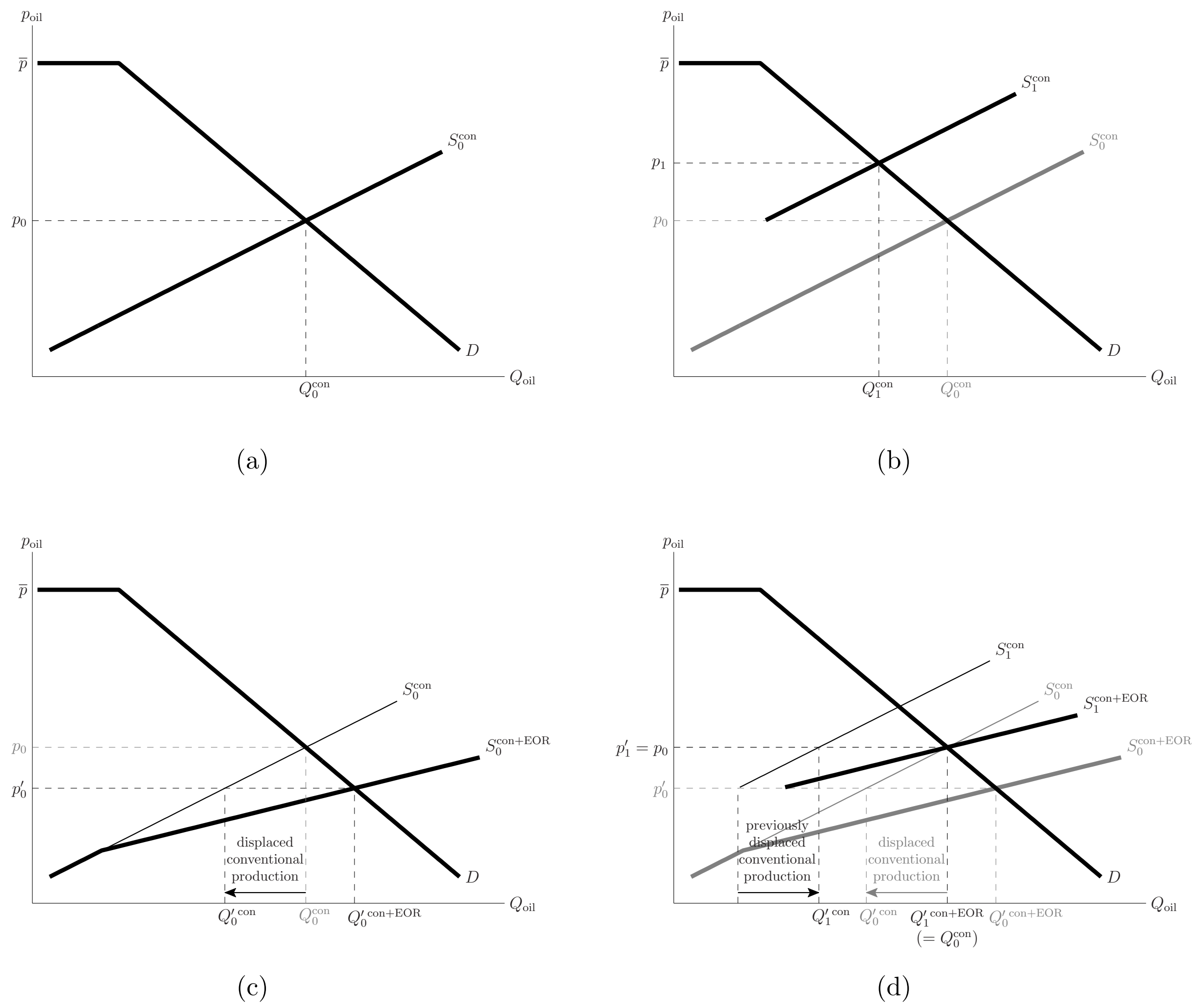

FIGURE 6. Stylized representation the world oil market at two periods in time, without (panels (a) and (b)) or with additional production from EOR (panels (c) and (d)).

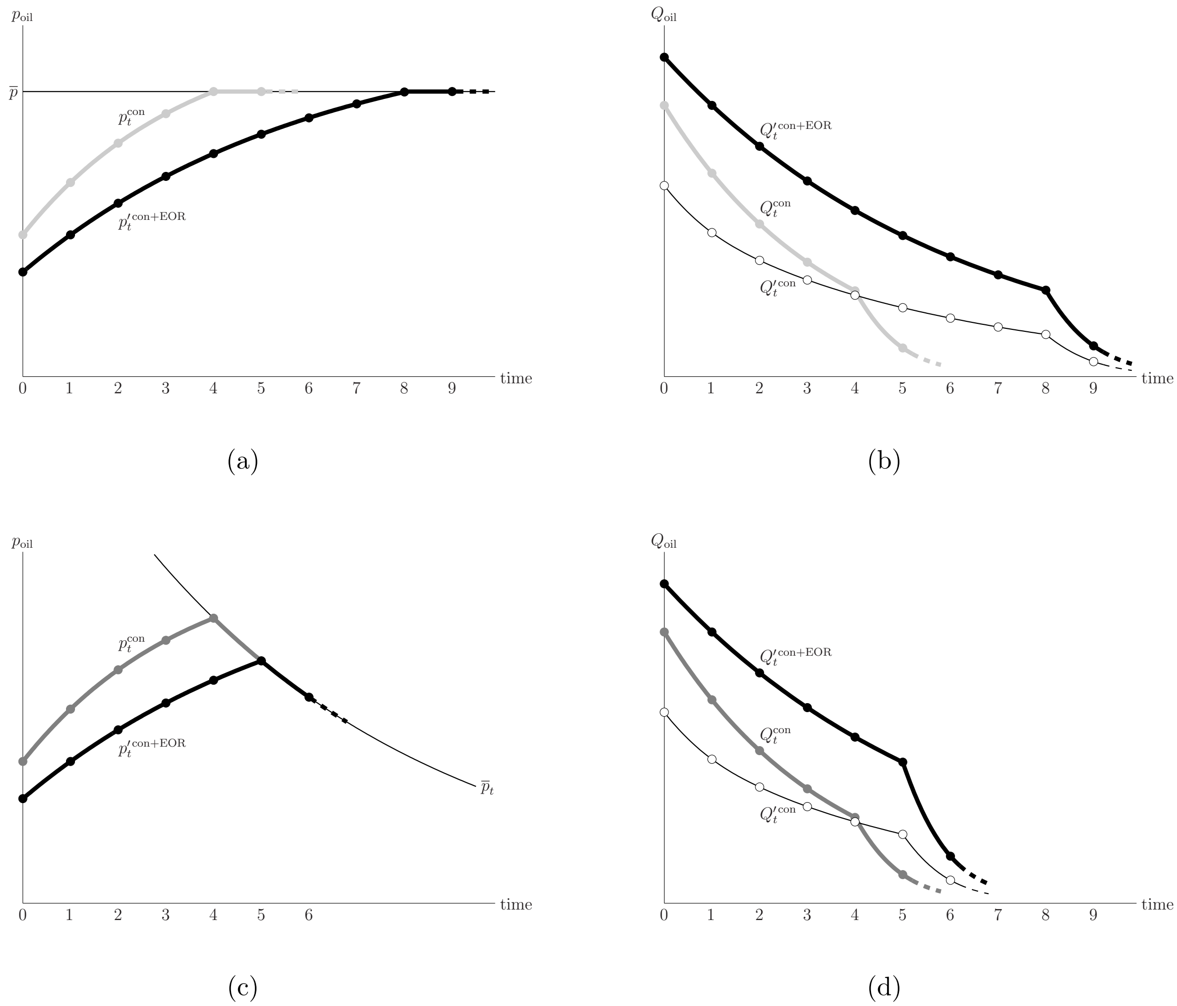

FIGURE 7. Oil-price and oil-production paths with and without EOR, when the price of the backstop technology is either constant (panels (a) and (b)) or falling over time (panels (c) and (d))