Hotelling's Rule in the limit: an agent-based exploration of the model space

David S. Dixon*

April 23, 2012

Hotelling's Rule is the observation that the exploitation of a nonrenewable resource can only be economically efficient if the resource owner's marginal profit increases at the prevailing discount rate. This has been a perennial topic in the literature of resource economics since the 1970s, with some authors extending the theory and others analyzing empirical data. This paper reports on the results from using agent-based modeling to assess the consequences of relaxing the optimality constraint to explore the ways in which the outcome space converges on Hotelling's Rule in the limit. The agent-based model (ABM) in this paper has one choice variable: increase, decrease, or maintain the current production level - based on one rule: choose the change in production level that maximizes estimated discounted profit. The results, based on a costless technology and a stylized demand function from Hotelling, indicate that total discounted profit has low sensitivity to deviations from the optimum. In extending the basic Hotelling model to stylized production technologies with cost, the simple ABM falls short of the optimum by as much as ten percent, depending on the magnitude and whether the cost is fixed, marginal, or based on the resource stock level. The optimization errors of the ABM are similar to the errors of a human production planner with incomplete information. The ABM also exhibits emergent collusion-like and Cournot-like behaviors when extended to a small oligopoly market.(JEL Q32)

*Department of Economics, MSC 05 3060, 1 University of New Mexico, Albuquerque, NM 87131-0001, ddixon@unm.edu. The author acknowledges the invaluable guidance and assistance of Janie Chermak, the comments and suggestions of Robert Solow, Gérard Gaudet and Meghan Hutchins, and the technical assistance of Denise Stark of Aptech. The MASON and GAUSS code for these models is available from http://www.unm.edu/~ddixon.

Keywords: Hotelling's Rule; agent-based modeling; agent-based computation economics.

"They're more what you'd call guidelines than actual rules."

-Hector Barbosa in *Pirates of the Caribbean:*

The Curse of the Black Pearl

With regard to Hotelling's Rule, Captain Barbosa might well have said "it's more what you'd call an outcome than an actual rule." Hotelling (1931) observes that the owner of a finite natural resource is indifferent to either exploiting it or leaving it in situ when the marginal profit increases at the prevailing interest rate. The rationale is that, if the return is lower than this, the resource owner will shift assets to a better performing investment. If the return is greater, the owner will leave the resource where it is as it appreciates faster than other investments. In other words, the resource will not be produced at all unless it can be produced at a rate that returns the prevailing interest rate.

Hotelling noted that this sets the upper limit on a monopoly producer's profit from production. To explain this, it is necessary to appeal to the law of demand. For a downward-sloping demand curve, the monopoly producer can only increase the price by decreasing the production level. Hotelling's Rule says that if the producer decreases production at a rate that makes the marginal profit change by the interest rate, total profit from the resource will be maximized. If the production level is changed in any other way, total profit will be less than the maximum.

Hotelling's Rule is often called *the r-percent rule* and paraphrased as "the price must increase by r percent," r referring to the interest rate.[1] The producer profit, or scarcity rent, is paid by the consumer, and is also called user cost. If profit is increasing by r-percent, then user cost is also increasing by r-percent. Using dynamic optimization, Hotelling shows that the r-percent rule is the outcome of the producer maximizing profit, rather than a rule for the producer to follow.

The responsibility for naming it a rule may fall on Robert Solow, who first used the term "Hotelling's rule" in his Ely Lecture presented to the 1973 conference of the American Economic Association (Solow, 1974, p 12). In reference to this, Solow reflects that "Hotelling's concept is not a 'rule' at all in the appropriate sense. It doesn't enjoin anything. Phelps's Golden Rule is and does. Hotelling's principle is a description of what a foresighted competitive market would do, under simple conditions. Neither Phelps's nor Hartwick's rule has that property. They have to be imposed."[2] Nonetheless, the term "Hotelling's Rule" appears in countless texts and papers and is well-known - even beloved - to generations of natural resource economists.

[1]Strictly speaking, price is not identical to marginal profit, particularly in a monopoly market.
[2]From a personal communication dated 25 May 2010.

1 Background

Gaudet (2007) presents an excellent historical and contextual background for Hotelling's 1931 paper, from which the points in this paragraph are excerpted. The potential exhaustion of natural resources was a politically charged topic in the early twentieth century. Early economic models of exhaustible natural resources were based on static equilibrium, motivating Hotelling to develop a dynamical model. The result was a first-order optimization, with reference to specific cases requiring calculus of variations. Though straightforward by modern standards, the approach was mathematically sophisticated for the time, so the key finding languished until interest in exhaustible resources reemerged in the 1970s.

Devarajan and Fisher (1981) review the first fifty years of theoretical developments based on the Hotelling model. Much of that work is shown to pertain to issues that Hotelling raised but did not pursue, such as the effects of cumulative production and uncertainty in stock size. Krautkraemer (1998) adds another decade and a half and includes developments in the econometric search for evidence of Hotelling's Rule. Another ten years are added to the Hotelling time line by Livernois (2009). For the purposes of this paper, there are two main themes of interest in the Hotelling's Rule literature: theoretical efforts to broaden the scope of Hotelling's Rule, and econometric efforts to find evidence of the Hotelling's Rule outcome.

1.1 The evolving theory of nonrenewable resources

Hotelling's Rule says that scarcity rents must be increasing by r-percent, which implies, ceteras paribus, that price should be increasing by r-percent. There are, however, theoretical bases for decoupling the trend in market price from the trend in scarcity rents. Solow (1974, p. 3) suggests that if extraction costs fall by more than scarcity rents increase, the trend in market price may be downward. He notes, however, that eventually scarcity rent will dominate market price. Krautkraemer (1998) presents theoretical extensions to the Hotelling model to take into account variable stock levels due to exploration, cost of capital, capacity constraints, ore quality, and market imperfections.

Heal (1976) presents a model in which stock effects produce a declining resource value. Levhari and Liviatan (1977) explore the conditions under which a resource becomes economically nonviable but not physically depleted, a situation that arises with cumulative production (stock) costs. Livernois and Martin (2001) illustrate circumstances in which market prices rise while scarcity rents decline to zero because of resource degradation.

Pindyck (1978) and Livernois and Uhler (1987) suggest that bringing new deposits into production as the result of exploration produces a U-shaped price path. Slade (1982) proposes a U-shaped price trend due to technological advances early in production and scarcity rents late in production. Similarly, Dasgupta and Heal (1980) and Arrow and Chang (1982) propose models of exploration that produce a saw-tooth price curve. Cairns and Van Quyen (1998) propose a model that combines exploration and stock effects in which the price trend is downward for most of the stock lifetime, but rises to the choke price at the end.

Much of the preceding theoretical development is motivated by studies of various resource markets in which no long-term price increase is evidenced. Each represents a major undertaking to explore a simple extension of the Hotelling model, with the exception of Cairns and Van Quyen (1998), which incorporates two simple extensions.

Hotelling presents a perfectly competitive model and a monopoly model, then makes suggestions as to what happens when it is a duopoly market. He describes what, in game theory parlance, is called the cooperation-defection problem. Hotelling notes that with an exhaustible resource, the defector has an incentive to raise the price, unlike in static equilibrium in which the defector has cause to lower price (and increase sales). With an exhaustible resource, the defector knows that if the competitor doesn't match the higher price, the defector will be at an advantage later as the resource becomes more scarce.

Salant (1976) presents a highly notional model of prices in an oligopoly market for an exhaustible resource. An intriguing outcome of this model is that, for a market where some competitors form a cartel and the rest do not, a portion of the scarcity rents is transferred away from the cartel to be shared by the competitive fringe.

Stiglitz (1976) observes that a monopoly in an exhaustible resource market has less market power than in the market for a non-exhaustible resource. This point is supported by Lewis et al. (1979) and confirmed by Pindyck (1987). The result is challenged by Gaudet and Lasserre (1988), however, who point out the Stiglitz non-exhaustible model assumes infinite input capacity, while inputs in the exhaustible models are, by definition, limited. By treating the monopoly and competitive models as having the same input capacity constraint, Gaudet and Lasserre show that the monopolist's market power is the same whether the resource is exhaustible or not.

The principal barrier to theoretical expansion of the Hotelling model is that dynamic optimization is difficult to characterize in general terms. Hotelling found it necessary to resort to specific demand functions in order to explore implications of the base, costless model. (Livernois, 2009) notes that the Hotelling model becomes complex when extended to include factors like resource degradation. Krautkraemer (1998) is able to make general statements about the stock cost term by breaking the base case into a benefit term and a cost term, but does not carry the theoretical development into the extensions.

Agent-based modeling can be a laboratory for exploring and experimenting with proposed extensions to the Hotelling model. Because agent-based modeling is a causal framework, the researcher must express the behavior in terms of what an agent - a resource producer, for example - would do under conditions that arise in simulation. For an exploration model in which the resource deposits are discovered in decreasing quality (Pindyck, 1987, Livernois and Uhler, 1987), the new deposits come into existence sequentially: each deposit is exploited until its costs equal those of the next most costly (lesser quality) deposit, then production shifts to that one. In other models the quality of newly discovered deposits is random (Swierzbinski and Mendelsohn, 1989), which has the effect of flattening the price curve. Agent-based modeling is well suited to comparing the outcomes of these two models.

1.2 Studies of the markets for nonrenewable resources

A decade before the birth of modern environmental and natural resource economics in the 1970s, Barnett and Morse (1963) finds that, over the period 1870 to 1957, resource prices show no discernible trend, despite continuous and often rapid increases in their production. More then a decade and a half later, Smith (1979) confirms these findings using more sophisticated techniques on data from 1900 to 1973. Fisher (1981, p. 102-103) concurs that there are no discernible trends in overall resource prices in the results of Barnett and Morse (1963), but notes that factor costs fell more rapidly than resource prices during this period, so there is evidence for some increase in rents during that time. Alternatively, Brown and Field (1978) assert that Barnett and Morse neglected to include some factor costs, notably transportation. Brown and Field point out that the first three-quarters of the twentieth century - the period covered by the study - was a period of dramatic technical and social change.

Halvorsen and Smith (1991) note that "the principle obstacle to empirical tests of the theory of exhaustible resources has been data availability." Since actual marginal cost and, therefore, actual user cost, are not observable, many of the market studies are essentially assessments of the suitability of proxies for these.

Nordhaus (1973, p. 566) observes that, on the average, scarcity rents on energy resources were "quite modest," on the order of one dollar per barrel of petroleum.[3] He notes that the exception was petroleum itself, for which market prices were 2.4 times his calculated optimal price, which includes scarcity rent. A year later, Nordhaus (1974) found no trend in mineral prices between 1900 and 1970. Nordhaus is responsible for the term "backstop technology," referring to nuclear power as the technology that would, ultimately, limit the maximum price on petroleum and eliminate energy scarcity (Nordhaus, 1973, p. 532).

Some researchers have found evidence of U-shaped price curves, such as Slade (1982), who postulated that the price trend was due to falling input costs early in production and rising scarcity costs at the end of resource lifetime. Berck and Roberts (1996) examine a larger data set that includes most of the data used by Slade. They find that price predictions depend on whether prices are modeled as trend-stationary or as difference-stationary. They also find that trend-stationary models predict rising resource prices while difference-stationary models are ambiguous.

Heal and Barrow (1980) develop an arbitrage model which is an attempt to detect, directly, the Hotelling outcome. They assert that, in an efficient resource market, "there will be a strong association between the rates of change of resource prices and the rates of return on other assets." They found, however, that changes in interest rates were the relevant explanatory variables, not the interest rate levels. They conclude that simple equilibrium theory is inadequate to the complexity of the problem.

Stollery (1983) estimates the cost function for the price-leader in the nickel market and infers user cost by subtracting marginal cost from market price. User cost is found to be

[3] For the years 1970-1973, crude oil prices averaged less than three dollars per barrel in current dollars. http://www.eia.doe.gov/aer/txt/ptb1107.html (accessed 14 April 2011)

low initially, but to increase by 15 percent. Farrow (1985), using proprietary mine data, presents a model to estimate the in situ value of the stock and compares that with the price trend. He does not find evidence of Hotelling's Rule.

Halvorsen and Smith (1984) introduce duality to estimate in situ stock value for Canadian metal mines and find that user cost decreased considerably. Chermak and Patrick (2001) also appeal to duality to estimate in situ stock value and find that, for 29 natural gas wells, the trend is consistent with Hotelling's Rule.

(Agostini, 2006) found no evidence of U.S. copper companies exercising oligopoly market power before 1978, when U.S. copper went on the world market. He suggests a possible explanation is that the copper firms did exercise market power for brief periods, but limited prices during periods of high demand as a barrier to new entrants.

The preceding studies are a sample of the variety of markets and methodologies applied to this problem, and the results are often intriguing, but seldom conclusive. "However, the ability of the theory of exhaustible resources to describe and predict the actual behavior of resource markets remains an open question."(Halvorsen and Smith, 1991)

1.3 Agent-based computational economics

In the social sciences, one early adopter of agent-based modeling is Robert Axelrod (1997), whose models started out as an extension of earlier work on the Prisoner's Dilemma in game theory (1987). Agent-based modeling flourished as computation became faster, cheaper, and widely available. Many of the behavioral models developed theoretically in the preceding decades, such as Axelrod's and those of Thomas Schelling (1978) were well suited to agent-based modeling. Epstein and Axtell, both jointly and separately, develop agent-based models including an adaptation of a cultural transmission model by Axelrod (Axtell et al., 1996), the emergence of classes (Axtell et al., 1999) and civil violence (Epstein, 2002). McFadzean et al. (2001) introduce agent-based modeling as a computational laboratory for trade networks, and Tesfatsion (2001) applies the approach to an adaptive search model of the labor market. The dynamic and emergent behaviors of agents in combat are examined in Reynolds and Dixon (2001) and Dixon and Reynolds (2003), while the latter also models how a national bond market crisis spreads globally. Gilbert and Troitzsch (2005) provide an overview of various topics in modeling and simulation of social systems, including agent-based modeling. Tesfatsion (2006) provides examples of agent-based models that correspond to and often extend traditional economic models.

A note about the initials ABM. They are used to refer to the methodology of agent-based modeling, or to an agent-based model. That is "we will use ABM to explore Hotelling's Rule by constructing multiple ABMs, each representing a different cost structure." In this paper, agent-based modeling is referenced in its entirety, while the initials ABM are reserved for the models.

Although many ABMs are ad hoc computer programs, groups within the agent-based community have developed programs to automate the modeling and simulation process to some extent. MASON, a project at George Mason University, is an example, and is used for the modeling and simulation presented in the following chapters. For more details on MASON, see the website.[4] The MASON code for this paper is available for download.[5]

2 Theory

"Problems of exhaustible assets cannot avoid the calculus of variations" noted Hotelling (1931, p. 140), which may have been why his work languished until the 1970s (Gaudet, 2007, p. 1035). Since its advent, optimal control theory (Pontryagin, 1959, Pontryagin et al., 1962) has become a standard tool for dynamic optimization in economics. Chiang (1992) notes that, unlike calculus of variations, optimal control can be used with functions that are piecewise continuous or that have corner solutions. The introductory optimal control problem in Chiang (1992) is the Hotelling model. Caputo (2005) observes that optimal control theory is more conducive to economic theory and intuition than calculus of variations.

The following is an optimal control development of the Hotelling monopoly model. The terminology and some notation derive from Kamien and Schwartz (1981), in particular, the use of $m(t)$ as the current value multiplier. The notation for partial derivatives is borrowed from Caputo (2005), and the economic interpretations are influenced by Krautkraemer (1998).

The Hotelling monopoly model begins with a known fixed stock x_0 of a nonrenewable resource. The problem for the resource owner is to determine a production path $q(t)$ that maximizes present value total net profit over the productive lifetime, T, of the resource. In general, net profit $\pi\left(q\left(t\right),x\left(t\right),t\right)$ is a function of production level, remaining stock level $x(t)$ and time. Assuming a constant discount rate r, the optimal control problem is

$$\max_{q} J\left(q\left(t\right),x\left(t\right)\right) = \int_0^T e^{-rt}\pi\left(q\left(t\right),x\left(t\right),t\right)dt \tag{1}$$

subject to these constraints

$$\begin{aligned}
\dot{x}\left(t\right) &= -q\left(t\right) \\
x\left(0\right) &= x_0 \\
x\left(t\right) &\geq 0 \qquad (2) \\
q\left(t\right) &\geq 0 \qquad (3) \\
x_0 &\geq \int_0^T q\left(t\right)dt \qquad (4)
\end{aligned}$$

[4] http://www.cs.gmu.edu/~eclab/projects/mason/ (accessed 25 June 2010)
[5] http://www.unm.edu/~ddixon (accessed 28 June 2010)

The first constraint is also the state equation and will be discussed subsequently. The next two constraints are the initial and terminal boundary conditions on the stock variable. The fourth constraint ensures that production is never negative, and the fifth ensures that total production never exceeds total resource stock.

The current-value Hamiltonian to maximize (1) with resource stock costate variable $m(t)$ is defined as

$$\mathcal{H}\left(q\left(t\right),x\left(t\right),t,m\right) \equiv \pi\left(q\left(t\right),x\left(t\right),t\right)-m\left(t\right)q\left(t\right) \tag{5}$$

The current-value formulation has the advantage that, since the ABM will be making decisions based on current state values, a direct comparison can be made between the state of the ABM at time t and the optimal state of the Hamiltonian at time t. The costate variable $m(t)$ is interpreted as the current-value shadow price of the resource stock at time t. This is the user cost of the next unit of remaining stock to be extracted.

The first order necessary conditions include the state equation

$$\dot{x}\left(t\right) = -q\left(t\right) \tag{6}$$

which imposes the dynamical constraint that the remaining stock be reduced at the rate of production, where $\dot{x}\left(t\right)$ is the time derivative of $x\left(t\right)$. The first order necessary costate equation is

$$\dot{m}\left(t\right) = rm\left(t\right)-\frac{\partial\mathcal{H}\left(q\left(t\right),x\left(t\right),t,m\right)}{\partial x\left(t\right)} \tag{7}$$

where $\dot{m}\left(t\right)$ is the time derivative of $m\left(t\right)$. If the Hamiltonian has no stock effect (no $x\left(t\right)$ dependency), this equation requires that the shadow price increase at the rate of the discounted shadow price, thus ending at some maximum value. Depending on its sign, the stock effect may accelerate or decelerate the increase in the shadow price, or, for a sufficiently positive stock effect, cause the shadow price to decrease over time.

The first order necessary optimality condition is

$$\frac{\partial\mathcal{H}\left(q\left(t\right),x\left(t\right),t,m\right)}{\partial q\left(t\right)} = 0 \tag{8}$$

which is the condition for static optimum, requiring that the Hamiltonian be maximized at all times. The transversality condition on the state variable is that

$$e^{-rT}m\left(T\right)\geqq 0,\quad x\left(T\right)\geqq 0,\quad e^{-rT}m\left(T\right)x\left(T\right)=0 \tag{9}$$

which constrains the ending shadow price to be non-negative and the ending stock level to be non-negative, but requires that the shadow value $(m \times x)$ of the ending stock must be zero. That is, either the stock is physically depleted and $x(T) = 0$, or the present value ending shadow price $e^{-rT}m(T)$ is zero. The latter condition can arise irrespective of the ending shadow price if the terminal time T can be infinite. For nonzero ending stock and finite T, that the shadow price goes to zero is intuitive, since terminating while there is remaining stock implies that it is not economic to extract the next unit of stock. If the stock variable is constrained to end at some value, the transversality condition on the Hamiltonian is that

$$e^{-rT}\mathcal{H}(T) = 0 \tag{10}$$

This ensures that the stock variable is stationary at the terminal time T (Chiang, 1992, p. 182).

The following relations are introduced for notational simplicity

$$\begin{aligned}
\pi_q(q(t), x(t), t) &= \frac{\partial}{\partial q}\pi(q(t), x(t), t) \\
\pi_x(q(t), x(t), t) &= \frac{\partial}{\partial x}\pi(q(t), x(t), t) \\
\mathcal{H}_q(q(T), x(T), T, m) &= \left.\frac{\partial \mathcal{H}(q(t), x(t), t, m)}{\partial q}\right|_{t=T}
\end{aligned}$$

Substituting for the Hamiltonian in (7)

$$\dot{m}(t) = rm(t) - \pi_x(q(t), x(t), t) \tag{11}$$

so that $\pi_x(q(t), x(t), t)$ is the Hamiltonian stock affect to which the previous remarks apply. That is, the rate at which the shadow price changes is either accelerated or decelerated by the $\pi_x(q(t), x(t), t)$ depending on its sign and, if $\pi_x(q(t), x(t), t) = 0$, shadow price increases at the discount rate.

Substituting for the Hamiltonian in (8)

$$\begin{aligned}
\pi_q(q(t), x(t), t) - m(t) &= 0 \\
\pi_q(q(t), x(t), t) &= m(t) \qquad (12)
\end{aligned}$$

which establishes the link between marginal profit and shadow price. Substituting (12) into (11) to eliminate $\dot{m}(t)$, then dividing by $\pi_q(q(t), x(t), t)$

$$\frac{\dot{\pi}_q(q(t), x(t), t)}{\pi_q(q(t), x(t), t)} = r - \frac{\pi_x(q(t), x(t), t)}{\pi_q(q(t), x(t), t)} \tag{13}$$

which is the general expression of Hotelling's Rule. For a production technology that has no dependence on the stock level, this simplifies to

$$\frac{\dot{\pi}_q\left(q\left(t\right),x\left(t\right),t\right)}{\pi_q\left(q\left(t\right),x\left(t\right),t\right)} = r \tag{14}$$

which is the relation first articulated by Hotelling. It states that, on the optimal production path, the percent change in marginal net profit is equal to the discount rate.

2.1 The Hotelling monopoly demand function

Hotelling's inverse demand function for the monopoly market Hotelling (1931, sec. 4) is

$$p = \left(1-e^{-Kq}\right)/q \tag{15}$$

This is a stylized demand not linked to any real market. It is not known why Hotelling chose it, but it has two pedagogical strengths. First, it yields tractable expressions for revenue and marginal profit and secondly, it has no finite static maximum with respect to q. That is, it can only be maximized in the dynamic context. Additionally, profit increases monotonically with q, a characteristic that is exploited in the models for which profit must be held above cost.

Equation (15) is the inverse demand function used throughout this paper, in the theoretical development and in the agent-based models. For Hotelling's costless production technology, the optimal production path is easily solved in closed form. However, for the nonzero-cost technologies, it is necessary to appeal to numerical solutions. In those cases, and in the agent-based models, the parameter values used are

$$\begin{aligned} K &= 5 \\ r &= (1+0.1)^{1/365.25}-1 \approx 2.16\times 10^{-4} \\ x_0 &= 100 \end{aligned}$$

These are all stylized and unitless values. K is the choke price and this value is chosen to be consistent with the other models used by Hotelling. The discount rate is ten percent per annum and is expressed in daily terms for use in and comparison with the agent-based models. The initial stock x_0 is chosen so that simulations of the agent-based models are long enough to be instructive yet short enough to be repeated many times.

2.2 A costless monopoly model

Costless production technology means that there is no cost term, so that profit is equal to revenue

$$\pi(q(t)) = p(q(t))\,q(t) \tag{16}$$

An optimizing costless monopoly producer, facing the full demand function, determines the optimal production path $q(t)$ based on the discount rate r, the extent of stock x_0, and any other boundary conditions. For the costless model, the stock is physically depleted at time T. Given a specific inverse demand function, the procedure is:

1. Solve (12) for the production path $q(t)$ in terms of $m(t)$.
2. Solve (11) to get $m(t)$.
3. Replace $m(t)$ in $q(t)$ and solve (10) to get $q(T)$ in terms of T.
4. Integrate (4) using the equality condition (because the stock is physically depleted) to get T in terms of initial stock x_0.
5. Solve $q(0)$ to get the initial production level.

For any reasonable inverse demand function (or approximation thereof) the terminal time T will be finite.

With the inverse demand function in Section 2.1, from step 1

$$\pi(q(t)) = 1 - e^{-Kq(t)} \tag{17}$$
$$\pi_q(q(t)) = Ke^{-Kq(t)} \tag{18}$$
$$m(t) = Ke^{-Kq(t)} \tag{19}$$
$$q(t) = \frac{\ln(K/m(t))}{K} \tag{20}$$

Equation (19) indicates that shadow price varies inversely with production level and has a maximum of K, the choke price.

Note that since there is no $x(t)$ term in (17), $\pi_x = 0$, so that from step 2

$$\begin{aligned} \dot{m}(t) &= rm(t) \\ m(t) &= m_0 e^{rt} \end{aligned} \tag{21}$$

where m_0 is the initial shadow price. The shadow price increases over time, which is intuitive, given that, under production, the resource becomes more scarce over time.

From step 3

$$q(t) = \frac{1}{K}\left(\ln\frac{K}{m_0} - rt\right) \tag{22}$$

If the stock is to be physically depleted, the transversality condition (10)

$$\mathcal{H}(T)e^{-rT} = 0$$

applies. For a finite lifetime T, this means that

$$\mathcal{H}(T) = 0 = 1 - e^{-Kq(T)} - m_T q_T$$

for which $q(T) = 0$ is the solution. Now (22) can be solved for T

$$\begin{aligned} 0 &= \frac{1}{K}\left(\ln\frac{K}{m_0} - rT\right) \\ T &= \frac{\ln\left(K/m_0\right)}{r} \end{aligned} \tag{23}$$

Substituting this back into (22) gives the production path

$$q(t) = \frac{r}{K}(T-t) \tag{24}$$

Equation (23) shows that higher discount rates promote more rapid depletion, which is expected, as a higher discount rate reduces the value of the resource in the future. Equation (24) shows that production follows a straight-line descending path that is increasingly steep as the discount rate increases, as expected. The production level is also inversely proportional to the choke price K. This is related to the inverse relationship of production with shadow price: since shadow price is increasing toward K, production is decreasing proportional to its inverse.

From step 4

$$x_0 = \int_0^T q(t)\,dt = \int_0^T \frac{r}{K}(T-t)\,dt = \frac{r}{2K}T^2 \tag{25}$$

$$T = \sqrt{\frac{2Kx_0}{r}} \tag{26}$$

Equation (26) makes explicit what was implied before: a higher discount results in a shorter resource lifetime. Also, resource lifetime is proportional to the square root of its extent.

Finally, from step 5

$$q(0) = \tfrac{r}{K}(T-0) = \sqrt{\frac{2rx_0}{K}} \tag{27}$$

The initial production level is proportional to the lifetime of the resource, which is proportional to the square root of its extent. The initial production level is proportional to the square root of the discount rate, moving production earlier as the discount rate increases, as expected. The inverse square root relation to choke price K is related to the inverse relationship between production and shadow price, as discussed previously.

Since $x(T) = 0$, the transversality condition (9) imposes no constraint on $m(t)$, which is

$$m(t) = Ke^{-r(T-t)} \tag{28}$$

That is, the shadow price of the remaining stock increases to the choke price K as the stock is physically depleted. Note also from (24) that

$$\dot{q} = -\frac{r}{K} \tag{29}$$

which is constant and negative. As mentioned previously, production follows a straight-line descending path. Finally, note that

$$\frac{\dot{\pi}_q(q(t),t)}{\pi_q(q(t),t)} = \frac{-K^2\dot{q}e^{-Kq(t)}}{Ke^{-Kq(t)}} = r$$

which is Hotelling's Rule.

The production path given by (24) maximizes present-value net profit over the lifetime of the resource. This is, by definition, the most profit a producer can ever get with this production technology and this demand function. Integrating (17) over the stock lifetime T, the theoretical maximum profit, therefore, is

$$\Pi_{max} = \int_0^T \left(1 - e^{-Kq(t)}\right) e^{-rt} dt = \frac{1}{r}\left[1 - e^{-rT}(1 + rT)\right] \tag{30}$$

For purposes of comparison with other models that cannot be solved in closed form, using the values from Section 2.1, the stock lifetime is 1958 days, and $\Pi_{max} = 358.33$.

2.3 A fixed cost model

Natural resource production often incurs fixed cost, including capital costs, leases or other per-period fees or taxes. Extractive industries tend to require large capital investments, and capital can be regarded as a quasi-fixed cost (Young, 1992). Hsiao and Chang (2002) have a groundwater optimization model of in which well-drilling is a fixed cost.

Consider a fixed, per-period cost c_0, so that net profit is

$$\pi\left(p\left(q\left(t\right)\right),q\left(t\right)\right) = p\left(q\left(t\right)\right)\cdot q\left(t\right)-c_0$$

Because the cost is not dependent on $q\left(t\right)$ or $x\left(t\right)$, this cost does not affect the dynamical constraints, appearing only in the solutions to the boundary conditions.

In terms of the demand function in Section 2.1 this is

$$\pi\left(q\left(t\right)\right) = 1-e^{-Kq(t)}-c_0$$

One characteristic of the revenue part of this is that it is monotonically increasing with $q\left(t\right)$. It can be anticipated, therefore, that there is some minimum production level, q_{min}, below which net profit is negative. Net profit is non-negative as long as

$$c_0 \leq 1-e^{-Kq(t)}$$

so that

$$q_{min} = \frac{1}{K}\ln\left(\frac{1}{1-c_0}\right) \tag{31}$$

Marginal profit is positive as long as production remains above this level. For nonzero c_0, terminal production $q(T)$ cannot be zero. Clearly, q_{min} is zero for $c_0 = 0$.

If the stock is to be physically depleted, the transversality condition (10)

$$\mathcal{H}(T)e^{-rT} = 0$$

applies. For a finite lifetime T, this means that

$$1-e^{-Kq_T}-c_0-m_Tq_T = 0$$

Using (19) to substitute m_T

$$\begin{aligned} e^{-Kq_T}\left(1-Kq_T\right) &= 1-c_0 \\ e^{Kq_T} &= \frac{1+Kq_T}{1-c_0} \end{aligned} \tag{32}$$

Equation (32) must be solved numerically.[6] The solution is shown in Figure 1. The figure shows that $q_T > q_{min}$ for all costs, so that the $q_T > q_{min}$ constraint is non-binding. Figure

Figure 1: Fixed cost model - numerical solutions for terminal production level q_T. The optimal production path starts at production level $q(0)$ and decreases continuously to the terminal production level, q_T. The production level is, at all times, well above the zero profit production level, q_{min}. Terminal time T, which decreases with increasing cost, is also shown. The bottom graph shows total net profit, which is also user cost, as a function of the fixed cost rate.

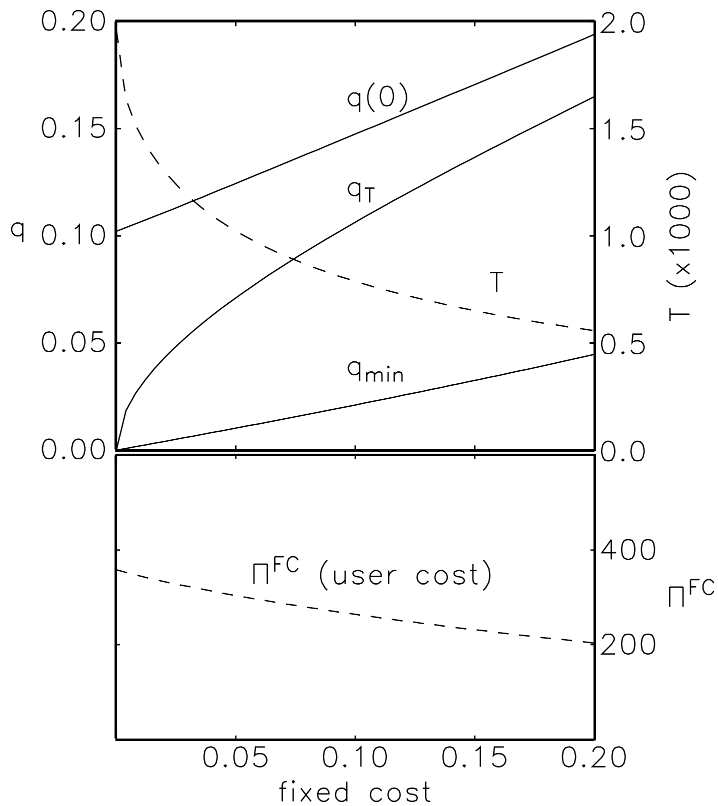

1 also shows that terminal time T decreases sharply as cost increases, and that total net profit (user cost) decreases steadily over the cost range.

Note, in Figure 1, that $q(0)$ is slightly steeper than q_{min}, while q_T is asymptotically parallel to $q(0)$. Thus, as the fixed cost increases, more of total production is pushed toward the present.

Replacing m_0 in (21) with $m_T e^{-rT}$

$$m(t) = m_0 e^{rt} \quad = \quad m_T e^{-rT} e^{rt}$$

then using (19)

$$m(t) \quad = \quad K e^{-Kq_T} e^{-r(T-t)}$$

so that

$$q(t) = \frac{1}{K} \ln \left[\frac{K}{K e^{-Kq_T} e^{-r(T-t)}} \right] \quad = \quad q_T + \frac{r}{K}(T-t) \tag{33}$$

Thus, the initial production level is

$$q(0) \quad = \quad q_T + \frac{rT}{K} \tag{34}$$

which is also shown in Figure 1.

Stock lifetime T is calculated from

$$\begin{aligned} x_0 &= \int_0^T q(t)\, dt \\ &= q_T T + \frac{rT^2}{2K} \end{aligned}$$

so that

$$T \quad = \quad \frac{K}{r} \left(\sqrt{q_T^2 + \frac{2rx_0}{K}} - q_T \right) \tag{35}$$

[6] The GAUSS code for numerical solutions is available from http://www.unm.edu/~ddixon (last accessed 28 February 2010)

which is shown in Figure 1. Substituting (35) back into (34)

$$q(0) = \sqrt{q_T^2 + \frac{2rx_0}{K}} \tag{36}$$

Total profit is

$$\begin{aligned}\Pi^{FC} &= \int_0^T \left(1 - e^{-Kq(t)} - c_0\right) e^{-rt} dt \\ &= \frac{1}{r}\left[1 - c_0 - e^{-rT}\left(1 - c_0 + rTe^{-Kq_T}\right)\right] \end{aligned}\tag{37}$$

which reduces to (30) for $c_0 = 0$ (for which $q_T = 0$). This is also shown in Figure 1.

2.4 A marginal cost model

Extractive technologies, like most production technologies, incur costs that are proportional to the level of production. Scott (1967) uses a quarrying example to illustrate that economy of scale considerations at low levels of production, and problems of marketing, delivery and storage at high levels of production, lead to a U-shaped marginal cost curve. Cobb-Douglas models in which production level appears are found in econometric models of nickel (Stollery, 1983) and copper (Young, 1992), for example. Conrad and Clark (1987, p. 165) give an example of a linear marginal cost associated with disposal of pollutants. For simplicity, this model considers a stylized linear marginal cost with marginal cost c_1, so that the net profit function is

$$\pi\left(p\left(q\left(t\right)\right), q\left(t\right), t\right) = p\left(q\left(t\right)\right) \cdot q\left(t\right) - c_1 \cdot q\left(t\right)$$

In terms of the demand function in Section 2.1 this is

$$\pi\left(q\left(t\right)\right) = 1 - e^{-Kq(t)} - c_1 q\left(t\right) \tag{38}$$

The transversality condition depends on whether or not $q(t)$ can go to zero when $t = T$. There is no minimum production level q_{min} as long as the cost goes to zero faster than the revenue. This is the case as long as

$$e^{-Kq(t)} \leq 1 - c_1 q\left(t\right) \tag{39}$$

Figure 2 shows graphs of the left-hand and right-hand sides of (39) for the parameter values presented in Section 2.1. The graphs show that, for all marginal costs lower than the choke

Figure 2: The marginal cost model - theoretical values for minimum and maximum production for selected marginal costs.
Profit is positive whenever the dashed line is below the solid line for a given marginal cost. For all costs less than 5 (the choke price), profit is non-negative in the vicinity of $q = 0$. The point where the dashed line crosses the solid line is q_{max} for that marginal cost. Thus, profit is non-negative over $0 \leq q \leq q_{max}$ for all marginal costs less than 5, and profit is zero for all marginal costs greater than or equal to 5.

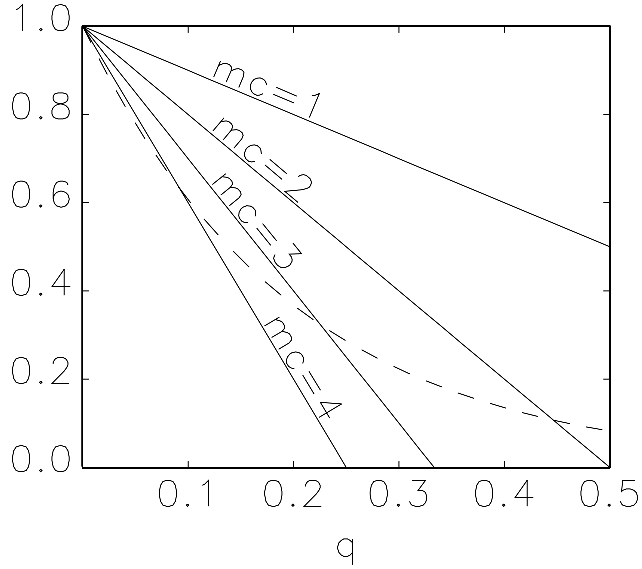

price ($c_1 < K$), profit is non-negative as $q(t)$ goes to zero. There is, however, a maximum production level constraint, q_{max}, above which profit is negative. The locus of points at which the dashed line intersects the solid lines defines the values of q_{max}. It will be shown that this constraint is not binding, however.

The solution proceeds as for the costless monopoly model, with the same form for the shadow price (19). Thus

$$q(t) = \frac{1}{K}\left(\ln \frac{K}{m_0 e^{rt} + c_1}\right) \tag{40}$$

The transversality condition $\mathcal{H}(T) = 0$ is satisfied when $q(T) = 0$, so

$$\begin{aligned} 0 &= \frac{1}{K}\left(\ln \frac{K}{m_0 e^{rT} + c_1}\right) \\ m_0 &= (K - c_1)\, e^{-rT} \\ q(t) &= \frac{1}{K}\left(\ln \frac{K}{(K - c_1)\, e^{-r(T-t)} + c_1}\right) \end{aligned} \tag{41}$$

Unlike the costless and fixed cost models, the rate of change in the production path is not constant, since

$$\begin{aligned} e^{-Kq(t)} &= \frac{(K-c_1)\,e^{-r(T-t)}+c_1}{K} \\ -K\dot{q}e^{-Kq(t)} &= \frac{r\,(K-c_1)\,e^{-r(T-t)}}{K} \qquad (42) \\ \dot{q}e^{-Kq(t)} &= -\frac{r}{K}\left[\frac{(K-c_1)\,e^{-r(T-t)}+c_1}{K}-\frac{c_1}{K}\right] \\ \dot{q} &= -\frac{r}{K}\left[1-\frac{c_1}{K}e^{Kq(t)}\right] \qquad (43) \end{aligned}$$

The terminal time T is found by integrating

$$x_0 = \int_0^T \frac{1}{K}\ln\frac{K}{(K-c_1)\,e^{-r(T-t)}+c_1}dt \qquad (44)$$

$$Kx_0 = \int_0^T \ln\frac{K}{(K-c_1)\,e^{-r(T-t)}+c_1}dt$$

Equation (44) is solved numerically in GAUSS using the parameters from Section 2.1.[7] The numerical solution for T as a function of marginal cost is shown in Figure 3. Once T for a given marginal cost is known, the initial production level is determined from

$$q\,(0) = \frac{1}{K}\left(\ln\frac{K}{(K-c_1)\,e^{-rT}+c_1}\right) \qquad (45)$$

Values of $q(0)$ corresponding to the numerical solutions for T are also shown in Figure 3. Also shown in the figure are the ranges of the rate of change in production, $\dot{q}(t)$, computed from (43). Note that $\dot{q}(t)$ starts negative and becomes more negative over the course of production, so that $q(0)$ is the maximum production level. It is clear from this plot that q_{max} is always above $q(0)$, which is always above $q(t)$, so that the maximum constraint never holds. The GAUSS procedures for solving T and computing total profit, $\Pi^{MC}$are included in the Appendices.

To compute percent change in marginal net profit, substitute for the exponential on the left-hand side of (42) and simplify

[7] The GAUSS code for numerical solutions is available from http://www.unm.edu/~ddixon (last accessed 28 February 2010)

Figure 3: Marginal cost model - numerical solutions for terminal time T.
$T$is solved numerically from equation (44). Initial production level $q(0)$ is solved using T. Also shown is the production maximum q_{max} computed from equation (39) using the equality condition. The lower graph shows the production rate of change as a function of marginal cost, with the arrow depicting the trajectory over time for a specific marginal cost. Also shown in the bottom plot is total profit as a function of marginal cost.

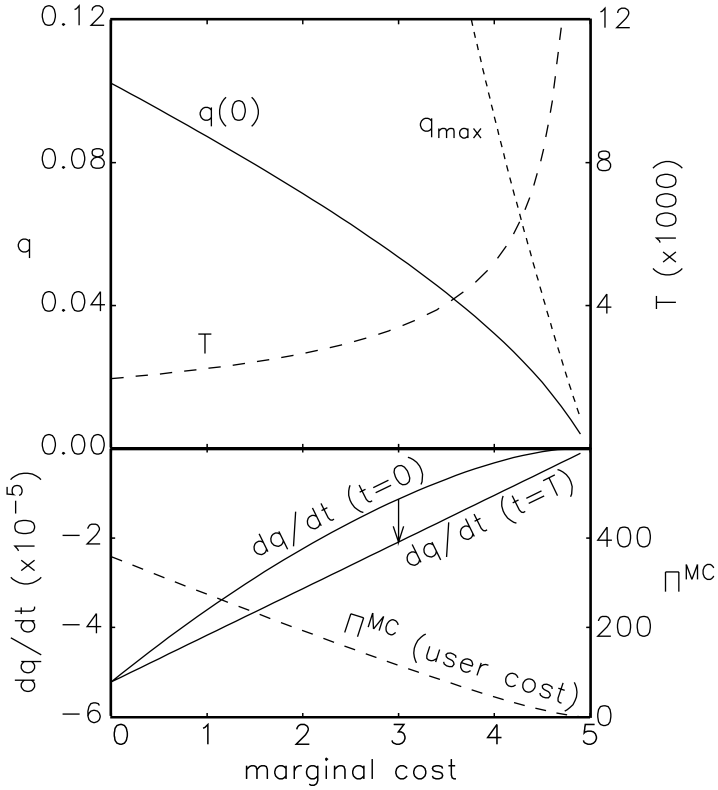

$$
\begin{aligned}
-K\dot{q}\frac{(K-c_1)\,e^{-r(T-t)}+c_1}{K} &= \frac{r\,(K-c_1)\,e^{-r(T-t)}}{K} \\
\dot{q} &= -\frac{r}{K}\frac{(K-c_1)\,e^{-r(T-t)}}{(K-c_1)\,e^{-r(T-t)}+c_1} \\
&= -\frac{r}{K}\left(1+\frac{c_1}{K-c_1}e^{r(T-t)}\right)^{-1}
\end{aligned}
$$

From this it is obvious both that the magnitude of the rate of change decreases with increasing c_1, and that the magnitude increases over time for a given c_1. Finally,

$$
\begin{aligned}
\frac{\dot{\pi}_q}{\pi_q} &= \frac{-K^2\dot{q}e^{-Kq(t)}}{Ke^{-Kq(t)}-c_1} \\
&= -K\dot{q}\frac{K}{K-c_1e^{Kq(t)}} \\
&= -K\left[-\frac{r}{K}\left[1-\frac{c_1}{K}e^{Kq(t)}\right]\right]\left[\frac{K}{K-c_1e^{Kq(t)}}\right] \\
&= r\left[\frac{K-c_1e^{Kq(t)}}{K}\right]\left[\frac{K}{K-c_1e^{Kq(t)}}\right] \\
&= r
\end{aligned}
$$

which is Hotelling's Rule.

2.5 A stock cost model

Stock costs - costs associated with cumulative production - are mentioned specifically by Hotelling (1931, p. 152) as a detail omitted from his model. Stock effects appear in many forms in natural resource production models. Lecomber (1979, p 54) sites the examples of decreasing pressure over the lifetime of an oil well, increased transportation costs as a mine becomes deeper, and a reduction in yield as the quality of ore decreases.[8] Like marginal cost models, stock cost models are often quadratic or in Cobb-Douglas form (Young, 1992).

The functional forms of stock effects in general vary broadly. In fishery models, for example, the stock variable may appear in the growth function as second-degree polynomials (Hanley et al., 1997, sec. 7.4). In econometric analysis of oil production in the U.K., Pesaran (1990) finds that production cost is inversely proportional to remaining stock. Pindyck (1978) presents a production model that includes growth from exploration, and finds an inverse relation between exploration and stock. Slade (1982) finds evidence of a

[8]Slade (1984) also points out that yield in copper mining depends on price: when the price is high, more expensive processing is used, which increases the yield.

cost curve that is U-shaped in cumulative production. Tietenberg and Lewis (2000, p. 149) present a resource model for which the stock cost is linear with cumulative production.

For simplicity, this model employs a stock cost that is linear in the stock variable $x(t)$. The stock cost c_S is

$$c_s(x(t)) = c_2(x_0 - x(t)) \tag{46}$$

where c_2 is the marginal cost of stock depletion. Note that

$$x_0 - x(t) = \int_0^t q(t')dt'$$

so that c_S is identical to a cost based on cumulative production.

The net profit function is

$$\pi(p(q(t)), q(t), t) = p(q(t)) \cdot q(t) - c_2(x_0 - x(t))$$

which, for the inverse demand function in Section 2.1, is

$$\pi(q(t), x(t)) = 1 - e^{-Kq(t)} - c_2[x_0 - x(t)] \tag{47}$$

Coming into the profit function via the state variable $x(t)$ means that

$$\pi_x = c_2 \tag{48}$$

Unlike the preceding models, the general form of Hotelling's Rule (13) applies rather than (14).

With cost based on cumulative production, it is possible for marginal cost to exceed marginal profit as the stock diminishes. Again it is necessary to invoke the non-negative profit constraint, but unlike the fixed cost and marginal cost models, this constraint can be binding.

If production is to halt when cost exceeds revenue, the producer will optimize such that (47) is non-negative at all times. For some values of c_2, this can be maintained until the stock is physically depleted, and terminal shadow price can be positive. In other cases, however, this results in production ceasing before the stock is physically depleted, so that $x(T) > 0$. In this case, the transversality condition (9) requires that $m(T) = 0$.

The first order necessary condition for $\mathcal{H}_q$ proceeds as for the costless model up to (20). From the first order necessary condition for $\mathcal{H}_x$ (11)

$$\dot{m}(t) = rm(t) - c_2$$

This differential equation has the solution

$$m(t) \;=\; \frac{c_2}{r}\left(1-e^{rt}\right)+Ce^{rt}$$

where C is a constant of integration. Solving for C using the yet-to-be-determined terminal shadow price m_T,

$$m(t) \;=\; \frac{c_2}{r}\left[1-e^{-r(T-t)}\right]+m_T e^{-r(T-t)} \tag{49}$$

Note that the static solution to the differential equation, in which $rm(t)=c_2$, is recovered for $m_0=m_T=c_2/r$. It will be seen that this static condition represents the transition from decreasing to increasing production path.

Assuming that T is noninfinite, there are two forms for the terminal Hamiltonian, depending on whether $x(T)=0$ or $x(T)>0$. For $x(T)=0$, $m(T)>0$, so

$$\mathcal{H}(T)=1-e^{-Kq(T)}-c_2x_0-m_Tq(T) \;=\; 0$$

which yields

$$e^{-Kq(T)}=\frac{1-c_2x_0}{1+Kq(T)} \tag{50}$$

This has a unique solution for $q(T)$ given c_2, as long as $c_2<\frac{1}{x_0}$. This is solved numerically in GAUSS.[9]

The condition $x(T)>0$ arises because marginal profit becomes negative before the stock is physically depleted. Marginal profit going to zero implies also that shadow price of the next unit of resource is zero. That is, $m(T)=0$, which is the transversality condition for $x(T)>0$. Profit going to zero provides an additional constraint on $q(T)$,

$$1-e^{-Kq(t)}-c_2\left[x_0-x(T)\right]=0 \tag{51}$$

The first order necessary condition (19) implies that, if $m(T)=0$, then $e^{-Kq(T)}=0$, so that (51) becomes

$$x(T)=x_0-\frac{1}{c_2} \tag{52}$$

Clearly, this only holds for $c_2\geq\frac{1}{x_0}$. Thus, $c_2=\frac{1}{x_0}$ marks the transition between physical depletion of the stock with a non-zero terminal shadow price, and economic depletion, with some physical stock remaining and a zero terminal shadow price. That is

[9] The GAUSS code for numerical solutions is available from http://www.unm.edu/~ddixon (last accessed 28 February 2010)

$$
\begin{array}{lcll}
m(T) > 0 & , & x(T) = 0 & \text{for} \quad c_2 < \frac{1}{x_0} \\
m(T) = 0 & , & x(T) > 0 & \text{for} \quad c_2 > \frac{1}{x_0}
\end{array}
$$

Finally, in this regime, the terminal Hamiltonian is

$$
\mathcal{H}(T) = 1 - c_2 (x_0 - x(T)) \quad = \quad 0
$$

which is satisfied by the terminal stock level (52). Finally, terminal time T is found by integrating

$$
x_0 - x(T) \quad = \quad \int_0^T q(t)\, dt \tag{53}
$$

where

$$
q(t) \quad = \quad \frac{1}{K} \ln \frac{K}{\frac{c_2}{r}\left[1 - e^{-r(T-t)}\right] + m_T e^{-r(T-t)}} \tag{54}
$$

The numerical solutions for the $x(T) = 0$ regime involve solving for $q(T)$ using (50), computing $m(T)$ from (19), then numerically integrating (54) to find the T that solves (53) with $x(T) = 0$. For the $x(T) > 0$ regime, m_T is assumed zero, and (54) is integrated to find the T that solves (53) where $x(T)$ is found using (52).

On a final note, (13) implies that the percent change in marginal net profit changes over time, since π_x is constant while π_q, equation (18), is a function of time . The percent change in marginal net profit is positive for

$$
r \quad > \quad \frac{\pi_x}{\pi_q} = \frac{c_2}{K e^{-Kq(t)}}
$$

percent change in marginal net profit is computed from the derivative of π_q with respect to time

$$
\frac{\dot{\pi}_q}{\pi_q} \quad = \quad \frac{\frac{d}{dt} K e^{-Kq(t)}}{K e^{-Kq(t)}} = -K\dot{q} \tag{55}
$$

where $\dot{q}$ is the time rate of change in production. This is found by taking the derivative of (19) with respect to time

Figure 4: Stock cost model - theoretical values for terminal time, initial production level, starting and ending shadow price, ending stock level, producer profit, and user cost as a function of stock cost.

The top plot shows the numerical solutions for initial production $q(0)$, starting shadow price $m(0)$ and ending shadow price $m(T)$ as a function of stock cost parameter c_2. The middle plot shows the numerical solutions for terminal stock $x(T)$, total producer profitΠ^{SC}, and user cost as a function of c_2. The bottom plot shows termination time T, percent change in marginal net profit and percent change in marginal net profit plus stock cost. This last value should equal the discount rate (see equation 13).

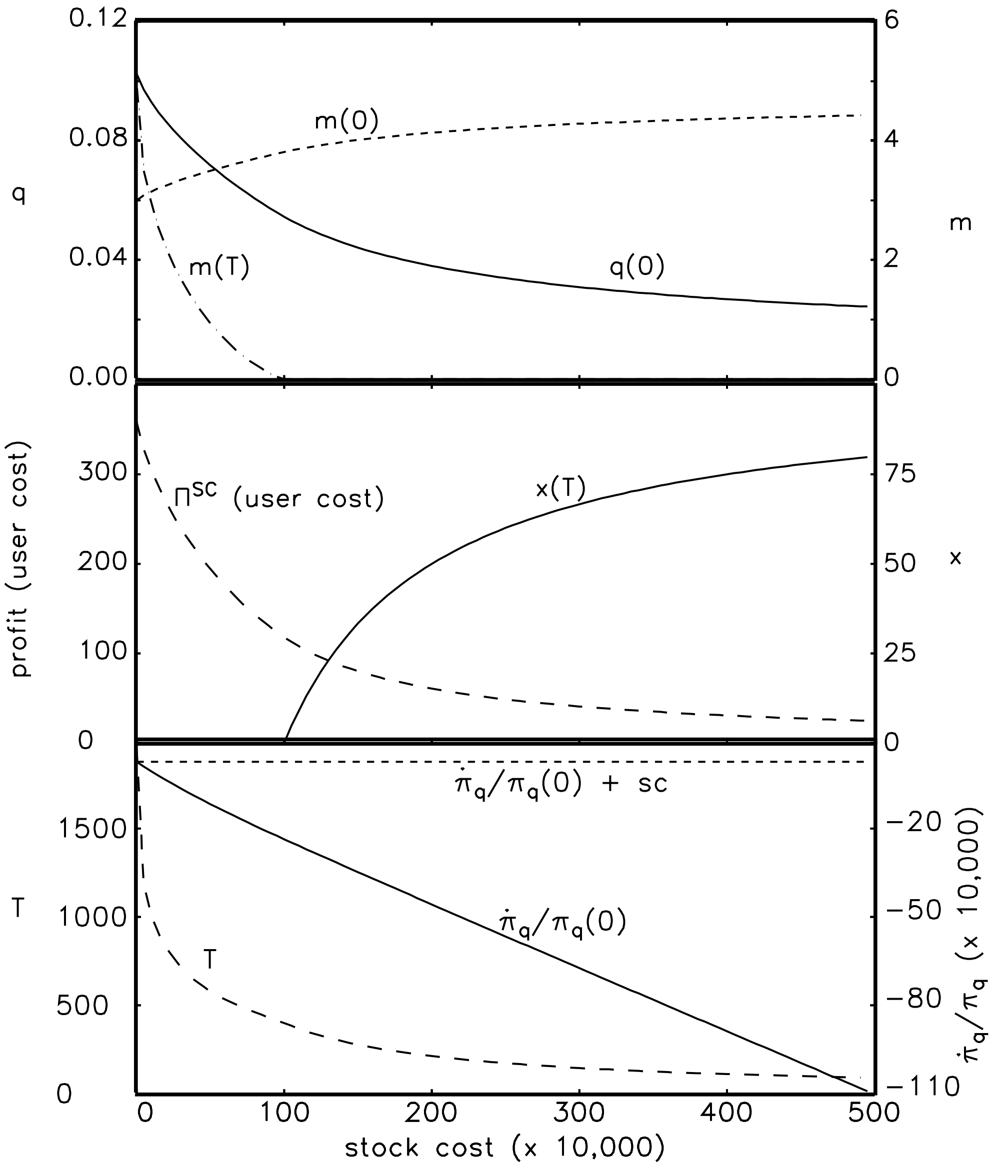

$$\begin{aligned}\frac{d}{dt}Ke^{-Kq(t)} &= \frac{d}{dt}m(t)\\ -K^2\dot{q}e^{-Kq(t)} &= \dot{m}\end{aligned}$$

Substituting (12) into the left hand side and (11) into the right hand side

$$\begin{aligned}-Km(t)\,q\dot{(t)} &= rm(t) - c_2\\ q\dot{(t)} &= -\frac{1}{K}\left(r - \frac{c_2}{m(t)}\right)\\ &= -\frac{r}{K}\left(1 - \frac{c_2}{rK}e^{kq(t)}\right)\end{aligned}$$

For small enough c_2, the production path will be downward sloping. For higher c_2, it will be upward sloping. Substituting back into the percent change in marginal net profit (55)

$$\frac{\dot{\pi}_q}{\pi_q} = r - \frac{c_2}{m(t)} \tag{56}$$

Finally, replacing c_2 from (48) and $m(t)$ from (12),

$$\frac{\dot{\pi}_q}{\pi_q} = r - \frac{\pi_x}{\pi_q}$$

This is Hotelling's Rule for nonzero stock cost as seen in (13). $\dot{\pi}/\pi$ in Figure 4 is computed using (56).

2.6 Oligopoly models

Perhaps the most straightforward definition of an oligopoly market is in terms of what it is not. It is not a monopoly market - there is more than one producer. Nor is it a competitive market, if a competitive market is defined as one in which there is a large number of producers, no one of which can affect the market equilibrium when acting independently. There are only a few ways in which an oligopoly producer can affect equilibrium, however, each depending on the reaction of the rest of the producers in the market. For example, total profit is maximized in a monopoly market, so if all the oligopolists can agree to hold their combined production to the monopoly level, the average profit per producer is maximum. This is a collusion market. At the other extreme, they can engage in price competition, driving the price down to marginal cost and eliminating economic profit altogether, and possibly driving higher-cost producers out of the market. This is the outcome of the price-competition, or Bertrand, oligopoly model. The other possible outcomes are modeled based on quantity competition (Cournot oligopoly model), market leadership (Stackelburg oligopoly model) or product differentiation (Bertrand oligopoly with product differentiation). These models have the distinction of giving the producers levels of profit intermediate between collusion and perfect competition.

Qualitatively, the expectations of an oligopoly market are:

- If the total production path is similar to the monopoly production path, it is a collusive market
- If total production is high and market price trends down to marginal cost then price competition is occurring
- If total production is higher than monopoly but lower than price competition, then there is production-level cooperation (Cournot or Stackelburg) or price competition with product differentiation (Bertrand).

2.7 Summary of the theoretical models

This section develops optimal control solutions to the production paths for the Hotelling costless model plus extensions for fixed cost, marginal cost, and stock cost production technologies. For the inverse demand function in Section 2.1, only the costless model can be solved in closed form. The others are solved numerically using the parameter values in Section 2.1.

Table 1 compares the production paths for the models in terms of initial production level $q(0)$, the slope of the production path, the terminal production level q_T, and the percent change in marginal profit $\dot{\pi_q}/\pi_q$. Clearly, for the fixed cost model, the initial production level increases with cost. For the marginal cost model, Figure 3 shows that initial production level decreases with cost. For the stock cost model, Figure 4 shows that initial production level also decreases with cost.

For the fixed cost model, the downward slope of the production path is identical to the costless model. For the marginal cost model, the production path becomes less steep with increasing cost. For the stock cost model, the slope also becomes less steep with increasing cost, becoming positive for $c_2 > 1/x_0$. This is illustrated by the curves for starting and ending shadow price, $m(0)$ and m_T in Figure 50. Recall that the production path trends in the opposite direction of the shadow price. In the figure, shadow price trends upward when $m(0)$ is below m_T, and downward otherwise. Thus, the production path is increasing for $c_2 > 1/x_0$. In all other cases and all other models, the production path trends downward. Note also that in the stock cost model the shadow price eventually descends to zero as cost increases, and that at costs above this, the final stock reserve is non-zero. These are the costs at which the stock is economically depleted before it is physically depleted.

For the costless and marginal cost models the ending production level is zero. For the fixed cost model, Figure 1 shows that the ending production level trends upward with increasing cost, nearly parallel to the starting production level. Thus, despite the fact that

Table 1: Production path comparison.

model	$q(0)$	slope	q_T	$\dot{\pi_q}/\pi_q$
costless	$\sqrt{\frac{2rx_0}{K}}$	$-\frac{r}{K}$	0	r
fixed cost	$\sqrt{q_T^2 + \frac{2rx_0}{K}}$	$-\frac{r}{K}$	> 0*	r
marginal cost	$\frac{1}{K}\left(\ln \frac{K}{(K-c_1)e^{-rT}+c_1}\right)$	$-\frac{r}{K}\left[1 - \frac{c_1}{K}e^{Kq(t)}\right]$	0	r
stock cost				
$c_2 < \frac{1}{x_0}$	$\frac{1}{K}\ln \frac{K}{\frac{c_2}{r}[1-e^{-rT}]+m_T e^{-rT}}$	$-\frac{r}{K}\left(1 - \frac{c_2}{rK}e^{Kq(t)}\right)$	> 0**	$r - c_2/\pi_q$
$c_2 \geq \frac{1}{x_0}$	$\frac{1}{K}\ln \frac{K}{\frac{c_2}{r}[1-e^{-rT}]}$	$-\frac{r}{K}\left(1 - \frac{c_2}{rK}e^{Kq(t)}\right)$	$\to \infty$***	$r - c_2/\pi_q$

* Solved numerically from (32)

** Solved numerically from (50)

*** Truncated at $q_T \ll \infty$ by the numerical integration (53)

the production path has the same downward slope as the costless model, the production level starts and ends higher as cost increases. The higher production levels result in more rapid physical depletion of the stock. For the stock cost model the ending production level increases from zero as c_2 increases, going to infinity for $c_2 \geq 1/x_0$. Were there a closed-form solution for this model, an additional capacity constraint would have to be added, but the numerical solution terminates when cumulative production reaches x_0, before reaching the instantaneous infinite production level.

3 The Hotelling ABMs

There are two general definitions of an agent in economics. The principal-agent model defines an agent as someone who acts on the part of another((Varian, 1992, p. 441)). In other cases an agent is defined as a self-interested actor with an endogenous state (e.g. income and wealth), as well as a representation of behavior (e.g. utility function for a consumer, production function for a producer)[10]. Though typically implicit in economic models, there also may be exogenous state variables (e.g. GDP) and behaviors (e.g. a supply or demand function.) In agent-based modeling, the model is a computer program representation of the agents. In its broadest usage, agent-based modeling also includes computer simulations of a model in which agents interact via their behaviors and modify their own states, the states of other agents, or the exogenous state (Epstein and Axtell, 1996, p. 4). Implicit in

[10]The term is not defined, but definition can be inferred from the literature. See, for example, Hartley (1996).

agent-based modeling is that the agents are autonomous (Tesfatsion, 2006, p. 843).

The agents in the ABM in this paper have no information about the demand function itself, each determining autonomously its own optimal production path. The extent of the resource is known exactly, but the market structure and demand function are unknown. These are adaptive agents for which the behavioral rule is a heuristic to continually adjust the production level so that estimated total profit is maximized. Profit estimates are based on the observed market response to changes in production level. In addition to the costless basic models, there are models with non-zero cost which may be constant (per period), marginal (per unit production) or cumulative (proportional to the stock level).

The models are constructed and the simulations run using the MASON agent-based modeling and simulation library and framework. They are based on and incorporated into a set of programs included with MASON to demonstrate the MASON Console environment. The Console provides a general graphical interface for editing model parameters, running simulations and for viewing model variables as time-series graphs or numerical tables.

The graphical interface for each Hotelling model provides the ability to select a demand function and one of the various market models described in the following sections, allowing the user to change the number of producers in the market. Once a specific demand function and market model has been selected, the user can change model-specific variables, such as the mean and standard deviation for cost variables, for example. While running, the simulation displays a custom window showing real-time plots of current profit, current percent change in marginal profit, stock level and production level. Each model optionally writes a file of key dynamical variables. These files were used to produce the plots presented in Section 4. Images of the interface and results windows are included in the Appendices.

Each model has two types of agent: a market agent and a producer agent. In a given model there is a single market agent and one or more producer agents. The simulation is initialized with Monte Carlo draws for the stochastic variables, then the simulation proceeds, one time-step at a time, until all producers have stopped. The producers stop either because the resource stock level is zero, or profit in the current period is negative.[11] Because the agents are autonomous, the simulation behaviors are mediated by information that is communicated between agents, specifically between each firm agent and the market agent. These exchanges occur as four distinct actions during each time-step, as illustrated in Figure 5. The size of a simulation time-step is arbitrary, though the default discount rate is assumed daily and compounds to ten percent per annum. Changing the discount rate, via the GUI, changes the implied time-step.

[11] In some models, profit goes negative even though an alternative production level would produce positive profit. A more advanced heuristic could explore alternative production levels to determine if this is the case, but the simple heuristic does not. This is not dissimilar to a situation in which the owner of the resource prefers to shut down leaving a small reserve rather than take the risk of incurring further negative profits while searching for a profitable production path.

Figure 5: The details of a time-step
The dashed box represents a single simulation time-step, during which four distinct actions mediate the exchange of information between each producer agent and the market agent.

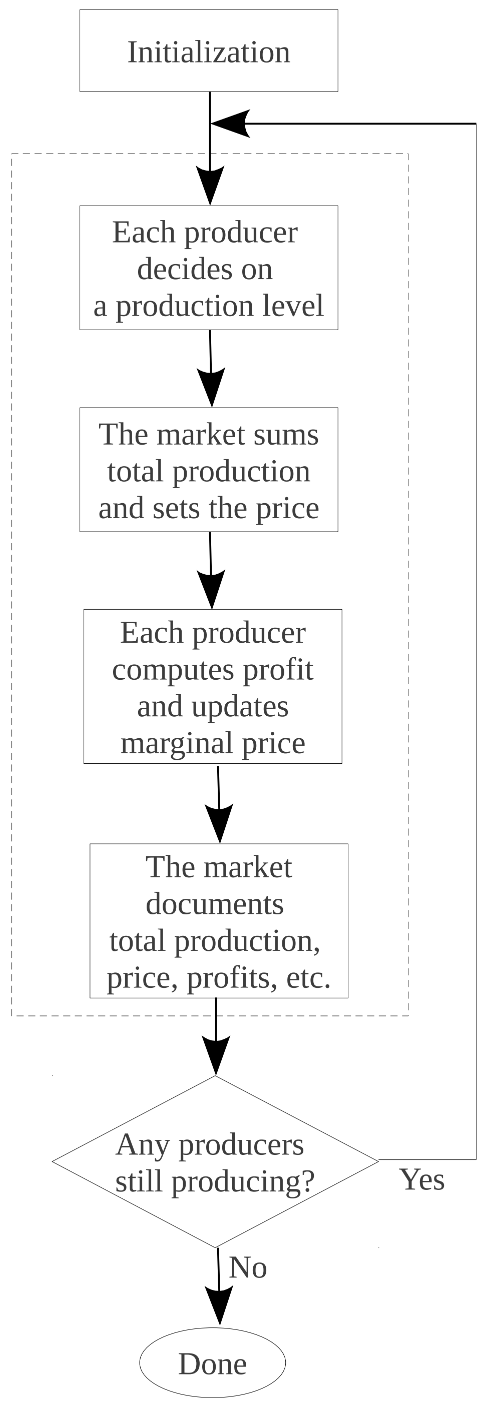

3.1 The market agent

The market agent, called Market, is assigned a demand function and controls any market information provided to the producers. Both the specific market agent and the demand function are user-selectable: changing from one model to another is simply a matter of changing the market agent and/or the market agent's demand function. The market agent represents a specific market structure and production technology. For example, there is a costless market agent, a fixed-cost market agent, an oligopoly market agent and so on.

3.2 The producer agent

In contrast, the producer agent, called Firm, is the same for all models. The producer agent computes profit, marginal price, marginal profit and the percent change in marginal profit for each time period. This information is used by the producer agent at the beginning of each time period to compute the next production level and is collected at the end of each time period by the market agent for the real-time displays and, optionally, written to a file for post-processing.

Although the producer is the optimizing agent, the algorithms to compute future profit are contained in the market agent so that the details can vary depending on the market model. A simple heuristic is used for optimization and is described in the next section.

3.3 A simple optimization heuristic

The first step in designing the model is to find an optimization heuristic that is as simple as possible while reproducing plausible behavior. Algorithmic simplicity contributes to robustness, that is, the ability to produce consistent behavior under the planned variety of production technologies and market structures. For example, in a tournament of bidding algorithms, Rust et al. (1992) found that the simplest algorithms consistently beat the more complex. Another advantage to simplicity is analytic transparency. The difficulty in associating specific outcomes with specific behaviors increases as the complexity of the algorithms increases. From an experimental control perspective, it is also easier to detect, explain and compute the impact of algorithmic artifacts for a simple algorithm. Algorithmic artifacts may results from the size of the simulation time-step, the size of changes in production level, or numerical errors in calculating profit, cost or production level changes. A possible added benefit of a simple algorithm is shorter computation times, since proper Monte Carlo sampling calls for large numbers of simulations.

The core of the heuristic is a simple estimation of future profit. The heuristic uses the profit estimation in two different ways, depending on the phase of the simulation. The phases are:

1. Increase production level from zero until estimated future profits begin to fall. This is called the *ramp-up* phase.

Figure 6: An example of the initial production level error.
At the beginning, the agent starts increasing production by Δq. In this example, the heavy line is the Hotelling's Rule optimal production path. The production level after the first increment is too low, and after the second increment it is also too low, but after the third increment, the production level is too high. At this point, the agent begins a constant downward production path but, because the initial production level is above the optimum, the stock is depleted more quickly than optimal, and the stock is depleted sooner than the Hotelling's Rule optimum.

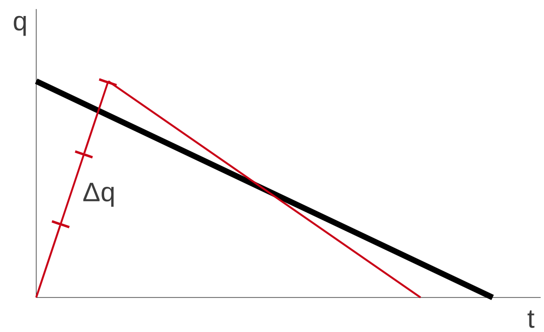

2. Estimate future profits in each time period based on the three candidate strategies described below and adopt the strategy that maximizes future profit. This is the *optimization* phase.

Each phase encompasses a number of time-steps. The ramp-up phase is not intended to simulate a real-world process, but it is a way for the heuristic to reach an efficient initial production level autonomously. The ramp-up phase provides an opportunity to estimate marginal profit for use by the optimization heuristic. The increments in production level during the ramp-up phase are coarse in order to keep the phase brief, but the coarseness makes it unlikely that an agent will reach the theoretically optimum production level exactly. This is illustrated in Figure 6. The error is useful in exploring the consequence of setting the initial production level sub-optimally.

For both the ramp-up phase and the optimization phase, future profit is estimated for three strategies, one with constant decreasing production, one which maintains the current production level, and one with constant increasing production. The decreasing strategy assumes a constant decrease of Δq^* each time-step, which will produce a straight-line decreasing production path that goes to zero when the resource stock is physically depleted. This is a simple geometric calculation that uses only information available to the agent. The increasing strategy assumes a constant increase per period that is one percent of the production level in the current period. This, too, is a simple calculation using only information available to the agent. The one percent increment is arbitrary, it is intended to be small, thus preventing large swings in production level. It is also advantageous that it be different in magnitude from the decreasing strategy, reducing the likelihood of non-damping oscillations.

3.4 The use of discrete summations

The following sections present the calculations used by the agents to determine the best optimization strategy. In contrast to the integrals presented in Section 2, these are discrete summations, and the derivatives are all discrete (e.g.Δq, Δp, $\Delta \pi$). This reflects the fact that the simulation itself employs discrete time. A producer agent has very little information and estimates future profit by counting up the discounted profit per period until the stock is physically depleted. Further advantage is taken of the fact that the producer agent's strategies all assume constant changes to the production level Δq, so that the amount of the change itself comes out of the summations.

The discrete summations also have the advantage of allowing for variable time steps. The simulations presented here all assume planning on a daily basis, which is may not be realistic. A real-world planner may reassess the production plan once a quarter or once per year. To simulate these planning periods, the user of these simulations need only change the discount rate from daily to quarterly or annual.

3.5 The heuristic algorithm

Estimated future profit is computed with all production in the future and profit discounted accordingly. For a costless model, the future profit calculation is the summation

$$\Pi_\tau = \sum_{i=0}^{\tau-1} q_i p_i (1+r)^{-i}$$

where

$$\begin{aligned} \Pi_\tau &= \text{estimated future profit} \\ q_i &= \text{production level in period } i \\ p_i &= \text{price in period } i \\ r &= \text{discount rate} \\ \tau &= \text{remaining lifetime of the stock} \end{aligned}$$

With a constant change in production level $\triangle q$ – which can be negative, positive, or zero – the production level in period i is

$$q_i = q_n + i\triangle q$$

where q_n is the base production level, meaning the production level at the time the estimate is being computed, and i enumerates the production periods into the future. The production period i starts at zero for the current production period n. Because the inverse demand function is unknown to the agent, price is estimated based on the most recent marginal price

$$p_i = p_n + \triangle q_i \left(\frac{\triangle p}{\triangle q}\right)_{i-1}$$

where p_n is the price in the current period, and

$$\left(\frac{\triangle p}{\triangle q}\right)_{i-1} = \frac{p_{i-1} - p_{i-2}}{q_{i-1} - q_{i-2}} \tag{57}$$

is the estimated marginal price based on price and production level changes between the previous two periods. Estimated future profit becomes

$$\begin{aligned} \Pi_\tau(\triangle q) &= \sum_{i=0}^{\tau-1} (q_n + i\triangle q)\left[p_n + i\triangle q\left(\frac{\triangle p}{\triangle q}\right)_{i-1}\right](1+r)^{-i} \\ &= q_n p_n A + \triangle q\left[p_n + q_n\left(\frac{\triangle p}{\triangle q}\right)_{i-1}\right]B + (\triangle q)^2\left(\frac{\triangle p}{\triangle q}\right)_{i-1} C \end{aligned} \tag{58}$$

where

$$A = \sum_{i=0}^{\tau-1}(1+r)^{-i} = \frac{1+r}{r}\left[1-(1+r)^{-\tau}\right] \tag{59}$$

$$B = \sum_{i=0}^{\tau-1} i(1+r)^{-i} = \frac{1}{r}\left[A - \frac{\tau}{(1+r)^{\tau-1}}\right] \tag{60}$$

$$C = \sum_{i=0}^{\tau-1} i^2(1+r)^{-i} = \frac{1}{r}\left[2B + A - \frac{\tau^2}{(1+r)^{\tau-1}}\right] \tag{61}$$

The lifetime of the remaining stock comes from the constraint that total production equal total current stock

$$\begin{aligned} x_n = \sum_{i=0}^{\tau-1}(q_n + i\triangle q) &= \tau q_n + \triangle q\frac{\tau(\tau-1)}{2} \\ \tau &= \frac{-\left(q_n - \frac{\triangle q}{2}\right) + \sqrt{\left(q - \frac{\triangle q}{2}\right)^2 + 2x_n\triangle q}}{\triangle q} \end{aligned} \tag{62}$$

where x_n is the reserve stock in the current period. There exists some minimum constant production change Δq^* for which the total remaining stock is exhausted, at which point production goes to zero. The lifetime in this case is constrained by

$$q_n + \sum_{i=0}^{\tau-1} \Delta q^* = 0 \tag{63}$$

and the constraint that total production equal the current stock by

$$\sum_{i=0}^{\tau-1} (q_n + i\Delta q^*) = x_n \tag{64}$$

Solving (63) for Δq^* and substituting into (64)

$$\tau = \frac{2x_n}{q_n} - 1 \tag{65}$$

Substituting this back into (63),

$$\Delta q^* = -\frac{q_n^2}{2x_n - q_n} \tag{66}$$

This is the lowest (most negative) $\triangle q$ that will result in a straight-line decreasing production path for which production goes to zero as the stock is physically depleted. This also satisfies the constraint that the term in the radical in (62) be non-negative.

3.6 Costless model

Initially, consistent with Hotelling, costless production is considered. The introduction of nonzero cost will be presented in Section 3.7. Production decisions use a heuristic that estimates future profit as described in Section 3.3.

The theoretical maximum profit is shown in (30). The equivalent summation expression is

$$\Pi_{max} = \sum_{i=0}^{T-1} \left(1 - e^{-Kq_i}\right)(1+r)^{-i} = \frac{1+r}{r}\left[1 - (1+r)^{-T}\right] - \frac{(1+r)^{-T} - e^{-rT}}{e^r (1+r)^{-1} - 1} \tag{67}$$

With the values given in Section 2.2, the discrete $\Pi_{max} = 358.53$, as compared to the continuous $\Pi_{max} = 358.33$, the difference being due to numerical errors. In practical terms, the monopoly models are special cases of the oligopoly models, so simulation results of the monopoly models are discussed in their respective oligopoly sections in Section 4.

A producer agent has no knowledge of the demand function, and can only infer it from the observed behavior of the market. Namely, the change in price that results from changes in the production level. The agent has two behavioral rules, corresponding to the two phases of the heuristic:

Rule 1. In the ramp-up phase, increase the production level from zero until estimated total profit is positive and begins to decrease. The default monopoly ramp-up rate is an increase of 0.01 units of production per period.

Rule 2. In the optimization phase, in each period, estimate future profit based on the three production strategies, then execute the strategy that maximizes future profit. Production ceases if, after the ramp-up phase, profit becomes negative.

3.7 Production technologies with nonzero cost

The production technology models with nonzero cost are the costless model with a nonzero cost term. This does not require a change in the producer agent, which is implemented with cost variables, all of which were zero for the costless model.

3.7.1 Fixed cost model

Recall from Section 2.3 that, for the fixed cost model, there is a minimum production level q_{min} below which profit is negative. For the discrete calculations used by the heuristic, the constraint (63) becomes

$$q_n + \sum_{i=0}^{\tau-1} \Delta q^* \quad = \quad q_{min} \tag{68}$$

which, when substituted into equation (64), means that equation (66) becomes

$$\triangle q^* \quad = \quad \frac{q_n^2 - q_{min}^2}{2x_n - q_n + q_{min}} \tag{69}$$

This is the minimum (most negative) change in production that results in a straight-line decreasing production path that reaches q_{min} at the moment the stock is physically depleted. The heuristic determines q_{min} by increasing production starting from zero and recording the production level at which profit becomes positive.

The estimate of future profit is

$$
\begin{aligned}
\Pi_\tau^{FC} &= \Pi_\tau^{NC} - c_0 \sum_{i=0}^{\tau-1} (1+r)^{-i} \\
&= \Pi_\tau^{NC} - c_0 A
\end{aligned}
$$

where Π_τ^{NC} is the no-cost future profit estimate (58) and A is from (59).

A fixed cost does not appear in π_q or π_x, so according to equation (13), there is no effect on the optimal percent change in marginal profit. The optimal production path is affected, however, since the initial production level $q(0)$ and the terminal production level q_T both increase with cost, as shown in Figure 1.

3.7.2 Marginal cost model

With no minimum production level constraint, the marginal cost model is identical to the costless model. The change comes in the estimate of future profit

$$
\begin{aligned}
\Pi_\tau^{MC} &= \Pi_\tau^{NC} - c_1 \sum_{i=0}^{\tau-1} (q_n + i\Delta q)(1+r)^{-i} \\
&= \Pi_\tau^{NC} - c_1 \left[q_n \sum_{i=0}^{\tau-1} (1+r)^{-i} + \Delta q \sum_{i=0}^{\tau-1} i (1+r)^{-i} \right] \\
&= \Pi_\tau^{NC} - c_1 (q_n A + \Delta q B)
\end{aligned}
$$

where Π_τ^{NC} is the costless future profit estimate (58) and A and B are from equations (59) and (60).

3.7.3 Stock cost model

With the addition of a stock cost as in (46), the future profit estimate becomes

$$
\begin{aligned}
\Pi_\tau^{SC} &= \Pi_\tau^{NC} - \sum_{i=0}^{\tau-1} c_2 (x_0 - x_i)(1+r)^{-i} \\
&= \Pi_\tau^{NC} - c_2 x_0 \sum_{i=0}^{\tau-1} (1+r)^{-i} + c_2 \sum_{i=0}^{\tau-1} x_i (1+r)^{-i} \\
&= \Pi_\tau^{NC} - c_3 \sum_{i=0}^{\tau-1} (1+r)^{-i} + c_2 \sum_{i=0}^{\tau-1} (x_n + \Delta x_i)(1+r)^{-i} \qquad (70)
\end{aligned}
$$

where $c_3 \equiv c_2 x_0$. The discrete form of equation (6) is $\Delta x = -q$. Assuming that q is changing by the constant increment Δq, then $\Delta x_i = -(q_n + i\Delta q)$. Now, (70) becomes

$$
\begin{aligned}
\Pi_\tau^{SC} &= \Pi_\tau^{NC} - (c_3 - c_2 x_n) \sum_{i=0}^{\tau-1} (1+r)^{-i} - c_2 \sum_{i=0}^{\tau-1} (q_n + i\Delta q)(1+r)^{-i} \\
&= \Pi_\tau^{NC} - (c_3 - c_2 x_n + c_2 q_n) \sum_{i=0}^{\tau-1} (1+r)^{-i} - c_2 \Delta q \sum_{i=0}^{\tau-1} i (1+r)^{-i} \\
&= \Pi_\tau^{NC} - [c_3 - c_2 (x_n - q_n)] A - c_2 B \Delta q
\end{aligned}
$$

where Π_τ^{NC} is the costless future profit estimate (58) and A and B are from (59) and (60).

3.8 Other sources of uncertainty

If the producer is not certain of the extent of the resource x_0, the consequent error in the lifetime of the stock will affect estimates of future profits. This, in turn, may affect the production strategy selected by the heuristic outlined above. Although the heuristic can adjust the rate of change as the stock is depleted, the total profit is sensitive to an error in the initial production level. An initial quantity that is too high will, in general, result in the resource being depleted too quickly, leaving unrealized profit in the future. An initial quantity that is too low will, in general, result in the resource being depleted too slowly, with unrealized profit in the present.

Errors in x_0 are similar to errors in the initial production level, so Monte Carlo sampling in the neighborhood of the initial production level will give an indication of sensitivity to errors in x_0. The relation between initial production level and initial stock is given by equation (27). The coarseness of the heuristic strategy serves as a proxy for errors in computing optima, including errors in x_0. This is illustrated in the discussion in Section 4.4.

Other sources of uncertainty in the interest rate, in the demand function, and in the production technology cost function could be explored in a similar manner. Uncertainty in the demand function can take on various forms, the simplest being random errors in constants and systematic errors in functional form. In the former, a sufficiently large sample reveals a constant variance while, in the latter, a large sample reveals variance that changes over the range of production. Uncertainty enters the cost function in ways similar to the demand function. Of particular interest are cases in which the production planner incorrectly assumes costless production. These issues are beyond the scope of this paper.

4 Simulation results

The ABMs are oligopoly models for which the monopoly results are special cases. The oligopoly simulations have one market agent and one or more producer agents, depending on the number of producers in the market. The number of producers is a user-set variable in the GUI. Each producer is unaware of the others. In terms of the ABM architecture, this is done by not providing any communication between producer agents.

The ensemble simulations are models in which there are multiple producer agents and multiple market agents. Each pair of producer agent and market agent behaves like a monopoly with dedicated stock and a dedicated market. Ensembles are a way to collect data about large number of monopolists while running only one simulation. In these models, the monopolists are all different because each one has been given production technology cost parameters drawn at random from statistical distributions. This method of mapping the parameter space onto the outcome space is called Monte Carlo sampling.

The first sections will discuss the costless oligopoly models. The monopoly model for each production technology is presented as an oligopoly with one producer. The following sections will address the ensemble models. The last section will examines the efficiency of the heuristic by introducing intentional error into the initial production level.

The models in the following sections are intentionally wrong. In these models the producer always behaves as though it is a monopoly market, even though the models include oligopolies of two to six producers. Given that a real-world producer is likely to make mistaken assumptions about the market structure, the object is to assess the worst case outcome of assuming no competitors whatsoever.

4.1 The costless oligopoly model

For the models in this section, the total stock is constant as the number of producers increases. For a duopoly, each producer begins with a stock of $x_0/2$, and with ten producers, each producer begins with a stock of $x_0/10$. The duopoly production path in Figure 10 reflects a collusion-like outcome, as do models for N=3 and N=4. That is, the producer agents arrive at what looks like a collusive market structure using only the optimization heuristic. For an oligopoly of five producers, however, something completely different occurs, as seen in Figure 9. In this model, after the ramp-up phase, some producers begin reducing production while others continue unchanged.

This behavior is an emergent property of the heuristic. Initially, the total stock of 100 is distributed among the five producers in near equal amounts, with small random deviations. In this model, the initial stock allocations are 20.11, 20.09, 19.90, 19.88, and 20.02 for Firm 1 through 5, respectively. All five firms conclude the ramp-up phase on day 53. At this point, Firms 1 and 2, with the largest allocations, select a decreasing strategy, while the rest of the firms select zero change strategies. That an individual producer chooses a flat production strategy is not unexpected. However, the decreases in production by the two largest producers cause price increases that are sufficient for the remaining producers to estimate increasing profits at constant production levels for the duration of their stocks. That is, Firms 3, 4, and 5 are self-optimizing to higher total production than the collusive outcome (as per Cournot-Nash equilibrium) by maintaining constant production levels while Firms 1 and 2 decrease theirs.

The results for models with from one to ten producers are shown in Figure 10. As the number of firms increases, the production paths present evidence of what Hotelling calls the "retardation of production under monopoly" (Hotelling, 1931, sec. 7) in that the lifetime of the stock decreases as the number of producers increases. The models with five, seven and ten producers show the Cournot-like outcome, while the rest show the collusion-like outcome. All of these models appear to reach the theoretically optimal profit. These models show that total collusion-level profits are not necessarily an indicator of collusion. Note also, in Figure 10, the discontinuities in the percent change in marginal profit curves where producers change strategies at the end stock life.

4.2 Accuracy and precision of the heuristic

The artifacts discussed in the preceding two sections beg the question of errors in the heuristic ramp-up and their impact on total profit. Figure 7 shows the simulation results from fifty Monte Carlo samples of the monopoly model. For each simulation, the fifth ramp-up increase was given a small adjustment $\epsilon \sim N(0, 0.00001)$. The left-hand plot shows the relationship between optimal initial production level and total profit. The upper curve represents samples which reached the ramp-up cutoff on the ninth day, the lower curve on the tenth day. The lower curve is essentially an extension of the upper: at the far right-hand side of the upper curve, it is no longer optimal to end ramp-up on the ninth day, so the next increment is to end ramp-up on the tenth day at the far left-hand end of the lower curve. The day on which ramp-up ends affects total profit in that the fewer days lost to non-optimum ramp-up production levels, the higher the total profit. The curves end abruptly at the low end because any lower sample ended ramp-up on the previous day. The right hand plot shows the relationship between percent change in marginal profit and total profit. The distribution of points is nearly identical, reflecting a nearly linear relationship between initial production level and the percent change in marginal profit. From Hotelling's Rule, maximum profit should occur for an initial production level of 0.1022 and a percent change in marginal profit of 0.000261. The curves peak very near these values. The key outcome of these plots is that the distributed error in the heuristic initial production level results in a error range of 0.45 out of about 358, or less than 0.14 percent.

Figure 7 also illustrates the impact of error in the estimate of x_0. According to equation (27), a one percent error in x_0 translates into a one-half percent error in initial production q_0, and from equation (30), that, in turn, translates into an error of 0.84 percent in the theoretical maximum profit. For the heuristic, however, Figure 7 shows a 10 percent error in q_0 leading to a 0.11 percent error in total profit. This is almost two orders of magnitude less impact than theory predicts. The heuristic lessens the impact on total profit from an underestimate of the stock by making Bayesian updates to the production level in each time-step.

Figure 7: Monte Carlo samples of heuristic outcomes.

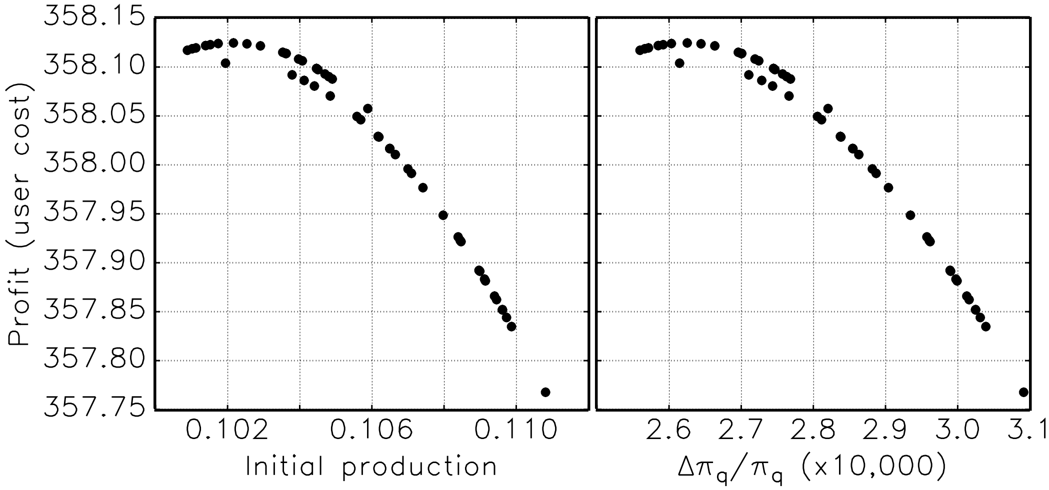

4.3 Nonzero-cost models

Three cost functions are examined: a fixed (per day) cost, a marginal (per unit) cost, and a stock (cumulative production) cost. For the fixed cost and marginal cost models, the non-negative profit restriction only comes into play when cost exceeds any feasible level of revenue, in which case the ramp-up phase ends with negative profit and the simulation terminates. In stock cost models, however, cost increases over time, and may exceed any feasible level of revenue before the stock is physically depleted. This is explored further in the section on the stock cost model.

For the inverse demand function in Section 2.1 there are no closed-form optimal control solutions when marginal or cumulative costs are included. The heuristic, however, considers only per period profit, marginal profit and marginal price. Qualitatively, Monte Carlo sampling should show a trend toward the costless behavior as cost decreases. That is, for the fixed and marginal cost models, percent change in marginal profit should approach equation (14). For the stock cost models, percent change in marginal profit should approach equation (13).

4.3.1 Fixed cost model

The fixed cost model is the costless model with a fixed cost $c_0 \sim N\,(0.012, 0.0016)$ truncated such that $c_0 \geq 0$. The mean and standard deviation were chosen such that $\pm 3\sigma = 0.24$ spans the revenue range of the heuristic's first six ramp-up iterations based on equation (17), which is now revenue rather than profit. Time-series plots for fifty Monte Carlo monopoly simulations are shown in Figure 12. For many of the samples, the heuristic selects the flat production level strategy, resulting in a zero percent change in marginal profit until the final few time-steps. In other samples, the heuristic selects the decreasing strategy followed by a flat strategy, resulting in percent change in marginal profit curves that sweep upward then fall to zero. The diamonds in the percent change in marginal profit plot indicate sharp spikes (to several hundreds on the scaled y-axis), indicating that the heuristic selects a rapid taper down to zero production over the final two or three time-steps as stock reaches physical depletion. Note that optimal percent change in marginal cost is 2.6 in this graph.

The heavy dashed line in the production path plot indicates theoretical q_{min} from equation (31). This is the envelope for optimal terminal production levels.

The results of the Monte Carlo monopoly simulations are summarized in Figure 15. Note that the results cluster based on the heuristic q_{min}, and that q_{min} is in multiples of 0.01. This is because 0.01 is the default ramp-up production change and q_{min} is recorded by the heuristic during the ramp-up phase. The inset graph compares the heuristic values for q_{min} with the theoretical values from (31). The graphs show a distinct trend in percent change in marginal profit with regard to fixed cost. The dashed lines are the theoretical values discussed in Section 2.3.

4.3.2 Marginal cost model

The marginal cost model is the costless model with a constant marginal cost $c_1 \sim N\left(1.0, 0.110889\right)$ truncated such that $c_1 \geq 0$. The time-series data for the marginal cost model are shown in Figure 13. The plot shows only three unique production paths because the initial value is weakly dependent on marginal cost, and the step size in the ramp-up heuristic is much greater than the variations between starting production levels. The time-series data also reveal that percent change in marginal profit is always close to the optimal 2.6, and trends toward it in the course of the simulation.

The results from fifty Monte Carlo monopoly simulations are shown in Figure 16. In these models, percent change in marginal profit is sensitive to the initial production level. A percent change in marginal profit below optimum implies that the initial production level is too low, preventing the heuristic from optimally reducing production each time step. Similarly, a percent change in marginal profit above optimum implies an initial production level that is too high. In all cases, the percent change in marginal profit trends downward.

4.3.3 Stock cost model

The stock cost model is the costless model. In this model, the optimum production path can be either decreasing or increasing over time, depending on the value of the stock cost parameter c_2. Also, for higher values of c_2 it is expected that production will stop before the stock is physically depleted. These characteristics are discussed in Section 2.5.

Before Monte Carlo sampling the parameter space, it is instructive examine the classes of behavior anticipated. Figure 14 shows a time-series plot for five representative values of stock cost. The stock cost parameter c_2 is shown as "sc" in the legend. The zero cost model is included for comparison and is identical to the costless monopoly model in Section 4.1. The model with $c_2 = 0.001$ is included to show the deviation of a small cost from the costless model. The model with $c_2 = 0.005$ is included to show the behavior in the vicinity of transition from decreasing production to increasing production discussed in Section 2.5 in the discussion of Figure 4. The model with $c_2 = 0.010$ is $1/x_0$, which is the cost at which the producer will begin to leave some of the stock unproduced as discussed in Section 2.5, and the model with $c_2 = 0.013$ is included to show the behavior well into the regime for which the physical stock is greater than zero when production stops.

Interesting to note here is that, for small stock cost (much less than 0.10) or large stock cost (significantly greater than 0.10), the production paths are fairly smooth. In the vicinity of 0.10, however, the heuristic makes frequent changes resulting in a highly volatile percent change in marginal profit. The percent change in marginal cost plot has been smoothed with a boxcar length of 100 days. Note also that, for smaller stock costs, the percent change in marginal profit turns toward negative infinity very quickly, and, at least for 0.005, and then comes down from positive infinity thereafter. This arises because, in this period, the production level is very close to the zero marginal profit regime. Marginal profit appears in the denominator of percent change in marginal profit, hence the switch between negative and positive infinity as marginal profit crosses through zero.

Figure Figure (17) shows the outcome space results from 50 Monte Carlo samples with stock cost taken from $c_2 \sim N(0.01, 0.000016)$, truncated such that $c_2 \geq 0$. Very little can be drawn from the percent change in marginal profit plot other than to note, as in the time-series plot, that percent change in marginal profit can be highly dynamic. The profit and user cost plot shows that the heuristic is near optimal at zero stock cost and when stock cost is 0.10. The error in the heuristic is especially large in the range between zero and 0.10. The initial production level plot provides a clue as to why: initial production levels are consistently too high, resulting in lower profits and an inefficiently rapid reduction in stock. Finally, the upper right plot shows the consequences: a much shortened lifetime for stock costs below 0.10, and too much stock left unexploited for stock costs above 0.10. The dashed lines are the theoretical values, as discussed in Section 2.5.

4.4 Theoretical efficiency

To put the foregoing results into perspective, Figure 8 shows the effect of error in the theoretical costless model. These are families of curves of total net profit from the theoretical costless monopoly model in Section 2.2. The initial production level q_0 is varied plus or minus 15 percent about the optimum of 0.1022. The production path slope is also varied so that percent change in marginal profit varies plus or minus 15 percent about the optimum of 2.61×10^{-4}.

The widely spaced diagonal lines (on the left in the left figure, on the right in the right figure) result from the initial production level being so low that the downward slope of the production path reduces production to zero before the stock is physically depleted. This represents a gross error in production planning and not one likely to be seen in real-world applications.

The closely spaced nearly horizontal lines in both figures (on the right in the left figure, on the left in the right figure) represent the performance of the Hotelling Rule optimal production path under slight deviations in q_0 and $\Delta\pi_q/\pi_q$. The inverted triangles point to the optimal control solution. The close spacing of these lines is an indication that, even for the theoretical solution, there is only a small penalty for small errors in initial production or the slope in the production path. That is, although Hotelling's Rule is the optimum, total discounted profits are only weakly affected from small deviations from Hotelling's Rule. There is less than one percent error in profit for errors of plus on minus twenty percent in initial production and plus or minus ten percent in percent change in marginal profit.

Figure 8: Errors in the theoretical model.
The optimal initial production level and optimal $\Delta\pi/\pi_q$ are indicated with an inverted triangle.

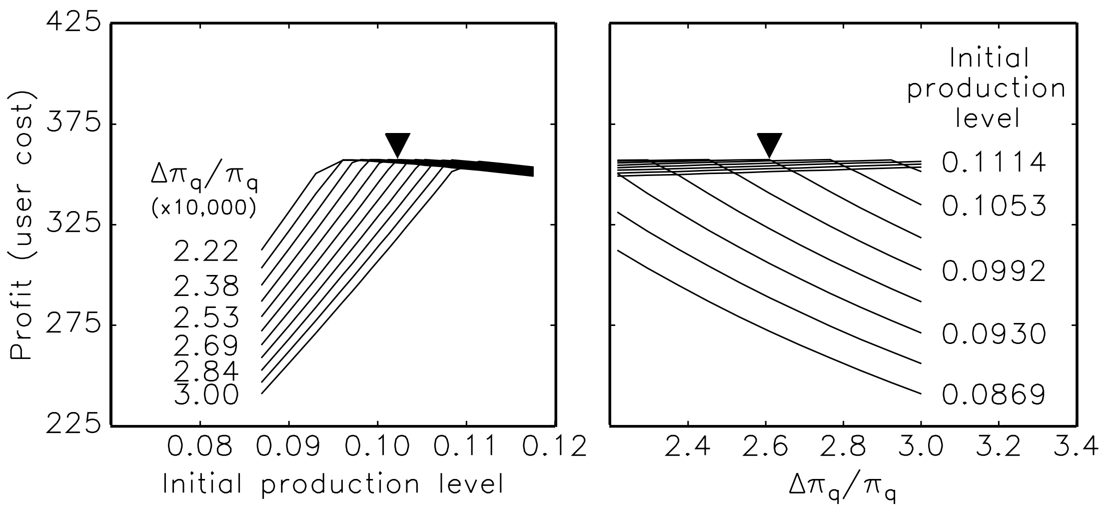

5 Conclusion

The optimization heuristic demonstrated in these models is extremely simple and inelegant. Yet the performance is effectively optimal in costless monopolies, and acceptably efficient in oligopoly markets in which there are no competitive behaviors. The oligopoly models demonstrated a tendency to reach a collusion-like market structure solely through self-interested optimization in some cases, and to reach a mixed-strategy Cournot-Nash-like equilibrium at other times.

The heuristic is somewhat less nimble in the face of non-zero costs, but net profits are within a few percent of optimum in the models tested. In general in these models, percent

Figure 9: Costless model - time-series. See the discussion in Section 4.1.

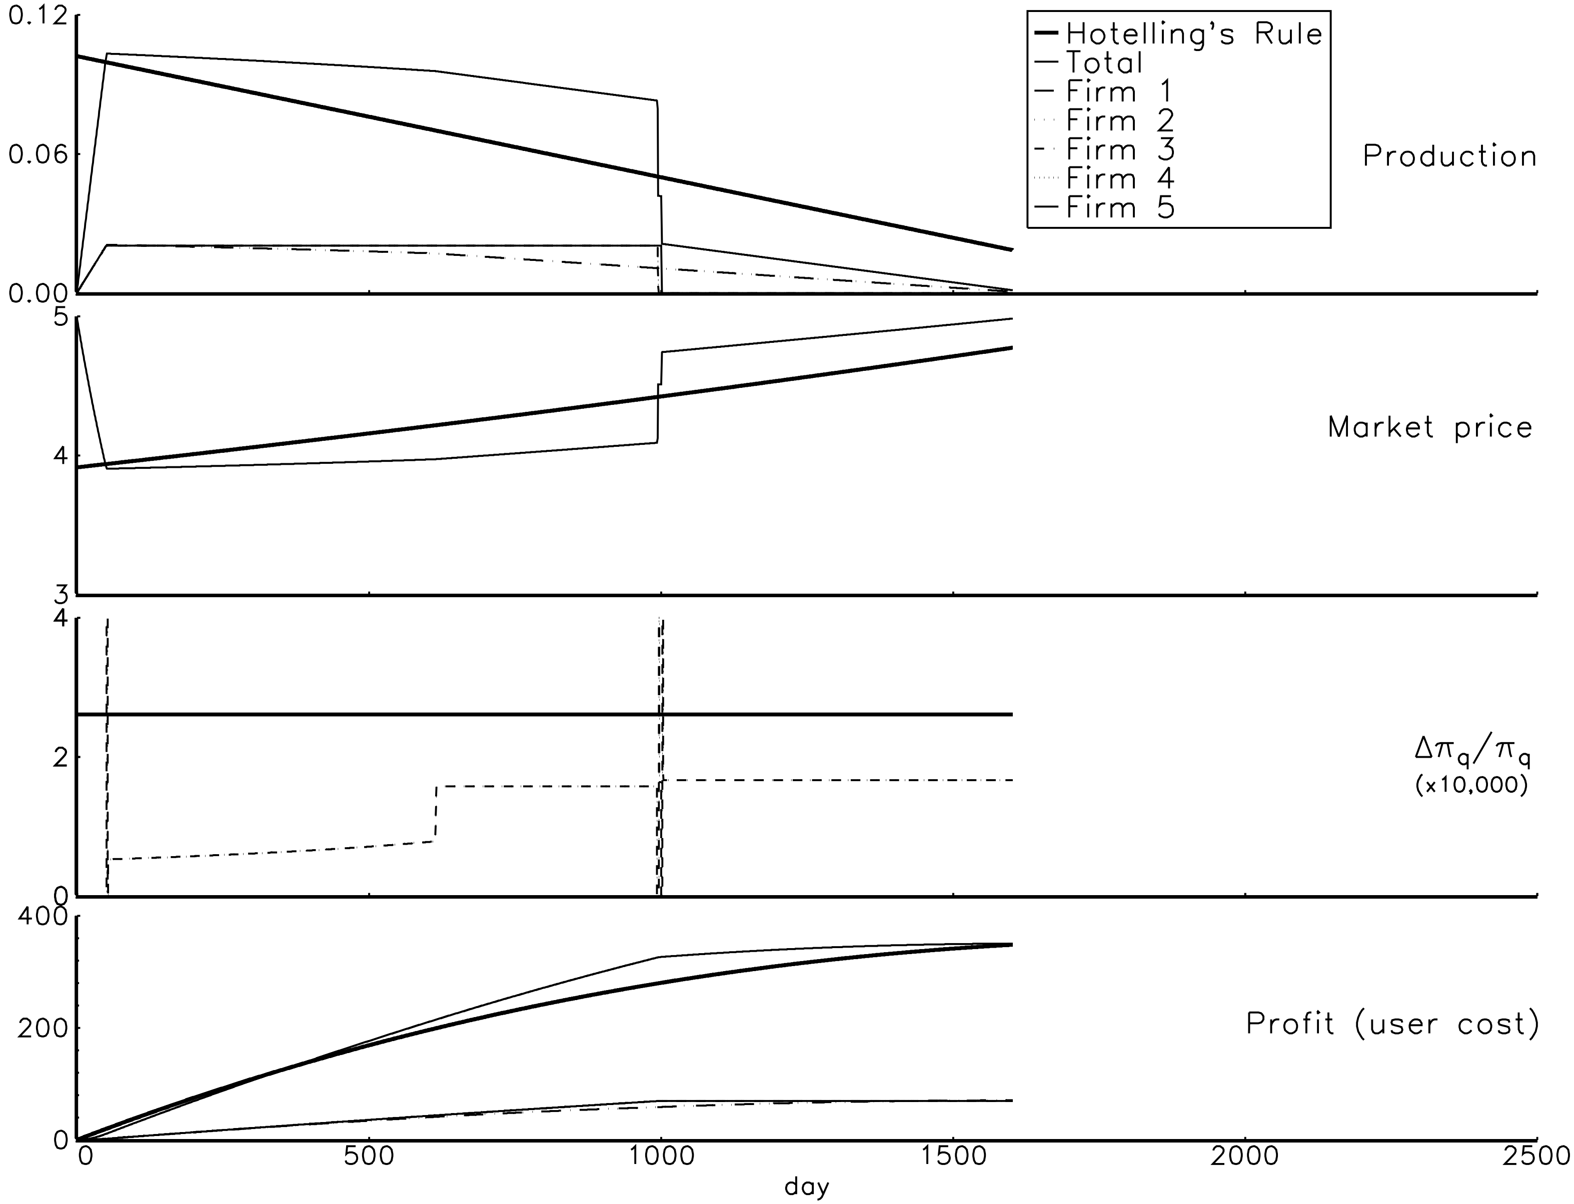

Figure 10: Costless model - 10 producers. See Section 4.1.

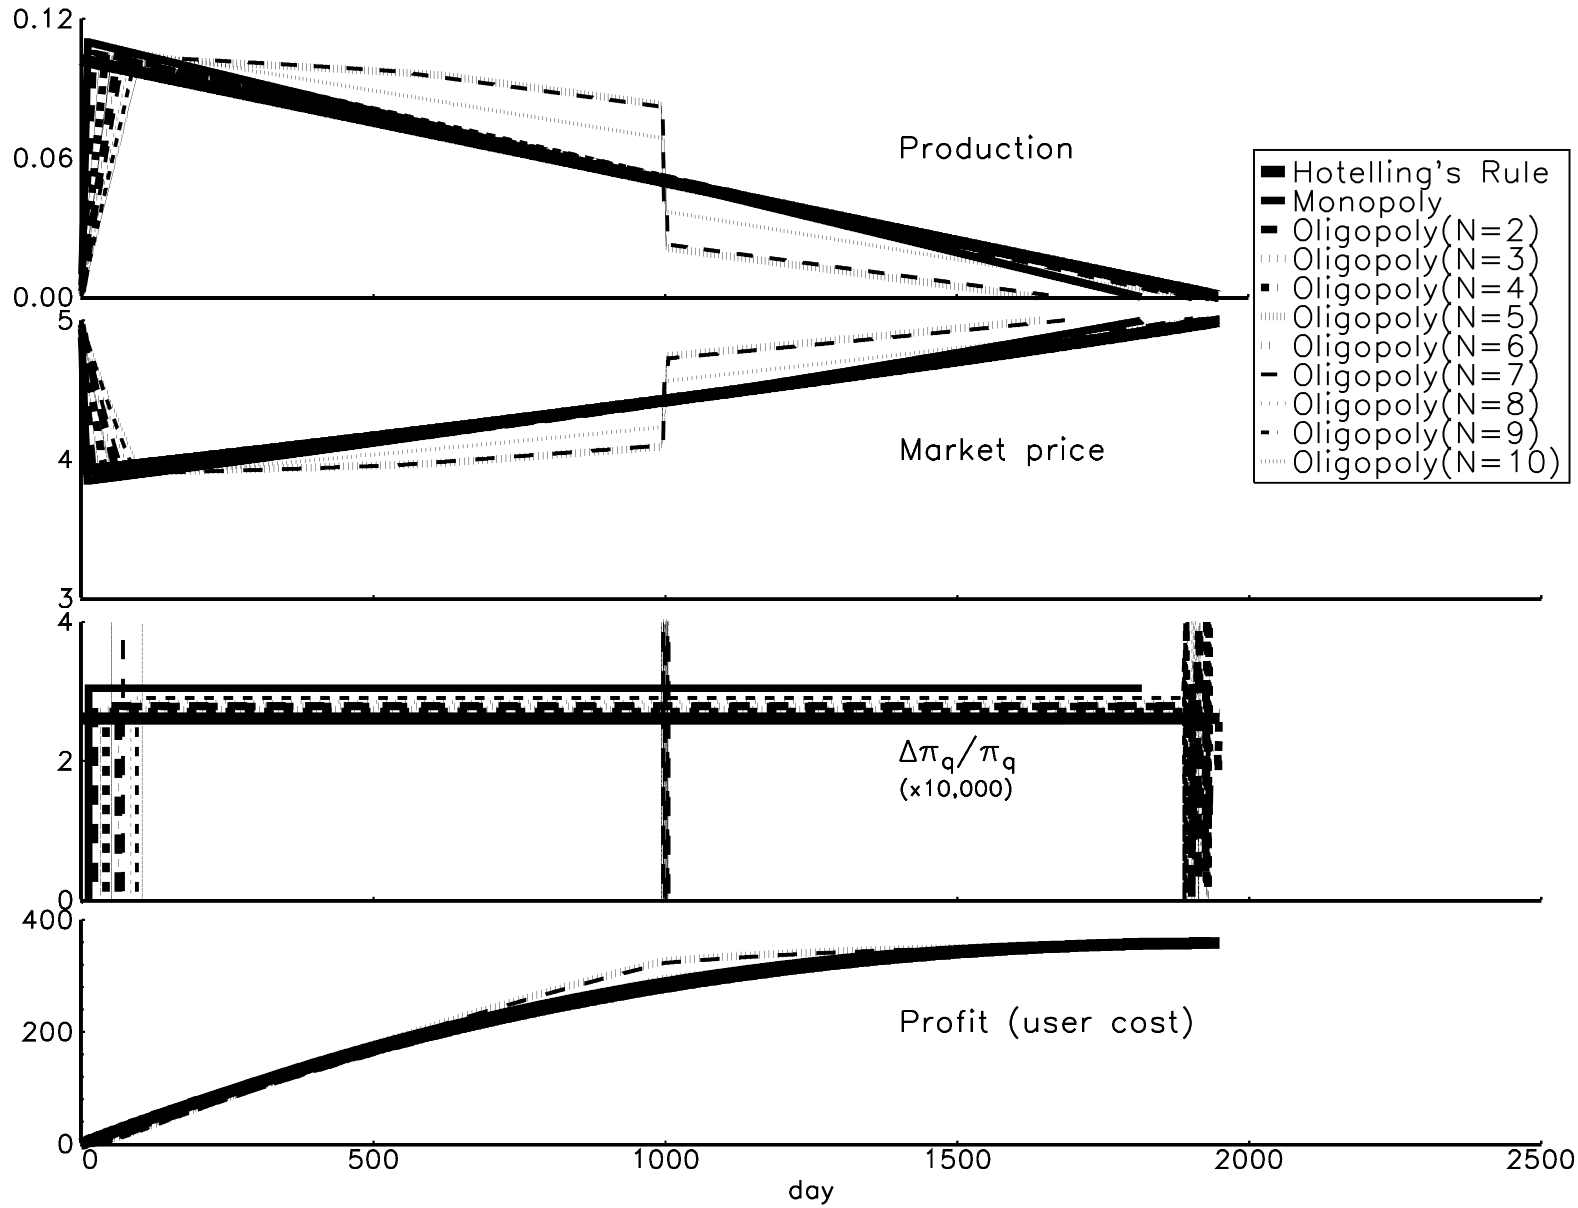

Figure 11: Collusion model - time-series. See the discussion in Section ??.

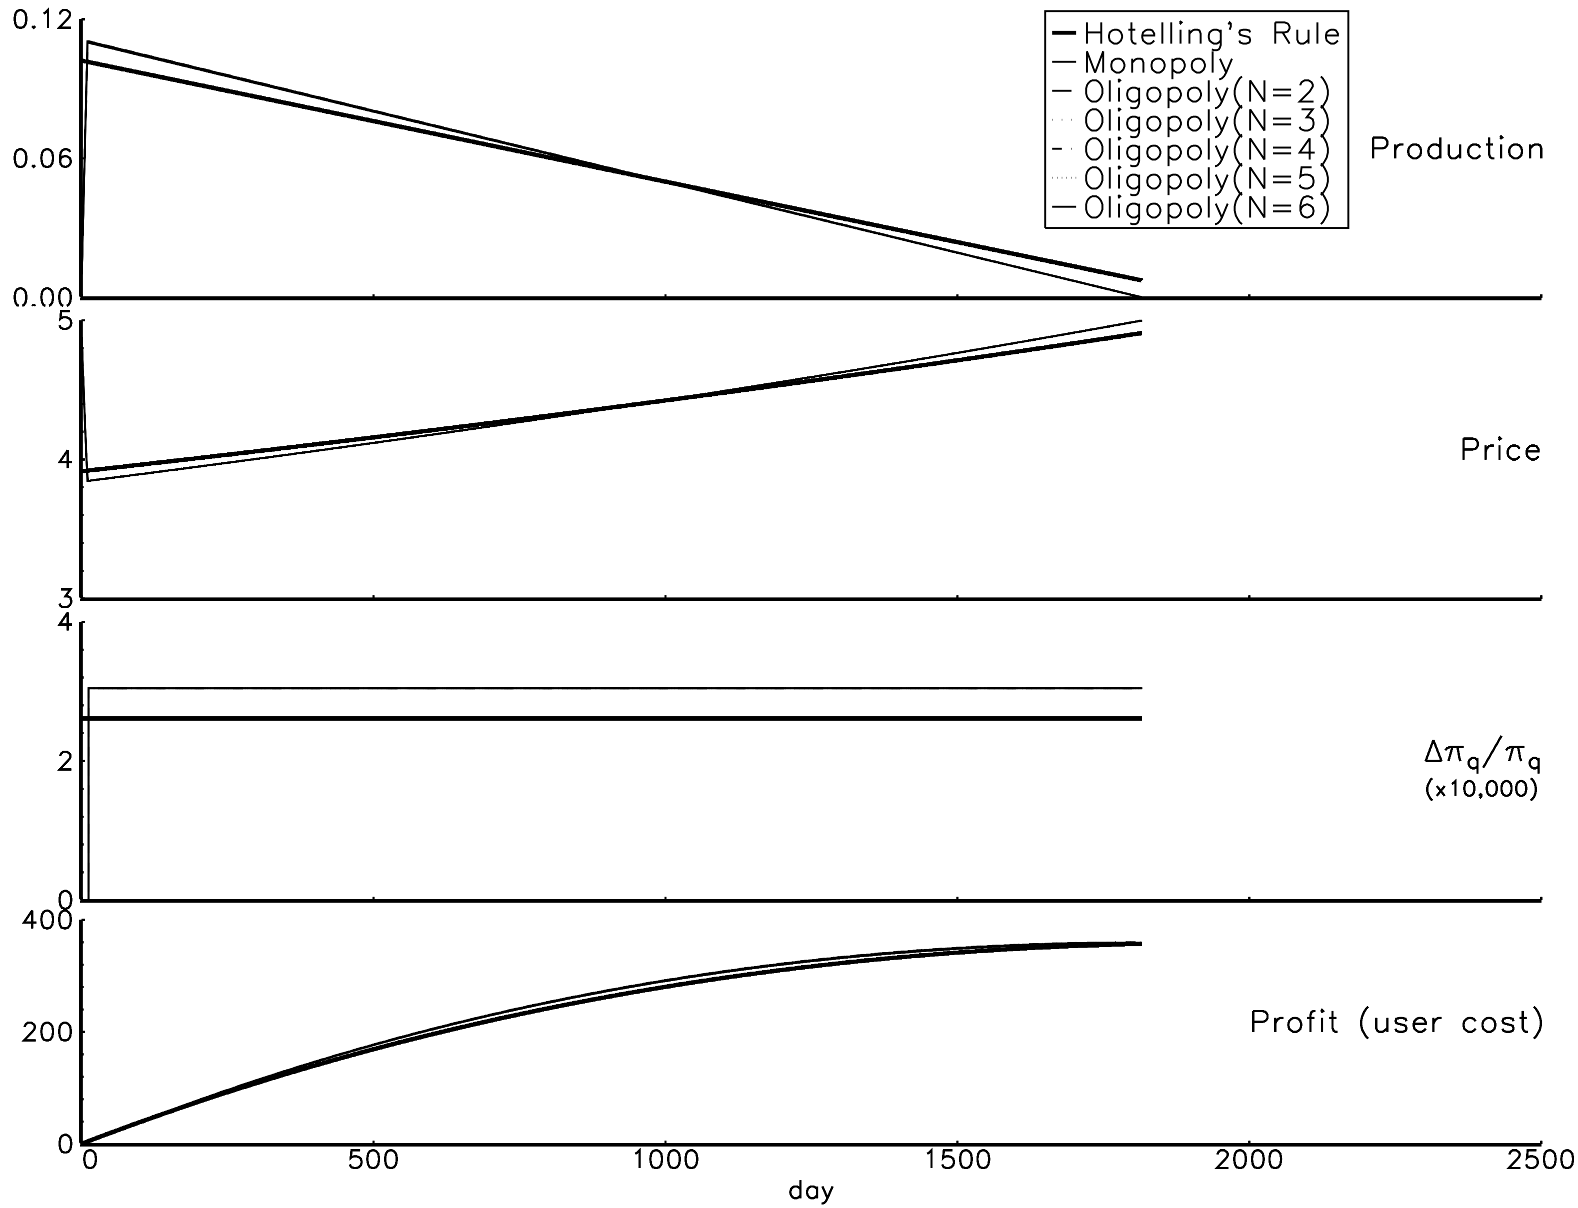

Figure 12: Fixed cost model - time-series. See the discussion in Section 4.3.1.

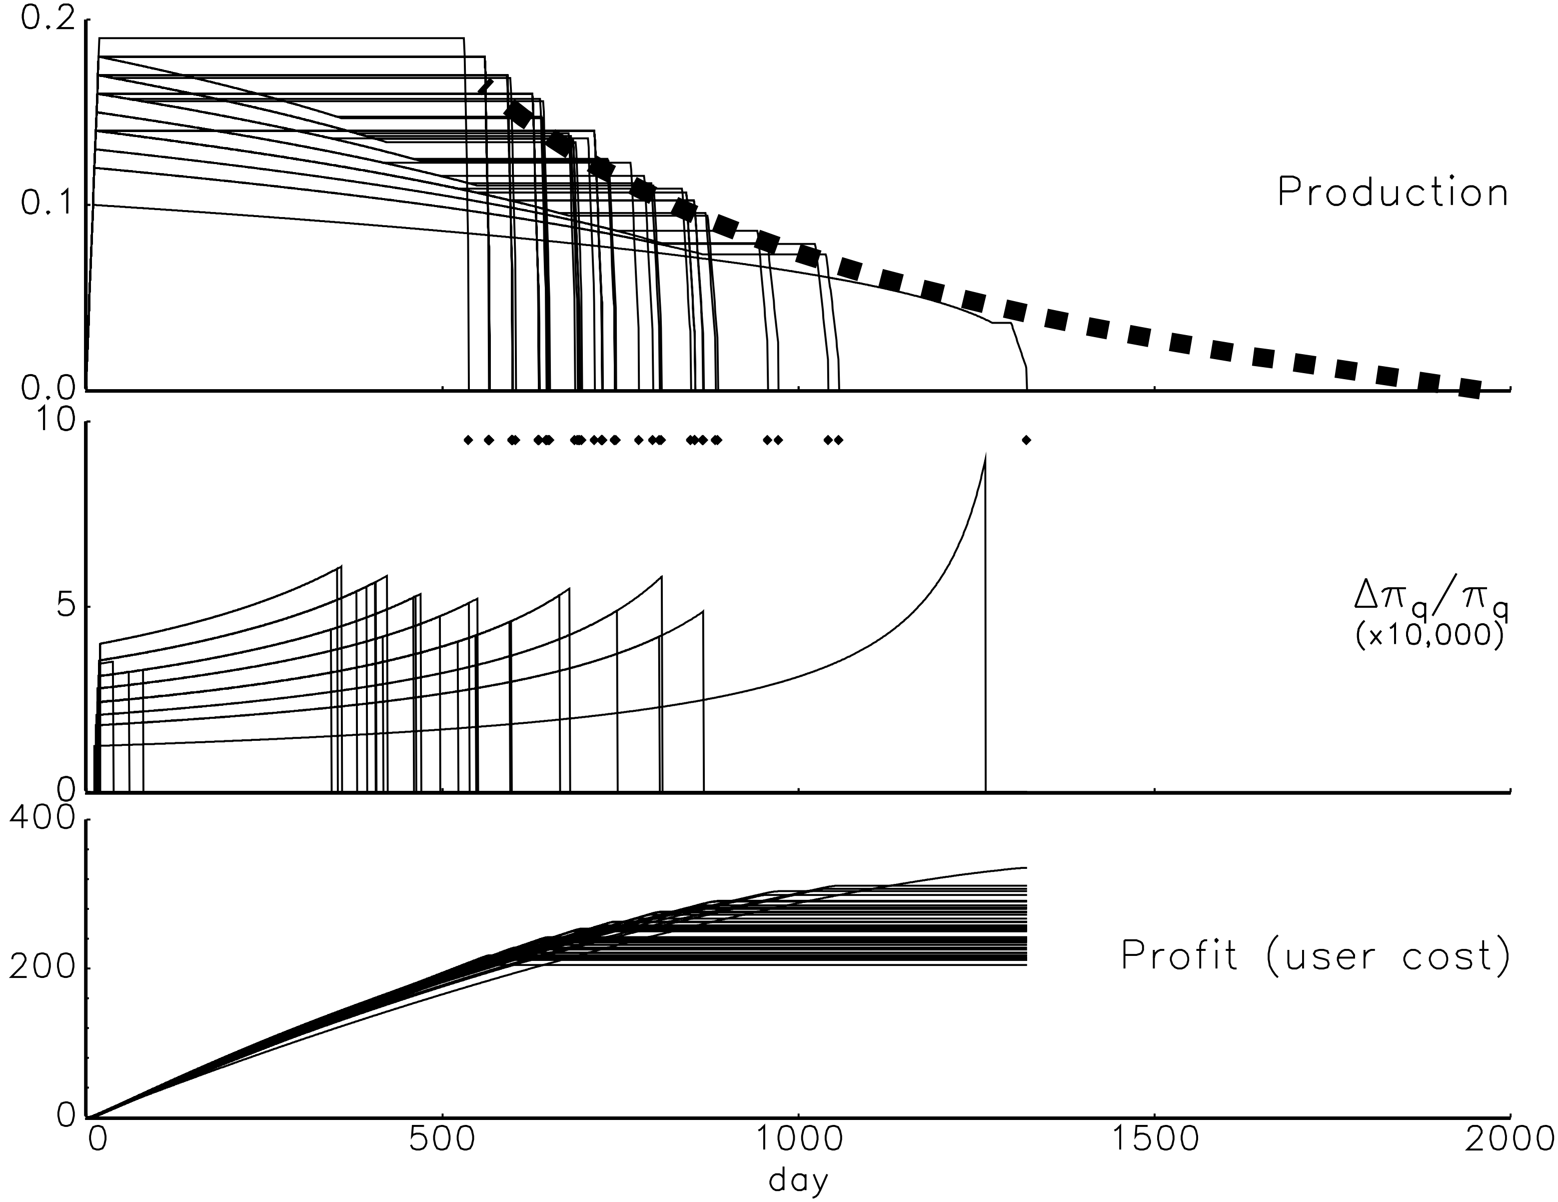

Figure 13: Marginal cost model - time-series. See the discussion in Section 4.3.2.

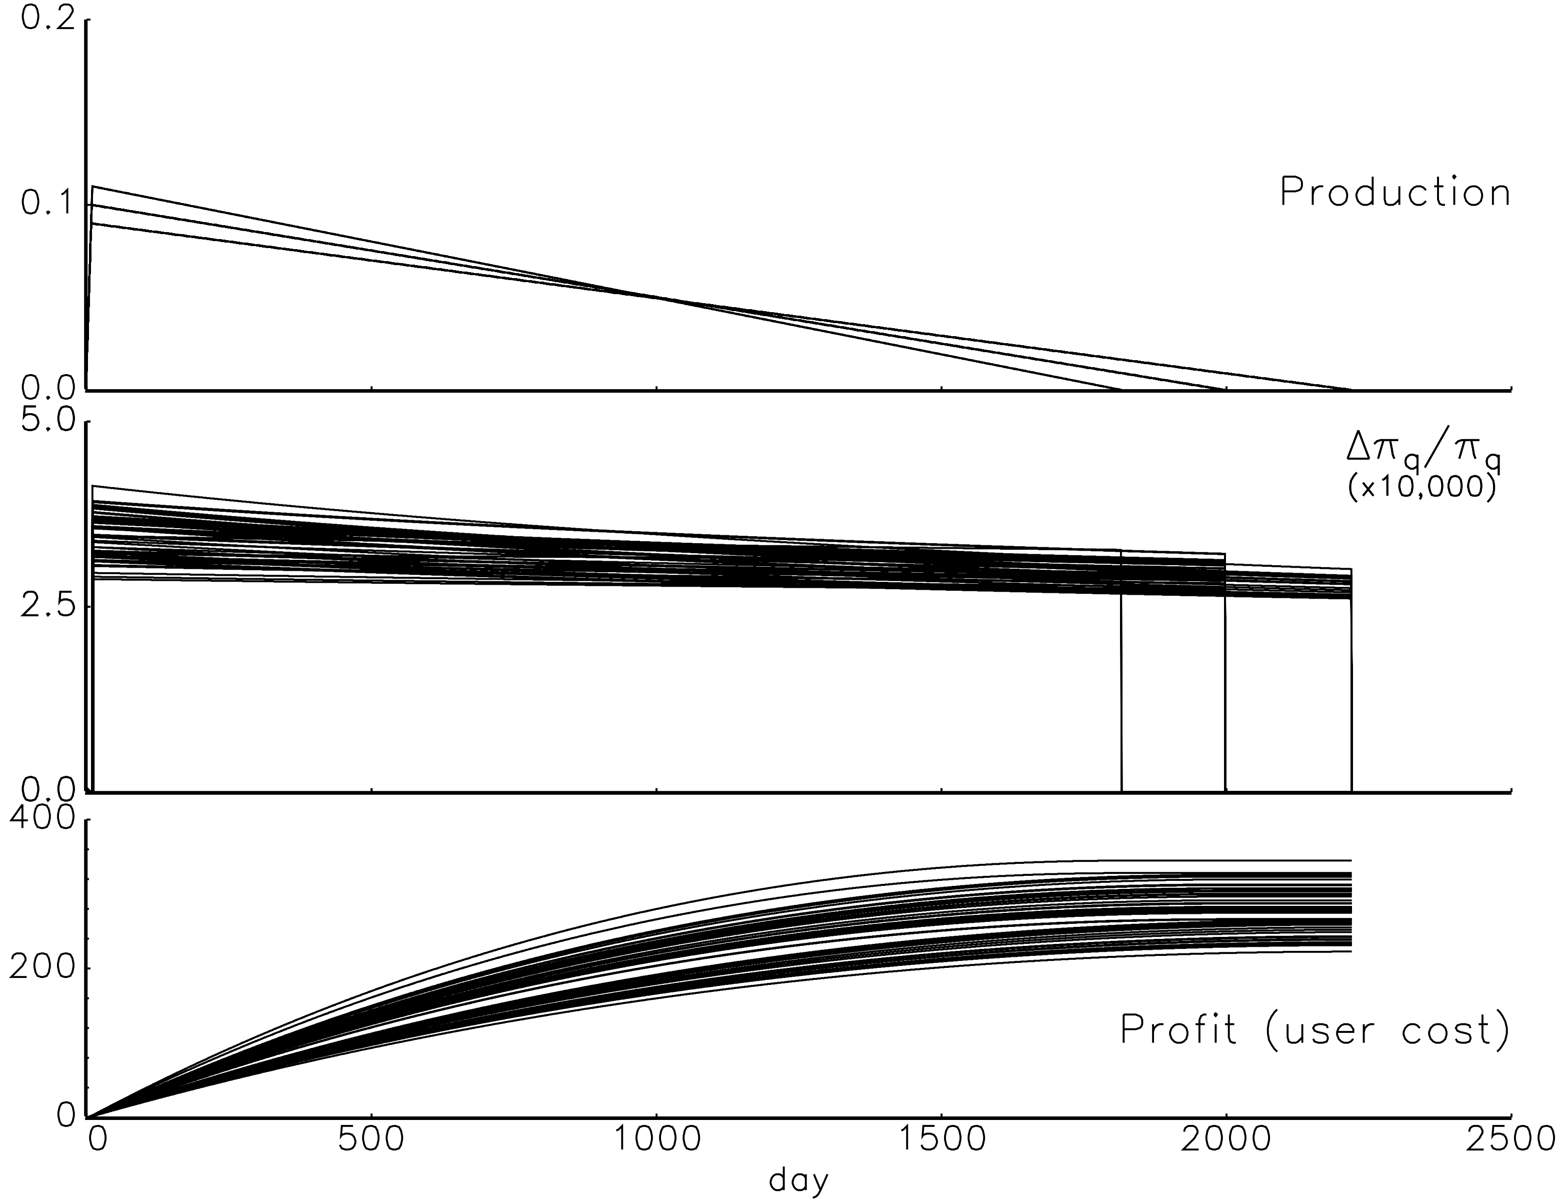

Figure 14: Stock cost model - time-series. See the discussion in Section 4.3.3.

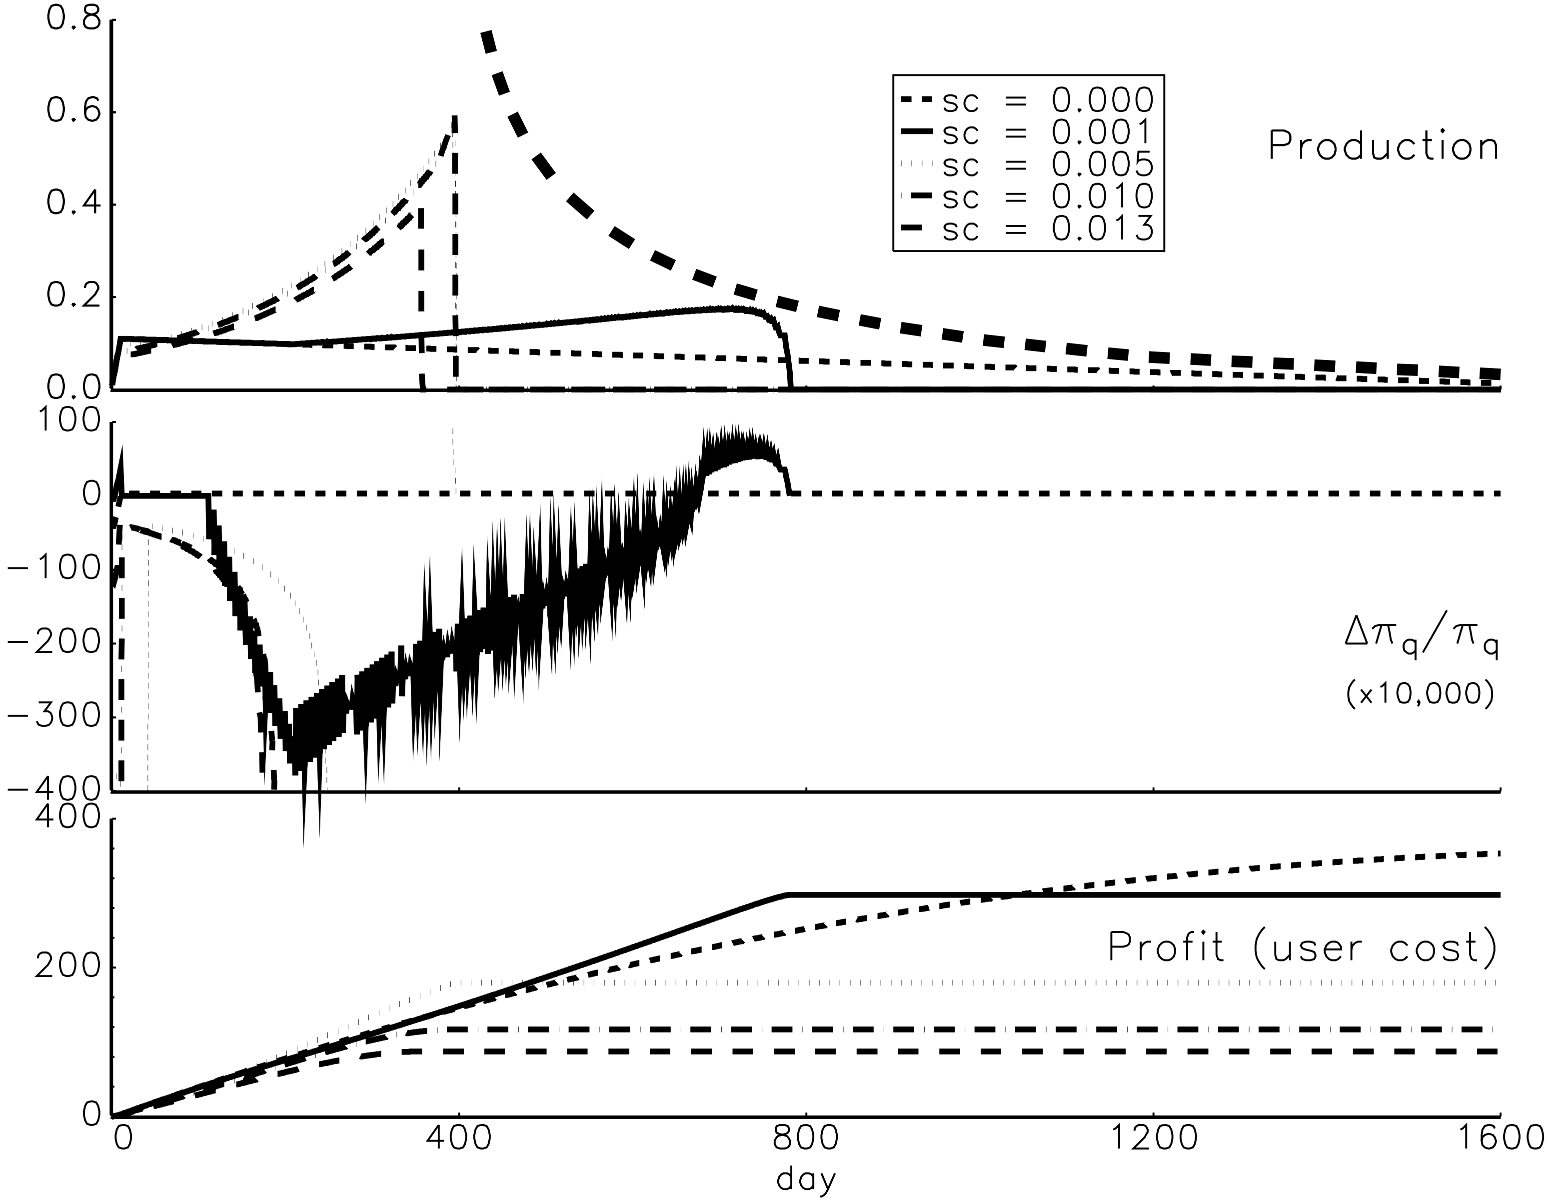

Figure 15: The fixed cost model - outcome space.
In the upper-left plot, the dashed line indicates optimum percent change in marginal profit. The arrows indicate the movement of the percent change in marginal profit during the simulation. The numbers below groups of arrows indicate the day number when the ramp-up phase was completed. The symbols reflect the heuristic value of q_{min} for each sample. The symbols correspond to values of q_{min} which are shown in the legend. The inset graph compares simulation stock lifetimes (dots) with theoretical (dashed line).

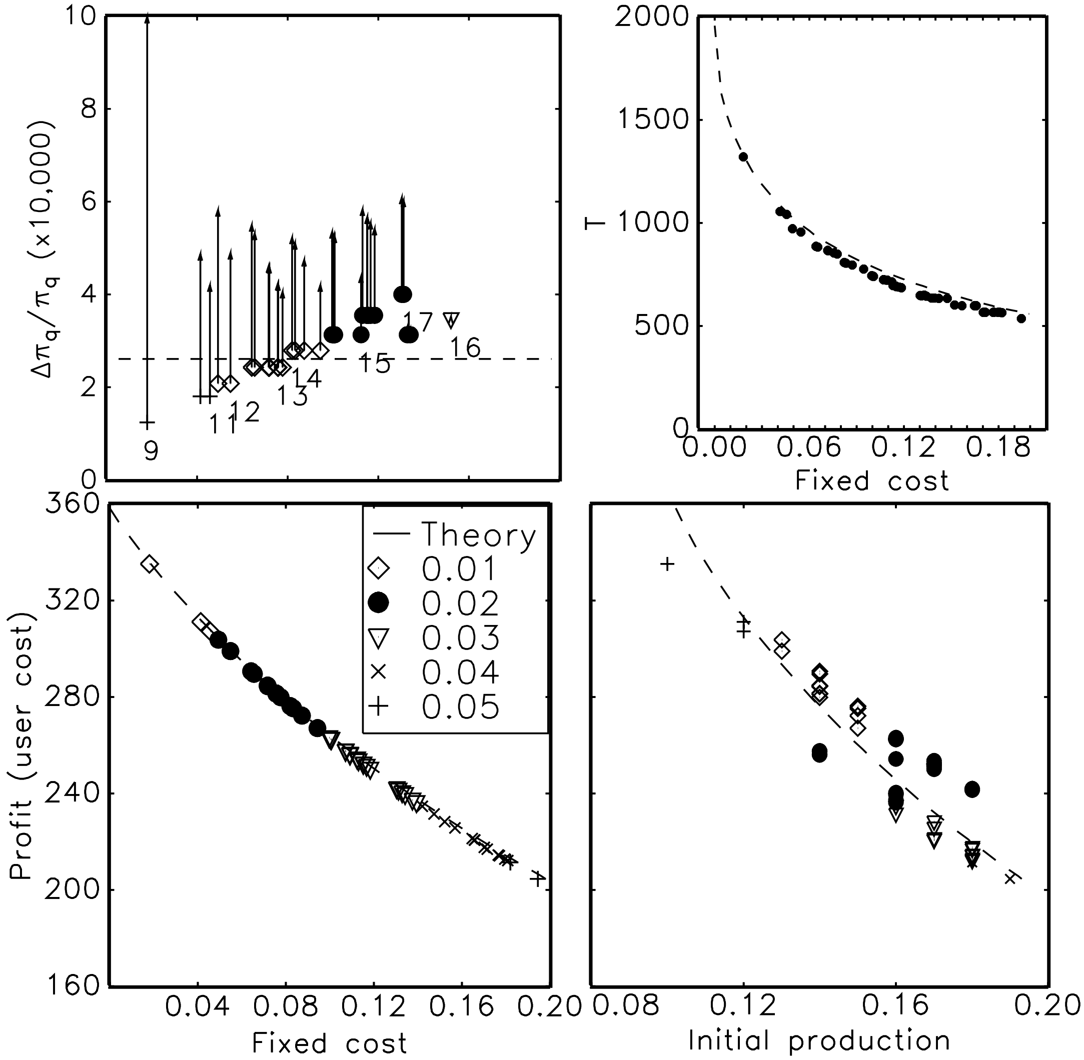

Figure 16: Marginal cost model - outcome space.
Marginal cost models (constant unit cost). The arrows depict the trajectory in the course of the simulation, starting at the X and ending at the arrowhead. The dashed lines indicate the theoretical optima.

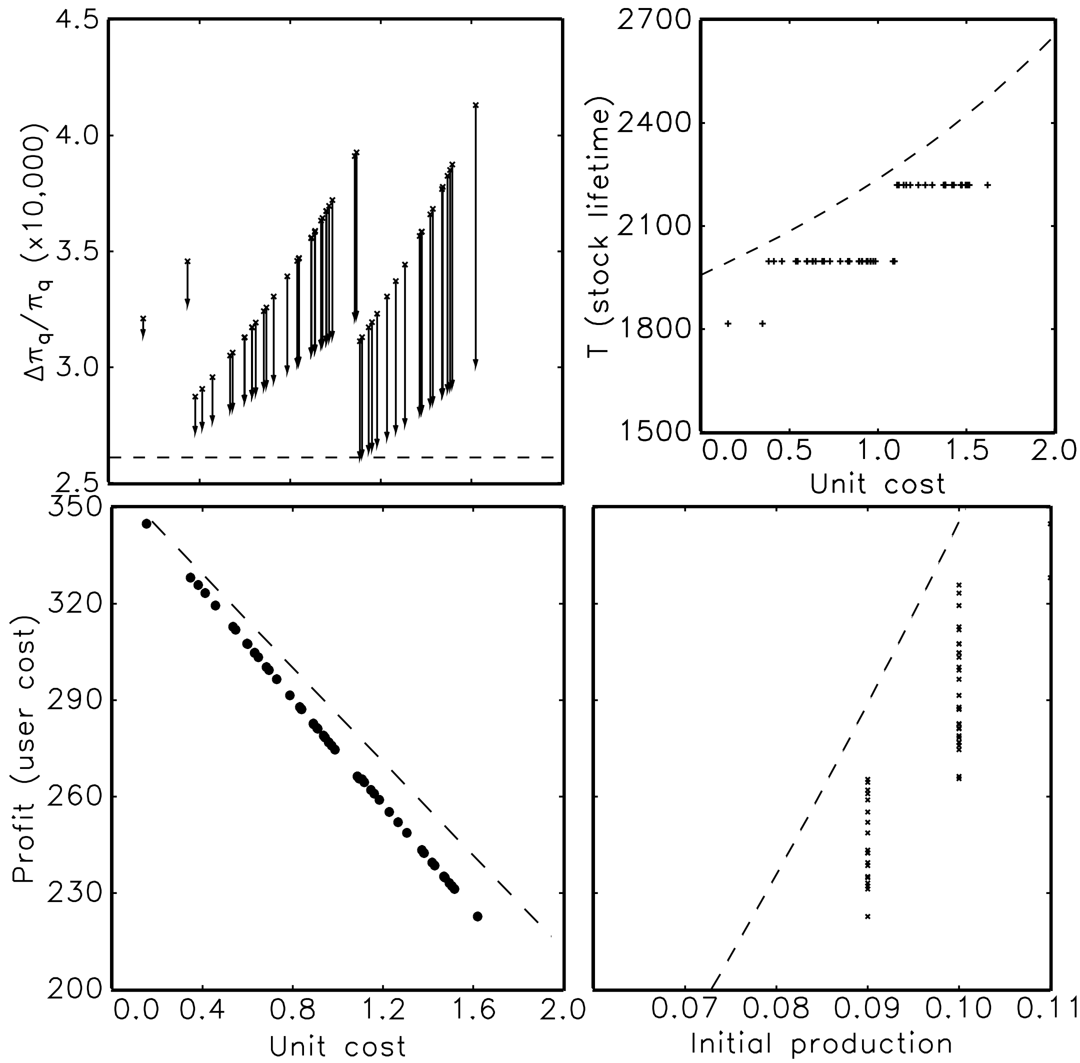

Figure 17: Stock cost model - outcome space.

The arrows depict the percent change in marginal profit trajectory in the course of the simulation, starting at the symbol and ending at the arrowhead. Note that y-values on this graph are 10,000 times those of the others. The inset plots stock lifetime and ending stock level versus marginal cost of stock.

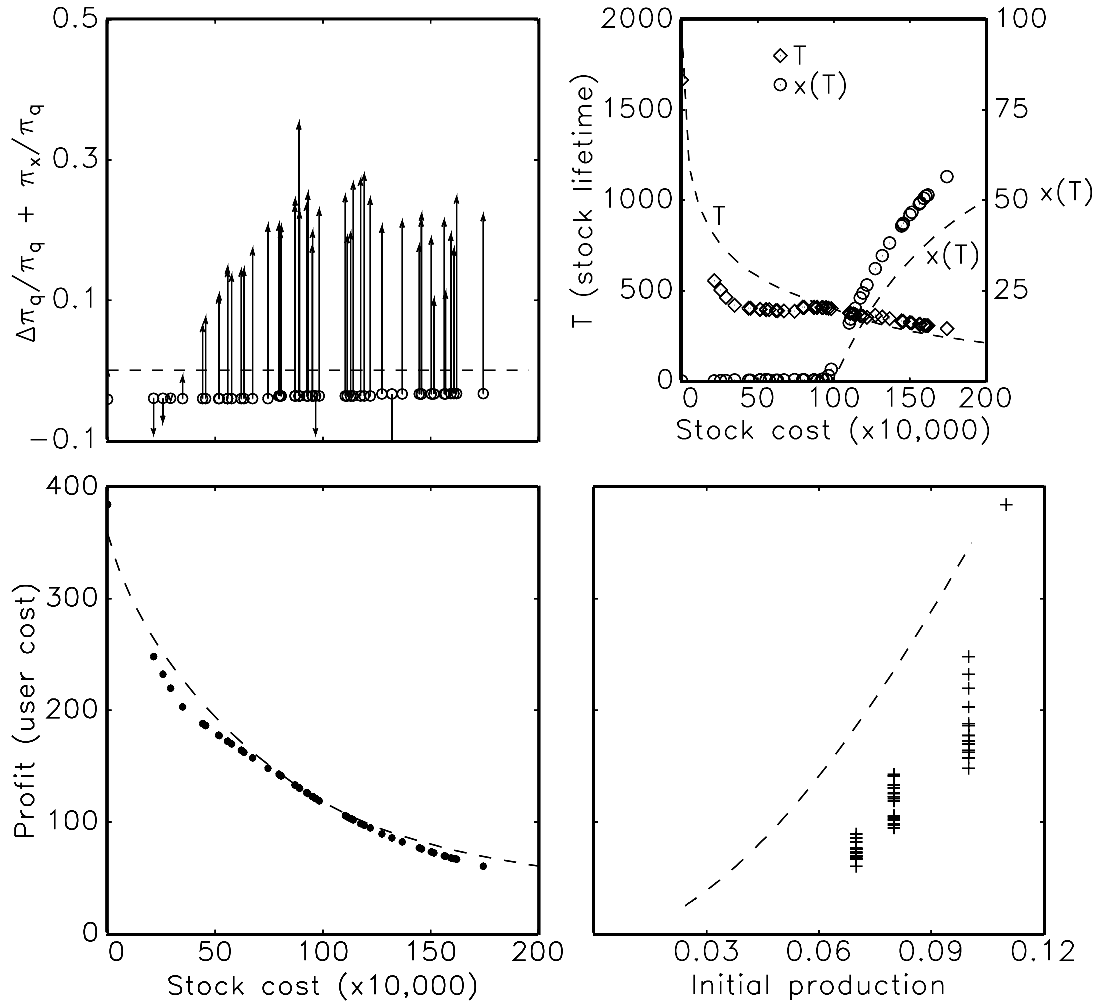

change in marginal profit can vary over the lifetime of the stock, approaching r percent only for a costless monopoly. Percent change in marginal profit is well above r percent for fixed cost, and greater than or less than r percent for unit and stock costs. These results indicate that percent change in marginal profit is not a sensitive metric for optimality. A market with even minor inefficiencies can deviate significantly from the r-percent rule. Conversely, a large deviation from the r-percent rule does not necessarily predict significant inefficiency.

The relative effectiveness of such a simple heuristic implies that producers may be able to profitably exploit natural resources with very little direct knowledge of the demand function or of the market structure. Examined here is profitability, but, in principle, the heuristic should perform similarly for a more general social welfare, for instance. That will be the topic for a forthcoming paper.

References

Claudio Agostini. Estimating market power in the US copper industry. *Review of Industrial Organization*, 28(1):17 – 39, 2006. ISSN 0889938X.

K.J. Arrow and S. Chang. Optimal pricing, use, and exploration of uncertain natural resource stocks. *Journal of Environmental Economics and Management*, 9(1):1–10, 1982. ISSN 0095-0696.

R. Axelrod. The evolution of strategies in the iterated prisoner's dilemma. *Genetic algorithms and simulated annealing*, pages 32–41, 1987.

R.M. Axelrod. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton Univ Pr, 1997.

R. Axtell, R. Axelrod, J.M. Epstein, and M.D. Cohen. Aligning simulation models: A case study and results. *Computational & Mathematical Organization Theory*, 1(2):123–141, 1996. ISSN 1381-298X.

R.L. Axtell, J.M. Epstein, and H.P. Young. The emergence of classes in a multiagent bargaining model. *Social dynamics*, pages 191–211, 1999.

H.J. Barnett and C. Morse. *Scarcity and growth: The economics of natural resource availability*, volume 11. Johns Hopkins University for Resources for the Future, 1963.

Peter Berck and Michael Roberts. Natural resource prices: Will they ever turn up? *Journal of Environmental Economics and Management*, 31(1):65 – 78, 1996. ISSN 0095-0696. doi: DOI: 10.1006/jeem.1996.0032.

G.M. Brown and B.C. Field. Implications of alternative measures of natural resource scarcity. *The Journal of Political Economy*, 86(2):229–243, 1978. ISSN 0022-3808.

R.D. Cairns and N. Van Quyen. Optimal Exploration for and Exploitation of Heterogeneous Mineral Deposits. *Journal of Environmental Economics and Management*, 35(2):164–189, 1998. ISSN 0095-0696.

M.R. Caputo. *Foundations of dynamic economic analysis: optimal control theory and applications*. Cambridge Univ Pr, 2005. ISBN 0521603684.

J.M. Chermak and R.H. Patrick. A Microeconometric Test of the Theory of Exhaustible Resources. *Journal of Environmental Economics and Management*, 42(1):82–103, 2001.

Alpha C. Chiang. *Elements of dynamic optimization*. McGraw-Hill, 1992. ISBN 0070109117.

J.M. Conrad and C.W. Clark. *Natural resource economics: notes and problems*. Cambridge Univ Pr, 1987. ISBN 0521337690.

P. Dasgupta and G.M. Heal. *Economic theory and exhaustible resources.* Cambridge Univ Pr, 1980.

S. Devarajan and A.C. Fisher. Hotelling's "Economics of exhaustible resources": Fifty years later. *Journal of Economic Literature*, 19(1):65–73, 1981.

D.S. Dixon and W.N. Reynolds. The BASP agent-based modeling framework: Applications, scenarios and lessons learned. In *System Sciences, 2003. Proceedings of the 36th Annual Hawaii International Conference on*, pages 8–pp. IEEE, 2003. ISBN 0769518745.

J.M. Epstein. Modeling civil violence: An agent-based computational approach. *Proceedings of the National Academy of Sciences of the United States of America*, 99(Suppl 3): 7243, 2002.

J.M. Epstein and R. Axtell. *Growing artificial societies: social science from the bottom up.* The MIT Press, 1996. ISBN 0262550253.

Scott Farrow. Testing the efficiency of extraction from a stock resource. *The Journal of Political Economy*, 93(3):452–487, 1985. ISSN 00223808.

A.C. Fisher. *Resource and environmental economics.* Cambridge Univ Pr, 1981. ISBN 0521285941.

G. Gaudet and P. Lasserre. On comparing monopoly and competition in exhaustible resource exploitation. *Journal of Environmental Economics and Management*, 15(4): 412–418, 1988. ISSN 0095-0696.

Gérard Gaudet. Natural resource economics under the rule of Hotelling. *Canadian Journal of Economics Revue Canadienne dEconomique*, 40:1033–1059, 2007.

G.N. Gilbert and K.G. Troitzsch. *Simulation for the social scientist.* Open Univ Pr, 2005.

R. Halvorsen and T.R. Smith. On measuring natural resource scarcity. *The Journal of Political Economy*, 92(5):954–964, 1984. ISSN 0022-3808.

R. Halvorsen and T.R. Smith. A test of the theory of exhaustible resources. *The Quarterly Journal of Economics*, 106(1):123, 1991. ISSN 0033-5533.

N. Hanley, J.F. Shogren, B. White, et al. *Environmental economics in theory and practice.* Macmillan, 1997. ISBN 0333582357.

J.E. Hartley. Retrospectives: The origins of the representative agent. *The Journal of Economic Perspectives*, 10(2):169–177, 1996. ISSN 0895-3309.

G. Heal. The relationship between price and extraction cost for a resource with a backstop technology. *The Bell Journal of Economics*, 7(2):371–378, 1976. ISSN 0361-915X.

G. Heal and M. Barrow. The relationship between interest rates and metal price movements. *The Review of Economic Studies*, 47(1):161–181, 1980. ISSN 0034-6527.

Harold Hotelling. The economics of exhaustible resources. *Journal of Political Economy*, 39(2):137–75, 1931.

C.T. Hsiao and L.C. Chang. Dynamic optimal groundwater management with inclusion of fixed costs. *Journal of water resources planning and management*, 128:57, 2002.

M.I. Kamien and N.L. Schwartz. *Dynamic optimization: the calculus of variations and optimal control in economics and management*. North Holland Amsterdam, 1981.

Jeffrey A. Krautkraemer. Nonrenewable resource scarcity. *Journal of Economic Literature*, 36(4):2065 – 2107, 1998. ISSN 00220515.

R. Lecomber. *Economics of natural resources*. MacMillan, London, 1979.

David Levhari and Nissan Liviatan. Notes on Hotelling's economics of exhaustible resources. *The Canadian Journal of Economics / Revue canadienne d'Economique*, 10(2): 177–192, 1977. ISSN 00084085.

T.R. Lewis, S.A. Matthews, and H.S. Burness. Monopoly and the rate of extraction of exhaustible resources: Note. *The American Economic Review*, 69(1):227–230, 1979. ISSN 0002-8282.

J. Livernois. On the empirical significance of the Hotelling rule. *Review of Environmental Economics and policy*, 3(1):22, 2009. ISSN 1750-6816.

J. Livernois and P. Martin. Price, scarcity rent, and a modified r per cent rule for nonrenewable resources. *Canadian Journal of Economics/Revue canadienne d'économique*, 34(3):827–845, 2001. ISSN 1540-5982.

J.R. Livernois and R.S. Uhler. Extraction costs and the economics of nonrenewable resources. *The Journal of Political Economy*, 95(1):195–203, 1987. ISSN 0022-3808.

D. McFadzean, D. Stewart, and L. Tesfatsion. A computational laboratory for evolutionary trade networks. *Evolutionary Computation, IEEE Transactions on*, 5(5):546–560, 2001. ISSN 1089-778X.

W.D. Nordhaus. The allocation of energy resources. *Brookings Papers on Economic Activity*, 1973(3):529–576, 1973.

W.D. Nordhaus. Resources as a Constraint on Growth. *The American Economic Review*, 64(2):22–26, 1974. ISSN 0002-8282.

M.H. Pesaran. An econometric analysis of exploration and extraction of oil in the UK Continental Shelf. *The Economic Journal*, 100(401):367–390, 1990. ISSN 0013-0133.

R.S. Pindyck. The optimal exploration and production of nonrenewable resources. *The Journal of Political Economy*, 86(5):841–861, 1978. ISSN 0022-3808.

R.S. Pindyck. On monopoly power in extractive resource markets. *Journal of Environmental Economics and Management*, 14(2):128–142, 1987. ISSN 0095-0696.

LS Pontryagin. Optimal control processes. *Uspekhi Mat. Nauk*, 14(1):85, 1959.

LS Pontryagin, VG Boltyanskiı, RV Gamkrelidze, and EF Mishchenko. The mathematical theory of optimal processes, Fizmatgiz, Moscow 1961. *English transl., Interscience, New York*, 1962.

W.N. Reynolds and DS Dixon. Archimedes: A Prototype Distillation. *Maneuver Warfare Science*, 2001.

J. Rust, J. Miller, and R. Palmer. Behavior of trading automata in a computerized double auction market. *The double auction market: Institutions, theories, and evidence*, pages 155–198, 1992.

S.W. Salant. Exhaustible resources and industrial structure: A Nash-Cournot approach to the world oil market. *The Journal of Political Economy*, 84(5):1079–1093, 1976. ISSN 0022-3808.

T.C. Schelling. *Micromotives and Macrobehavior*. New York: Norton, 1978.

A.D. Scott. The theory of the mine under conditions of certainty. *Extractive Resources and Taxation*, pages 25–62, 1967.

M.E. Slade. Trends in natural-resource commodity prices: an analysis of the time domain. *Journal of Environmental Economics and Management*, 9(2):122–137, 1982. ISSN 0095-0696.

M.E. Slade. Tax policy and the supply of exhaustible resources: Theory and practice. *Land Economics*, 60(2):133–147, 1984. ISSN 0023-7639.

V.K. Smith. Natural resource scarcity: a statistical analysis. *The Review of Economics and Statistics*, 61(3):423–427, 1979. ISSN 0034-6535.

Robert M. Solow. The economics of resources or the resources of economics. *The American Economic Review*, 64(2):1–14, 1974. ISSN 00028282.

J.E. Stiglitz. Monopoly and the rate of extraction of exhaustible resources. *The American Economic Review*, 66(4):655–661, 1976. ISSN 0002-8282.

Kenneth R. Stollery. Mineral depletion with cost as the extraction limit: A model applied to the behavior of prices in the nickel industry. *Journal of Environmental Economics and Management*, 10(2):151 – 165, 1983. ISSN 0095-0696.

J.E. Swierzbinski and R. Mendelsohn. Exploration and exhaustible resources: the micro-foundations of aggregate models. *International Economic Review*, 30(1):175–186, 1989. ISSN 0020-6598.

L. Tesfatsion. Structure, behavior, and market power in an evolutionary labor market with adaptive search. *Journal of Economic Dynamics and Control*, 25(3-4):419–457, 2001. ISSN 0165-1889.

L. Tesfatsion. Agent-based computational economics: A constructive approach to economic theory. *Handbook of computational economics*, 2:831–880, 2006. ISSN 1574-0021.

T.H. Tietenberg and L. Lewis. *Environmental and natural resource economics*. Addison-Wesley, 2000. ISBN 0321031288.

H.R. Varian. Microeconomic analysis. *3rd. ed. ed. New York, London: WW Norton & Company*, 1992.

D. Young. Cost specification and firm behaviour in a Hotelling model of resource extraction. *Canadian Journal of Economics*, 25(1):41–59, 1992.