The Long-Run Evolution of Energy Prices

Robert S. Pindyck*

In this paper I examine the long-run behavior of oil, coal, and natural gas prices, using up to 127 years of data, and address the following questions: What does over a century of data tell us about the stochastic dynamics of price evolution, and how it should be modelled? Can models of reversion to stochastically fluctuating trend lines help us forecast prices over horizons of 20 years or more? And what do the answers to these questions tell us about investment decisions that are dependent on prices and their stochastic evolution?

INTRODUCTION

Energy producers and consumers regularly attempt to forecast prices of oil, coal, and other resources over time horizons as long as twenty or thirty years. Producers make these forecasts for general purposes of strategic planning, and for specific purposes of evaluating investment decisions, e.g., related to resource exploration, reserve development, and production. Industrial consumers, such as petrochemical companies or electric utilities, make these forecasts for the same kinds of reasons—oil, coal, and natural gas are important input costs that can affect investment decisions (e.g., an oil- versus coal-fired power plant for an electric utility), or even the choice of products to produce (e.g., a set of chemicals or the processes used to produce those chemicals).

Ideally, we would like to be able to explain energy prices in structural terms, i.e., in terms of movements in supply and demand, and variables that determine supply and demand. After all, it is these movements that cause prices to fluctuate in the first place. However, structural models are not always useful for long-run forecasting, in part because it is difficult to forecast the explanatory

*   Massachusetts Institute of Technology, Sloan School of Management, 50 Memorial Drive, E52-450, Cambridge, MA 02139-4307, USA. E-mail: rpindyck@mit.edu

variables in such models (e.g., investment and production capacity, inventory levels, and determinants of demand) over long horizons. Structural models are usually better suited at providing insight into the causes of short- or intermediate-run fluctuations of prices and other variables.[1]

As a result, industry forecasts of energy prices over long time horizons are often extrapolations in which prices are assumed to grow in real terms at some fixed rate. The rate of growth might reflect some notion of resource depletion and/or technological change. Sometimes the price is assumed to grow from its current level, consistent with the view that it follows a random walk with drift. Alternatively, the price might be assumed to revert to a trend line that grows or declines. A forecast of this sort would be consistent with the notion that the resource is produced and sold in a competitive market, so that price should revert towards long-run marginal cost, which is likely to change only slowly. This would imply that price shocks are temporary, i.e., that over sufficiently long horizons, prices are mean-reverting rather than random walks.

Whether such approaches to long-run forecasting are reasonable depends on the stochastic process that the price follows. Identifying this process for the price is also important for firms making investment decisions. For example, is it reasonable to model the price process as a geometric Brownian motion or some related random-walk process, or is it better described as a mean-reverting process? This question is relevant because investments are typically irreversible and hence have option-like characteristics. Simple net present value rules are based only on expected future prices—second moments do not matter for NPV assessments of investment projects. But this is not true when investment decisions involve real options, as is the case when the investment is irreversible. Then second moments matter very much, so that an investment decision based on a mean-reverting process could turn out to be quite different from one based on a random walk.[2]

In this paper I examine the long-run behavior of oil, coal, and natural gas prices in the United States without any attempt at structural modelling. Instead, I focus on alternative stochastic processes that might be consistent with this long-run behavior. I first consider whether prices are mean-reverting. I present unit root tests, but argue that even for time series spanning a century, they are likely to be inclusive. Nonetheless, those tests, along with variance ratio tests, suggest that prices are indeed mean-reverting, but the rate of mean reversion is slow (so that for purposes of making investment decisions, one

---

1. On the other hand, short-run fluctuations (over time horizons of a month or two) will be driven in part by changes in inventory demand and supply, which complicates structural modelling. See the discussion in Pindyck (1994).

2. See the discussion in Dixit and Pindyck (1994). Baker, Mayfield, and Parsons (1998) provide examples of different models of the price process and their implications for valuation and investment decisions.

could just as well treat the price of oil as a geometric Brownian motion, or related random-walk process.) Also, the trends to which prices revert are themselves fluctuating over time.

The idea that prices revert to trends that move over time is not new. Perron (1989) developed a stochastic switching model that allows for discrete shifts in the slope or level of the trend line. Applying the model to data on real GNP, he found two events that seem to represent permanent changes in the underlying process, the Great Crash (which shifted the trend line downward) and the 1973 oil price shock (which changed the slope of the trend line). Videgaray (1998) examined the possibility of such a structural change in the context of the price of crude oil. Working with about 120 years of data (the same oil price data that I use here), he applied Perron's method to the estimation of stochastic switching models of price, and found a structural change around 1973. However, such stochastic switching models increase the possibilities for "data snooping" (for any long time series, it is likely that one will always find one or more "structural changes"), and they provide no explanation for what such a structural change means or why it occurred. For example, in the case of Perron's results, it is not clear why a sharp increase in oil prices should have resulted in a one-time change in the slope of the GNP trend line, especially given that by the late 1980s real oil prices had fallen back below their 1972 levels.

As I will show below, the behavior of real energy prices suggests reversion to trend lines with slopes and levels that are both shifting continuously and unpredictably over time, so that each price follows a multivariate stochastic process. The shifts themselves may be mean-reverting, but ignoring them (or assuming instead that there were only a few discrete shifts in the trend line) is misleading, and can lead to suboptimal forecasts. I will also show that a multivariate model with continuous fluctuations in the trend line slope and level is consistent with basic models of exhaustible resource production.

A simple way to incorporate these features is through a Kalman filter model. Here, the trend line slope and level are treated as state variables that evolve stochastically, and that cannot be observed directly. Instead, their (changing) values are estimated recursively over the sample horizon, along with any fixed parameters. Thus, the Kalman filter is a type of time-varying parameter estimator. Its advantage in our context is that it is forward looking, in that state variable estimates are continuously updated as new data become available. The Kalman filter has been applied recently to modelling and forecasting financial asset prices over relatively short horizons; I will argue that it is well suited to energy and other commodity prices over long horizons.[3]

---

3. For applications in finance, see Campbell, Lo, and MacKinlay (1997) and Lo and Wang (1995). For an introduction to the Kalman filter, see Harvey (1989).

In the next section I present data on real oil, coal, and natural gas prices, and show how quadratic trend lines fitted to these data move as the sample period is lengthened. I also present unit root and variance ratio tests that suggest reversion to a trend line, but one that is fluctuating. In Section 3, I show how stochastic fluctuations in the level and slope of the trend are consistent with a basic model of exhaustible resource production. I also lay out a multivariate Ornstein-Uhlenbeck model that incorporates such fluctuations. Section 4 presents Kalman filter estimates of this model for each of the three prices, along with a set of forecasts. Section 5 concludes.

2. MOVEMENTS IN REAL ENERGY PRICES

I examine the real prices of crude oil and bituminous coal over the 127-year period 1870-1996. Natural gas is a much "newer" resource, so data is available beginning only in 1919, a time horizon of about 75 years. The data for nominal oil and coal prices for 1870 to 1973, and for natural gas for 1919 to 1973, were obtained from Manthy (1978) and from the U.S. Department of Commerce, *Historical Statistics of the United States* (1975), and are annual averages of producers' prices in the United States. I updated these series through 1995 using data from the U.S. Energy Information Agency, and for 1996 using data from *The Wall Street Journal*. I then deflated these nominal series to 1967 dollars using the Wholesale Price Index for all commodities, obtained from *Historical Statistics* through 1970, and the Producer Price Index for all commodities from 1970 onwards. Finally, I took the natural logarithm of each of the deflated series.

Figures 1, 2, and 3 show the log real prices of each of the three resources. Observe that both oil prices and coal prices generally fell from 1870 through 1900, a period during which the production of these resources developed on a large scale.[4] From 1900 until the oil shock of the early 1970s, these prices fluctuated considerably, but generally stayed close to an average value (in 1967 dollars) of about $3.50 per barrel for oil, and $4 per ton for coal. This might suggest that these prices are mean-reverting, and we will explore this shortly. From 1973 to 1981, the prices of all three fuels increased dramatically, but by the mid-1980s, oil and coal prices had returned to levels not much above those of thirty to eighty years earlier, while natural gas prices remained high.

---

4. The first commercial oil well went into operation in the United States in 1859 in Titusville, PA.

**Figure 1. Log Price of Crude Oil and Quadratic Trend Lines**

**Figure 2. Log Price of Bituminous Coal and Quadratic Trend Lines**

Figure 3. Log Price of Natural Gas and Quadratic Trend Lines

For each resource, I fit the log real price series to a quadratic time trend, first using all of the data in the sample, and then using data only through 1960, through 1970, and through 1980.[5] (In each case I ran an OLS regression of the log price on a constant, time, and time squared.) Each fitted trend equation was used to forecast prices through the year 2000, and the various trend lines and forecasts are shown in each of the figures. Not surprisingly, these fitted trend lines (and the resulting forecasts) move considerably as the sample period to which they are fit is lengthened. Furthermore, although the magnitudes of the shifts vary, there is no single point in time for any resource at which shifts can be exclusively localized.[6]

These graphs suggest two basic characteristics of long-run price evolution. First, the log real price of each resource seems to be mean-reverting to a quadratic trend line, although the rate of mean reversion is slow, taking up to a decade to occur. Second, the trend line itself fluctuates as the sample is extended. Below, I explore these two characteristics further; later I will discuss how they can be incorporated in a forecasting equation.

---

5. A quadratic U-shaped trend line is consistent with models of exhaustible resource production that incorporate exploration and proved reserve accumulation over time, as well as technological change. See Pindyck (1978, 1980).

6. This may be hard to see from the figures. However, even when the trend line regressions are repeatedly reestimated with one year increments in the sample period, the trend parameters and resulting forecasts continually shift.

Although the graphs appear to exhibit mean reversion, we might test for this by running augmented Dickey-Fuller unit root tests, as is often done. However, it is important to be clear about the limitations of such tests.

For example, as Froot and Rogoff (1995) have illustrated in the context of exchange rates, the identification of a unit root is likely to require data covering very long time horizons. To see this, suppose that the detrended logarithm of price around its mean follows an AR(1) process:

$$p_t = \rho p_{t-1} + \epsilon_t$$

where $0 \leq \rho < 1$, and ϵ_t is a white noise process. Then the asymptotic standard deviation of ρ is given by:

$$\mathrm{s.d.}(\rho) = \left[\frac{1 - \rho^2}{T} \right]^{1/2}$$

where T is the number of observations. Now, suppose we run a Dickey-Fuller unit root test to determine whether $\rho \leq 1$. To reject the hypothesis that $1 - \rho = 0$ at the 5 percent level, we would need a t-statistic on $1 - \hat{\rho}$ of at least 2.89. Thus, we would need

$$(2.89)^2 \leq \frac{T(1 - \rho)^2}{1 - \rho^2}$$

or equivalently, $T \geq (8.352)(1 - \rho^2)/(1 - \rho)^2$.

Now, consider the price of oil. Figure 1 suggests that this price is mean-reverting, and that the half-life of the process is around five years. In this case, the true value of ρ would be $(0.5)^{1/5} = 0.87$. Plugging this into the equation for T, we find that we would need at least 120 years of data to reject a unit root at the 5 percent level. And if prices reverted even a bit more slowly, so that the true value of ρ was 0.9, we would need about 159 years of data. Even if the half-life of the process were three years, so that $\rho = (0.5)^{1/3} = 0.794$, we would need 73 years of data.

I nonetheless ran augmented Dickey-Fuller unit root tests on the full sample of data for each price and on various subsamples of the data. The tests involve regressions of the form:

$$\Delta p_t = (\rho - 1)p_{t-1} + \sum_{k-1}^{N} \alpha_k \Delta p_{t-k} + c_0 + c_1 t, \tag{1}$$

where p is the log real price, and N is the number of lags. I include 1, 2, and 4 lags in these tests, and I also include a time trend to allow for gradual shifts in the "normal" price. The test is based on the MacKinnon (1991) critical values for the t-statistic on $p - 1$.

The results of these tests, along with the estimated values of p, are shown in Table 1. Because of the lagged values of Δp used in the test, the full sample covers 1875 to 1996 (and 1924 to 1996 for natural gas). The test is also run for the subsamples 1900-1996, 1925-1996, 1940-1996, and 1950-1996. Observe that a unit root can be rejected only for oil, and even then only for samples extending back to at least 1900. Of course this failure to reject may be due to the sample size issue raised above. (Note that for coal and natural gas, the estimated values of p are around 0.9 or greater, suggesting that a substantially longer time series would be needed to reject.) Alternatively, these results may be explained by shifts in the slope of the trend line. In any case, a failure to reject a unit root does not imply an acceptance of a unit root; it simply leaves the question open.

Table 1. ADF Unit Root Tests and AR(1) Parameters

Resource	Sample	$N = 1$		$N = 2$		$N = 4$	
		t	ρ	t	ρ	t	ρ
Oil	1875-1996	3.859**	0.772	3.368*	0.792	2.786	0.828
	1900-1996	3.172*	0.809	3.394*	0.788	3.335*	0.781
	1925-1996	2.406	0.818	2.723	0.787	2.444	0.792
	1940-1996	2.231	0.819	2.565	0.780	2.221	0.776
	1950-1996	1.947	0.827	2.241	0.787	1.870	0.788
Coal	1875-1996	2.639	0.875	2.340	0.883	2.779	0.850
	1900-1996	2.079	0.901	1.582	0.920	2.545	0.860
	1925-1996	1.095	0.952	0.960	0.956	2.464	0.885
	1940-1996	1.001	0.955	0.747	0.964	2.354	0.888
	1950-1996	0.982	0.955	0.722	0.965	2.501	0.883
Nat. Gas	1924-1996	1.736	0.963	1.705	0.963	2.041	0.954
	1925-1996	1.860	0.960	1.809	0.960	1.993	0.954
	1940-1996	2.838	0.897	2.777	0.894	3.170	0.876
	1950-1996	2.376	0.885	2.260	0.883	3.119	0.819

Note: N is the number of lags included in the augmented Dickey-Fuller unit root test, t is the t-statistic on $(\rho-1)$, and *, ** indicate that a unit root can be rejected at the 10% and 5% levels respectively. Also shown is the estimated value of ρ for each regression.

Rather than focusing on tests of whether or not prices follow random walks, it may be more informative to address the *extent* to which price shocks are temporary (i.e., dissipate as price mean reverts) or permanent. Variance ratio tests of the sort used by Cochrane (1986) and Campbell and Mankiw (1987) are informative in this regard. Such tests are based on the fact that if price follows a random walk, i.e., is not stationary, then the variance of k-period differences should grow linearly with k. Hence the ratio

$$R_k = \frac{1}{k} \frac{Var(p_{t+k} - p_t)}{Var(p_{t+1} - p_t)} \tag{2}$$

should approach 1 as k increases. On the other hand, if price follows a stationary (mean-reverting) process, the variance of k-period differences will approach an upper limit as k grows, so that this ratio will fall to zero as k increases. More generally, the ratio provides a measure of the extent to which price shocks are persistent, or equivalently, the relative importance of any random walk component of price.

Figures 4, 6 and 6 plot the variance ratio R_k for each of the three log real prices. Observe that for oil and coal, this ratio declines to a value of about 0.1 to 0.2 as k approaches 25, and then levels out. This means that any random walk (i.e., permanent) component of price shocks is small, so that shocks are mostly transitory. This is consistent with a model in which price is slowly mean-reverting, as suggested by the patterns in Figures 1 and 2.

For natural gas, on the other hand, the ratio R_k increases with k, reaching a level above 2. This is inconsistent with mean reversion, and even with a geometric Brownian motion process. However, this pattern may simply reflect the shorter time series for natural gas prices, along with the high degree of curvature of the quadratic trend line. (The short time series and highly curved trend line would also account for the unit root test statistics shown in Table 1.)

In summary, the price trends shown in Figures 1, 2 and 3 and the variance ratios shown in Figures 4, 5 and 6 suggest that even with 120 years of data, unit root tests are unlikely to give us much information about the stochastic processes that best represent long-run price evolution. As Perron (1989) has shown, even a single discrete shift in the trend line can bias unit root tests, and Figures 1, 2 and 3 indicate that the trend lines for these resources fluctuate continuously over time. In the next two sections, I show that continuous and unpredictable fluctuations in the level and slope of the trend line are also predicted by the theory of depletable resource production, and I lay out a framework for estimating models that allow for such fluctuations.

Figure 4. Variance Ratio - Crude Oil

**Figure 5. Variance Ratio - Bituminous Coal**

Figure 6. Variance Ratio - Natural Gas

3. THE TREND LINE FOR A DEPLETABLE RESOURCE

Although OPEC has succeeded in pushing oil prices above competitive levels for periods of time, over the long run oil production has been largely competitive. The same is true for coal and natural gas. Hence we would expect the real prices of these resources to revert to long-run total marginal cost, i.e., a marginal cost that includes user costs associated with reserve accumulation and resource depletion. As shown below, even a simple model of depletable resource production predicts that this long-run marginal cost is a trend line that fluctuates over time.

Movements in Trend Level and Slope

For a depletable resource such as oil, natural gas, or copper, we would expect both the level of the log price trajectory and its slope to fluctuate over time in response to fluctuations in demand, extraction costs, and reserves. To see this, consider the basic Hotelling model of a depletable resource produced in a competitive market with a constant marginal cost of extraction, c. In this model the price trajectory is $dP/dt = r(P\text{-c})$. Hence the price level itself is given by

$$P_t = P_0' e^{rt} + c,$$ (3)

where $P_0' = P_0 - c$ is the net price. If the demand function is isoelastic, i.e., is of the form $Q_t = AP_t^{-\eta}$, the trajectory for the rate of production will be given by

$$Q_t = A(c + P_0' e^{rt})^{-\eta} .$$ (4)

In this case, we can find the initial net price P_0' by making use of the fact that cumulative production over the life of the resource must equal the initial reserve level, R_0:

$$R_0 = \int_0^\infty Q_t dt = \int_0^\infty A(c + P_0' e^{rt})^{-\eta} dt .$$ (5)

For arbitrary values of the elasticity of demand, η, this equation can be solved numerically for P_0'. For a unitary elasticity of demand ($\eta = 1$), we can solve it analytically:

$$R_0 = \frac{A}{rc} \log \frac{c + P_0'}{P_0'} ,$$

or

$$P_0' = \frac{c}{e^{rcR_0/A} - 1} .$$ (6)

Hence the price level at any time t is given by

$$P_t = c + \frac{ce^{rt}}{e^{rcR_0/A} - 1} ,$$ (7)

and the slopes of the price trajectory and log price trajectory are given by

$$\frac{dP_t}{dt} = \frac{rce^{rt}}{e^{rcR_0/A} - 1} \; ,$$ (8)

and

$$\frac{d\log P_t}{dt} = \frac{rc}{(e^{rcR_0/A} - 1)ce^{-rt} + c} \; .$$ (9)

From these equations it is easy to see that an upward shift in the demand curve, i.e., an increase in A, leads to an increase in the price level P_t, and an increase in the slopes of both the price trajectory and the log price trajectory. An increase in the level of extraction cost, c, leads to an increase in price, but a decrease in the slopes of the price and log price trajectories. Finally, if new discoveries result in an unexpected increase in the reserve level, R_0, this will cause a decrease in price, and will also lead to decreases in the slopes of the price and log price trajectories. For most depletable resources, one would expect demand, extraction costs, and reserves all to fluctuate continuously and unpredictably over time. Whether or not the processes that these variables follow are stationary is an open matter. But in either case, we would expect price to revert to a trend line with a level and slope that likewise fluctuate over time.[7]

If demand, extraction costs, and reserves change very infrequently but by large, discrete amounts, then a switching model of the sort estimated by Perron (1989) is appropriate as a description of price. But there is little evidence or reason to believe that this is the case for virtually any resource. In the case of oil, for example, estimated reserves, whether potential or proved, have fluctuated continuously over the past century, as has demand. Although OPEC's impact on oil prices changed significantly in 1973-74, its influence has waxed and waned since then.

It is important to note that this trend line to which price reverts, and which represents long-run total marginal cost, *is itself unobservable*. We might estimate the "parameters" of the trend line (and hence marginal cost itself) at

---

7. Of course, if demand, extraction costs, and reserves fluctuate stochastically through time, the simple Hotelling model above is incorrect. Producers will take these fluctuations into account when making their production decisions, so that the net price, $P - c$, need not rise at the rate of interest. Nonetheless, prices and expected rates of change of prices will fluctuate over time as demand, costs, and reserves fluctuate. See Pindyck (1980) for a model that illustrates this.

any point in time using data up to that point, but those parameters (and hence the corresponding level of marginal cost) will change over time. Thus if we want to forecast price under the belief that it reverts to long-run marginal cost, we must also estimate marginal cost and its trend through time.

4. A MODEL OF PRICE EVOLUTION

We have seen that a model of long-run commodity price evolution should incorporate two key characteristics: (i) reversion to an unobservable long-run total marginal cost, which follows a trend; and (ii) continuous random fluctuations in both the level and slope of that trend. A continuous time model that has these characteristics is a version of the multivariate Ornstein-Uhlenbeck process.

Suppose, first, that the log price follows a simple trending Ornstein-Uhlenbeck (OU) process. If the trend is quadratic, that process can be written in continuous time as:

$$d\bar{p} = -\gamma \bar{p} dt + \sigma dz , \qquad (10)$$

where $\bar{p} = p - \alpha_0 - \alpha_1 t - \alpha_2 t^2$ is the detrended price. In terms of the price level itself, this is equivalent to:

$$dp = [-\gamma (p - \alpha_0 - \alpha_1 t - \alpha_2 t^2) + \alpha_1 + 2\alpha_2 t] dt + \sigma dz . \qquad (11)$$

Note that the parameter γ describes the rate of reversion to the (fixed) trend line. If $\gamma = 0$, the log price follows an arithmetic Brownian motion (so price is a geometric Brownian motion), and the variance ratio would approach 1.

The multivariate Ornstein-Uhlenbeck process is a simple generalization of eqn. (10). In the bivariate case, we could write the process as:

$$d\bar{p} = (-\gamma \bar{p} + \lambda x) dt + \sigma dz_p , \qquad (12)$$

where x is itself an OU process:

$$dx = -\delta x dt + \sigma_x dz_x \qquad (13)$$

where dz_p and dz_x may be correlated. Eqns. (12) and (13) simply say that $\bar{p}$ reverts to $(\lambda/\gamma)x$ rather than 0, and x is mean-reverting around 0 if $\delta > 0$ and

a random walk if $\delta = 0$. In general, the variable x could be an observable economic variable, or it could be unobservable.[8]

In discrete time, this process would be given by the following equations:

$$\bar{p}_t = \alpha_p \bar{p}_{t-1} + \lambda x_{t-1} + \epsilon_{p,t} \tag{14}$$

$$x_t = \alpha_x x_{t-1} + \epsilon_{x,t} \ , \tag{15}$$

where $\epsilon_{p,t}$ and $\epsilon_{x,t}$ are normally distributed with mean 0, and with some covariance. Eqns. (14) and (15) describe a bivariate AR(1) process. As discussed below, if x were unobservable, this process could be estimated using the Kalman filter.

I consider a slightly more general multivariate version of this process that allows for fluctuations in both the level and slope of the trend. In particular, suppose that in continuous time the process for the log price is:

$$d\bar{p} = (-\gamma \bar{p} + \lambda_1 x + \lambda_2 yt)dt + \sigma dz_p \ , \tag{16}$$

with

$$dx = -\delta_1 x dt + \sigma_x dz_x \tag{17}$$

$$dy = -\delta_2 y dt + \sigma_y dz_y \tag{18}$$

Writing eqn. (16) in terms of the price p rather than the detrended price $\bar{p}$, we have:

$$dp = [-\gamma(p - \alpha_0 - \alpha_1 t - \alpha_2 t^2) + \alpha_1 + 2\alpha_2 t + \lambda_1 x + \lambda_2 yt]dt + \sigma dz_p \ . \tag{19}$$

Combining terms, we can write this as:

8. For a detailed discussion of this process, see Lo and Wang (1995).

$$dp = (-\gamma p - \alpha_0' - \alpha_1' t - \alpha_2' t^2 + \lambda_1 x + \lambda_2 yt)\, dt + \sigma dz_p \ . \tag{20}$$

Eqns. (20), (17), and (18) describe a process in which the log price reverts to a trend line with level and slope that fluctuate stochastically, and which may or may not be observable.[9]

These equations imply the following discrete-time model, which will be the basis of the empirical work that follows:

$$p_t = \rho p_{t-1} + b_1 + b_2 t + b_3 t^2 + \phi_{1t} + \phi_{2t} t + \epsilon_t \tag{21}$$

$$\phi_{1t} = c_1 \phi_{1,t-1} + \nu_{1t} \tag{22}$$

$$\phi_{2t} = c_2 \phi_{2,t-1} + \nu_{2t} \tag{23}$$

I will treat ϕ_{1t} and ϕ_{2t} as unobservable state variables. This is appropriate, since marginal cost at any point in time, the resource reserve base, and the demand parameters are all unobservable.

If we make the further assumption that the distribution of the error terms, ϵ_t, ν_{1t}, and ν_{2t} is multivariate normal and that ϵ_t is uncorrelated with ν_{1t} and ν_{2t}, then a natural estimator of this system of equations is the Kalman filter. The Kalman filter is a procedure that calculates maximum likelihood estimates of the parameters along with optimal (minimum mean-square error) estimates of the state variables ϕ_{1t} and ϕ_{2t} in this case). It is a forward looking procedure, in that it recursively estimates the values of the state variables at each point in time using all of the information available up to that point, and hence is well suited for forecasting.[10]

One issue that arises in the use of the Kalman filter is that initial estimates for the parameters and state variables are needed to begin the recursion. Typically OLS estimates of the parameters are used for this purpose, obtained by assuming that the state variables are constant parameters. (In the

---

9. In principle, this model could be generalized further by adding a third state variable that captures fluctuations in the curvature of the trend line, but without a great deal of data, the estimation of such second-order effects is unlikely to be feasible.

10. See Harvey (1989) for a detailed discussion of the Kalman filter. For the estimations in this paper, I use the implementation of the Kalman filter in EViews 3.0.

case of eqn. (20) these constants would be zero.) Initial values of the state variables are then obtained using these OLS estimates along with the first several data points for the observable variables. However, this initialization can be sensitive to these first few data points, particularly if the covariance of ν_{1t} and ν_{2t} must also be estimated. To simplify matters, I therefore assume that the errors in the state equations, ν_{1t}, and ν_{2t}, are uncorrelated.

5. ESTIMATION AND FORECASTING

I first estimate eqns. (21), (22) and (23) for each of the three resources using the full data set, i.e., 1870-1996 for oil and coal, and 1913-1996 for natural gas. In each case I drop the term $b_3 t^2$ in order to get convergence and stable estimates of the unobservable state variables. The Kalman filter yields estimates of each of the parameters c_1, c_2, etc., along with estimates of the state variables ϕ_{1t} and ϕ_{2t} for each year t after starting five years after the beginning of the sample. Note that each state variable estimate, e.g., $\hat{\phi}_{1t}$, is based on data only up to time t.

In principle, the state variables could follow stationary or nonstationary processes, i.e., the coefficients on $\phi_{1,t-1}$ and $\phi_{2,t-1}$ in eqns. (22) and (23) could be less than or equal to 1. For oil, I obtained a coefficient on $\phi_{1,t-1}$ very slightly above 1. For coal, the estimated coefficient on $\phi_{1,t-1}$ was slightly below 1, but the estimated coefficient on $\phi_{2,t-1}$ was about 1.03, so that the recursive estimates $\hat{\phi}_{2t}$ exploded. Hence I reestimated the model for coal, constraining the coefficient on $\phi_{2,t-1}$ to be 1.0. Finally, in the case of natural gas, I could not obtain convergence when there were two state variables in the model. The best results were obtained by dropping the state variable ϕ_{1t} and retaining only ϕ_{2t}.

Estimation results for the full sample are shown in Table 2. That table shows the estimates of the parameters, along with the final year (1996) estimates of the state variables, $\hat{\phi}_{1T}$ and $\hat{\phi}_{2T}$. Note that ϕ_{1t} is nonstationary for oil, both ϕ_{1t} and ϕ_{2t} are nonstationary or close to nonstationary for coal, and ϕ_{2t} is close to nonstationary for natural gas. To examine the forecasting performance of these models, I also estimated them for sample periods ending before 1996 (in particular, 1970, 1980, and for oil, 1981). For these shorter samples, the coefficients on $\phi_{1,t-1}$ and $\phi_{2,t-1}$ were always less than (but still close to) 1.

As explained above, because of the forward-looking nature of the Kalman filter, these models (or at least models of this type) should be well suited for forecasting, and I turn to this below. I focus mostly on crude oil, but I also present forecasts for coal and natural gas.

Table 2. Kalman Filter Estimates

	Oil	Coal	Nat. Gas
c_1	-0.7771	0.2053	-0.8776
	(0.0568)	(12.117)	(0.1339)
c_2	0.8041	0.0144	0.6872
	(0.0528)	(0.0723)	(0.0431)
c_3	1.0009	0.9961	-
	(0.0008)	(0.0468)	
c_4	0.5418	1.000	0.9910
	(1.833)	-	(0.0022)
ϕ_{1T}	1.1075	0.5199	-
	(0.0155)	(0.0797)	
ϕ_{2T}	2.96 x 10^{-9}	0.00657	0.0181
	(3.65 x 10^{-6})	(8.57 x 10^{-9})	(0.00057)

Note: Standard errors in parentheses. Estimates for oil and coal use data for 1870-1996, and for natural gas, 1919-1996. Model is:

$$p_t = c_1 + c_2 p_{t-1} + \phi_{1t} + \phi_{2t} t + \epsilon_t$$
$$\phi_{1t} = c_3 \phi_{1, t-1} + \nu_{1t}$$
$$\phi_{2t} = c_4 \phi_{2, t-1} + \nu_{2t}$$

Values for ϕ_{1T} and ϕ_{2T} are 1996 (final year) Kalman filter estimates of these unobservable state variables. For coal, c_4 is constrained to 1.0. For natural gas, only unobserved state variable is ϕ_{2t}.

Crude Oil

I estimated eqns. (21), (22), and (23) using data for the subsamples 1870-1970, 1870-1980, 1870-1981, and the full sample 1870-1996. (I included 1870-1981 because oil prices peaked in real terms towards the end of that year.) In each case, I used the resulting model, i.e., the estimates of the parameters, along with the final year estimates of the state variables, $\hat{\phi}_{1T}$ and $\hat{\phi}_{2T}$, to forecast the log price out to the year 2010.

These forecasts are shown in Figure 7, along with the actual series for the log price. Observe that the forecasts, which begin at dates spanning a range of 26 years, all converge to a narrow band for the years 2000 to 2010. Thus the forecasts give little weight to the unusually large variations in price that occurred during the 1974 to 1987 period.

Figure 7. Forecasts of the Log Oil Price

One might think that with more than a century of data, these forecasts simply reflect a trend line that is not much affected by the inclusion or exclusion of the 1974 to 1987 price variations. That is not the case, however, as is shown by the trend line forecasts in Figure 1. (Those trend lines end at the year 2000, but it is clear that even at that point they diverge from the forecasts in Figure 7.) Of course those trend line forecasts include no mean reversion. One could argue that the simple alternative to eqns. (21), (22), and (23) is a model of mean reversion around a fixed trend line, e.g., a simple trending Ornstein-Uhlenbeck process in continuous time. (Although the trend in such a model is assumed fixed, different sample periods, e.g., 1870-1970 versus 1870-1980, will result in different trend lines.)

Figure 8 compares the forecasts from the Kalman filter estimation of eqns. (21), (22), and (23) to forecasts based on mean reversion to a fixed linear trend, i.e., a model of the form:

$$p_t = \rho p_{t-1} + b_1 + b_2 t + \epsilon_t \ . \tag{24}$$

In Figure 8, the alternative models are estimated using data first through 1981, and then through 1996, and in each case forecasts are generated out to the year 2010. (The Kalman filter forecasts are labelled KF and the forecasts based on eqn. (24) are labelled MR.)

Figure 8. Log Oil Price Forecasts - Kalman Filter and Mean Reversion to Fixed Trend


Note that for both the 1981 and 1996 dates, the model with a stochastically varying trend produces much lower forecasts than the alternative model with mean reversion to a fixed trend. The reason is that the latter model gives more weight to the 1974-1985 data. Unusually high real oil prices during that period increase the slope of the trend line (whether estimated through 1981 or 1996) to which price is assumed to revert. For forecasts out to 2010, mean reversion is almost complete, so the high forecasted prices simply reflect this trend line.

On the other hand, the model of eqns. (21), (22), and (23) allows the trend line to fluctuate, so that the high prices during 1974-1985 correspond in part to a trend line with a slope that is unusually high. This slope is predicted to decline to historical values, so that the forecasts of price are accordingly lower. (The level of the trend line is high as well, but as Table 2 shows, the state variable corresponding to this level is estimated to be a random walk, and thus does not revert to an average historical value.)

Figure 9 illustrates this for the 1981 forecasts. It shows the estimated values of the unobservable state variables $\hat{\phi}_{1t}$ and $\hat{\phi}_{2t}$ up to 1981, along with forecasts of these variables for the years 1982 to 2010. Note that ϕ_{1t}, the state variable representing the level of the trend line, is forecasted to grow steadily, which is a reflection of the fact that it has been estimated to be a random walk.

However, ϕ_{2t}, the state variable representing the slope of the trend line, is forecasted to revert to a level close to zero.

Figure 9. State Variables - Estimates and 1981 Forecasts

How useful is this model as a forecasting tool? As Figure 8 shows, the forecasts it produces using data through 1981 are much closer to actual oil prices over the next 15 years than are forecasts based on mean reversion to a fixed trend line. Of course this might be just luck—if oil prices rise substantially during the first decade of the 21st century, the model's 1996 forecasts will not have performed as well as the simpler alternative. (However, at the time of this writing, the model's 1996 forecasts seem much closer to the consensus view about oil markets.) But putting aside its forecasting performance over the past two decades, the model captures in a nonstructural framework what basic theory tells us should be driving price movements.

Coal and Natural Gas

I also estimated eqns. (21), (22), and (23) for the log price of coal and the log price of natural gas, again using data through 1970, through 1980, and through 1996. I again used the resulting models to forecast the log price out to the year 2010. The forecasts are shown in Figures 10 and 11 along with the actual series for the log prices.

These forecasts look very different from those for crude oil. The forecasts of the log oil price all converged to a narrow band for the years 2000 to 2010, whereas these forecasts diverge considerably. The reason for this is that all of the state variables for coal and natural gas are estimated to be random walks or something very close to a random walk. (The estimates in Table 2 are for the full sample, but the estimates of c_3 and c_4—the AR(1) parameters for the state variables ϕ_{1T} and ϕ_{2T}—for the shorter samples are very close to those in the table.) In the case of coal, both c_3 and c_4 are equal to or very close to 1, and for natural gas, where there is only one unobservable state variable (ϕ_{2t}), the coefficient c_4 is again very close to 1.

Thus, the forecasting performance of this model is poor for coal and natural gas. This is the case not so much because the 1970 and 1980 forecasts do not closely replicate the data through 1996, but more because the model does not describe movements in the trend lines for these resources that are consistent with theory. Again, the problem is in the estimates of the state equations (22) and (23), and may be due to the sensitivity of the Kalman filter estimates to the initialization. Using a different sample period, e.g., dropping the first 10 years of data, produces estimates that yield better forecasts.

Figure 10. Forecasts of the Log Coal Price

Figure 11. Forecasts of the Log Gas Price

6. IMPLICATIONS FOR INVESTMENT DECISIONS

So far we have focussed on the problem of long-term forecasting. Putting that aside, what do the results above tell us about the kinds of models of price evolution that should be used for evaluating irreversible investment decisions? Much of the literature on "real options" makes the convenient assumption that output price, input cost, or some other relevant stochastic state variable follows a geometric Brownian motion (GBM). A good example is Paddock, Siegel, and Smith (1988). If the true process is a multivariate Ornstein-Uhlenbeck process, or even a simple trending OU process, how far off might one go by assuming that price follows a GBM instead?

Lo and Wang (1995) have calculated call option values for stocks with prices that follow a trending Ornstein-Uhlenbeck process, and compared these to the values obtained from the Black-Scholes model (which is based on a GBM for the stock price). They show that Black-Scholes can over- or underestimate the correct option value, but generally the size of the error is small, at least relative to errors that would be tolerable in real option applications. (They find errors on the order of 5 percent, which would be significant for a financial option. In the case of a capital investment decision, there are enough other uncertainties regarding the modelling of cash flows that an error of this size is unlikely to be important.)

Of course financial options typically have lifetimes of a few months to a year, and Lo and Wang focussed on the value of the option rather than the optimal exercise point. Real options are typically much longer lived, and determining the optimal exercise point (i.e., the investment rule) is usually more important than valuing the option (i.e., the investment opportunity) itself.

Dixit and Pindyck (1994) solve for optimal investment rules when a fixed capital expenditure I results in a project worth V, and V follows a stochastic process. They considered mean-reverting processes for V (as well as a GBM), and showed that if the rate of mean reversion is fast, the optimal investment rule will depend strongly on the value $\bar{V}$ (assumed fixed) to which V reverts. The dependence is much weaker, however, if the rate of mean reversion is very slow. In the case of energy prices, Figures 1, 2, and 3, along with the estimates of ρ in Table 1, suggest that the rate of mean reversion (whether to a fixed or stochastically fluctuating trend line) is slow, suggesting that for many applications, the GBM assumption may not be bad.

Even if mean reversion is slow, the GBM assumption will be appropriate only if the implied volatility is relatively constant. (If volatility is in fact fluctuating stochastically, then models that assume it is constant can lead to significant errors in the optimal investment rule.) We can address this issue with the long data series that we have used above. Suppose we assume that the real price of a particular resource follows a geometric Brownian motion, i.e.,

$$dP = \mu P dt + \sigma P dz . \tag{25}$$

We can use our data to estimate the mean expected growth rate μ and the annual volatility σ. Estimates of these parameters for various subsamples of the data are shown in Table 3.

Note that over the full sample and most of the subsamples the mean rates of growth of the real prices of oil and coal have been close to zero, and for natural gas have been in the range of 2 to 4 percent per year. The stability of these growth rates is consistent with a model of slow mean reversion. The estimated annual standard deviation, σ, has been even more stable. For oil, it stays between 15 and 20 percent for the full sample and for all of the subsamples. For coal, the annual standard deviation has been between 7 and 9 percent, and for natural gas between 11 and 14 percent. These numbers suggest that for irreversible investment decisions for which energy prices are the key stochastic state variables, the GBM assumption is unlikely to lead to large errors in the optimal investment rule.

Table 3. Annual Percentage Growth Rates and Standard Deviations

Sample	Oil		Coal		Nat. Gas	
	μ	σ	μ	σ	μ	σ
1870-1996	-0.0027	0.2072	0.0011	0.0873		
1880-1996	0.0064	0.1794	0.0028	0.0876		
1890-1996	0.0052	0.1784	0.0021	0.0817		
1900-1996	0.0018	0.1702	0.0027	0.0852		
1910-1996	0.0111	0.1740	0.0039	0.0854		
1920-1996	0.0050	0.1712	0.0027	0.0833	0.0210	0.1193
1930-1996	0.0069	0.1683	0.0025	0.0701	0.0188	0.1079
1940-1996	0.0096	0.1603	-0.0014	0.0705	0.0274	0.1099
1950-1996	0.0067	0.1725	-0.0077	0.0716	0.0424	0.1144
1960-1996	0.0091	0.1762	-0.0038	0.0768	0.0395	0.1259
1970-1996	0.0144	0.2069	-0.0030	0.0894	0.0490	0.1451

Note: μ is the sample mean and σ the sample standard deviation of annual price changes.

CONCLUSIONS

I have argued that the theory of depletable resource production and pricing, and the actual behavior of real prices over the past century, both imply that nonstructural forecasting models should incorporate mean reversion to a stochastically fluctuating trend line. That trend line reflects long-run (total) marginal cost, which is unobservable. Hence these models are naturally estimated using Kalman filter methods.

These models seem promising as a forecasting tool, even though the results in this paper were mixed. The models performed well in forecasting oil prices, but less well for coal and natural gas. The difficulties with respect to coal and natural gas may be due to problems of initialization, and the sensitivity of the estimates to the first few data points. (I obtained very different estimates for coal by beginning the sample in 1880 or 1890, instead of 1870.) But overall, the promise of these models derives largely from the fact that they capture in a nonstructural framework what basic theory tells us should be driving price movements.

The framework presented here could be extended (although how easily is not clear). A natural extension would be to estimate these price models for different energy resources as a system, accounting for cross-equation errors. In theory, the Kalman filter could be used to estimate a larger system of this sort, but in practice it may be difficult. I found the estimation of a model with more than two unobservable state variables to be infeasible, and I even found it difficult to increase the number of exogenous variables in the price equation. On the other hand, this may simply reflect limitations in the algorithm that I used. The pursuit of such extensions certainly seems warranted.

REFERENCES

Baker, Malcolm P., E. Scott Mayfield, and John E. Parsons (1998). "Alternative Models of Uncertain Commodity Prices for Use with Modern Asset Pricing Methods." *The Energy Journal* 19(1): 115-148.

Campbell, John, and N. Gregory Mankiw (1987). "Are Output Fluctuations Transitory?" *Quarterly Journal of Economics* 102: 857-880.

Campbell, John Y., Andrew W. Lo, and A. Craig MacKinlay (1997). *The Econometrics of Financial Markets*. Princeton: Princeton University Press.

Cochrane, John H. (1988). "How Big is the Random Walk in GNP?" *Journal of Political Economy* 96: 893-920.

Dixit, Avinash K., and Robert S. Pindyck (1994). *Investment Under Uncertainty*. Princeton, New Jersey: Princeton University Press.

Froot, Kenneth A., and Kenneth Rogoff (1995). "Perspectives on PPP and Long-Run Real Exchange Rates." *Handbook of International Economics, Vol. 3*, Gene Grossman and Kenneth Rogoff, eds., Amsterdam: North Holland.

Harvey, Andrew C. (1989). *Forecasting, Structural Time Series Models, and the Kalman Filter*. Cambridge, England: Cambridge University Press.

Lo, Andrew W., and Jiang Wang (1995). "Implementing Option Pricing Models When Asset Returns are Predictable." *Journal of Finance* (March) 50: 87-129.

MacKinnon, James (1991). "Critical Values for Cointegration Tests," in R. F. Engle and C. W. J. Granger, eds., *Long-Run Economic Relationships*. Oxford: Oxford University Press.

Manthy, Robert S. (1978). *Natural Resource Commodities: A Century of Statistics*. Baltimore: Johns Hopkins Press.

Paddock, James L., Daniel R. Siegel, and James L. Smith (1988). "Option Valuation of Claims on Real Assets: The Case of Off-Shore Petroleum Leases." *Quarterly Journal of Economics* (August) 103: 478-508.

Perron, Pierre (1989). "The Great Crash, the Oil Price Shock, and the Unit Root Hypothesis." *Journal of Political Economy* (December) 57: 1361-1401.

Pindyck, Robert S. (1978). "The Optimal Exploration and Production of Nonrenewable Resources." *Journal of Political Economy* (October) 86: 841-861.

Pindyck, Robert S. (1980). "Uncertainty and Exhaustible Resource Markets." *Journal of Political Economy* 88: 1203-1225.

Pindyck, Robert S. (1994). "Inventories and the Short-Run Dynamics of Commodity Prices." *The RAND Journal of Economics* (Spring) 25: 141-159.

U.S. Department of Commerce (1975). *Historical Statistics of the United States.*

Videgaray, Luis (1998). "The Fiscal Response to Oil Shocks," Ph.D. dissertation, M.I.T. Department of Economics, (June).