Cournot equilibrium as emergent behavior in a nonrenewable resource agent-based model

David S. Dixon*

June 20, 2011

In a simple agent-based model of a small oligopoly nonrenewable natural resource model, the agents, communicating solely through the market price, sometimes exhibit collusion-like behavior, sometimes Cournot-like behavior. The collusion-like behavior is shown to arise when differences between the agents are small. Conversely, the Cournot-like behavior is shown to result from differences in production decisions based on differences in the agents.

Close examination of the Cournot-like behavior indicates that the outcome results from a misinterpretation of market price response. This motivates investigation into how additional qualitative information about the market leads to quantitative improvements in the estimated marginal price.

(JEL Q32)

Keywords: Agent-Based Modeling; Agent-Based Computation Economics; Natural Resource Modeling; Natural Resource Production.

An agent-based model (ABM) of a nonrenewable resource market presented in Dixon (2010) exhibits mixed behavior for a small oligopoly. The Monte Carlo simulations sometimes show a collusion-like outcome, with all agents adopting the same production path and dividing the monopoly rent. Other simulations show a Cournot-like outcome, where higher total production and lower rent than the collusion-like outcome. Preliminary examination of the dynamics reveals that, while the larger producers adopt a decreasing production path similar to the collusion-like outcome, the smaller ones follow a flat production path. In these cases, total production decreases and marginal profit increases over time, consistent

*Department of Economics, MSC 05 3060, 1 University of New Mexico, Albuquerque, NM 87131-0001, ddixon@unm.edu. The MASON code for these models is available from http://www.unm.edu/~ddixon.

with Hotelling's Rule. Each agent is making a production decision based on the decisions of the others, a condition for Cournot equilibrium. In this model, the decisions of the other agents are aggregated and communicated solely through the market price. Once each agent has selected a production path that is either decreasing or flat, there is no incentive to change, a condition for Nash equilibrium.

This paper examines the details of the decisions that result in either a collusion-like or a Cournot-like outcome. These are decisions in the computational sense - they involve simple calculations and basic logic. The timing of these decisions, meaning the times at which the logical conditions change, affects the overall rate of production and ultimate depletion of the resource. The relationship between timing and the characteristics of the Cournot-like outcome are examined. The paper concludes with a discussion of ways to improve an agent's decision-making with minimal increase in model complexity. That is, addressing the question of what additional information is available to the agents without expanding the model or assuming omniscience.

1 The oligopoly model

The basic ABM is intended to examine the behaviors of producer agents in comparison with Hotelling's Rule (Hotelling, 1931). Much of the discussion of the ABM is taken from Dixon (2011). The 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.[1] 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.[2] 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, each agent determines a production path heuristically. The heuristic

[1] *Production path* refers to either the time-series plot of production level or to the equation or heuristic that produces it.

[2] The ABM uses the inverse demand function from Hotelling (1931, Sec. 4):

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

with the choke price $K = 5$.

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, which includes discounting future profit. This decision is made in each production period (days) and amounts to Bayesian updating.

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. Each simulation is a Monte Carlo sample of the stochastic variable space. 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] The producer agents are autonomous and do not communicate among themselves: the only communication is in the form of the market price, which reflects aggregated production in the previous period. This occurs over four distinct periods during each time-step, as illustrated in Figure 1. The size of a simulation time-step is arbitrary, though the default discount rate is assumed daily and compounds to ten percent per annum. Changing the discount rate would change the implicit time-step.[4]

At the beginning of each simulation, there is a short period during which the producer agent searches for the optimal starting production level. This is called the *ramp-up* phase. The number of time-steps depends on the size of the initial production increments. The basic monopoly model reaches the optimal production level in about ten time-steps, and runs to completion in 1958 time-steps. The oligopoly model examined in this paper varies the ramp-up increments so the results are quite different, as shown in the next section.

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 monopoly model production path is very similar to the optimal (Hotelling's Rule) production path. The size of the discrete ramp-up increment results is a small error, however, so that the production path begins about one percent too high at the end of ramp-up, and ends a fraction of a percent too soon. This error is illustrated in Figure 2 with the error greatly exaggerated. The error in the monopoly ABM results in a total profit that is only a fraction of a percent lower than the Hotelling's Rule optimum, however.

The preceding discussion of the monopoly model is relevant to the oligopoly model for two reasons. First, the oligopoly model is identical to the monopoly model, but with more than one producer. That is, the behaviors are not modified to address the presence of competitors in the market. The total resource stock is divided among the producers to simplify comparison between models with different numbers of producers. The division of the resource stock is nearly even with small stochastic variation in order to distinguish the

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

[4]The ABM is implemented using the MASON platform. The interface provides the capability for a use to change a number of model parameters, including the interest rate.

Figure 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.

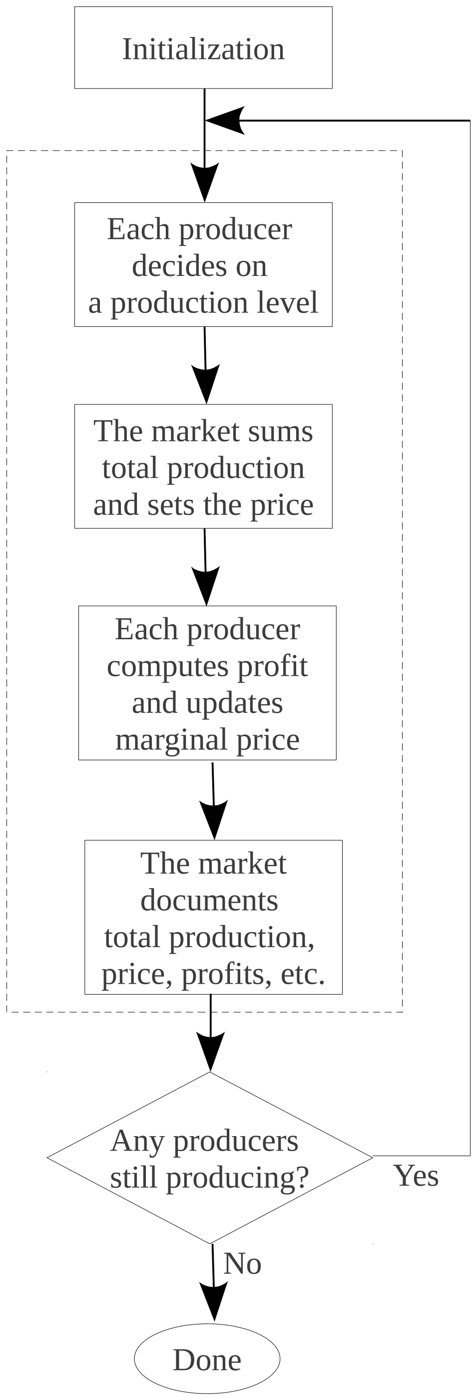

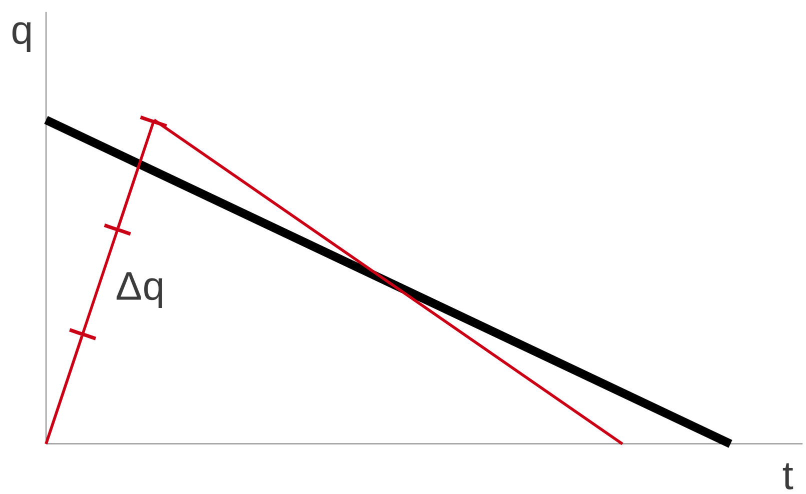

Figure 2: An example of the initial production level error.
The heavy black line is the Hotelling's Rule optimum production path. The thinner red line is the production path as determined by the ABM. At the beginning, the agent starts increasing production by Δq. In this example, the heavy line is the 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.

producers. The ramp-up increment is also reduced proportional to the average share of the total resource, so that the ramp-up in aggregate production for multiple producers is very similar to the ramp-up for a monopoly producer.

The second reason for presenting the details of the monopoly model is that collusion in an oligopoly model should look like a monopoly market. That is, if the producers are colluding to maximize total rent, the aggregate production path should be identical to the monopoly production path.

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.

The aggregate production paths for markets with from 1 to 10 producers is shown in Figure 4. Two distinctly different types of paths are seen here. Those for 2, 3, and 4 producers are similar to the monopoly path, reflecting a collusion-like outcome. The paths for 6, 8 and 9 producers are similar. However, the aggregate production paths for markets of 5, 7 and 10 producers exhibit a stepped production path that is quite different from the monopoly path. The outcomes for N=(2, 3, 4, 6, 8, 9) is collusion-like, because true collusion cannot occur in these models. That is, the agents do not communicate directly and, therefor, cannot collude. They are all self-optimizing based on exactly the same information: the market price. That they make the same decisions is not surprising.

The outcomes for N=(5, 7, 10) are Cournot-like because each agent appears to be self-optimizing in consideration of the production decisions of the others, but the agents are not able to know the production decisions of the others directly. The only information available is the market price response to the aggregate production. This is an emergent outcome of the heuristic and is discussed qualitatively in Dixon (2010). That discussion is summarized in the next few paragraphs and concludes this section.

The individual production paths for the five-producer market are shown in Figure 4. In this simulation, after the ramp-up phase, some producers begin reducing production while others continue unchanged. Initially, the total stock of 100 is distributed among the five producers in near equal amounts, with small random deviations. The five producers begin the ramp-up phases together, but some reach the optimal production path one time-step before the others. As a result, some first select a decreasing production path, while others select a flat (constant) production path. All firms experience increasing marginal profit, since market price is increasing due to decreasing aggregate production. The increase in marginal profit is sufficient to make the current strategy, whether decreasing or flat, to appear optimal. The producers maintaining constant or nearly constant production deplete the stock sooner than the others, resulting in the step in the aggregate production paths in Figure 4. Ultimately, the total profit for either strategy is nearly the same, as seen by the convergence of the total profit curves in Figure 4.

Figure 3: Oligopoly models - 1 to 10 producers.

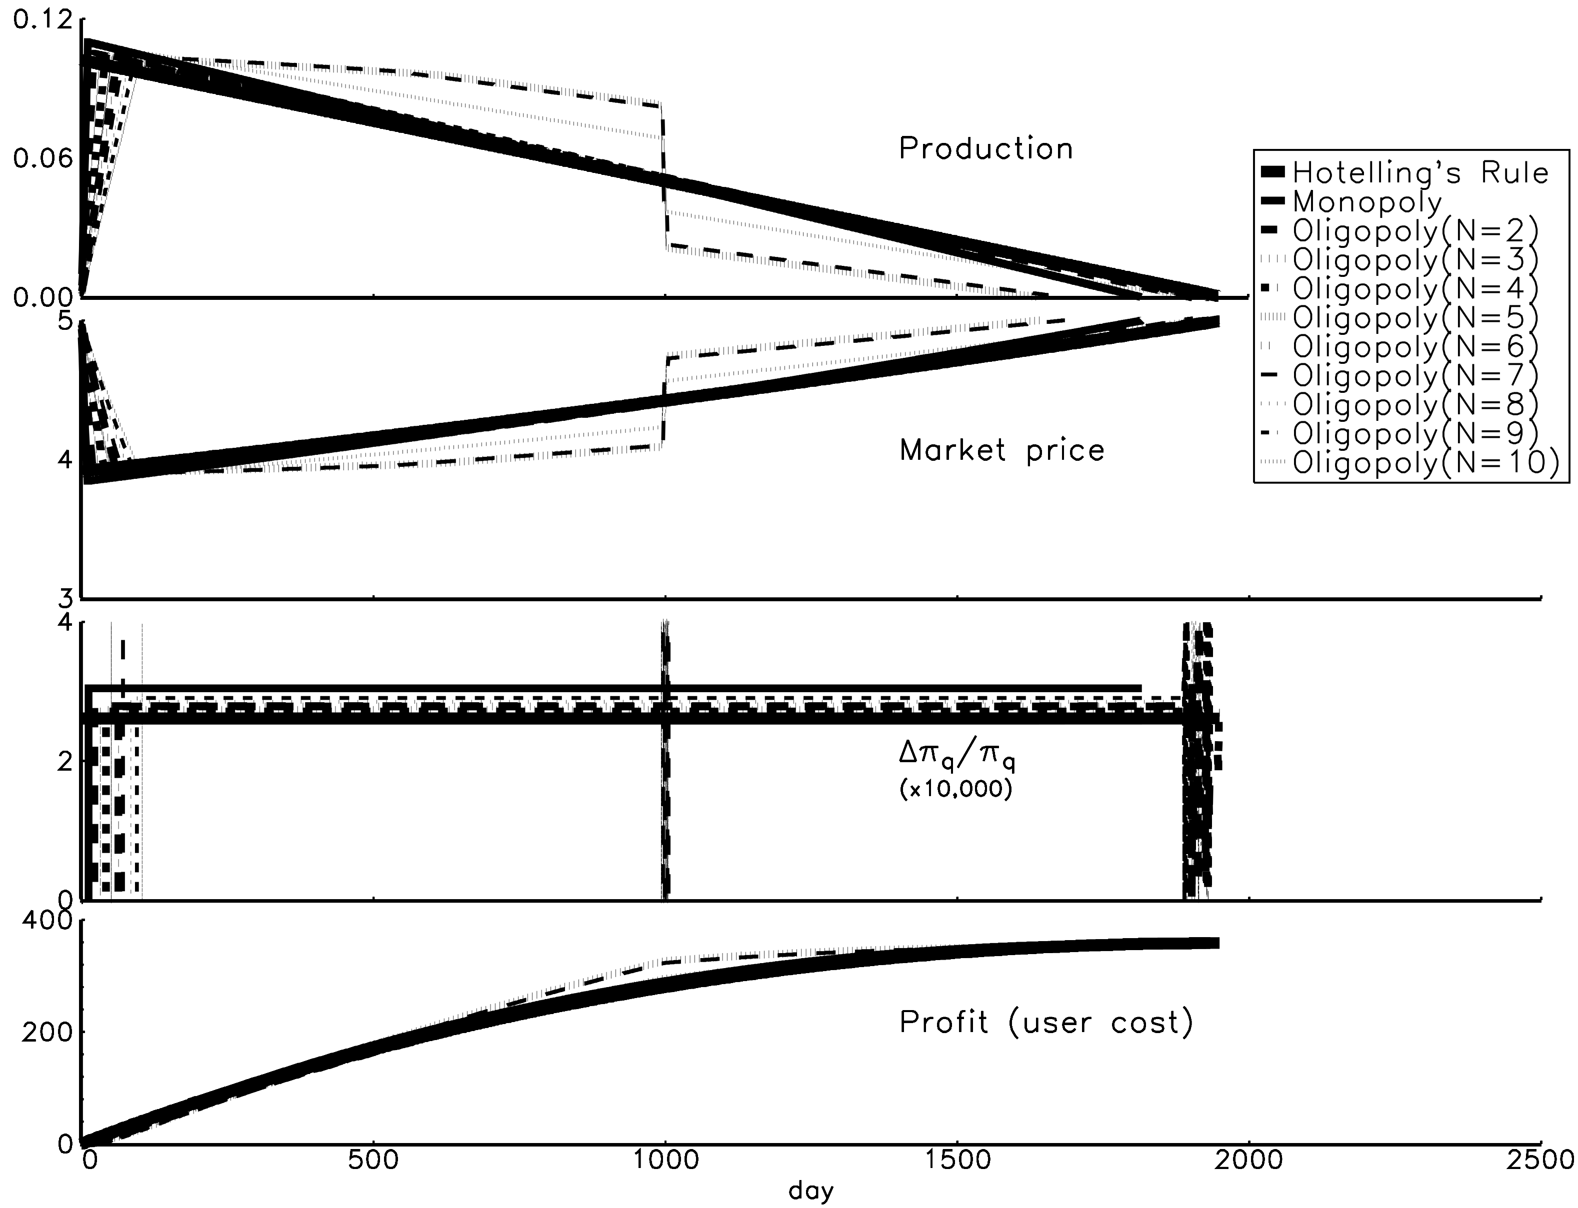

Figure 4: Adaptive agent pooled stock models - time-series. See the discussion in Section ??.

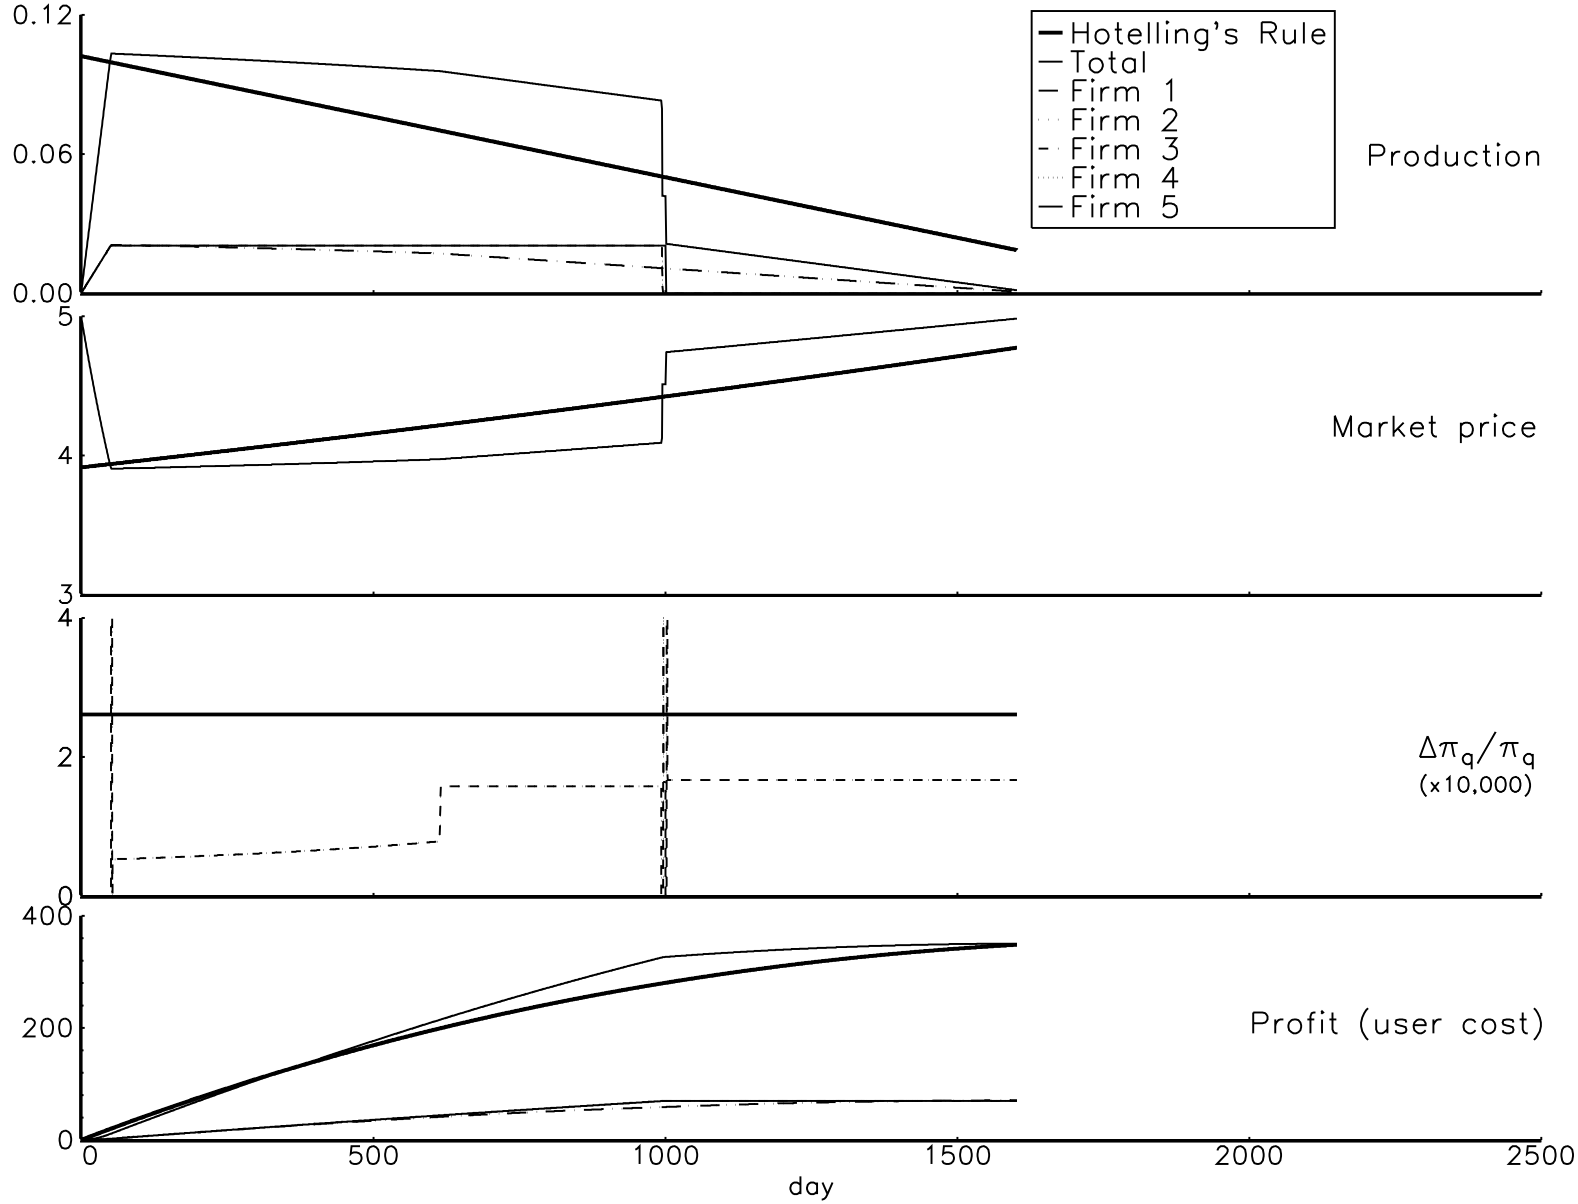

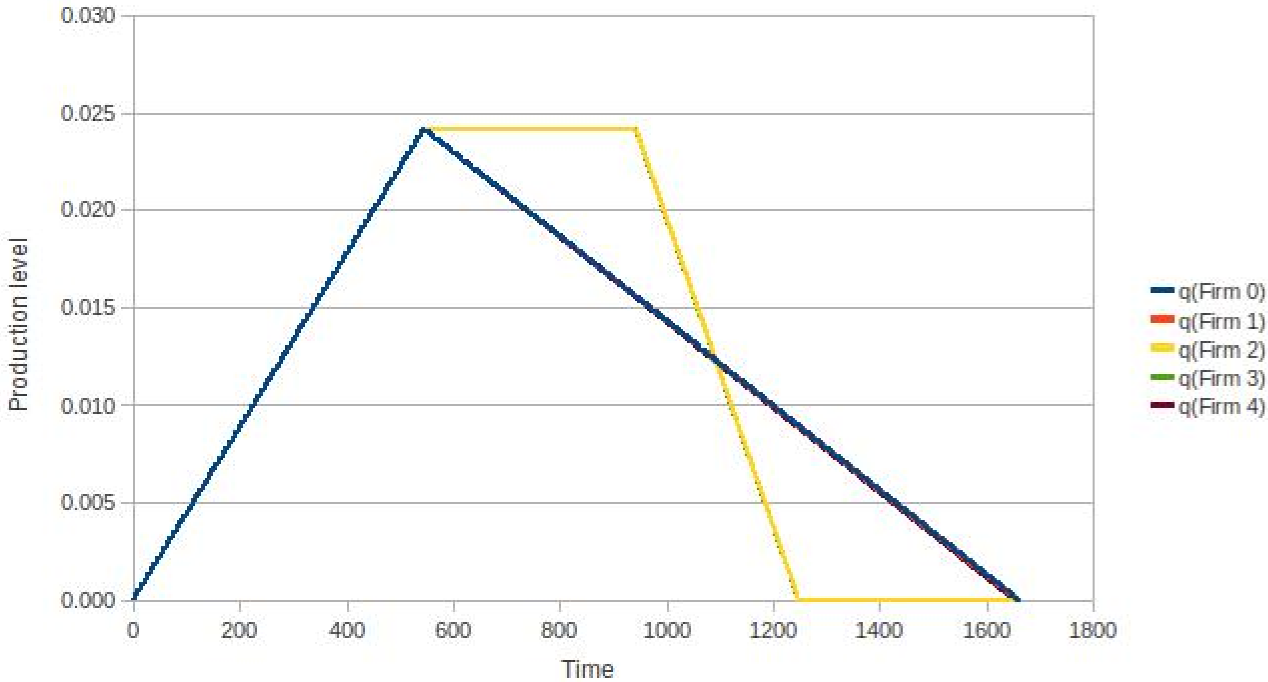

Figure 5: Time-series for Simulation #4.

2 The details of Cournot-like outcome

The preceding section reviews the initial results from the oligopoly ABM. This paper examines the details of the production decisions that lead to the Cournot-like outcomes. The original ABM is modified to introduce stochastic variation in the ramp-up increment for the N=5 oligopoly model. This variation leads to some simulations with collusion-like outcomes and some with Cournot-like outcomes. The characteristic that distinguishes the producers is the peak production level at which each producer switches to either a decreasing production path or a flat production path. This information for 25 simulations is summarized in Table 1. The table shows that the collusion-like outcome results when the variation in peak production level between the producers is less than 10^5.

A time-series plot of Simulation #4 is shown in Figure 5. The plots for Producers 1 and 4 are identical to and hidden by the plot for Producer 0, and the plot for Producer 3 is identical to and hidden by the plot for Producer 2.

The details of peak production are shown in Figure 6. In this plot, dashed lines have been used to show the individual producers. This shows that producers 2 and 3 (the two producers with the smallest stocks) reach peak production in period 544, at which point they begin a decreasing production path. In the next time period, however, they select a flat production path.

Figure 7 is the same as Figure 6 with price included. This shows that, when producers 2

Table 1: Initial results from varying ramp-up increment.

Simulation #	Minimum peak production level	Peak production level range	Outcome
0	0.02075	0.00031	Cournot
1	0.02094	$< 10^5$	collusive
2	0.02102	$< 10^5$	collusive
3	0.02092	$< 10^5$	collusive
4	0.02416	0.00004	Cournot
5	0.02060	$< 10^5$	collusive
6	0.02137	0.00002	Cournot
7	0.02092	$< 10^5$	collusive
8	0.02126	0.00018	Cournot
9	0.02084	$< 10^5$	collusive
10	0.02127	0.00017	Cournot
11	0.02068	0.00033	Cournot
12	0.02147	0.00015	Cournot
13	0.02146	$< 10^5$	collusive
14	0.02075	$< 10^5$	collusive
15	0.02082	$< 10^5$	collusive
16	0.02063	$< 10^5$	collusive
17	0.02068	0.00030	Cournot
18	0.02355	0.00006	Cournot
19	0.02079	$< 10^5$	collusive
20	0.02056	0.00040	Cournot
21	0.02095	0.00022	Cournot
22	0.02066	0.00002	Cournot
23	0.02425	0.00004	Cournot
24	0.02387	0.00005	Cournot

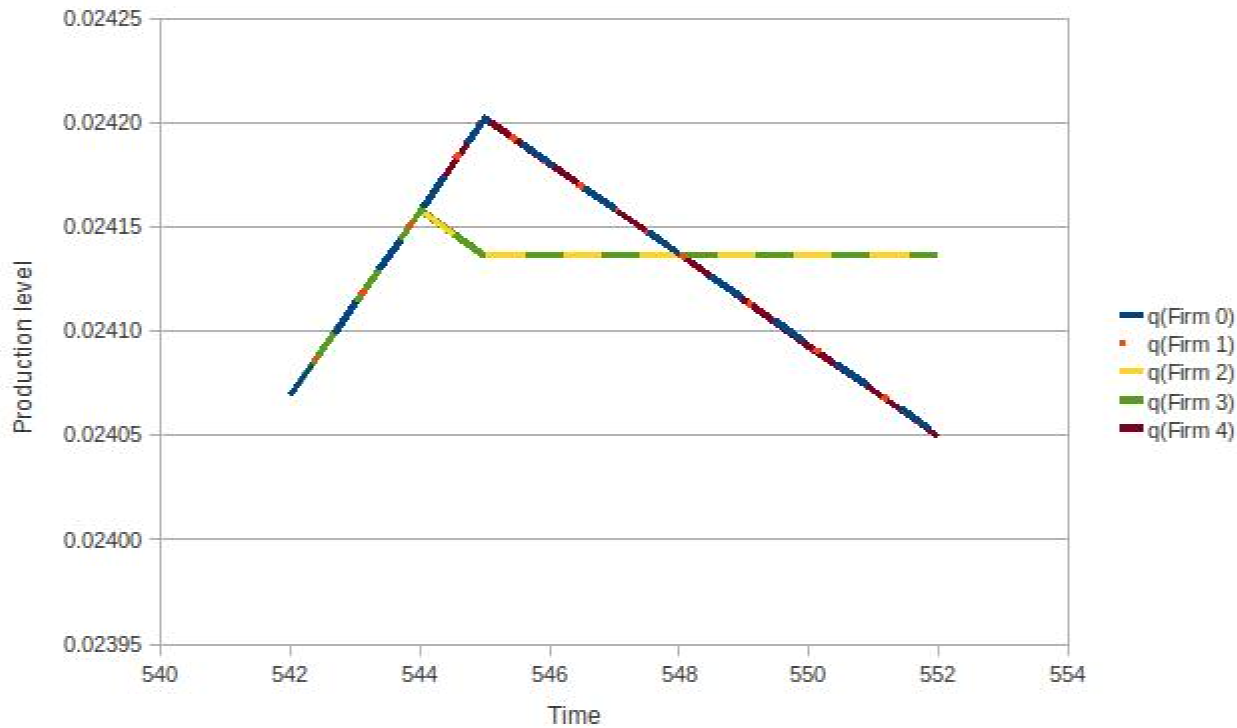

Figure 6: Peak production detail for Simulation #4.

Table 2: The Cournot-like time-line.

Time-step	Producers	
543	All	The heuristic determines that ramp-up should continue.
544	2,3	The heuristic determines that a decreasing production path is optimal.
	0, 1, 4	The heuristic determines that ramp-up should continue.
545	2,3	Even with the decrease in production by producers 2 and 3, aggregate production increased over the previous time-step. As a result, marginal price is positive, leading the heuristic to determine that a flat production path is optimal. With a flat production path, the heuristic does not update marginal price, meaning the previous, positive value is used.
	0, 1, 4	The heuristic determines that a decreasing production path is optimal.
546 onward	All	Based on previous marginal price calculations, producers continue to maintain the previous production path.

and 3 reduced production, aggregate production still increased, albeit by somewhat less. As a result, price continued to fall, which producers 2 and 3 interpreted as a positive marginal price. The positive marginal price causes the heuristic to choose a flat production path. The time-line for the heuristic decisions is outlined in Table 2.

The miscalculation of marginal price that occurs at time-step 545 depends only on a) one or more producers reaching peak production and decreasing before the other(s), and b) aggregate production increasing despite the decrease by the earlier producers. That is, the magnitude of the differences in stock level between the producers is irrelevant as long as it is sufficient to cause them to reach peak production at different time steps.

3 Additional information for producers

It is not unreasonable for a producer to know how many competitors there are in the market. This information alone, however, does nothing to improve the marginal price estimate. It may be reasonable to reject a non-negative marginal price, since a positive marginal price is contrary to the law of demand and a zero marginal price, even if correct, is not useful. Suppose, however, that a producer is able to make a qualitative estimate of the production paths of the other producers. That is, a tally of how many are increasing production, decreasing production, or maintaining current production. For example, counting the truckloads of ore departing a competitor's mine, or the number of oil wells in operation, or evidence from a factor market such as skilled labor. At the least, this could mitigate the

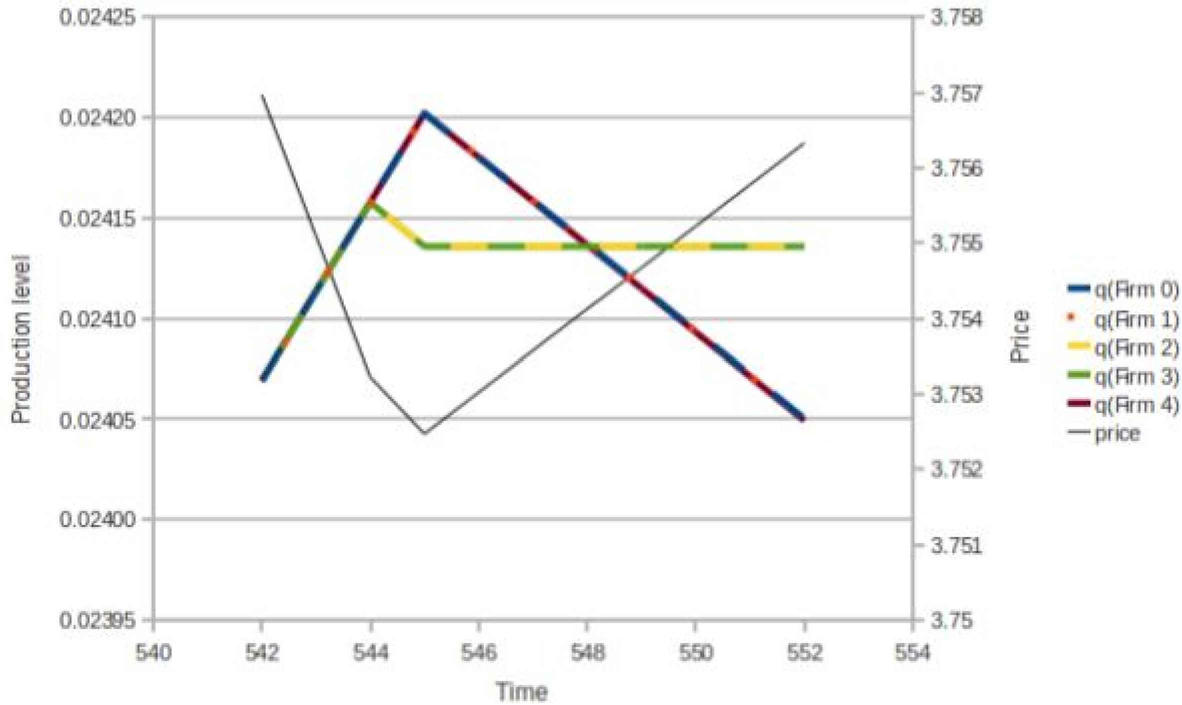

Figure 7: Peak production detail for Simulation #4 with price.

error that arose in the previous section.

The point was made that, for this particular inverse demand function and set of parameters, the cost of the error was small. That may not always be the case, but there is another reason for reducing the error in estimating the marginal price. That is, the agents will require a better estimate of the marginal price in order to form an intentional collusive or Cournot market. In other words, a producer's signal can neither be sent nor received unless both parties know what the correct self-optimizing production level should be.

The proposed additions to the heuristic are:

1. Determine the production level direction of the market and assess its strength by assigning a score of -1, 0 or +1 to each producer that, over the last period, decreased, maintained, or increased production, respectively. The scores are summed and divided by the number of producers, and this is the market *direction*. The sign of *direction* indicates the trend in aggregate production, and the magnitude of *direction*, which ranges from -1 to +1, gives the strength of the trend.

2. Remember the previous value for marginal price, and combine it with the current calculation. The sign of the current calculation is inverted if the producer is moving opposite to the market. That is, if *direction* is positive, and the producer is decreasing production, the sign on the calculated marginal price for this producer is reversed.

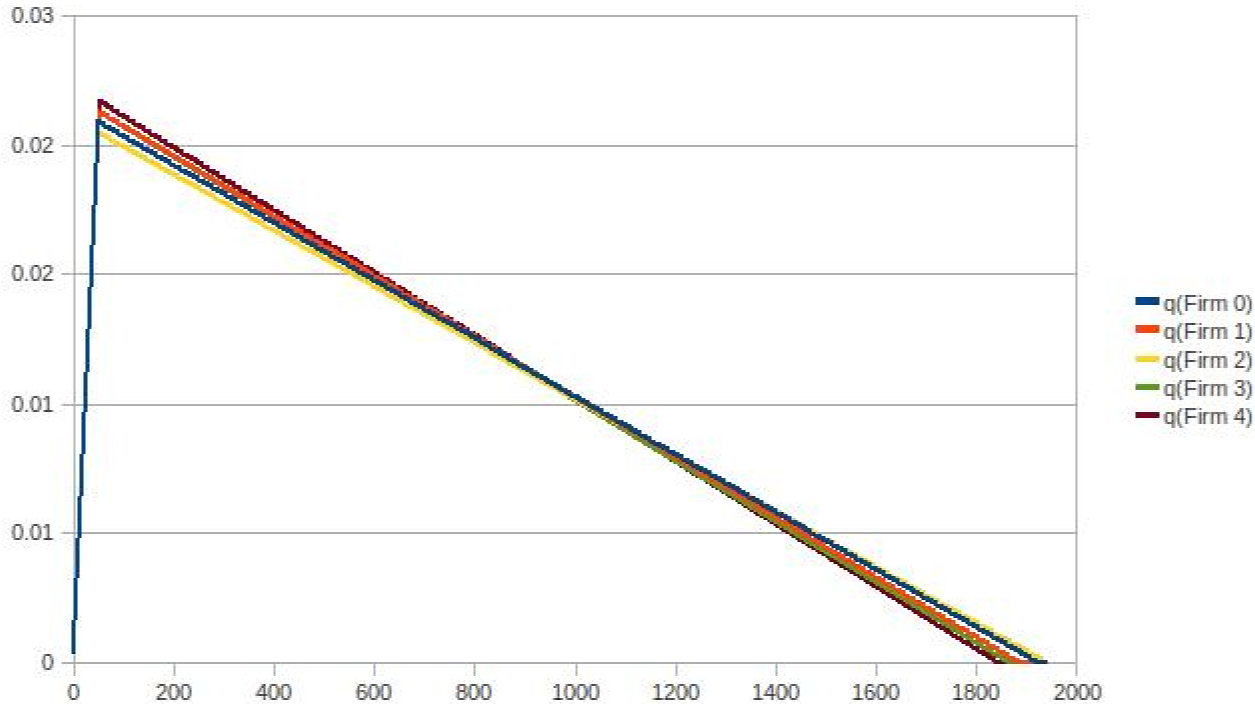

Figure 8: Better marginal price.

> The current calculation is given a weight proportional to *direction*, the previous period's marginal price is given a weight (1 - *direction*), and the two are summed to give the next marginal price.

With this minor change to the heuristic, the Cournot-like behavior in the oligopoly model disappears for all N up to 10. The behavior is distinguishable from the collusion-like outcomes because the production paths are all slightly different, as seen in Figure 8.

4 Conclusion

The collusion-like and Cournot-like behaviors reported in Dixon (2010) are shown to be a specific error in the way that agents compute marginal price. An addition to the optimization heuristic for producer agents is proposed that a) estimates the overall movement of the market and b) does a weighted average of the current period calculated marginal price and that of the previous period. This minor change eliminates the computational error, which, alas, also eliminates the emergent Cournot-like behavior. A clearer estimate of marginal price, however, opens the door for signaling between producers, making it possible to explore overt collusive and Cournot equilibria in a future paper.

References

References

David S. Dixon. Hotelling's Rule In the Limit: An Agent-Based Exploration of the Model Space. 85th Annual Conference of the Western Economics Association International, 2010.

David S. Dixon. Computational Methods in Natural Resource Economics: Agent-Based Modeling and Hotelling's Rule. Unpublished Ph.D. dissertation, June 2011.

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