Optimal World Oil Extraction:
Calibrating and Simulating the Hotelling Model

C.-Y. Cynthia Lin[1]

Preliminary—Comments welcome

Abstract

This paper uses data on world oil price and consumption to calibrate a Hotelling model of optimal resource extraction with unlimited potential reserves when costs exhibit stock effects. Numerical solutions are generated for various specifications of the elasticity of demand for both isoelastic demand and linear demand under each of two possible market structures: perfect competition and monopoly. My simulations enable me to examine the following questions. First, how well does the simple Hotelling model appear to explain historical data? And second, how are future world oil prices and extraction rates predicted to evolve? From among the various specifications I tried, the model best fits actual data when the oil market is perfectly competitive and when demand is inelastic. Under these assumptions, real oil price should fall in the range $117-160/barrel over the years 2000-2010 and extraction should be roughly constant at 59.8 million barrels/day over the entire simulation period 1857-2300. However, even the best-fit specification fails to adequately explain the data, which suggests that a richer theoretical model may be needed.

JEL Classification: Q31, C61, E27, D4

First draft: April 20, 2004
This draft: November 12, 2004

[1]Department of Economics, Harvard University; cclin@fas.harvard.edu. I would like to thank Gary Chamberlain for his advice and guidance throughout this project. This paper also benefited from discussions with Martin Weitzman, William Hogan and Howard Stone. The time series data used in this study were acquired with the help of Brian Greene and with funds from the Littauer Library at Harvard University. I am indebted to Bijan Mossavar-Rahmani (Chairman, Mondoil Corporation) and William Hogan for arranging for me to visit Apache Corporation's headquarters in Houston and a drilling rig and production platform offshore of Louisiana, and for their support of my research, and I thank the Repsol YPF - Harvard Kennedy School Energy Fellows Program for providing travel funds. I thank Mark Bauer (Reservoir Engineering Manager, Apache), Robert Dye (VP, Apache), Steve Farris (President, CEO & COO, Apache), Richard Gould (VP, Wells Fargo), Paul Griesedieck (Manager, Apache), Thomas Halsey (Corporate Strategic Research, ExxonMobil Research and Engineering), Becky Harden (Land Manager, Apache), David Higgins (Director, Apache), Jon Jeppesen (Sr. VP, Apache), Adrian Lajous (President, Oxford Institute for Energy Studies), Kenneth McMinn (Offshore District Production Manager, Apache), Kregg Olson (Director, Apache), and Bob Tippee (Editor, Oil and Gas Journal) for enlightening discussions about the petroleum industry. I thank Derrick Martin for flying me offshore by helicopter, Mike Thibodeaux for giving me a tour of the production platform, Joey Bridges for giving me a tour of the drilling rig, and especially Billy Ebarb (Production Superintendent, Apache) for accompanying me throughout the entire offshore trip. William Horvath (Energy Information Administration) provided data that might be useful for future versions of this paper. Aloulou Fawzi (Energy Information Administration) and Jeff Obermiller (American Petroleum Institute) provided leads for possible data sources. I received financial support from an EPA Science to Achieve Results (STAR) graduate fellowship, a National Science Foundation (NSF) graduate research fellowship, and a Repsol YPF - Harvard Kennedy School Pre-Doctoral Fellowship in energy policy. All errors in this draft (and there are many!) are my own.

1 Introduction

The problem of optimal nonrenewable resource extraction was first examined by Hotelling (1931), whose basic model predicted that the shadow price of the resource stock, which is an economic measure of the scarcity of the resource, should grow at the rate of interest. Since then, economists have expanded Hotelling's basic theoretical framework to allow for more realistic features such as increasing extraction costs (Hanson, 1980; Solow & Wan, 1976), unlimited potential reserves (Pindyck, 1978), market imperfections (Khalatbari, 1977; Stiglitz, 1976; Sweeney, 1977), technological progress (Lin, 2004b,c), and uncertainty (Hoel, 1978; Pindyck, 1980).

In addition to enriching the theory, economists have also taken the Hotelling model to data. For example, much work has also been done in attempt to measure the shadow price of the resource (see e.g., Devarajan & Fisher, 1982; Halvorsen & Smith, 1984; Lasserre, 1985). Some studies have also attempted to use data to test Hotelling's theory. Empirical tests of the dynamic efficiency conditions of the Hotelling model under perfect competition have been applied to data on a hard rock mining firm (Farrow, 1985), on Canadian metal mining firms (Halvorsen & Smith, 1991), and on Canadian copper mining firms (Young, 1992). However, owing in part to limitations on the cost data needed, such tests have yet to be applied to the world oil industry.[2] Moreover, an empirical test of dynamic efficiency of world oil extraction may also need to relax the assumption of perfect competition.[3]

In this paper, I use data on world oil price and consumption to calibrate a Hotelling model of optimal oil extraction when reserves are unlimited and when costs depend on the cumulative stock extracted. I use the calibrated model to simulate solution trajectories under both perfect competition and monopoly for various specifications of the elasticity of demand for both isoelastic demand and linear demand.

My paper follows most closely the work of Pindyck (1978), who first develops a theoretical model that allows for unlimited potential reserves and that requires resource producers to simultaneously determine optimal rates of exploration and production, and who then examines numerically the characteristics of the competitive and monopoly solutions to his model using data for oil in the Permian region of Texas over the period 1965-1974.

Though this paper also examines both competitive and monopoly solutions and allows for unlimited reserves, it differs from that of Pindyck (1978) in several ways. First, the data I use to calibrate my model is for world oil over the period 1965-2001, which spans both a wider geographic area and a longer period of time than Pindyck's data does. Second, instead of Pindyck's more complex—and perhaps more realistic—model that allows producers to choose the rates of both exploration and extraction, I use Hotelling's more basic model of allowing producers

[2] Miller and Upton (1985) use data on U.S. domestic oil- and gas-producing companies to test another reduced-form implication of a Hotelling model. This reduced-form implication, which they term the "Hotelling Valuation Principle", is that the value of a unit of reserves in the ground is the same as its current value above the ground less the marginal costs of extracting it.

[3] In future work, I hope to develop an analogous test for world oil.

to choose extraction rates only. I thus can gauge whether Hotelling's simpler model is sufficient to explain the historical data. Third, while Pindyck assumes the demand is linear, I examine the both the linear demand case and the isoelastic demand case. The fourth way in which I expand Pindyck's work is that I present my results under a range of values of the demand elasticity instead of just choosing one.[4] A fifth difference is that, while this paper compares the predicted trajectories for both price and extraction with actual data and varies the parameters in attempt to match the data, Pindyck only compares the optimal values of well drilling and price to their historical values for the one set of parameters he used.[5]

My simulations enable me to examine the following questions. First, *how well does the simple Hotelling model appear to explain historical data?* And second, *how are future world oil prices and extraction rates predicted to evolve?*

From among the various specifications I tried, the model best fits actual data when the oil market is perfectly competitive and when demand is inelastic. Under these assumptions, real oil price should fall in the range $117-160/barrel over the years 2000-2010 and extraction should be roughly constant at 59.8 million barrels/day over the entire simulation period 1857-2300. However, even the best-fit specification fails to adequately explain the data, which suggests that a richer theoretical model may be needed.

The balance of this paper proceeds as follows. In Section 2, I present my basic Hotelling model. In Section 3, I describe my data. In Section 4, I explain the functional form assumptions and calibration methods used for my simulations. My results are presented in Section 5. Section 6 concludes.

2 The Basic Hotelling Model

In this section, I present my theoretical model of optimal nonrenewable resource extraction under both perfect competition and monopoly. The notation follows closely that used by Weitzman (2003).

2.1 The General Framework

Suppose there are T oil markets indexed by $t = 1, ..., T$. For each time t, the supply of oil is given by $E(t)$, the total extraction flow in units of oil per unit time at time t.[6] Let $X(t)$ denote the total cumulative stock of oil extracted at time t:

[4] Possible extensions of my work would vary either the base year on which some parameters are calibrated or the discount rate, or the terminal date of the simulations.

[5] Pindyck also compares the optimal and actual values to the myopic values that would occur if future depletion were ignored but the reserve-production ratio were maintained at its optimal level.

[6] I assume that, at any given time t, all the oil extracted at time t is sold on the market at time t.

$$X(t) = X(0) + \int_0^t E(\tau) d\tau \quad , \tag{1}$$

where the initial stock $X(0)$ is taken as given.

The market price of oil at time t is $P(t)$. The demand for oil when the market price is P is given by the demand function $D(P)$. Markets are assumed to clear, which means that, at each time t, the price $P(t)$ acts to equate supply and demand:

$$E(t) = D(P(t)) \quad \forall t. \tag{2}$$

The cost of extracting E units of oil when the total stock already extracted is X is given by $C(X, E)$. I use the term "stock effects" to refer to the dependence of extraction cost on the stock X of reserve extracted. There are several possible reasons why this dependence is positive. First, extraction costs may increase with the cumulative stock extracted if the resource needed to be extracted from greater depths as it was being depleted. Second, costs may increase if well pressure declined as more of the reserve was depleted. Third, since different grades of oil may differ in their extraction costs, and since the cheaper grades are likely to be mined to exhaustion before the more expensive grades are mined, the cost of extraction may increase as the cheaper grades are exhausted, and therefore as the total stock already extracted increased.

Let $p(t)$ denote the non-negative current-value shadow price measuring the value of a unit of reserve at time t. This shadow price is known by a variety of terms, including "marginal user cost", because it measures the opportunity cost of extracting the resource; "in situ value", because it measures the marginal value of leaving an additional unit of resource in the ground; "scarcity rent", because it is an economic measure of scarcity, and "dynamic rent", to reflect the difference between price and marginal extraction cost (Krautkraemer, 1998; Weitzman, 2003). As the reader will soon see, the shadow price plays a crucial role in the Hotelling model.

The competitive interest rate is ρ.

2.2 Perfect Competition

When the oil market is perfectly competitive, extraction $E(t)$ represents the total amount of oil extracted by all the firms in the market at a given point in time.[7]

[7] I ignore any common access problems that may arise in perfect competition. In other words, I assume, as does Pindyck (1978), that there is a large number of identical firms that all ignore each other, or, equivalently, that a social planner or a state-owned company has sole production rights and sets a competitive price.

The total benefits $U(\cdot)$ from oil at time t is given by the area under the demand curve:

$$U(E(t)) = \int_0^{E(t)} D^{-1}(x)dx. \tag{3}$$

This area measures the gross consumer surplus, and is a measure of the consumers' total money-metricized willingness-to-pay. As shown in Weitzman (2003), using the area under the demand curve in place of revenue yields the same outcome as a perfectly competitive market.[8] For mathematical simplicity, I thus choose to model the perfectly competitive firm's maximization problem using the area under the demand curve.

The per-period net benefit $G(X, E)$ from extracting E units of oil when the total stock already extracted is X is given by total benefits minus total costs:

$$G(X, E) = U(E) - C(X, E). \tag{4}$$

The social planner's optimal control problem, which yields the same solution as would arise in perfect competition, is to choose her extraction profile $\{E(t)\}$ to maximize the present discounted value of her entire stream of net benefits, given her initial stock $X(0)$ and given the relationship between her extraction $E(t)$ and the cumulative stock extracted $X(t)$, and subject to the constraints that both extraction and stock are nonnegative. Her problem is thus given by:

$$\begin{aligned} \max_{\{E(t)\}} & \int_0^{\infty} \left(U(E(t)) - C(X(t), E(t))\right) e^{-\rho t} dt \\ \text{s.t.} \quad & \dot{X}(t) = E(t) \qquad : q(t) \\ & E(t) \geq 0 \\ & X(t) \geq 0 \\ & X(0) = X_0 \quad , \end{aligned} \tag{5}$$

where $q(t) \leq 0$ is the multiplier associated with the equation of motion for the total stock $X(t)$ of oil extracted. The absolute value of this multiplier is precisely the shadow price $p(t)$ of the reserve:

$$p(t) = |q(t)|. \tag{6}$$

[8] This is because $P(t) = U'(E(t))$, so that the first-order conditions for the social planner's problem are the same as those that arise in perfect competition.

From the Maximum Principle, the first-order necessary conditions for a feasible trajectory $\{X^*(t), E^*(t)\}$ to be optimal are:[9]

$$[\#1]: \qquad p(t) = P(t) - \frac{\partial C\left(X(t), E(t)\right)}{\partial E} \tag{7}$$

$$[\#2]: \qquad \overset{\bullet}{p}(t) = -\frac{\partial C\left(X(t), E(t)\right)}{\partial X} + \rho p(t) \tag{8}$$

$$[\#3]: \qquad \lim_{t\to\infty} p(t)X(t)e^{-\rho t} = 0 \tag{9}$$

Condition [#1] states that, at each time t, the shadow price $p(t)$ must equal the competitive market price $P(t)$ minus the marginal cost of extraction $\frac{\partial C(X(t),E(t))}{\partial E}$; this condition is needed to ensure static optimality at each point in time. Condition [#2] governs how the shadow price $p\left(t\right)$ must evolve over time; conditions [#1] and [#2] combined are needed to ensure intertemporal optimality over all finite subperiods. Condition [#3], the transversality condition, is required for the solution to be dynamically optimal over the entire infinite horizon (Weitzman, 2003).

One can use the constraints, market clearing condition and first-order conditions to reformulate the Hotelling problem into the following ordinary differential equation boundary value problem:

$$\begin{array}{lcc} \text{differential equations} & : & \begin{array}{c} \frac{d}{dt}P(t) = \frac{d}{dt}\frac{\partial C(X(t),D(P(t)))}{\partial E} - \frac{\partial C(X(t),D(P(t)))}{\partial X} + \rho\left(P(t) - \frac{\partial C(X(t),D(P(t)))}{\partial E}\right) \\ \frac{d}{dt}X(t) = D(P(t)) \end{array} \\ \text{boundary conditions} & : & \begin{array}{c} X(0) = X_0 \\ \lim_{t\to\infty}\left(P(t) - \frac{\partial C(X(t),D(P(t)))}{\partial E}\right)X(t)e^{-\rho t} = 0 \end{array} \end{array} \tag{10}$$

The solution to the boundary value problem (10) is equivalent to that of the optimal control problem (5); I thus derive solutions to the latter problem by solving the former.

2.3 Monopoly

When oil is produced by a single monopolist rather than a multitude of perfectly competitive firms, the total benefits of oil production no longer equal the area under the demand curve, but rather total revenue instead. Total revenue $\Phi(\cdot)$ at time t is given by:

[9] If then per-period net benefit function $G(X, E)$ is concave in both X and E, then, since the control set $\{E \mid E \geq 0\}$ is convex, the first-order conditions are both necessary and sufficient for an optimum (Weitzman, 2003).

$$\Phi(E(t)) = D^{-1}(E(t)) \cdot E(t). \tag{11}$$

As a consequence, the monopolist's per-period profit $G(X, E)$ is given by:

$$G(X, E) = \Phi(E) - C(X, E), \tag{12}$$

and his optimal control problem is given by:

$$\begin{aligned} \max_{\{E(t)\}} & \int_0^{\infty} \left(\Phi(E(t)) - C(X(t), E(t))\right) e^{-\rho t} dt \\ \text{s.t.} \quad & \dot{X}(t) = E(t) \quad : q(t) \\ & E(t) \geq 0 \\ & X(t) \geq 0 \\ & X(0) = X_0 \ , \end{aligned} \tag{13}$$

which yields the following first-order conditions:

$$[\#1] \quad : \quad p(t) = \Phi'(E(t)) - \frac{\partial C\left(X(t), E(t)\right)}{\partial E} \tag{14}$$

$$[\#2] \quad : \quad \dot{p}(t) = -\frac{\partial C\left(X(t), E(t)\right)}{\partial X} + \rho p(t) \tag{15}$$

$$[\#3] \quad : \quad \lim_{t \to \infty} p(t) X(t) e^{-\rho t} = 0 \tag{16}$$

In order to ensure static optimality, condition [#1] requires that in monopoly, unlike in perfect competition, the shadow price $p(t)$ must equal marginal revenue $\Phi'(E(t))$ minus marginal cost $\frac{\partial C(X(t),E(t))}{\partial E}$ at every time t. Conditions [#2] and [#3] are the same in monopoly as in perfect competition.

As before, one can use the constraints, market clearing condition and first-order conditions to reformulate the Hotelling problem into an ordinary differential equation boundary value problem, where now the problem is given by:

$$
\begin{aligned}
\text{differential equations} \quad : \quad & \frac{d}{dt}\Phi'(D(P(t)) = \frac{d}{dt}\frac{\partial C(X(t),D(P(t)))}{\partial E} - \frac{\partial C(X(t),D(P(t)))}{\partial X} + \rho\left(\Phi'(D(P(t))) - \frac{\partial C(X(t),D(P(t)))}{\partial E}\right) \\
& \frac{d}{dt}X(t) = D(P(t)) \\
\text{boundary conditions} \quad : \quad & X(0) = X_0 \\
& \lim_{t\to\infty}\left(\Phi'(E(t)) - \frac{\partial C(X(t),D(P(t)))}{\partial E}\right)X(t)e^{-\rho t} = 0
\end{aligned}
\tag{17}
$$

The solution to the boundary value problem (17) is equivalent to that of the optimal control problem (13).

Having explained my theoretical model, I now describe the data I use both for calibrating my model and for assessing its validity.

3 Data

In order to calibrate my theoretical model and assess its validity, I use annual data spanning the years 1965-2001. For oil price P, I use the real annual spot price for crude oil, averaged over the Brent, Dubai, and West Texas Intermediate (WTI) prices. This average price time series was obtained from the World Bank and deflated to 1982-1984 U.S. dollars using the consumer price index (CPI).[10] For oil quantity E, I use world oil consumption as reported by BP.[11] Table 1 provides summary statistics for my price and quantity data.

TABLE 1. Summary statistics

Variable	mean	s.d.	min	max	trend
real world oil price (1982-1984$/barrel)	16.00	11.03	3.12	44.75	0.13 (0.17)
world oil consumption (million barrels/day)	59.33	11.54	31.23	75.45	0.99 * (0.07)

Notes: The trend is the coefficient on year when the variable is regressed on year and a constant. Significance codes: * 0.1 % level.

As can be seen from Table 1, real world oil price has no significant trend over 1965-2001. The trendless nature of oil price over time is in accordance with many empirical studies (see Krautkraemer, 1998, & references therein).[12] In contrast, world oil consumption is increasing over 1965-2001 by nearly 1 million barrels/day each year.

[10] I use a U.S. deflator rather than a world deflator because the original nominal time series was in current U.S. dollars.

[11] As a possible extension to my paper, I could construct a longer time series by combining my 1965-2001 world oil consumption data with 1860-1948 world oil production data from Zimmermann (1951, p. 513), or I could use the 1951-2001 world oil production data generously given to me by William Horvath of the Energy Information Administration (EIA).

[12] Economists have found the trendless nature of oil prices puzzling. See Lin (2004b) and Lin (2004c) for theoretical expositions of this puzzle and attempts to reconcile the puzzle with theory.

4 Calibration

In this section, I describe the functional form and parameter assumptions I use to calibrate my model.

For my cost function $C(X, E)$, I use a cost function estimated by Chakravorty, Roumasset and Tse (1997). Their estimate was derived from world data on proven and estimated reserves and on extraction costs compiled by the East-West Center Energy Program. After trying a variety of functional forms, the marginal cost function they found to fit the data was of the form:

$$\frac{\partial C(X,E)}{\partial E} = c_1 e^{c_2 X} \tag{18}$$

where, when X is in units of 10^{15} British thermal units (Btu) and costs are in units of dollars per million Btu, the parameter values are given by $c_1 = 0.1774$ and $c_2 = 0.000217$.[13] I therefore use as my cost function:[14]

$$C(X,E) = c_1 e^{c_2 X} E \ . \tag{19}$$

For demand, a crucial parameter is the elasticity of demand ε. Possible values for the elasticity, according to various empirical studies and surveys, are shown in Table 2; they range from -0.5 to -0.0.

TABLE 2. Estimates of the demand elasticity ε

Estimate of ε	Source
-0.49 to -0.45	Berndt & Wood (1975, p. 265)
-0.5	Edmonson (1975, p. 172)
0.0	Lin (2004a)[15]
-0.3	Nordhaus (1980, p. 347)
-0.5 to -0.1	Pindyck (1978, p. 857)

I run my model under different values of ε, mostly in the range $\varepsilon \in [-2, 0]$, except as noted below.

I use two different forms for the demand function $D(\cdot)$. In the first form, I assume that demand is isoelastic:

$$D(P) = d_1 P^{-d_2}. \tag{20}$$

where the parameter d_2 is the absolute value of the elasticity of demand:

$$d_2 = |\varepsilon| \ . \tag{21}$$

[13] Because my data is in terms of barrels rather than mmBtu, I convert these parameters using the conversion factor: 5.8004 mmBtu = 1 barrel (USGS, 2004).

[14] I ignore any possible constant term when integrating the marginal cost function to yield the total cost function.

[15] These are from the 2SLS results in Lin (2004a).

Isoelastic demand functions are common in theoretical Hotelling models, as they can lend themselves to analytic solutions (see e.g., Lin 2004c; Stiglitz, 1978). For the monopoly case, the shadow price is positive only if $\varepsilon < -1$, which is less elastic than the range of the estimates in Table 2. Thus, in order for any extraction to occur under monopoly when demand is isoelastic, the demand function needs to be more elastic than reported in previous studies. Since d_1 satisfies:

$$d_1 = E \cdot P^{d_2}, \tag{22}$$

I calibrate d_1 by using data on world oil consumption and world oil price for a given base year τ for E and P, respectively.

The second form of the demand function that I use is linear demand:

$$D(P) = d_1 - d_2 P, \tag{23}$$

when the slope d_2 as a function of the elasticity of demand is given by:

$$d_2 = |\varepsilon| \cdot \frac{E}{P} \tag{24}$$

and the intercept d_1 as a function of the slope is given by:

$$d_1 = E + d_2 P \ . \tag{25}$$

Pindyck (1978) uses a linear demand function in his simulations. I calibrate d_1 and d_2 by using data on world oil consumption and world oil price for a given base year τ for E and P, respectively.

I use $\rho = 0.05$ as my market interest rate.

According to Zimmermann (1951, p. 512), the first year in which world petroleum production data is available is 1857, when 2000 barrels of oil were produced. I therefore use 1857 as my initial year, corresponding to $t = 0$, and set my initial cumulative stock of extracted resource, X_o, to be 2000 barrels.[16] I assume that the final year, corresponding to $t = \infty$, occurs in 2300. For calibrating the demand function, I use 1981 as my base year τ.[17]

[16] I also run simulations in which I instead pin down the initial price $P(0)$ to equal to actual real world oil price at time t. Since the first year of price data in my data set in 1965, for these simulations $t = 0$ corresponds to1965, when the real world oil price was \$4.51/barrel. However, because pinning down the initial price rather than initial stock often leads to solutions with negative stocks, and because the qualitative features of the results are robust to the type of initial condition chosen, the results of the simulations in which an initial condition was imposed on price are not reported here.

[17] This particular base year was chosen because I wanted to use a year after the 1973 Arab oil embargo and becuase 1981 corresponds to the first year of a monthly data set on world oil price and quantity that I have and might later use. For the most part, however, the choice of 1981 for a base year was fairly arbitrary.

5 Simulation Results

In this section, I present my results under the two market structures of perfect competition and monopoly for both isoelastic demand and linear demand, and compare my results to the actual data. Although my model generates a solution for the years 1857-2300, I use my model's simulated solution for the years 1965-2001, the period spanned by my actual data, when comparing my model to data.

5.1 Isoelastic Demand

5.1.1 Results under perfect competition when demand is isoelastic

In analyzing the results under perfect competition when demand is isoelastic, I first attempt to discern which elasticity in the range $\varepsilon \in [-2, 0]$ generates trajectories for market price and for extraction that best fit the data. I use three different measures to examine the fit of the model.

My first measure of fit is based on the summary statistics. A solution fits the data well if its summary statistics are similar to those of the actual data. Figures 1 and 2 plot the means, standard deviations, maxima, and minima over the years 1965-2001 of the optimal values for market price $P(t)$ and extraction $E(t)$, respectively, as a function of the demand elasticity ε. Error bars denote the standard deviation. For comparison, the dotted lines in the two figures show the same summary statistics over the same years for the actual world oil price and world oil consumption, respectively. For all values of the demand elasticity considered, the model's optimal market price lies above the actual price levels. The model's optimal extraction falls in the range of actual data when demand is less elastic, but is lower than actual demand when demand is more elastic. When demand is inelastic, the model matches the mean extraction well, but fails to capture the variance in actual extraction. Thus, when a comparison of summary statistics is used as a measure of fit, $\varepsilon = 0$ appears to best fit the data.

In addition to a summary statistic comparison, my second measure of fit is the mean squared error (MSE). A lower MSE indicates a better fit. Figure 3 presents the MSE between the model's solution market price and actual market price for the period 1965-2001. The MSE is lowest at $\varepsilon = 0$ and at $\varepsilon = -2$. Figure 4 presents the MSE between the model's solution extraction and actual world oil consumption for the period 1965-2001. Unlike for market price, the MSE for extraction is lower when demand is less elastic. Using MSE as a measure of fit, either $\varepsilon = 0$ or $\varepsilon = -2$ appears to best fit the data.

My third measure of fit is correlation. A higher correlation indicates a better fit. Figure 5 presents the correlation coefficient between the my model's solution and actual data over 1965-2001 for both market price and extraction as a function of the demand elasticity. The price correlation is low and decreases slightly as demand becomes more inelastic. Predicted extraction is highly negatively correlated with actual extraction and invariant to elasticity except when $\varepsilon = 0$, in which case the correlation is negligible (correlation = -0.00). Using correlation as a measure of fit, $\varepsilon = 0$ appears to best fit the data.

Based on my three measures of fit, it appears that perfect competition with isoelastic demand yields results that best fit the data when the demand elasticity ε is either 0 or -2.

Figures 6 and 7 plot the solution trajectories for market price $P(t)$, cumulative extraction $X(t)$, shadow price $p(t)$, and extraction $E(t)$ when the demand elasticity ε is 0 and 2, respectively. For both elasticities, market price is increasing while extraction is weakly decreasing, and both trajectories are monotonic. In contrast, actual world oil price is trendless but highly volatile, actual world oil consumption is increasing, and neither trajectory is strictly monotonic. The model appears to best fit the data when $\varepsilon = 0$. For both elasticities, the shadow price eventually goes to zero so that at terminal time T, the resource is no longer scarce.

How much oil will eventually be extracted? Figure 8 plots the total stock extracted $X(T)$ as a function of demand elasticity. The less elastic the demand, the more oil will be extracted in total. If $\varepsilon = 0$, a total of 9.67 trillion barrels of oil will eventually be extraction; if $\varepsilon = -2$, a total of 4.07 trillion barrels will be extracted.

I now compare my results from perfect competition when demand is isoelastic to those from monopoly when demand is isoelastic.

5.1.2 Results under monopoly when demand is isoelastic

As explained above, when demand is isoelastic, a monopolist would only extract from a reserve when $\varepsilon < -1$. The results in this section are thus generated for $\varepsilon \in [-2.0, -1.1]$.

I first find the elasticities that yield results that best fit the data under my three measures of fit: summary statistic comparison, MSE and correlation. As seen in Figure 9, the model yields market price profiles that best match the summary statistics of actual data for the years 1965-2001 when $\varepsilon = -2.0$. As seen in Figure 10, extraction best matches the summary statistics of actual data when $\varepsilon = -1.1$. If MSE is the measure of fit, then market price is best fit to data when $\varepsilon = -2.0$ (Figure 11), and extraction is best fit to data when $\varepsilon = -1.2$ (Figure 12). If correlations are used as a measure of fit, all elasticities yield results that fit the data equally well (Figure 13), with predicted market price weakly correlated with actual data and predicted extraction strongly anti-correlated with actual data. Thus, the model appears to best fit the data when $\varepsilon = -1.1$, $\varepsilon = -1.2$ and $\varepsilon = -2.0$, though the fit is poor in all cases.

Figures 14, 15 and 16 present the results under the best-fit elasticities of $\varepsilon = -1.1$, $\varepsilon = -1.2$ and $\varepsilon = -2.0$, respectively. As with perfect competition under isoelastic demand, monopoly under isoelastic demand yields increasing market price and decreasing extraction. As in perfect competition, the shadow price eventually goes to zero so that at terminal time T, the resource is no longer scarce.

Figure 17 plots the total stock extracted as a function of elasticity. In contrast to perfect competition, the total stock extracted increases as demand becomes more elastic. Moreover, for each elasticity, the monopolist extracts a smaller amount in total than does a perfectly competitive industry. Just as in a standard static market, a monopolist produces less for each time-t market than would a perfectly competitive industry in order to increase prices and thus his revenue.[18]

In comparing the best-fit scenarios for perfect competition with those for monopoly, both when demand is isoelastic, it appears that the model is best fit by perfect competition with inelastic demand. I now examine the results when demand is linear rather than isoelastic.

5.2 Linear Demand

5.2.1 Results under perfect competition when demand is linear

How well does perfect competition fit the data when demand is linear, and under which demand elasticity is the data best fit? Figures 18 and 19 plot the summary statistics for the predicted price and extraction, respectively, for the years 1965-2001, and compares them to the analogous statistics for the actual data. The predicted market price is too high for all the elasticities, while extraction best fits the range of the actual data when demand is inelastic (i.e., $\varepsilon = 0$). When MSE is used to measure fit, market price is best matched to data when $\varepsilon = -2$ (Figure 20), while extraction is best matched to data when $\varepsilon = 0$ (Figure 21). When correlation is used to measure fit, the fit of market price to data decreases slightly as demand becomes more inelastic, while extraction is best fit to data when $\varepsilon = 0$ (Figure 22). Thus, perfect competition under linear demand appears to best fit the data when $\varepsilon = 0$ or $\varepsilon = -2$.

Figures 23 and 24 plot the solution trajectories for the best-fit scenarios of $\varepsilon = 0$ and $\varepsilon = -2$, respectively. In both scenarios, market price is weakly increasing and extraction is weakly increasing. As with isoelastic demand, perfect competition best fits the data when $\varepsilon = 0$. Also as before, the shadow price goes to zero.

Figure 25 plots the total stock extracted as a function of elasticity. As before, more total stock is extracted when demand is less elastic. When compared with the analogous isoelastic demand results, the total stock extracted is weakly lower when demand is linear than when it is isoelastic for any given elasticity.

[18] In general, a monopolist would extract more slowly than would a perfectly competitive industry. One exception is that in the case of isoelastic demand with zero extraction costs, monopoly and perfect compeition yield identical solutions for both price and extraction (Stiglitz, 1976).

5.2.2 Results under monopoly when demand is linear

I now examine the fit of the model under monopoly and linear demand.[19] According to the summary statistics, model price best fits the data when $\varepsilon = -2$ (Figure 26). Model extraction best fits the mean of the data when $\varepsilon = -0.7$ and the variance of the data when $\varepsilon = -2$ (Figure 27). According to MSE, model price best fits data when $\varepsilon = -2$ (Figure 28), while model extraction best fits the data when $\varepsilon = -0.7$ (Figure 29). According to correlations, all elasticities yield roughly the same fit to data (Figure 30). The model thus appears to best fit data when $\varepsilon = -0.7$ or $\varepsilon = -2$.

Figures 31 and 32 plot the solution trajectories for the best-fit scenarios of $\varepsilon = -0.7$ and $\varepsilon = -2$, respectively. For both scenarios, market price is weakly increasing while extraction is weakly decreasing. As with the previous specifications, the shadow price eventually goes to zero so that at terminal time T, the resource is no longer scarce.

Figure 33 plots the total stock extracted as a function of elasticity. As before, total stock extracted increases as demand becomes more inelastic, and the monopolist extracts less than the perfectly competitive industry for any given elasticity under linear demand.

From among the various specifications I tried, the model best fits actual data when the oil market is perfectly competitive and when demand is inelastic. For perfect competition and inelastic demand, the model predictions for price and extraction for the years of actual data are compared with actual data in Figures 34 and 35 for isoelastic and linear demand, respectively. Under these assumptions, real oil price should fall in the range \$117-160/barrel over the years 2000-2010 and extraction should be roughly constant at 59.8 million barrels/day over the entire simulation period 1857-2300. Even for the best-fit specification, however, the simple Hotelling model fails to adequately explain the data.

6 Conclusion

This paper uses data on world oil price and consumption to calibrate a Hotelling model of optimal resource extraction with unlimited potential reserves when costs exhibit stock effects. Which specification for demand (isoelastic or linear), market structure (perfect competition or monopoly) and demand elasticity yields results that best fit the historical data? From among the various specifications I tried, the model best fits actual data when the oil market is perfectly competitive and when demand is inelastic. Under these assumptions, real oil price should fall in the range \$117-160/barrel over the years 2000-2010 and extraction should be roughly constant at 59.8 million barrels/day over the entire simulation period 1857-2300.

[19] For these simulations, I vary ε from -0.1 to -2.0 since I was unable to generate a solution for the case $\varepsilon = 0$.

So, how well does the simple Hotelling model presented appear to explain historical data? Unfortunately, the answer is somewhat disappointing. Based on my three measures of fit (summary statistic comparison, mean squared error and correlation), none of the simulations appear to adequately explain the historical data. In particular, the theoretical model fails to capture two salient features of data. The first feature that is not captured by the theoretical model is the highly volatile but roughly trendless nature of market prices. The second feature that is not captured by the model is the increasing trend in extraction.

Thus, it seems that a basic Hotelling model with unlimited potential reserves and costs that exhibit stock effects fails to explain the historical data on world oil prices and world oil consumption. In order to better reconcile the theory with data, one may therefore need to augment the basic model. Possible modifications include a different estimate of the cost function, technological progress, oligopoly, time-varying demand, and uncertainty; such models will be the subject of future work.

References

[1] Berndt, E.R., & Wood, D.O. (1975). Technology, prices, and the derived demand for energy. The Review of Economics and Statistics, 57 (3), 259-268.

[2] Chakravorty, U., Roumasset, J., & Tse, K. (1997). Endogenous substitution among energy resources and global warming. The Journal of Political Economy, 105 (6), 1201-1234.

[3] Deffeyes, K.S. (2001). Hubbert's peak: the impending world oil shortage. Princeton, NJ: Princeton University Press.

[4] Devarajan, S., & Fisher, A.C. (1982). Exploration and scarcity. The Journal of Political Economy, 90 (6), 1279-1290.

[5] Edmonson, N. (1975). Real price and the consumption of mineral energy in the United States, 1901-1968. The Journal of Industrial Economics, 23 (3), 161-174.

[6] Farrow, S. (1985). Testing the efficiency of extraction from a stock resource. The Journal of Political Economy, 93 (3), 452-487.

[7] Halvorsen, R., & Smith, T.R. (1984). On measuring natural resource scarcity. The Journal of Political Economy, 92 (5), 954-964.

[8] Halvorsen, R., & Smith, T.R. (1991). A test of the theory of exhaustible resources. The Quarterly Journal of Economics, 106 (1), 123-140.

[9] Hanson, D.A. (1980). Increasing extraction costs and resource prices: Some further results. The Bell Journal of Economics, 11 (1), 335-342.

[10] Hoel, M. (1978). Resource extraction, uncertainty, and learning. The Bell Journal of Economics, 9 (2), 642-645.

[11] Hotelling, H. (1931). The economics of exhaustible resources. The Journal of Political Economy, 39 (2), 137-175.

[12] Khalatbari, F. (1977). Market imperfections and the optimum rate of depletion of natural resources. Economica, 44 (176), 409-414.

[13] Krautkraemer, J.A. (1998). Nonrenewable resource scarcity. Journal of Economic Literature, 36 (4), 2065-2107.

[14] Lasserre, P. (1985). Discovery costs as a measure of rent. The Canadian Journal of Economics, 18 (3), 474-483.

[15] Lin, C.-Y.C. (2004a). Efficient and consistent estimation of static supply and demand in the world oil market: A dry hole? Mimeo. Harvard University.

[16] Lin, C.-Y.C. (2004b). Hotelling revisited: Oil prices and endogenous technological progress. Mimeo. Harvard University.

[17] Lin, C.-Y.C. (2004c). Steady-state growth in a Hotelling model of resource extraction. Mimeo. Harvard University.

[18] Miller, M.H., & Upton, C.W. (1985). A test of the Hotelling valuation principle. The Journal of Political Economy, 93 (1), 1-25.

[19] Nordhaus, W.D. (1980). Oil and economic performance in industrial countries. Brookings Papers on Economic Activity, 1980 (2), 341-399.

[20] Pindyck, R.S. (1978). The optimal exploration and production of nonrenewable resources. The Journal of Political Economy, 86 (5), 841-861.

[21] Pindyck, R.S. (1980). Uncertainty and exhaustible resource markets. The Journal of Political Economy, 88 (6), 1203-1225.

[22] Solow, R.M., & Wan, F.Y. (1976). Extraction costs in the theory of exhaustible resources. The Bell Journal of Economics, 7 (2), 359-370.

[23] Stiglitz, J.E. (1976). Monopoly and the rate of extraction of exhaustible resources. The American Economic Review, 66 (4), 655-661.

[24] Sweeney, J.L. (1977). Economics of depletable resources: Market forces and intertemporal bias. The Review of Economic Studies, 44 (1), 124-141.

[25] Tietenberg, T.H. (1996). Environmental and natural resource economics (4th ed.). New York: Harper-Collins.

[26] U.S. Geological Survey [USGS]. (2004). United States energy and world energy production and consumption statistics. [Online: web]. Cited 7 April 2004. URL: http://energy.cr.usgs.gov/energy/stats_ctry/Stat1.html

[27] Weitzman, M.L. (2003). Income, wealth, and the maximum principle. Cambridge, MA: Harvard University Press.

[28] Young, D. (1992). Cost specification and firm behaviour in a Hotelling model of resource extraction. The Canadian Journal of Economics, 25 (1), 41-59.

[29] Zimmermann, E.W. (1951). World resources and industries: A functional appraisal of the availability of agricultural and industrial materials. New York: Harper and Brothers, Publishers.

FIGURE 1.

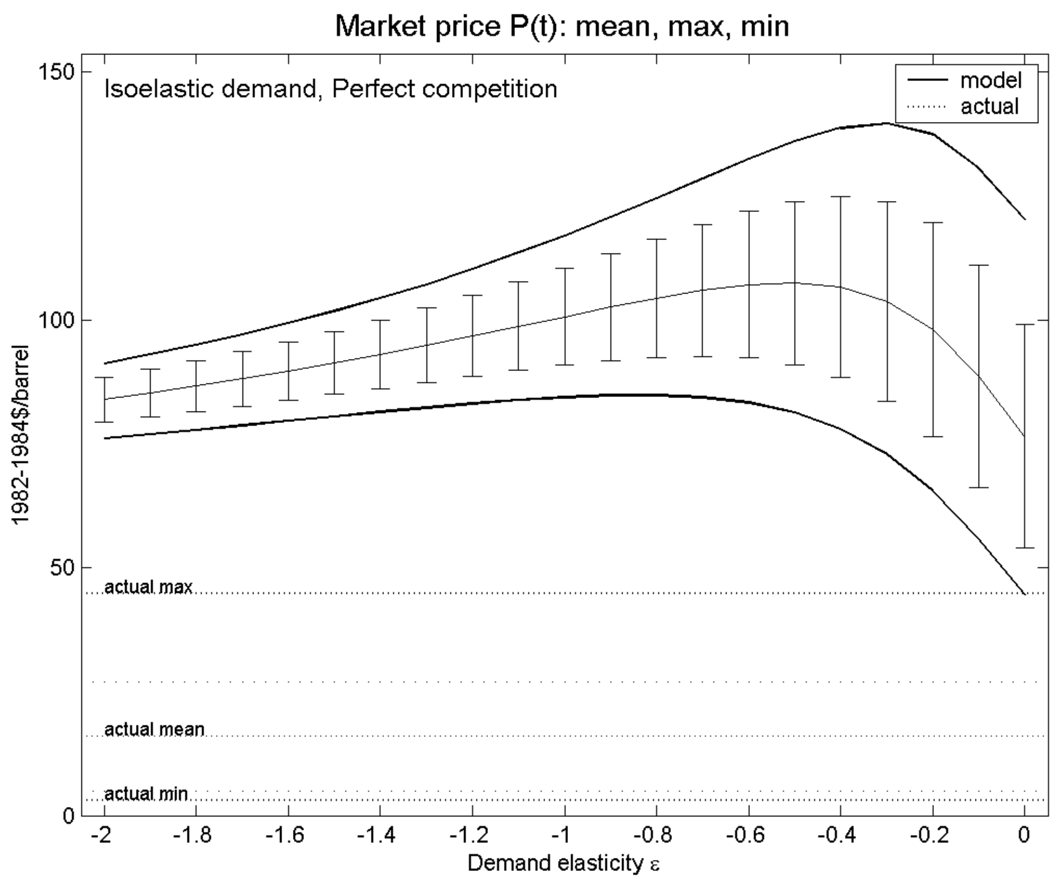
Market price P(t): mean, max, min
Isoelastic demand, Perfect competition
model
actual
actual max
actual mean
actual min
1982-1984$/barrel
Demand elasticity ε
150
100
50
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 2.

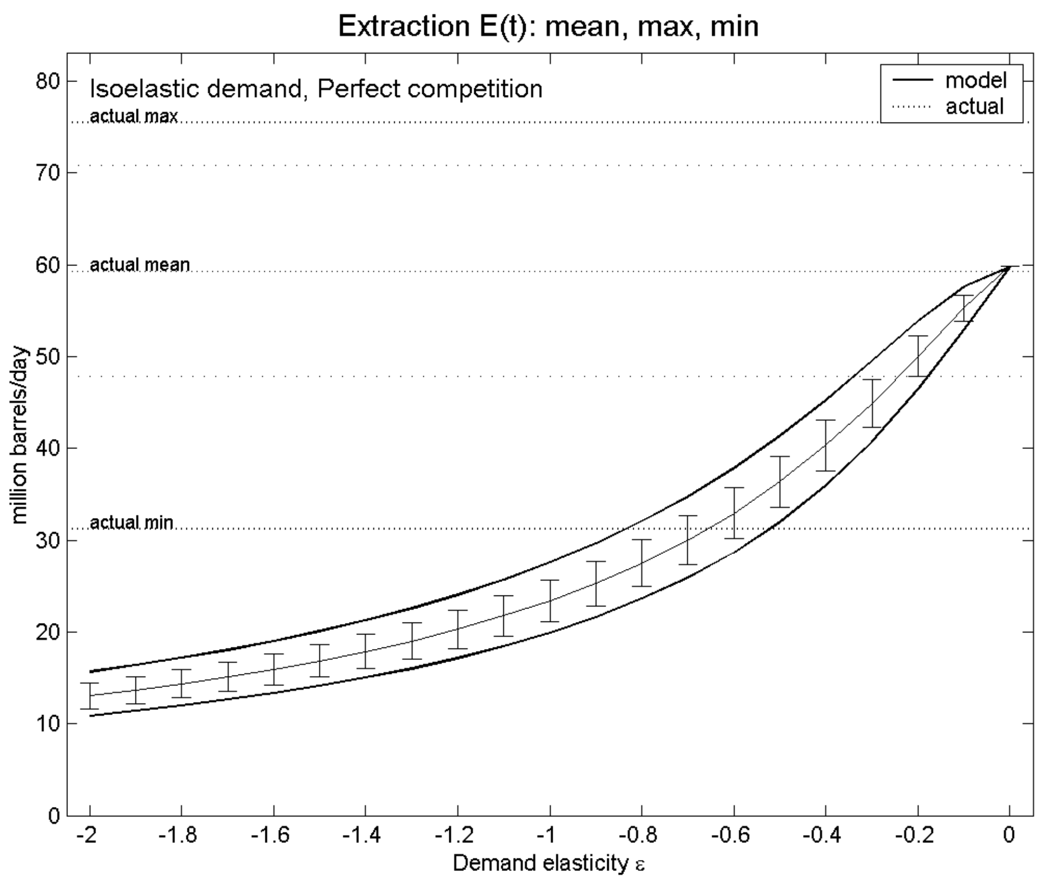
Extraction E(t): mean, max, min
Isoelastic demand, Perfect competition
model
actual
actual max
actual mean
actual min
million barrels/day
Demand elasticity ε
80
70
60
50
40
30
20
10
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 3.

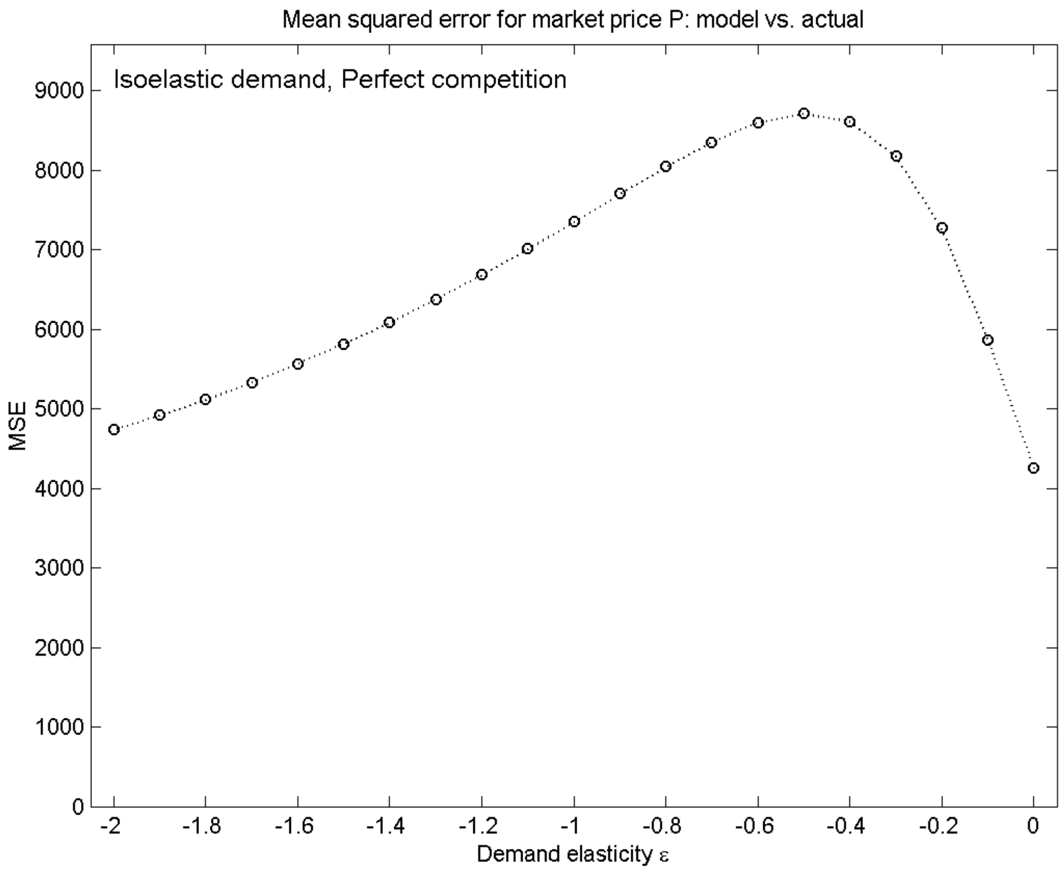
Mean squared error for market price P: model vs. actual
Isoelastic demand, Perfect competition
MSE
Demand elasticity ε
0
1000
2000
3000
4000
5000
6000
7000
8000
9000
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 4.

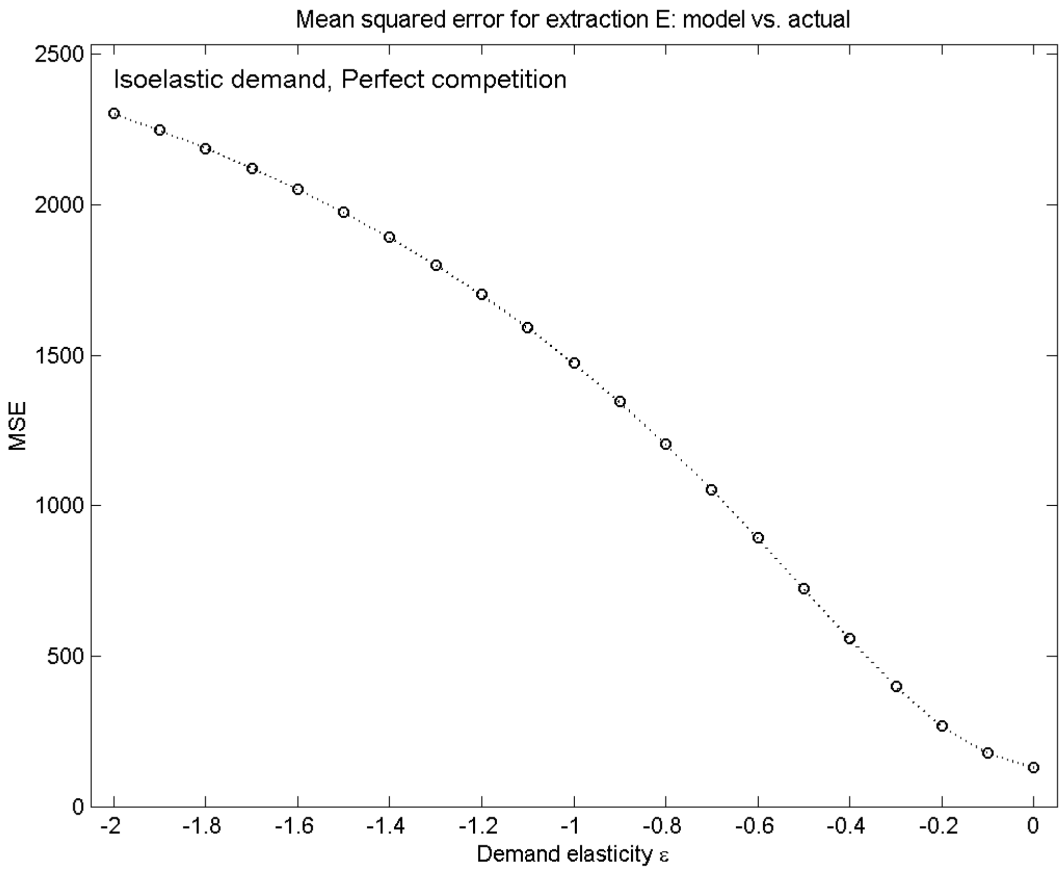
Mean squared error for extraction E: model vs. actual
Isoelastic demand, Perfect competition
MSE
0
500
1000
1500
2000
2500
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0
Demand elasticity ε

FIGURE 5.

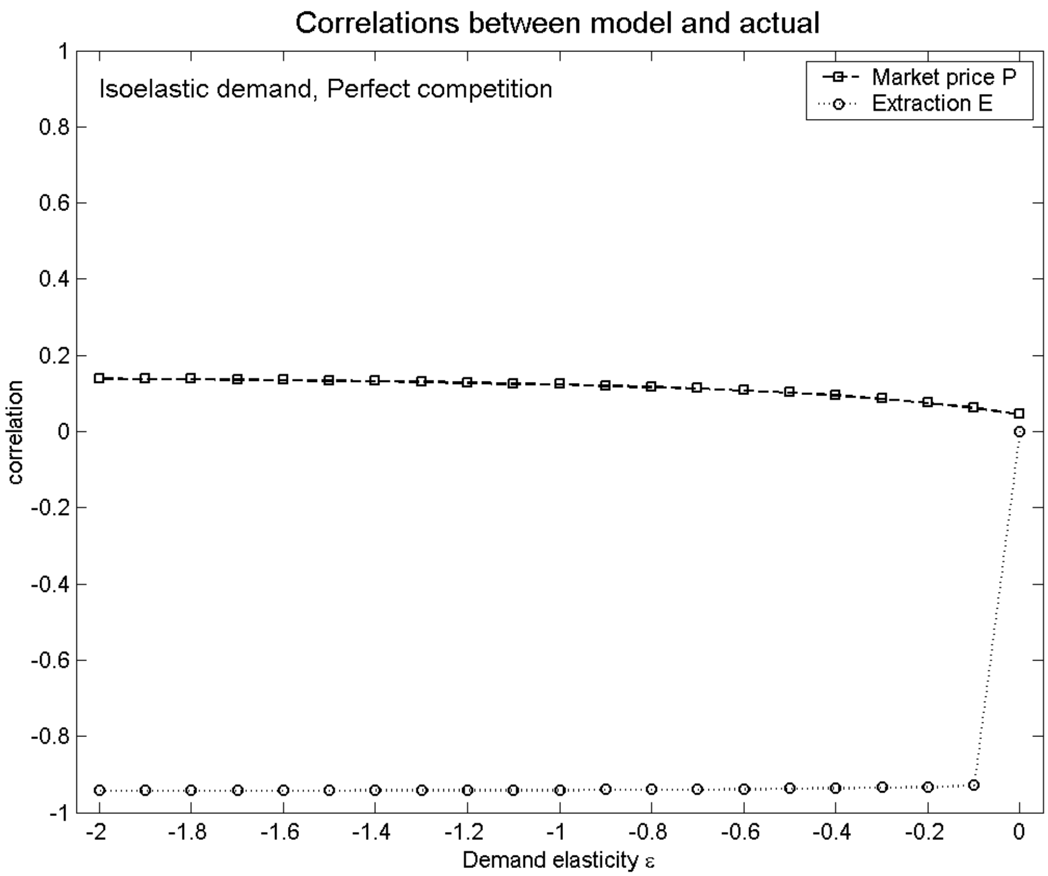
Correlations between model and actual
Isoelastic demand, Perfect competition
Market price P
Extraction E
correlation
Demand elasticity ε
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 6.

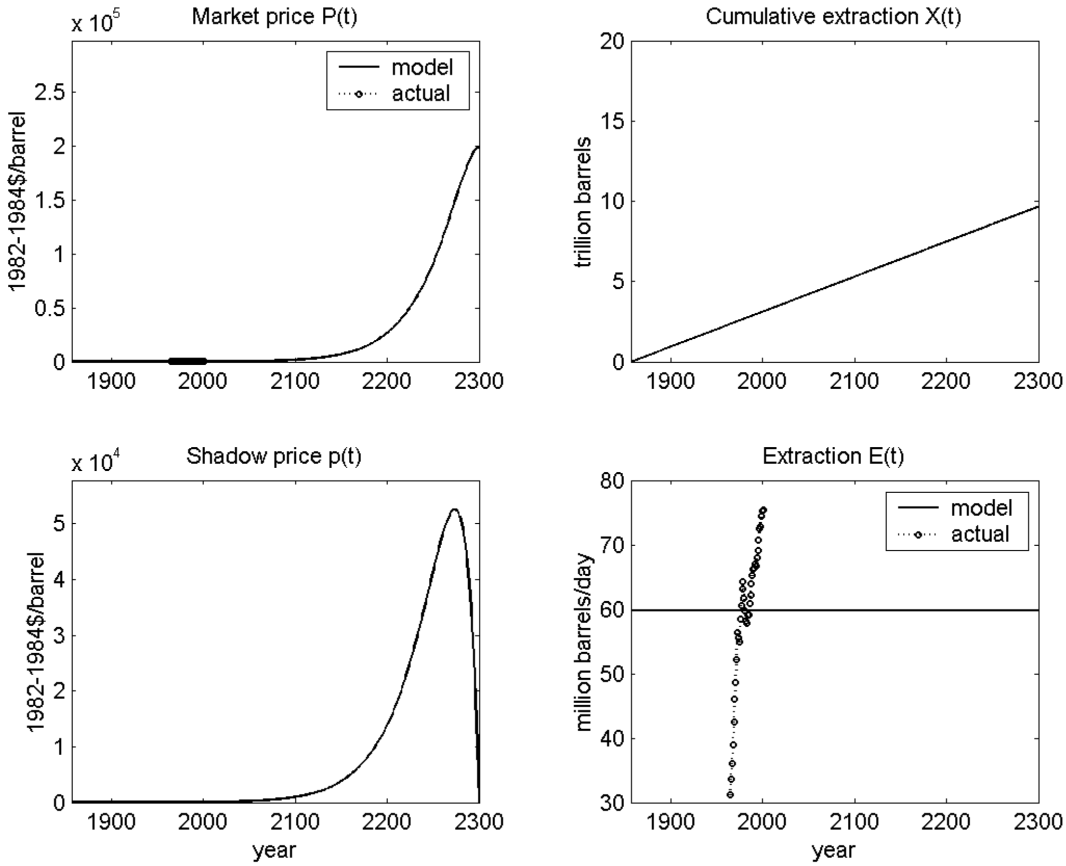

Perfect competition with isoelastic demand when demand elasticity ε is 0.

FIGURE 7.

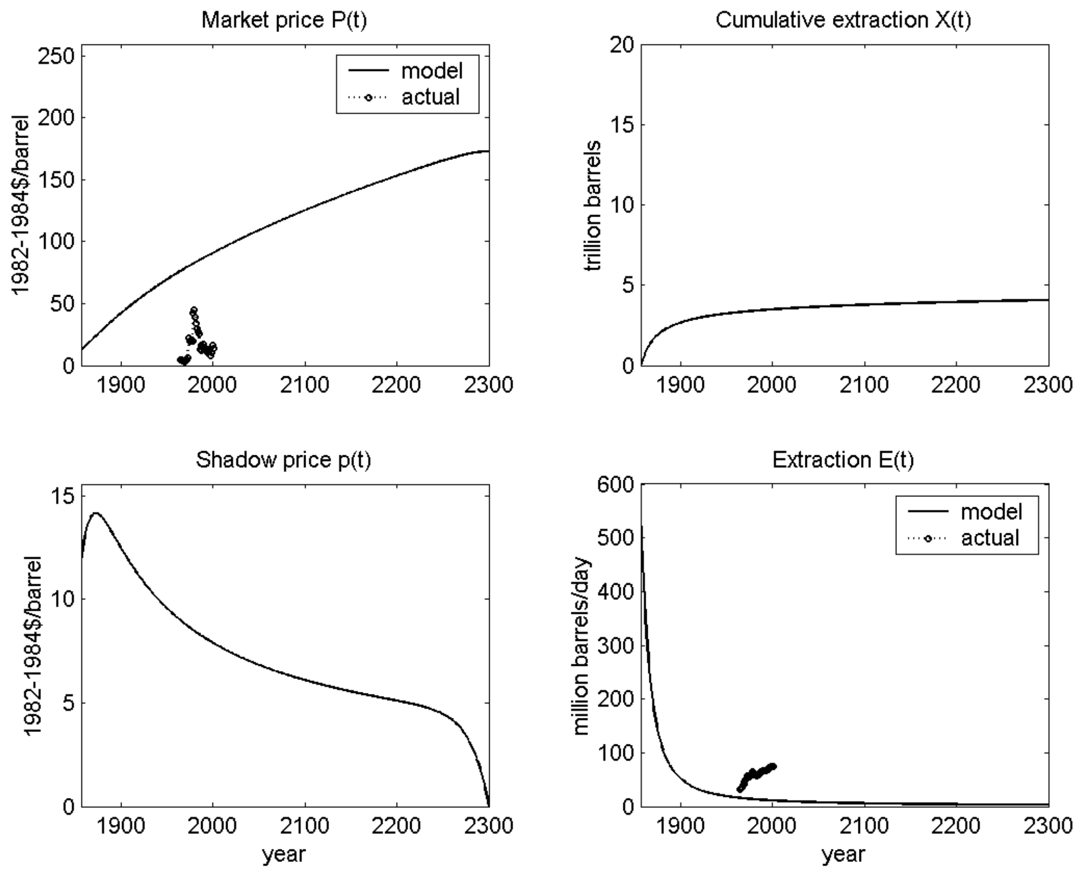

Perfect competition with isoelastic demand when demand elasticity ε is –2.0.

FIGURE 8.

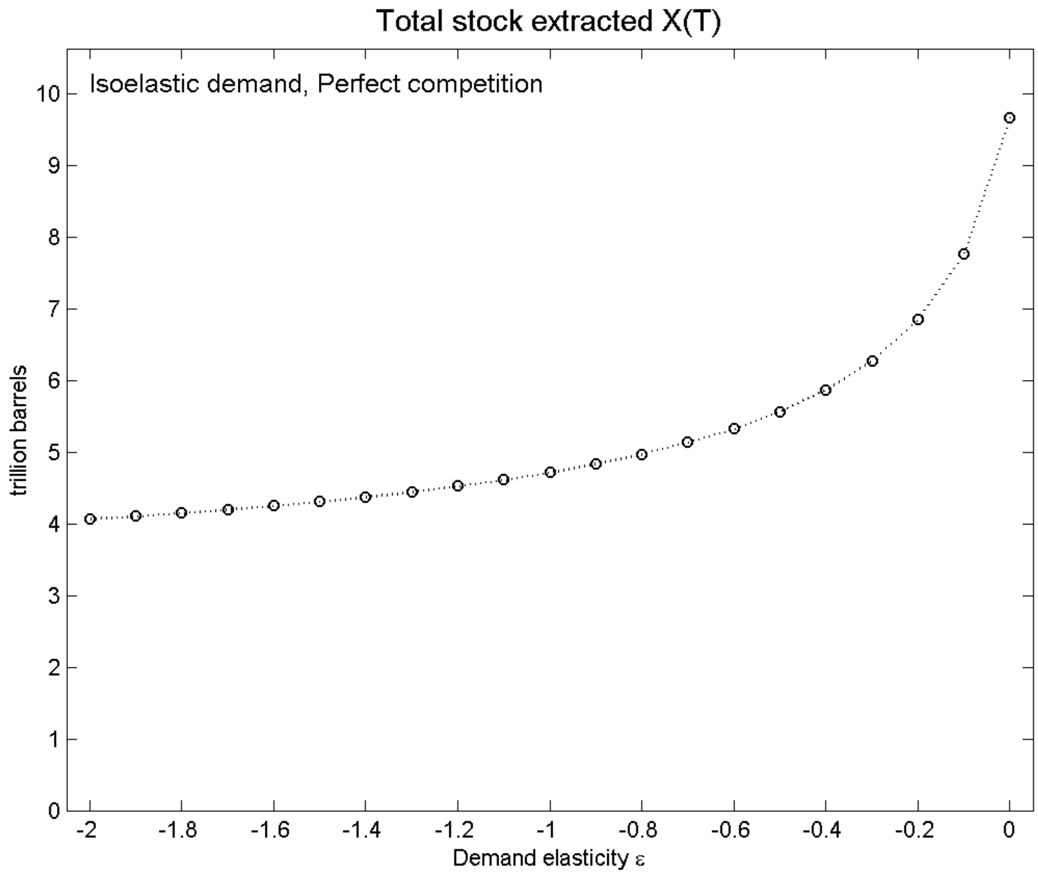
Total stock extracted X(T)
Isoelastic demand, Perfect competition
trillion barrels
Demand elasticity ε
0
1
2
3
4
5
6
7
8
9
10
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 9.

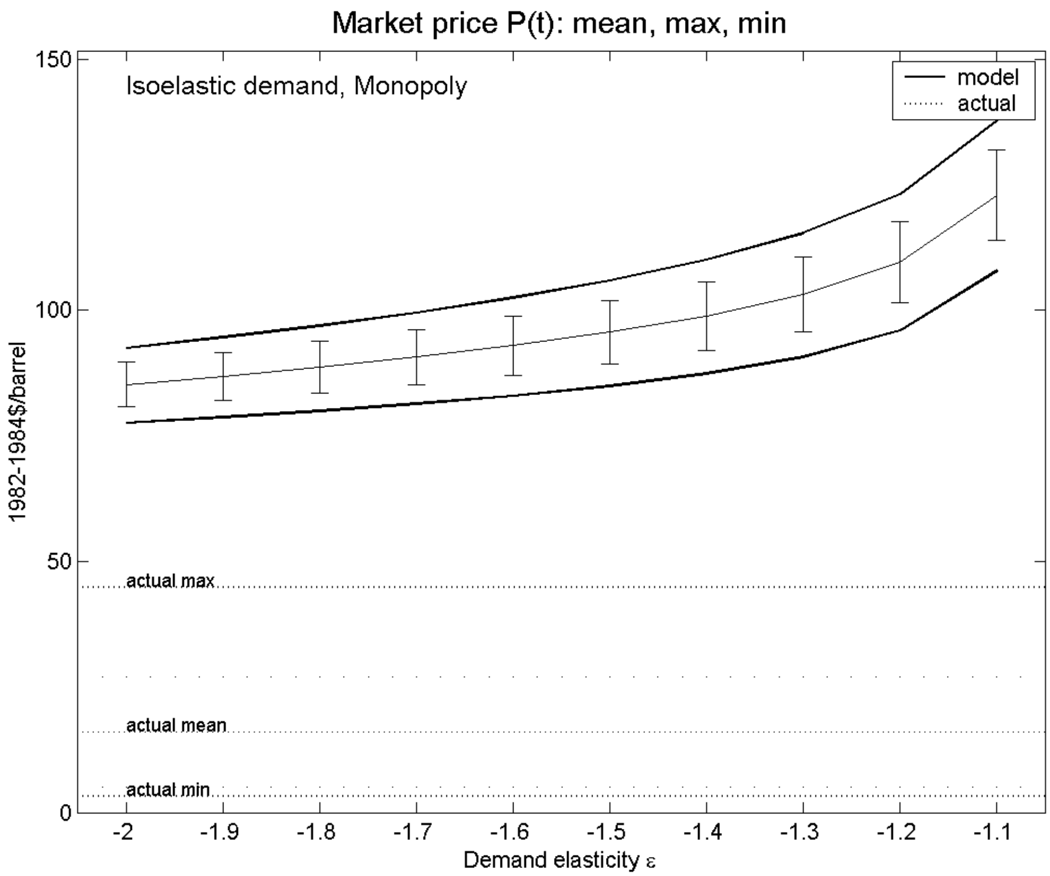
Market price P(t): mean, max, min
Isoelastic demand, Monopoly
model
actual
150
100
50
0
1982-1984$/barrel
actual max
actual mean
actual min
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1
Demand elasticity ε

FIGURE 10.

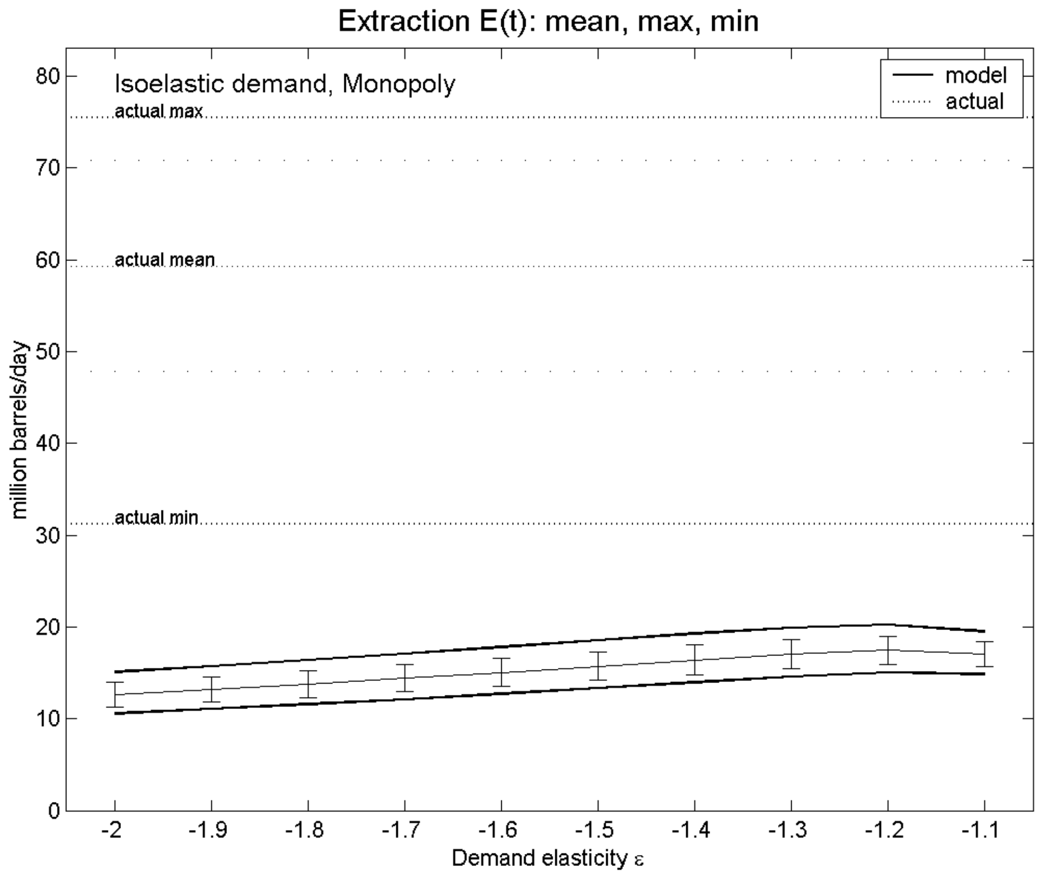
Extraction E(t): mean, max, min
Isoelastic demand, Monopoly
model
actual
actual max
actual mean
actual min
million barrels/day
Demand elasticity ε
80
70
60
50
40
30
20
10
0
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1

FIGURE 11.

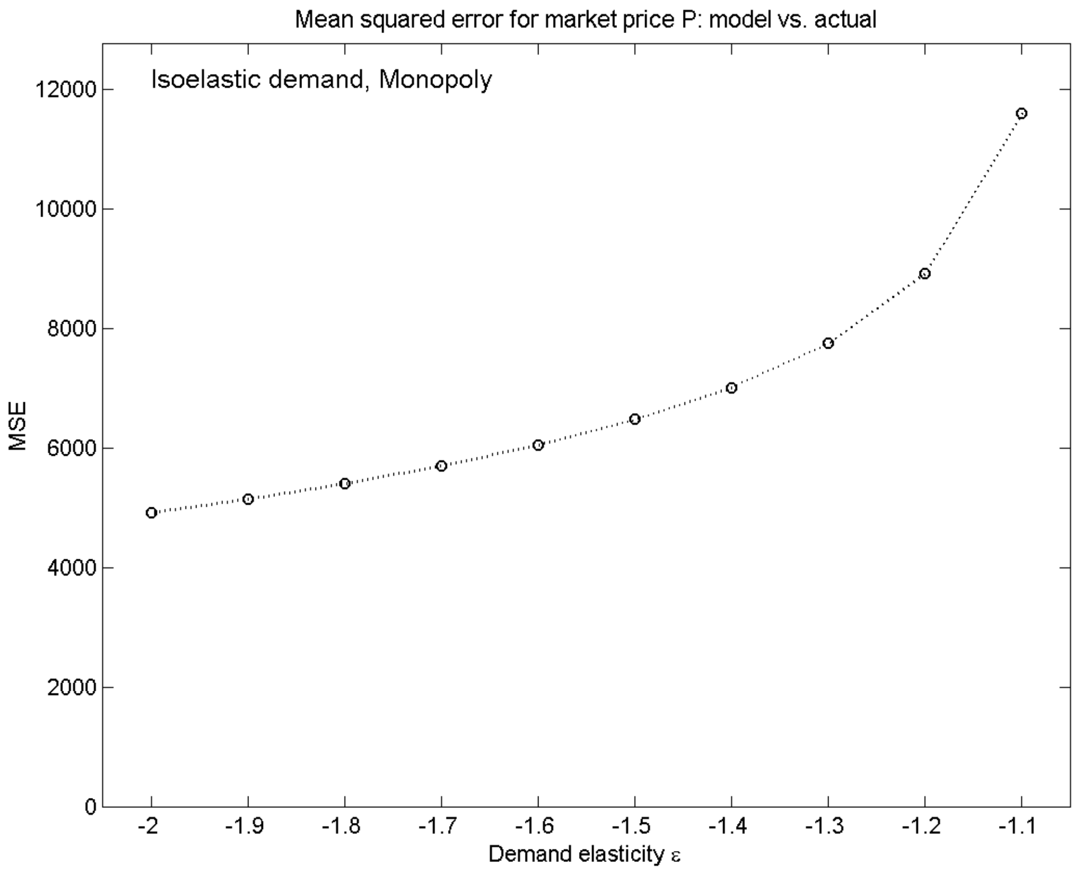
Mean squared error for market price P: model vs. actual
Isoelastic demand, Monopoly
MSE
Demand elasticity ε
0
2000
4000
6000
8000
10000
12000
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1

FIGURE 12.

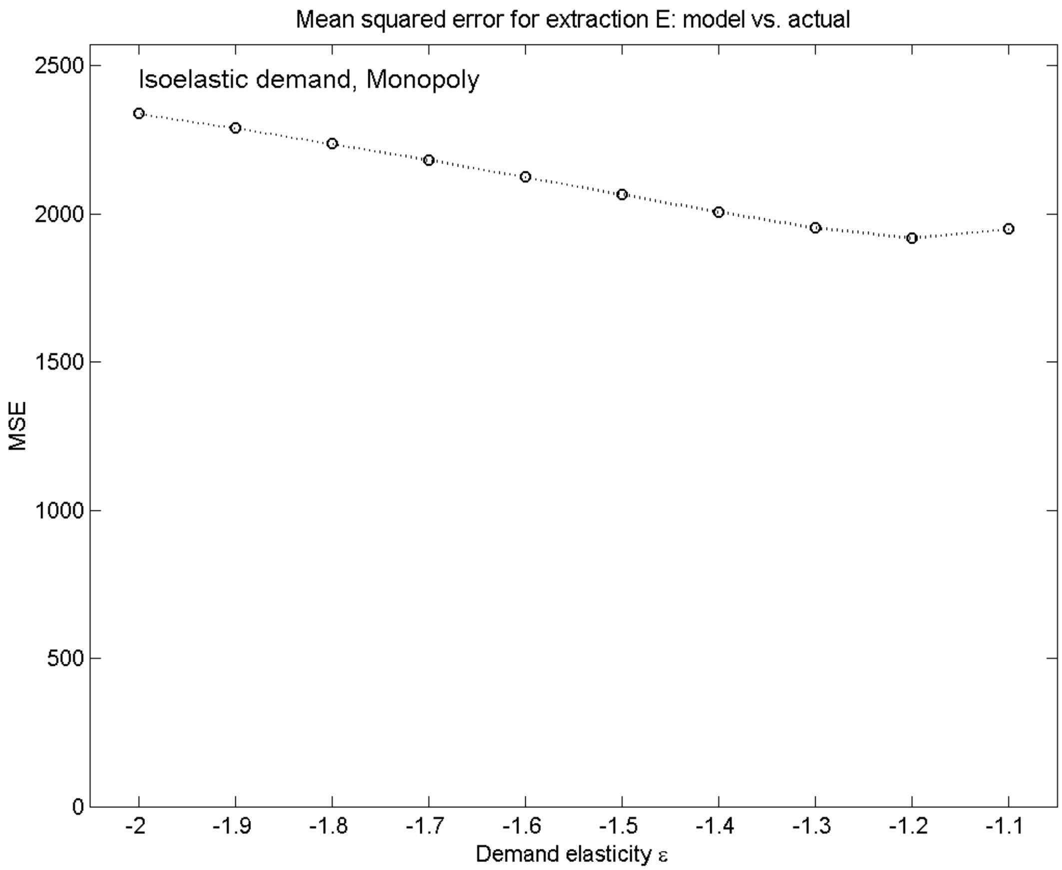
Mean squared error for extraction E: model vs. actual
Isoelastic demand, Monopoly
MSE
Demand elasticity ε
2500
2000
1500
1000
500
0
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1

FIGURE 13.

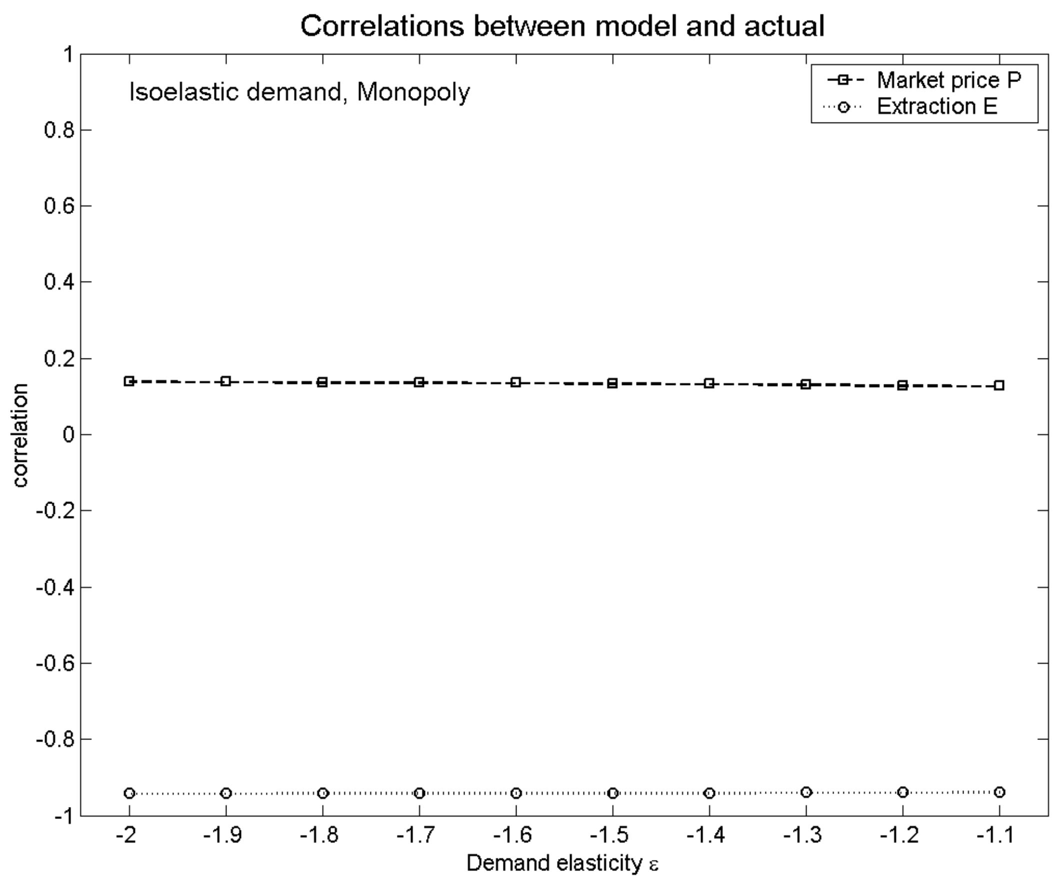
Correlations between model and actual
Isoelastic demand, Monopoly
Market price P
Extraction E
correlation
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1
Demand elasticity ε

FIGURE 14.

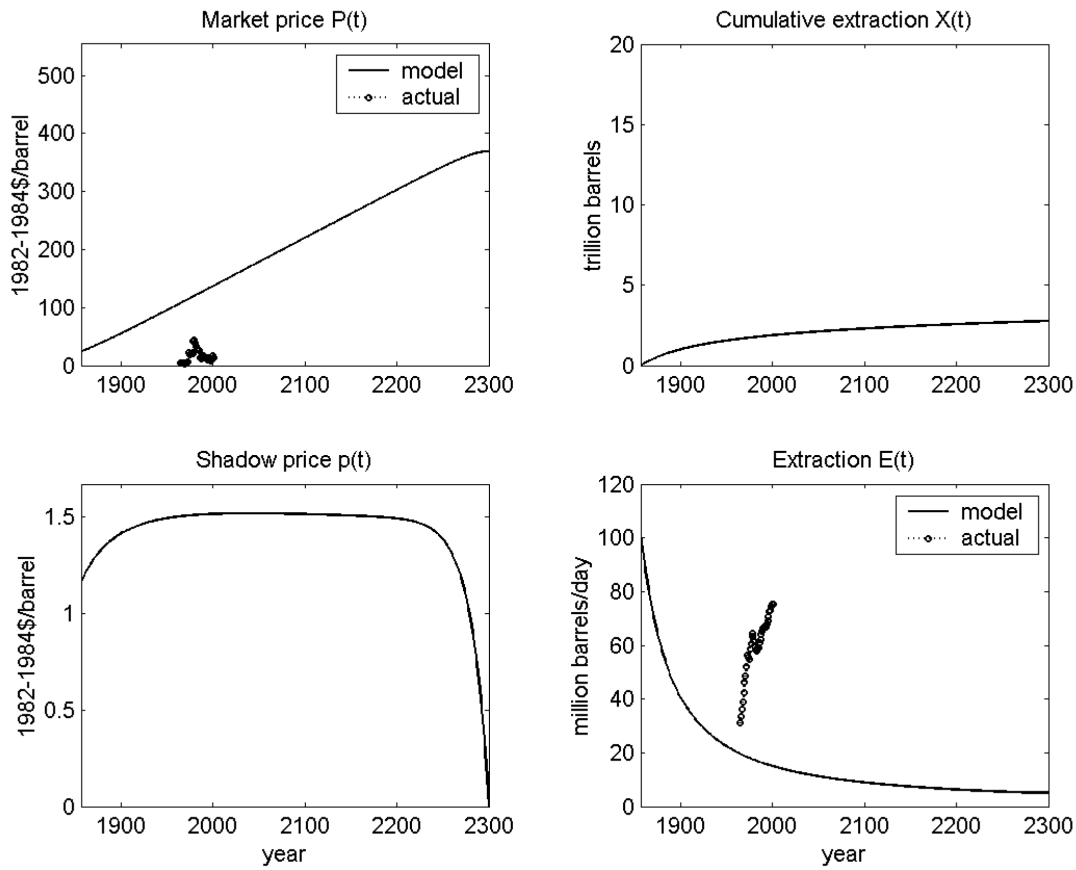

Monopoly with isoelastic demand when demand elasticity ε is –1.1.

FIGURE 15.

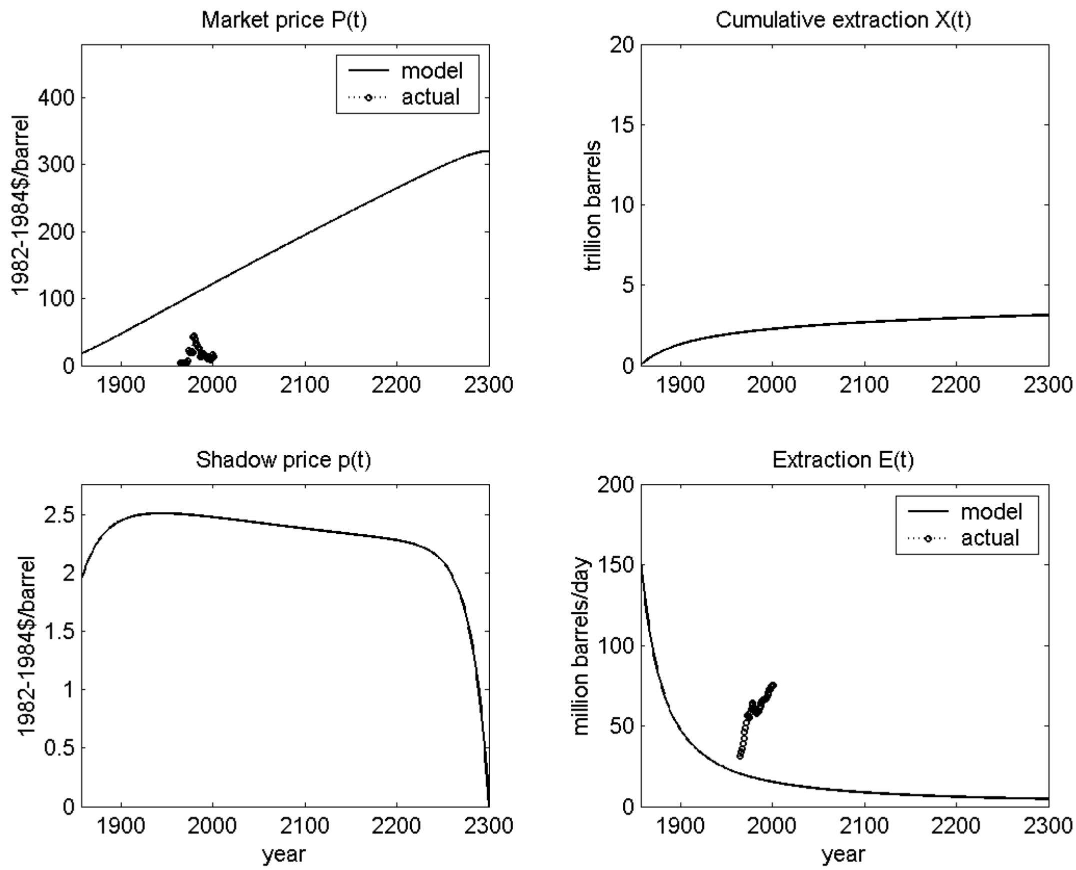

Monopoly with isoelastic demand when demand elasticity ε is –1.2.

FIGURE 16.

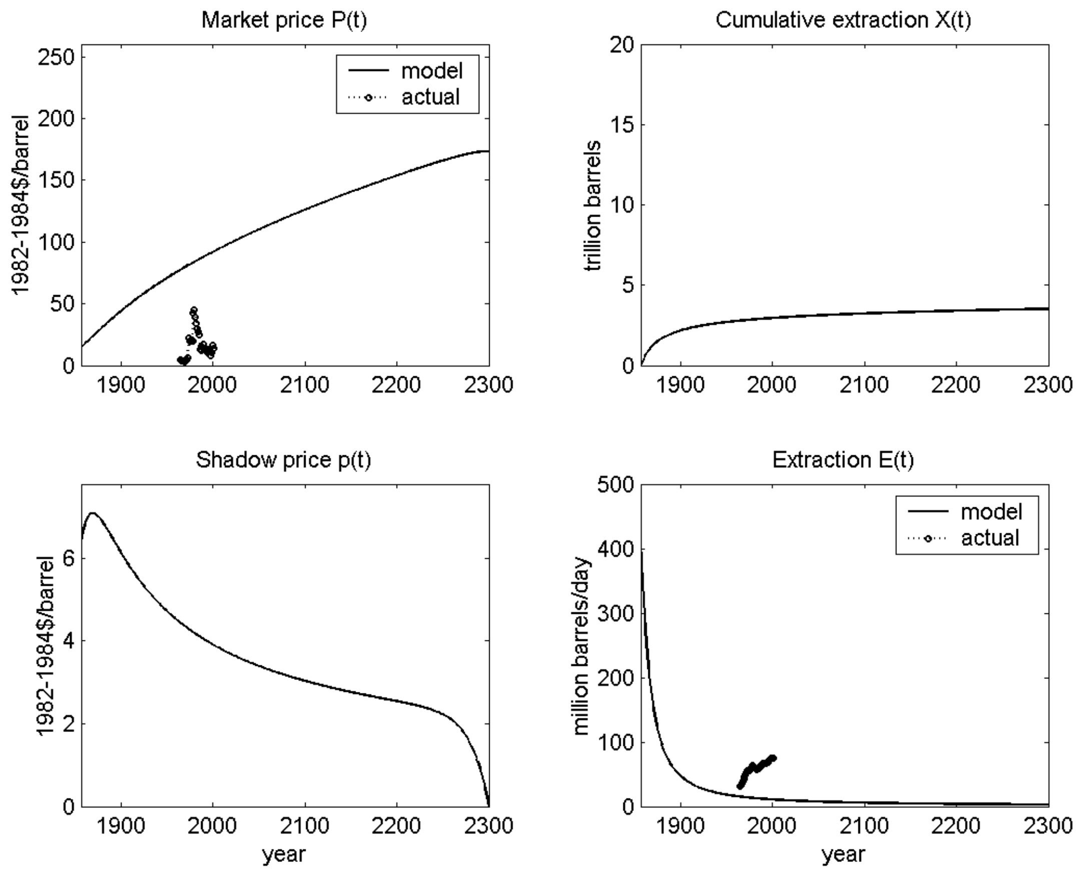

Monopoly with isoelastic demand when demand elasticity ε is –2.0.

FIGURE 17.

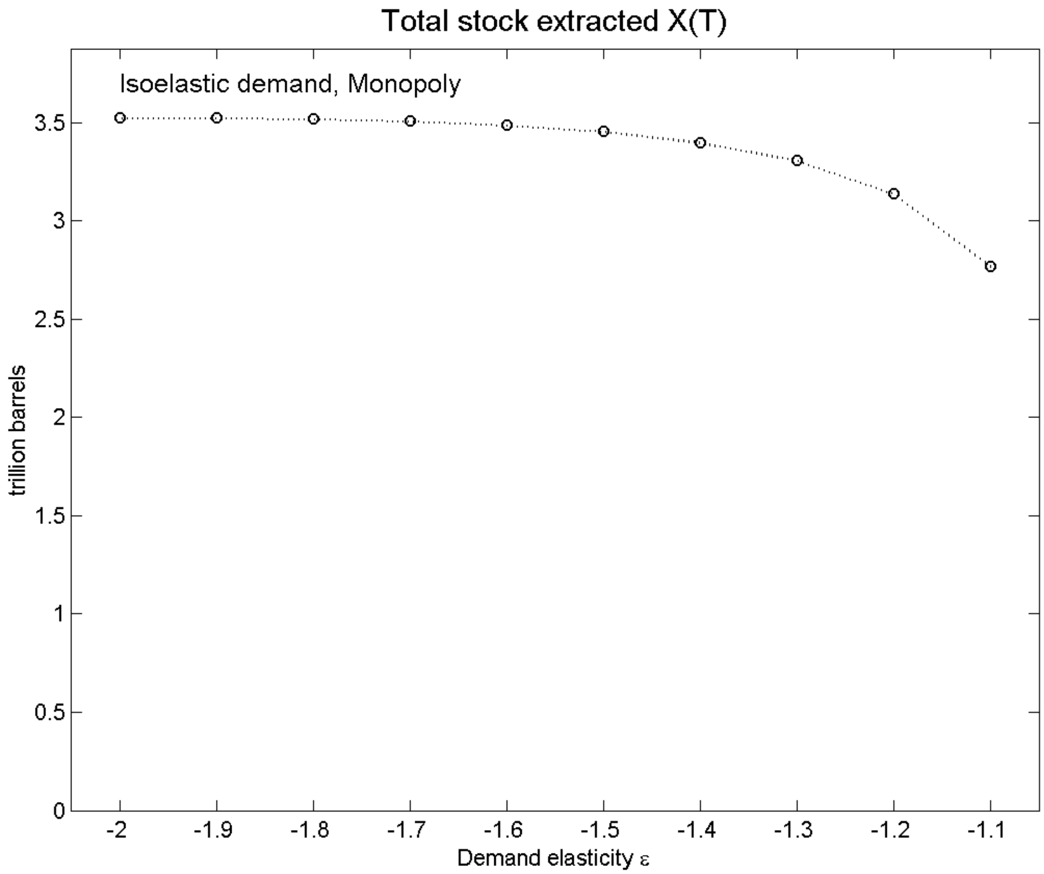
Total stock extracted X(T)
Isoelastic demand, Monopoly
trillion barrels
Demand elasticity ε
0
0.5
1
1.5
2
2.5
3
3.5
-2
-1.9
-1.8
-1.7
-1.6
-1.5
-1.4
-1.3
-1.2
-1.1

FIGURE 18.

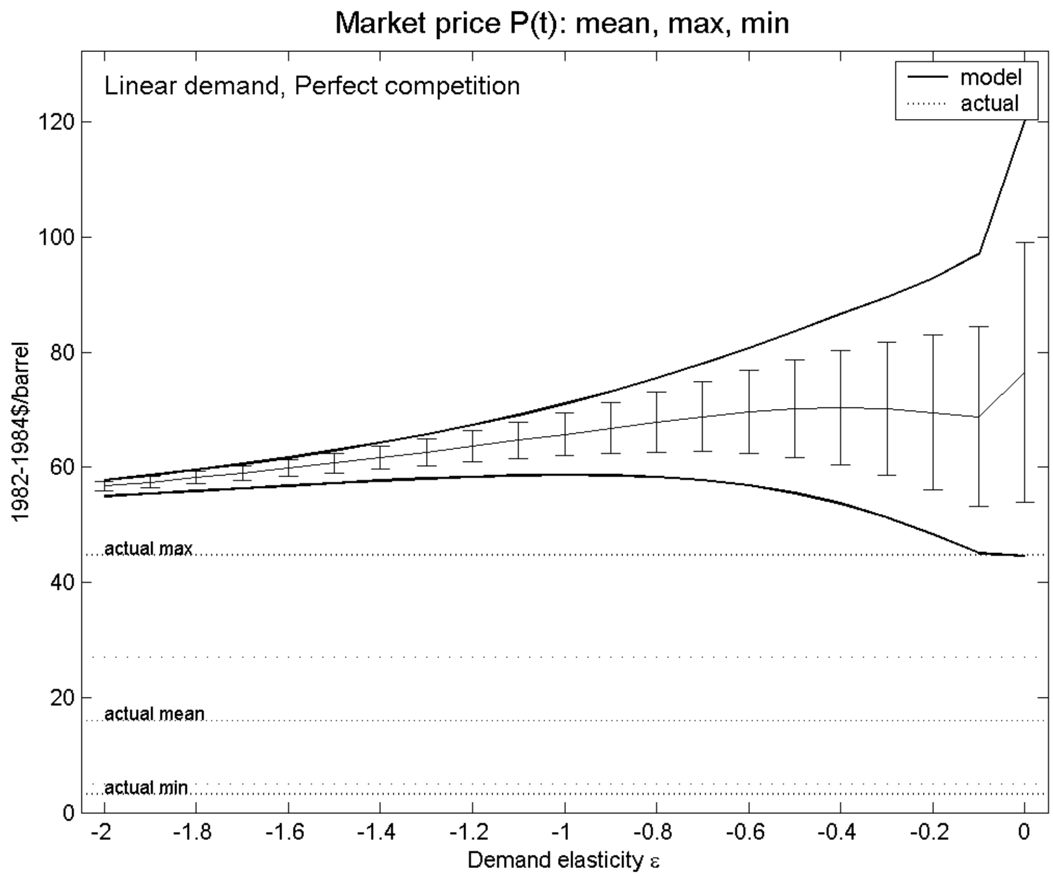
Market price P(t): mean, max, min
Linear demand, Perfect competition
model
actual
1982-1984$/barrel
120
100
80
60
40
20
0
actual max
actual mean
actual min
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0
Demand elasticity ε

FIGURE 19.

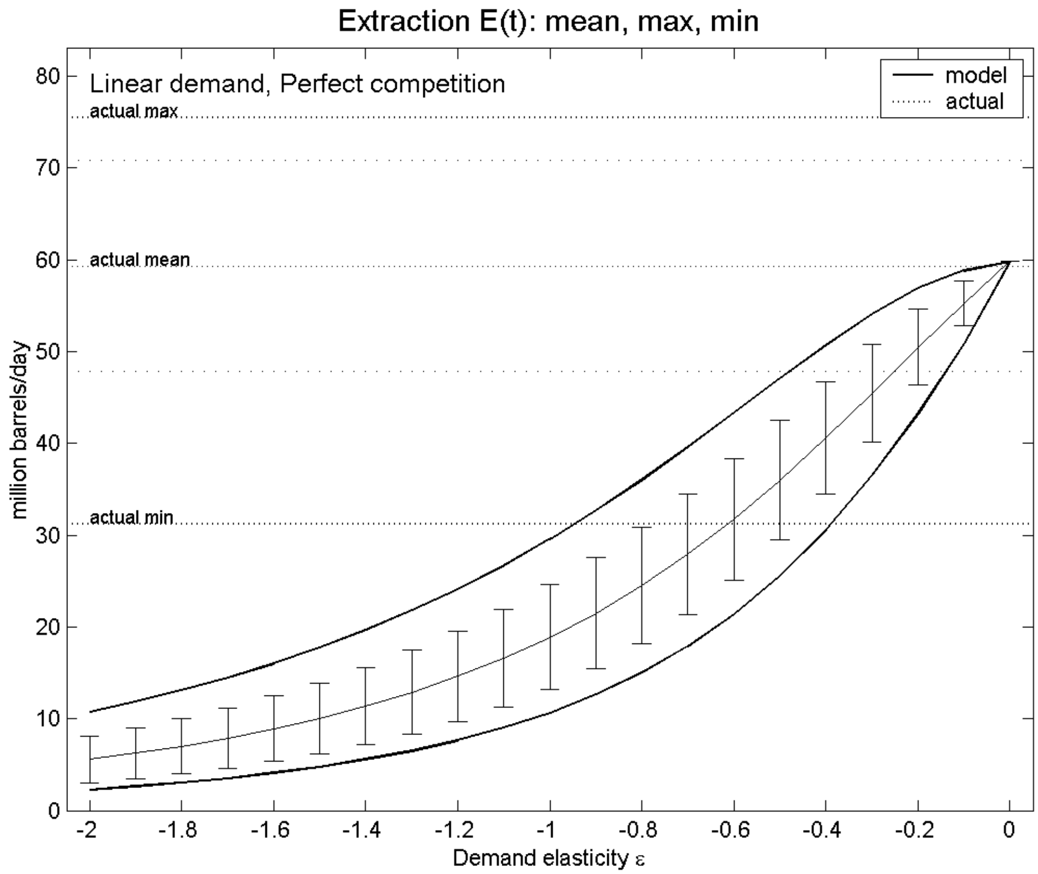
Extraction E(t): mean, max, min
Linear demand, Perfect competition
model
actual
actual max
actual mean
actual min
million barrels/day
Demand elasticity ε
80
70
60
50
40
30
20
10
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 20.

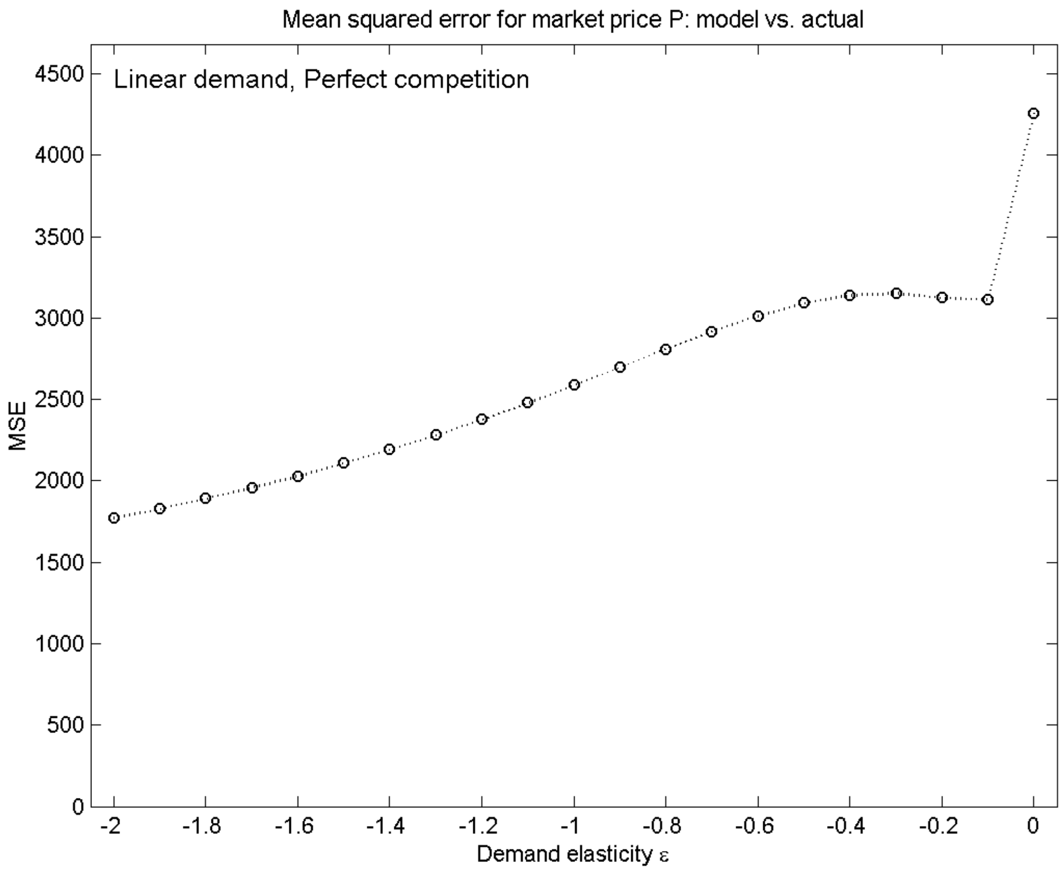
Mean squared error for market price P: model vs. actual
Linear demand, Perfect competition
MSE
4500
4000
3500
3000
2500
2000
1500
1000
500
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0
Demand elasticity ε

FIGURE 21.

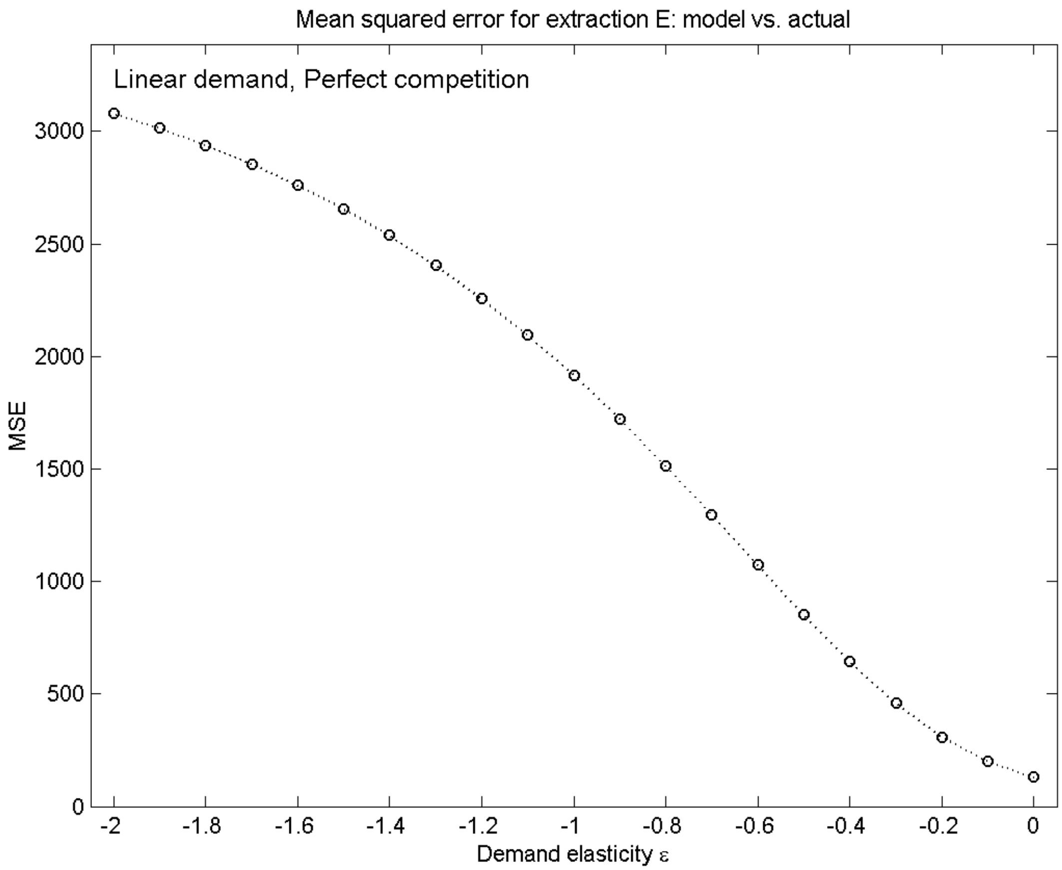
Mean squared error for extraction E: model vs. actual
Linear demand, Perfect competition
MSE
3000
2500
2000
1500
1000
500
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0
Demand elasticity ε

FIGURE 22.

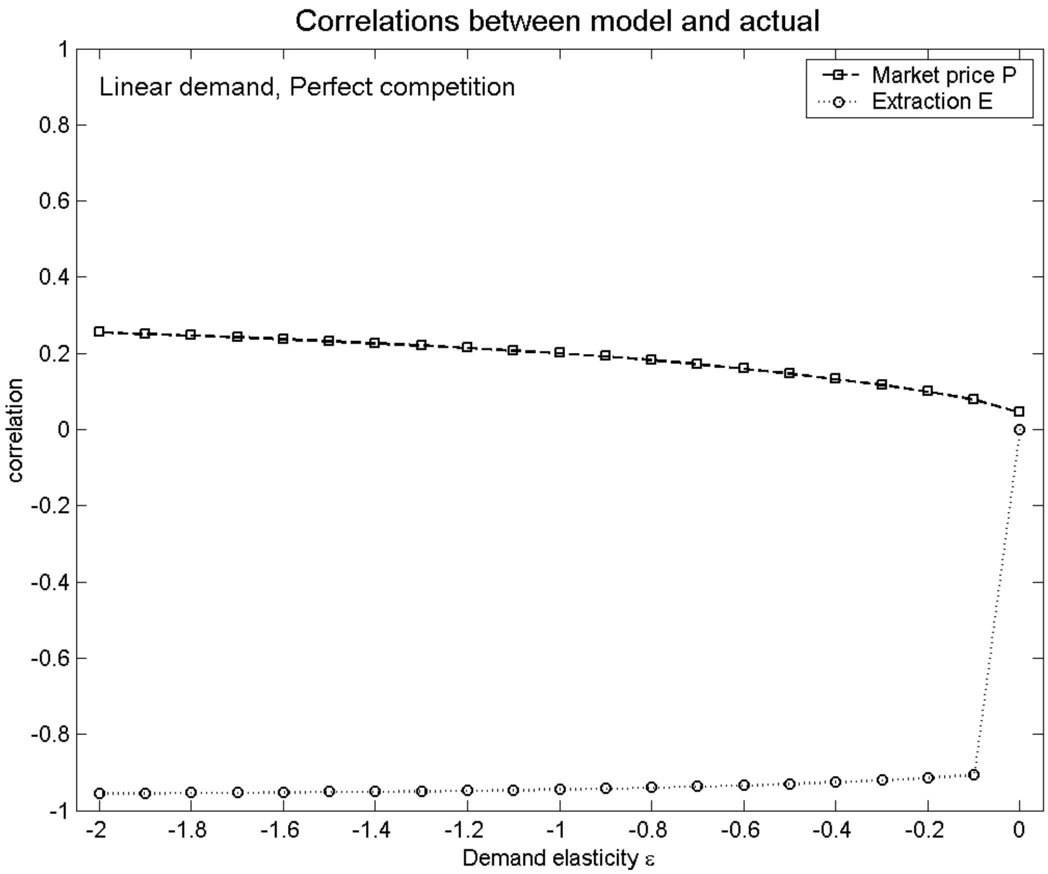
Correlations between model and actual
Linear demand, Perfect competition
Market price P
Extraction E
correlation
Demand elasticity ε
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 23.

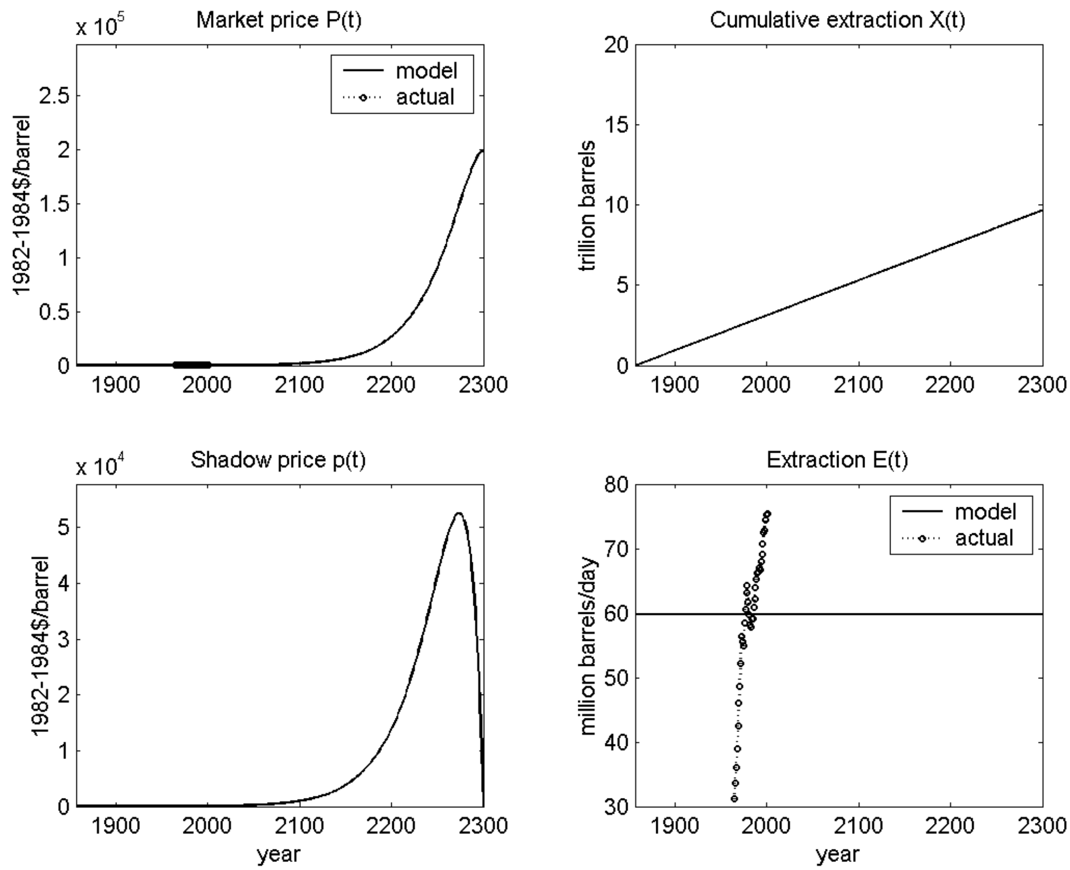

Perfect competition with linear demand when demand elasticity ε is 0.

FIGURE 24.

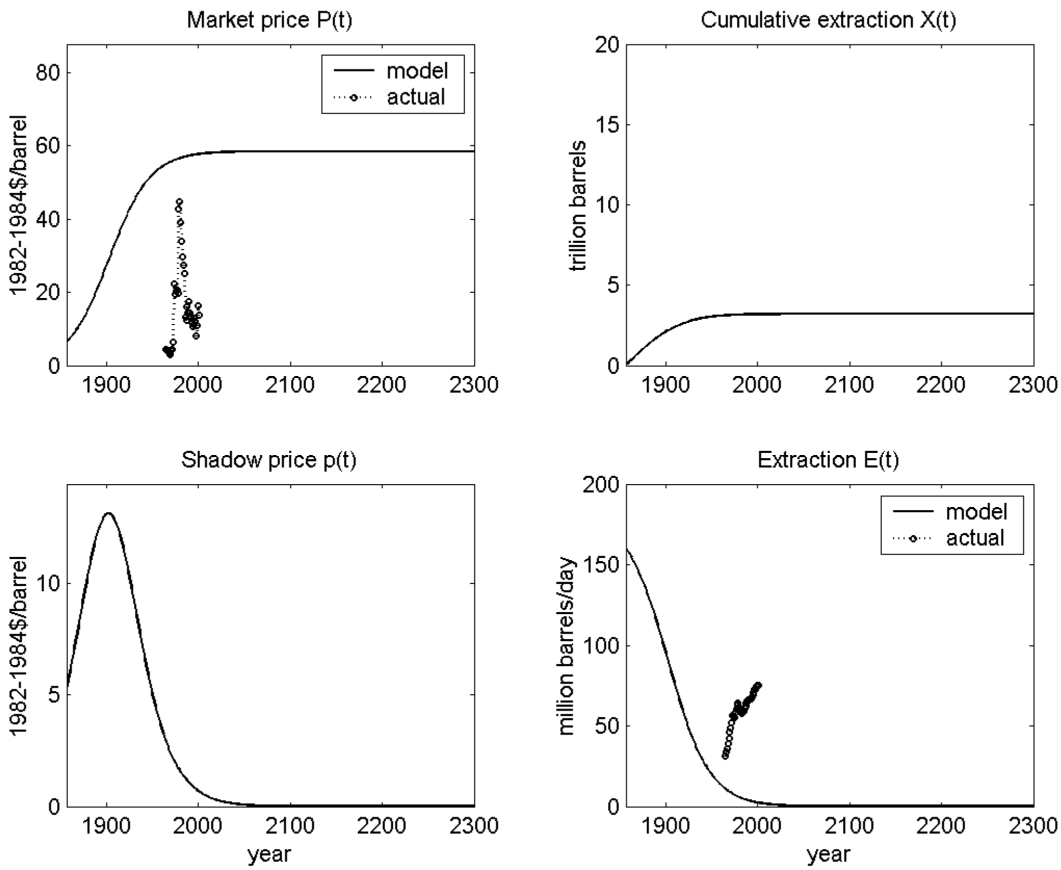

Perfect competition with linear demand when demand elasticity ε is –2.0.

FIGURE 25.

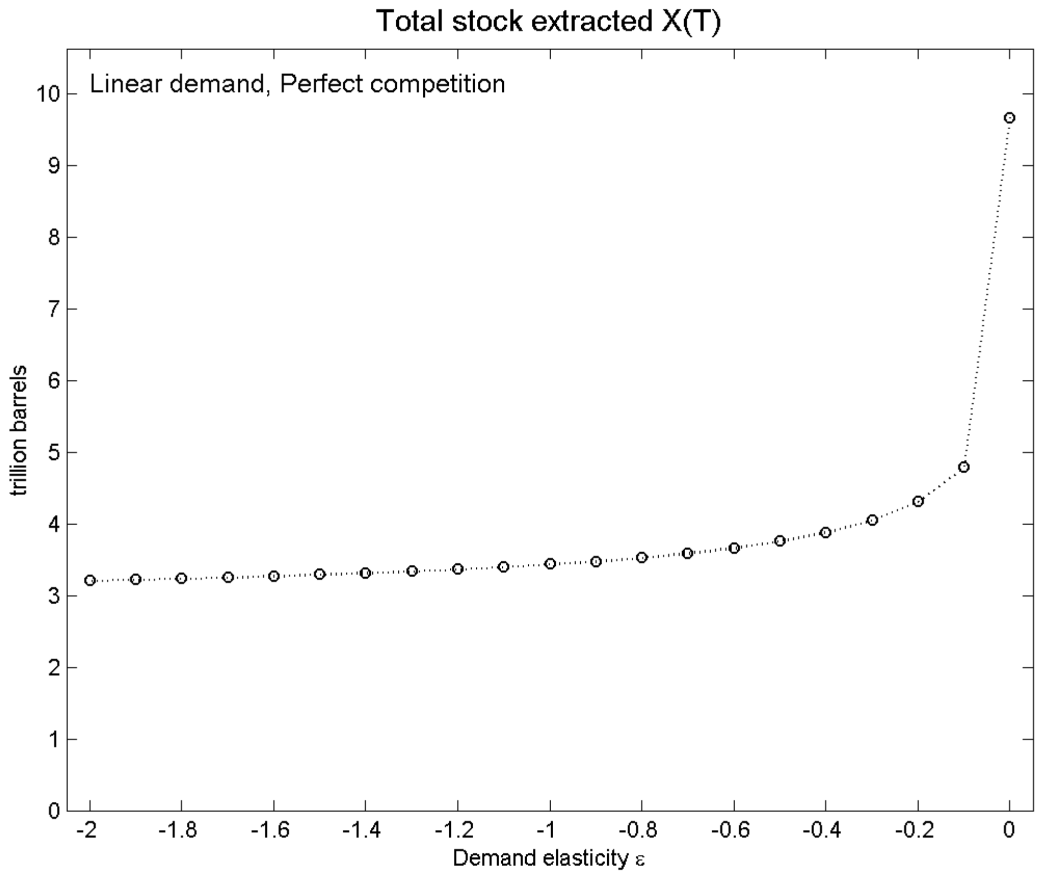
Total stock extracted X(T)
Linear demand, Perfect competition
trillion barrels
Demand elasticity ε
0
1
2
3
4
5
6
7
8
9
10
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
0

FIGURE 26.

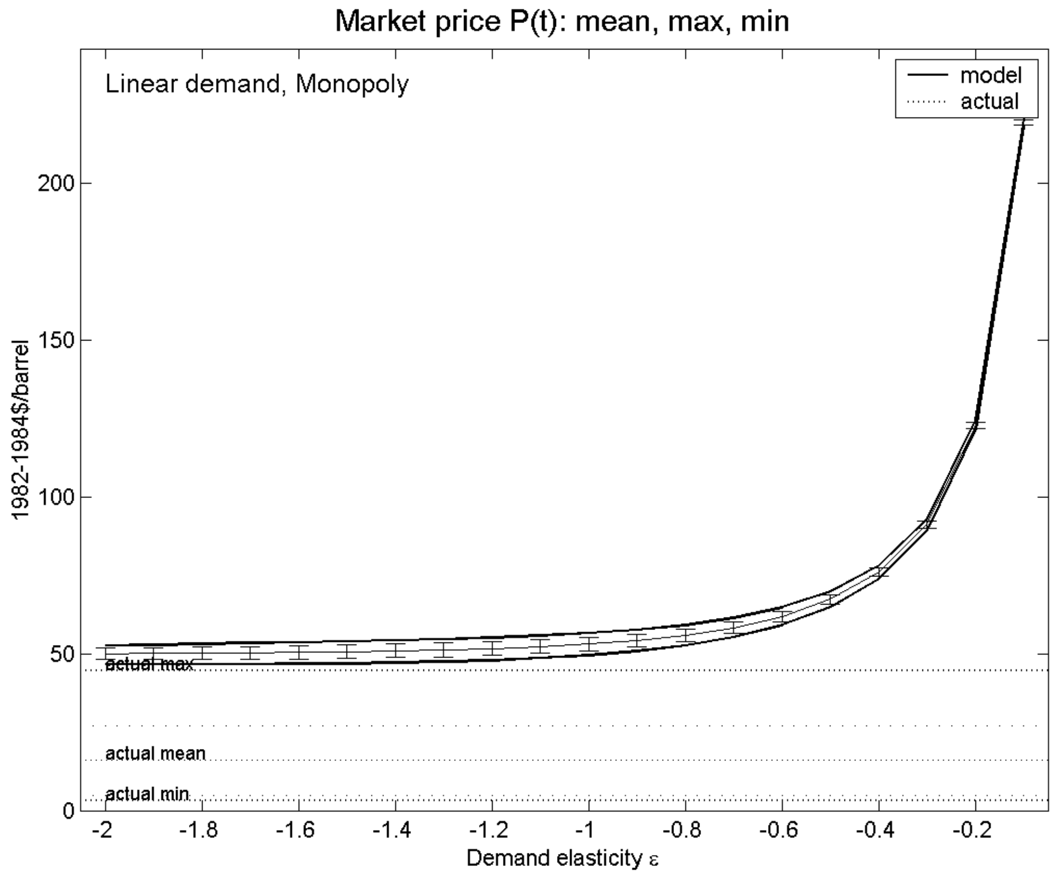

Market price P(t): mean, max, min
Linear demand, Monopoly
model
actual
1982-1984$/barrel
200
150
100
50
0
actual max
actual mean
actual min
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
Demand elasticity ε

FIGURE 27.

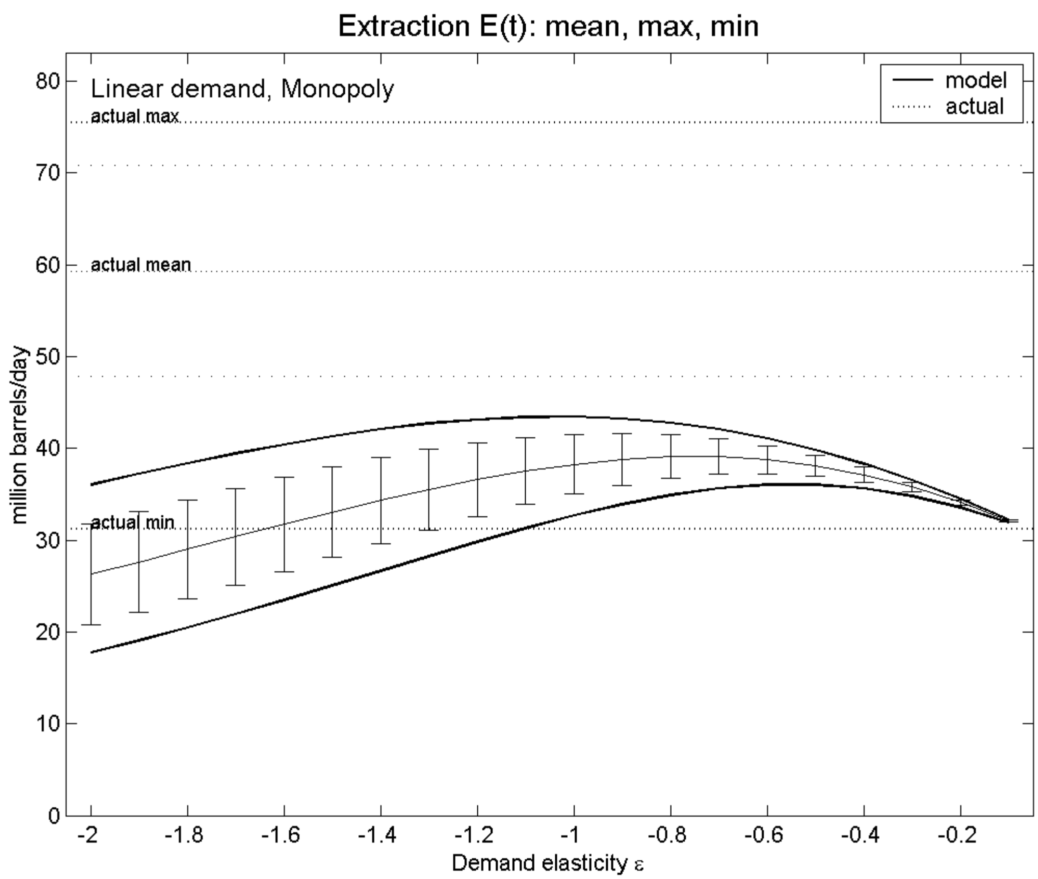

Extraction E(t): mean, max, min
Linear demand, Monopoly
model
actual
actual max
actual mean
actual min
million barrels/day
Demand elasticity ε
80
70
60
50
40
30
20
10
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2

FIGURE 28.

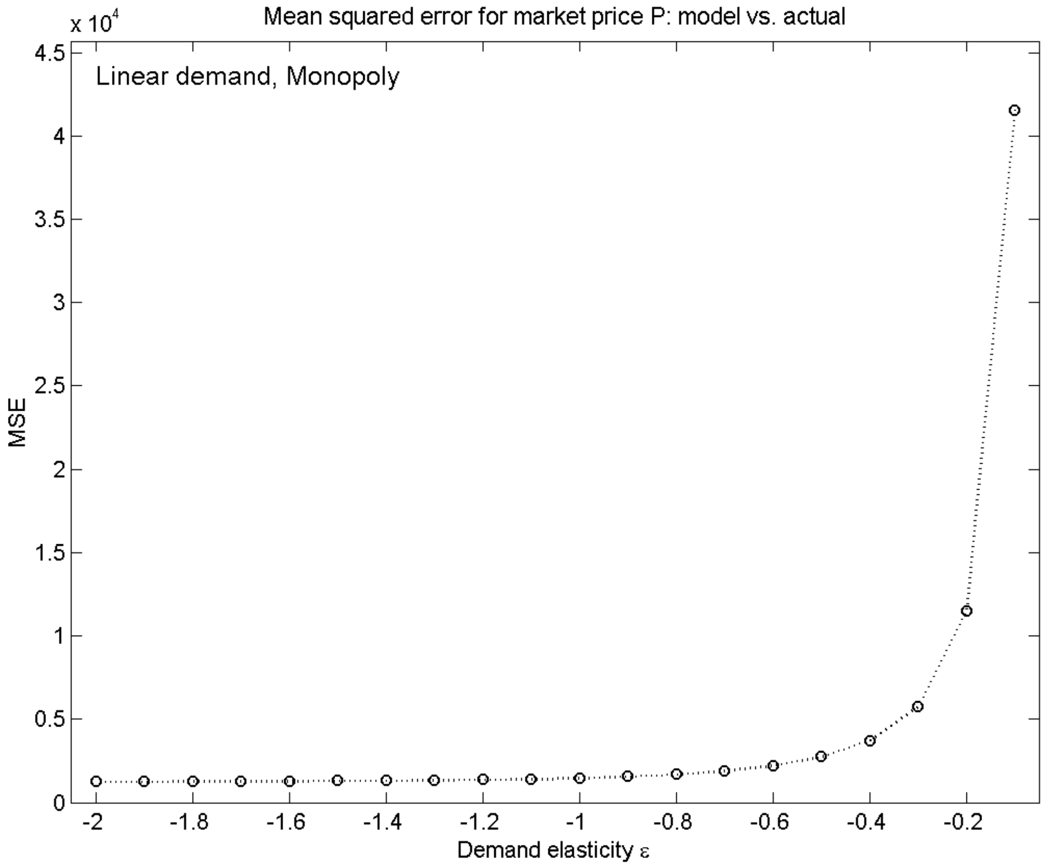
Mean squared error for market price P: model vs. actual
x 10^4
Linear demand, Monopoly
MSE
4.5
4
3.5
3
2.5
2
1.5
1
0.5
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
Demand elasticity ε

FIGURE 29.

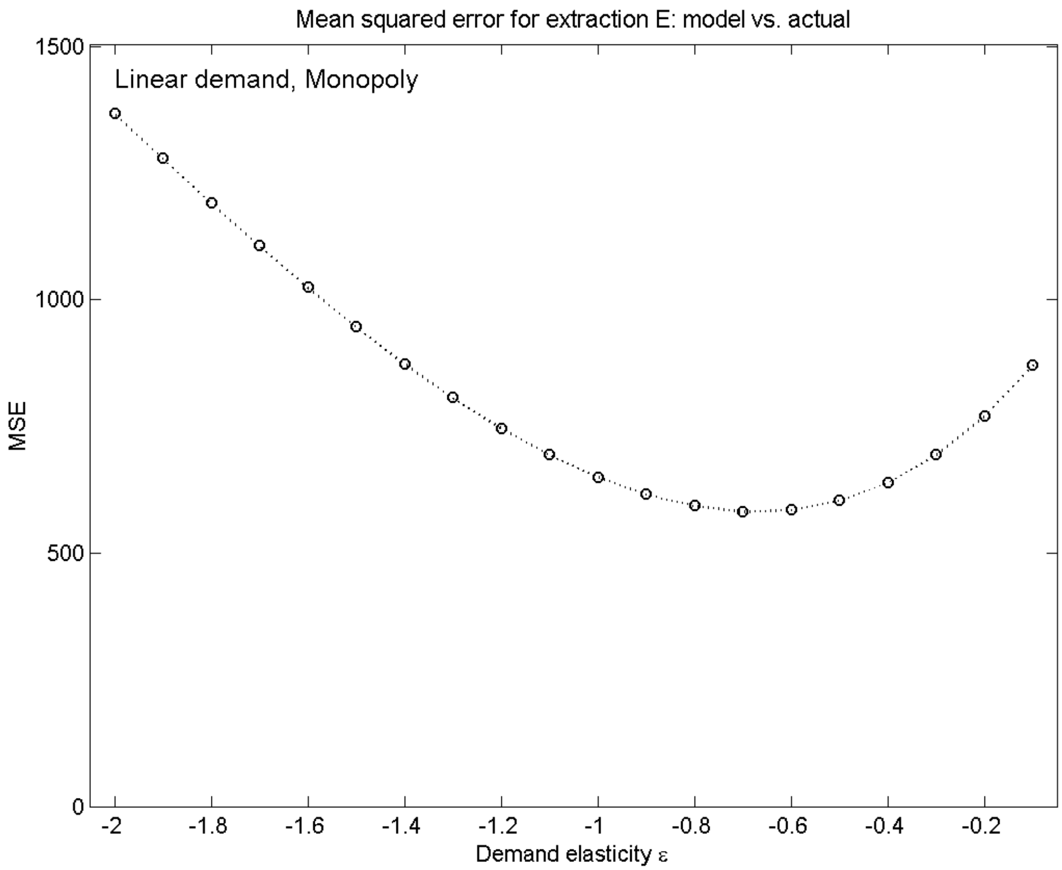

Mean squared error for extraction E: model vs. actual
Linear demand, Monopoly
MSE
Demand elasticity ε
1500
1000
500
0
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2

FIGURE 30.

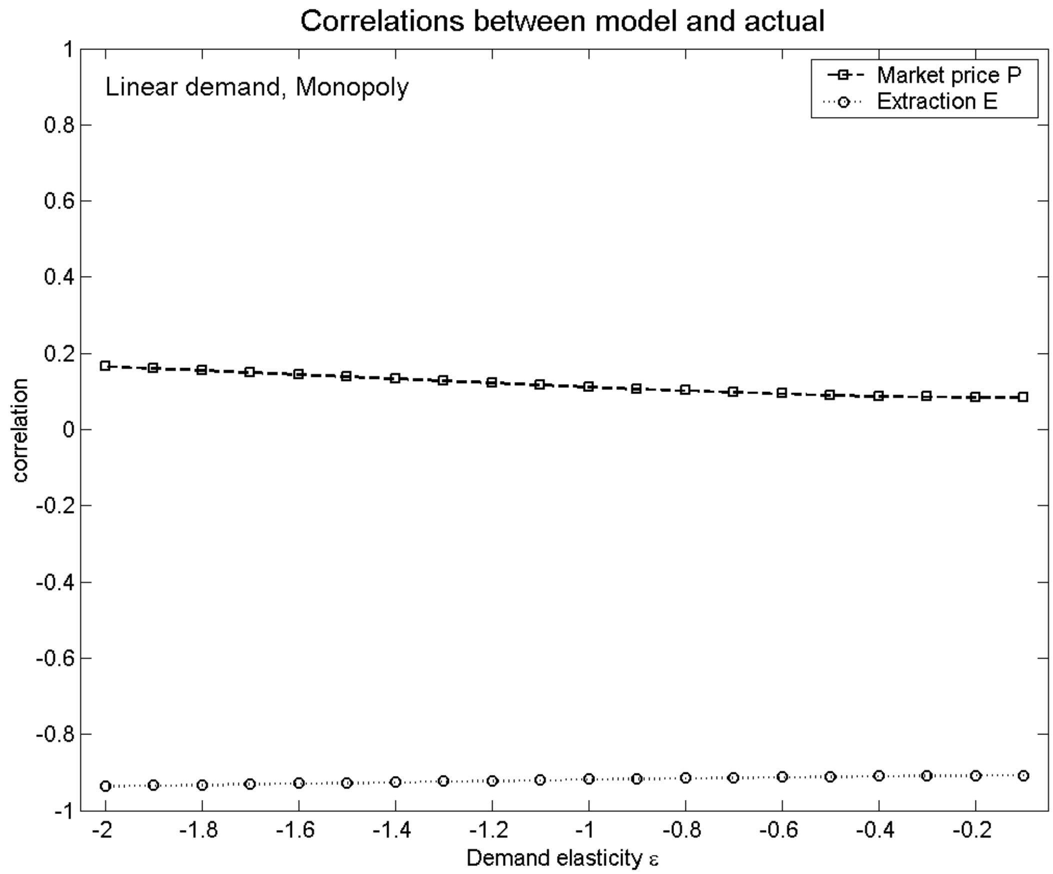

Correlations between model and actual
Linear demand, Monopoly
Market price P
Extraction E
correlation
1
0.8
0.6
0.4
0.2
0
-0.2
-0.4
-0.6
-0.8
-1
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2
Demand elasticity ε

FIGURE 31.

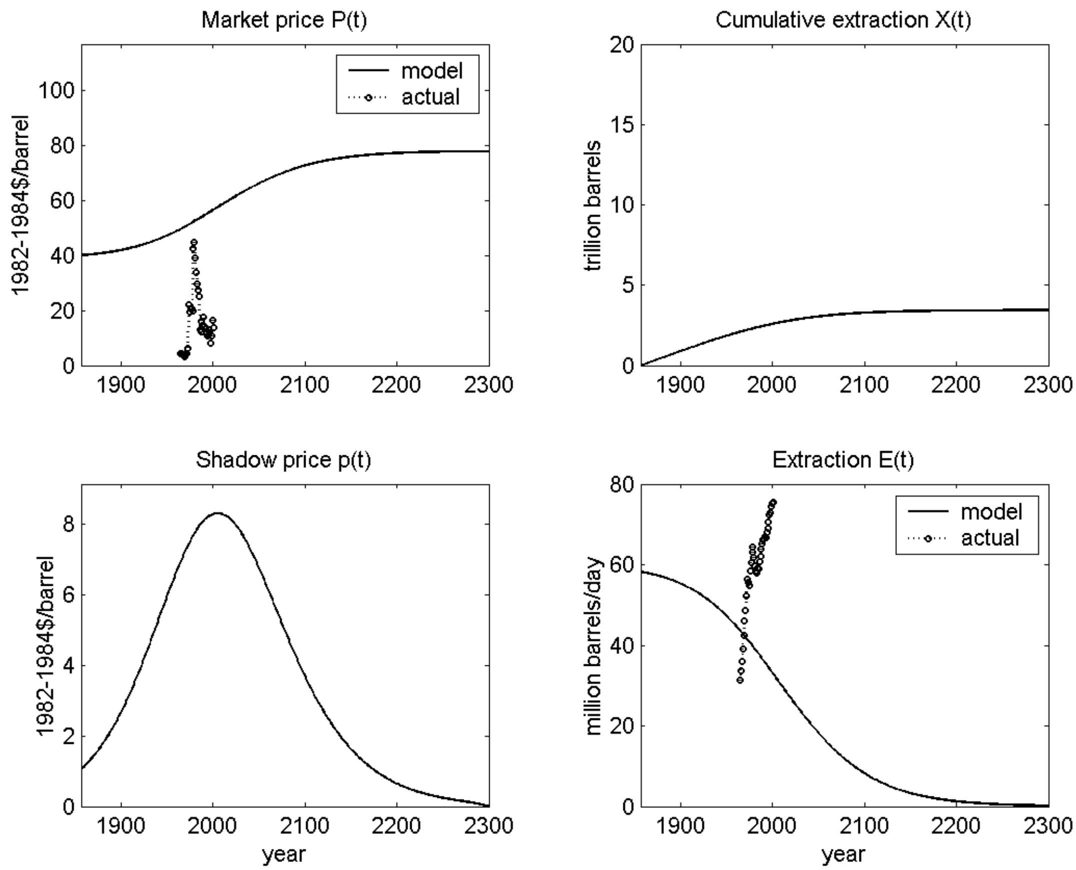

Monopoly with linear demand when demand elasticity ε is –1.0.

FIGURE 32.

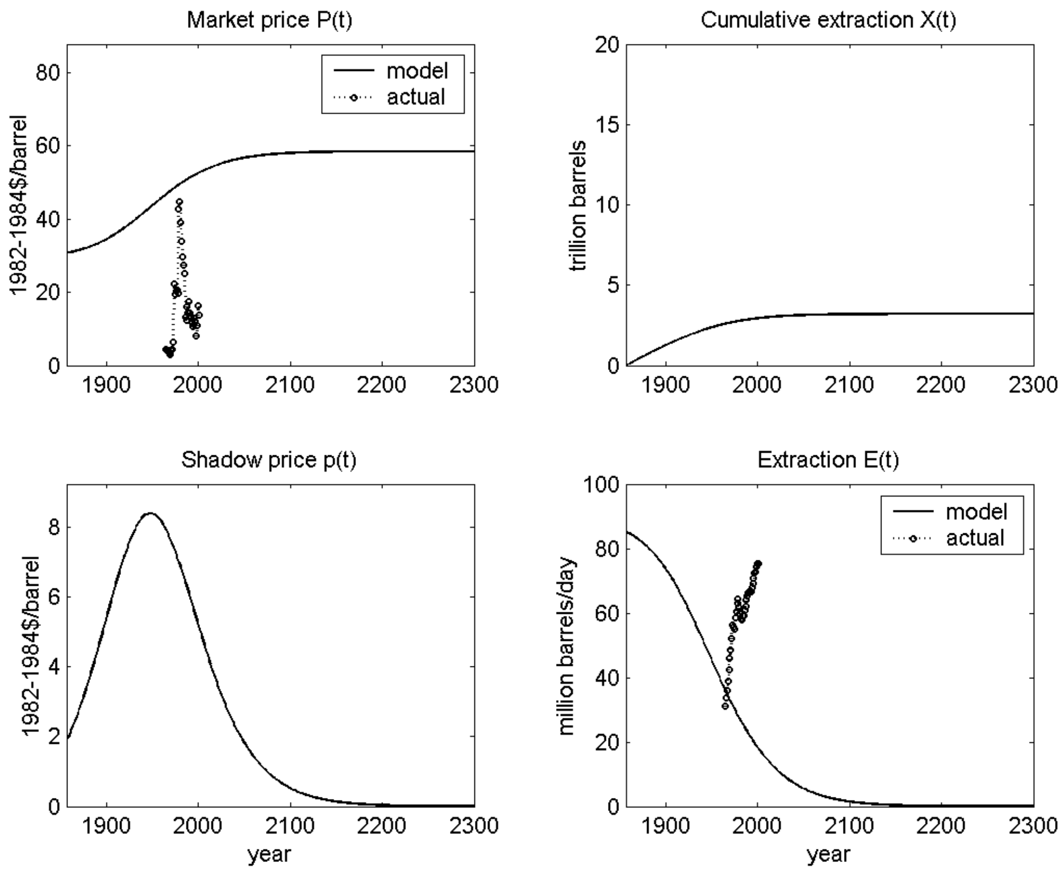

Monopoly with linear demand when demand elasticity ε is –2.0.

FIGURE 33.

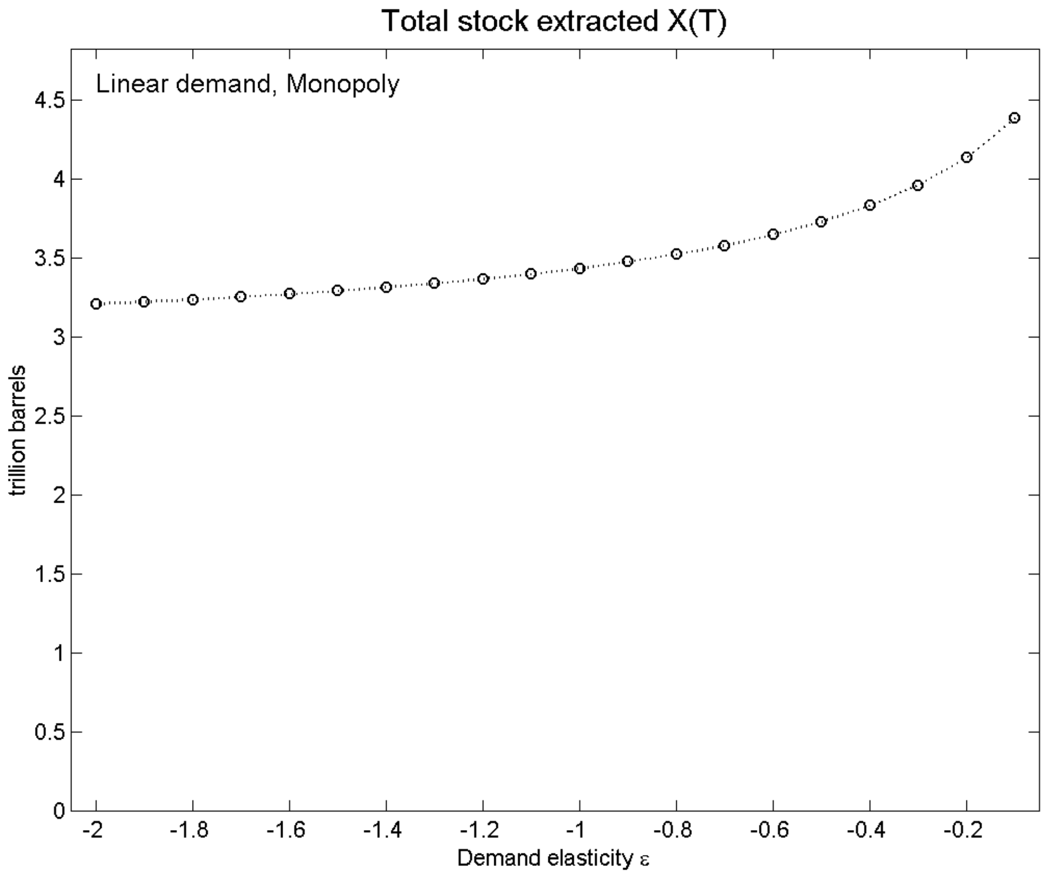
Total stock extracted X(T)
Linear demand, Monopoly
trillion barrels
Demand elasticity ε
0
0.5
1
1.5
2
2.5
3
3.5
4
4.5
-2
-1.8
-1.6
-1.4
-1.2
-1
-0.8
-0.6
-0.4
-0.2

FIGURE 34.

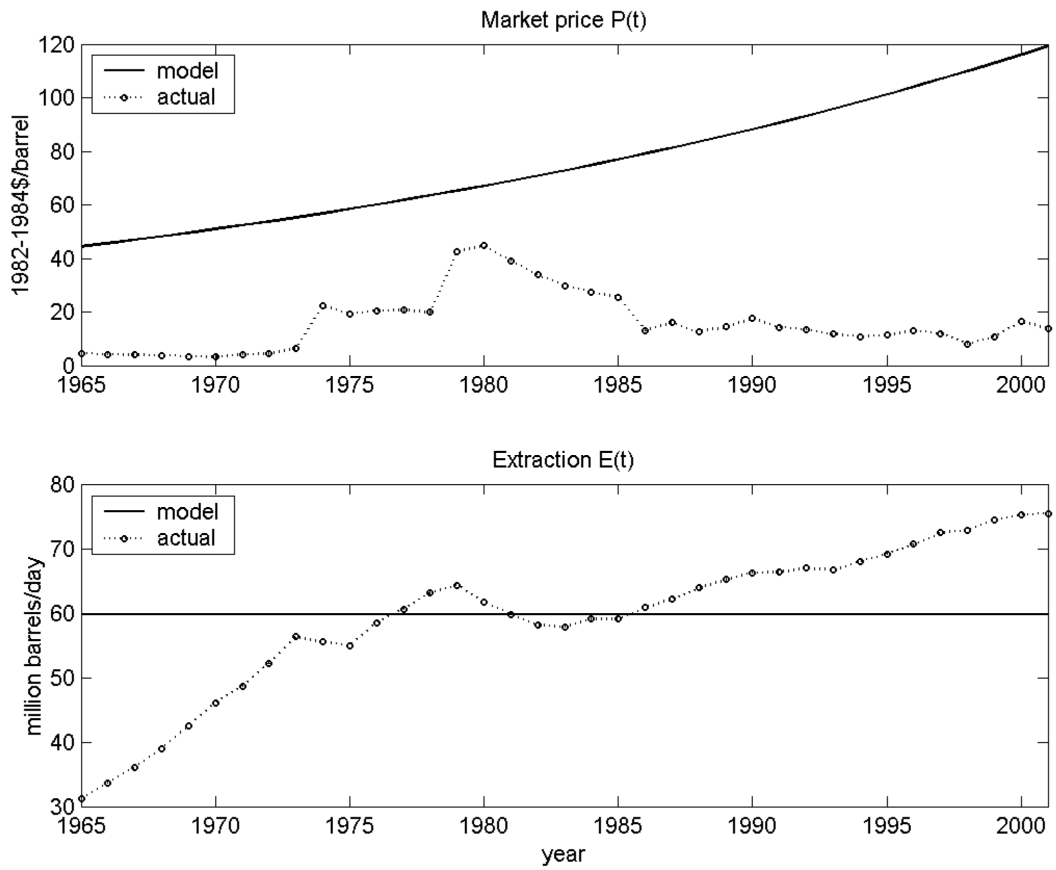

Perfect competition with isoelastic demand when demand elasticity ε is 0.

FIGURE 35.

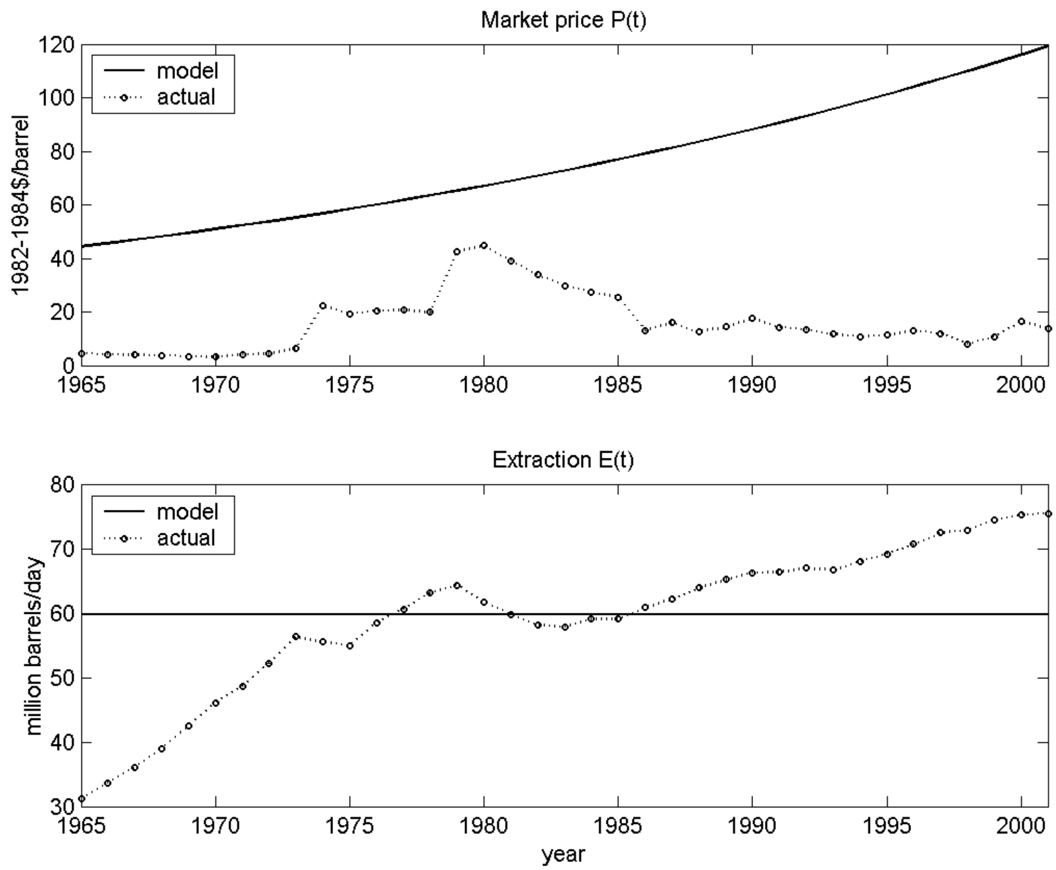

Perfect competition with linear demand when demand elasticity ε is 0.