Flexibility and Endogenous Innovation

Robert M. Solow

ABSTRACT. The trail I want to follow leads from Stigler to Mansfield with few, if any, intermediate stops. That should sound like a pretty good genealogy.

I

Long ago (1939 to be exact) George Stigler published an important article called "Production and Distribution in the Short Run" (*Journal of Political Economy*, XLVII (1939), 305–27). I want to extract just one simple but fruitful point from many. Imagine a firm that is about to choose among several available technologies, build the suitable plant, or otherwise fix its method of production for some appreciable time to come. Focus on two alternatives. One of them yields an average cost curve that has a very low unit cost at its minimum point, but rises sharply at smaller or greater outputs. The other can not achieve such a low unit cost anywhere. Its minimum average cost, which we can take for simplicity (and to avoid irrelevant questions about scale) to occur at the same output as the first, is noticeably higher; but the average cost curve is much flatter, so that it achieves a lower unit cost than the first technology for outputs that are some distance from the least-cost level.

It makes sense to say that the second technology is more flexible than the first. A firm choosing between them would take the first technology if it were pretty sure that it would be producing an amount close to the least-cost output most of the time. A firm with greater uncertainty about its output pattern over time would be more likely to prefer the second, more flexible, technology. Of course the firm *chooses* its output each period, so that it is not really satisfactory to make the distinction in those terms. What one really has in mind is that the flexible technology is better for a firm facing highly variable demand conditions; the sharply targeted technology is more suitable for a firm that expects to face a stable demand.

That is simple and straightforward, and does not need even an example to clinch the point. But I want to provide an example anyway, precisely because it is elementary-textbook stuff, and can serve as an analogy when we come to a more complicated setting. I do not even have to make the contrast between the two cost curves very pronounced, though of course it could be. So, suppose the targeted technology has a total cost curve $t(x) = a + bx^2$, where x is output. Thus the average and marginal cost curves are $a(x) = a/x + bx$ and $m(x) = 2bx$. The least-cost output is $x = (a/b)^{1/2}$, and at that output unit cost is $2(ab)^{1/2}$. Figure 1 gives the familiar picture: $a(x)$ goes to infinity as x vanishes, and is asymptotic to the ray bx as x gets large.

For the flexible technology, I might as well go whole hog and make average and marginal cost constant at c. So $t(x) = cx$. It is only necessary that $c > 2(ab)^{1/2}$, or else the flexible technology dominates its rival. For an elementary reason that will be obvious in a moment, I suppose that the flexible technology has an inviolable capacity limit at $x = M$. So average cost is horizontal for $x < M$ and turns vertical at $x = M$.

Now suppose the firm is a price-taker. If it adopts the targeted technology, it will not produce unless p exceeds $2(ab)^{1/2}$, but at higher prices it will produce $x = p/2b$ and earn positive profits of $p^2/4b - a$. If it uses the flexible technology, it will not produce at all unless p exceeds c, but at higher prices it will produce at capacity and earn profits $(p-c)M$. (Of course this is why we need a finite capacity. If the flexible

Massachusetts Institute of Technology,
Department of Economics
50 Memorial Drive
Cambridge, MA 02142
U.S.A.

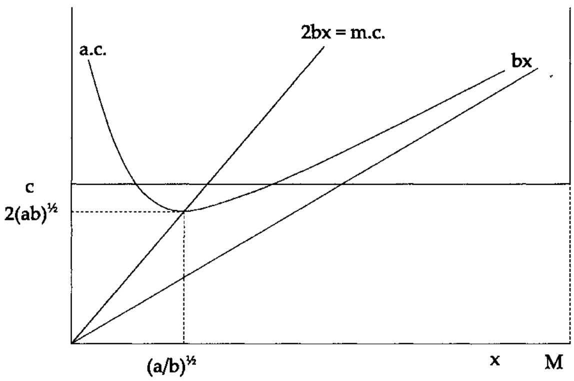

Figure 1. Targeted and flexible static average and marginal costs curves.

technology came with a flattish saucer-shaped average-cost curve instead of a horizontal line, there would be more arithmetic to do but the broad picture would be similar. It is not worth the trouble.)

The firm's decision problem turns on its *ex ante* probability distribution of the market price. (Under imperfect competition, the firm would have to have an *ex ante* joint distribution over at least two parameters of its demand curve, one for level and one for slope or elasticity.) Under risk neutrality, it will choose the technology giving the higher expected profit. This cannot depend solely on the variability of price, because any location parameter of the price distribution will also be relevant. (For example, if the largest possible price exceeds $2(a/b)^{1/2}$ but falls short of c, the flexible technology loses regardless of variability.) The point of the exercise, however, is not any such simple theorem, but rather a feeling for the way the story goes.

There is no need actually to carry out the expected-profit calculations. The situation can be easily read from Figure 2. We already know the *profit function* of each of the eligible technologies. For the targeted technology it is $\pi(p) = p^2/4b-a$ if $p > 2(ab)^{1/2}$, and zero otherwise. It is plotted in Figure 2. For the flexible technology, $\pi(p) = (p-c)M$ if $p > c$ and zero otherwise, and it is plotted in the same diagram. The advantage of one choice over the other is easily read off for each possible price. Any notion of likely price behavior over time can be translated into a choice of one technology over the other.

There is one qualification, however. From the diagram, it looks as if the profit function for the targeted technology will eventually overtake the other, so it will be preferred for all very high prices as well as for low prices. But this apparent paradox is an artifact of the arbitrary capacity limit M imposed on the flexible technology. From general duality theory, or direct calculation, we know that $\pi'(p) = x$; the slope of the profit function at any price is the profit-maximizing output at that price. Consider the point in the diagram where the targeted profit function is momentarily parallel to the flexible one. That is to say, its slope is M (and increasing). At that price, a firm operating the targeted technology would already be producing and selling M units of output. So the targeted technology's profit function only begins to catch up with its rival at the rival technology's capacity output. If the makeshift flexible technology, with a flat average cost and an absolute capacity limit, were replaced with the shallow saucer-shaped average cost curve it is intended to suggest, the capacity limit could be dispensed with. The apparent anomaly would disappear and the Stigler calculation would be routine.

II

Examples like the one just examined can be generated from a standard two-factor production function with one factor fixed and the other variable. The details may not always be so convenient, but the qualitative properties can be preserved. In fact, the quadratic-total-cost case

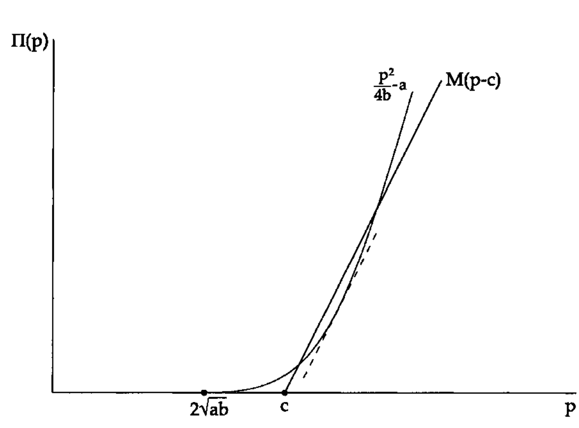

Figure 2. Profit functions associated with targeted and flexible cost curves.