

8-31-2011

Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule

David S. Dixon

Follow this and additional works at: https://digitalrepository.unm.edu/econ_etds

Recommended Citation

Dixon, David S.. "Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule." (2011). https://digitalrepository.unm.edu/econ_etds/4

This Dissertation is brought to you for free and open access by the Electronic Theses and Dissertations at UNM Digital Repository. It has been accepted for inclusion in Economics ETDs by an authorized administrator of UNM Digital Repository. For more information, please contact disc@unm.edu.

David S. Dixon

Candidate

Department of Economics

Department

This dissertation is approved, and it is acceptable in quality and form for publication:

Approved by the Dissertation Committee:



, Co-Chairperson

Kustine Swinsud

, Co-Chairperson



Budley T. Collins

Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule

by

David S. Dixon

B.A., Physics, University of Vermont, 1985

M.S., Physics, University of New Mexico, 1993

DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Economics

The University of New Mexico
Albuquerque, New Mexico

July, 2011

Dedication

This is dedicated to my patient, supportive, and always cherished wife, daughter and son.

Acknowledgments

I owe a great debt to my graduate advisor and committee co-chair Janie Chermak, whose influence began long before there was a dissertation, and whose expertise and sage advice made it possible to reach any conclusion. Also a tremendous influence both before and during this effort was my committee co-chair, Kristine Grimsrud, whose devotion to rigor continues to inspire me. I thank Kate Krause for introducing me to experimental economics and game theory in that context, and for serving on my committee. Also, thanks to committee member Brad Cullen, who was patient with the arcane aspects of this approach. Thanks also to Jenn Thacher for introducing me to the topic of Hotelling's Rule. I want to thank Gérard Gaudet for helping me understand the origins of the term "Hotelling's Rule". And especially, I wish to acknowledge the graciousness of Robert Solow in responding to my questions about his role in all this.

Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule

by

David S. Dixon

ABSTRACT OF DISSERTATION

Submitted in Partial Fulfillment of the
Requirements for the Degree of

Doctor of Philosophy
Economics

The University of New Mexico
Albuquerque, New Mexico

July, 2011

Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule

by

David S. Dixon

B.A., Physics, University of Vermont, 1985

M.S., Physics, University of New Mexico, 1993

Ph.D., Economics, University of New Mexico, May 2011

Abstract

This dissertation presents a prototype model and methodology for validating a simple agent-based model against the basic Hotelling monopoly model and a few basic extensions. Hotelling's Rule identifies the expected behavior of a market for a nonrenewable resource. The statement is simple - marginal profit will increase at the prevailing rate of interest - but the implications are far-reaching and not broadly understood. Agent-based modeling is a computer modeling and simulation methodology. It has its origins in biology and physics, but has become a powerful tool in the social sciences for examining systems in which the well-understood behaviors of individuals result in unanticipated outcomes.

Validation of the basic Hotelling monopoly model is a necessary step in wider acceptance of agent-based modeling as a predictive and analytical tool in natural resource economics. An agent-based model is a valid predictive tool if, given rules to express preferences, it is possible to predict the large-scale outcomes of the choices made by individuals over time. An

agent-based model is a valid analytical tool if it provides a means to explore the behaviors that lead to known results, much like nonlinear regression.

This simple agent-based model is found to be valid for the basic Hotelling monopoly model. The agent-based model is validated with caveats for the Hotelling monopoly model with extensions to include basic production technologies. The caveats are based on small deviations, the magnitudes of which depend on the specific form of costs associated with a production technology. It is argued that those deviations are not unlike the deviations that a human would make.

An extension of the Hotelling monopoly model to a small oligopoly exhibits emergent cooperation-like properties, despite the absence of explicit interagent communication. Depending on the number of producers and the initial distribution of resource stocks, the behavior is either collusion-like or Cournot-like. The Cournot-like outcome occurs when only some of the producers lower production, resulting in a rise in the market price, which causes the other producers to experience a Hotelling's Rule increase in marginal profit without reducing their own production. This continues at each time step, so that the latter producers maintain a constant, higher production level while the others continue to decrease production.

The outcomes of the agent-based models are reassessed with the costs previously arising from production technologies replaced by taxes associated with fiscal policies. Each fiscal regime is evaluated in terms of its efficacy and the unintended consequences of the policy. The policy goals examined are preservation for future generations, internalization of externalities (Pigouvian taxation), and revenue generation. This comparison uses data from the agent-based models, examining agent error as one of the unintended consequences of fiscal policy.

The basics of agent-based modeling are presented. This methodology is suited to problems in which the immediate preferences of the agents can be stated as equations or rules but complexity in their interactions with the environment or each other make predicting the outcomes difficult.

Contents

List of Figures	xiv
List of Tables	xvi
Glossary	xvii
1 Introduction	1
1.1 The inverse demand function and its parameters	4
1.2 A simple ABM and Hotelling's outcome	4
1.3 The ABM and production technologies	7
1.3.1 Lease and capital costs	8
1.3.2 Marginal costs	9
1.3.3 Stock and cumulative production costs	9
1.4 The ABM in an oligopoly market	11
1.5 The ABM and the consequences of taxation	13
1.6 Chapter organization	14

2	Background	16
2.1	The evolving theory of nonrenewable resources	18
2.2	Studies of the markets for nonrenewable resources	22
2.3	The basics of agent-based modeling	24
2.4	The MASON agent-based modeling platform	29
3	Theory	32
3.1	Basic theory	33
3.1.1	A benefit-cost treatment	37
3.2	The Hotelling monopoly demand function	37
3.3	A costless monopoly model	39
3.4	A fixed cost model	43
3.5	A marginal cost model	47
3.6	A stock cost model	51
3.7	Oligopoly models	58
3.8	Summary	59
4	The Hotelling Agent-based Models	64
4.1	ABM architecture	64
4.2	Simulation architecture	65
4.2.1	The market agent	67
4.2.2	The producer agent	67

4.3	A simple optimization heuristic	68
4.3.1	The use of discrete summations	69
4.3.2	The heuristic algorithm	70
4.3.3	The inverse demand function and its parameters	73
4.4	Monopoly models	73
4.5	The computational agent model	74
4.6	The adaptive agent models	75
4.7	Production technologies with nonzero cost	76
4.7.1	Fixed cost model	76
4.7.2	Marginal cost model	77
4.7.3	Stock cost model	77
4.8	Other sources of uncertainty	78
5	The Hotelling Simulations	81
5.0.1	The inverse demand function and its parameters	82
5.1	The computational agent models	82
5.2	The adaptive agent models	84
5.2.1	A pooled-stock oligopoly model	86
5.2.2	Accuracy and precision of the heuristic	87
5.3	A collusion model	89
5.4	Nonzero-cost models	89

5.4.1	Fixed cost model	90
5.4.2	Marginal cost model	91
5.4.3	Stock cost model	92
5.5	Theoretical efficiency	93
6	Analysis of the Simulation Results	106
6.1	The efficacy of the heuristic	107
6.2	The penalty for mistaking the number of competitors	109
6.3	The emergent oligopoly models	110
6.4	Overt collusion is indistinguishable from optimization	112
6.5	The nonzero cost models	113
6.5.1	The fixed cost model	113
6.5.2	The marginal cost model	114
6.5.3	The stock cost model	114
6.6	Summary	117
7	Policy Implications	118
7.1	Policy objectives	119
7.1.1	Policies to forestall depletion	121
7.1.2	Policies to internalize externalities	122
7.1.3	Taxes as revenue	123
7.2	The agent-based models in the policy context	123

7.2.1	The inverse demand function and its parameters	124
7.2.2	Franchise taxes	124
7.2.3	Unit severance taxes	125
7.2.4	Ad valorem taxes	126
7.2.5	Royalties	126
7.2.6	Cumulative production fees/bonds/taxes	127
7.3	The efficiency of fiscal regimes	128
8	Conclusions and Future Directions	135
8.1	The rationale for agent-based modeling	136
8.2	Reaching the objectives	136
8.2.1	The ABM framework	137
8.2.2	Validating the basic Hotelling ABM	138
8.2.3	Exploring emergence	139
8.3	Future directions	140
	Bibliography	141
A	GAUSS procedure to solve for q_{max} in the marginal cost model	151
B	GAUSS procedure for numerical solutions to marginal cost model	154
C	GAUSS procedure for numerical solutions to the stock cost model	159

D The MASON Hotelling package

List of Figures

1.1	An example of the initial production level error.	6
3.1	Fixed cost model - numerical solutions for terminal production level q_T . . .	45
3.2	The marginal cost model - theoretical values for minimum and maximum production for selected marginal costs.	61
3.3	Marginal cost model - numerical solutions for terminal time T	62
3.4	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.	63
4.1	The details of a time-step	80
5.1	Monte Carlo samples of heuristic outcomes.	88
5.2	Errors in the computational agent model.	94
5.3	Computational agent models - time-series.	95
5.4	Adaptive identical stock models - time-series.	96
5.5	Adaptive agent pooled stock models - time-series.	97
5.6	Adaptive agent conserved stock models - 10 producers.	98

5.7	Collusion model - time-series.	99
5.8	Fixed cost model - time-series.	100
5.9	Marginal cost model - time-series.	101
5.10	Stock cost model - time-series.	102
5.11	The fixed cost model - outcome space.	103
5.12	Marginal cost model - outcome space.	104
5.13	Stock cost model - outcome space.	105
7.1	Franchise taxes.	130
7.2	Severance taxes.	131
7.3	Royalties.	132
7.4	Cumulative taxation.	133
7.5	Fiscal policy efficiency and efficacy	134
D.1	Hotelling's Rule Model in the MASON	167
D.2	Hotelling's Rule Model control panel.	168
D.3	An oligopoly market with five firms.	169
D.4	A running simulation.	170

List of Tables

3.1	Production path comparison.	59
5.1	Comparison of the computational agent models and the adaptive agent models ($\Delta\pi_q/\pi_q$ values are multiplied by 10,000).	85
7.1	Natural resource taxes as revenue: states for which natural resource taxes constitute more than ten percent of total tax revenue.	123

Glossary

Abbreviations and specialized terms

ABM	<i>Agent-based model</i> or <i>agent-based modeling</i> . Agent-based modeling is a technique for making models - primarily computer models - in which the individual agent (e.g. people, firms, countries) are represented in terms of their state (<i>wealth</i> , for example) and behaviors (<i>buy more stuff</i> , for example). The initials ABM are also used in reference to an agent-based model. The ambiguity in usage is usually resolved by context, but in this dissertation, the initials ABM are used only to refer to a model.
Agent	An agent is an actor in an ABM. It could be an individual, a firm, a country, or an aggregation of other agents (Epstein and Axtell (1996, p. 4)).
Artifact	The Random House Dictionary defines artifact as “a spurious observation or result arising from preparatory or investigative procedures.” ¹ In the context of modeling and simulation, an artifact is typically an outcome that is the result of the way the simulation was done rather than a characteristic of the underlying model.
Emergent	An emergent behavior is one that arises from the the interactions or aggregate behavior of agents. It is not programmed into the model and often

¹<http://dictionary.reference.com/browse/artifact> (accessed 20 April 2010).

comes as a surprise to the modeler. *Emergence* refers to the presence of one or more emergent properties, or to the capability to produce emergent properties. The term appears throughout the literature on complexity, see Epstein and Axtell (1996, p. 33) for example.

Monte Carlo This is the city in Monaco famous for its casino. The term is used to refer to a technique for solving numerical problems for which an exact computation either does not exist or is computationally expensive (Metropolis and Ulam, 1949). The problematic calculation is replaced with a stochastic representation of its solution space. This latter step typically involves *distribution sampling*, a term with which Monte Carlo sampling is sometimes confused (Hendry, 1984). A simulation of an agent-based model in which certain variables are sampled from distributions is a Monte Carlo sample.

Chapter 1

Introduction

“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 unless 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 the prevailing interest rate. 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.

Hotelling's observation applies to the conditions imposed on the original model: a producer of an exhaustible resource with perfect information and a costless technology. Hotelling's primary goal was to illustrate how dynamic equilibrium produces a result that is not evident using the conventional (at that time) methods of static equilibrium. Hotelling's methodology has stood the test of time, although calculus of variations, to which Hotelling appeals for dynamic optimization, has since been supplanted by optimal control theory.

This dissertation has three objectives: to establish a framework for validation of an agent-based modeling approach to nonrenewable resource production planning and analysis; to use

¹Strictly speaking, price is not identical to marginal profit, particularly in a monopoly market.

²From a personal communication dated 25 May 2010.

this framework to validate a simple optimizing agent-based model (ABM); and to explore any new aspects of the problem that are brought to light by agent-based modeling.

The ABM developed specifically for the Hotelling monopoly model is extremely simple. It will be shown that, when incorporating highly stylized production technologies, the same ABM makes slight deviations from the optimal production. These deviations result in total profits that are less than optimal by from less than one percent to up to ten percent, depending on how the technology comes into the cost equation. I argue that the errors³ made by the agent in the model are similar to the errors that a human would make in the face of uncertainties in demand or resource extent. That argument, however, is based on elementary principles of agent-based modeling, which will be introduced and discussed in Chapter 3.

The ABM will also be extended to an oligopoly market where it will be shown that, despite the absence of communication between the agents, the production paths mimic either collusive or Cournot equilibrium depending on initial conditions. Finally, the results from the ABM simulations are re-examined as stylized fiscal policies. The efficiencies and unintended consequences of the policies are discussed in conjunction with the kinds of errors made by the ABM.

The preceding points are presented in greater detail in the following sections. The overarching idea is that, much as Hotelling's Rule sets the upper limit on a producer's profit, this simple ABM presents the worst that an optimizing producer might do. Certainly a more sophisticated ABM - or a human - could make better production path decisions. The intent of

³The term *error* is used throughout the following discussion of the ABM results. This is not error in the econometric sense, but an ABM production decision that deviates from the Hotelling's Rule optimum production path. Since the Hotelling's Rule production path results in the maximum possible total profit, the effect of that deviation can be measured in terms of the resulting shortfall in total profit. Quantitative error, then, is that shortfall as a percentage of the Hotelling's Rule optimal total profit. Because the decision rules applied by the ABM are intended to mimic a human decision process (however naive), the term *error* also reinforces that parallel without being overly anthropomorphic.

agent-based modeling is to have the software agents make decisions using the same information and rules that are available to a human. Agent-based modeling, however, requires much less time and expense than experiments with humans and allows sampling of a much larger parameter space than is available through data collection. Agent-based modeling presents the possibility of exploring dozens of production technologies, to consider uncertain and possibly changing resource extents, and to incorporate stochastic uncertainties in demand and interest rates. In addition to what-if analysis for production planning, agent-based modeling provides a laboratory for testing theories for why resource producers follow production paths that appear to deviate from Hotelling's Rule. It is hoped that, by validating the approach, this dissertation sets the stage for these and other uses of agent-based modeling in resource economics.

1.1 The inverse demand function and its parameters

All of the ABMs in this dissertation use the inverse demand function introduced by Hotelling (1931, sec. 4) for a monopoly producer. The details are presented in Section 3.2. Similarly, the parameter values for all simulations in this dissertation are from Section 3.2.

1.2 A simple ABM and Hotelling's outcome

The background and theory of Hotelling's Rule are presented in Chapters 2 and 3, respectively. The basic model is of a monopoly producer of a nonrenewable resource. The producer employs a technology that is costless, or for which the cost is absorbed in the price. That is, the technology imposes no costs that are dependent on time, the production level, or the amount of remaining stock. Hotelling's Rule states that if the producer is maximizing total profit, the marginal profit will increase at the interest rate. For a normal demand function, that means decreasing production each period at a rate that causes the percent change in

marginal profit to equal the interest rate. Given a specific demand function and a final state of stock depletion, this determines the production path. The initial stock level is the only additional information needed to determine the initial production level.

The process described above is straightforward, given knowledge of the demand function. If the demand function is differentiable and the constraint equation is integrable, the production path can be expressed in closed form. Otherwise, the production path will be expressed as a summation or, in the case that the demand function is not differentiable, as a complicated numerical expression. In any case, the optimal production path is explicit.

For the ABM, the production path is determined heuristically. An overview of agent-based modeling is given in Section 2.3, but for the moment, think of an agent as a robot or android. That is, an actor that can only do what it is told. In the monopoly ABM, there is a single producer agent, and this agent is given one task: extract a nonrenewable resource at a rate that maximizes total profit over the lifetime of the resource. The details of the agent's behaviors are discussed in Chapter 4 and can be summarized by noting that the agent uses basic arithmetic to decide whether to increase, decrease or maintain the current production level based on estimated total profit. This decision is made in each production period (days) and amounts to Bayesian updating.

At the beginning of each simulation, there is a short period during which the producer agent searches for the optimal starting production level. On average, the search takes approximately ten time-steps. For comparison, the most basic model in this dissertation completes in 1958 time-steps. Once the producer agent finds the optimal production level, it begins adjusting production to maximize total profit. Aside from the brief period at the beginning, the production path is very similar to the optimal (Hotelling's Rule) production path, being about one percent too high at the initial downturn, and ending a fraction of a percent too soon. This error is illustrated in Figure 1.1 with the error greatly exaggerated.

Total profit for the ABM is a fraction of one percent lower than the optimal production path. However, the ABM's percent change in marginal profit is nearly fifteen percent higher

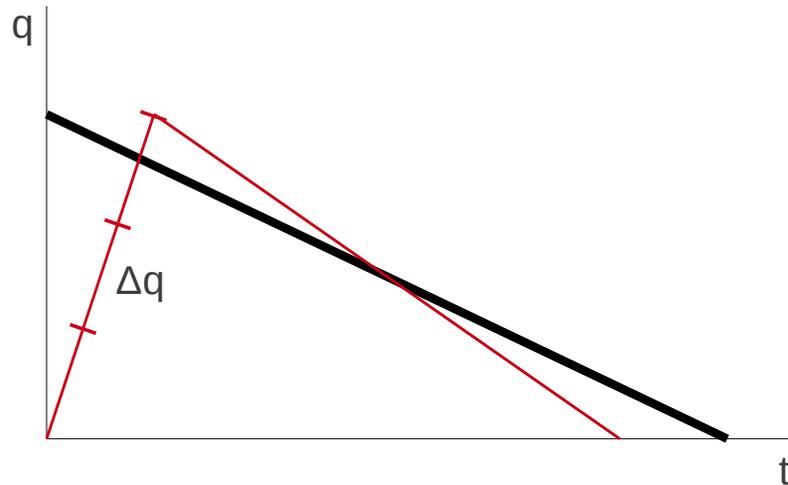


Figure 1.1: An example of the initial production level error.

The heavy black line is the Hotelling's Rule optimum production path. The thinner red line is the production path as determined by the ABM. 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.

than the r-percent rule. It will be shown that, although this differs from Hotelling's Rule, the consequences are very small.

The preceding results, discussed in detail in Chapter 6, are the first step in demonstrating the utility of agent-based modeling in natural resource economics. In addition, the results also show that, although the r-percent rule is optimal, production paths can differ from the r-percent rule and still produce nearly optimal total profit. The low sensitivity to deviations from the r-percent rule has important implications for studies in actual resource markets.

The ABM in these models is intentionally simplistic. However, real producers may only have a general idea of the demand function they face, have sparse and out-of-date information

on competitors or product substitutes, and contend with a multitude of other unknown or uncertain factors. It is possible that these producers would use crude approximations not unlike this one.

This dissertation focuses entirely on Hotelling's monopoly model (Hotelling, 1931, sec. 4), in which the rate-of-change of marginal profit is endogenous. In the competitive model (Hotelling, 1931, sec. 2), the corresponding price rate-of-change is imposed exogenously. In the latter case, Hotelling shows that resource owners would either divest themselves when price change is less than the interest rate, or forestall production when price change is greater than the interest rate. By establishing the use of agent-based modeling for the monopoly Hotelling model, the framework in this dissertation serves as starting point for developing a model that endogenizes the price mechanism of a competitive market.

1.3 The ABM and production technologies

Production technologies come into dynamic optimization problems in the form of costs. The production model in Section 1.2 employs a costless production technology. That is, there are no costs that depend on time, production level, or stock. Real production technologies may have costs associated with all three, often in combination. For example, there may be lease costs that affect the optimal production path. Most production technologies have costs that depend on the production level, and these costs may be polynomial or exponential in the production level. In most resources there are also stock-dependent costs. These are costs associated with extraction becoming more expensive, such as the deepening of a mine or decreasing pressure in a well. Another source of stock cost is the tendency to exploit higher grade ore first, causing the yield to decrease as the stock is depleted.

For simplicity, this ABM assumes that costs are separable and examines each one individually. To the extent that separability holds, the effect of a technology which imposes costs from more than one regime is likely to be a superposition. Nonseparable costs are not

considered in these models.

The following sections summarize the behavior of this ABM in each cost regime. In brief, the extent to which the ABM producer deviates from the optimal production path varies with cost structure. For error measured in terms of deviation from optimal, the ABM producer is unaffected by a fixed cost. Under a marginal cost, the ABM producer makes a small error that increases to about eight percent with cost level. With a stock cost the error is as high as ten percent and depends on whether any stock will be left unproduced. In the marginal cost and stock cost models, it is argued that part of the error is an artifact of the simulation, and part represents an issue also faced by a human production planner. That is, a small error in estimating future price leads to a small production path error early on, with ramifications over the lifetime of the resource. The ABM in this case can make daily adjustments to the production level, making up, partially, for earlier errors. A human production planner may not have such flexibility.

1.3.1 Lease and capital costs

Extractive industries tend to require large capital investments, and capital can be regarded as a quasi-fixed cost (Young, 1992). A lease incurs a fixed cost irrespective of the amount of resource extracted. Intuitively, a fixed cost provides an incentive to accelerate depletion of the resource stock in order to reduce total cost, and Section 3.4 shows that this is what is expected theoretically. In terms of total profit, the simulation result from the monopoly ABM is indistinguishable from the theoretical result. This is not surprising, since, theoretically, the optimum is a straight line descending production path with the same slope as the costless model.

1.3.2 Marginal costs

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. For simplicity, marginal cost is linear in production level for the monopoly ABM model with marginal cost. Because the ABM producer overshoots the optimal initial production level, as explained in Section 1.2, there is also an error in the marginal cost model. In this model, however, the size of the error increases with increasing cost. The effect is linear and proportional to the cost, amounting to an error in total profit from zero, at zero cost, to eight percent at the high end of the range of costs simulated. In all cases the agent picks in initial production level that is too high, resulting in the resource being depleted in less time than optimal.

1.3.3 Stock and cumulative production costs

Costs that depend on the stock level are typically costs related to the cumulative production. 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.⁴ Like marginal cost models, stock cost models are often quadratic or in Cobb-Douglas form (Young, 1992). For simplicity, stock cost is linear in stock level for the monopoly ABM model with stock cost. The effect is similar to that in the marginal cost model: in the cases where the stock is completely depleted, the ABM producer selects a production path that is higher than optimum, therefore depleting the stock sooner than optimal. However, with the increasing stock cost arises the possibility that at

⁴Slade (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.

some time before the stock is depleted, cost exceeds revenue, and marginal profit becomes negative. This is the *critical stock cost* - the stock cost below which the resource will be physically depleted, and above which production will cease before the stock is physically depleted. When marginal profit becomes negative, the simulation terminates, leaving up to sixty percent of the stock unproduced for the highest stock cost simulated. The range of stock costs simulated is chosen so that the critical stock cost is approximately mid-range.

Like the marginal cost model, deviation from the optimal production path increases as stock cost increases. Unlike the marginal cost model, however, the deviation is not linear with cost and actually improves as stock cost approaches the critical stock cost. The total profit for the ABM is approximately equal to the theoretical optimum at the critical stock cost, and differs by about ten percent at about half of the critical stock cost. ABM performance again begins to fall below the theoretical optimum at costs above the critical stock cost.

The performance of the ABM at costs below the critical stock cost is related to the coarseness of the production change strategy. The straight line decreasing production path optimizes total profit for zero cost, but is too steep for non-zero stock cost, with the error increasing as stock cost increases. At about one-half of the critical stock cost, the ABM producer begins switching to a level production path, which is too level at that stock cost, but as the optimal production path continues to level off, moving closer to the level ABM production path until the critical stock cost is reached.

Above the critical stock cost, the range over which the production level yields non-negative marginal profit narrows as the stock cost increases. The narrowing of the range increases the likelihood that an error in the production level will result in negative marginal profit, triggering termination of the simulation. As a result, production ceases earlier than optimal, and this error increases with increasing stock cost.

Overall, then, the performance of the ABM worsens as stock cost increases from zero, then improves as stock cost reaches critical stock cost, then worsens again as stock increases above critical cost.

Like the marginal cost model, the errors in the stock cost model are due to the initial production being too high, for the reasons given in Section 1.2. For the stock cost model, that error increases with increasing stock cost only to the point when the heuristic begins choosing the flat production path. Then, as the stock cost increases, the error decreases until the stock cost reaches the critical stock cost. At the critical stock cost, the optimal production path would be flat, so there is very little error in the production path chosen by the heuristic. Above the critical stock cost, the error in the initial production level results in production levels that reach negative marginal profit too soon, with this error increasing as stock cost increases.

1.4 The ABM in an oligopoly market

The oligopoly ABM is the monopoly ABM with more than one producer. This has some important implications. The producers don't interact directly, affecting each other only through the market price, which is based on total production. Having no information about the other producers, each producer estimates the marginal price assuming a monopoly market. This is intended to a) keep the ABM as simple as possible, and b) illustrate a worst case outcome. That is, any real-world producer will know if there are competitors, but a real-world producer is likely to make some erroneous assumptions about market structure. Mistakenly assuming a monopoly market is likely to be the worst case producer error.

The market in the oligopoly is responding to the total production. Each ABM producer, assuming itself a monopoly, estimates marginal price based on its own production changes. For example, if there are five identical producers and they all decrease production level in the same period by the same amount, each will overestimate the marginal price by a factor of five. This is canceled out, however, because each producer underestimates the required production change based on the overestimated marginal price. The change to total production is the sum of the changes made by all five producers, and the production change made by each producer has the desired effect after all.

If the producers are not identical, however, they will not be making margin price calculations at the same time and, therefore, will not be adjusting production levels simultaneously. In the oligopoly ABM, asymmetry among the producers is introduced by giving each one an initial resource stock from a tight distribution about a mean. The resource stock level is used to estimate the straight-line descending production path, so that each producer will have, in principle, a different initial production path. In practice, however, the step sizes in the ramp-up to the production path are large compared with the distribution of production paths, so that multiple producers are likely to find the same production path. With five producers, they typically fall into two groups, one at a production path slightly above the other. What happens next depends on the relative stock levels of the group members.

If there isn't much difference in the stock levels of the producers, they will behave much as they would with uniform stock levels. That is, changing production simultaneously, overestimating marginal price and underestimating production levels. These producers will all follow the straight-line descending production path, appearing to collude to divide the monopoly rent.

If, on the other hand, some of the producers reach their production path a period earlier than the others, they will begin on the straight-line descending path one period earlier than the others. The others have seen their marginal profit increase without reducing their production, so they will maintain production. In the next period, the bigger producers again decrease production, causing price to increase, leading the smaller producers to mistakenly conclude that maintaining their production level is increasing their marginal profit. This continues until the smaller producers deplete their stocks, which happens well before the others since they've been producing at a constant higher level. At this point the remaining producers see a spike in marginal price, since the total market production has decreased drastically, but it doesn't affect their behavior: they continue along the straight-line descending production path until their stocks are also depleted.

The latter outcome is a Cournot equilibrium, since each producer is making production decisions based on the production decisions of the others (as seen through the marginal

price), and total production is greater than the collusive level, but considerably less than the perfectly competitive level.

In either case, the outcome has the appearance of cooperative behavior, yet the producers have no means to cooperate. This has limited implications for real-world producers, since the likelihood is small that many producers will have the same technology at the same production levels and with nearly identical stock levels. The collusion-like outcome is validated by converging to the costless monopoly results. The Cournot-like behavior, though unexpected, illustrates the possibility for discovering new behaviors with agent-based models.

1.5 The ABM and the consequences of taxation

The models summarized in Section 1.3 examine costs as a consequence of production. If, on the other hand, the costs are treated as taxes on an otherwise costless technology, total welfare now encompasses the effectiveness of the taxes in achieving their fiscal goals. In all cases, taxes result in a transfer of producer surplus to public welfare. When equilibrium price or quantity is affected by taxation, total welfare must take deadweight loss into account.

The effects of the different cost models as taxes are based on the results discussed in Section 1.3. In addition to the fixed cost, marginal cost and stock cost models, an additional cost model is introduced in this discussion: royalties, which are levied as a percentage of profit. Suppose the goal of a tax policy is to extend the viable lifetime of a nonrenewable resource, or to mitigate a pollution externality, or to create revenue to fund public education. These ABMs show which policies slow production and which accelerate it, which reduce pollution externalities and which worsen them, and how much deadweight loss each policy incurs. These results are consistent with the literature and extend the theoretical models by illustrating the additional impact of the planning errors discussed in Section ??.

1.6 Chapter organization

The preceding is a high-level overview, the details of which are explored in the following chapters. Chapter 2 presents the related research and places this work in the context of the literature. Section 2.3 introduces the basics of agent-based modeling. Agent-based modeling is not new to the field of Economics but neither is its use widespread. Section 2.4 also includes a high-level overview of MASON, the ABM framework used for these models, which provides the basic tools for simulation and data visualization.

Chapter 3 reviews the theory behind Hotelling's Rule and presents the theoretically expected outcomes when introducing different production technologies. The basic approach is through optimal control theory, carrying the general solutions as far as they can be taken, then proceeding with a specific demand function from Hotelling. This chapter focuses entirely on the mathematical results: their interpretation will come into play in the discussions of Chapter 6.

Chapter 4 defines and describes the specific ABM used in this study. The behavior of interest initially is profit maximization given a finite resource, and in this model the producer agent is endowed with the heuristic decision process described in Section 4.3. The optimization heuristic is a simple set of rules and some basic mathematics with which a producer can make a crude estimate of total profit given a small set of production paths. The details of the actual simulation process are also presented in Chapter 4.

The simulation results are shown in Chapter 5, primarily in the form of graphs. Some of the results are presented as time-series plots, which is natural given dynamic simulation, but much of the results of interest occur in phase space: a space defined by the objective (profit maximization), the theoretical constraint (interest rate) and the behavioral outcome variable (percent change in marginal profit).

Chapter 6 discusses the results presented in Chapter 5 and compares them with the theoretical expectations developed in Chapter 3. The behavioral interpretation of the math-

ematics is combined with the economics interpretation of the simulations.

Chapter 7 examines the effect of taxes in the context of the ABM. In general, taxes take the form of the various cost structures presented in the preceding chapters, where a cost for the producer is income for the policy-maker.

Chapter 8 summarizes the relationship between theoretical and practical (meaning computer simulated) outcomes. This chapter will also discuss some of the ways in which the ABM models can be extended to explore additional aspects of natural resource production.

Chapter 2

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. Modern natural resource economics arose in response to the Club of Rome report “Limits to Growth” (Meadows et al., 1972), which was interpreted as predicting depletion of petroleum by 1992 and other key resources within 100 years.¹

Literature relating population growth to limited resources dates back, at least, to Jonathan Swift’s satire “A Modest Proposal for Preventing the Children of Poor People in Ireland From Being a Burden on Their Parents or Country, and for Making Them Beneficial to the Publick,” published anonymously in 1729. Nearly 70 years later, also publishing anonymously,

¹The report estimates the lifetime of known reserves at the time of writing, noting that new reserves will be found, but that these resources are, ultimately, finite.

Thomas Malthus introduced exponential growth models in “An Essay on the Principle of Population, as it affects the future improvement of society with remarks on the speculations of Mr. Godwin, M. Condorcet, and other writers.” Malthus’s developed his now famous economic models to refute the writings of William Godwin and the Marquis de Condorcet that implied there were no limits to growth in Europe.

Between Swift and Malthus, Adam Smith, in 1776, published “The Wealth of Nations,” setting the stage for market-based analysis of exhaustible resources. Then, 19 years after Malthus, David Ricardo published his “Principles of Political Economy and Taxation,” introducing the notion of economic profits in the form of rent. Ricardian rent reflects that, for two pieces of land that are indetical in every other way (including marginal cost), a landowner receives additional rent for the more productive land. In “Principals of Political Economy,” John Stuart Mill, in 1848, integrated Adam Smith’s free markets, malthusian limits on resources, and ricardian rents with a model of the role of technological advancement.

By the time of the Club of Rome report, economists were prepared to respond to it, and at the core of that response was Hotelling’s 1931 paper, the central point of Robert Solow’s 1973 Ely Lecture to the American Economics Association (Solow, 1974). Solow’s paper, along with a number of other papers from 1974, are introduced in the following sections. There are a number of historical reviews of the literature descending from Hotelling. 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).

There are two broad areas of interest in the Hotelling’s Rule literature: theoretical efforts to widen the scope of Hotelling’s Rule, and econometric efforts to find evidence of the Hotelling’s Rule outcome. Section 2.1 reviews developments in the theory, while Section 2.2 examines the evidence for scarcity rents in nonrenewable resource markets.

2.1 The evolving theory of nonrenewable resources

The essence of Hotelling's finding is that in any market for a nonrenewable resource, monopoly to competitive, market prices will exceed marginal costs by the scarcity rent. Furthermore, economically efficient production of the resource will require that the rent, and therefore, net price, will increase at the discount rate. Were the latter not the case, resource owners would either divest if prices rise slower than this, or let the resource stock appreciate in situ if prices rise more quickly.

Resource rents are made up of ricardian rent plus scarcity rent. Suppose that there are two pieces of land that are identical in every way (including marginal cost), differing only in the quality of a resource they hold. The land owner may receive rents (payments above marginal cost) for both pieces of land due to the presence of the resource. The owner will receive a larger rent for the piece of land with higher quality resource, and this is ricardian rent. If the resource is being extracted, the rent will increase by r -percent, and this increase is scarcity rent. That is, ricardian rent reflects the fact that the land has a valuable attribute, and scarcity rent reflects that fact that the attribute can be used up.

Many important aspects of natural resource markets follow from Hotelling's Rule. Not only is it optimal for the owner of a resource to profit from its scarcity, but it is the only way that the resource will be produced. The economic profit that arises because a resource is becoming more scarce is called *scarcity rent*, because the market price is greater than the marginal cost. Scarcity rent is also often called the *user cost* because it reflects an additional cost that is related to using up the resource. That is, it is the opportunity cost of forgoing the future benefit, including appreciation, of the next unit being produced.

Another consequence of Hotelling's Rule is that higher quality resources will be extracted before lower quality. Consider the previous example of two identical pieces of land that differ only in the quality of resource they contain. Presumably, the higher-quality deposit results in lower cost per ton of processed ore. As the lower cost (higher quality) ore is produced, the user cost (scarcity rent) increases until total cost equals that of the higher cost (lower

quality) ore. Then, the lower quality ore will be produced.

The production of nonrenewable resources is made complex by the existence of multiple grades (or costs) of deposits, changes in known reserves (discovery), trends in substitutes and in backstop technologies. These complexities tend to weaken or obscure the coupling between prices and scarcity rent. 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) explore 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. The evidence for these price curves is discussed in the next section.

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 in a duopoly market (Hotelling, 1931, sec. 15). He notes that with an exhaustible resource, where each competitor is optimizing based on the assumed optimization of the others, there is an incentive for a competitor to unilaterally raise its price by a small amount.

Consider, for example, a duopoly at Cournot equilibrium. For a product that is not exhaustible, there is an incentive for a producer to lower its price slightly, taking sales from the competitor unless the competitor reciprocates, ultimately driving profits to zero. For a nonrenewable resource, however, lowering price and taking sales early in the lifetime of the stock for a small increase in profit leaves the higher-price producer in the position of having additional stock late in the lifetime of the stock, when profits are higher. Thus, with a nonrenewable resource there is a disincentive to cheat by lowering price.

If, from the initial Cournot equilibrium, the cheater increases price slightly, there are two possible outcomes. If the competitor reciprocates, market price goes immediately to the higher price, and they both enjoy increased profit (until they reach the monopoly equilibrium). If the competitor does not reciprocate, the competitor is forced into the role of the price-lowering cheater described in the previous paragraph.

Salant (1976) presents a 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). This result is challenged by Gaudet and Lasserre (1988), however, who point out the Stiglitz 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) makes 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.

At present, there is no literature on agent-based models of nonrenewable resources. Agent-based modeling can be a laboratory for exploring and experimenting with proposed extensions to the Hotelling model. Agent-based modeling is a causal framework, meaning that 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. In a hypothetical agent-based model, the producer agent would be given a behavioral rule to "produce whichever deposit is least costly to produce," which would work for either class of exploration model. In this hypothetical agent-based model, a parameter would determine whether new deposits are discovered in decreasing quality or random quality.

A limitation to extending the theory of Hotelling's model to oligopoly markets is that a market structure must be assumed. For the agent-based models in this dissertation, an oligopoly market is the same as a monopoly market with additional agents. The market structure is, therefore, an endogenous feature of the simulation, with the outcome depending on the kinds of communication the modeler has given the agents.

Agent-based modeling in support of theoretical development is said to have an objective of “qualitative insight and theory generation” (Tesfatsion, 2006). In the sense of any model, it is *predictive*. The modeler defines the behavior space (rules and initial conditions) and explores the outcome space.

2.2 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 than 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. "In a world of rapid technological change ... unit extraction cost will fail to increase even though a natural resource is becoming more scarce." (Brown and Field, 1978, p. 231)

Halvorsen and Smith (1991) note that “the principal obstacle to empirical tests of the theory of exhaustible resources has been data availability.” Since mine owners do not reveal actual marginal costs, so that actual user cost is not observable, many of the market studies are essentially assessments of the suitability of proxies for these.

Nordhaus (1973) compares the price paths of different sources and grades of energy resources with the price paths of the good produced using them. Nordhaus (1973, p. 566) observes that, on the average, scarcity rents on energy resources were "quite modest," or

“about one dollar per barrel of petroleum.”² 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 scarcity rents 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 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

²For 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)

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. Halvorsen and Smith (1991) conclude that "the ability of the theory of exhaustible resources to describe and predict the actual behavior of resource markets remains an open question."

Agent-based modeling can be used as a method for nonlinear regression. Given a set of target outcomes (market data), a hypothetical ABM can systematically vary behavioral rules and initial conditions until simulation outcomes converge to the target outcomes. For example, in the Agostini (2006) model, the hypothetical ABM producers might switch from profit maximizing to risk averse as the price approaches a threshold. The threshold itself can be fuzzy in this hypothetical model, meaning it has a stochastic distribution which may, in turn, be dependent on price history, trends in substitutes or complements, and so on.

Agent-based modeling to explore behavior space in order to explain outcomes is said to have an objective of "empirical understanding" (Tesfatsion, 2006), and such models are referred to as "generative" (Epstein and Axtell, 1996). It is *analytical*: the modeler defines the target outcome space, then varies the behavior space (rules and initial conditions) until outcomes fall within the target outcome space.

2.3 The basics of agent-based modeling

In economics there are two different uses of the term *agent*. In the principal-agent problem, an agent is someone who does something for the principal (Varian, 1992, p. 441). Alternatively, 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). Though typically implicit in economic models, there may also 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 each agent. In agent-based modeling, the term endogenous is used to imply autonomy. By contrast, the term exogenous is used to imply that agents are being told what to do. In a simple ABM of flocking birds, for example, if each bird decides on the best distance to be from all nearby birds, that is an endogenous behavior rule. The fact that they want to be together at all, however, may be imposed, making that an exogenous behavior rule. 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.

The last statement begs some definitions. Practitioners of modeling and simulation will often use the terms *model* and *simulation* in ways that make them sound interchangeable (Department of Defense, 1998, p. 138). They are not, however, and confused readers are in good company on this topic. Just as in economics or architecture, a computer model is an *abstraction* of reality. That is, it is not a complete description of reality, for many features of reality have been omitted (because they're thought to be irrelevant) or abstracted (in detail to reduce complexity in parameter space or in behavior, or to reduce complexity in outcome space). By analogy, an architectural model may omit the details of subsurface soil (parameter space), use decals to represent windows and doors (behavior space), and have removable parts to expose functionality (outcome space).

Simulation is what a model does, or what you do with it. As an example of the first meaning, a model may include some number of birds and the simple rules they follow: stay close together, but not too close. Simulation would involve running a computer program that start all the birds in some position, gives them each a chance to react according to their rule, then records their new positions, and does this over and over for some period of

time. Depending on the rules, the motion (as portrayed by dots on a screen, for example) may look like flocking, or it may look like swarming. As an example of the second meaning of simulation, you might use that model "to simulate a flock of starlings", or "to simulate crowds moving through a theme park." In other words, simulation is an abstraction of what it *does*, much like modeling is an abstraction of what it is.

In any model, there are things about the agents that distinguish them from each other and cause them to interact with each other or their environment. For example, in a basic economic model, an agent may have a variable *Wealth* and a variable *Preference*, which represents a preference for guns or butter as a real number between 0 and 1, with 0 for guns and 1 for butter. The variables *Wealth* and *Preference* are a representation of an agent's *state*. This particular model is not interesting without a means for agents to change state, and that is accomplished through behavioral rules. Something like "buy guns and butter in the ratio reflected by *Preference* until *Wealth* is depleted." In general, what is interesting in a behavioral model like this is its development over time. That is, 1) each agent acts on its state, 2) the action results in a change of state, 3) the new state results in new actions. The time evolution of a model is performed in simulation, which simulates the passage of time and any other exogenous changes while the actions of the agents bring about any endogenous changes. The initial and terminal conditions are also aspects of simulation, so that the same model may be used to simulate prosperity or poverty (average *Wealth* is either large or small).

A model may include a large number of these agents, all with different preferences. The number and types of agents, the things they are able to affect, and the rules they use to affect those changes are all part of the model. What happens in the course of the agents changing is the outcome of simulation.

Early agent-based models attempted to replicate simplified insect behaviors, such as group movement (Reynolds, 1987) or an ant colony's quest for food (Epstein and Axtell, 1996). In the former, optimizing interactions between individuals leads to new behaviors like flocking or swarming, while in the latter, self-optimization by individuals leads to positive

group outcomes. In either case, the individual actors are the agents (each bee, bird, or ant), and each agent has simple rules for its own behavior. The group behavior is not programmed, and is said to be an *emergent* outcome of the model. It is the goal of autonomous agent-based models to discover or explain emergent behaviors or outcomes. Other agent-based models are computational models: they simulate large ensembles of individually optimized equations. Both uses arise in this study, with autonomous agents as profit-maximizing producers, and computational agents as their Hotelling's-Rule-following counterparts.

In the social sciences, one early adopter of agent-based modeling is Axelrod (1997b), whose models started out as an extension of earlier work on the Prisoner's Dilemma in game theory (Axelrod, 1987). Agent-based modeling arose as computation became faster, cheaper, and widely available. Many of the behavioral models developed theoretically in the preceding decades, such as those of Axelrod (1987) and those of Schelling (1978) were well suited to agent-based modeling. Epstein and Axtell, both jointly and separately, developed 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 presents models of 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) presents 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 dissertation, agent-based modeling is referenced in its entirety, while the initials ABM are reserved for a model.

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. One approach is to provide a complete agent-based modeling and simulation program with which the user may create an agent-based model and run simulations. These programs work much like spreadsheet or painting software. An example is the NetLogo project at the Center for Connected Learning and Computer-Based Modeling. A simple NetLogo model can be constructed and simulated quickly, and there is large body of example models as well as an extensive user community. Though it is technically not programming, writing rules for behavior in NetLogo's unique language becomes very much like writing computer programs.

The other approach to agent-based software is to provide a library of programs to manage the modeling details of how to store and change agent state and how to communication between agents, as well as simulation issues such as keeping (and advancing) system time, managing simulation events, displaying state in real time, and saving simulation results. These libraries are written in a specific computer language such as Java or Objective-C, and a user is expected to do some programming to create the necessary agents and to program any behaviors required. MASON, a project at George Mason University, is an example, and is used for the modeling and simulation presented in the following chapters. More details on MASON are presented in the next section.

As an example of a simple ABM, consider the segregation model of Schelling (1971). This model was originally developed with coins and a checkerboard, but has since been a popular model for software agents. The general model is a checkerboard where each cell is occupied by an agent of one type (a copper coin, for example), an agent of another type (a silver coin, for example), distributed randomly, and there are a few unoccupied squares. Each agent has a preference for a certain percentage of neighbors to be of the same type. Each agent is given an opportunity to move into an unoccupied squares. The emergent behavior is that even with a mild preference for like neighbors (33 percent for example), in only a few turns the board is segregated into clusters of only copper coins or only silver coins. In this

model, there are two agent types, each agent has a state that is the percentage of neighbors of the same type, and each agent has a behavioral rule: move to an unoccupied square if the percentage of like neighbors is less than the preference (e.g. 33 percent).³

2.4 The MASON agent-based modeling platform

MASON (Luke et al., 2003) is a project of two organizations at George Mason University, the Evolutionary Computation Laboratory and the George Mason University Center for Social Complexity.⁴ MASON is one of a class of agent-based modeling and simulation software that provides a set of programs to support modeling and simulation, but the user must do some programming to make it work. This class of software is often referred to as a *toolkit*, and the overall software as an agent-based modeling *tool*. MASON is written entirely in the Java programming language, requiring a MASON modeler to be able to program in Java.

Most agent-based modeling tools use discrete-event simulation, meaning that time passes in discrete units and some assumptions are made about the intervening period. In some cases it is assumed that state variables evolved continuously during that time (bodies in motion tend to stay in motion, for example). In other cases, the discrete time unit is intrinsic to state, such as interest that is compounded daily. As with other discrete-event tools, the central feature of MASON is a means for scheduling time steps. In this sense, MASON regards each agent as an event on a schedule. Typically, after completing any actions for a time step, agents either put themselves back, or are put back, on the schedule for the next time step. MASON also provides a means to initialize agents at the start of a simulation,

³Strictly speaking, the original is a *cellular automata* model, but it is easily cast from a model of the state of a square on the checkerboard to a decision of a coin to move into one, making it an agent-based model. The Schelling Segregation Model is a common example in agent-based modeling texts and courses. See, for example, Clark (1991) or Laurie and Jaggi (2003).

⁴Information and software downloads can be found at <http://cs.gmu.edu/~eclab/projects/mason/> (last accessed 21 May 2011)

and a means to stop the simulation when a particular state occurs (such as an agent running out of money or a certain time having elapsed).

The MASON Hotelling model developed for this dissertation can be downloaded⁵ and anyone with the MASON software can compile it and have it running in a minute or two. MASON provides a variety of predefined graphical user interface (GUI) windows, some of which are shown in the Appendices. The MASON windows (GUI components) serve two purposes: 1) to provide a way to display features of a model and changes of state during a simulation, and 2) to provide a means to set or change the parameters of a simulation. In this sense, the act of model-building is extended a bit beyond the Java programming phase by giving the user some limited means to affect the model or its simulation. The models described in Chapter 4 can be selected, initialized and run with the information provided in the following paragraphs.

The interactive aspect of the MASON Hotelling model provides the non-programmer user a means to participate in defining the model. Specifically, although each model of a market structure was programmed in Java, the user is able to select from among twenty market models. Similarly, the user can select from two demand functions.⁶ Additionally, the user can select the number of producers and the interest (discount) rate before starting the simulation. Once the simulation is initialized, the user can modify the state of any agent, such as starting production level, initial stock, and costs.

In more general agent-based modeling terms, selecting among the market models amounts to selecting among behavior rules, while setting the demand function and other values is a matter of setting simulation parameters. It would take a very long time to explore all the permutations available at the time of writing, and the number of additional market models

⁵<http://www.unm.edu/~ddixon/ACE> (last accessed 21 May 2011)

⁶There are 20 market models (all of which are included in this dissertation) and 2 inverse demand functions (of which only one is used in this dissertation) as of May 2011. (The unused demand function is the competitive market model from Hotelling (1931, sec. 2)).

and demand functions is theoretically unlimited.⁷

⁷A comment like this is not without caveats. Performance is a decreasing concern as computational capabilities increase rapidly, but the program could reach the point where just selecting among the options could take a very, very long time.

Chapter 3

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).

3.1 Basic theory

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 $\pi(q(t), x(t), t)$ is a function of production level, remaining stock level $x(t)$ and time. Assuming a constant discount rate r , the optimal control problem is

$$\max_q J(q(t), x(t)) = \int_0^T e^{-rt} \pi(q(t), x(t), t) dt \quad (3.1)$$

subject to

$$\begin{aligned} \dot{x}(t) &= -q(t) \\ x(0) &= x_0 \\ x(t) &\geq 0 \end{aligned} \quad (3.2)$$

$$q(t) \geq 0 \quad (3.3)$$

$$x_0 \geq \int_0^T q(t) dt \quad (3.4)$$

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 (3.1) with resource stock costate variable $m(t)$ is defined as

$$\mathcal{H}(q(t), x(t), t, m) \equiv \pi(q(t), x(t), t) - m(t)q(t) \quad (3.5)$$

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, the value of the next unit of remaining stock to be extracted.

The first order necessary conditions include the state equation

$$\frac{\partial \mathcal{H}(q(t), x(t), t, m)}{\partial m} = \dot{x}(t) = -q(t) \quad (3.6)$$

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

$$\dot{m}(t) = rm(t) - \frac{\partial \mathcal{H}(q(t), x(t), t, m)}{\partial x(t)} \quad (3.7)$$

where $\dot{m}(t)$ is the time derivative of $m(t)$. If the Hamiltonian has no stock effect (no $x(t)$ dependency), this equation requires that the shadow price increase at the discount rate. That is, the shadow price increases at the same rate of an alternative investment. 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

$$\frac{\partial \mathcal{H}(q(t), x(t), t, m)}{\partial q(t)} = 0 \quad (3.8)$$

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

$$m(T) \geq 0, \quad x(T) \geq 0, \quad m(T)x(T) = 0 \quad (3.9)$$

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 (mx) of the terminal stock must be zero. That is, either the stock is physically depleted and $x(T) = 0$, or the ending shadow price $m(T)$ is zero. For nonzero ending stock, the requirement 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

$$e^{-rT}\mathcal{H}(T) = 0 \quad (3.10)$$

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 (3.7)

$$\dot{m}(t) = rm(t) - \pi_x(q(t), x(t), t) \quad (3.11)$$

so that $\pi_x(q(t), x(t), t)$ is the Hamiltonian stock effect to which the previous remarks apply. That is, the rate at which the shadow price changes is either accelerated or decelerated by

$\pi_x(q(t), x(t), t)$ 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 (3.8)

$$\begin{aligned}\pi_q(q(t), x(t), t) - m(t) &= 0 \\ \pi_q(q(t), x(t), t) &= m(t)\end{aligned}\tag{3.12}$$

which establishes the link between marginal profit and shadow price. Substituting (3.12) into (3.11) to eliminate $\dot{m}(t)$, then dividing by $\pi_q(q(t), x(t), t)$

$$\frac{\dot{\pi}_q(q(t), x(t), t)}{\pi_q(q(t), x(t), t)} = r - \frac{\pi_x(q(t), x(t), t)}{\pi_q(q(t), x(t), t)}\tag{3.13}$$

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

$$\frac{\dot{\pi}_q(q(t), x(t), t)}{\pi_q(q(t), x(t), t)} = r\tag{3.14}$$

which is the relation first articulated by Hotelling. It states that, along the optimal production path, the percent change in marginal net profit is equal to the discount rate. Thus, for a marginal net profit increase at a rate below r , it is optimal for the owner to extract the resource immediately and invest the returns with a return of r . However, for a marginal net profit increase at a rate above r , it is optimal for the owner to leave the resource in situ and let it appreciate faster than the alternative investment.

3.1.1 A benefit-cost treatment

In the preceding discussion, in Hotelling's original models, and for the purposes of this dissertation, the object of (3.1) is maximization of the producer's profit. Krautkraemer (1998) takes a more general, benefit-cost approach that is instructive when expanding the scope to include total welfare or social benefit.

If a resource yields gross social benefit $B(q(t), x(t), t)$ and the production technology incurs costs $C(q(t), x(t), t)$, equation (3.5) becomes

$$\mathcal{H}(q(t), x(t), t, m) \equiv [B(q(t), x(t), t) - C(q(t), x(t), t)] - m(t)q(t) \quad (3.15)$$

In this case, the optimality condition (3.8) is

$$\begin{aligned} B_q(q(t), x(t), t) - C_q(q(t), x(t), t) - m(t) &= 0 \\ B_q(q(t), x(t), t) &= C_q(q(t), x(t), t) + m(t) \end{aligned} \quad (3.16)$$

This condition requires that, at every moment in time, the marginal benefit from producing the resource exactly equals the marginal cost, which includes the user cost. Since the user cost represents the cost of forgoing future benefit, this is benefit in the present that comes at the expense of benefit in the future. When the same party is benefited in either case, it represents an intertemporal preference. When they are different parties it represents an externality.

3.2 The Hotelling monopoly demand function

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

$$p = (1 - e^{-Kq}) / q \quad (3.17)$$

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, revenue

$$p(q)q = 1 - e^{-Kq}$$

increases monotonically with q , a characteristic that is exploited in the models for which revenue must be held above cost. Because the revenue function increases monotonically, there is a unique minimum production level q_{min} which corresponds to a minimum revenue.

Equation (3.17) is the inverse demand function used throughout this dissertation, 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.

3.3 A costless monopoly model

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

$$\pi(q(t)) = p(q(t))q(t). \quad (3.18)$$

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 (after Hotelling (1931, sec. 4)) is:

1. Solve (3.12) for the production path $q(t)$ in terms of $m(t)$.
2. Solve (3.11) to get $m(t)$.
3. Replace $m(t)$ in $q(t)$ and solve (3.10) to get $q(T)$ in terms of T .
4. Integrate (3.4) 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 3.2, from step 1

$$\pi(q(t)) = 1 - e^{-Kq(t)} \quad (3.19)$$

$$\pi_q(q(t)) = Ke^{-Kq(t)} \quad (3.20)$$

$$m(t) = Ke^{-Kq(t)} \quad (3.21)$$

$$q(t) = \frac{\ln(K/m(t))}{K} \quad (3.22)$$

Equation (3.21) 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 (3.19), $\pi_x = 0$, so that from step 2

$$\begin{aligned} \dot{m}(t) &= rm(t) \\ m(t) &= m_0 e^{rt} \end{aligned} \quad (3.23)$$

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

$$q(t) = \frac{1}{K} \left(\ln \frac{K}{m_0} - rt \right) \quad (3.24)$$

If the stock is to be physically depleted, the transversality condition (3.10)

$$\mathcal{H}(T) = 0$$

applies. This means that

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

for which $q(T) = 0$ is the solution. Now (3.24) can be solved for T

$$\begin{aligned} 0 &= \frac{1}{K} \left(\ln \frac{K}{m_0} - rT \right) \\ T &= \frac{\ln(K/m_0)}{r} \end{aligned} \quad (3.25)$$

Substituting this back into (3.24) gives the production path

$$q(t) = \frac{r}{K} (T - t) \quad (3.26)$$

Equation (3.25) shows that higher discount rates promote more rapid depletion, which is expected, as a higher discount rate increases the opportunity cost of not extracting. Equation (3.26) 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

$$x_0 = \int_0^T q(t) dt = \int_0^T \frac{r}{K} (T - t) dt = \frac{r}{2K} T^2 \quad (3.27)$$

$$T = \sqrt{\frac{2Kx_0}{r}} \quad (3.28)$$

Equation (3.28) 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

$$q(0) = \frac{r}{K}(T - 0) = \sqrt{\frac{2rx_0}{K}} \quad (3.29)$$

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 (3.9) imposes no constraint on $m(t)$, which is

$$m(t) = Ke^{-r(T-t)} \quad (3.30)$$

That is, the shadow price of the remaining stock increases to the choke price K as the stock is physically depleted. Note also from (3.26) that

$$\dot{q} = -\frac{r}{K} \quad (3.31)$$

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

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

which is Hotelling's Rule.

The production path given by (3.26) 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 (3.19) over the stock lifetime T , the theoretical maximum profit, therefore, is

$$\Pi_{max} = \int_0^T (1 - e^{-Kq(t)}) e^{-rt} dt = \frac{1}{r} [1 - e^{-rT} (1 + rT)] \quad (3.32)$$

For purposes of comparison with other models that cannot be solved in closed form, using the values from Section 3.2, the stock lifetime is 1958 days, and $\Pi_{max} = 358.33$.

3.4 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

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

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

Inserting the demand function in $p(q(t))$ (Section 3.2) this is

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

One characteristic of the revenue part of this is that it is monotonically increasing with q .

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

$$q_{min} = \frac{1}{K} \ln \left(\frac{1}{1 - c_0} \right) \quad (3.33)$$

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 $c_0 = 0$.

The transversality condition (3.10)

$$\mathcal{H}(T) = 0$$

applies, which means that

$$1 - e^{-Kq_T} - c_0 - m_T q_T = 0$$

Using (3.21) to substitute m_T

$$\begin{aligned} e^{-Kq_T} (1 - Kq_T) &= 1 - c_0 \\ e^{Kq_T} &= \frac{1 + Kq_T}{1 - c_0} \end{aligned} \quad (3.34)$$

Equation (3.34) must be solved numerically. The solution is shown in Figure 3.1. The figure shows that $q_T > q_{min}$ for all costs, so that the $q_T > q_{min}$ constraint is non-binding. Figure

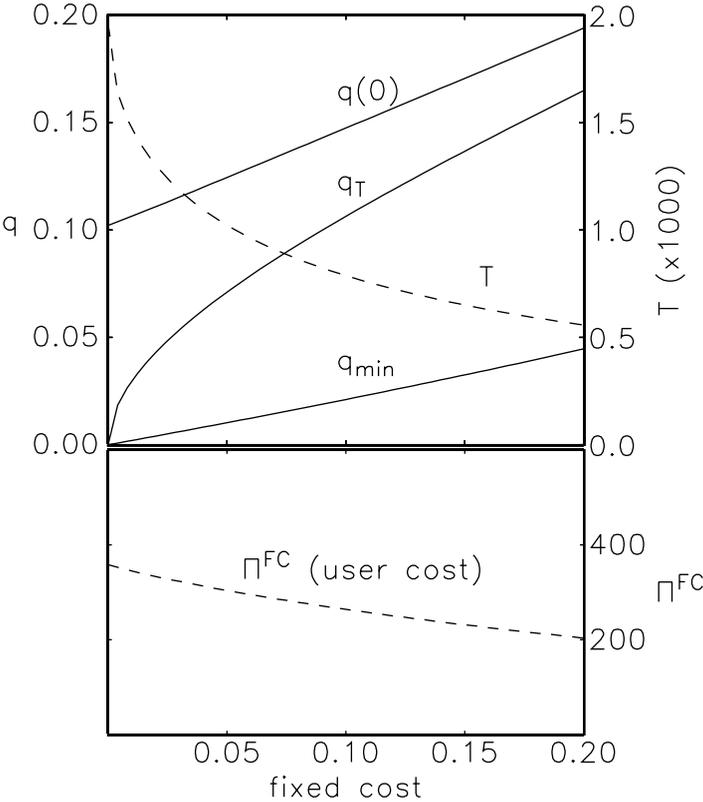


Figure 3.1: Fixed cost model - numerical solutions for terminal production level q_T . The optimal production path starts at production level $q(0)$ and decreases continuously to the terminal production level, q_T . The production level is, at all times, well above the zero profit production level, q_{min} . 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.

3.1 also shows that terminal time T decreases sharply as cost increases, and that total net profit (user cost) decreases steadily over the cost range. These are the theoretical outcomes against which the ABMs are to be compared.

Note, in Figure 3.1, 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 (3.23) with $m_T e^{-rT}$

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

then using (3.21)

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

so that

$$q(t) = \frac{1}{K} \ln \left[\frac{K}{K e^{-Kq_T} e^{-r(T-t)}} \right] = q_T + \frac{r}{K} (T - t) \quad (3.35)$$

Thus, the initial production level is

$$q(0) = q_T + \frac{rT}{K} \quad (3.36)$$

which is also shown in Figure 3.1.

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

$$T = \frac{K}{r} \left(\sqrt{q_T^2 + \frac{2rx_0}{K}} - q_T \right) \quad (3.37)$$

which is shown in Figure 3.1. Substituting (3.37) back into (3.36)

$$q(0) = \sqrt{q_T^2 + \frac{2rx_0}{K}} \quad (3.38)$$

Total profit is

$$\begin{aligned} \Pi^{FC} &= \int_0^T (1 - e^{-Kq(t)} - c_0) e^{-rt} dt \\ &= \frac{1}{r} [1 - c_0 - e^{-rT} (1 - c_0 + rT e^{-Kq_T})] \end{aligned} \quad (3.39)$$

which reduces to (3.32) for $c_0 = 0$ (for which $q_T = 0$). This is also shown in Figure 3.1, and the GAUSS code to solve for q_T and to compute $q(0)$, T and total profit is included in the Appendices.

3.5 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

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

In terms of the demand function in Section 3.2 this is

$$\pi(q(t)) = 1 - e^{-Kq(t)} - c_1q(t) \quad (3.40)$$

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

$$e^{-Kq(t)} \leq 1 - c_1q(t) \quad (3.41)$$

Figure 3.2 shows graphs of the left-hand and right-hand sides of (3.41) for the parameter values presented in Section 3.2. The graphs show that, for all marginal costs lower than the choke 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 (3.21). Thus

$$q(t) = \frac{1}{K} \left(\ln \frac{K}{m_0 e^{rt} + c_1} \right) \quad (3.42)$$

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

$$\begin{aligned} 0 &= \frac{1}{K} \left(\ln \frac{K}{m_0 e^{rT} + c_1} \right) \\ m_0 &= (K - c_1) e^{-rT} \\ q(t) &= \frac{1}{K} \left(\ln \frac{K}{(K - c_1) e^{-r(T-t)} + c_1} \right) \end{aligned} \quad (3.43)$$

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

$$\begin{aligned} e^{-Kq(t)} &= \frac{(K - c_1) e^{-r(T-t)} + c_1}{K} \\ -K\dot{q}e^{-Kq(t)} &= \frac{r(K - c_1) e^{-r(T-t)}}{K} \end{aligned} \quad (3.44)$$

$$\begin{aligned} \dot{q}e^{-Kq(t)} &= -\frac{r}{K} \left[\frac{(K - c_1) e^{-r(T-t)} + c_1}{K} - \frac{c_1}{K} \right] \\ \dot{q} &= -\frac{r}{K} \left[1 - \frac{c_1}{K} e^{Kq(t)} \right] \end{aligned} \quad (3.45)$$

The terminal time T is found by integrating

$$x_0 = \int_0^T \frac{1}{K} \ln \frac{K}{(K - c_1) e^{-r(T-t)} + c_1} dt \quad (3.46)$$

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

Equation (3.46) is solved numerically in GAUSS using the parameters from Section 3.2. The numerical solution for T as a function of marginal cost is shown in Figure 3.3. Once T for a given marginal cost is known, the initial production level is determined from

$$q(0) = \frac{1}{K} \left(\ln \frac{K}{(K - c_1) e^{-rT} + c_1} \right) \quad (3.47)$$

Values of $q(0)$ corresponding to the numerical solutions for T are also shown in Figure 3.3. Also shown in the figure are the ranges of the rate of change in production, $\dot{q}(t)$, computed from (3.45). 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, Π^{MC} are included in the Appendices.

To compute percent change in marginal net profit, substitute for the exponential on the left-hand side of (3.44) and simplify

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

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.

3.6 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.¹ 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 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

$$c_s(x(t)) = c_2(x_0 - x(t)) \quad (3.48)$$

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

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

¹Slade (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.

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 3.2, is

$$\pi(q(t), x(t)) = 1 - e^{-Kq(t)} - c_2[x_0 - x(t)] \quad (3.49)$$

The derivative of the Hamiltonian is no longer zero

$$\frac{\partial \mathcal{H}(q(t), x(t), t, m)}{\partial x(t)} = \pi_x = c_2 \quad (3.50)$$

so that, unlike the preceding models, the general form of Hotelling's Rule (3.13)

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

applies rather than (3.14)

$$\frac{\dot{\pi}_q(q(t), x(t), t)}{\pi_q(q(t), x(t), t)} = r.$$

With cost based on cumulative production, it is possible for marginal cost to exceed marginal revenue 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 (3.49) 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 (3.9) requires that $m(T) = 0$.

The first order necessary condition for \mathcal{H}_q proceeds as for the costless model up to (3.22). From the first order necessary condition for \mathcal{H}_x (3.11)

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

or

$$\dot{m}(t) - rm(t) = -c_2 \tag{3.51}$$

Solving for the homogeneous part

$$\begin{aligned} \frac{dm(t)}{dt} - rm(t) &= 0 \\ m(t) &= Ce^{rt} \end{aligned} \tag{3.52}$$

which is equivalent to

$$m(t)e^{-rt} = C.$$

Differentiating the term on the left

$$\frac{d}{dt} [m(t)e^{-rt}] = e^{-rt} [\dot{m}(t) - rm(t)] \tag{3.53}$$

Multiplying both sides of 3.51 by e^{-rt} and integrating

$$\int [\dot{m}(t) - rm(t)] e^{-rt} dt = -c_2 \int e^{-rt} dt.$$

Substituting the term in the left integral with 3.53 and integrating the term on the right

$$\int \frac{d}{dt} [m(t) e^{-rt}] dt = \frac{c_2}{r} [e^{-rt} - 1] + C.$$

Multiplying both sides by e^{rt}

$$m(t) = \frac{c_2}{r} [1 - e^{rt}] + Ce^{rt}$$

Solving for C using the yet-to-be-determined terminal shadow price m_T ,

$$\begin{aligned} m_T &= \frac{c_2}{r} [1 - e^{rT}] + Ce^{rT} \\ C &= m_T e^{-rT} - \frac{c_2}{r} [e^{-rT} - 1] \end{aligned}$$

so that

$$\begin{aligned} m(t) &= \frac{c_2}{r} [1 - e^{rt}] + \left\{ m_T e^{-rT} - \frac{c_2}{r} [e^{-rT} - 1] \right\} e^{rt} \\ &= \frac{c_2}{r} - \frac{c_2}{r} e^{rt} + m_T e^{-rT} e^{rt} - \frac{c_2}{r} e^{-rT} e^{rt} + \frac{c_2}{r} e^{rt}. \end{aligned}$$

and finally

$$m(t) = \frac{c_2}{r} [1 - e^{-r(T-t)}] + m_T e^{-r(T-t)}. \quad (3.54)$$

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

$$\mathcal{H}(T) = 1 - e^{-Kq(T)} - c_2 x_0 - m_T q(T) = 0$$

which yields

$$e^{-Kq(T)} = \frac{1 - c_2 x_0}{1 + Kq(T)} \quad (3.55)$$

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.

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 $q(T)$,

$$1 - e^{-Kq(t)} - c_2 [x_0 - x(T)] = 0 \quad (3.56)$$

The first order necessary condition (3.21) implies that, if $m(T) = 0$, then $e^{-Kq(T)} = 0$, so that (3.56) becomes

$$x(T) = x_0 - \frac{1}{c_2} \quad (3.57)$$

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

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

Finally, in this regime, the terminal Hamiltonian is

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

which is satisfied by the terminal stock level (3.57). Finally, terminal time T is found by integrating

$$x_0 - x(T) = \int_0^T q(t) dt \quad (3.58)$$

where

$$q(t) = \frac{1}{K} \ln \frac{K}{\frac{c_2}{r} [1 - e^{-r(T-t)}] + m_T e^{-r(T-t)}} \quad (3.59)$$

The numerical solutions for the $x(T) = 0$ regime involve solving for $q(T)$ using (3.55), computing $m(T)$ from (3.21), then numerically integrating (3.59) to find the T that solves (3.58) with $x(T) = 0$. For the $x(T) > 0$ regime, m_T is assumed zero, and (3.59) is integrated to find the T that solves (3.58) where $x(T)$ is found using (3.57). The GAUSS procedures for solving $q(T)$ and T and for computing m_0 , m_T , x_T , and total profit, Π^{SC} are included in the Appendices.

On a final note, (3.13) implies that the percent change in marginal net profit changes over time, since π_x is constant while π_q , equation (3.20), is a function of time. The percent change in marginal net profit is positive for

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

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

$$\frac{\dot{\pi}_q}{\pi_q} = \frac{\frac{d}{dt}Ke^{-Kq(t)}}{Ke^{-Kq(t)}} = -K\dot{q} \quad (3.60)$$

where \dot{q} is the time rate of change in production. This is found by taking the derivative of (3.21) with respect to time

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

Substituting (3.12) into the left-hand side, and (3.60) into the right-hand side,

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

The slope of the production path changes over time, can be positive for some values of c_2 , and becomes infinite when $m(T)$ is zero. These outcomes have the following economic interpretations. For a small stock cost, the shadow price increases as the stock is physically depleted, as before, but the final shadow price is lower if the stock cost is higher. At the critical stock cost, the terminal shadow price is zero and the instantaneous slope of the production path is infinite. This reflects the fact that the opportunity cost of the next unit of production, were there one, would be zero. For stock costs above the critical stock cost, the shadow price becomes zero before the stock is physically depleted, with more stock left at higher stock costs. In this case the opportunity cost of the remaining stock is zero.

3.7 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 perfectly competitive market, if a perfectly 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 an example of a collusion. 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).

Table 3.1: 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				
$c_2 < \frac{1}{x_0}$	$\frac{1}{K} \ln \frac{K}{\frac{c_2}{r}[1-e^{-rT}] + m_T e^{-rT}}$	$-\frac{r}{K} \left(1 - \frac{c_2}{rK} e^{Kq(t)} \right)$	$> 0^{**}$	$r - c_2/\pi_q$
$c_2 \geq \frac{1}{x_0}$	$\frac{1}{K} \ln \frac{K}{\frac{c_2}{r}[1-e^{-rT}]}$	$-\frac{r}{K} \left(1 - \frac{c_2}{rK} e^{Kq(t)} \right)$	$\rightarrow \infty^{***}$	$r - c_2/\pi_q$

* Solved numerically from (3.34)

** Solved numerically from (3.55)

*** Truncated at $q_T \ll \infty$ by the numerical integration (3.58)

3.8 Summary

This chapter 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 3.2, only the costless model can be solved in closed form. The others are solved numerically using the parameter values in Section 3.2.

Table 3.1 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 $\dot{\pi}_q/\pi_q$. For the fixed cost model, the initial production level increases with cost. For the marginal cost model, Figure 3.3 shows that initial production level decreases with cost. For the stock cost model, Figure 3.4 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 3.55. 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 3.1 shows that the ending production level trends upward with increasing cost, nearly parallel to the starting production level. Thus, despite the fact that 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.

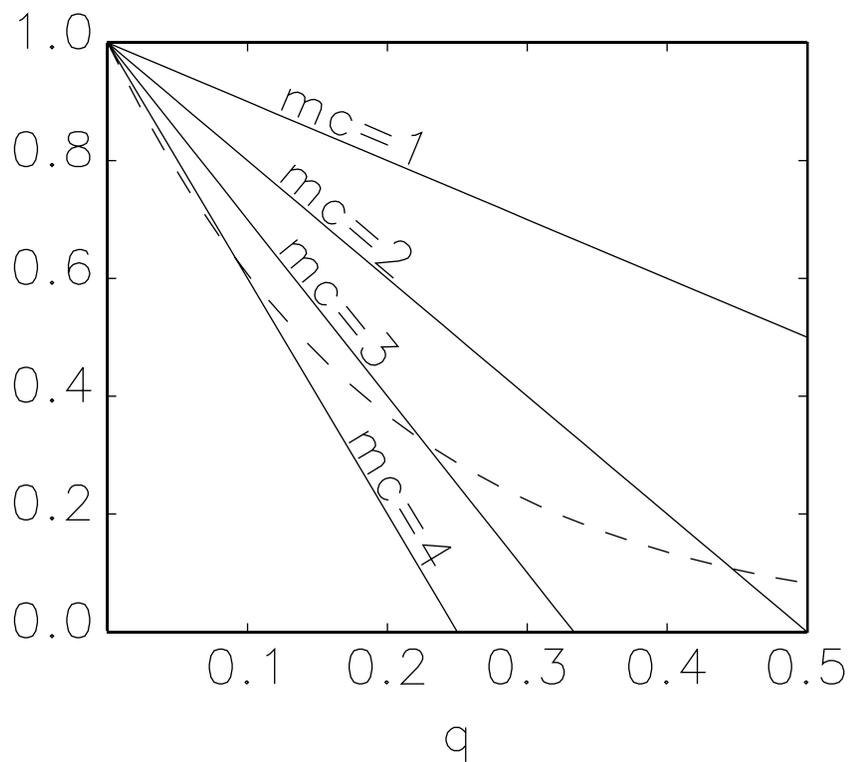


Figure 3.2: The marginal cost model - theoretical values for minimum and maximum production for selected marginal costs.

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.

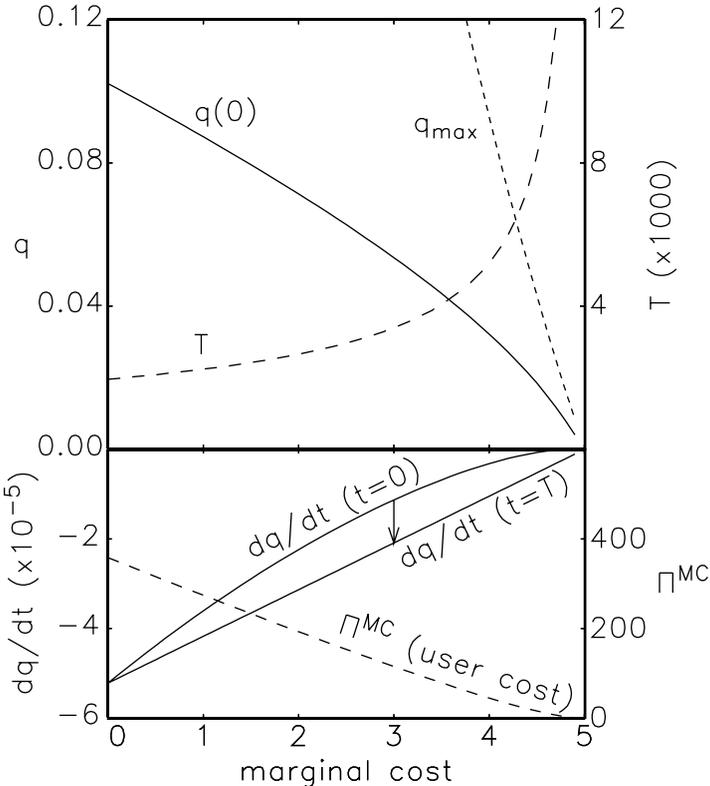


Figure 3.3: Marginal cost model - numerical solutions for terminal time T . T is solved numerically from equation (3.46). Initial production level $q(0)$ is solved using T . Also shown is the production maximum q_{max} computed from equation (3.41) using the equality condition. The lower graph shows 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.

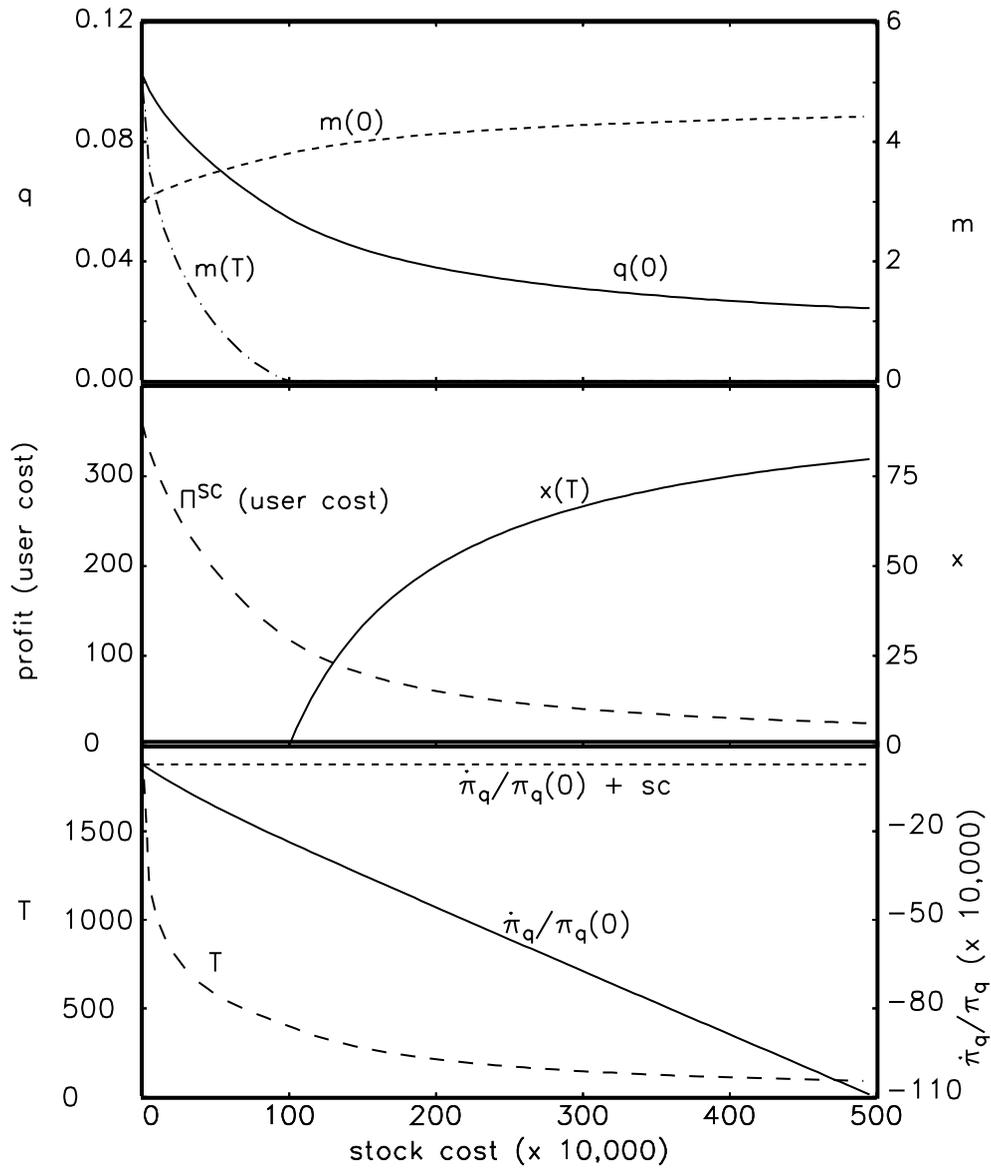


Figure 3.4: Stock cost model - theoretical values for terminal time, initial production level, starting and ending shadow price, ending stock level, producer profit, and user cost as a function of stock cost.

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 profit Π^{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 3.13).

Chapter 4

The Hotelling Agent-based Models

The mathematical models used in this dissertation for Hotelling's Rule are presented in Chapter 3. This chapter will present the agent-based models (ABMs) that correspond to those mathematical models. This will, necessarily, include some discussion of the overall ABM architecture, as well as some discussion of how the models will behave in simulation. The details of the simulations will be discussed in Chapter 5. Chapter 6 will tie together the mathematical models, the ABMs, and the simulation results.

4.1 ABM architecture

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)¹. Though typically implicit in economic models, there may also be exogenous state variables (e.g. GDP) and behaviors (e.g. a supply or demand

¹The term is not defined, but definition can be inferred from the literature. See, for example, Hartley (1996).

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

In an ABM, it can be assumed that an agent is capable of solving the equations in the Chapter 3. This is the basis for the *computational agent* ABMs in Section 4.5. For the monopolist in a computation model, the initial production level is an initial condition and the only required behavioral rule is to decrease production such that marginal profit increases by r percent each time period.

The rest of the ABMs used in this dissertation suppose that the producers 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).

4.2 Simulation architecture

The models are constructed and the simulations run using the MASON² 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

²<http://cs.gmu.edu/~eclab/projects/mason/> (last accessed on 23 May 2011)

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, including monopoly and oligopoly variations of the computation agent model, the adaptive agent costless model, models with different distributions of initial stock, an overt collusion model, and Monte Carlo ensembles of costless, fixed-cost, marginal-cost, stock-cost and ad valorem cost monopolies. These models are described in the following sections. In the oligopoly models, the user is able 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 and initial stock levels. 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 Chapter 5. Images of the interface and results windows are included in the Appendices.

Each model has two agent types: 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 to produce values 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.³ 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 4.1. The size of a simulation time-step is arbitrary, though the default discount rate

³In 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.

is assumed daily and compounds to ten percent per annum. Changing the discount rate, via the GUI, changes the implied time-step.

4.2.1 The market agent

The market agent, called Market, is assigned a demand function and controls any market information provided to the producers. For the computational agent models, this means that the market agent computes the initial production level for each producer. 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.

4.2.2 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 production level for the time period. At the end of the time period, the agent subtracts that production level from the remaining stock, and the market agent collects the production level, remaining stock, profit and costs for the real-time displays. The user may choose to also write these data to a file for post-processing.

Although the producer is the optimizing agent, production level computation is done by the market agent on behalf of the producer agent so that the details can vary depending on the market model. For example, in the computational agent model, the production path is in closed mathematical form, while in the rest of the models, a simple heuristic is used for optimization. The heuristic is described in the following section.

4.3 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 result 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 the present value of total profit over the lifetime of the resource. 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 the estimated present value of total profit begins to fall. This is called the *ramp-up* phase. This phase begins at time zero and is repeated at every time step until the heuristic determines that production is on or near the optimal production path.
2. Estimate the present value of total profit in each time period based on the three candidate strategies: increase production, reduce production, or maintain current production. The strategy that yields the highest present value total profit is implemented. This is the *optimization* phase. The phase begins once the heuristic is on or near the optimal production path, and is repeated every time step for the lifetime of the

resource.

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 introduces a source of error that 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.

4.3.1 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 Chapter 3, 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.

Just as the continuous solutions presented in Chapter 3 would serve as the limits on discrete theoretical models, they provide the limits on the ABM results as well. Furthermore, the continuous solutions provide insights into the qualitative effects of parameters, such as initial stock and costs, that otherwise may be obscure in the discrete theoretical models. Finally, the potential exists in the ABMs to decouple the planning period and the production period, producing daily but planning quarterly, for example. Ultimately, however, the continuous solutions set the theoretical limits on the simulation results.

4.3.2 The heuristic algorithm

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

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

where

- Π_τ = estimated total profit
- q_i = production level in period i
- p_i = price in period i
- r = discount rate
- τ = remaining lifetime of the stock

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

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

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 choice of value for Δq is discussed at the end of this section. 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

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

where p_n is the price in the current period, and

$$\left(\frac{\Delta p}{\Delta q} \right)_{i-1} = \frac{p_{i-1} - p_{i-2}}{q_{i-1} - q_{i-2}} \quad (4.1)$$

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

$$\begin{aligned} \Pi_\tau(\Delta q) &= \sum_{i=0}^{\tau-1} (q_n + i\Delta q) \left[p_n + i\Delta q \left(\frac{\Delta p}{\Delta q} \right)_{i-1} \right] (1+r)^{-i} \\ &= q_n p_n A + \Delta q \left[p_n + q_n \left(\frac{\Delta p}{\Delta q} \right)_{i-1} \right] B + (\Delta q)^2 \left(\frac{\Delta p}{\Delta q} \right)_{i-1} C \end{aligned} \quad (4.2)$$

where

$$A = \sum_{i=0}^{\tau-1} (1+r)^{-i} = \frac{1+r}{r} [1 - (1+r)^{-\tau}] \quad (4.3)$$

$$B = \sum_{i=0}^{\tau-1} i(1+r)^{-i} = \frac{1}{r} \left[A - \frac{\tau}{(1+r)^{\tau-1}} \right] \quad (4.4)$$

$$C = \sum_{i=0}^{\tau-1} i^2(1+r)^{-i} = \frac{1}{r} \left[2B + A - \frac{\tau^2}{(1+r)^{\tau-1}} \right] \quad (4.5)$$

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

$$\begin{aligned} x_n = \sum_{i=0}^{\tau-1} (q_n + i\Delta q) &= \tau q_n + \Delta q \frac{\tau(\tau-1)}{2} \\ \tau &= \frac{-(q_n - \frac{\Delta q}{2}) \pm \sqrt{(q_n - \frac{\Delta q}{2})^2 + 2x_n \Delta q}}{\Delta q} \end{aligned} \quad (4.6)$$

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. For a given current production level q_n and stock level x_n , there is only one Δq^* and one τ that satisfies this. The lifetime in this case is constrained by

$$q_n + \sum_{i=0}^{\tau-1} \Delta q^* = 0 \quad (4.7)$$

and the constraint that total production equal the current stock by

$$\sum_{i=0}^{\tau-1} (q_n + i\Delta q^*) = x_n \quad (4.8)$$

Solving (4.7) for Δq^* and substituting into (4.8)

$$\tau = \frac{2x_n}{q_n} - 1 \quad (4.9)$$

Substituting this back into (4.7),

$$\Delta q^* = -\frac{q_n^2}{2x_n - q_n} \quad (4.10)$$

This is the lowest (most negative) Δq 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 (4.6) be non-negative.

4.3.3 The inverse demand function and its parameters

For all of the ABMs discussed in this dissertation, the inverse demand function from Hotelling and discussed in Section 3.2 is used. The parameter values from Section 3.2 are also used in every simulation. As a consequence, any two graphs throughout this dissertation can be compared to assess the impacts of the respective production technology or market structure. Although the simulations are, by necessity, numerical, it is hoped that this similarity produces results that are qualitatively general, given the demand function.

4.4 Monopoly models

Initially, consistent with Hotelling, costless production is considered. The introduction of nonzero cost will be presented in Section 4.7. In the computational agent models, production decisions are based on optimization constrained by the inverse demand function. That is, the producer agent implements the theoretical optimal control solution. In the adaptive agent models, production decisions use a heuristic that estimates future profit as described in Section 4.3.

The theoretical maximum profit is shown in (3.32). The equivalent summation expression is

$$\Pi_{max} = \sum_{i=0}^{T-1} (1 - e^{-Kq_i}) (1 + r)^{-i} = \frac{1+r}{r} \left[1 - (1+r)^{-T} \right] - \frac{(1+r)^{-T} - e^{-rT}}{e^r (1+r)^{-1} - 1} \quad (4.11)$$

With the values given in Section 3.3, 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 Chapter 5.

4.5 The computational agent model

The computational agent models are used to benchmark the adaptive agent models. That is, these models are only used to show the theoretical optimal control solution as it would be implemented by perfectly optimizing agents. These are not behavioral models *per se*: the agents compute the production path using the solution to the constrained optimization equations. Because they are not adaptive, the computational agent models are extremely sensitive to numerical errors in simulation, as will be shown in Section 5.5.

In the computational agent models, the agent computes the first-period production level from (3.29). The agent has only two behavioral rules:

1. Adjust the production level to increase the marginal profit each period at the interest rate. For the demand function in Section 3.2, this is equivalent to reducing production by r/K (3.31).
2. Maintain the production level above zero until the stock is depleted.

The second rule is necessary because of numerical errors and the discrete nature of agent-based simulation. That is, since the resource-owner is adjusting quantity on a per-period basis (as opposed to continuously), the production level may reach zero while there is stock

remaining. This is prevented by foregoing the reduction in a given period if the amount reduced is greater than or equal to the current production level until there is no stock remaining. Because these simulations have fine granularity (time step size is one day), this will typically only affect the final period. This is not entirely an artifact of simulation, however: end of life considerations may affect real-world production decisions if planning periods are long. In these cases, Hotelling's Rule behavior may not be evident at end of life even if all other conditions are met. Aside from the terminal boundary condition, the production path should be nearly identical to the optimal (3.26) and total profit within numerical error of the theoretical maximum (3.32).

4.6 The adaptive agent models

All of the models except the benchmark computational agent models are adaptive agent models. In the adaptive agent models, the 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:

1. In the ramp-up phase, increase the production level from zero until estimated total profit begins to decrease. The default ramp-up rate is an increase of 0.01 units of production per period.
2. In the optimization phase, in each period, each agent estimates total profit based on the three production strategies, then executes the strategy that maximizes estimated total profit.

4.7 Production technologies with nonzero cost

The production technology models with nonzero cost are the costless adaptive agent models with a nonzero cost term. This does not require a change in the producer agents, which are implemented with cost variables, all of which were zero for the costless adaptive agent model. Since revenue may now become less than cost, the behavioral rules are modified to avoid negative profits. This is discussed in Chapter 5 with regard to specific cost functions.

4.7.1 Fixed cost model

Recall from Section 3.4 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 (4.7) becomes

$$q_n + \sum_{i=0}^{\tau-1} \Delta q^* = q_{min} \quad (4.12)$$

which, when substituted into equation (4.8), means that equation (4.10) becomes

$$\Delta q^* = \frac{q_n^2 - q_{min}^2}{2x_n - q_n + q_{min}} \quad (4.13)$$

This is the largest possible reduction in production that results in a straight-line decreasing production path that reaches q_{min} at the moment the stock is physically depleted. The adaptive agent 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 (4.2) and A is from (4.3).

A fixed cost does not appear in π_q or π_x , so according to equation (3.13), 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 3.1.

4.7.2 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 (4.2) and A and B are from equations (4.3) and (4.4).

4.7.3 Stock cost model

With the addition of a stock cost as in (3.48), the future profit estimate becomes

$$\begin{aligned}
\Pi_{\tau}^{SC} &= \Pi_{\tau}^{NC} - \sum_{i=0}^{\tau-1} c_2 (x_0 - x_i) (1+r)^{-i} \\
&= \Pi_{\tau}^{NC} - c_2 x_0 \sum_{i=0}^{\tau-1} (1+r)^{-i} + c_2 \sum_{i=0}^{\tau-1} x_i (1+r)^{-i} \\
&= \Pi_{\tau}^{NC} - c_3 \sum_{i=0}^{\tau-1} (1+r)^{-i} + c_2 \sum_{i=0}^{\tau-1} (x_n + \Delta x_i) (1+r)^{-i} \tag{4.14}
\end{aligned}$$

where $c_3 \equiv c_2 x_0$. The discrete form of equation (3.6) is $\Delta x = -q$. Assuming that q is changing by the constant increment Δq , then $\Delta x_i = -(q_n + i\Delta q)$. Now, (4.14) 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 (4.2) and A and B are from (4.3) and (4.4).

4.8 Other sources of uncertainty

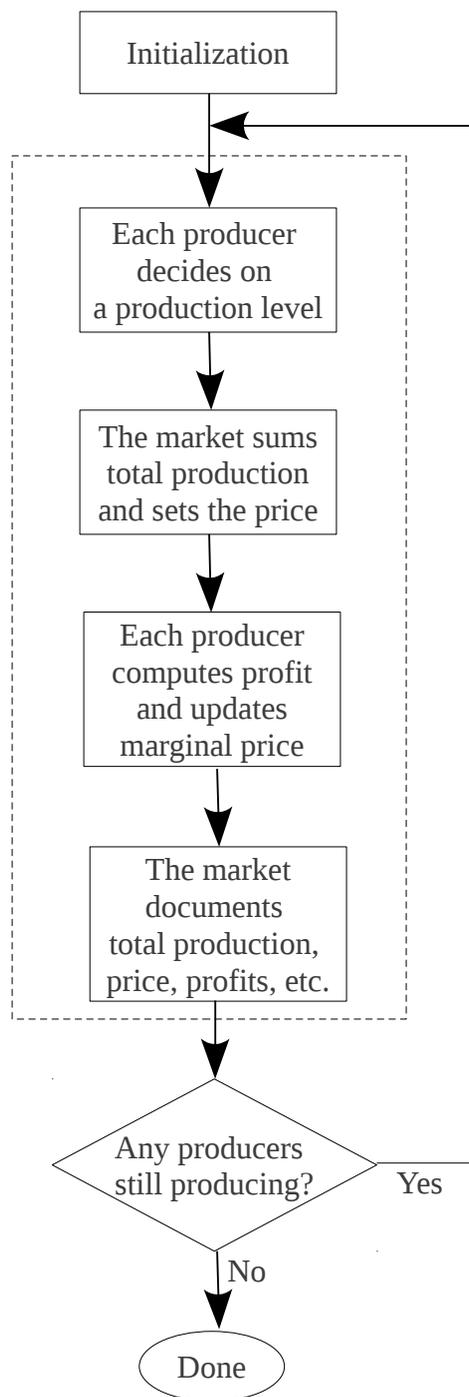
If the producer is not certain of the extent of the resource x_0 , the consequent error in the lifetime of the stock will affect estimates of future profits. This, in turn, may affect the production strategy selected by the heuristic outlined above. Although the heuristic can adjust the rate of change 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 initial stock level, or initial reserves 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 (3.29). 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 5.5.

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. These issues are beyond the scope of this dissertation.

Figure 4.1: 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.



Chapter 5

The Hotelling Simulations

The mathematical models for Hotelling's Rule are presented in Chapter 3. The corresponding ABMs are described in Chapter 4. This chapter will present the simulation results and discuss their characteristics. The discussion of how the simulation results compare with the theoretical models is saved for Chapter 6, which will tie together the mathematical models, the ABMs and the simulation results.

The ABMs are oligopoly models for which the monopoly results are special cases. The oligopoly ABMs 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. In all oligopoly models except the explicit collusion model, 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 models 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. Monte Carlo sampling is preferred to stepping over the range of values in regular intervals, which can produce spurious trends that are artifacts of the size of the interval. Also, by sampling at random intervals, the Monte Carlo technique is less likely to skip over small intervals in which there are unusual outcomes.

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 examine the efficiency of the heuristic by introducing intentional error into the initial production level.

The models in Sections 5.1 and 5.2 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. The object is to compare the erroneous results from the computational agent models and those of the adaptive agent models. The total discounted profit is compared of each model is compared with

5.0.1 The inverse demand function and its parameters

Recall from Chapter 4 that all ABMs use the inverse demand function from Section 3.2. Similarly, the parameter values for all simulations in this chapter are from Section 3.2. The graphs in the chapter can be compared directly, so that, given the demand function, the qualitative differences - in terms of the relative impacts of different technologies and the effect of added producers - can be stressed.

5.1 The computational agent models

In the computational agent model, only the costless model is presented, since the production path cannot be solved in closed form for the production technologies with nonzero cost.

The computational agent model serves two different purposes. The monopoly computational agent model reproduces the theoretically optimal production path and is a benchmark for the adaptive agent models. The oligopoly computational agent models, on the other hand, represent the worst case, meaning the greatest possible error an optimizing agent is likely to make in the presence of competitors. Recall that the producer agents in these models are unaware of each other, and each producer agent sets a production plan as though it were a monopoly market. The computational agents set the initial production level based on there being no other product in the market. This is wrong by a factor of the number of producers in the market. The adaptive agents, however, revise the production plan in each period and, as a result, the production paths converge toward the optimal for the respective market size.

Figure 5.3 shows the time-series results of simulations of the monopoly and oligopoly computational agent models with from two to six producers. The theoretical optimum is shown as the heavy solid line. With only one producer, which is following the theoretically optimum production path, the theoretical and monopoly ABM curves are identical, as expected. All models have identical production paths because all computational agents set the production path based on the optimal control solution for a monopoly. In all cases, the monopoly curve is completely obscured by the theoretical optimum curve.

Output per producer is identical in each model, so total production in the oligopoly models is equal to the production level in the monopoly model times the number of producers. The price curves reflect the higher levels of production as the number of producers increases in each model. The percent change in marginal profit scales with the number of producers, as well. This is a consequence of each producer optimizing under the assumption of being a monopoly. The percent change in marginal profit for the monopoly model is equal to the discount rate, while it is the discount rate times the number of producers in the oligopoly models. Finally, producer profit decreases with the number of producers, reflecting the lower prices with higher total production levels. In essence, these results are showing that, with an incorrect assumption about the market structure, an agent that strictly follows Hotelling's Rule performs very poorly. The following section will show how an adaptive agent, which

appears to follow Hotelling's Rule when the conditions are optimal, is also able to make production path corrections that reduce the penalty for an incorrect assumption of the market structure.

5.2 The adaptive agent models

As discussed in Section 4.6, the adaptive agents use a simple heuristic to select an optimal production path. Figure 5.4 shows the the time-series results from the adaptive agent model corresponding to the computational agent model in Section 5.1. In these models, each producer has the same beginning stock. Unlike the computational agent model, where the production path is fixed for all producers, the heuristic leads to different production paths depending on the number of producers in the market. This is a results of each producer agent using an estimate of the marginal price (4.1) based on assuming that the entire market response was due to that individual's production changes. The production path shows the error due to the coarseness of the ramp-up production increases. The monopoly curve starts somewhat above the Hotelling's Rule curve, and ends somewhat before. The heuristic stopped the ramp-up phase at the first time-stamp in which the decreasing strategy was most profitable, but the large ramp-up increase in production caused it to over-shoot the optimum. Note that the price curves fan out, unlike the convergence to the choke price seen in Figure 5.3. This is a result of the daily Bayesian updates which allow the heuristic to adapt to changes in the market due to all the producers, even though the producers are unaware of each other. The percent change in marginal profit $\Delta\pi_q/\pi_q$ is 3.03×10^{-4} for the monopoly producer, which is somewhat higher than the optimum of 2.61×10^{-4} . This is also an artifact of the coarseness of the ramp-up increments, which place initial production above the optimum, permitting a steeper rate of decrease over the lifetime of the stock. As the ramp-up increment is decreased, the error is reduced somewhat, but persists because of the discrete time steps of the simulation.

The $\Delta\pi_q/\pi_q$ curve in Figure 5.4 reveals that, for $N=4, 5$ and 6 , the producers alter-

Table 5.1: Comparison of the computational agent models and the adaptive agent models ($\Delta\pi_q/\pi_q$ values are multiplied by 10,000).

		computational agent models		adaptive agent models	
	N	Π_{total}	$\Delta\pi_q/\pi_q$	Π_{total}	$\Delta\pi_q/\pi_q$
Theoretical optimum		358.53	2.61	358.53	2.61
Monopoly	1	358.07	2.61	357.82	3.04
Oligopoly	2	306.31	5.22	313.16	2.82
Oligopoly	3	265.42	7.83	283.17	2.71
Oligopoly	4	232.72	10.43	258.51	1.83
Oligopoly	5	206.23	13.04	240.41	1.86
Oligopoly	6	184.51	15.65	225.36	1.89

nated between the decreasing strategy and the constant production strategy for the first approximately 100 days. For the monopoly producer, total discounted profit is within a few hundredths of a percent of the optimum 358.53. Table 5.1 compares the total profit and percent change in marginal profit $\Delta\pi_q/\pi_q$ for the computational agent model and the adaptive agent model. Note that, for N=6, the computational agent model achieves only 51 percent of optimal profit, whereas the adaptive agent model captures nearly 63 percent of optimal profit.

These models illustrate that it is less costly to mistakenly assume monopoly market power under heuristic optimization than when applying Hotelling's Rule explicitly. No real-world producer is likely to be unaware of competitors. Rather, these models serve as a worst case proxy for errors in assumptions about market structure in general. The simulation artifacts in the adaptive agent model are not unlike errors that are likely to happen in reality. Real-world production planners may be subject to uncertainty in stock or interest rates. or stickiness in wages or prices.

5.2.1 A pooled-stock oligopoly model

For the adaptive agent models in Section 5.2, each producer agent is endowed with x_0 initial stock. For example, the total stock in the $N=6$ model is $6x_0$. 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 $x_0/10$. This model is slightly different from the previous in two ways. First, the ramp-up increment is divided by the number of producers, and second, initial stock levels are given a small random variation. The ramp-up increment is scaled because, with an initial stock of x_0 the default increment of 0.01 is small compared with the optimal starting production level of 0.1022. For an initial stock of $x_0/5$, however, an increment of 0.01 is large compared with the starting production level of approximately 0.01. The random variation is introduced to avoid any artifacts due the ratio between total stock and ramp-up increment.

The adaptive agent duopoly production path in Figure 5.5 reflects a collusion-like outcome, as do models for $N=3$ and $N=4$. That is, the adaptive producer agents arrive at a collusive market structure using only the optimization heuristic. For an oligopoly of five producers, however, something completely different occurs, as seen in Figure 5.5. 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. These sum to 100, as in all the pooled-stock models. The five firms begin the ramp-up phases together, and all five reach the optimal production path after 53 iterations. 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, as shown by the small steps at the beginning of the percent change in marginal profit curves in Figure 5.4.

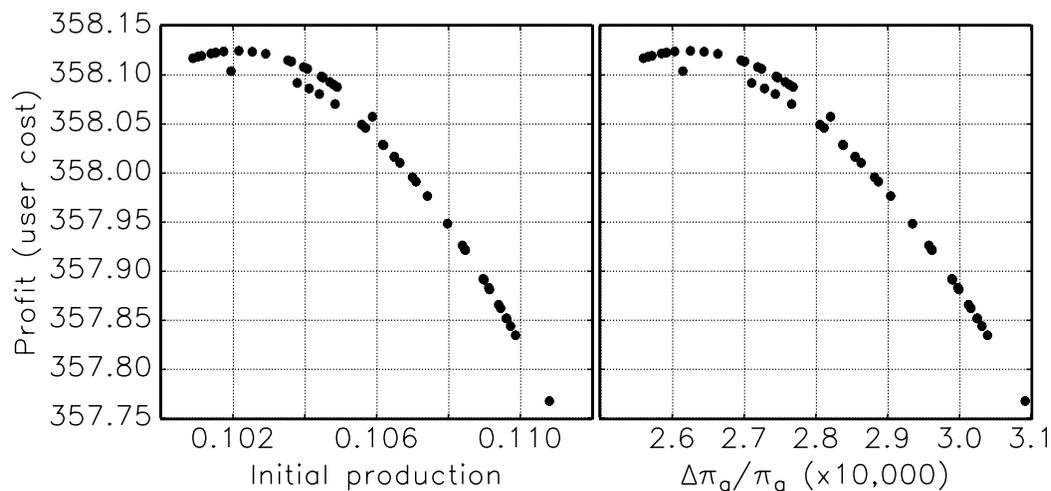
This shows that the estimate of total profit is sensitive to small differences in total reserves for this particular distribution of stocks. The smaller producers find the level production path optimal because the decreases in production by the two largest producers cause price increases that are sufficient for the remaining producers to estimate increasing profits at constant production levels. 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 adaptive agent models with from one to ten producers are shown in Figure 5.6. For a single producer, this model is effectively identical to the monopoly producer in Figure 5.4. 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 outcome, while the rest show the collusion 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 5.6, the discontinuities in the percent change in marginal profit curves where producers change strategies at the end stock life.

5.2.2 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 5.1 shows the simulation results from fifty Monte Carlo samples of the adaptive agent 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

Figure 5.1: Monte Carlo samples of heuristic outcomes.



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 5.1 also illustrates the impact of error in the estimate of x_0 . According to equation (3.29), a one percent error in x_0 translates into a one-half percent error in initial production q_0 , and from equation (3.32), that, in turn, translates into an error of 0.84 percent in the theoretical maximum profit. For the heuristic, however, Figure 5.1 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.

5.3 A collusion model

The collusion model is an extension of the pooled-stock model in Section 5.2.1. This is the only one in which the producer agents communicate with each other. The initial stock is divided among the producers with small random fluctuations, as discussed in Section 5.2.1. The producer with the greatest stock is considered the market leader. The market leader has the same behaviors as the producers in Section 5.2.1, assuming that it is in a monopoly market. At the point in a time-step when the producers decide on a production level (see Figure 4.1), all the follower producers mimic the leader's production changes, scaled by their relative stock levels. See Figure 5.7.

This model does not address issues associated with cheating in a collusive market, but cheating is a simple extension of the current model. The ability for producers to cheat would be introduced at the point of the calculation to scale the leader's production changes. For a cheating follower, this could be a small increase above the scaled production share. For a cheating leader, it would be a matter of the leader communicating one production level change and implementing another.

5.4 Nonzero-cost models

Three cost functions are examined: a fixed (per day) cost, a marginal (per unit) cost, and a stock (cumulative production) cost. The possibility that cost may exceed revenue imposes the following expansions of the behavioral rules (changes in *italic*):

1. In the ramp-up phase, increase the production level from zero until estimated total

profit *is positive and* begins to decrease.

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 profit becomes negative.*

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 3.2 there are no closed-form optimal control solutions when marginal or cumulative costs are included. The adaptive agent 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 (3.14). For the stock cost models, percent change in marginal profit should approach equation (3.13).

5.4.1 Fixed cost model

The fixed cost model is the pooled-stock adaptive agent 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 (3.19), which is now revenue rather than profit. Time-series plots for fifty Monte Carlo monopoly simulations are shown in Figure 5.8. 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 (3.33). This is the envelope for optimal terminal production levels.

The results of the Monte Carlo monopoly simulations are summarized in Figure 5.11. 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 (3.33). 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 3.4.

5.4.2 Marginal cost model

The marginal cost model is the costless model with a constant marginal cost $c_1 \sim N(1.0, 0.110889)$ truncated such that $c_1 \geq 0$. The time-series data for the marginal cost model are shown in Figure 5.9. 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 5.12. 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.

5.4.3 Stock cost model

The stock cost model is the adaptive agent model with pooled stock. 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 3.6.

Before Monte Carlo sampling the parameter space, it is instructive to examine the classes of behavior anticipated. Figure 5.10 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 5.2.1. 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 3.6 in the discussion of Figure 3.4. 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 3.6, 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 (5.13) 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 3.6.

5.5 Theoretical efficiency

To put the foregoing results into perspective, Figure 5.2 shows the effect of error on the theoretical costless model. These are families of curves of total net profit from the theoretical costless monopoly model in Section (3.3). 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 or minus twenty percent in initial production and plus or minus ten percent in percent change in marginal profit.

Figure 5.2: Errors in the computational agent model.

The optimal initial production level and optimal $\hat{\pi}_q/\pi_q$ are indicated with an inverted triangle.

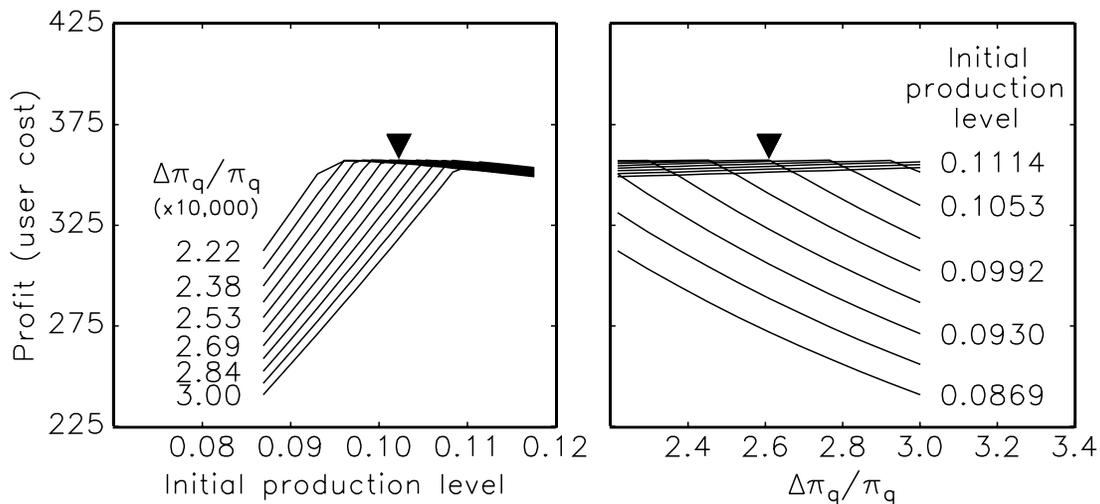


Figure 5.3: Computational agent models - time-series. See the discussion in Section 4.5.

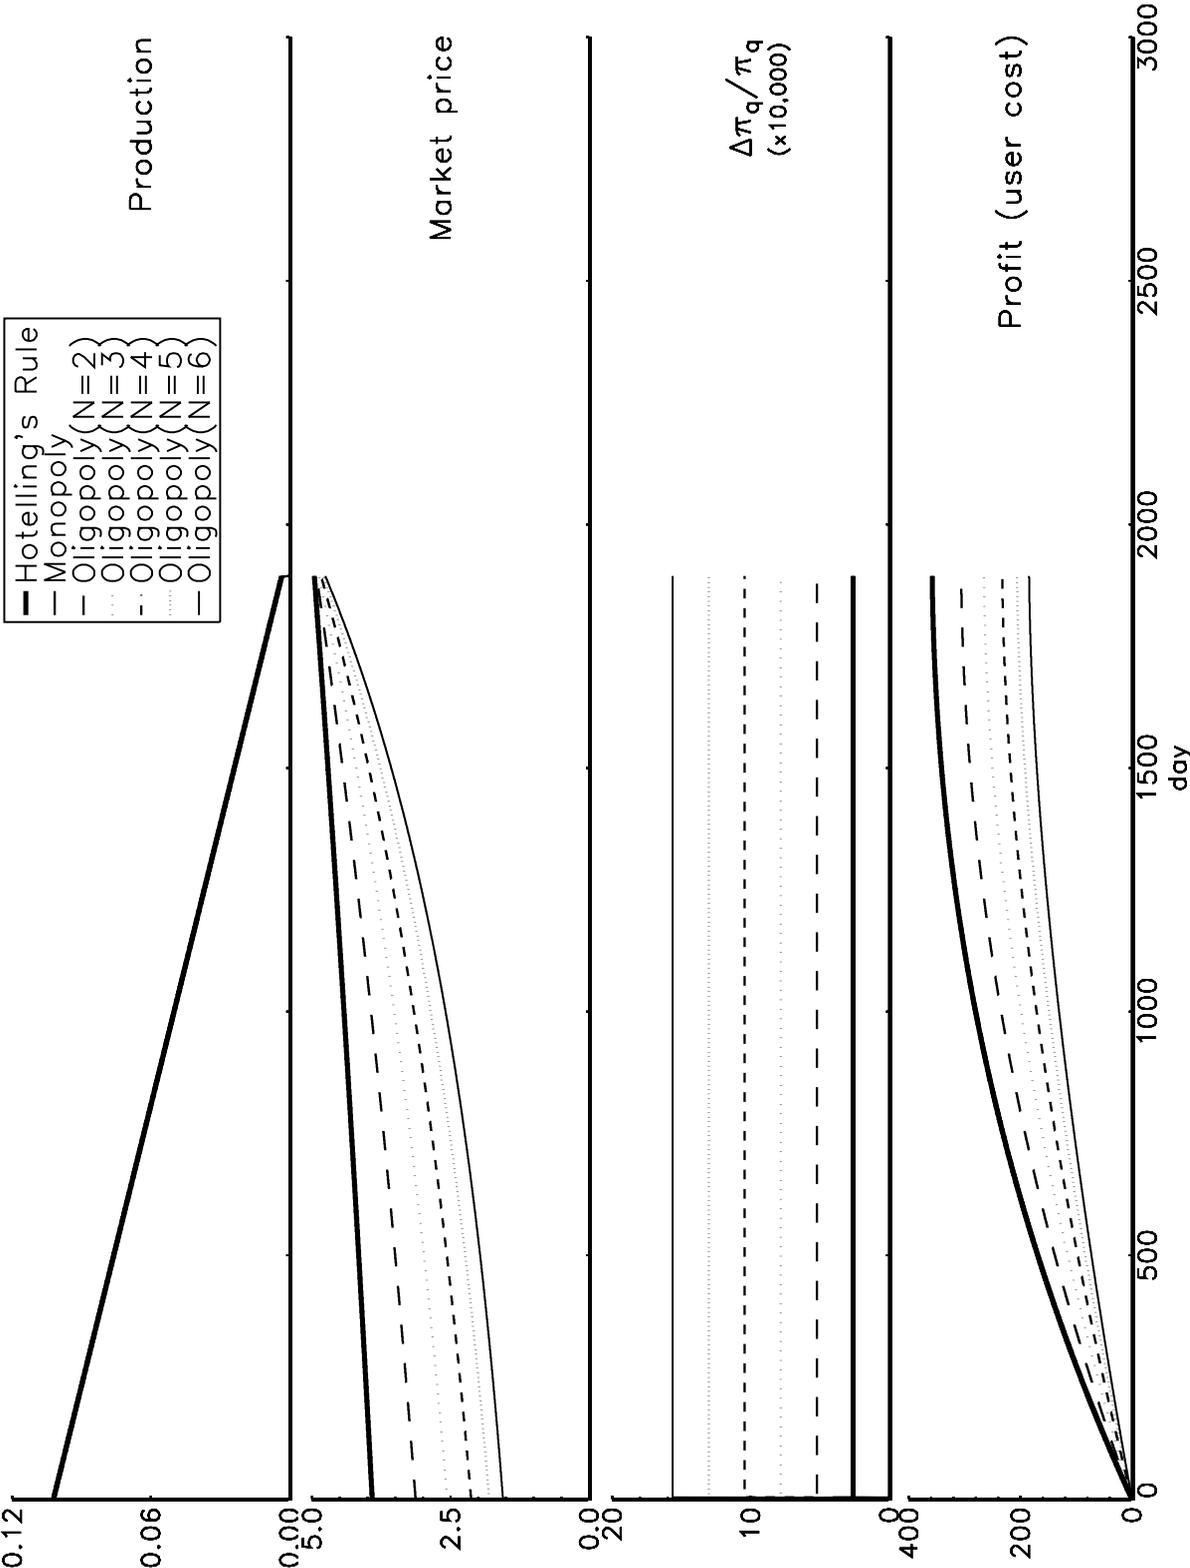


Figure 5.4: Adaptive identical stock models - time-series. See the discussion in Section 5.2.

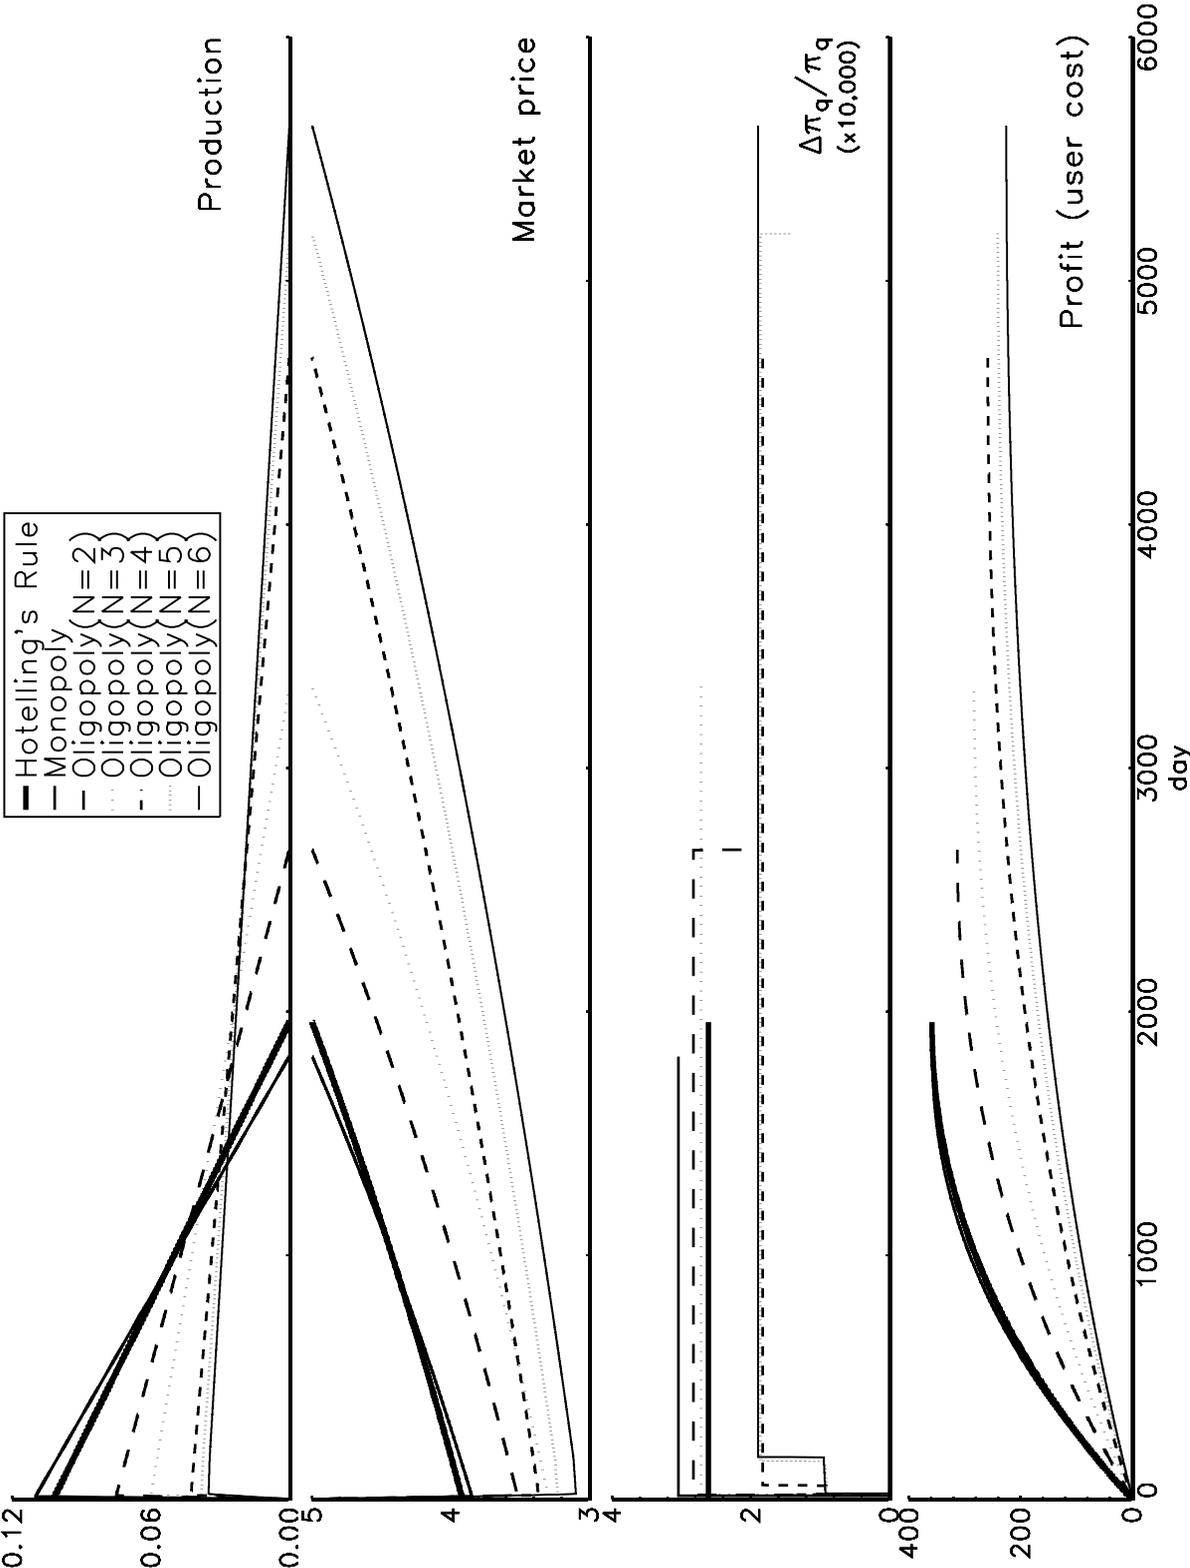


Figure 5-5: Adaptive agent pooled stock models - time-series. See the discussion in Section 5.2.1.

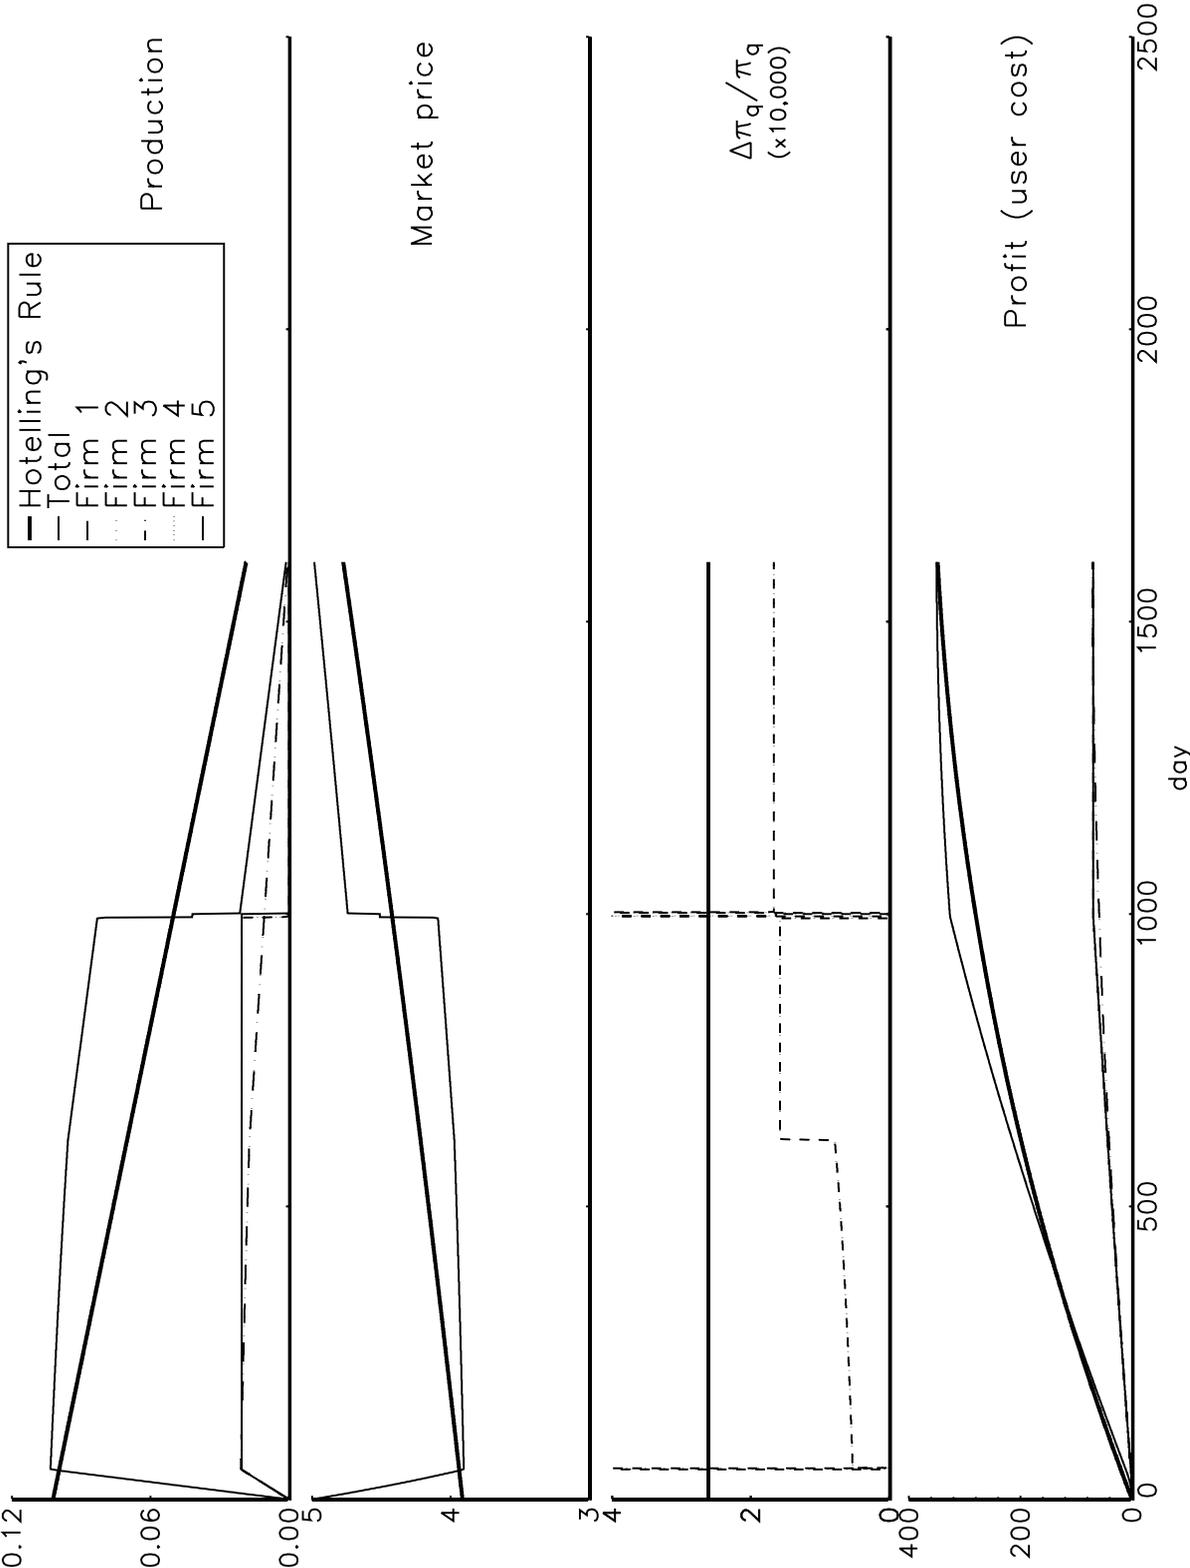
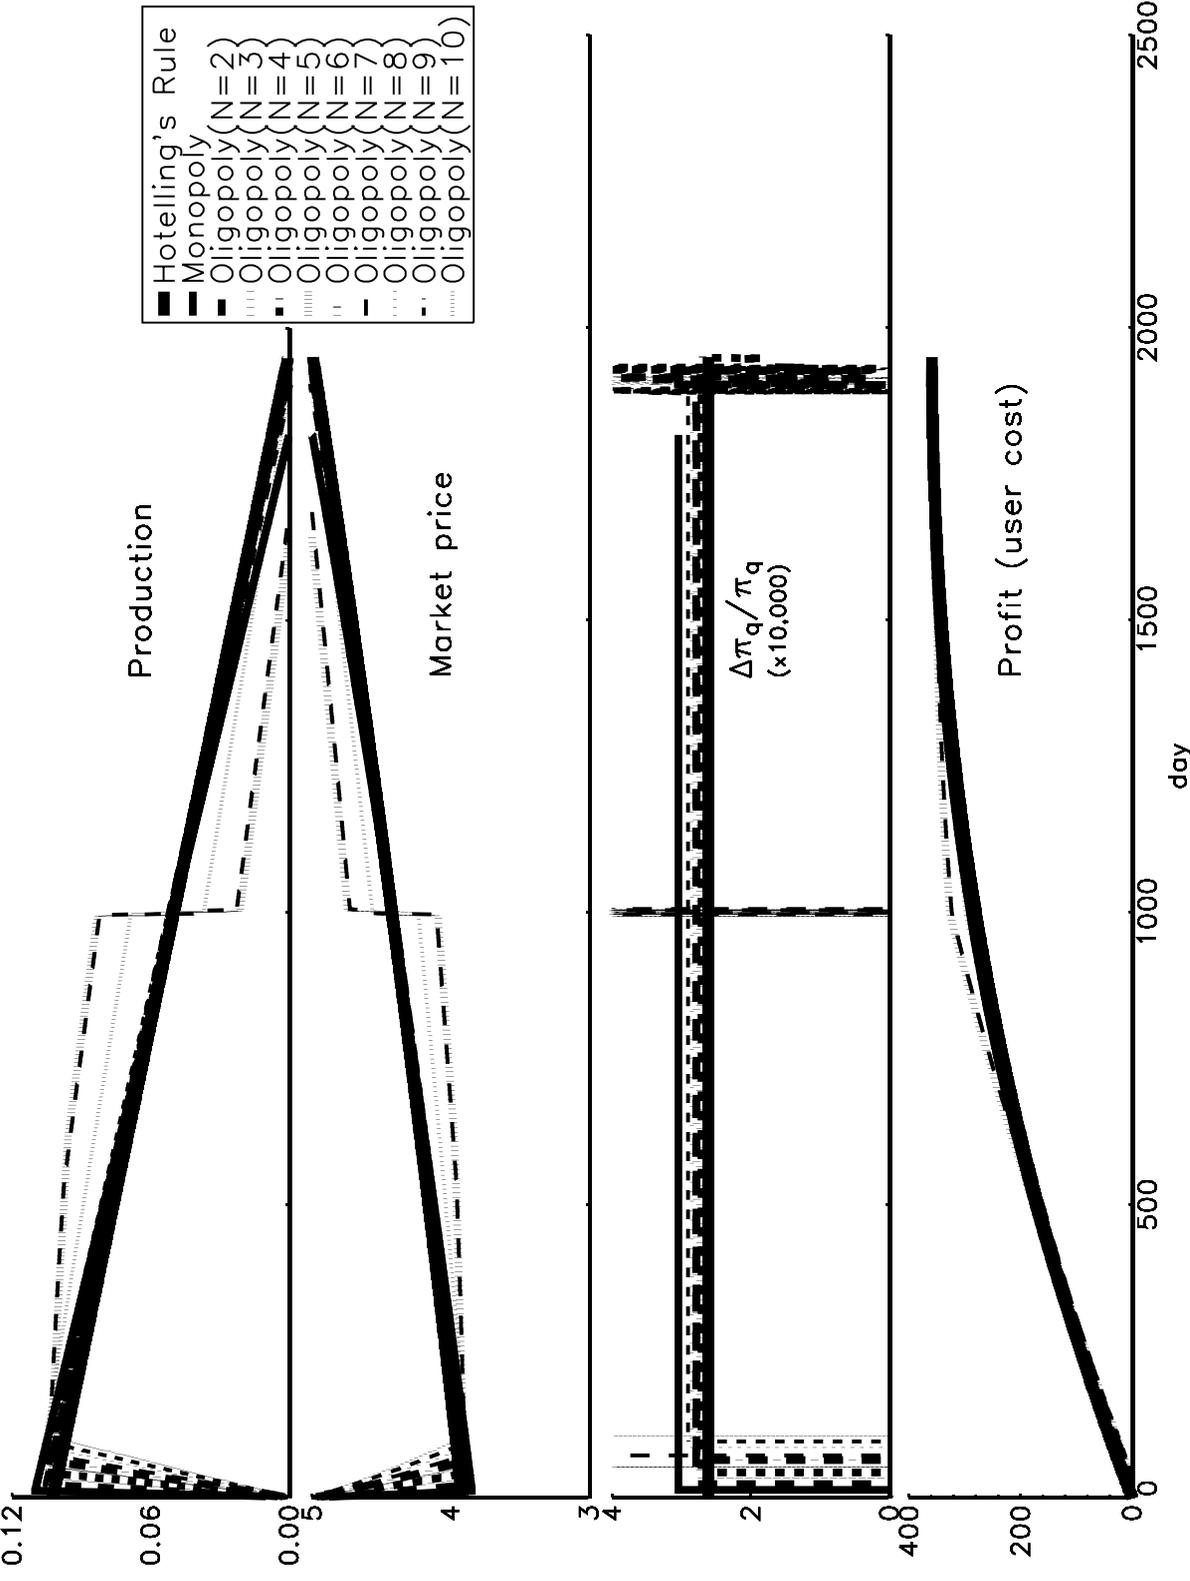


Figure 5.6: Adaptive agent conserved stock models - 10 producers. See the caption on Figure 5.5 and Section 5.2.1.



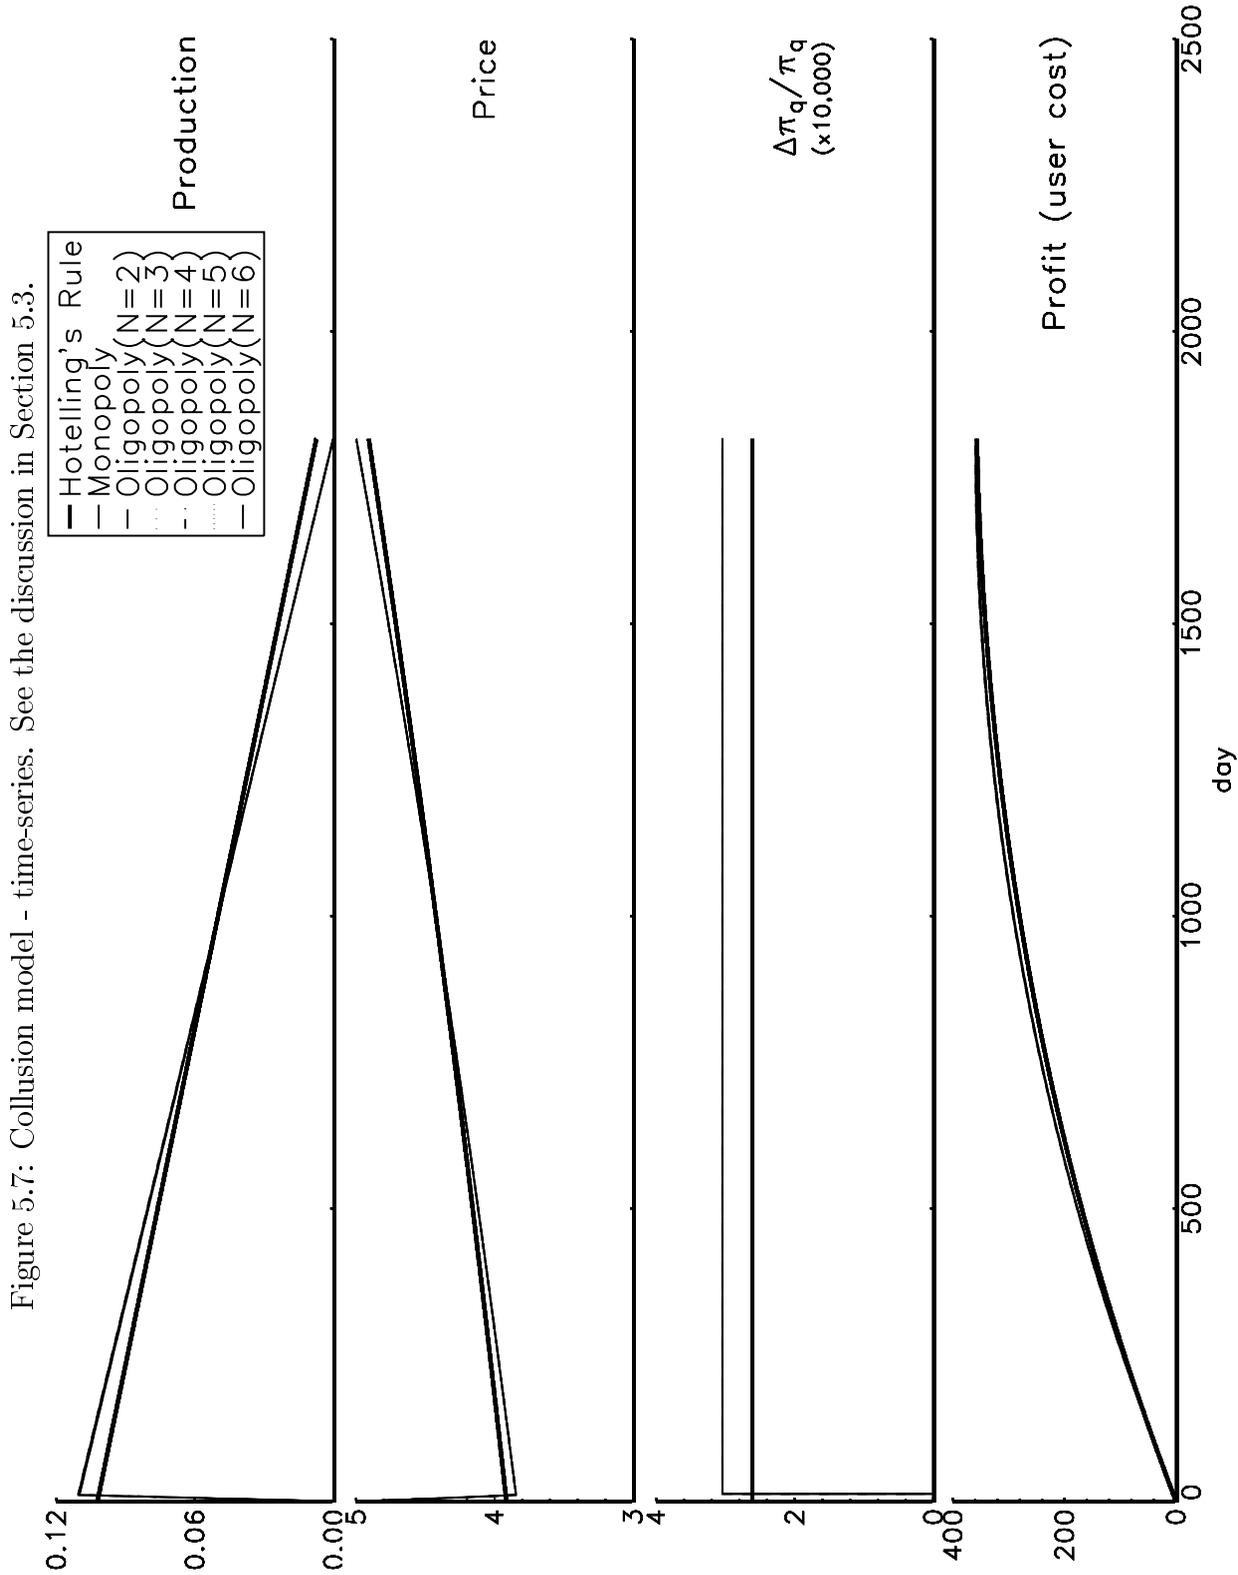


Figure 5.8: Fixed cost model - time-series. See the discussion in Section 5.4.1.

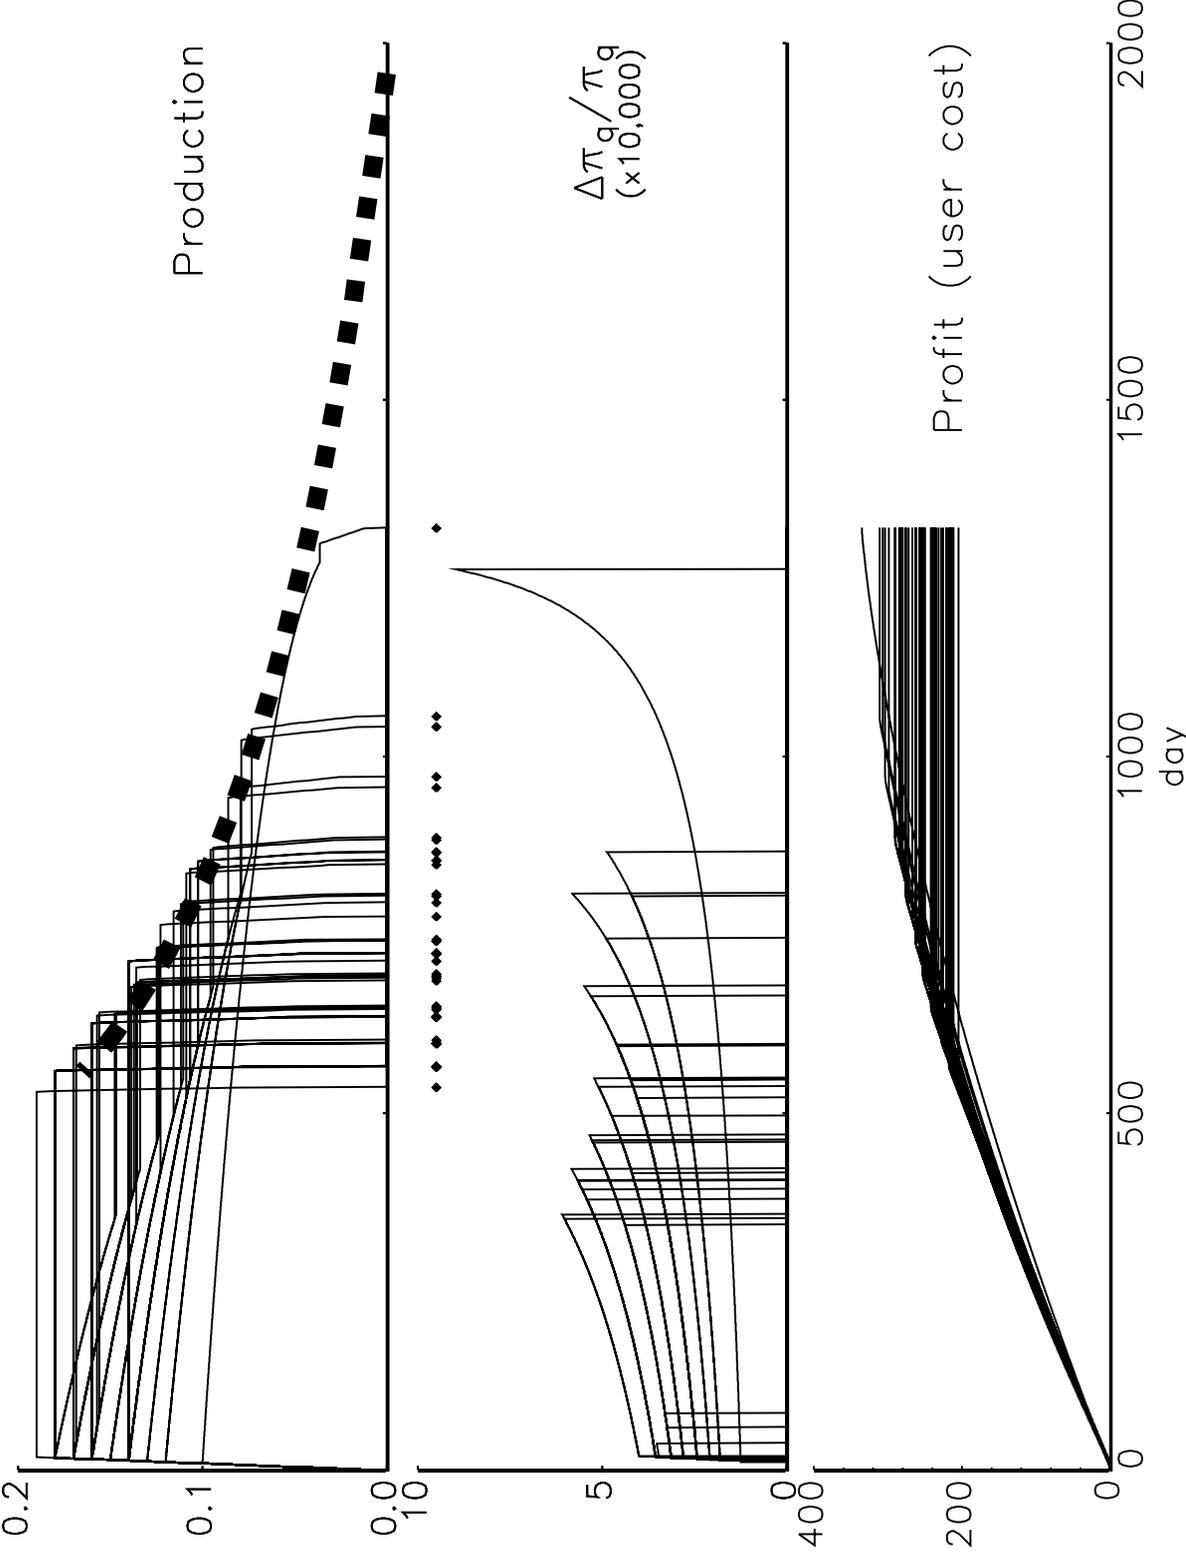


Figure 5.9: Marginal cost model - time-series. See the discussion in Section 5.4.2.

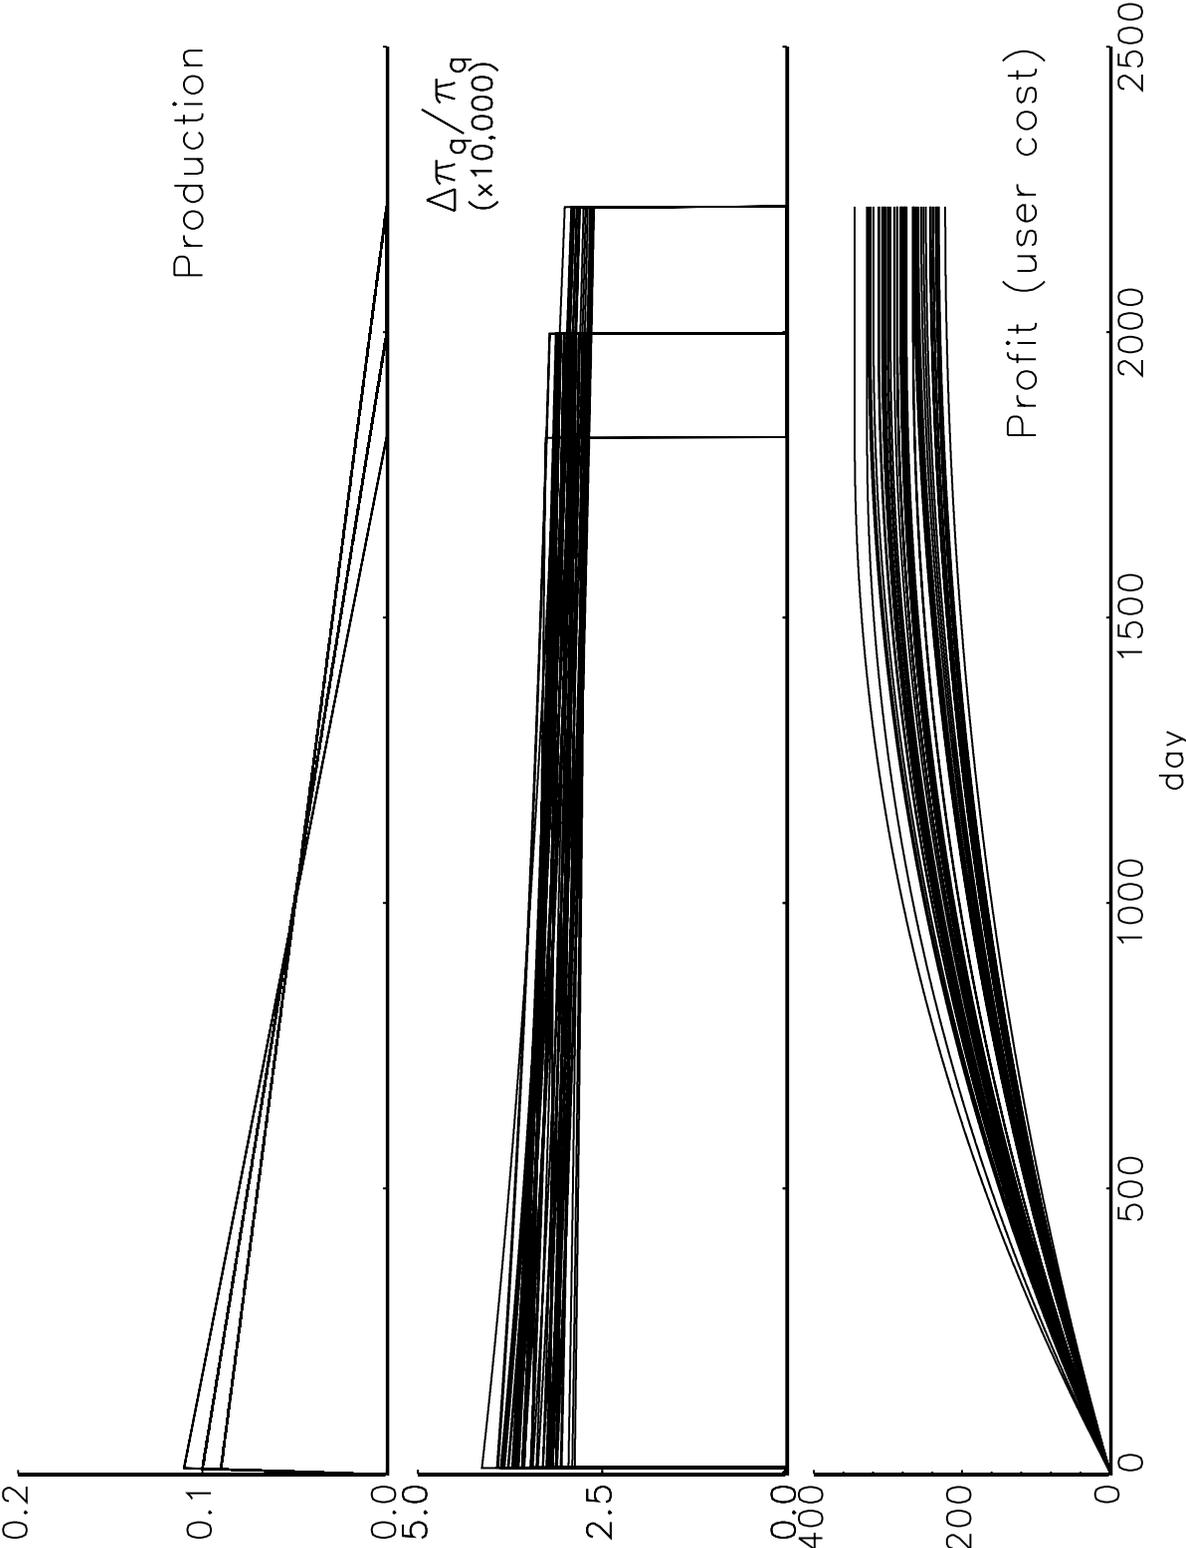


Figure 5.10: Stock cost model - time-series. See the discussion in Section 5.4.3.

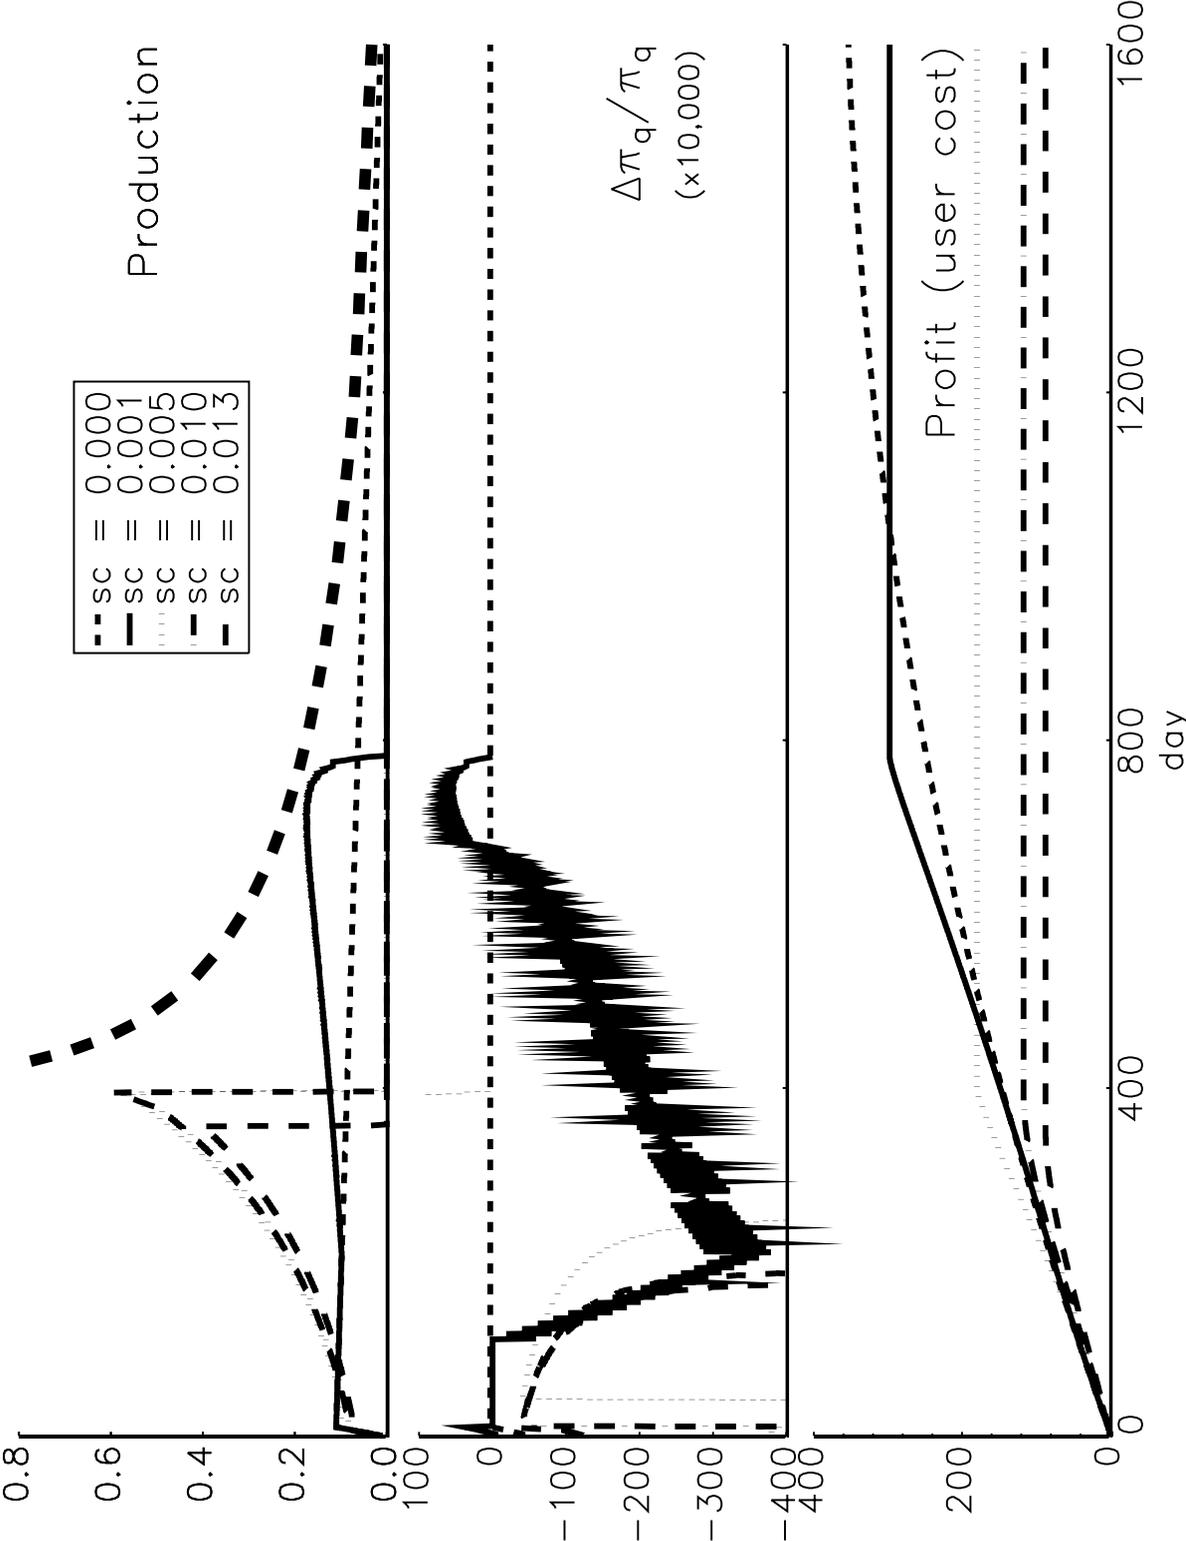


Figure 5.11: The fixed cost model - outcome space.

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).

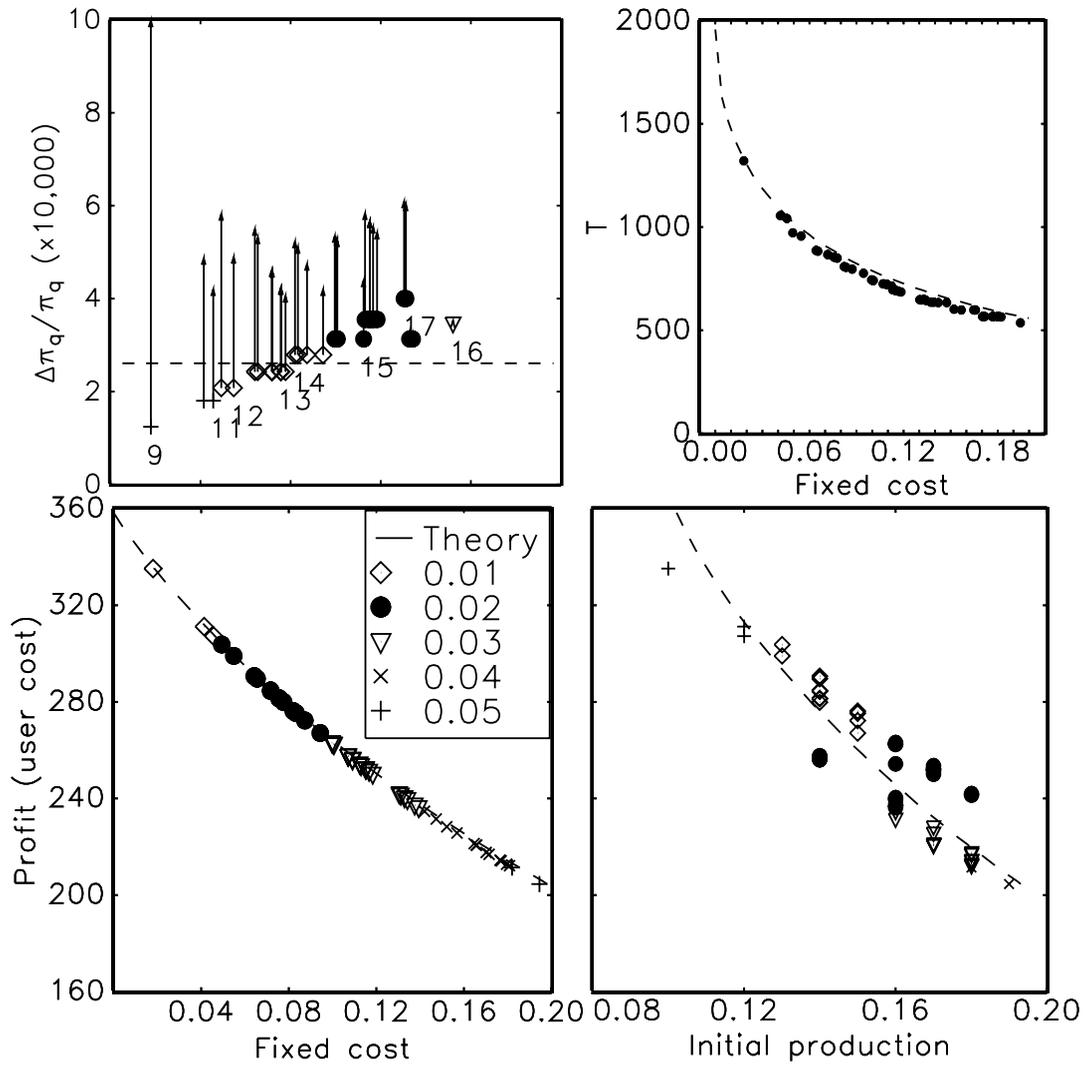


Figure 5.12: Marginal cost model - outcome space. Marginal cost models (constant unit cost). The arrows depict the trajectory in the course of the simulation, starting at the X and ending at the arrowhead. The dashed lines indicate the theoretical optima.

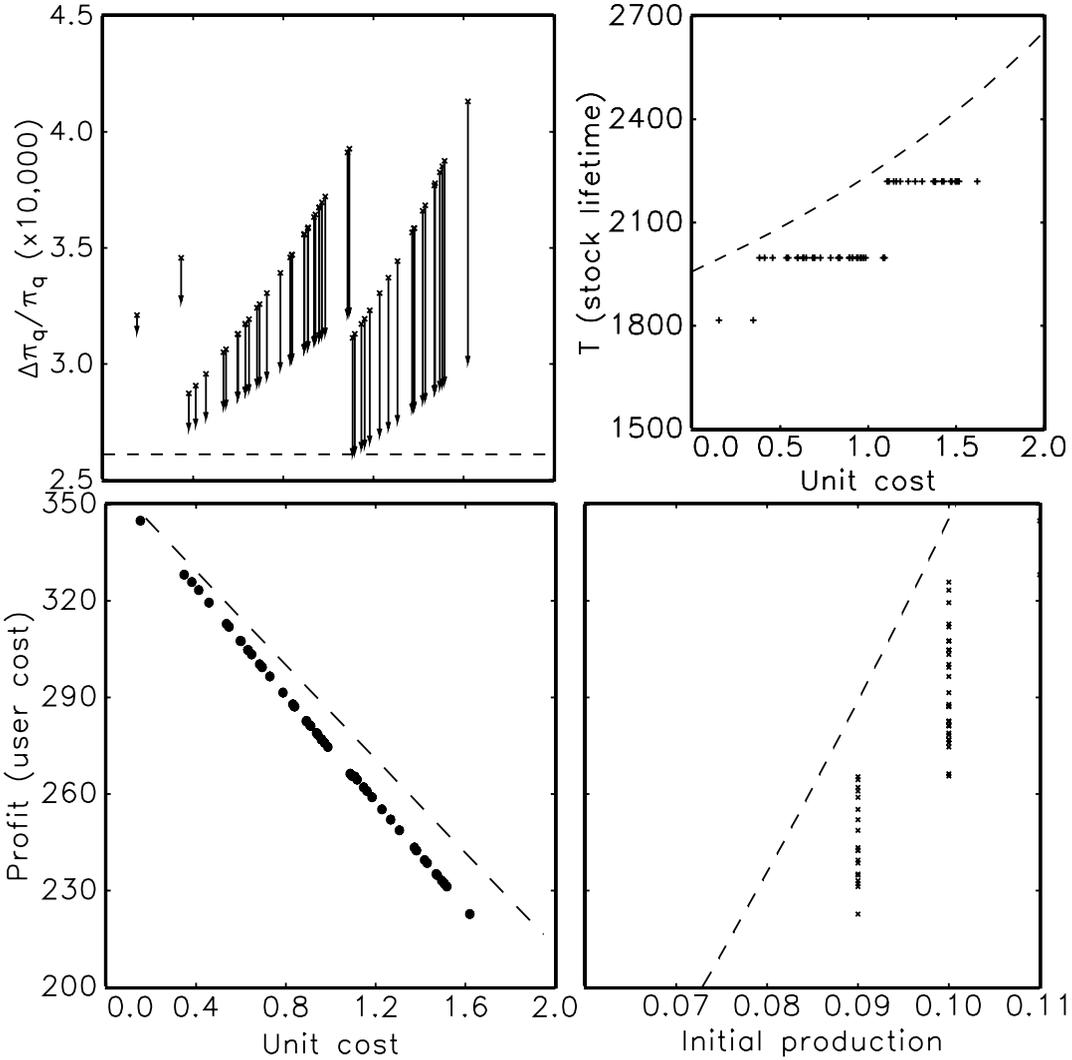
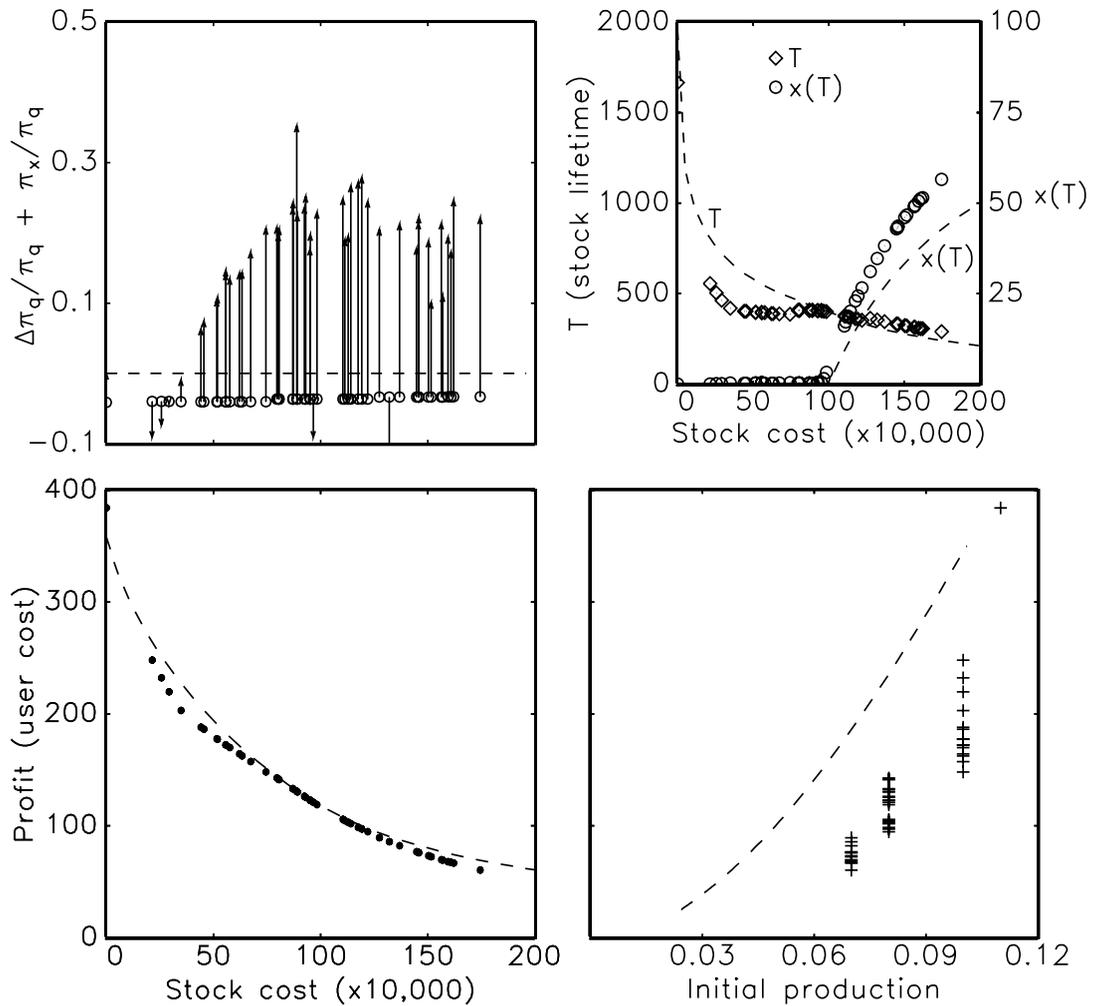


Figure 5.13: Stock cost model - outcome space.

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.



Chapter 6

Analysis of the Simulation Results

The mathematical models for Hotelling's Rule are presented in Chapter 3. The corresponding ABMs are described in Chapter 4. Chapter 5 presents the simulation results and their characteristics. This chapter will attempt to tie together the mathematical models, the ABMs and the simulation results.

The theoretical model developed in Chapter 3 shows that Hotelling's Rule - increasing percent change in marginal profit by r -percent each period - is incontrovertibly the production path to maximize present value profit given the assumptions of the model. The import of those assumptions is, ultimately, the subject of the ABM results in Chapter 5. In general, however, it is observed in the ABM results that the returns on effort to optimize are small, as discussed in Section 5.5, which is pursued further in Section 6.1. That is, although the r -percent rule is the very best strategy to maximize profit, other strategies may produce results that look as good from a reasonable distance. Support for this observation is given in the following sections discussing the individual models.

The ABM optimizing heuristic is described in detail in Section 4.3, but it can be summarized in the following way. The program makes a day-by-day decision to continue, increase, or reduce the production level based on a simple estimate of total profit. The only knowledge of how the market will react to those changes is an estimate based on previous days: how

much the price went down when production was increased, and *vice versa*. A heuristic is often referred to as a *rule of thumb*. Heuristics can be very sophisticated, but this one is not. However, as the results show, it doesn't need to be. The results consistently show, however, an artifact - or fingerprint - of the heuristic. Because production changes are made in steps rather than continuously, the heuristic typically finds production paths that are not-quite-optimal. This characteristic is used to examine the sensitivity of the outcome to the step size.

6.1 The efficacy of the heuristic

The inverse demand function in Section 3.2 is the basic demand function used by Hotelling to demonstrate the r-percent rule for an optimizing monopolist. As a costless model it is appealing because the profit (3.19) and marginal profit (3.20) functions are straightforward to derive and simple to manipulate. There is no static optimum for a monopolist facing the full market demand. That is, there is no finite optimal production path without the constraints imposed in dynamic optimization.

The solution to the dynamic constraint $m(t)$ in (3.23) is called the shadow price. It represents the price (or opportunity cost) of the next unit to be extracted (or which would be extracted if one remained, in the case of completely exhausted stock). For a natural resource, the shadow price is often called the user cost. The shadow price and optimal production level are inversely linked: in the optimal solution for the costless model, as the stock is exhausted, the shadow price edges up to the choke price, which pushes the optimal production level down toward zero. For all but some of the stock cost models, production level and shadow price each reach their final value just as the stock is depleted.

For reasons explained in Chapter 2, the percent change in marginal profit is the central characteristic of interest in the Hotelling models. The simulation results presented in Chapter 5, however, reveal that the optimizing producer behaviors outlined by Hotelling are far more

important than the r -percent outcome. The evidence for this is presented in the following paragraphs.

For the costless monopoly model, the optimal production level decreases at a constant rate, which is r/K , the discount rate divided by the choke price. Consider only the monopoly outcome, the thin solid line in Figure 5.4, and the Hotelling's Rule theoretical optimum, which is the thick solid line. The ABM production path starts slightly higher and is a little steeper than optimum. Because of the size of the increments in the ramp-up phase, the heuristic will not, in general, exactly match the optimal production path. The effect on total producer profit (or user cost) is very slight - the thinner ABM curve in the bottom plot of Figure 5.4 is barely distinguishable from the thick Hotelling optimum.

The steeper production path selected by the heuristic results in a higher percent change in marginal profit,¹ as illustrated in the third plot of Figure 5.4. Note that although there is only a slight difference between the heuristic and Hotelling in the production path, resulting in a nearly imperceptibly lower producer profit, the percent change in marginal profit for the heuristic is nearly 20 percent higher. This is the first evidence that a deviation from Hotelling's Rule of the percent change in marginal profit is not a reliable indicator that a producer is not operating optimally.

While it is correct to say that the error introduced by the size of production increments is an artifact of simulation, it also reflects the kinds of errors made in real-world production. A production planner will have incomplete or uncertain information on the demand function, future discount rates, and the extent of the resource. Additionally, factor inputs or output requirements may be sticky - they can't be changed infinitesimally or on short notice.

To understand how well the heuristic performs in the absence of these artifacts, examine Figure 5.1. For this plot, the basic heuristic is not changed - the increments are still large - but the base value is nudged by a small amount that is randomly distributed. This means

¹The percent change in marginal profit is labeled $\Delta\pi_q/\pi_q$ in the graph as a reminder that this is calculated from discrete changes in the simulation and is not a differential with respect to time.

that, within 50 Monte Carlo samples, some production paths are likely to be nearly identical to the Hotelling's Rule path, and others off by more than the error previously described in Figure 5.4. The samples in which the initial production level are closest to the theoretical optimum of 0.1022 result in the highest producer profit in Figure 5.1. That peak corresponds to a percent change in marginal profit of 0.000261, also the theoretical optimum.

6.2 The penalty for mistaking the number of competitors

The theoretical Hotelling's Rule curves in the adaptive agent model shown in Figure 5.4 are identical to the monopoly curves in the computational agent models in Figure 5.3. Aside from the monopoly ABM curve in 5.4, the remaining curves are generated from models with more than one producer. In both figures, the computational agent model and the adaptive agent model, these curves reflect the behavior of producers who behave as though it is a monopoly market even though it is not. For the computational agent model, that means reducing the production level by r/K each period. For the adaptive agent model, it means a producer assumes that only its change in production results in observed price changes.

Table 5.1 makes it clear that the error of assuming a monopoly market when there are competitors is more costly in the computational agent model. With six competitors, the monopoly assumption reduces total profit by about 18 percent. The percent change in marginal profit for the adaptive agent model is between 0.0002 and 0.0003, while for the computational agent model it ranges from 0.00261 to nearly 0.0015. The adaptive behavior of the heuristic is able to compensate somewhat for a mistake in assuming the size of the market.

While it is not likely that a real-world producer will mistakenly assume a monopoly market, errors in assumptions about the market structure are likely. These plots quantify the worst case consequences of an error in estimating market structure. It is also evident that daily Bayesian updates by the heuristic can compensate somewhat for the error.

6.3 The emergent oligopoly models

The theory of oligopoly models is described qualitatively in Section 3.7. Three different approaches to an oligopoly market are taken in the ABMs: adaptive with identical stock, adaptive with conserved stock, and adaptive collusion.

In the adaptive identical stock model, each producer treats the market as though it were a monopoly, and each producer has an initial stock of 100. That is, for the duopoly model, the total market stock is 200, while the total stock is 500 for five producers. This model is discussed in the Sections 6.1 and 6.2, and serves primarily as a comparison with the computational agent model in Figure 5.3. In the computational agent model, each producer also assumes a monopoly and blindly follows Hotelling's Rule by reducing production by r/K each period, which is why all production paths overlap in the top plot of Figure 5.3. Because total production increases with the number of producers, however, prices are lower as the number of producers increases, as seen in the second plot. For N producers, the reduction in total production per day is Nr/K , so percent change in marginal profit increases with the number of producers, as seen in the third plot. Finally, the bottom plot shows that profit per producer decreases as the number of producers increases. By comparison, Figure 5.4 shows that the heuristic in the adaptive identical stock model is able to adjust somewhat, so that percent change in marginal profit does not scale by the number of producers, as seen in the plot, and profit per producer is not quite as low as for the computational agent model. As discussed in Section 6.2, the mistake of assuming monopoly powers in an oligopoly market is costly if one optimizes analytically, but much less so if one optimizes heuristically.

In order to make a direct comparison between the production paths of monopoly and oligopoly markets, the adaptive conserved-stock model maintains a total market stock of 100 that is divided roughly equally among the producers. Thus, in the duopoly, each producer begins with a stock of approximately 50, while in the five-producer model the initial stock is about 20 for each producer. The division of initial stock is given a very small random variation so that the individual producers are distinguishable. This model led to two different

outcomes, depending on the number of producers. This is reflected in two different curve shapes in the upper plot of Figure 5.5, which shows individual outputs for an oligopoly of five producers. The nearly identical diagonal lines give the appearance that some producers are colluding to share monopoly profits. The other curves, however, show that some producers elected to maintain nearly constant production levels until their stocks were physically depleted. Figure 5.6 shows that this is not unique for five producers, it is also seen with seven and with ten producers. Each of these outcomes is surprising. Figure 5.6 shows that, for five, seven and ten producers.

The collusive outcome is surprising because of the simple nature of the heuristic. The producers don't communicate and, therefore, cannot collude. The only evidence one producer has that there are others is in the observed marginal price. That is, a producer assumes that its own change in production is solely responsible for any change in price. This results in a distorted sense of the demand function which leads to the monopoly outcome only because all of the producers are making the same decisions at the same time.

The non-collusion outcome is the unanticipated consequence of the randomized distribution of stock levels. The producers with the largest stock reach the time to begin reducing production before the smaller producers. This happens because the target reduction rate (4.10) is inversely proportional to current stock level. The profitable production level is slightly lower for higher stock levels. For some divisions of the total stock, that small difference is enough to let the larger producers switch from the ramp-up phase to the optimizing phase one time step sooner than the the smaller producers.

When the larger producers begin reducing their production levels, this causes an increase in the market price, so the smaller producers see an increase in their profit without reducing their own production levels. This only occurs in the models for five, seven and ten producers, in which the variation in initial stock levels was great enough to lead to a timing difference among producers. The drop at about Day 1000 is when the producers that did not reduce their production levels physically depleted their stocks.

There are three things that suggest that this is a Cournot-Nash outcome:

1. The producers optimize production levels based on the production decisions of others
2. Once the producers have elected their strategies, for producers that decide to decrease production, changing to a strategy of level production would reduce estimated profits. For the producers that decide *not* to change production, decreasing production would reduce estimated profits. In this way, a given distribution of strategies is a Nash equilibrium.
3. Total production is greater than for a monopoly or collusion market, but less than perfect competition. Similarly, total profit is less than a monopoly, but more than perfect competition.

6.4 Overt collusion is indistinguishable from optimization

In contrast to the preceding model, in the collusion model the producers do communicate. This model also distributes the total stock evenly but with random variations. The producer with the largest stock is then considered the market leader (see Section 5.3). The market leader optimizes using the heuristic, and all the rest of the producers do whatever the leader does, scaled by their relative stock levels. That is, if the leader lowers production by 0.01, another producer with a stock that is 90 percent of the leader's will lower production by 0.009. This approach results in lock-step production identical to a monopoly as shown in Figure 5.7. Again, the heuristic tends to find an initial production level slightly higher than theoretically optimal (the thick line in Figure 5.7) resulting in a production path slightly steeper than optimal. Qualitatively, the overt collusion production path in Figure 5.7 is indistinguishable from the collusion-like outcomes (the unbroken diagonal lines) in Figure

Figure 5.6. Based on production path alone, it is not possible to discern intentional collusion from summed individual profit maximization.

6.5 The nonzero cost models

The performance of the ABM heuristic in the face of nonzero costs varies widely based on the cost structure. For a fixed cost, the accuracy of the heuristic is very close to Hotelling's Rule. For a marginal cost, the heuristic has decreasing accuracy with increasing cost. For the stock cost model, the accuracy varies considerably over the cost range.

6.5.1 The fixed cost model

For the fixed cost model, the theoretically optimal production path (3.35) decreases continuously from $q(0)$ to q_T . The fixed cost model has an ending production level q_T that increases with cost. If the cost is greater than zero, the production level must not go to zero, as this will result in negative profit, as discussed in Section 3.4. The time-series plot of the fixed cost model in Figure 5.8 shows that the heuristic tends to decrease too rapidly, then level off at q_T . For this reason, as the upper-left plot in Figure 5.11 shows, percent change in marginal profit values start lower than Hotelling's Rule and trend to well above it. Even so, total profit is within one percent of the theoretical given a level of fixed price, as seen in the lower-left plot in Figure 5.11.

What is clear from the production paths in the top-right plot in Figure 5.11 is that a fixed cost accelerates the depletion of the stock. This is intuitive: for a per period cost, the total cost is less if the stock is depleted more quickly. This shows that a producer that is simply maximizing profit will achieve nearly optimal profit in the face of a fixed cost. Consumers will see the same user costs, but over a shorter period of time, as the lifetime of the resource is reduced as fixed costs increase.

6.5.2 The marginal cost model

The production path for the marginal cost model (3.43) ranges from flat, when $c_1 = K$, to decreasing with the same slope as the costless model, $-r/K$, when the cost is zero. The production paths in the plot at the top of Figure 5.12 appear to be too steep, resulting in percent change in marginal profit values somewhat higher than Hotelling's Rule, as seen in the middle plot in Figure 5.12. This is borne out by the upper-left plot in Figure 5.12, showing that starting percent change in marginal profit values ranging from less than 0.00030 to as high as 0.00040. All of them trend downward toward the optimum of 0.00026 in the course of simulation. The clusters in these curves are an artifact of the large ramp-up that set initial production in multiples of 0.01. This is seen in the lower-right plot in Figure 5.12, where the starting production levels fall into groups, with the least cost sample in each group at the top, near the optimum. This error, which contributes to the error in the slope of the production path, leads to decreasing accuracy with increasing cost, as seen in the lower-left plot in Figure 5.12. The error appears linear, at a rate of about five percent per unit marginal cost.

The upper-right plot in Figure 5.12 shows that the lifetime of the stock is extended by a marginal cost, as anticipated from microeconomic equilibrium theory. This plot also shows that the lowest cost sample from each cluster for a given initial production level is the most accurate. This shows that a producer that is simply maximizing profit will forgo some profit as marginal costs increase, benefiting consumers with slightly lower user costs and a prolonged resource lifetime.

6.5.3 The stock cost model

The production paths in the top plot of Figure 5.10 are an illustration of the complexity of this cost model compared with the fixed cost and marginal cost models. The zero cost plot is a monotonically descending curve, identical to the costless production path, as it should

be. A small stock cost, 0.001, produces the solid curve, which starts downward, then bows up, finally curving rapidly downward to zero as the stock is depleted. The remaining paths curve upward monotonically and end abruptly: these are the production paths for which stock cost eventually equals revenue. In theory, these production paths are seen for costs of 0.01 or greater, yet here it appears for a stock cost of 0.005. The reason for this is found in the upper-right plot in Figure 5.13. Note that for 50 on the cost axis (corresponding to a stock cost of 0.005), the stock lifetime is well below optimum, yet the stock is fully depleted. It appears that the heuristic treats this like a case where cost will eventually exceed revenue, keeping the production level high until the stock is depleted. The upper-left plot in Figure 5.13 shows that percent change in marginal profit starts large and negative (note that the scale of this plot is 10,000 times greater than in the preceding cost models). This means that production starts out increasing at a rapid rate (driving percent change in marginal profit downward at a rapid rate). Then, as the stock is depleted, the percent change in marginal profit becomes large and positive (at the arrow head) as the production level curves rapidly downward. The lower-right plot in Figure 5.13 shows that, in addition to production being inefficiently high, initial production levels start out well above theoretical optimum as well.

The middle plot in Figure 5.10 shows that the heuristic oscillates rapidly between increasing and decreasing production when the stock cost is 0.001. Note that these curves are smoothed over 100 periods. The actual excursions in percent change in marginal profit are orders of magnitude greater than they appear here. This is the cost at which the optimal production path begins leaving some of the stock behind, as seen in the middle plot of Figure 3.4.

The smoothness of the curves resembling an upper case letter "D" at other costs is also the result of smoothing. These curves show evidence that the production path crosses over from negative marginal profit to positive marginal profit. Marginal profit appears as the denominator in percent change in marginal profit, which goes toward negative infinity as marginal profit approaches zero from the negative side, then percent change in marginal profit decreases from infinity as marginal profit moves away from zero in the positive direction.

This is evidenced by the 0.005 curve between days 310 and 400. Note that this cost is close to the cost at which $m(0)$ and m_T cross in the top plot in Figure 3.4, which is the cost at which the optimal production path goes from downward sloping to upward sloping.

The dynamics of these plots reveal an inaccuracy of the coarse strategies available to the heuristic. When the optimal production path is nearly flat, as with a cost of 0.001, the heuristic switches rapidly between stationary, decreasing and increasing strategies. The coarse strategies lead to inefficiently high production levels at costs below 0.01, and to an error in terminating too early for costs above 0.01, as seen by the T plot in Figure 5.13. The latter is a threading-the-needle problem: as stock cost rises toward revenue limits, the heuristic has a narrowing range over which to adjust strategies. One error can result in a negative marginal profit for a period, which is sufficient to terminate production, as per the modified Rule 2 in Section 5.4.

This should not be viewed as strictly a computational limitation of this heuristic. Real production levels are *sticky*: they cannot be varied infinitesimally because of labor, transportation, sales, and other constraints. A human mine operator may realize that a minor adjustment to the production plan will move marginal profit back into the black, but only for a short time, as this is end of life for this resource. Or, the mine operator may be under corporate restrictions that make that determination irrelevant.

One feature of interest in this model is the performance of the heuristic at a stock cost of 0.01. This is the stock cost at which stock cost equals revenue just as the stock is depleted. It is also, as evidenced by the plots in Figure 5.13, the nonzero stock cost for which the heuristic is most accurate. This is the only stock cost for which stock lifetime, ending stock level, and producer profit are all equivalent to the optimum. This is despite having an initial production level that is much higher than optimum, as shown in the lower-right plot in Figure 5.13.

All of the preceding discussion of heuristic errors aside, producer profit deviates by less than ten percent at the worst. This is equivalent to a ten percent reduction in user cost, but

at the cost of a nearly one-third reduction in the resource lifetime.

6.6 Summary

The latter models in Section 6.3 show that, in an oligopoly market with producers optimizing independently, the result is either collusion-like behavior, or Cournot-like behavior, depending on the market structure. This implies that, in the absence of a smoking gun, a collusive market outcome is not *prima facie* evidence of collusion in an oligopoly resource market.

The nonzero cost models are fixed cost, linear marginal cost, and linear stock cost. These represent highly stylized production technologies, though they have real-world parallels. An example of a fixed cost is drilling water wells (Hsiao and Chang, 2002). An example linear marginal cost is the transport of pollutants (Conrad and Clark, 1987). Linear stock costs are seen in some mineral production models (Tietenberg and Lewis, 2000).

For the nonzero cost models, the behavior of the independently optimizing producer is qualitatively consistent with the Hotelling's Rule optimum. The performance of the adaptive producer varies depending on cost structure, ranging from uniformly accurate under a fixed cost, a small inaccuracy that increases with cost for a marginal cost structure, and a widely varying accuracy with a stock cost structure. Although lost producer profits translate into reduced user costs, it is universally at the expense of foreshortening the resource lifetime.

Chapter 7

Policy Implications

Whether privately or publicly held, nonrenewable resources impact public welfare and, therefore, often come under public policy purview. Those who formulate policy typically have multiple and sometimes conflicting goals. Of interest to policy-makers is the relative effectiveness of the tools at their disposal given a natural resource, its ownership, and the market in which it's traded.

Agent-based modeling is well suited to representing the disparate interests of resource owners, consumers, taxpayers, politicians, environmentalists, regulators, and other parties affected by or otherwise interested in nonrenewable resources. Although an ABM that fully endogenizes all of these interests is beyond the scope of this dissertation, the simulation results presented in Chapter 5 make it possible to explore an important consideration for policy-makers: the unintended consequences of taxation.

Chapter 3 presents costs as a consequence of production technologies. Suppose, instead, that those cost functions represent fiscal regimes: E.g. the fixed cost is a franchise tax, the marginal cost is a severance tax, and the stock cost is an environmental restoration bond. Given a set of stylized policy goals, which of these fiscal regimes best accomplishes those goals? Which minimizes market distortion in the form of deadweight loss?

Despite a highly simplistic optimization strategy, the ABM discussed in the preceding chapters is effective at finding production paths that are optimum in terms of maximizing producer profit. In Chapter 6 it is shown that some cost structures can result in a transfer of producer profits to consumer surplus through lower user cost. If, in addition, the producer cost is also tax revenue, the net impact of a tax on public welfare may be neutral or even positive. The net impact depends on the definition of public welfare.

Optimality is in the eye of the beholder: what is optimal for a mine's holding company may not be optimal in terms of public welfare. For the policy-maker interested in maximizing public welfare, the preferred mechanism is to give the resource owner incentives to move the production plan closer to the socially optimal path. The challenge for the policy-maker is to eschew the unintended consequences that may arise from the unseen - and unseeable - measures of optimality employed by the firm or firms in the regulated industry.

This chapter examines the simulation results in Chapter 5 in terms of net impact to public welfare. The agent-based models are presented in Chapter 4. Despite a highly simplistic optimization strategy, the basic ABM is effective at finding production paths that are optimum in terms of maximizing producer profit. This is regarded as the baseline *laissez faire* model with no taxation. Then, the cost structures considered in Chapters 3 and 4 are viewed as taxes. The results discussed in Chapter 6 are revisited here to provide insights into the impacts that can be expected from different fiscal regimes.

7.1 Policy objectives

Public economic policy concerns itself, generally, with social welfare - the well-being of society as a whole. Natural resources certainly impact the economic components of consumer and producer surplus. From a public policy point of view, social welfare often incorporates more than economic surplus. For example, a mining industry may create negative externalities in the form of pollution. Often, public fiscal policies are intended to transfer some of

the public cost of pollution - illness, peeled paint, etc. - to the producer with a Pigouvian tax, thus internalizing that cost. In classic microeconomic equilibrium theory, the increased cost of taxation results in a reduced equilibrium quantity, thereby also reducing the amount of pollution produced. This is not necessarily the outcome for a nonrenewable resource, however.

Many nonrenewable natural resources are publicly owned. Mineral rights, for example, are typically publicly owned and leased to private producers. In these cases, public policy strives to balance current demand for the resource with the wellbeing of future generations (Lecomber, 1979). An important consideration introduced by Hotelling and revisited by subsequent authors is that there is no reason to believe that society discounts the wellbeing of future generations at the same rate that the producer discounts future profit. This represents a significant mismatch between optimality for the producer and the social optimum. For a number of U.S. States, natural resources are a major source of public revenue (See Table 7.1). When big producer profits also mean big royalties, optimality for the producer may well align with the social optimum. Public preferences vary over time, as will the productivity of resource owners, so that public and private optima are likely to be in disequilibrium in general, with only brief periods of concord.

The broad goal of nonrenewable resource policies is to align privately optimal resource depletion with the social optimum (Burness, 1976). The policy mechanisms include sector-specific rules, such as limitations and quotas, and taxation (Hartwick and Olewiler, 1986). This chapter will consider taxation as a policy tool to bring a resource owner's privately optimal production path in line with the social optimum.

Considerations that may be included in the socially optimal production path include forestalling resource depletion (Hotelling, 1931) and internalizing externalities (Dasgupta and Heal, 1980, p. 52). Taxation as a policy tool is also a revenue source (Stiglitz, 2000, p. 718) and this is also a consideration in social optimality. These are discussed in the following sections.

7.1.1 Policies to forestall depletion

The Conservationists in the United States at the turn of the twentieth century promoted the efficient development of nonrenewable resources, which typically meant slowing their exploitation (Hotelling, 1931, Gaudet, 2007). Since that time, this has been a main theme in U.S. natural resource policy, much of which includes government ownership of resources. Implicit in this argument is that the welfare of future generations is no less important than the current generation. That is, public welfare is not discounted. This is a concern expressed by Hotelling and reiterated by Solow, Krautkraemer, and Gaudet. The other argument of the Conservationists is that competition for nonrenewable resources is wasteful, and point also illustrated by Hotelling. That there are efficiencies inherent in putting resources under a single owner (Lecomber, 1979, p. 113) is used to support arguments for government control of nonrenewable resources.

Much of the lexicon of the Conservation Movement appears in contemporary sustainability discussions. For example, the term “future generations” appears frequently in European Union policy statements about sustainability¹.

It should be noted that policies to slow nonrenewable resource production are at odds with other policies to protect or promote the industries that develop them. Many extractive industries are subject to government subsidies that effectively accelerate the depletion of nonrenewable resources. Extractive industries are also subsidized by government-funded research, which reduces uncertainty and lowers development costs, thereby accelerating depletion (Lecomber, 1979, p. 119).

¹See, for example, the position statements of the European Union on sustainable use of natural resources, <http://ec.europa.eu/environment/natres/> (accessed 25 February 2011)

7.1.2 Policies to internalize externalities

The extraction and production of many nonrenewable resources leads to externalities in the forms of air pollution, water pollution, scenic degradation, or other amenity effects (Hanley et al., 1997). Additionally, U.S. States that are resource rich may seek compensation for the exportation of those resource (Burness, 1976, p. 294). Policies to internalize externalities are typically Pigouvian taxes (Dasgupta and Heal, 1980, p. 52). For unlimited production, these taxes add to the cost of production, shifting the supply curve upward and lowering the market-clearing equilibrium quantity. This is a textbook use of Pigouvian taxes (Pindyck and Rubinfeld, 2005, Ch. 18). For a nonrenewable resource, however, taxes will reduce rents, but will not necessarily reduce the equilibrium quantity, and may increase it. This is the result with a fixed tax rate, such as a franchise tax, which affects the producer's optimal production path in the same way as the fixed cost shown in Section 3.4.

Krautkraemer (1998) makes a distinction between flow externalities and stock externalities. Flow externalities are those that arise because of the level of production, such as air pollution. Stock externalities are those that arise from the cumulative effects of production, such as site degradation and atmospheric accumulation of greenhouse gases. Flow externalities represent a per-unit social cost, while stock externalities represent a cumulative production social cost. While it may appear appropriate that flow externalities be addressed by marginal taxes and that stock externalities be addressed by cumulative production taxes, the relative efficiencies of these fiscal regimes indicate otherwise. It is shown in the following sections that marginal taxation results in minimal market distortion while a cumulative production tax incurs a highly variable deadweight loss. When planning errors are taken into consideration (see Section ?? for a discussion of error in this context), there are cumulative production tax regimes that have both low deadweight loss and low vulnerability to planning error. This is discussed in Section 7.3.

Table 7.1: Natural resource taxes as revenue: states for which natural resource taxes constitute more than ten percent of total tax revenue.

State	Percent of total tax revenue
Alaska	77.3%
Wyoming	43.6%
North Dakota	34.3%
New Mexico	19.2%
Montana	14.5%
Oklahoma	13.1%

Source: U.S. Census Bureau

7.1.3 Taxes as revenue

For all the justification that may have been given upon their introduction, governments now rely on resource taxes and fees as a significant and irreplaceable source of revenue. Thirty-four U.S. States levy resource taxes, which make up nearly two percent of all state tax revenues.² For the six states listed in Table 7.1, however, natural resource taxes constitute more than ten percent of total tax revenue. Policy-makers in these states may find it difficult to balance the policy objectives of conservation and reduced externalities when such a significant fraction of tax revenue is at stake.

7.2 The agent-based models in the policy context

Chapters 3 and 4 examined the behavior of the simple ABM for fixed costs, marginal costs, and stock costs. In this chapter, those cost structures are viewed both as costs to the producer and as revenue to the fiscal authority. This will provide a means to investigate the efficiency of the fiscal policies in terms of transferring producer surplus to social welfare, where social welfare is composed of consumer surplus and tax revenue. The results in Chapters 5 and 6 are revisited here to provide insights into the impacts that can be expected from the fiscal policy objectives discussed in Section 7.1.

²From the U.S. Census Bureau

Tax policies are designed to use the market mechanism to align the owner's optimal production path with the social optimum. An ABM of the scope described in the opening section of this chapter would be required to examine the benefits and drawbacks of restrictions and quotas, so this is not examined in this dissertation. What are examined are taxes that impact profit in the form of fixed costs, per-unit costs, royalties, or stock (cumulative production) costs. The ABM results from each cost model are presented and discussed in the policy context.

To recap the behavior of the ABM, in each production period, the agent has the choice to reduce production, increase production, or maintain the current level. The amount by which the agent can reduce production is limited to a reduction that will exactly exhaust the total resource. The increment by which the agent can increase production is one percent. The choice to reduce, increase or maintain production is based on a crude estimate of lifetime income from the resource. Projected discounted future income is based on a rough estimate of the demand function based on the intertemporal change in price.

7.2.1 The inverse demand function and its parameters

All of the ABMs discussed in the chapter come from Chapter 4, and they all use the inverse demand function from Section 3.2. Similarly, the parameter values for all simulations in this chapter are from Section 3.2. The graphs in the chapter can be compared directly, and should be considered general results given the demand function.

7.2.2 Franchise taxes

A franchise tax is a fixed cost incurred by the producer per day even if there is no production. A mineral rights lease is an example of a franchise tax. The State of Idaho, for example, auctions oil and gas leases for starting bids of \$0.25 per acre, and the successful bidder also pays \$1.00 annual rental per acre. Idaho also assesses an additional \$1.00 per

acre annual penalty if the lease is not producing after six years.³

The stylized fixed cost monopoly model from Section 5.4.1 is used to simulate a franchise tax. Figure 7.1 shows the result of 50 Monte Carlo simulations over fixed tax rates distributed over $N(0.12, 0.0016)$. This range provides samples from near costless to 0.20, the rate at which the production path flattens out. Higher tax rates result in more rapid depletion of the resource, with the commensurate loss in total revenue. From the point of view of slowing the rate of depletion, this policy is an abject failure. For the same reason, it fails as a means to slow production to reduce a pollution externality. It is also inefficient at transferring producer profit to tax revenue. This is evidenced by the downward slope of the total revenue curve in the upper right-hand plot in 7.1, which represents increasing deadweight loss, as presented in Section 7.3. This is consistent with the findings of Burness (1976, Table I). Recall from the discussion in Chapter 5 that there is no significant planning error with a fixed cost.

7.2.3 Unit severance taxes

Some U.S. States assess severance taxes based on the quantity of the resource extracted. The State of Ohio, for example, assesses a tax of \$0.10 per barrel for oil, \$0.09 per ton for coal, and \$0.025 per 1,000 cubic feet of natural gas.⁴ Burness (1976, Table II) finds that unit severance taxes will extend the resource lifetime in a competitive market, but have no effect on a monopolist's output. Also, in a monopoly market, taxes are paid entirely from producer profits.

The stylized marginal cost monopoly model from Section 5.4.2 is used to simulate a severance tax. Figure 7.2 shows the result of 50 Monte Carlo simulations over unit tax rates distributed over $N(1.0, 0.110889)$. This range provides samples from near costless to 2.00

³http://www.idl.idaho.gov/bureau/minerals/min_leasing/leasing.htm (accessed 21 February 2011)

⁴<http://codes.ohio.gov/orc/5749> (accessed 21 February 2011)

per unit, which, at the optimal starting production level, is about 0.20 per period (similar to the fixed rate in the previous section). In this model, the higher tax rate puts weak upward pressure on the stock lifetime, while efficiently transferring producer profit into tax revenue. As a policy to prolong the lifetime of the resource, this is mildly successful. As a means to lower production levels to reduce pollution externalities, it is not effective. It is an efficient means to generate tax revenue, however. The weak effect on resource lifetime places this outcome between the competitive and monopoly models of Burness. Deadweight loss, implied by the slope of the total revenue curve, is small, especially considering that planning error increases with the marginal tax rate (see Chapter 5).

7.2.4 Ad valorem taxes

Other U.S. States assess severance taxes based on the market value of the resource. The State of New Mexico, for example, imposes an oil and natural gas tax of 3.75 percent of assessed value based on the price received.⁵ An ad valorem tax has the effect of increasing the price in much the same way as a severance tax. Burness (1976) finds that an ad valorem tax has the same effect as a unit tax.

7.2.5 Royalties

While an ad valorem tax is applied to revenue, a royalty is applied to net profit. For oil and gas produced in the Gulf of Mexico, the U.S. government collects royalty payments ranging from 12 and a half percent for onshore and deep water wells (greater than 400 meters) and 16 and two-thirds percent for shallow water wells.⁶

The royalty cost model is not developed in Section 5.4.3, but the derivation is straight-

⁵<http://www.tax.newmexico.gov/All-Taxes/Pages/Oil-and-Gas-Production-Taxes.aspx> (accessed 21 February 2011)

⁶<http://www.doi.gov/budget/2009/09Hilites/BH019.pdf> (accessed 21 February 2011)

forward. Consider the Hamiltonian for ad valorem tax rate ρ

$$\mathcal{H}(q(t), x(t), t, m) = \pi(q(t), x(t), t) - \rho\pi(q(t), x(t), t) - m(t)q(t) \quad (7.1)$$

The production path is functionally identical to a costless model with a constant multiplier $1 - \rho$

$$\mathcal{H}(q(t), x(t), t, m) = (1 - \rho)\pi(q(t), x(t), t) - m(t)q(t)$$

Tax revenue is total revenue times ρ and the optimal production path is identical to the costless model for this demand function.

The royalty model is a modified version of the stylized costless monopoly model from Section 5.2.1. The model is modified to compute a cost in each period that is a fixed fraction of revenue in the current period. This model is used to simulate the producer response to a royalty tax. Figure 7.3 shows the result of 50 Monte Carlo simulations over royalty rates distributed over $N(0.1, 0.00110889)$. This range provides samples from near costless to a 20 percent tax rate. In this model, the tax rate has no effect on the stock lifetime, and efficiently transfers producer profit into tax revenue. Thus, as a policy to prolong the lifetime of the resource or to lower production levels to reduce pollution externalities, it is not effective. It is an efficient means to generate tax revenue. Here, again, planning errors are small.

7.2.6 Cumulative production fees/bonds/taxes

Cumulative production taxes might represent the costs of site cleanup and mitigation, which increase as more earth is displaced or enhanced recovery techniques are employed. Under some circumstances, this fiscal policy may resemble the reclamation bonds required under the Surface Mining Control and Reclamation Act of 1977 (McDaniel, 1977). Heaps (1985) finds that cumulative production taxation will result in higher rates of extraction but over shorter periods of time, resulting in less total resource being extracted.

The stylized stock cost monopoly model from Section 5.4.3 is used to simulate a cumulative production tax. Figure 7.4 shows the result of 50 Monte Carlo simulations over cumulative production tax rates distributed over $N(0.005, 0.000009)$. This range provides samples from near costless, to a cost that results in shutdown after producing only 50% of the physical stock. There is a distinct kink in the curves at a tax rate of about 100×10^{-4} (the critical point). Below this rate, the agent is able to optimize normally. At the critical point, future profits become negative, and the agent increases production until cost exceeds revenue. At this point, the agent ceases production, even if there is remaining stock. The stock remaining is shown in the upper right-hand graph. These results are consistent with those of Heaps.

As a policy tool, cumulative production taxation increases the rate of depletion, as shown by the first downward section of the upper-left-hand graph (pre-critical-point). The transfer of producer profit to tax revenue is highly inefficient, particularly in the pre-critical-point regime, as shown by the rapid drop in the total revenue curve. Much of the initial drop is due to planning error. Similarly, there is increasing planning error above the critical point. There is very little planning error at the critical point, though there is significant deadweight loss here.

7.3 The efficiency of fiscal regimes

Figure 7.5 compares the efficiency of the fiscal regimes in terms of the deadweight losses they create, and the efficacy of the fiscal regimes in terms of slowing production and forestalling depletion. Deadweight loss in this case is presented as a fraction of the theoretical maximum producer profit as determined from Hotelling's Rule for a costless monopoly producer found in Section 5.2.1. The formula is

$$DWL_{i,j} = 1 - \frac{\Pi_{ij} + C_{ij}}{\Pi^{max}}$$

where

$$\begin{aligned} DWL_{ij} &= \text{deadweight loss for policy } i \text{ and rate } j \\ \Pi_{ij} &= \text{producer profit under policy } i \text{ at rate } j \\ C_{ij} &= \text{total tax revenue under policy } i \text{ at rate } j \\ \Pi^{max} &= \text{Hotelling's Rule maximum producer profit} \end{aligned}$$

The results from a simple optimizing agent-based model (ABM) indicate that cumulative production taxation is largely counter-productive, while unit severance tax accomplishes efficient transfer of profits to taxes and extends the life of a nonrenewable resource. In the case of pollution by-product externalities, the two-pronged goal of internalizing costs and reducing excess supply is only achieved with unit severance taxes. These results are consistent with theoretical findings (Heaps, 1985, Burness, 1976).

Figure 7.1: Franchise taxes.

The per-period cost leads to accelerated depletion. Higher tax rates lead to higher levels of production which, in turn, lead to lower total revenue. See the discussion in Section 7.2.2.

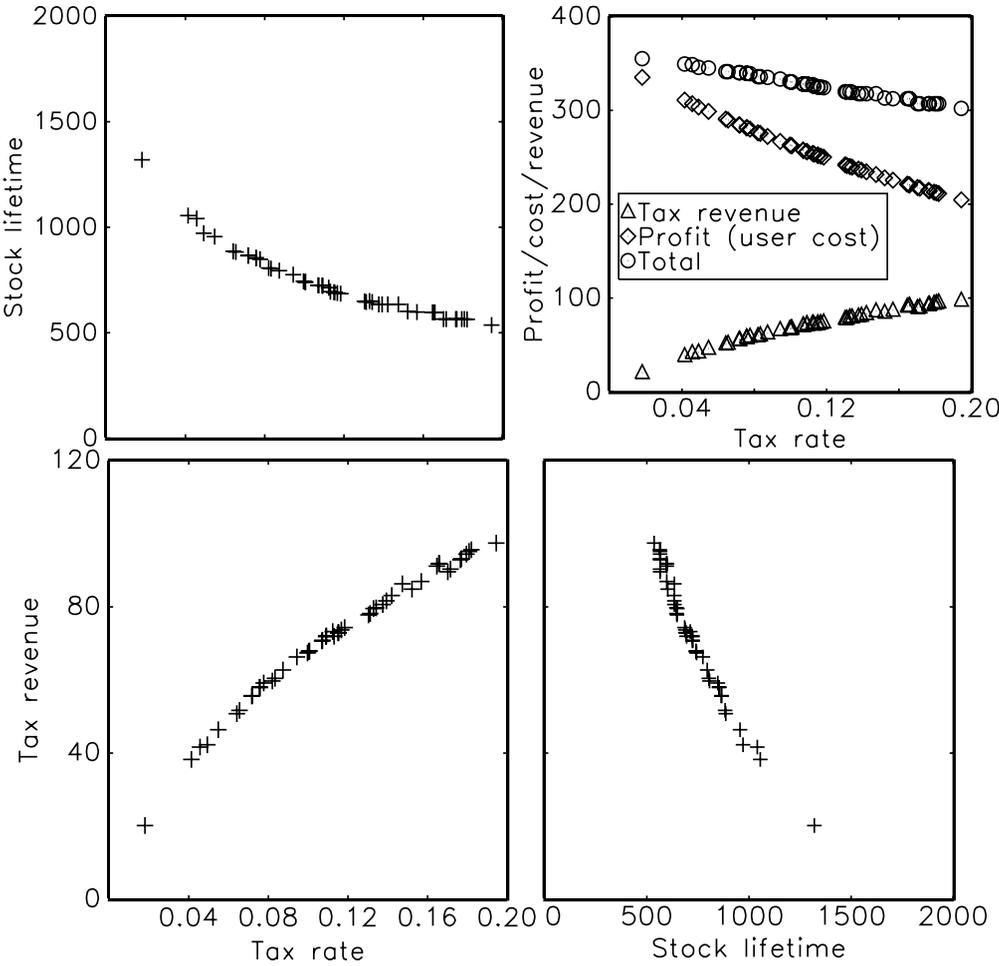


Figure 7.2: Severance taxes.

Higher tax rates extend the lifetime of the stock slightly. Total revenue is conserved: taxation is a direct transfer from the producer to the taxing authority. See the discussion in Section 7.2.3.

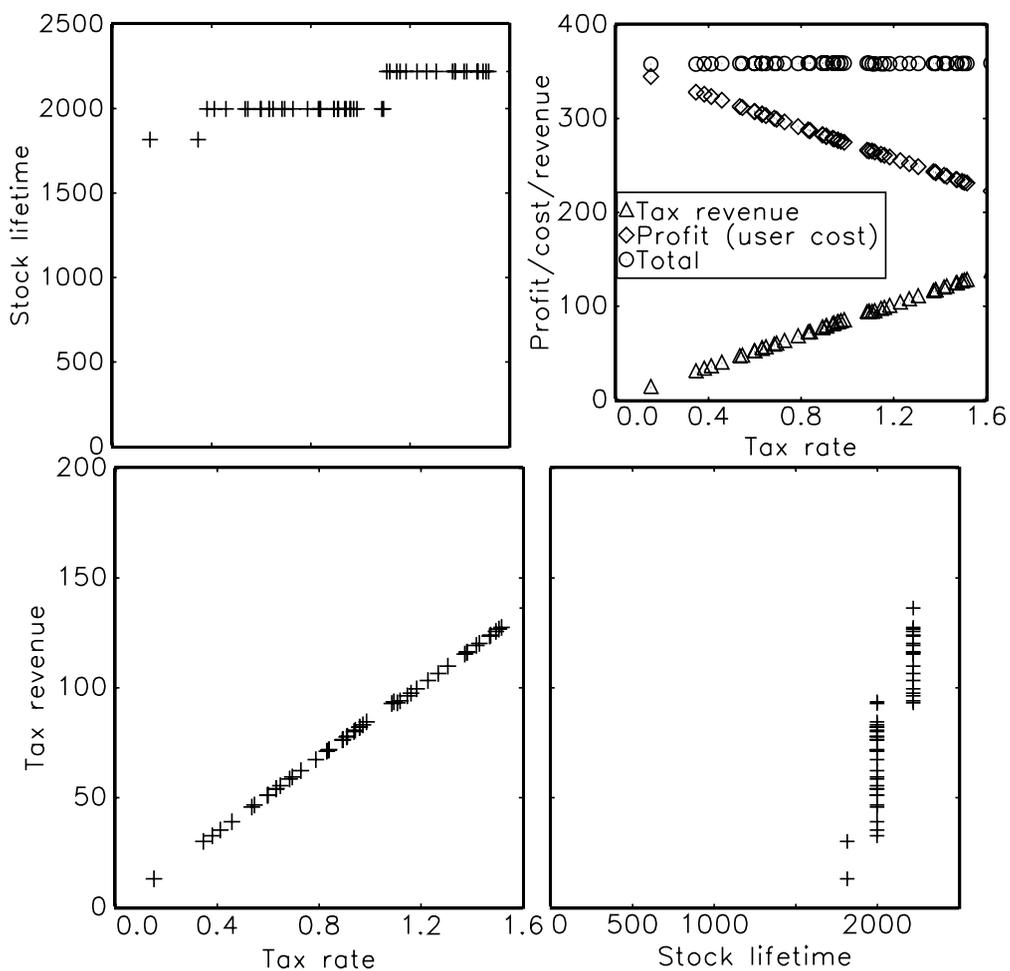


Figure 7.3: Royalties.

The rate is a fraction of total revenue. The lifetime of the stock is unaffected, as is total revenue. The tax is simply a transfer from the producer to the taxing authority. See the discussion in Section 7.2.5.

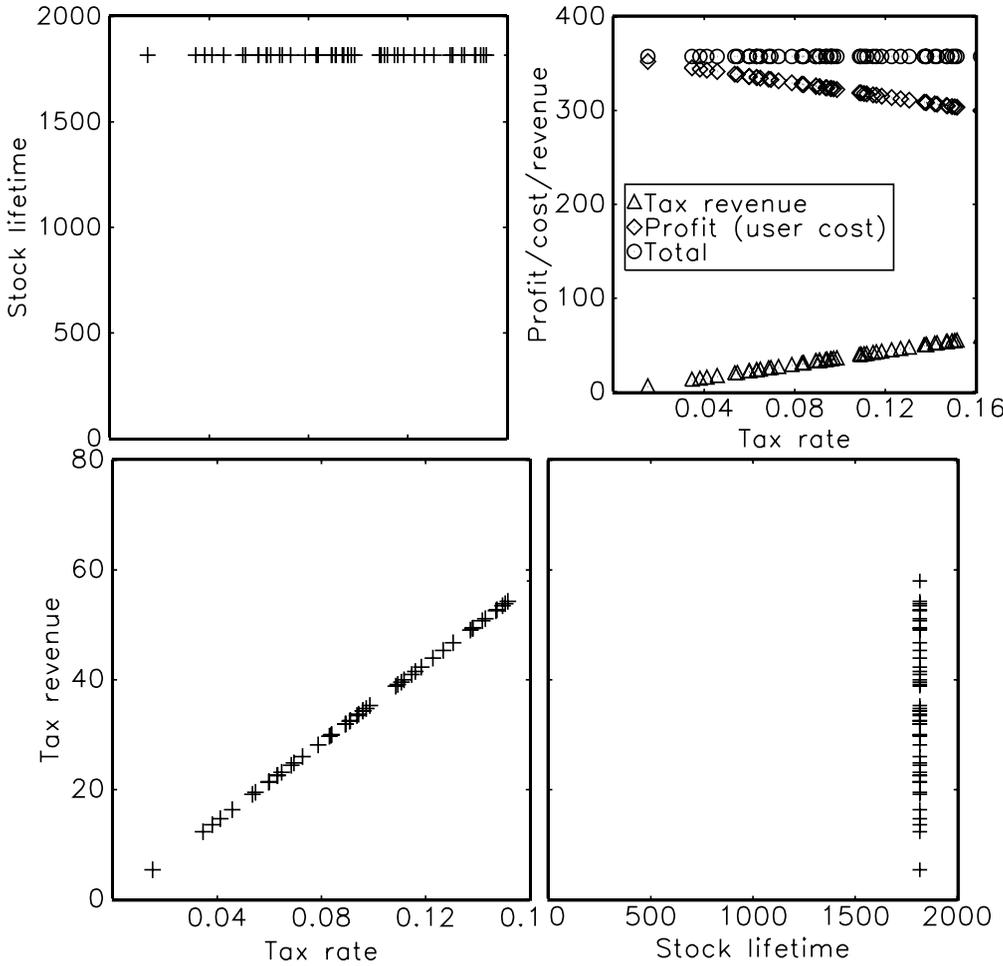


Figure 7.4: Cumulative taxation.

The tax is on cumulative production per period. The rate is multiplied by 10,000. The rate of 100 is the point at which terminal cost equals terminal marginal revenue. At higher cumulative production tax rates, some stock remains, being too costly to produce. See the discussion in Section 7.2.6.

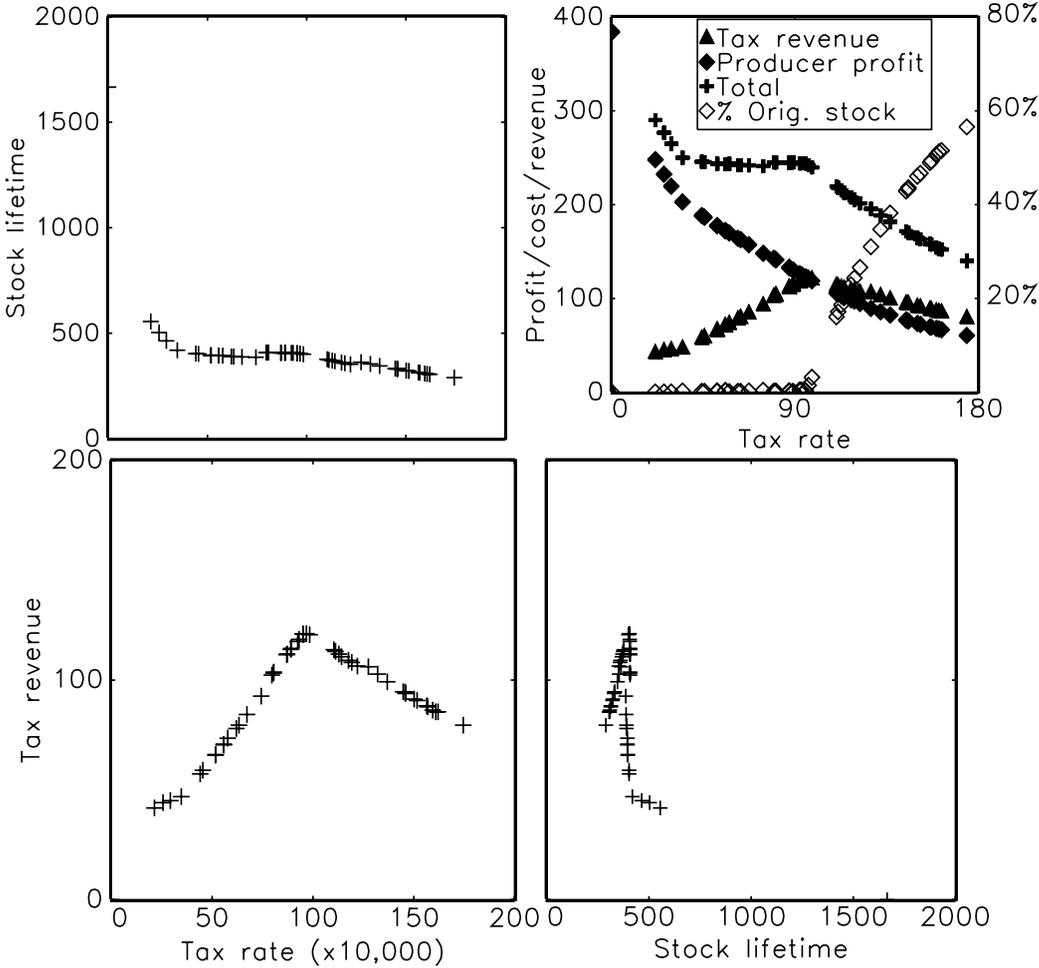
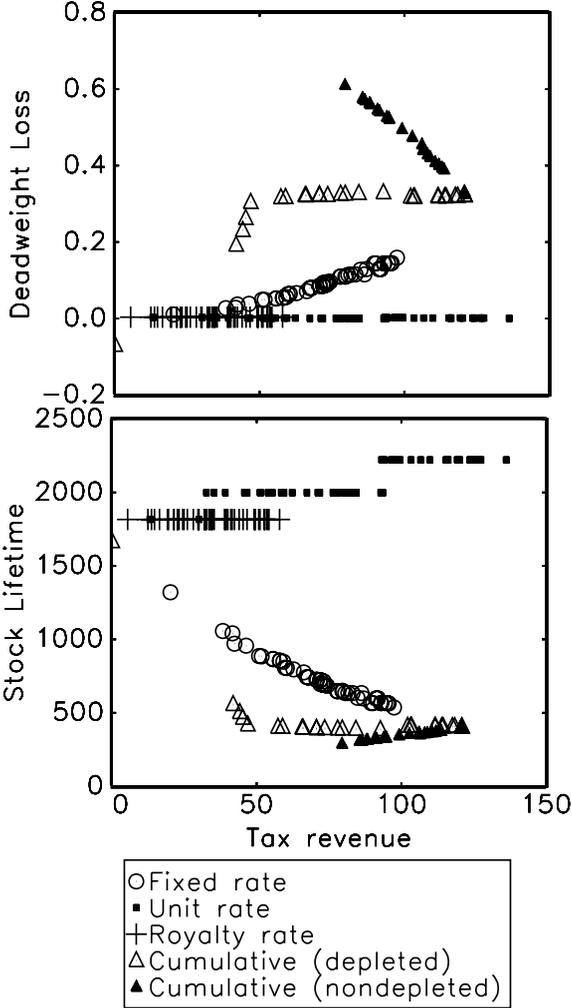


Figure 7.5: Fiscal policy efficiency and efficacy

Deadweight loss is expressed as a fraction of the maximum (Hotelling's Rule) producer profit. The cumulative production tax regime is broken into physical depletion (no stock remaining at the end) and non-depletion (the stock becomes too costly to produce and is not exploited). See the discussion in Section 7.3.



Chapter 8

Conclusions and Future Directions

Were this study to be distilled into a single graphic, it might be Figure 5.1. This figure says that, for this model, Hotelling got it right, but that getting it wrong only costs 0.14 percent for the inverse demand function in Section 3.2. In essence, this is saying that constrained optimization yields the highest profit, but a crude approximation can come very close to the optimum. If there is an information cost associated with computing the optimal path, the resource owner may find that the marginal benefit is less than the marginal cost. These results, as interesting as they may be, are not why this dissertation could be summed up by Figure 5.1. This figure says three things: 1) the ABM results fall in a distribution in the vicinity of the optimal path; 2) the mean of that distribution is very close to the theoretical optimum; and 3) the distribution of errors in the ABM can be quantified from this figure.

The importance of quantifying the performance of an ABM is central to the objectives of this dissertation. This is the topic of Section 8.2.

8.1 The rationale for agent-based modeling

The richness of agent-based modeling lies in the ability to assign to agents actions that are expressed in ways similar to how real-world actors behave. It would be difficult to express, in mathematic terms, outcomes from a population of heterogeneous consumers whose behaviors are based on a sequence of decisions relating to nonlinear individual preferences. Constructing such an agent-based model, however, is straightforward. Furthermore, exploring deviations from the base behaviors may be a simple task in an ABM yet could require a whole new mathematical model. Additionally, simulations of ABMs frequently reveal emergent outcomes - outcomes that are not a direct result of agent rules - while mathematical models do not.

The ABMs in this dissertation are extremely simple by design. In addition to traditional arguments for parsimony, the simplest model can be thought of as the worst case. If adding sophistication improves the accuracy of the model, then the model may be that much closer to reproducing the behavior of interest. If, alternatively, more sophisticated rules worsen the accuracy, those rules are easily jettisoned.

8.2 Reaching the objectives

The three objectives of this dissertation were introduced in Chapter 1. They are 1) to establish a framework for validation of an agent-based modeling approach to nonrenewable resource production planning and analysis; 2) to use this framework to validate a simple optimizing ABM; and 3) to explore any new aspects of the problem that are brought to light by agent-based modeling. Each of these is discussed in the following sections.

8.2.1 The ABM framework

The procedures outlined in the preceding Chapters can be easily extrapolated as a general approach. In this way, it is a prototype for more rigorous validation of agent-based modeling for nonrenewable resource production planning. The procedure outlined is to 1) develop the underlying theory to the extent possible, 2) implement an ABM with behaviors to correspond with the theoretical model, and 3) quantify the differences between theoretical and ABM outcomes.

Developing the theoretical model may require making simplifying assumptions or imposing implied constraints on the problem space, and it is important to make note of these. Unless the modeler implements the behavior directly, an ABM will not observe assumptions like non-negative production or the law of demand. Some solutions to the theoretical model may have to be completed numerically, as is the case in this dissertation. The range of the parameter space for the numerical solutions is another assumption about the theoretical model. Baseline correspondence between theory and the ABM can be established only for an ABM that implements the assumptions of the theoretical model. The baseline correspondence is necessary to evaluate the performance of the ABM outside of the initial parameter range, if that is a goal of the research.

There are no formal definitions for the development of ABMs like the theory of computability in mathematics and computer science (Turing, 1936). Some authors present guidelines on building ABMs (Arthur, 1993, 2006, Axelrod, 2006, 1997a, Epstein, 2006, Axelrod and Tesfatsion, 2006). The baseline ABM is one that is intended to produce the same results as the theoretical model. If the theory is sound and the ABM represents the behaviors that are appropriate for the theoretical model, Monte Carlo sampling the simulation results provides a distribution how the ABM deviates from theory.

Monte Carlo sampling is a common means for collecting simulation data from an ABM (Cook and Skinner, 2005). Hutchins (2010) adopts a hypothesis testing approach to compare the distribution of ABM simulations with a theoretical model. This approach emphasizes the

backwards-induction nature of validating and verifying computer models and simulations. For the ABMs that are extensions of the baseline, confidence estimates can be bootstrapped from the characteristics of the fit to theory in the baseline model.

The application of rigorous statistical and analytical techniques is out of the scope of this dissertation. In order to focus on the process, the assessments of the accuracy and precision of the ABM simulations are primarily qualitative with occasional quantitative points of reference. An application of this approach to a real-world production technology with a real demand function and actual parameters would certainly merit more exhaustive analysis.

8.2.2 Validating the basic Hotelling ABM

The measures used to validate the basic Hotelling ABM are primarily qualitative, as discussed in the preceding section. This is consistent with the highly stylized demand function and production technologies. The preceding two sentences must be associated as a caveat with any statement in this dissertation about the validation of the models. In this sense the process represents qualitative validation.

It is important to recall that, although the goal is to reproduce Hotelling's Rule outcomes using the ABM, Hotelling's Rule is not encoded in the model. The ABM is simply a producer making estimates of total profit based on simple changes to the production path. In the following paragraphs, the term accuracy is used as a measure of the extent to which the behavioral model in the ABM reproduces the outcomes predicted by the mathematical model from Hotelling's Rule. There is no a priori means to estimate this accuracy quantitatively. Qualitatively, Hotelling's Rule is purported to be the outcome for just such a producer, so some similarity in the outcomes is expected.

The basic monopoly ABM reproduces the Hotelling theoretical outcome for the demand function and parameters in Section 3.2 with an accuracy within 0.84 ± 0.011 percent (see Section 5.2.2). This is not a statistically rigorous measure of the error, but given the stylized

nature of the mode, there is no reason to believe the error is significantly greater. With the all the aforementioned caveats, this ABM is validated for this application.

The other models in this dissertation are presented as illustrations of the kinds of errors that can be associated with ABMs and the extent to which they reflect errors made by real people. Qualitatively, however, errors in the fixed cost model appear to be on the order of the errors in the base model. The marginal cost and stock cost models, however, are candidates for the addition of more sophisticated rules. Both the marginal cost and stock cost models exhibit error due to the size of changes in production. This is a simulation artifact and is illustrated in Section 1.2 and discussed for the respective models in Sections 6.5.2 and 6.5.3. Both models would likely benefit from finer granularity in the production level choices. Additionally, the stock cost model could be improved with a simple trial-and-error rule to maintain positive profit as stock costs increase.

The Hotelling ABMs are assessed in terms of the effectiveness of specific stylized fiscal regimes in Chapter 7. This is an example of using agent-based modeling to explore unintended consequences. These outcomes cannot be generalized beyond the stylized demand function and fiscal regimes, although the results are qualitatively consistent with theory.

8.2.3 Exploring emergence

Emergent properties are essentially unexpected behaviors. The classic example is the model of a flock of birds. The individual agents - birds in this case - are given a simple rule: stay close to your neighbors but not too close. For some values of closeness, the motion takes on the appearance of a flock of geese. For other values of closeness, the motion looks like swarming insects. In either case, the group behavior is an emergent property, because it wasn't programmed in the model or the simulation.

Section 6.3 presents an oligopoly extension of the base Hotelling ABM model in which there are emergent outcomes. For some distributions of initial stock, the market behaves as

though the agents were colluding. For other distributions of the initial stock, the market behaves as though the agents have reached a Cournot-Nash equilibrium. For either case the model is identical and it has no way for one agent to know what any other agent is doing other than the aggregate effect on price.

Emergent outcomes are not uncommon in even very basic ABMs, but there is an element of serendipity to their discovery. This particular discovery suggests a more exhaustive investigation of the behaviors and a search for real-world counterparts. This, and other ideas, are presented in the next section.

8.3 Future directions

A few ideas for future work have already arisen in this chapter. It is the intent that this be the first of many steps in establishing the utility and validity of agent-based modeling in natural resource economics. In terms of the validation framework, there is much to be done in terms of formalizing the procedures. This should include the addition of formal tests like the hypothesis testing done by Hutchins (2010).

Some natural extensions of the Hotelling ABM include incorporating real-world demand functions and production functions, real costs and taxes composed from multiple regimes, renewable backstop technologies, overt competitive behaviors, cat-and-mouse games with regulators, and so on. Much of this is achievable with little or no modification to the MASON program developed for this project and available online at <http://www.unm.edu/~ddixon> . As is often the case with ABMs, the effort to analyze these results is likely orders of magnitude greater than the effort to create the models.

Bibliography

- 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.
- W.B. Arthur. On designing economic agents that behave like human agents. *Journal of Evolutionary Economics*, 3(1):1–22, 1993.
- W.B. Arthur. Out-of-equilibrium economics and agent-based modeling. *Handbook of computational economics*, 2:1551–1564, 2006. ISSN 1574-0021.
- R. Axelrod. The evolution of strategies in the iterated prisoner’s dilemma. *Genetic algorithms and simulated annealing*, pages 32–41, 1987.
- R. Axelrod. Advancing the art of simulation in the social sciences. *Complexity*, 3(2):16–22, 1997a.
- R. Axelrod. Agent-based modeling as a bridge between disciplines. *Handbook of computational economics*, 2:1565–1584, 2006. ISSN 1574-0021.
- R. Axelrod and L. Tesfatsion. Appendix AA Guide for Newcomers to Agent-Based Modeling in the Social Sciences. *Handbook of computational economics*, 2:1647–1659, 2006. ISSN 1574-0021.

- R.M. Axelrod. *The complexity of cooperation: Agent-based models of competition and collaboration*. Princeton Univ Pr, 1997b.
- 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.
- H.S. Burness. On the taxation of nonreplenishable natural resources. *Journal of Environmental Economics and Management*, 3(4):289–311, 1976.
- R.D. Cairns and N. Van Quyen. Optimal Exploration for and Exploitation of Heterogeneous Mineral Deposits* 1. *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* 1. *Journal of Environmental Economics and Management*, 42(1):82–103, 2001.
- Alpha C. Chiang. *Elements of dynamic optimization*. McGraw-Hill, 1992. ISBN 0070109117.

- W.A.V. Clark. Residential preferences and neighborhood racial segregation: a test of the Schelling segregation model. *Demography*, 28(1):1–19, 1991. ISSN 0070-3370.
- J.M. Conrad and C.W. Clark. *Natural resource economics: notes and problems*. Cambridge Univ Pr, 1987. ISBN 0521337690.
- D.A. Cook and J.M. Skinner. How to Perform Credible Verification, Validation, and Accreditation for Modeling and Simulation. *The Journal of Defense Software Engineering*, May, 2005.
- P. Dasgupta and G.M. Heal. *Economic theory and exhaustible resources*. Cambridge Univ Pr, 1980.
- Under Secretary of Defense for Acquisition Technology Department of Defense. *Department of Defense Modeling and Simulation (M&S) Glossary. DoD 5000.59-M*. Department of Defense of the United States of America, 1998.
- 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. Remarks on the foundations of agent-based generative social science. *Handbook of computational economics*, 2:1585–1604, 2006. ISSN 1574-0021.
- 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* 1. *Journal of Environmental Economics and Management*, 15(4):412–418, 1988. ISSN 0095-0696.
- G erard Gaudet. Natural resource economics under the rule of Hotelling. *Canadian Journal of Economics Revue Canadienne dEconomie*, 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.
- J.M. Hartwick and N.D. Olewiler. *The Economics of Natural Resource Use*. Harpercollins College Div, 1986. ISBN 0060426950.
- 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.

- Terry Heaps. The taxation of nonreplenishable natural resources revisited. *Journal of Environmental Economics and Management*, 12(1):14 – 27, 1985. ISSN 0095-0696. doi: DOI: 10.1016/0095-0696(85)90013-0.
- D.F. Hendry. Monte carlo experimentation in econometrics. *Handbook of Econometrics*, 2: 937–976, 1984.
- 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.
- Meghan Hutchins. *Exploring the Effects of Yard Management and Neighborhood Influence on Carbon Storage in Residential Subdivisions*. University of Michigan, 2010.
- 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.
- A.J. Laurie and N.K. Jaggi. Role of ‘vision’ in neighbourhood racial segregation: A variant of the Schelling segregation model. *Urban Studies*, 40(13):2687–2704, 2003. ISSN 0042-0980.
- 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’Economie*, 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 non-renewable 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.
- S. Luke, G.C. Balan, L. Panait, C. Cioffi-Revilla, and S. Paus. MASON: A Java multi-agent simulation library. In *Proceedings of Agent 2003 Conference on Challenges in Social Simulation*, volume 9, 2003.
- J.A. McDaniel. Surface Mining Control and Reclamation Act of 1977: An Analysis, *The Harv. Envtl. L. Rev.*, 2:288, 1977.
- 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.
- D.H. Meadows, D.L. Meadows, J. Randers, and W.W. BEHRENS III. *The limits to growth. A report for the Club of Rome's project on the predicament of mankind*. Universe Books, 1972.
- N. Metropolis and S. Ulam. The monte carlo method. *Journal of the American Statistical Association*, 44(247):335–341, 1949.
- 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.
- Robert S. Pindyck and Daniel L. Rubinfeld. *Microeconomics*. Pearson Prentice Hall, Upper Saddle River, NJ, 6. ed., internat. ed. edition, 2005. ISBN 0131912070.
- 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 Boltyanskii, RV Gamkrelidze, and EF Mishchenko. The mathematical theory of optimal processes, Fizmatgiz, Moscow 1961. *English transl., Interscience, New York*, 1962.
- C.W. Reynolds. Flocks, herds and schools: A distributed behavioral model. In *Proceedings of the 14th annual conference on Computer graphics and interactive techniques*, pages 25–34. ACM, 1987. ISBN 0897912276.
- 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. Dynamic models of segregation. *The Journal of Mathematical Sociology*, 1(2):143–186, 1971. ISSN 0022-250X.

- 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.
- J.E. Stiglitz. *Economics of the public sector*. Norton, New York, 2000.
- 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* 1. *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.
- A.M. Turing. On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society Series*, 2(42):230–265, 1936.
- 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.

Appendices

Appendix A

GAUSS procedure to solve for q_{max} in the marginal cost model

```

#include inthp.sdf
#include qnewtonmt.sdf
library pgraph;
/*****
* QNewton
*
* Returns the fitness of the
* the current values of
*
* mc = marginal cost
* K = choke price
* r = daily interest rate
*
* This finds the value of Tmax that
* makes total production equal to x0.
*****/
proc QNewton(struct PV Qparam, struct DS qNewtonData);
local Qvar, mc, K, leftSide, rightSide, qDelta;
// get the Q parameter
Qvar = pvUnpack(Qparam,"Q");
// put it into the data object
qNewtonData.dataMatrix[2] = Qvar;
// get the parameter(s)
mc = qNewtonData.dataMatrix[1];
K = qNewtonData.dataMatrix[3];

```

```

// compute the left and ride sides
leftSide = exp(-K*Qvar);
rightSide = 1 - mc * Qvar;
// the fitness is the difference squared
qDelta = leftSide - rightSide;
if (Qvar <= 0 and mc < 5);
// steer it away from zero
qDelta = 100;
endif;
retp(qDelta ^ 2);
endp;
/*****
* QmaxFuncC
*
* Returns the maximum production
* level for a given marginal cost
*
* mc = marginal cost
* K = choke price
* r = daily interest rate
*
* Computes
*
* Qmax = maximum production level
*
*****/
proc QmaxFuncC(struct DS *qFuncData, mc);
external struct PV Qparam;
external struct QNewtonmtControl qnewControl;
external struct QNewtonmtOut QmaxOpt;
local Qmax;
// put mc into the data
qFuncData->dataMatrix[1] = mc;
// solve for optimal Qmax
QmaxOpt = QNewtonmt(&QNewton, Qparam, *qFuncData, qnewControl);
// put Tmax into the data
Qmax = pvUnpack(QmaxOpt.par, "Q");
retp(mc~Qmax);
endp;
/* DECLARE STRUCTURES */
struct DS qData;
struct DS *pqData;

```

```

struct PV Qparam;
struct QNewtonmtControl qnewtControl;
struct QNewtonmtOut QmaxOpt;
/* ASSIGN VALUES */
mc = 1; // base value
K = 5; // constant
r = 1.1^(1/365.25)-1; // constant
Qmax = 100.0; // starting value
/* CREATE CONTROL STRUCTURES */
qnewtControl = QNewtonmtControlCreate;
Qparam = pvCreate;
Qparam = pvPack(Qparam, Qmax, "Q");
pqData = &qData;
qData.dataMatrix = mc | Qmax | K | r;
qnewtControl.output = 1; /* print results */
print "mc = " qData.dataMatrix[1];
print "Qmax = " qData.dataMatrix[2];
print "K = " qData.dataMatrix[3];
print "r = " qData.dataMatrix[4];
// a vector of unit cost values from 0 to 5.0
c = seqa(0, 0.1, 50);
// a vector of profit solutions for each value of the unit cost
piV = {};
for n(1,50,1);
Qmax = 100.0; // starting value
piV = piV | QmaxFuncC(pqData, c[n]);
endfor;
print piV;
outfile = "unitqmax";
let varnames = unit_cost q_max;
success = saved (piV, outfile, varnames);

```

Appendix B

GAUSS procedure for numerical solutions to marginal cost model

```

#include inthp.sdf
#include qnewtonmt.sdf
library pgraph;
/*****
* QOpt
*
* Returns the optimal production path
* at time t given the values of
*
* mc = marginal cost
* Tmax = terminal time of production
* K = choke price
* r = daily interest rate
*
*****/
proc QOpt(struct DS *qOptData, t);
local mc, Tmax, K, r;
mc = qOptData->dataMatrix [1];
Tmax = qOptData->dataMatrix [2];
K = qOptData->dataMatrix [3];
r = qOptData->dataMatrix [4];
retp(1/K*(ln(K) - ln((K - mc)*exp(-r*(Tmax - t)) + mc)));
endp;
/*****
* QTotal
*

```

```

* Returns the total production over
* the period (0,Tmax) given the values of
*
* mc = marginal cost
* Tmax = terminal time of production
* K = choke price
* r = daily interest rate
*
* This is used to solve for Tmax
* by finding the Tmax that makes
* total production equal to total
* stock.
*****/
proc QTotal(struct DS *qTotalData);
local Tmax,lim,qTot;
external struct inthpControl intControl;
Tmax = qTotalData->dataMatrix[2];
if (Tmax <= 0);
retp(0);
endif;
lim = Tmax | 0;
qTot = inthp4(&QOpt, qTotalData, intControl, lim);
retp(qTot);
endp;
/*****
* QNewton
*
* Returns the fitness of the
* the current values of
*
* mc = marginal cost
* K = choke price
* r = daily interest rate
* x0 = total stock
*
* This finds the value of Tmax that
* makes total production equal to x0.
*****/
proc QNewton(struct PV Tparam, struct DS qNewtonData);
local Tmax, x0, qTot, qDelta;
// get the T parameter
Tmax = pvUnpack(Tparam,"T");

```

```

// put it into the data object
qNewtonData.dataMatrix[2] = Tmax;
// get the x0 parameter
x0 = qNewtonData.dataMatrix[5];
// compute total production given these parameters
qTot = QTotal(&qNewtonData);
// the fitness is the difference squared
qDelta = x0 - qTot;
retp(qDelta ^ 2);
endp;
/*****
* PiQ
*
* Returns the profit at time t
* given the current values of
*
* mc = marginal cost
* K = choke price
* r = daily interest rate
*
*****/
proc PiQ(struct DS *piQData, t);
local mc, K, r, q, piQ;
mc = piQData->dataMatrix[1];
K = piQData->dataMatrix[3];
r = piQData->dataMatrix[4];
q = QOpt(piQData, t);
piQ = (1 - exp(-K * q) - mc * q)*exp(-r * t);
retp(piQ);
endp;
/*****
* PiTotal
*
* Returns the total profit given
* the current values of
*
* mc = marginal cost
* Tmax = terminal time of production
* K = choke price
* r = daily interest rate
*
*****/

```

```

proc PiTotal(struct DS *piTotalData);
local Tmax,lim;
external struct inthpControl intControl;
Tmax = piTotalData->dataMatrix[2];
lim = Tmax | 0;
if (Tmax <= 0);
retp(0);
endif;
retp(inthp4(&PiQ, piTotalData, intControl, lim));
endp;
/*****
* PiFuncC
*
* Returns the total optimized
* profit as a function of mc,
* given the values of
*
* K = choke price
* r = daily interest rate
* x0 = total stock
*
* Computes
*
* Tmax = terminal time of production
*
*****/
proc PiFuncC(struct DS *piFuncData, mc);
external struct PV Tparam;
external struct QNewtonmtControl qnewControl;
external struct QNewtonmtOut TmaxOpt;
local Tmax,piTot;
// put mc into the data
piFuncData->dataMatrix[1] = mc;
// solve for optimal Tmax
TmaxOpt = QNewtonmt(&QNewton, Tparam, *piFuncData, qnewControl);
// put Tmax into the data
Tmax = pvUnpack(TmaxOpt.par,"T");
piFuncData->dataMatrix[2] = Tmax;
// compute total profit based on mc and Tmax
piTot = PiTotal(piFuncData);
retp(mc~Tmax~PiTot);
endp;

```

```

/* DECLARE STRUCTURES */
struct DS qData;
struct DS *pqData;
struct PV Tparam;
struct inthpControl intControl;
struct QNewtonmtControl qnewtControl;
struct QNewtonmtOut TMaxOpt;
/* ASSIGN VALUES */
mc = 1; // base value
Tmax = 2235.255; // starting value
K = 5; // constant
r = 1.1^(1/365.25)-1; // constant
x0 = 100; // constant
/* CREATE CONTROL STRUCTURES */
intControl = inthpControlCreate;
qnewtControl = QNewtonmtControlCreate;
Tparam = pvCreate;
Tparam = pvPack(Tparam, Tmax, "T");
pqData = &qData;
qData.dataMatrix = mc | Tmax | K | r | x0;
// qnewtControl.output = 1; /* print results */
print "PiZero = " PiFuncC(pqData, 0);
print "mc = " qData.dataMatrix[1];
print "Tmax = " qData.dataMatrix[2];
print "K = " qData.dataMatrix[3];
print "r = " qData.dataMatrix[4];
print "x0 = " qData.dataMatrix[5];
print "QTotal = " QTotal(pqData);
// a vector of unit cost values from 0 to 5.0
c = seqa(0, 0.04, 50);
// a vector of profit solutions for each value of the unit cost
piV = {};
for n(1,50,1);
piV = piV | PiFuncC(pqData, c[n]);
endfor;
// outfile = "unitnumerical3";
// let varnames = unit_cost t_max pi_total;
// success = saved (piV, outfile, varnames);
// regress Tmax as a function of cost and cost squared (because the residuals
_olsres = 1;
{ nam,m,b, stb,vc, std, sig, cx,rsq, resid,dbw } = ols("",piV[.,2],piV[.,1]~piV[.,
print resid;

```

Appendix C

GAUSS procedure for numerical solutions to the stock cost model

```

#include inthp.sdf
#include qnewtonmt.sdf
library pgraph;
/*****
* QOpt
*
* Returns the optimal production path
* at time t given the values of
*
*****/
proc QOpt(struct DS *qOptData, t);
local sc, Tmax, K, r, mT, deltaT, ert, m, q;
sc = qOptData->dataMatrix [1];
Tmax = qOptData->dataMatrix [2];
K = qOptData->dataMatrix [3];
r = qOptData->dataMatrix [4];
mT = qOptData->dataMatrix [7];
deltaT = Tmax - t;
q = 0;
if 0 == sc;
q = r/K*deltaT;
else;
ert = exp(-r * deltaT);
m = sc / r * (1 - ert) + mT * ert;
q = 1 / K * ln( K / m);
endif;

```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model160

```
retp(q);
endp;
/*****
* QTotal
*
* Returns the total production over
* the period (0,Tmax) given the
* current values of the parameters
*****/
proc QTotal(struct DS *qTotalData);
local Tmax,lim,qTot;
external struct inthpControl intControl;
Tmax = qTotalData->dataMatrix[2];
lim = Tmax | 0;
qTot = inthp4(&QOpt, qTotalData, intControl, lim);
retp(qTot);
endp;
/*****
* QCurrent
*
* Returns the current total production over
* the period (0,t) given the current values
* of the parameters
*****/
proc QCurrent(struct DS *qTotalData);
local Tmax,lim,qTot,t;
external struct inthpControl intControl;
t = qTotalData->dataMatrix[6];
qTot = 0;
if (0 < t);
lim = t | 0;
qTot = inthp4(&QOpt, qTotalData, intControl, lim);
endif;
retp(qTot);
endp;
/*****
* QNewtonT
*
* This finds the square of the
* difference between total production
* and x0-xT given the current values
* of the parameters.
```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model161

```

*****/
proc QNewtonT(struct PV Tparam, struct DS qNewtonData);
local Tmax, sc, K, x0, mT, dx, xT, qTot, qDelta;
// get the T parameter
Tmax = pvUnpack(Tparam,"T");
// put it into the data object
qNewtonData.dataMatrix[2] = Tmax;
// get the parameters
x0 = qNewtonData.dataMatrix[5];
xT = qNewtonData.dataMatrix[8];
// compute total production given these parameters
qTot = 0;
if (0 < Tmax);
qTot = QTotal(&qNewtonData);
endif;
// the fitness is the difference squared
qDelta = (x0 - xT - qTot)/x0;
retp(qDelta^2);
endp;
/*****
* QNewtonQT
*
* Returns the square of the difference
* between
*
*  $\exp(-K*qT)$ 
*
* and
*
*  $sc / K * qT$ 
*****/
proc QNewtonQT(struct PV Qparam, struct DS qNewtonData);
local qT, sc, K, r, ekt, linQ, qDelta, x0;
// get the Q parameter
qT = pvUnpack(Qparam,"QT");
// get the parameters
sc = qNewtonData.dataMatrix[1];
K = qNewtonData.dataMatrix[3];
r = qNewtonData.dataMatrix[4];
x0 = qNewtonData.dataMatrix[5];
// compute the exponential and linear terms
ekt = exp(-K * qT);

```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model162

```

linQ = (1 - sc * x0) / (1 + K *qT);
if (abs(ekt) < abs(linQ));
qDelta = (linQ - ekt)/ekt;
else;
qDelta = (ekt - linQ)/linQ;
endif;
// the fitness is the difference squared
retp(qDelta ^2);
endp;
/*****
* PiX
*
* Returns the present value of
* the profit at time t given
* the current values of the
* parameters.
*****/
proc PiX(struct DS *piQData, t);
local sc, K, r, mT, q, qCumulative, piX;
external struct inthpControl intControl;
sc = piQData->dataMatrix [1];
K = piQData->dataMatrix [3];
r = piQData->dataMatrix [4];
mT = piQData->dataMatrix [7];
piQData->dataMatrix [6] = t; // pass the current time to QCurrent
q = QOpt(piQData, t);
qCumulative = QCurrent(piQData); // this is the same as x0 - x
piX = (1 - exp(-K * q) - sc * qCumulative)*exp(-r * t);
retp(piX);
endp;
/*****
* CostX
*
* Returns the present value of
* the cost at time t given
* the current values of the
* parameters.
*****/
proc CostX(struct DS *piQData, t);
local sc, K, r, mT, q, qCumulative, costX;
external struct inthpControl intControl;
sc = piQData->dataMatrix [1];

```

*Appendix C. GAUSS procedure for numerical solutions to the stock cost model*163

```
K = piQData->dataMatrix [3];
r = piQData->dataMatrix [4];
mT = piQData->dataMatrix [7];
piQData->dataMatrix [6] = t; // pass the current time to QCurrent
q = QOpt(piQData, t);
qCumulative = QCurrent(piQData); // this is the same as x0 - x
costX = sc * qCumulative * exp(-r * t);
retp(costX);
endp;
/*****
* PiTotal
*
* Returns the total profit given
* the current values of the
* parameters.
*****/
proc PiTotal(struct DS *piTotalData);
local Tmax,lim,piTot;
external struct inthpControl intControl;
Tmax = piTotalData->dataMatrix [2];
piTot = 0;
if (0 < Tmax);
lim = Tmax | 0;
piTot = inthp4(&PiX, piTotalData, intControl, lim);
endif;
retp(piTot);
endp;
/*****
* TotalCost
*
* Returns the total cost given
* the current values of the
* parameters.
*****/
proc TotCost(struct DS *costTotalData);
local Tmax,lim,totCost;
external struct inthpControl intControl;
Tmax = costTotalData->dataMatrix [2];
totCost = 0;
if (0 < Tmax);
lim = Tmax | 0;
totCost = inthp4(&CostX, costTotalData, intControl, lim);
```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model164

```

endif;
retp(totCost);
endp;
/*****
* PiFuncC
*
* Returns the total optimized
* profit as a function of mc,
* given the current values of
* the parameters.
*****/
proc PiFuncC(struct DS *piFuncData, sc);
external struct PV QTparam;
external struct QNewtonmtControl qnewtControl;
external struct QNewtonmtOut TmaxOpt;
external struct QNewtonmtOut QTmaxOpt;
local qT, K, mT, Tmax, piTot, dx, x0, xT, r, eKT;
struct PV Tparam;
Tparam = pvCreate;
// the newton solver is especially sensitive to
// the starting value, so this is a rough hack
// at an estimate for Tmax
if (0 == sc);
// the right answer is 1957
Tparam = pvPack(Tparam,2000,"T");
else;
// this is based solely on the fact that sc = 0.0035 is
// problematic unless the starting estimate is close
// to the actual value (638)
Tparam = pvPack(Tparam,2/sc,"T");
endif;
// put sc into the data
piFuncData->dataMatrix[1] = sc;
K = piFuncData->dataMatrix[3];
r = piFuncData->dataMatrix[4];
x0 = piFuncData->dataMatrix[5];
////////////////////////////////////
//
// First, figure out qT, which requires numerical fitting if xT is going to b
//
if (0 == sc);
qT = 0;

```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model165

```

// calculate mT
mT = K * exp(-K * qT);
// calculate xT
xT = 0;
elseif (sc < 1/x0);
// solve for optimal qT
QTmaxOpt = QNewtonmt(&QNewtonQT, QTparam, *piFuncData, qnewtControl);
qT = pvUnpack(QTmaxOpt.par, "QT");
// calculate mT
eKT = exp(-K * qT);
mT = K * eKT;
// calculate xT
xT = 0;
else;
qT = 100;
// calculate mT
mT = 0;
// calculate xT
xT = x0 - 1 / sc;
endif;
// put it in the structure
piFuncData->dataMatrix [7] = mT;
piFuncData->dataMatrix [8] = xT;
piFuncData->dataMatrix [9] = qT;
////////////////////////////////////
//
// Second, find the Tmax that uses up the stock (down to xT if its nonzero)
//
// solve for optimal Tmax
TmaxOpt = QNewtonmt(&QNewtonT, Tparam, *piFuncData, qnewtControl);
// put Tmax into the data
Tmax = pvUnpack(TmaxOpt.par, "T");
piFuncData->dataMatrix [2] = Tmax;
// compute total profit based on sc and Tmax
piTot = PiTotal(piFuncData);
retp(sc~qT~mT~Tmax~xT~PiTot);
endp;
/* DECLARE STRUCTURES */
struct DS qData;
struct DS *pqData;
struct PV Tparam;
struct PV QTparam;

```

Appendix C. GAUSS procedure for numerical solutions to the stock cost model166

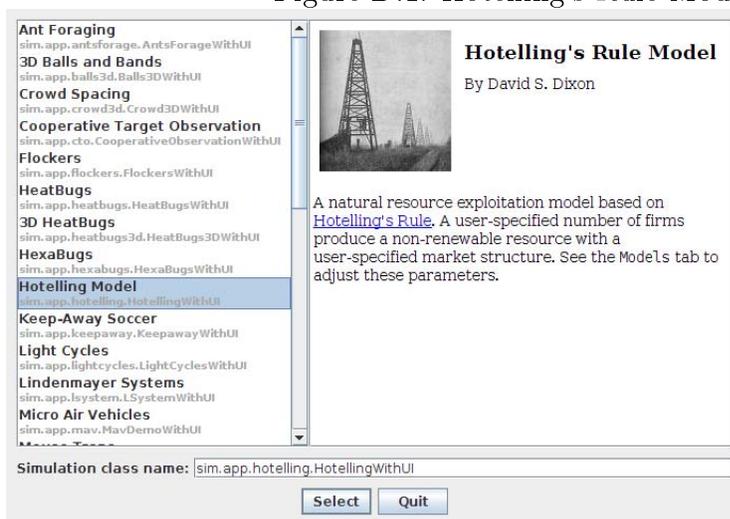
```
struct inthpControl intControl;
struct QNewtonmtControl qnewtControl;
struct QNewtonmtOut TMaxOpt;
struct QNewtonmtOut QTMaxOpt;
/* ASSIGN VALUES */
sc = 0.000026; // base value
Tmax = 600; // starting value
K = 5; // constant
r = 1.1^(1/365.25)-1; // constant
x0 = 100; // constant
t = 1;
mT = 0.01;
qT = 0.01;
xT = 0;
/* CREATE CONTROL STRUCTURES */
intControl = inthpControlCreate;
qnewtControl = QNewtonmtControlCreate;
Tparam = pvCreate;
Tparam = pvPack(Tparam, Tmax, "T");
QTparam = pvCreate;
QTparam = pvPack(QTparam, qT, "QT");
pqData = &qData;
qData.dataMatrix = sc | Tmax | K | r | x0 | t | mT | xT | qT;
// qnewtControl.output = 1; /* print results */
// a vector of stock cost values up to .05
c = seqa(0, 0.0005, 100);
print "sc\tqT\tmT\tTmax\txT\tPiTot";
// a vector of profit solutions for each value of the stock cost
piV = {};
for n(1,100,1);
piSet = PiFuncC(pqData, c[n]);
tCost = TotCost(pqData);
piV = piV | piSet~tCost;
endfor;
print piV;
outfile = "stocknumerical2";
let varnames = stock_cost q_t m_t t_max x_t pi_total cost_total;
success = saved (piV, outfile, varnames);
// regress Tmax as a function of cost and cost squared (because the residuals
// _olsres = 1;
// { nam,m,b,stb,vc,std,sig,cx,rsq,resid,dbw } = ols("",piV[.,2],piV[.,1]);
```

Appendix D

The MASON Hotelling package

MASON¹ is a Java software library to support discrete-event multiagent simulation. The MASON software distribution includes a collection of sample applications presented in the Console application, shown in Figure D.2. The panel on the left lists the sample ABMs provided in MASON. The panel on the right shows the title display for the Hotelling's Rule Model application created for this project.

Figure D.1: Hotelling's Rule Model in the MASON



¹<http://cs.gmu.edu/~eclab/projects/mason/>

When the Hotelling's Rule Model application is selected, the control panel is displayed in Figure D.2. Below the control panel is the market display, which allows the user to view statistics for specific firms. To the right of the control panel are plot panels displaying the dynamical variables in the simulation. The control panel allows the user to select the number of firms and the market model (monopoly, oligopoly, etc.) The simulation is started by clicking the triangular button at the bottom of the control panel.

Figure D.2: Hotelling's Rule Model control panel.

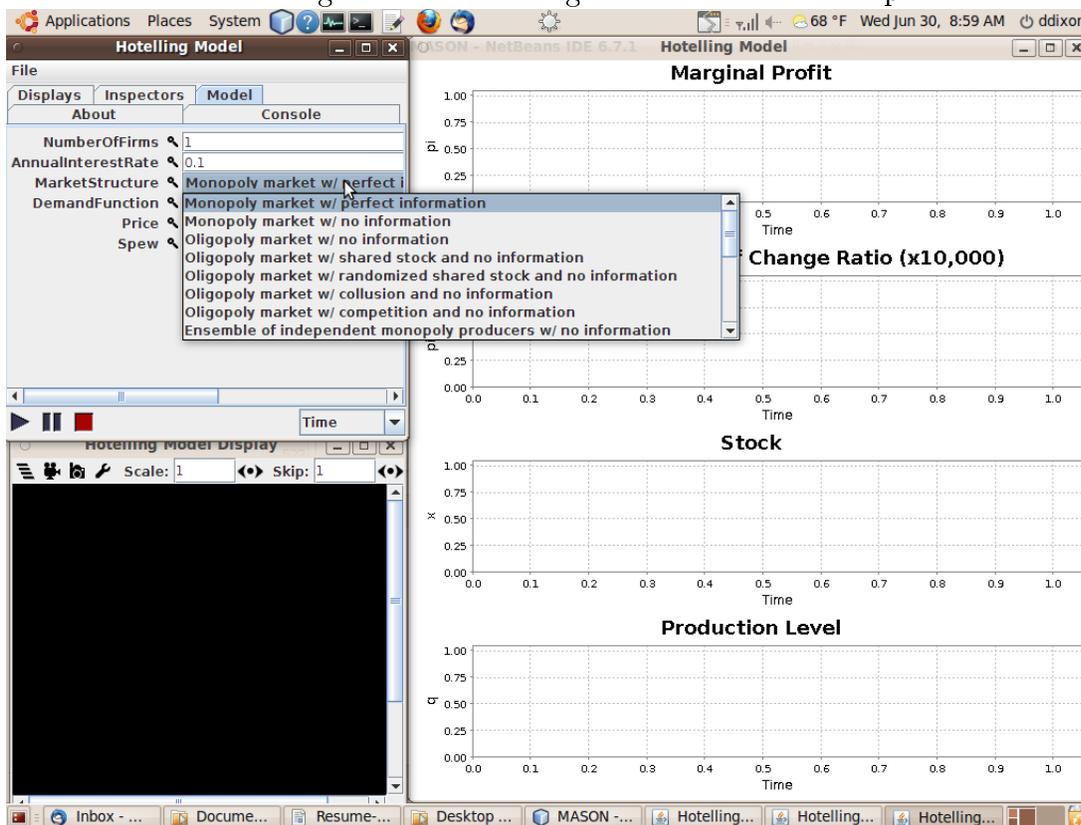


Figure D.3 shows the display after selecting an oligopoly market of five firms with shared stock. In the market display, the red dot represents the market and the white dots represent each of the firms, numbered from zero going clockwise from the top. By clicking on the red dot, the user can see current market statistics, such as current market price. By clicking on a white dot, the user can see current firm statistics, such as production level, stock level, costs, etc. Figure D.4 shows a running simulation.

Figure D.3: An oligopoly market with five firms.

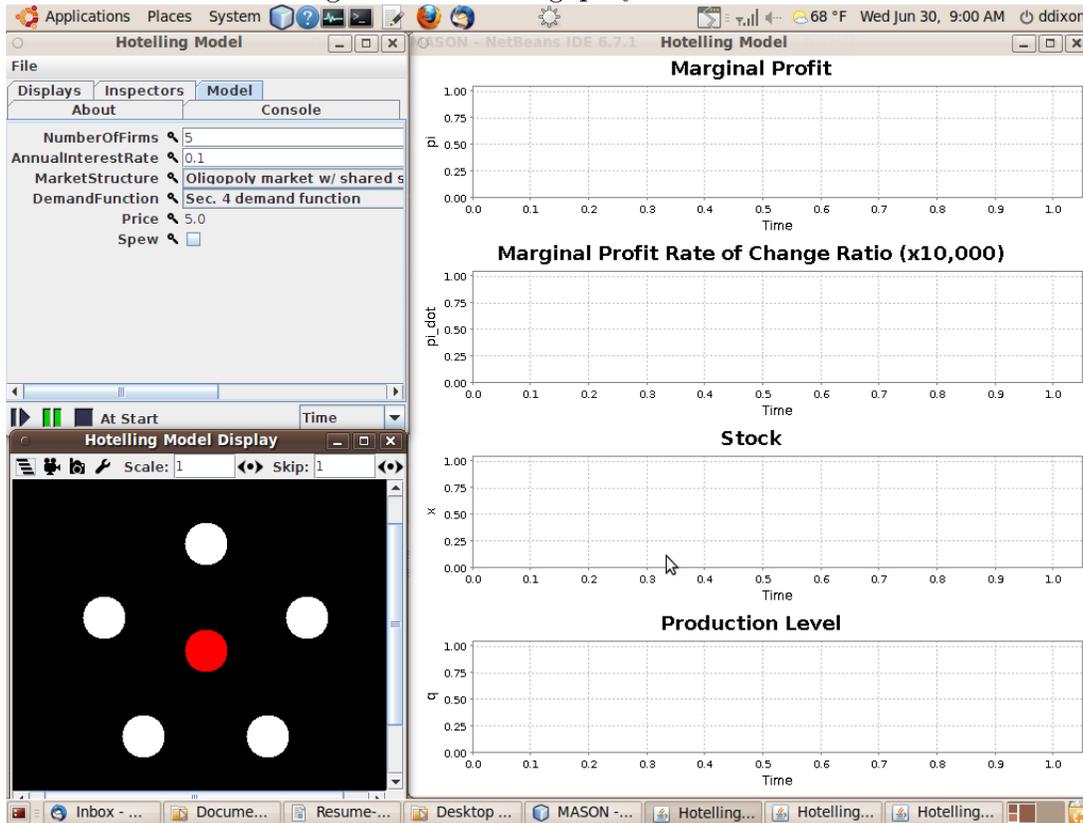


Figure D.4: A running simulation.

