An economic model of oil exploration and extraction

Alfred Greiner* Willi Semmler†

Abstract

In this paper we present empirical facts on oil exploitation and a model that can replicate some of these facts. In particular, we show that the time path of the oil price, on the one hand, and the extraction and discovery rate, on the other hand, seem to follow a U-shaped and an inverted U-shaped relationship, respectively, which is confirmed by simple non-parametric estimations. Next, we present a theoretical model where a monopolistic resource owner maximizes inter-temporal profits from exploiting a non-renewable resource where the price of the resource depends on the extraction rate and on cumulated past extraction. The resource is finite and only a part of the resource is known while the rest has not yet been discovered. The analysis of that model demonstrates that the extraction rate and the price of the resource show the empirically observed pattern if the stock of the initially known resource is small.

JEL: Q30, C61

Keywords: Non-renewable resources, optimal control, non-parametric estimation, oil production

*Department of Business Administration and Economics, Bielefeld University, P.O. Box 100131, 33501 Bielefeld, Germany

†Department of Economics, New School for Social Research, New School, 79 Fifth Ave, New York, NY 10003, USA

1 Introduction

Already more than 70 years ago Hotelling (1931) analyzed the problem of how a non-renewable resource should be optimally exploited. This is indeed an important issue from an economic point of view because non-renewable resources, such as oil, gas or ore, are important factors of production and exhaustible. Even if recycling can be used and allows to reuse a certain part of the resources, this does not change the fact that the major part of non-renewable resources is lost in the production process. With a limited amount of resources available, this implies that most of the resources used today are not available in the future.

Although Hotelling presented his model in the 30's of the last century, it was only in the 70's of the last century that his contribution received the attention it deserved. This is not too surprising since it was only in the early 70's that people became aware of the limited availability of natural resources, in particular of oil. Thus, the Club of Rome argued that the process of economic growth cannot continue given a finite amount of natural resources (cf. Meadows, 1972) and the two oil crises in the 70's lead to a drastic rise of the price for that resources. Therefore, economists spent great efforts in finding the optimal rate at which a non-renewable resource should be exploited (see e.g. Dasgupta and Heal, 1979, or Conrad and Clark, 1987).

The basic model presented by Hotelling assumes that the market for the exhaustible resource is perfectly competitive. A representative supplier of the resource solves an intertemporal optimization problem where the problem is to find the optimal rate of extraction given the price trajectory. The solution to that problem, which is equivalent to the social optimum, shows that the price of the resource grows at the interest rate that is used to discount profits. This rule is the so-called Hotelling's rule which is at the heart of the economics of non-renewable resources. While the price of the resource grows, the extraction rate monotonically declines over time.

The basic model presented by Hotelling can be varied in several directions. One obvious extension is to assume that the extraction of resources incurs costs. In this case, the basics of the Hotelling rule dose not change. Hence, in optimum the net price, i.e. price minus marginal cost of extraction, rises at the interest rate. If costs are taken into account, the time path of costs can be crucially as to the price path of the resource. If, for example, marginal costs decline over time due to technical progress, it may well be that the price of the resource first declines before it rises again (see e.g. Khanna, 2003). A fact, that seems to hold for some resources such as oil for example, as the next section will show. Another possibility to reconcile theory with first declining prices that rise at a later stage is to assume that exploratory efforts build up a stock of reserves that is used to satisfy demand for that resource, as suggested by Pindyck (1978). Then, given some technical conditions the price path first declines before it rises.

An important variation of the basic Hotelling model is to assume that the supplier of the resource can control the price at least to a certain degree. Thus, the 70's of the last century demonstrated that the assumption of perfect competition does not necessarily hold for the oil market. Hence, another variation of the Hotelling model consists in assuming that the market structure is given by a monopoly where the monopolist again maximizes inter-temporal profits. The result to this optimization problem gives the modified Hotelling's rule, stating that the net marginal revenue must rise at a rate equal to the interest rate. In addition, it can be shown that, with a monopolistic market structure, the resource is exhausted at a later point in time compared to the case of perfect competition. This is due to the monopolist offering a lower amount of the resource at a higher price compared to suppliers under perfect competition.

With the rising growth rates of some Asian economies in the last few years, in particular China and India, the demand for some resources has drastically increased, particularly for oil. Therefore, the question of how the price of oil and the production of that resource have evolved has again become a matter of interest both for policy makers as well as economists. In this paper we intend to contribute to that line of research. Thus, we present some facts as concerns the evolution of the oil price and as concerns the extraction and the discovery of new deposits and we present a theoretical model that is able to replicate some of these facts.

The rest of the paper is organized as follows. The next section presents facts on oil exploration and a simple non-parametric estimation of the evolution of the oil price, of U.S. oil production and of the discovery rate. Section 3 presents a theoretical model that can replicate some of the empirical facts. There, we first present the analytical model and, then, resort to numerical simulations in order to gain additional insight in the dynamics of the model. Section 4, finally, concludes.

2 Facts on oil exploration

In this section we summarize some important facts about the production of oil over the last century. The data are taken from BP (2007) and from the Energy Information Administration (2007). We begin with the evolution of the oil price.

2.1 History of the oil price

Figure 1 gives the price for oil over the last 140 years. The figure clearly demonstrates that both the nominal and the real price first declined before they began to rise.

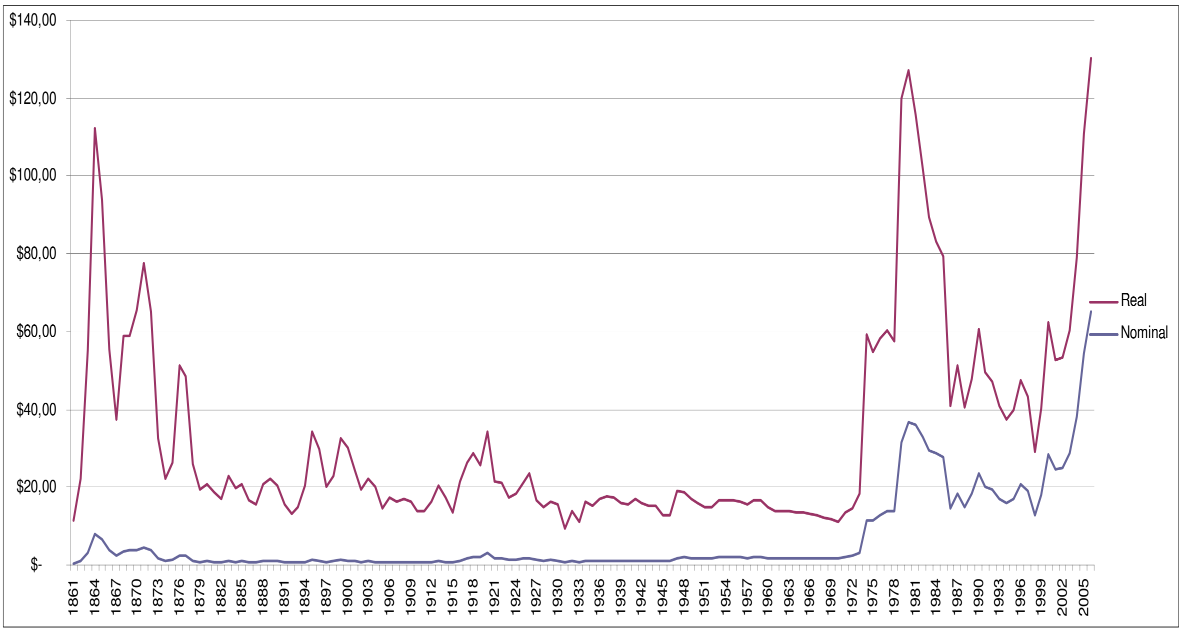

Figure 1: Nominal and real oil price 1861-2005.

Performing a simple p-spline estimation confirms the conjecture of a first decreasing and then increasing time path of the real oil price. The result of the estimation is given in table 2 in appendix A and shows that the non-linear term is highly significant.[1] Figure 2 shows the estimated function with the smoothing parameter set equal to 0.01.[2] The dotted lines give the 95% significance interval and the function is such that its average value is equal to zero.

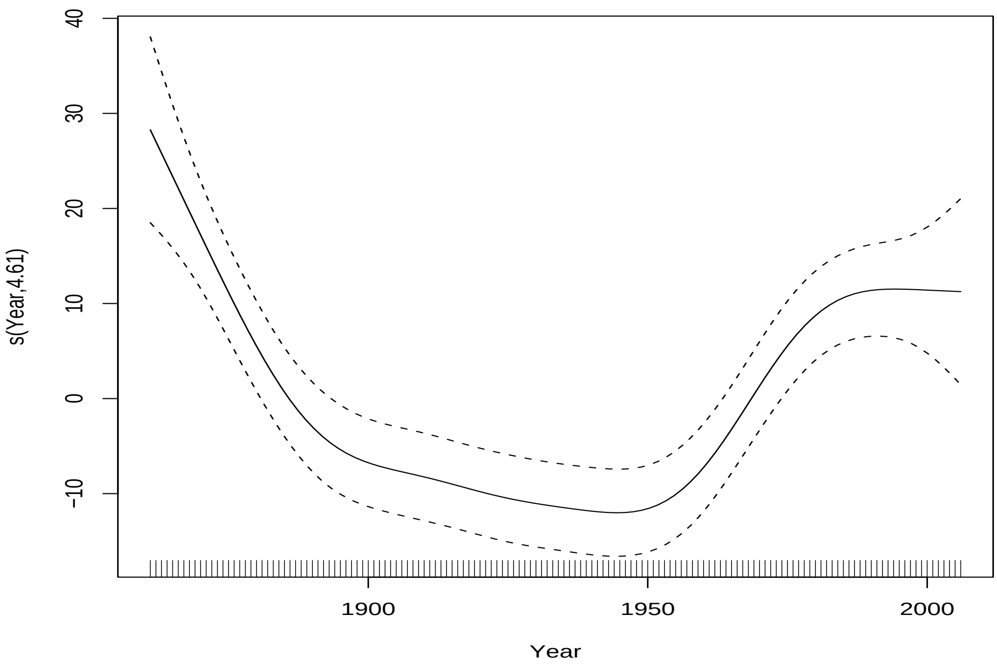

Figure 2: Estimation of the (real) oil price as a function of time.

[1]All estimations were done with the package mgcv, version 1.3-23, in R, version 2.5.0, that can be downloaded from http://www.r-project.org/. For a short introduction into p-spline estimation see Greiner (2008) and a more thorough treatment can be found in Ruppert and Wand (2003).

[2]Selecting the smoothing parameter data driven by resorting to the generalized cross validation criterion would give a value of $8.7\,10^{-5}$ and a more wiggly function without altering the initial decrease and final increase of the estimated function.

2.2 History of oil production and discovery rates

Since oil reserves are finite oil production will decline sooner or later. The Peak-Oil theory, developed by Hubbert (1956), states that oil production first rises and then declines, implying that oil production typically follows a bell-shaped curve. Current estimates of the peak in oil production posit that this point will be reached by about 2010 or has already been reached in 2006 (see e.g. Energy Watch Group, 2007).

Some countries definitely reached their maximal oil production already long ago. In the US, for example, the peak in oil production was reached in 1971, as shown in figure 3, where production first rises and then declines.

Figure 3: Annual US oil production from 1954-2007 (in Thousand Barrels).

Performing p-spline estimation gives the estimated curve as a function of time.[3] Figure 4 shows that the increase in oil production up to 1971 is followed by a decline, except for

[3]For this estimation, the smoothing parameter is selected by applying the GCV criterion. The estimation output is in table 3 in appendix A.

the first half of the 1980's where oil production temporarily rose.

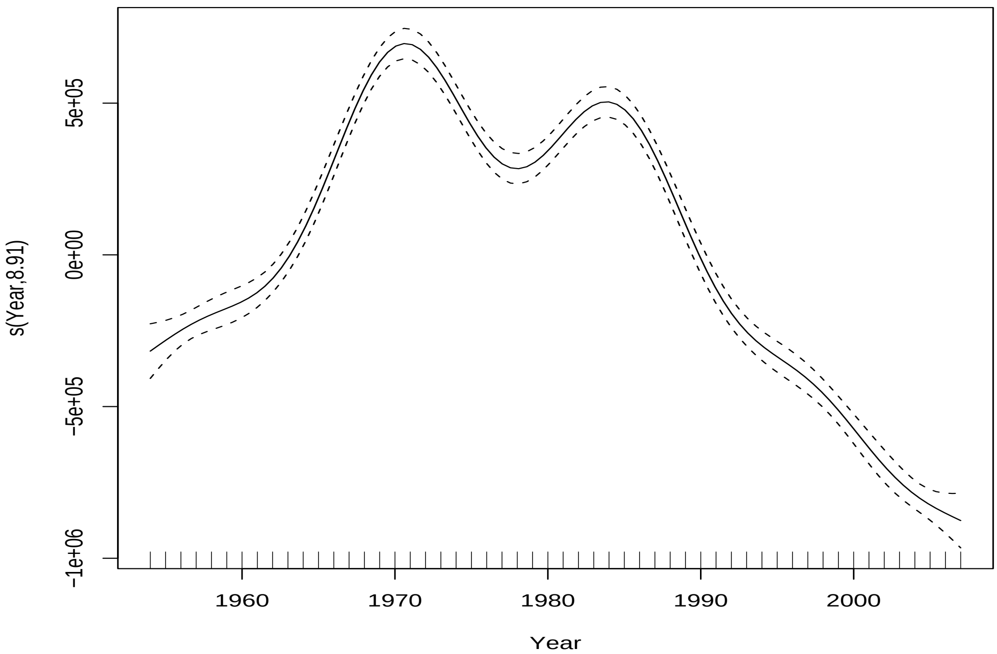

Figure 4: Estimation of the extraction rate as a function of time.

Analogously to the time path of oil production, the discovery rates of new oil fields first rose before they began to decline. Figure 5 shows the evolution of world-wide oil discovery rates starting from 1930 and also gives a projection for the next 40 years.

Again, we perform a p-spline estimation where we estimate the discovery rates as a nonlinear function of time, with data from 1930-2004.[4] As for the oil price and for the extraction rate, figure 6 confirms the picture suggested by the data. The evolution clearly follows an inverted U-shaped relationship.

[4]See table 4 in appendix A for the estimation results. Again, the GCV criterion has been used to select the smoothing parameter.

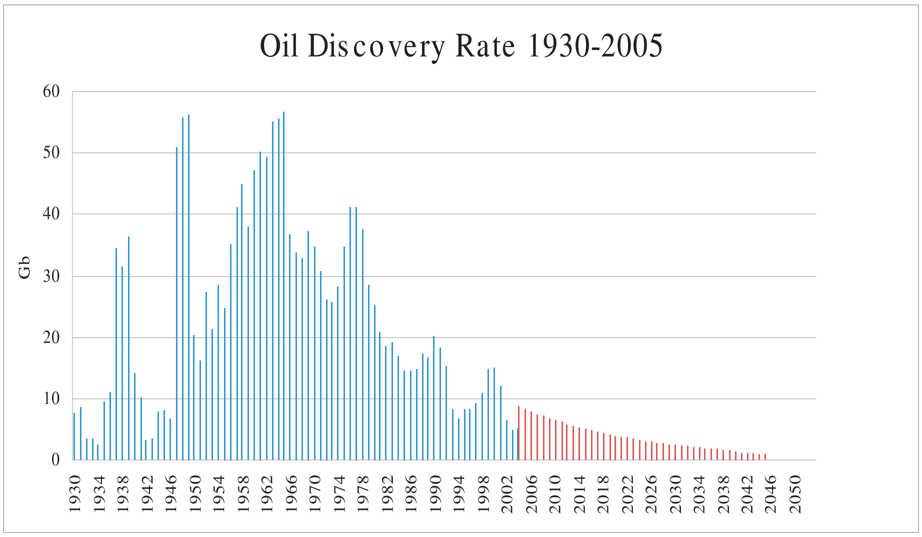

Figure 5: Oil discovery rate 1930-2005; data for 2006-2050 are estimated.

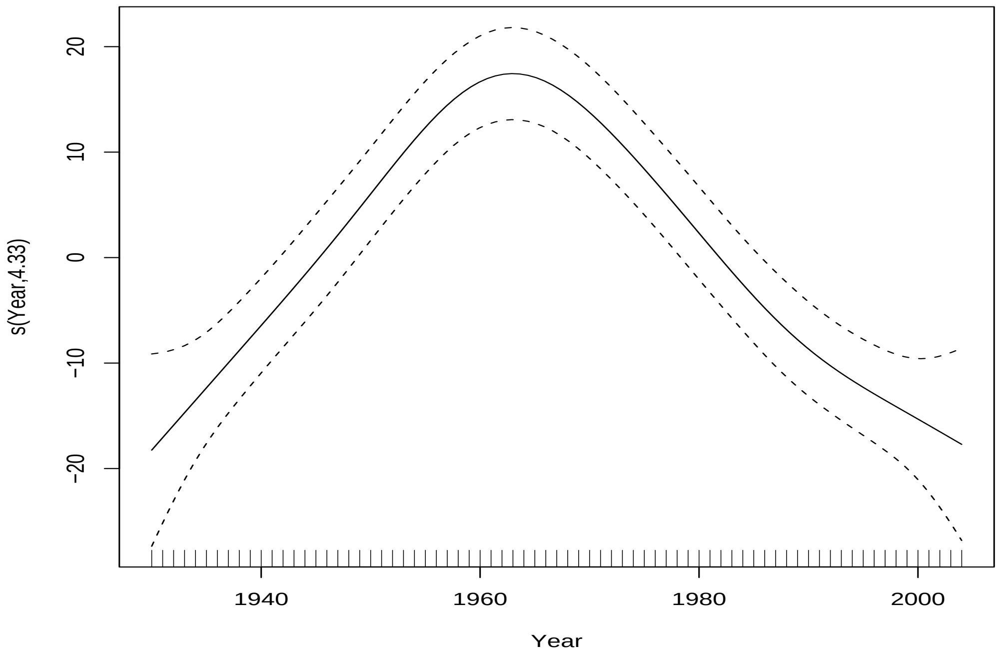

Figure 6: Estimation of the discovery rate as a function of time.

2.3 Current reserves and time to exhaustion

Next we will briefly summarize some facts on the current situation. We will discuss data on total current reserves, distribution of total reserves and time to exhaustion.

Important figures with respect to future oil supply are the annual production, in relation to the remaining reserves, and the time to exhaustion. Table 1 summarizes world petroleum reserves (first row), the ratio of annual production to reserves (second row) and the time to exhaustion in years (third row) for the year 2001.

	World	PG/ME	Iran	Iraq	Kuwait	SA	UAE	VEN	US	Russia
Reserves (% of world)	1.017,73	662,48 (65.09)	99.08 (9.74)	115,00 (11.3)	98,85 (9.71)	261,65 (25.71)	62,82 (6.17)	50,22 (4.93)	21,50 (2.11)	53,86 (5.29)
Prod./Res.	0.024		0.0123	0.00073	0.00067	0.0104	0.0117	0.0189	0.057	0.047
Exhaustion	41.6		81.3	140.1	149.1	96.1	85.5	52.9	17.5	21.0

Table 1: Petroleum reserves (Billion Barrels), ratio of production to reserves and time to exhaustion (years) in 2001 (annual data, PG/ME: Persian Gulf/Middle East, SA: Saudia Arabia, UAE: United Arab Emirates, VEN: Venezuela).

If we take the total oil reserves in the US alone, proven and unproven, at the current rate of extraction of oil through annual production, they will not last for more than a couple of decades. Taking world oil reserves as a whole, at current rates of production there are more than 40 years left before exhaustion. The time to exhaustion for individual countries is shown in figure 7.

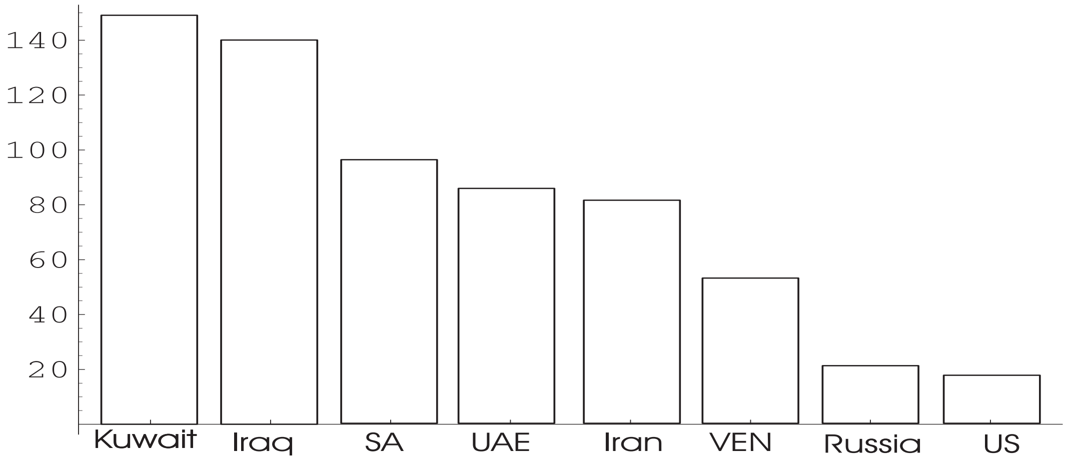

Figure 7: Time to exhaustion in years for 2001 (SA: Saudi Arabia, UAE: United Arab Emirates, VEN: Venezuela).

The important point of figure 7 is that the time to exhaustion of the Middle East oil resources are longest because they are under-exploited. Figure 7 in fact underlines the fact that the reserves of the advanced economies are rapidly shrinking and thus overexploited, while, on the other hand, those of the Middle East are both much larger and significantly under-exploited.

The reserves consist of proved reserves and other reserves (inferred, measured and indicated reserves) which will be added to the proved reserves in the future.[5] Recently, estimates of the recoverable reserves in Alberta, Canada and Venezuela have been greatly increased by using new methods. The resources are now deemed recoverable because there have been great advances in the technology of processing tar and petroleum sands. Furthermore, new oil exploration, for example in the Gulf of Mexico and in Colombia, have come up with significant new findings. However, the costs of recovering petroleum from oil sands may prove daunting; almost as much energy must be used in processing as the final product will contain. Moreover, the processes generate severe pollution. The new strikes in the Gulf are in exceptionally deep waters and have already run into problems.

[5]Of course, the exploration and production technology for petroleum is continuously changing, more rapidly than ever nowadays. This is the result of the so called 'DOFF' - the 'digital oil field of the future'. The DOFF is based on information, exploration and production far more exact and targeted.

Finally, all new unconventional petroleum will be expensive and it will be some time before any comes on line. On the other hand, rapid advances in petroleum exploration and recovery are possible in all oil rich regions. Yet, overall as section 2.2 showed, the discovery rates are unambiguously declining.

Substitutes of coal, natural gas and crude oil are nuclear and renewable power. Since the share of these alternative energy production sources has increased in the last four decades it seems likely that further research and development will enhance this trend.[6]

The above assessments are made under the assumption that only currently available technology is applied. But it seems to be likely that technological progress develops further substitutes and improves mining and refining methods or makes discoveries of new resources fields possible, and, therefore, enhances present estimations of available reserves.[7]

Overall, as we have discussed above the oil discovery rates are shrinking, the total reserves are declining, the time to exhaustion becomes shorter and the oil price shows an upward trend since the 1970's (slightly disrupted by the 1980's).

3 The model

The next issue is whether we can replicate the above mentioned facts by an economic model. To do so we present a model where we assume that the total stock of the resource consists of a certain part that is known and of a certain part that is unknown up to time t but can be discovered and, then, becomes known.

3.1 The analytical model

Assume that the goal is to optimally exploit a non-renewable resource $x(t)$, where t is the time argument. At time t, a certain part of the resource has already been discovered and is known, with $x^k(t)$ denoting the stock of the resource which is known. The rest of

[6] Also, an increased use of particularly plastics as a renewable substitute may relax the constraints that economic growth faces.

[7] Also, if nuclear fusion becomes practical in some decades new energy sources may be added.

the total amount of the resource has not yet been discovered and we denote by $x^n(t)$ the stock of the resource which is not known. Thus, we have $x(t) = x^k(t) + x^n(t)$.

At each point in time a certain part of the hidden resource is discovered with a certain rate $f \geq 0$. We assume that the rate at which the resource is discovered is a function which positively depends on the amount of the resource not yet discovered,[8] i.e. $f = f(x^n)$, with $f'(n) > 0$. This is certainly reasonable because the larger the amount of the resource not yet discovered, the easier it is to find new deposits. The more the resource is exploited the more difficult it becomes to detect new deposits.

The stock of the known resource is exploited and declines at the rate u which gives the amount of exploitation at each point in time. At the same time, it also rises because at each point in time a certain part of the resource is discovered and raises the stock of the known resource. The differential equation describing the time path of the known resource is written as,

$$\dot{x}^k = -u + f(x^n) = -u + f(x - x^k), \tag{1}$$

where we have used $x^n = x - x^k$.

It should be noted that the stock of the unknown resource declines at the rate at which new deposits are discovered. Thus, the evolution of the stock of the unknown resource is given by,

$$\dot{x}^n = -f(x^n)\,. \tag{2}$$

The total stock of the resource evolves according to

$$\dot{x} = \dot{x}^k + \dot{x}^n = -u. \tag{3}$$

Thus, in the long-run, i.e. for $t \to \infty$, $u = 0$ must hold because the resource is non-renewable and because the stock is finite.

Using equation (3) we can rewrite (1) as follows,

$$\dot{x}^k = -u + f\left(x_o - \int_{-\infty}^{t} u(\nu) d\nu - x^k\right) = -u + f\left(x_o - y - x^k\right), \tag{4}$$

[8]In the following we delete the time argument if no ambiguity arises.

with y cumulated past extraction and x_o the initially totally available resource that must exceed past extraction plus the stock of the known resource, i.e.

$$y(t) = \int_{-\infty}^{t} u(\nu)d\nu, \ x_o \geq y + x^k.$$

As to the optimization problem we assume that a monopolistic firm wants to maximize the discounted stream of profits resulting from exploiting the stock of the known resource. The demand curve is given by $p(u, y) > 0$ and is a strictly negative function of the amount of resources supplied at each point in time, u, and positively depends on cumulated past exploitation, y. With respect to u the demand curve is characterized by the usual assumptions, $p_u(\cdot) < 0$ and $p_{uu}(\cdot) < -2p_{uu}(\cdot)/u$. The latter assumption implies that the marginal revenue is a declining function of u. In addition, we posit that the price remains finite, $p(0, y) < \infty$.

As regards the derivative of the demand function with respect to cumulated past extraction, y, we posit $p_y(\cdot) \geq 0$. If the inequality sign is strict the price for the resource is the higher the more of the resource has already been extracted. This assumption is reasonable because, given the finiteness of the resource, the price will be higher when less of the resource is left. If the level of remaining reserves is declining the economic agents will expect declining supply in the future and thus will expect the price to be rising. Through the future markets oil will be bought in order to trade barrels of oil when the oil is in short supply leading to higher demand and a higher price.

The firm incurs costs resulting from exploiting the resource and we posit that it is the more expensive to exploit the resource the smaller the actual stock of the resource is. This holds because at first those deposits are exploited which are less costly. When more and more of the resource has been exploited those deposits have to be exploited which are less accessible and the exploitation of which is associated with higher costs. Thus, it becomes both more difficult to find new deposits as well as more expensive to extract the resource when a large fraction of the resource has already been discovered and exploited.

Costs of extracting the resource are given by $u\,C(x_o - y)$, with $0 \geq C_{(x_o - y)} > -\infty$. This implies that marginal costs are positive, as usual, and that costs are the smaller, the less the resource has been exploited. The latter assumption is intuitively plausible because extraction of resources starts with those deposits which can be exploited less costly. It should be noted that all of the known resource will be extracted if demand for $u = 0$ exceeds extraction costs, i.e. $p(0, y) \geq C(x_z^k) \geq C(x_o - y)$, with $x_z^k \geq 0$ that value of the known resource below which no new deposits are discovered. If this does not hold the resource may not be completely extracted because extraction costs become too high.

Denoting by $r > 0$ the constant discount rate, the optimization problem is formulated as

$$\max_u \int_0^\infty e^{-rt} \left(p(u, y) - C(x_o - y)\right) u \, dt, \tag{5}$$

subject to

$$\begin{aligned} \dot{x}^k &= -u + f\left(x_o - y - x^k\right), \qquad x^k(0) > 0 && (6) \\ \dot{y} &= u, \qquad y(0) \geq 0 && (7) \\ \lim_{t \to \infty} x^k &\geq 0 && (8) \end{aligned}$$

To get insight into the optimal solution, we formulate the current-value Hamiltonian $H(\cdot)$ which is written as

$$H = (p(u, y) - C(x_o - y))u + \lambda_1(f(x_o - y - x^k) - u) + \lambda_2 \, u \tag{9}$$

with λ_1 and λ_2 denoting costate variables or shadow prices of x^k and y.

Assuming that the demand for the resource is sufficiently high such that an interior solution exists, the necessary optimality conditions are obtained as

$$\begin{aligned} \frac{\partial H}{\partial u} &= \lambda_2 - \lambda_1 + u p_u(\cdot) + p(\cdot) - C(\cdot) = 0 && (10) \\ \dot{\lambda}_1 &= r\lambda_1 + \lambda_1 f'(\cdot) && (11) \\ \dot{\lambda}_2 &= r\lambda_2 + \lambda_1 f'(\cdot) - u \, C_{(x_o - y)}(\cdot) - u \, p_y(\cdot) && (12) \end{aligned}$$

Equation (11) shows that the shadow price of the resource grows at a rate larger r since $f'(\cdot) > 0$. For example, if $f(\cdot)$ is linear, λ_1 grows at the rate $r + \xi$, with $\xi > 0$ the slope of the function $f(\cdot)$.

In the following we further specify the function underlying our model. As concerns the demand function, we assume that it is given by

$$p(u, y) = \left(\frac{1}{\gamma + \eta u - \mu y}\right)^\alpha, \ \alpha > 0, \ \gamma > 0, \ \eta > 0, \ \mu \geq 0 \tag{13}$$

and the cost function is

$$C(x_o - y) = (\phi/2)\,(x_o - y)^{-2},\ \phi > 0 \tag{14}$$

The rate at which new deposits of the resource are discovered is linear in its argument,

$$f(x_o - y - x^k) = \xi\,(x_o - y - x^k - x_z^k),\ \xi > 0,\ x_z^k \geq 0 \tag{15}$$

For x_z^k sufficiently large so that $p(0, y) \geq C(x_z^k) \geq C(x_o - y)$ holds,[9] the known resource is completely exploited such that $(x^k)^\star = 0$. The steady state value for y, then, is given by $y^\star = x_o - x_z^k$ and $u^\star = 0$.

In order to gain insight into the transitional dynamics of our model we resort to numerical simulations in the next section.

3.2 Numerical results

Next we use the dynamic programming method by Grüne (1997) and applied by Grüne and Semmler (2004) to study the dynamics of the model with oil extraction. The dynamic programming method can explore the local and global dynamics by using a coarse grid for a larger region and then employing grid refinement for smaller region. Since it does not use first or second order Taylor approximations to solve for the local dynamics, dynamic programming can provide one with the truly global dynamics in a larger region of the state space.[10] The algorithm is explained in appendix B of the paper.

Using this algorithm we first solve the model of section 3 where we set $p_y(\cdot) = 0$, implying that the oil price only depends on the extraction rate, u, regardless of the level of the remaining oil resource. For the numerical solution of the model of section 3 we choose the following parameter values: $\alpha = 2$, $\xi = 0.5$, $r = 0.03$, $\gamma = 0.05$, $\phi = 4$, $\eta = 1$, and $\mu = 0$. Moreover, we set $x_o = 6$ and $x_z^k = 3$.

The numerical solution of the model of section 3.1, using the dynamic programming method of appendix B, gives us the results as shown in figures 8 and 9. Figure 8 shows

[9] This holds for $x_z^k \geq \gamma^{\alpha/2}\sqrt{\phi/2}$. With (14) the inequality must be strict for $C_{(x_o - y)} > -\infty$.

[10] For details on the analysis of why the dynamic programming algorithm is globally significantly more accurate than algorithms using the second order approximations see Becker et al. (2007).

the optimal trajectories in the state space of known, x^k, (horizontal axis) and already exploited, y, (vertical axis) resources and figure 9 gives the corresponding value function.

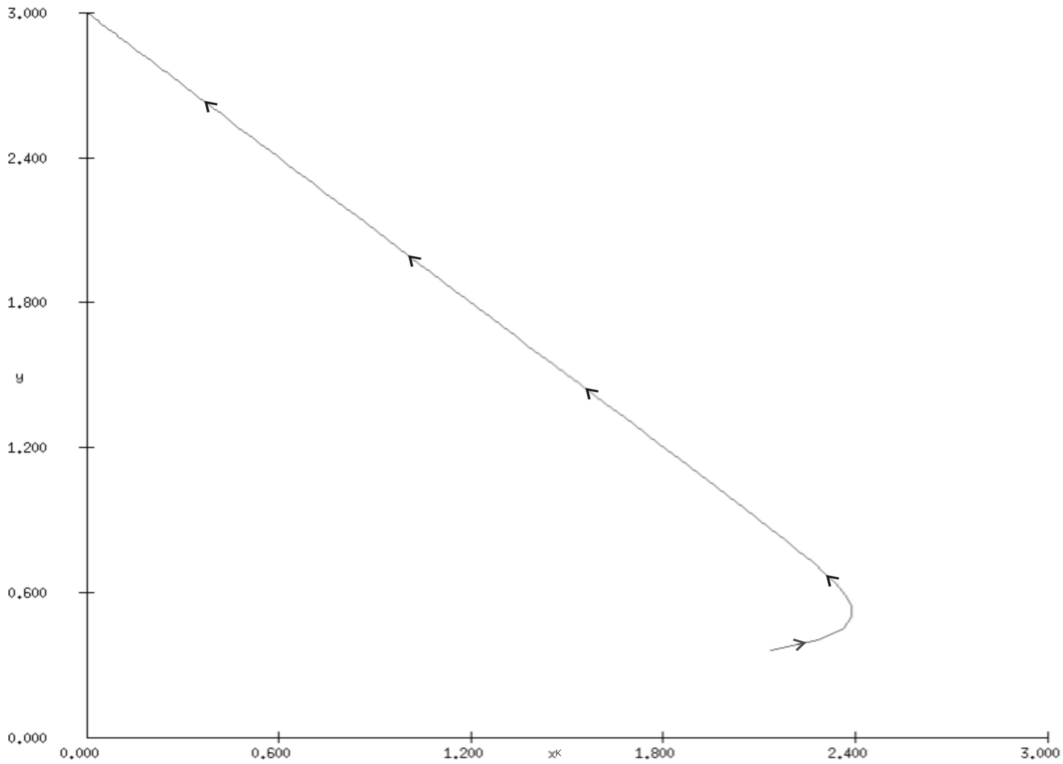

Figure 8: Optimal trajectories for the case of a monotonically decreasing extraction rate and monotonically increasing price.

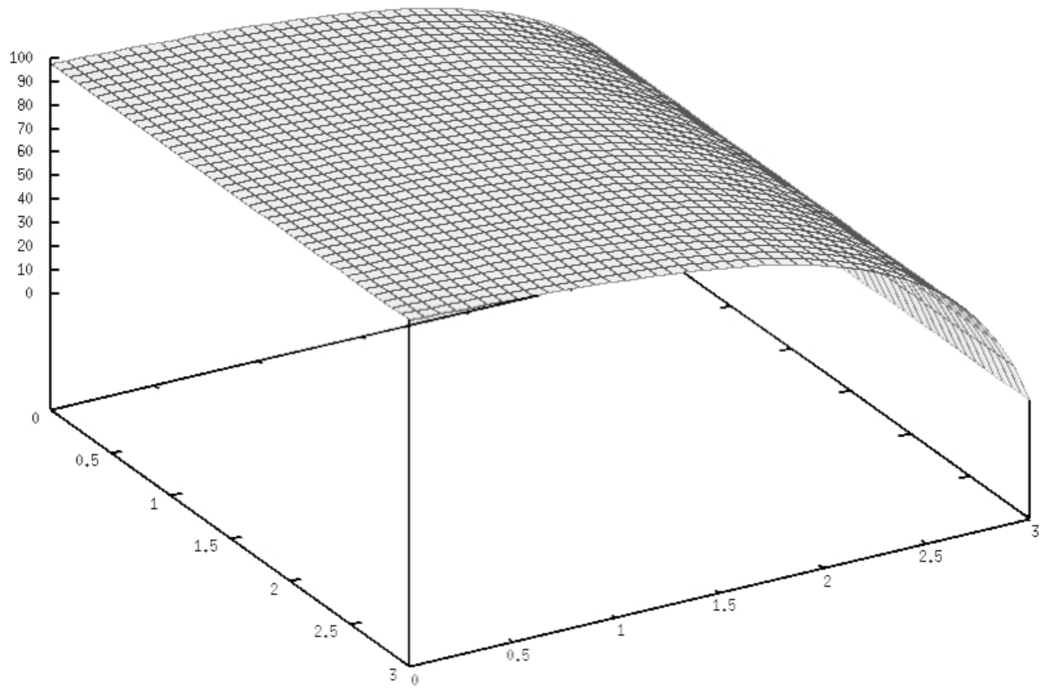

Figure 9: Value function for the case of a monotonically decreasing extraction rate and monotonically increasing price.

Figure 10 shows the trajectory in the $(p-u)$ state space. It can be realized that the oil production u monotonically declines while the price p monotonically rises.

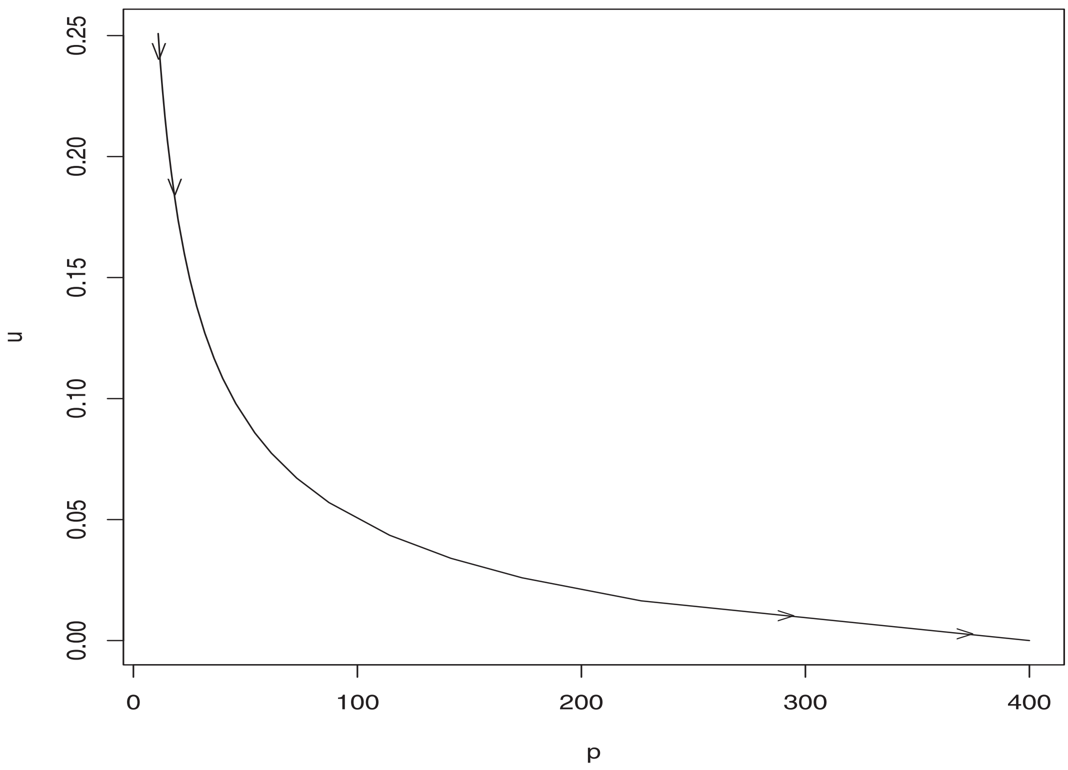

Figure 10: Optimal extraction rate and price in the $(p - u)$ state space.

In a next step we analyze our model assuming that cumulated past extraction positively affects the price of the resource, i.e. $p_y(\cdot) > 0$. To do so, we set $\eta = 4$ and $\mu = 0.05$. In that case, there is the negative effect of the level of already achieved extraction of reserves in the denominator of the demand function (13).

Figure 11 shows the two optimal trajectories in the state space of known, x^k, and already exploited, y, resources. The two trajectories correspond to different initial conditions for the known oil resource, $k^k(0)$. The lower trajectories represents the optimal extraction path for mostly unknown resources, whereas the upper trajectory shows the path of the oil resource when most of the oil resource is known.

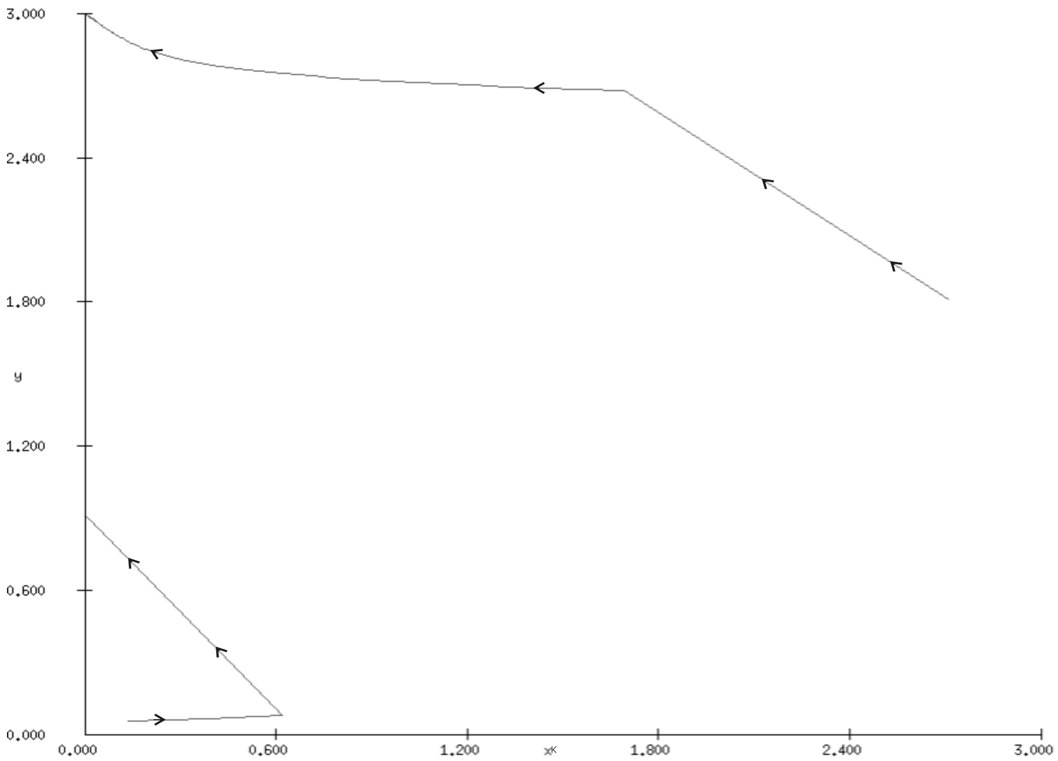

Figure 11: Optimal trajectories in case of a U-shaped extraction rate and price.

Figure 12 depicts the value function over the entire (x^k, y) state space. Now, the value function is quite bumpy reflecting the fact that the additional nonlinearity in the demand function makes the price and thus the pay off moving more irregularly.

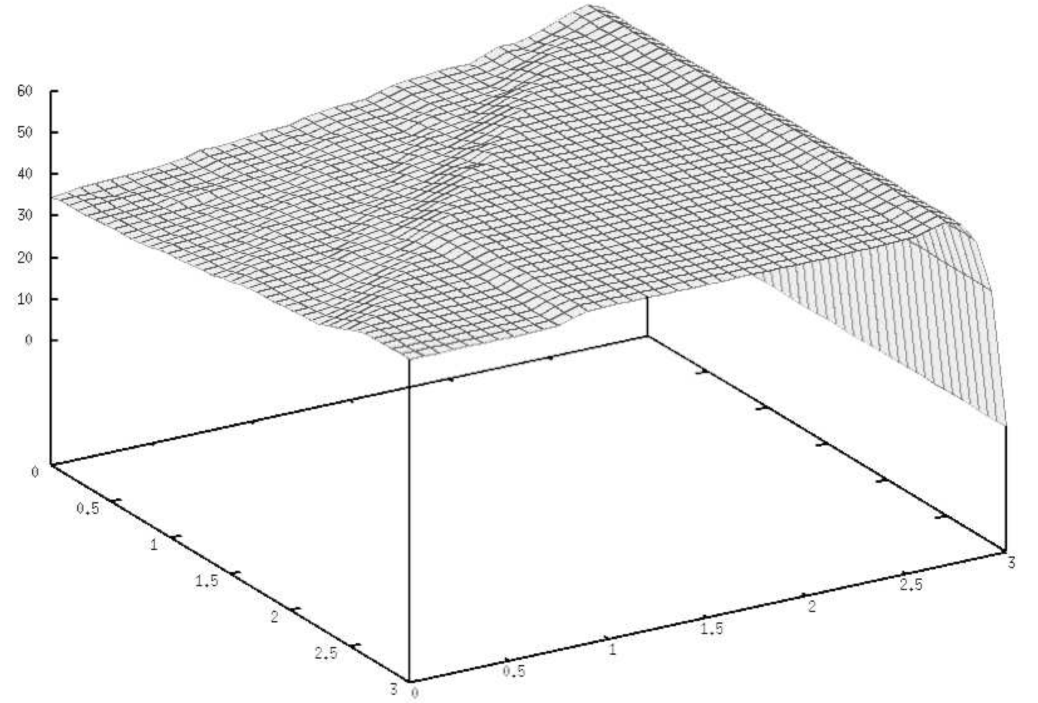

Figure 12: Value function in case of a U-shaped extraction rate and price.

We want to note that the implied time path for the optimal extraction rate of the upper trajectory depicted in figure 11 shows the same monotonic time path as for the model of section 3.1, with only the extraction rate in the demand function. We thus do not need to plot the price movement from it: the price will always increase since the optimal extraction rate monotonically decreases, as does the reaming oil resource.

The more interesting case is the path of the optimal extraction rate u and the corresponding price p when the initial stock of the known resource is small. Then, the optimal extraction rate is hump-shaped, first increasing then decreasing, and the price movement due to the eventually exhausted resource first falls and then rapidly rises, as shown in figure 13 and in figure 14.[11]

The observation that in the case of the upper trajectory of figure 11 the price will monotonically increase is a very plausible scenario since the total stock of the resource

[11]This corresponds to the actual price and oil production which are also hump-shaped, see figures 1-4.

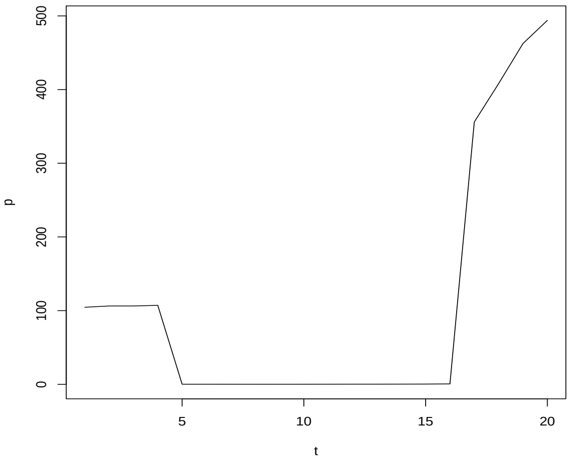

Figure 13: U-shaped price movement.

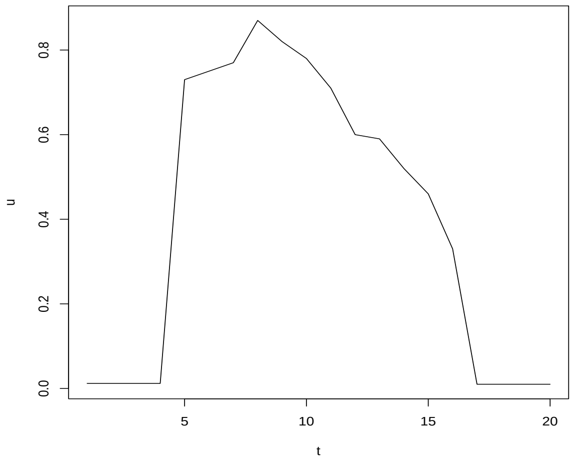

Figure 14: Optimal extraction rate.

is overwhelmingly known and does not have to be discovered. Thus, the function $f(\cdot)$ in equation (1) is not positive or only slightly positive, and the known resource does not rise. Therefore the extraction rate is not likely to rise, but rather to fall. In the case of the lower trajectory the oil resource is not known. It has to be discovered and its discovery adds to the total oil resources which increases first and then decreases. Hereby the optimal extraction rate will rise first and then fall. It is the latter effect that produces the U-shaped price movement as we can in fact also observe in the real data.

4 Conclusions

In this paper we have presented facts on oil exploitation as well as a theoretical model that can replicate some of these facts. In particular, we have shown that the time path of the oil price and the extraction rate seem to follow a U-shaped and an inverted U-shaped time path, respectively.

The model we presented consisted of a monopolistic owner of the resource who knows only a certain part of the total stock of the resource and who discovers new reserves at a certain rate. Assuming that the price of the resource depends on the current extraction rate and on cumulated extraction, we could show that the optimal oil extraction rate may follow an inverted U-shaped time path. The price of the resource in that case first declines and, then, rises again. A prerequisite for that outcome is that the initial stock of the known resource is small. If the stock of resources initially known is large, the extraction rate and the price of the resource follow monotonically declining and rising time paths, respectively.

References

Becker, S., Grüne, L., Semmler, W. (2007), "Comparing Accuracy of Second Order Approximation and Dynamic Programming." *Computational Economics,* Vol. 30, 65-91.

BP (2007) *Statistical Review of World Energy.* London.

Conrad, J.M., Clark, C.W. (1987) *Natural Resource Economics.* Cambridge University Press, Cambridge.

Dasgupta, P.S., Heal G.M. (1979) *Economic Theory and Exhaustible Resources.* Cambridge University Press, Cambridge.

Energy Information Administration (2007) *Annual Energy Report.* Report No DOE/EIA-0384(2007), Washington, D.C.

Energy Watch Group (2007) *Crude Oil. The Supply Shock.* EWG Series No 3/2007, Ottobrunn.

Greiner, A. (2008) "Estimating penalized spline regressions: Theory and application to economics." *Applied Economics Letters.* Forthcoming.

Grüne, L. (1997) "An adaptive grid scheme for the discrete Hamilton-Jacobi-Bellman equation." *Numerische Mathematik,* Vol. 75, 319-37.

Grüne, L., Semmler W. (2004) "Using Dynamic Programming with Adaptive Grid Scheme for Optimal Control Problems in Economics." *Journal of Economic Dynamics and Control,* Vol. 28, 2427-56.

Hotelling, H. (1931) "The economics of exhaustible resources." *Journal of Political Economy,* Vol. 39, 137-75.

Hubbert M.K., (1956) "Nuclear Energy and the Fossil Fuels." *Drilling and Production Practice,* American Petroleum Institute, San Antonio.

Khanna, N. (2003) "On the economics of non-renewable resources." in John M. Gowdy (ed.) *Economics Interactions with other Disciplines,* in *Encyclopedia of Life Support Systems (EOLSS)*, Oxford [http://www.eolss.net]

Meadows, D.L. (1972) *The Limits to Growth: A Report for the Club of Rome's Project on the Predicament of Mankind.* Universe Books, New York.

Pindyck, R.S. (1978) "The optimal exploitation and production of nonrenewable resources." *Journal of Political Economy,* Vol. 86, 841-61.

Ruppert R., Wand M.P., Carroll R.J. (2003) *Semiparametric Regression.* Cambridge University Press, Cambridge.

A Estimation results

	Parametric coefficients:		
	Estimate	Stand. error (t-stat)	p-value
(Intercept)	26.101	1.205 (21.65)	$< 2.2 \cdot 10^{-16}$
	Approximate significance of smooth terms:		
	edf	F	p-value
s(Year)	4.607	13.21	$3.82 \cdot 10^{-15}$
	R^2(adj) = 0.385	n = 146	

Table 2: Estimation results producing figure 2.

	Parametric coefficients:		
	Estimate	Stand. error (t-stat)	p-value
(Intercept)	$2.7 \cdot 10^{6}$	8704 (315.3)	$< 2 \cdot 10^{-16}$
	Approximate significance of smooth terms:		
	edf	F	p-value
s(Year)	8.908	305.9	$2 \cdot 10^{-16}$
	R^2(adj) = 0.981	n = 54	

Table 3: Estimation results producing figure 4.

	Parametric coefficients:		
	Estimate	Stand. error (t-stat)	p-value
(Intercept)	23.879	1.192 (20.04)	$< 2 \cdot 10^{-16}$
	Approximate significance of smooth terms:		
	edf	F	p-value
s(Year)	4.335	11.81	$4.72 \cdot 10^{-11}$
	R^2(adj) = 0.57	n = 75	

Table 4: Estimation results producing figure 6.

B Numerical solution method

We here briefly describe the dynamic programming algorithm as applied in Grüne and Semmler (2004) that enables us to numerically solve our dynamic model variants. The feature of the dynamic programming algorithm is an adaptive discretization of the state space which leads to high numerical accuracy with moderate use of memory.

Such algorithm is applied to discounted infinite horizon optimal control problems of the type introduced for the study of the global dynamics. In our model variants we have to numerically compute $V(x)$ for

$$V(x) = \max_u \int_0^\infty e^{-r} f(x,u)dt$$
$$\text{s.t. } \dot{x} = g(x,u)$$

where u represents the control variable and x a vector of state variables.

In the first step, the continuous time optimal control problem has to be replaced by a first order discrete time approximation given by

$$V_h(x) = \max_j J_h(x,u), \quad J_h(x,u) = h\sum_{i=0}^{\infty}(1-\theta h)f(x_h(i),u_i) \tag{A1}$$

where x_u is defined by the discrete dynamics

$$x_h(0) = x,\ \ x_h(i+1) = x_h(i) + hg(x_i, u_i) \tag{A2}$$

and $h > 0$ is the discretization time step. Note that $j = (j_i)_{i\in\mathbb{N}_0}$ here denotes a discrete control sequence.

The optimal value function is the unique solution of a discrete Hamilton-Jacobi-Bellman equation such as

$$V_h(x) = \max_j\{hf(x, u_o) + (1 + \theta h)V_h(x_h(1))\} \tag{A3}$$

where $x_h(1)$ denotes the discrete solution corresponding to the control and initial value x after one time step h. Abbreviating

$$T_h(V_h)(x) = \max_j\{hf(x, u_o) + (1 - \theta h)V_h(x_h(1))\} \tag{A4}$$

the second step of the algorithm now approximates the solution on a grid Γ covering a compact subset of the state space, i.e. a compact interval $[0, K]$ in our setup. Denoting the nodes of Γ by $x^i, i = 1, ..., P$, we are now looking for an approximation V_h^Γ satisfying

$$V_h^\Gamma(X^i) = T_h(V_h^\Gamma)(X^i) \tag{A5}$$

for each node x^i of the grid, where the value of V_h^Γ for points x which are not grid points (these are needed for the evaluation of T_h) is determined by linear interpolation. We refer to the paper cited above for the description of iterative methods for the solution of (A5). Note that an approximately optimal control law (in feedback form for the discrete dynamics) can be obtained from this approximation by taking the value $j^*(x) = j$ for j realizing the maximum in (A3), where V_h is replaced by V_h^Γ. This procedure in particular allows the numerical computation of approximately optimal trajectories.

In order the distribute the nodes of the grid efficiently, we make use of a posteriori error estimation. For each cell C_l of the grid Γ we compute

$$\eta_l := \max_{k\in c_l} \mid T_h(V_h^\Gamma)(k) - V_h^\Gamma(k) \mid$$

More precisely we approximate this value by evaluating the right hand side in a number of test points. It can be shown that the error estimators η_l give upper and lower bounds for the real error (i.e., the difference between V_j and V_h^Γ) and hence serve as an indicator for a possible local refinement of the grid Γ. It should be noted that this adaptive refinement of the grid is very effective for computing steep value functions and models with multiple equilibria, see Grüne and Semmler (2004).