Learning Curves and Cost: An Application of Optimal Control

The learning curve is a process by which a firm experiences falling average cost as output increases. The leaning curve is not merely some fixed cost effect. (A fixed cost effect means that the firm incurs expenses at the beginning of a project and the more it produces, the lower is average fixed cost.) The learning curve phenomenon means that the whole U-shaped average cost curve shifts down as accumulated output grows. There is an implicit time dimension to the problem.

Consider a non-storable good whose production is characterized by the learning curve phenomenon. Assume that demand for this good is constant through time. What we want to know is how a monopolist will behave in terms of price and quantity over time. This problem raises the issue of optimal control in a fairly simple setting, which shows the incredible complexity of optimal control problems. I think the best way to handle this problem is to do it in discrete time:

First consider a cost function in which the cumulative production in the past causes current cost to be lower. A parsimonious expression of this cost process is:

$$C_t = C(q_t, Q_t) = \frac{h(q_t)}{Q_t}, \text{ where } Q = \sum_{\tau=0}^{t-1} q_\tau \qquad 1$$

This says that cost falls in direct proportion to the amount of output that has been produced in the project life to date. Let $C_0=h(q)$. Note that given this cost function, if h is U-shaped, learning shifts it vertically downward.

The present value of profits can then be written as:

$$\Pi = \sum_{t=1}^{T} \frac{R(q_t) - C(q_t, Q_t)}{(1+r)^t} = \sum_{t=1}^{T} \frac{R(q_t) - \dfrac{h(q_t)}{\sum_{\tau=0,t-1} q_\tau}}{(1+r)^t}$$

Profits are then maximized by the choice of q_t in each of the T periods. There are T FOC all of which look like:

$$\frac{\partial \Pi}{\partial q_t} = \frac{R_{q_t} - C_{q_t}}{(1+r)^t} - \sum_{j=t+1}^{T} \frac{C_{Q_j} \dfrac{\partial Q_j}{\partial q_t}}{(1+r)^j} = 0 \qquad 2$$

Note that $\dfrac{\partial Q_j}{\partial q_t}$ is equal to 1 and C_Q is $-C(q, Q)/Q$ for all j. It is intriguing that given the fairly innocuous looking form of the cost process, the marginal effect of current production on future cost is total cost in the future period divided by cumulative output in the future period.

Equation (2) can be rewritten as:

$$R_{q_t} = C_{q_t} - \sum_{j=1}^{T-t} \frac{C_{t+j} / Q_{t+j}}{(1+r)^j}$$

which shows that the production rule in any period is to equate marginal revenue in that period with current marginal cost minus the discounted value of the effect of current production on future cost in each future period. It is informative to look at these FOC at a couple of points in time.

Consider the final time period:

$$\frac{\partial \Pi}{\partial q_T} = \frac{R_{q_T} - C_{q_T}}{(1+r)^T} = 0$$

This says that in the final period profits are maximized by equating current marginal revenues with current marginal costs.

Consider the period just prior:

$$\frac{\partial \Pi}{\partial q_{T-1}} = \frac{R_{q_{T-1}} - C_{q_{T-1}}}{(1+r)^{T-1}} + \frac{C_T / Q_T}{(1+r)^T} = 0$$

This can be rewritten as:

$$R_{q_{T-1}} = C_{q_{T-1}} - \frac{C_T / Q_T}{(1+r)}$$

Here the effect of current output on future cost is taken into account. The firm produces *beyond* the point where current marginal revenue is equal to marginal cost. The cost reducing effects of current output on future production are measured by C/Q and are discounted based on their one period delay.

As an example, consider the simple case where $h(q)$ is a linear function, aq. Let the discount rate be zero and let there be just two periods, 0 and 1. Given this function, marginal operating cost in the first period is a. However, the marginal opportunity cost of production is $[a–C_1/Q_1]$ which is $[a–aq_1/q_0^2]$. Marginal cost in the second period is a/q_0.

Figure 1 shows the equilibrium. Holding q_1 constant, we can graph the marginal opportunity cost of first period production as a function of first period production. The firm chooses its first period output where this cost equals marginal revenue. This is point x in the diagram. Given this level of first period production, marginal cost in the second period is determined. The firm chooses second period production where marginal revenue equals this marginal cost.[1] If marginal revenue intersects the two functions to the right of point z, then first period production is less than output in the second period, and price falls from the first to the second periods. On the other hand, if marginal revenue intersects to the left of z, the pattern is

[1] In fact, the two equations must be solved simultaneously; the description of the equilibrium in Figure 1 leaves out the iterations necessary to achieve this simultaneously solution. However, this does not affect the substantive conclusions.

reversed. The model allows for the learning curve effect to result in increasing or decreasing prices through time.

Figure 1: Learning Curve Effect with Linear Cost and Two Periods

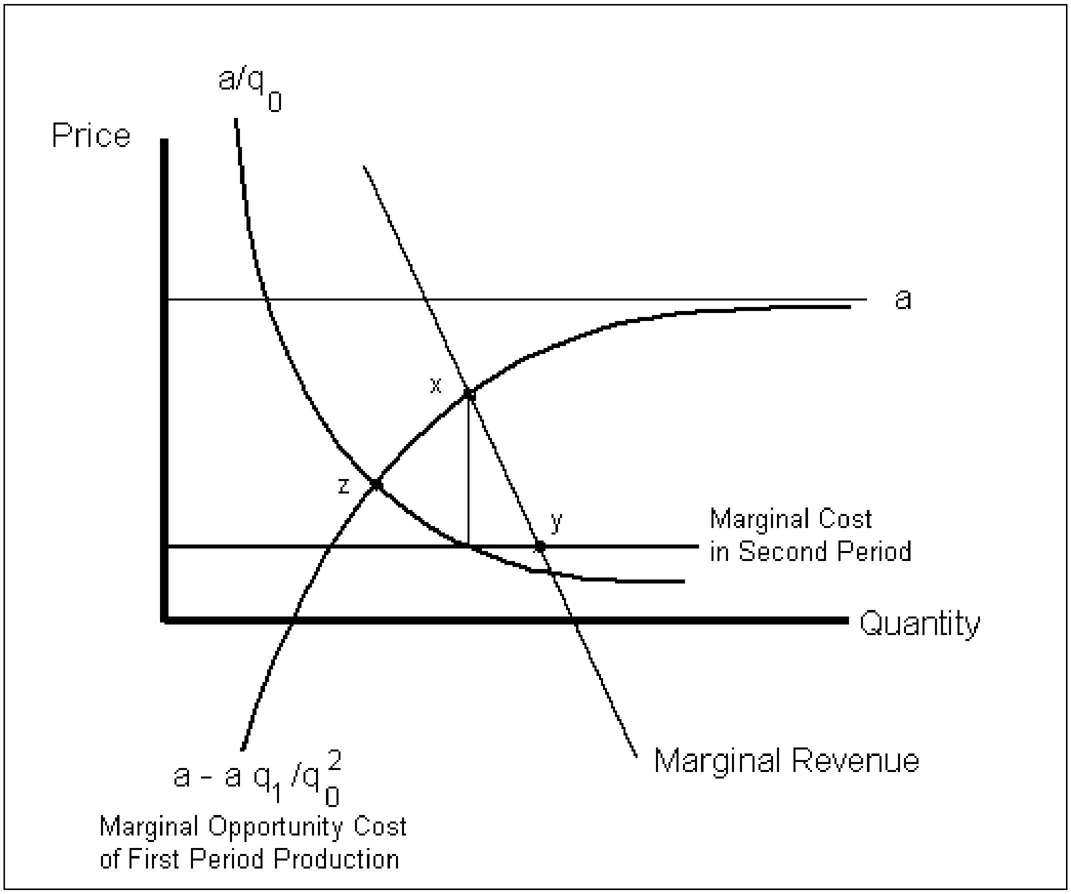

Competitive Equilibrium

Having described the learning curve in terms of monopoly behavior, we are driven to ask how such a phenomenon could exist in a competitive setting. Can there be a competitive market for a good whose production is affected by learning. Let's continue to describe the learning function as one that is linked to the cumulative production achieved to date.

There are two ways that we can imagine learning affecting a competitive industry. One is that the cost curves of the firms in the industry fall based on the total amount of output that has ever been produced by all the firms in the industry. This is simple and simplistic. Given this assumption the cost curves shift down and we have a standard competitive equilibrium in each period that occurs at a lower price.

Alternatively we can assume that learning depends on the level of output produced by each firm in the industry. For instance, a new firm has to begin its learning from scratch. This is a tricky question but the answer goes something like this: The number of firms necessary to drive price to min AC in period one may be too many or too few to equate price to min AC in period two. If too many, some exit (no problem if resources are mobile). If too few, then firms anticipate that zero profits in period one will imply excess profits in period two. Entrepreneurs in period one recognize and respond to this. Price is driven below min AC in period one in order to increase the number of firms in competition in period two. The same thing may be true between period two and three, and three and four, etc. The equilibrium is that the number of firms in the market producing in order to learn so that their costs are lower in the future is determined by discounting the economic profits at the expected competitive equilibrium in every period back to the initial period. Competition will force this discounted excess profit to equal zero in the beginning. Through time we see the competitive price falling and market output increasing just like happens in the case of the monopoly.

The best way to display this result is graphically in a two-period model with a discount rate of zero.

SAS program to generate Figure 1.

```
data one ;
x=15; a=5; b=-2;
pistar=0; q1star=0; q0star=0;
do q0=.01 to x by .01;
do q1=.01 to x by .01;
p0=x+b*q0;
p1=x+b*q1;
r1=p1*q1; r0=p0*q0;
c0=a*q0;
c1=a*q1/q0;
pi=r0+r1-c0-c1;
if  pi>pistar then do;
pistar=pi; q1star=q1; q0star=q0;
end;
end; end;
do q0=.01 to x by .01;
mc1=a/q0;
mc0=a-a*q1star/q0**2;
mr=x+2*b*q0;
output; keep q1star q0star q0 mc0 mc1 mr pistar
x a b;       end;
run;
proc means; run;
proc gplot;
symbol i=join;
plot (mc0 mc1 mr)*q0/overlay;
where mc0>-x and mc1<x and mr>-x;
run;
proc gplot;
symbol i=join;
plot (mc0 mc1 mr)*q0/overlay;
where q0>2 and q0<4;
run;
```