What's Experience Got to Do With It? Sources of Cost Reduction in a Large Specialty Chemicals Producer

Gavin Sinclair • Steven Klepper • Wesley Cohen
School of Technology, Purdue University, West Lafayette, Indiana 47907
College of Humanities & Social Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
College of Humanities & Social Sciences, Carnegie Mellon University, Pittsburgh, Pennsylvania 15213
gwsinclair@tech.purdue.edu • sk3f@andrew.cmu.edu • wc02@andrew.cmu.edu

Conventional learning curves relating unit cost to measures of production experience are estimated for 221 specialty chemicals produced by a Fortune 500 company. Detailed records on cost and R&D coupled with insights from company personnel are used to explain the variation across products in the rate of cost reduction. Products that exhibited the strongest relationship between unit cost and measures of production experience were subject to specific initiatives, particularly process R&D. The R&D was not, however, generally motivated or informed, by production experience. However, cumulative past output, the most commonly used measure of production experience, was related to expected future output, which conditioned the expected future returns from R&D and the choice of R&D projects. Thus, cumulative output was connected to unit costs through its role in conditioning incentives to undertake process R&D rather than as a proxy for production experience. This suggests that the strong relationship commonly found between unit cost and measures of production experience may reflect incentives to reduce cost as much as learning from production experience.
(*Organizational Learning*; *Learning Curves*; *Process Innovation*)

1. Introduction

It has repeatedly been found that a product's unit cost tends to fall by a constant percentage with each doubling of cumulative output or time in production. Not only does this hold for a wide range of products and methods of production (Dutton and Thomas 1984), but it persists when controls are introduced for the scale of production and the amount of effort devoted to process engineering and R&D (cf. Lieberman 1984, Adler and Clark 1991, Jarmin 1994b). The ubiquity of this relationship, known as the "learning" or "experience" curve, has led many to infer that production experience causes cost reduction. Building on the idea of progressive cost reduction as a by-product of production, economic theorists have claimed "learning by doing" to be a source of increasing returns and consequent first-mover advantages in models of firm output strategy, industry evolution, international trade, and national economic growth (Arrow 1962, Spence 1981, Fudenberg and Tirole 1983, Krugman 1987, Dasgupta and Stiglitz 1988, Lucas 1993, and Jarmin 1994a).

Although the learning metaphor is attractive for the direct link it provides between production and costs, as Lucas (1993) and others acknowledge, the statistical studies relating unit costs to measures of production experience, such as cumulative output, do not provide

Management Science © 2000 INFORMS
Vol. 46, No. 1, January 2000 pp. 28–45

0025-1909/00/4601/0028$05.00
1526-5501 electronic ISSN

any direct evidence about actual learning or about how production experience contributes to cost reduction. Numerous investigations, some involving case studies of particular production processes (e.g., Stern 1933, Hollander 1965, Enos 1967) and others examining the War World War II products that launched the learning curve literature (Argote et al. 1990, Bell and Scott-Kemis 1990, Mishina 1992), have probed the sources of cost reduction. The statistical studies take as their point of departure the close link between measures of experience and cost, exploring statistically the pathways by which experience could inform cost reduction. Alternatively, the case studies focus on the actions that reduce cost and find that cost reduction is realized chiefly through technical improvements developed by specialized engineering and R&D personnel. If these improvements were motivated by production experience, as the statistical studies would suggest, then controlling for engineering efforts and process R&D should largely account for the statistical relationship between measures of experience and unit cost. The statistical studies that introduce such measures, however, generally find them to have the wrong sign and only reinforce the effect of production experience on cost (Adler and Clark 1991, Mishina 1992, Jarmin 1994b). Thus, we are left with little consensus about why unit cost and measures of production experience are related and what role learning from production experience plays in cost reduction.

The purpose of this study is to probe the sources of the relationship between experience and manufacturing cost by analyzing in depth the process of cost reduction in the specialty chemicals division of a Fortune 500 company. A distinguishing feature of our analysis is the wealth of quantitative and qualitative information available to us. We had access to data on manufacturing costs, output, R&D, and related items for over 1,000 specialty chemicals products manufactured over a two and one half year period. Just as important, we had access to personnel and documentation on the firm's operational, managerial, and accounting practices. We used this information to estimate conventional learning curves, and exploited insights from the firm's personnel coupled with R&D and related records to explore why learning rates varied greatly across the set of products that were studied. We anticipated that if we could account for the sources of the variation in the learning rates then we could infer the principal factors connecting unit costs and production experience.

Specialty chemicals are produced using a capital-intensive method of production in which, for safety reasons, production workers are heavily circumscribed in the modifications they can make in the production process. Consequently, one route by which learning can occur, namely through production workers, was largely foreclosed. Nonetheless, for the specialty chemicals in which there was an incentive to economize on the use of production equipment, we found conventional patterns relating unit cost and production experience. Looking across products, we found the strongest relationship between unit costs and production experience for products subject to process R&D. Whether a product was subject to process R&D had little to do with experience gleaned from production, though. Moreover, the R&D itself was not generally informed by production experience but drew upon a well-developed science and engineering knowledge base. Further analysis suggested, however, a novel link between the most commonly used measure of production experience, cumulative output, and cost reduction. Cumulative output was related to expected future output, which conditioned the expected returns from R&D and in turn the choice of R&D projects. Consequently, the products with greater cumulative output were more likely to be subject to R&D and, hence, cost reduction. Thus, cumulative output was connected to unit costs through its role in conditioning *incentives* to engage in process R&D rather than operating as a proxy for experience in production.

Incentives to invest in process innovation have not been featured in the literature on the learning curve. The virtue of attention to incentives is that it casts purposeful R&D and engineering activities in a central role in reducing costs, thus linking learning to the larger literature on R&D and technological change. Indeed, the major conclusion of our study is that what looks like learning from production experience may reflect the outcome of an incentive-driven R&D driven process in which learning from production plays a minimal role. More generally, to the extent that production experience, in the form of cumulative output, is related to cost reduction, part of its role may come from its connection to incentives to conduct R&D rather than learning per se. This is consistent with Cohen and Klepper's (1996a, 1996b) findings concerning the source of the tight link between firm R&D spending and firm size. These findings suggest that the greater the level of output then the greater the base over which the knowledge produced from a given R&D effort can be applied and thus the greater the firm's incentives to invest in R&D.

The paper is organized as follows. Section 2 reviews past studies and highlights some of the key questions they raise. Section 3 describes the organization studied and the data. Section 4 reports the results of fitting learning curves for 221 specialty chemicals and explores the factors accounting for variation in the learning rates across products. Section 5 probes the role of R&D in cost reduction. Section 6 places our findings in the broader context of the questions raised in §2. Concluding remarks are offered in §7.

2. Past Findings and Current Questions

Studies of progressive cost reduction have been of two types, statistical and more qualitative case studies. The basic model employed in the statistical studies is:

$$C_t = \alpha Q_t^{\gamma} X_t^{\beta} \epsilon_t, \tag{1}$$

$$X_t = Q_{t-1} + X_{t-1}, \tag{2}$$

where C_t is average cost (or some variant, such as labor per unit of output or total factor productivity) in period t; Q_t is the quantity of output in period t; X_t is experience in period t, α, γ, and β are coefficients; and ϵ_t is a disturbance with a mean of one for all t. Equation (1) relates average cost to the quantity of output and past experience. Equation (2) specifies past experience as cumulative output through the last period. The basic model captures the idea that learning from production experience occurs through information gained from the ongoing process of production. Such information is either of immediate help in reducing costs, or provides feedback about deficiencies in production that conditions efforts by manufacturing or other units to modify the production process in order to reduce cost.

Many variants of the basic model have been estimated. Other determinants of average cost, such as process R&D and labor training, have been used and alternative measures and functional forms for experience have been tried. Nearly all the statistical studies find some measure of production experience to be related to average cost. Most use cumulative output as the measure of experience. Argote et al. (1990) generalize Equation (2) as $X_t = Q_{t-1} + \lambda X_{t-1}$ to allow the information generated by continued production to depreciate, where $1 - \lambda$ is the rate of depreciation. Mishina (1992) specifies Equation (2) as $X_t = \max\{Q_1, Q_2, \ldots, Q_{t-1}\}$ to capture the idea that new information used to reduce cost is only generated when the scale of production exceeds its prior maximum. Both studies find a modest improvement in fit with these alternative specifications. The predominant interpretation of the findings of the statistical studies is that learning from continued production is an important source of cost reduction.

Many statistical and case studies have probed the link between cost reduction and production experience. Considerable variation in learning rates across products (Dutton and Thomas 1984) and across producers of the same product (cf. Alchian 1963, Jarmin 1994a, 1994b) suggests that cost reduction results from something more complicated than production experience alone. Moreover, some of the studies cast doubt on the most direct route through which production experience is believed to reduce costs, namely through production workers. Lieberman (1984), for example, finds considerable rates of cost reduction associated with cumulative output in the nonlabor intensive production processes employed by commodity chemicals firms. Argote et al. (1990) find that labor turnover did not affect the average cost of Liberty ships even though turnover should disrupt cost reduction if labor experience is the key to increasing efficiency.

Numerous statistical and case studies suggest that what appears to be cost reduction as a result of continued production of a product actually reflects the application of expertise culled from sources other than experience in producing the affected product. For example, a number of learning curve studies reveal transfer of learning across related facilities, such as World War II shipyards (Argote et al. 1990), or across establishments owned by the same firm (Darr et al. 1995) or those in close proximity to one another (Jarmin 1994b). A number of case studies also highlight the role of expertise coming from training and broad technical experience. For example, detailed case studies of production, such as Stern (1933) for four tire plants and Hollander (1965) for Dupont's rayon plants, find that progressive cost reduction results largely from numerous incremental technical changes in the production process that are typically developed and implemented by specialized engineering and R&D personnel.[1] While the case studies are not incompatible with experience in production somehow conditioning efforts to improve the production process, they suggest that cost reductions are directly attributable to process innovation and not to production experience per se. On the other hand, the statistical studies of cost reduction that include quantitative measures for engineering efforts, training of labor, and R&D do not consistently find a link between these measures and unit costs. If anything, the inclusion of these measures strengthens the relationship between unit costs and cumulative experience (Adler and Clark 1991, Mishina 1992, Jarmin 1994b).

Thus, the literature leaves matters in an unsettled state. If cost reduction arises principally from investments in process innovation, why do variables that control for such investments not explain the observed cost reductions? Moreover, if the critical knowledge comes from any source other than continued production, how can the tight, robust relationship observed between cumulative output and cost reduction be explained? To consider these questions, we turn to our business unit.

[1] In their review of both the statistical and qualitative literature on cost reduction in World War II products, Bell and Scott-Kemis (1990) reach a similar conclusion.

3. Production Process and Data

Our analysis focuses on the specialty chemicals business unit of a Fortune 500 company. The business unit has two plants in the United States and small foreign operations. Only the two U.S. plants, which account for 90% of total volume, were included in the analysis.

The two plants produce over 1,000 specialty chemicals using a batch process. Raw materials are introduced in a reactor and then processed through one or more chemical steps, such as reaction or filtering. Samples are taken of the finished product to check it against specifications, and intermediate samples are often taken to ensure proper progress of reactions. Wastes from the production process must be treated before disposal. The finished product is pumped out and packaged into drums or stored in bulk. If there is any change in the specifications of a product, the business unit assigns it a new code number that identifies it as a new product. Products are classified into thirteen product lines and a miscellaneous category. The product lines are distinct in that each is produced in a specific area of the two plants based on the technology required to produce it, with production workers and even R&D personnel specializing by product line.

Production is scheduled according to an inventory tracking system. A product can be made in more than one reactor, so production scheduling takes into account reactor availability. Reactor sizes and batch sizes for a given product can vary considerably, with an important effect on unit costs. Apart from increases in batch size or changes in the reactor vessels, reductions in manufacturing costs come principally from production changes such as reducing reaction times or eliminating a process step, which conserve on the use of production equipment. Changes can be suggested by anyone, including production workers, marketing personnel, and technical personnel, although for safety reasons production workers are heavily circumscribed in the changes they can implement on their own. Many changes result from formal R&D projects. Every R&D project is prioritized based on technical feasibility and projected profit.

The data come principally from production log sheets. For each batch of each product, the log sheets record the date of the batch; the number of the batch; the quantity produced; the reactor(s) used; the procedure(s) used; the quantity of each raw material; the standard cost of each raw material; the time spent on each piece of equipment; the number of samples taken; the quantity of waste treated; and the standard cost of each piece of equipment, sample, and waste treatment. Standard costs for materials are established at the beginning of a year based on the prior year's unit cost and a forecast for the next year. Standard costs for equipment are computed by dividing the total cost of the equipment, including energy, labor, and depreciation by the number of manufacturing hours at 100% capacity utilization. Standard sample costs and waste treatment costs are computed similarly. Price data are not recorded for each batch. Instead, invoices are accumulated and an average monthly price is calculated for each product.

The business unit has maintained a consistent cost accounting system since January 1990, which begins our sample period; the period ends June 1992. During this 30-month period, the business unit manufactured 1,026 products in a total of 13,076 batches. There were 12 products with over 100 batches and 847 with less than 20 batches. To fit learning curves with some reliability, we analyzed only the 221 products with at least ten nontrial batches (trial batches are small, experimental batches that were analyzed separately). These 221 products accounted for 7,040 of the 13,076 batches, for an average of 31.86 batches per product. The starting date of production for each product was known. Fifty-eight of the products were produced less than five years, 59 were produced for over 10 years, and 15 products were introduced during the sample period.

4. Analysis of Learning Rates

The first step in analyzing learning rates across the 221 products was to fit learning curves for each product. We then analyzed the primary factors that accounted for variation across the products in estimated learning rates.

4.1. Estimation of Learning Curves

For each product, we fit the model described by Equations (1) and (2) using all the batches produced of the product as observations. The dependent variable, C_t, was the unit manufacturing cost of the batch produced in time t. This included the cost of all manufacturing steps, including labor and equipment, but excluded raw material costs.[2] The scale of output, Q_t, was the number of units of output of the batch. The measure of experience, X_t, was the time between when the product was first produced and the time of the batch. We also conducted the analysis with X_t measured in terms of past cumulative output and the number of past batches produced, which required estimation by the production personnel of the output and number of batches of each product prior to the sample period. We focus on the results with X_t measured as the time elapsed since production began because we did not need to estimate this variable and because it provided a slightly better fit than the alternative measures.

The model was estimated using ordinary least-squares regression after transforming equation (1) by taking logs of both sides. It was estimated in two ways. In the first, both γ and β were estimated. In the second, γ was constrained to equal -0.4 based on the six-tenths power rule. Like many types of equipment, the cost of chemical processing equipment rises less than proportionally with the capacity of the equipment, giving rise to scale economies in production. The six-tenths power rule predicts that total costs rise as a function of output raised to the 0.6 power, which implies a value of γ equal to -0.4. Table 1 reports the distributions for the 221 products of the unconstrained estimates of γ and β in Columns 1 and 2 and the constrained estimates of β in Column 3. For the unconstrained estimates, the number of significant estimates (0.10 level, two-tailed) and the mean and standard deviation of the R^2 for the 221 regressions is also reported.

Focusing first on the (unconstrained) estimates of γ in Column 1, nearly all the estimates are negative, indicating that unit costs were lower for bigger batches. The mean estimate of γ is -0.58, which is

[2] The various standard costs did not change appreciably over the sample period, so changes in manufacturing cost were driven primarily by changes in equipment usage, sampling, and waste treatment.

Table 1 Distribution of Estimates of the Model with γ Unconstrained and Constrained to Equal -0.4.

	221 Products, γ *Unconstrained*		Estimates of β with γ Constrained to -0.4							
Range	γ	β	221 Products	122 Products Made in Slack Area of Large Plant	99 Other Products	Seven Campaigned Products	13 Products Affected by Sampling Efforts	25 Products Subject to R&D	Six Products with 10 to 11 Batches	Remaining 48 Products with More than 11 Batches
<-1.4	11	16	21	2	19	3	4	9	1	2
-1.4–1.0	16	6	13	3	10	2	4	3	0	1
-1.0–0.6	87	17	16	5	11	2	0	5	0	4
-0.6–0.2	66	38	34	13	21	0	4	5	0	12
-0.2–0.2	28	52	45	31	14	0	1	1	0	12
0.2–0.6	6	30	36	23	13	0	0	2	1	10
0.6–1.0	3	23	21	19	2	0	0	0	0	2
1.0–1.4	2	19	16	11	5	0	0	0	2	3
>1.4	2	20	19	15	4	0	0	0	2	2
Mean	-0.58	0.14	0.02	0.42	-0.48	-1.41	-1.04	-1.20	0.89	0.02
Standard Dev.	0.59	1.06	1.12	0.91	1.14	0.65	0.67	1.22	1.50	0.70
#Sig (.10) Est.	145	59								
Mean R^2	0.41									
Stan. Dev. of R^2	0.26									

close to the -0.4 value implied by the six-tenths power rule. The range of estimates of γ, though, is quite wide, with the standard deviation of the estimates of γ slightly greater than the mean. This is not surprising given the small number of batches used to estimate the model for many of the products. Of the 221 estimates of γ, 145 are significant at the 0.10 level.

Unlike the estimates of γ, many of the (unconstrained) estimates of β in Column 2 of Table 1 are positive, with the mean estimate of β positive. This indicates that for a majority of the products, after controlling for scale effects, unit manufacturing cost actually rose over time. This is quite different from most estimates reported in the literature, which of course are for a very selected set of products. As indicated below, however, a large part of this is due to the presence of excess capacity in part of the large plant, which distorted accounting costs and undermined incentives to keep down costs for a majority of the products. The range of estimates of β is wide, with 67 of the estimates less than -0.2, indicating considerable cost reduction.[3] Of the 221 estimates of β, 59 are significant at the 0.10 level.

The mean R^2 of the regressions is 0.41. Thus, almost half the variation in unit manufacturing costs across batches can be accounted for by the batch volume and experience variables. As might be expected based on the estimates of γ and β, most of this is due to the batch volume variable; when it is omitted, the mean R^2 drops to 0.13. This is low compared to most other studies. It reflects the considerable random variation in unit costs from batch to batch. Most other studies abstract from this by computing unit costs over an arbitrary time interval such as a week or month, which would average out much of the variation in our sample at the batch level.

The distribution of the estimates of β for γ constrained to equal -0.4, which are reported in Column 3 of Table 1, are similar to the unconstrained estimates

[3] To calibrate this, a value of β of -0.32 corresponds to unit costs falling by 20% with each doubling of cumulative output.

of β. The biggest difference between the two sets of estimates is that the number of large negative estimates in the top two ranges increases when the constraint is imposed. This is reflected in the mean estimate of β, which falls from 0.14 to 0.02 when the constraint on γ is imposed. In the remainder of the paper, we analyze the estimates of β with the constraint on γ imposed. The results were quite similar using the unconstrained estimates (Sinclair 1994).

4.2. Explaining the Variation in the Estimates of β

Table 1 shows tremendous variation in the estimates of β across the 221 products. In light of the small number of batches for many of the products, one source of this variation is surely sampling error. To gain insight into whether there were systematic factors behind the variation, a preliminary analysis of the estimates was done. Regressions indicated that the estimates were significantly more negative for products produced in the smaller plant, for products produced in parts of the larger plant that were operating near 100% capacity, and for one particular product line (Sinclair 1994, p. 108).[4] We asked the personnel in the business unit for possible reasons for these patterns and whether other factors were relevant that we had not considered in our preliminary analysis. The personnel emphasized the importance of four factors, which are considered in turn.

The first factor was the location where the product was manufactured. During the sample period, the capacity of some parts of the large plant became heavily strained. Products produced in these areas received disproportionate attention from plant personnel. As a consequence, batches of products produced in the less strained slack area of the plant were often left in reactors well after reactions were complete to allow plant workers to attend to products in the strained area of the plant. Since charges increase with time in the reactor (even though the reactor may not be operating nor otherwise needed), the unit cost of products produced in the slack area of the large plant tended to rise over time. In total, there were 122 products produced in the slack area of the large plant. The batch logs indicated that the reactor times of many of these products did rise over time, causing the recorded costs of the products to rise.

The rise in recorded costs is an accounting phenomenon; the opportunity cost of using reactors with excess capacity is zero, hence economic costs for these products did not necessarily rise over time. Since the learning curves are estimated with the accounting data, however, the estimates of β for these 122 products would be expected to reflect the rise in their accounting costs over time. This is confirmed in Table 1, where the distribution of the estimates of β is reported separately for the 122 products (Column 4) and the other 99 products (Column 5). Consistent with the rise in costs reflected in the batch records, most of the estimates of β for the 122 products are positive and the mean estimate is 0.42, indicating that unit costs tended to rise over the sample period for these products. In contrast to the 122 products, most of the estimates of β for the other 99 products are negative, with a mean estimate of -0.48 and 35 of the estimates significant at the 0.10 level.[5] Compared to the modal estimate of β of -0.32 reported in the literature,[6] the estimates for the 99 products exhibit at least as fast a rate of cost reduction as the typical product reported in the literature. Thus, while our initial results for all 221 products are inconsistent with those found more broadly in the learning literature, that inconsistency was resolved once we accounted for the circumstances of production and the nature of the firm's cost accounting system.

The second factor noted by the business unit personnel explained the large negative estimates of β for

[4] The date the products were first manufactured was also included as an explanatory variable, but it was not significant. This suggests that the potential for learning in the older products was not exhausted prior to the sample period.

[5] The differences in the estimates of β for the two groups of products could have been addressed by estimating β for the 122 products using data on all costs other than time in the reactors. However, time in the reactors was the most important determinant of cost and cost reduction for the 99 products, so that estimates of β for the 122 products excluding time in the reactors would not have been comparable to estimates of β for the other 99 products.

[6] This estimate of β corresponds to unit costs falling by 20% with each doubling of output, which is the modal pattern discussed in the literature (cf. Dutton and Thomas 1984).

the one product line that stood out in our preliminary analysis. Over the sample period, a large volume product was introduced that required the biggest reactor, which displaced to a smaller reactor seven products in the group of 99 that fell within one product line. To minimize the effect on unit costs, these seven products were "campaigned," which involved producing them in consecutive batches in the same reactor to economize on the time required to clean the reactor (which is included in the cost of the reactor). After controlling for the change in reactor size through the batch size variable, this would be expected to cause the unit costs of the products to fall over time. This accords with the estimates of β for the seven products. The distribution of these estimates is reported in the sixth column of Table 1. All seven are in the top three intervals in the distribution, with a mean estimate of β of -1.41 compared to the mean estimate for the 99 products of -0.48.

The next factor mentioned by the business unit personnel was a business unit initiative motivated by the practice of quality circles. The business unit occasionally formed small teams to investigate and improve certain areas of the business. During the sample period, a team was formed of two R&D employees, the plant manager, and a business manager to look at ways to reduce the number of in-process samples that were taken during the manufacturing process to ensure proper progress of reactions. They developed a statistical process control to provide additional information about processing variables such as reactor temperature and pressure. They used this information to analyze the products which had the most in-process samples according to the batch records. It was found that samples for some of these products were unnecessary because the underlying reactions always progressed smoothly. Samples could be restricted to key decision times, which made it possible to reduce the number of in-process samples for 13 products in the group of 99. The batch logs indicate a sharp reduction in sampling costs for these 13 products after the new system was implemented. As would be expected, this is reflected in the estimates of β for these 13 products. The distribution of these estimates is reported in the seventh column of Table 1. The estimates are quite negative, with eight of the 13 estimates in the top two intervals and the mean estimate for the 13 products equal to -1.04 versus -0.48 for the 99 products.

The last and most important factor mentioned by the business unit was R&D. The technical reports of the business unit indicated 32 of the 221 products were affected by R&D efforts during the sample period. Twenty-five of them were in the group of 99, with none of the 25 affected by the campaigning and in-process sample initiatives. The distribution of the estimates of β for the 25 products is reported in the eighth column of Table 1. These estimates are again concentrated in the top intervals of the distribution, with a mean of -1.20 compared to the overall mean of the 99 products of -0.48.

In total, 45 of the 99 products were affected by the campaigning and in-process sample initiatives and by R&D. Among the remaining 54, six had only 10 or 11 batches produced during the sample period. This barely met the minimum of 10 required to be included in the sample, which is not many observations to estimate β. The distribution of the six estimates is reported in the ninth column of Table 1. Consistent with the estimates being subject to considerable sampling error, the six estimates were clustered at the extremes in the distribution, with four in the bottom two intervals and one in the top. We decided to separate these six products along with the other 45 subject to campaigning, the sampling initiative, and R&D. The distribution of the remaining 48 of the original 99 products is reported in the last column of Table 1. If cost reduction was confined to the products subject to campaigning, the sampling initiative, and R&D, then the distribution of the estimates for the other 48 products should be concentrated near zero. The mean of the 48 estimates is 0.02, with 34 of the estimates lying in the three middle intervals for β spanning -0.6 to 0.6. If the dispersion in the estimates is interpreted as a reflection of sampling error, the unit costs of the 48 products would appear to have remained constant over the sample period.

Thus, it appears that a considerable part of the variation in estimated learning rates for the products can be explained by four factors: a divergence between accounting and economic costs, campaigning, the sampling initiative, and R&D. The latter three factors appear to be the main driving forces behind progressive cost reduction. Among these three factors, R&D is the most significant. It affected more products than campaigning and the sampling initiative combined. Furthermore, the products affected by R&D exhibited very high rates of cost reduction, with the estimate of β less than -0.6 for 17 of the 25 products (versus 40 out of the full group of 99). Moreover, the new statistical control process developed in the sampling initiative was the result of a research effort in which samples were restricted to key decision times. Thus, for our sample of specialty chemicals, cost reduction was largely driven by R&D.

Production experience may have still played a central role in cost reduction if it influenced the products subjected to campaigning, the sampling initiative, and R&D, or if it shaped the nature of these activities. In the case of the seven campaigned products, production experience associated with the seven products was not significant. The impetus for the campaigning was the introduction of the new product that displaced the seven campaigned products, and campaigning itself is a widely used strategy to minimize costs. In the case of the sampling initiative, the impetus for the initiative was a business unit quality circle program rather than any experiences associated with the products themselves. The products chosen for the initiative were ones with the highest sampling costs based on the batch records, and thus production did play some role in the sampling initiative, but experience in producing the products played no further role in the design or execution of the initiative. Regarding the R&D products, we had considerable information about the choice of products subjected to R&D and the way the R&D operated. In the next section we exploit this information to explore the role played by production experience in reducing the cost of these products.

5. R&D

We investigated the R&D process using the detailed records of the business unit and guidance from business unit personnel, particularly the technical personnel and business managers. Of the 32 products in the original group of 221 that had been subjected to R&D, two types of R&D were conducted. One type focused on improving the production process while the other involved finding new applications for existing or new products. We expected the former type to contribute to greater cost reduction by its very nature, although applications projects sometimes motivated minor process work. Consistent with our expectations, the mean estimate of β was -1.254 for the 22 products subject to process R&D and -0.323 for the 10 products subject to applications R&D. Accordingly, we focused on the 22 products subject to process R&D. Among these 22 products, the estimate of β was negative for 19 of them. For these 19 products where R&D reduced costs, we investigated how R&D operated and how R&D projects were picked, concentrating especially on the role played by production experience.[7]

5.1. R&D and Cost Reduction

We attempted to discern the effects of R&D on costs by comparing batches just before and after the initiation of the R&D projects. This was not generally revealing. One reason was that the bulk of the R&D work was often performed well after the initiation of the projects,[8] and no record was maintained of the progress of the projects.[9] A second complication was that R&D did not generate a smooth process of continual improvements in production from batch to batch. Changes suggested by R&D were often disruptive because their implementation commonly created other problems that had to be solved before costs were reduced.[10] For example, in one instance R&D changes caused unexpected and undesirable changes in the

[7] The major objective of the other three projects was not cost reduction; one project was devoted to achieving a particular color and two projects developed a new process to circumvent a discontinued raw material.

[8] Indeed, for some of the products the R&D projects were actually initiated prior to the sample period but costs were not reduced until sometime into the sample period.

[9] The R&D personnel were also unable to reconstruct the timing of the R&D projects.

[10] Problems such as these can sometimes be avoided by first experimenting with pilot operations before transferring an R&D suggestion into the plant. For this business unit, however, it was not generally economical to experiment with pilot plant operations because of the small number of batches produced of its specialty

color of chemicals used in cosmetics. Further R&D then had to be done to restore the original color of the chemicals. Thus, during the adjustment process, unit costs of batches might actually rise before they would ultimately fall. The amount of time before such problems were resolved varied considerably across R&D projects depending on factors such as how frequently the batches of a product were produced and whether they were produced with dedicated equipment. Furthermore, during the adjustment process, R&D personnel often monitored production carefully. Consequently, production workers were particularly attentive to the batches, emptying production vessels immediately and taking other steps that kept down costs. When the R&D was complete, some of this attentiveness waned, causing costs to drift upward modestly. As the analysis indicates below, however, the permanent cost savings realized by R&D were still substantial.

To deal with the complications involving the effect of R&D on costs, we looked at the first and last batches in the sample of the products subjected to R&D to assess the impact of the R&D. In all but two of the 19 products the cost changes between the first and last batch were so sharp that the impact of the R&D was easily determined. We found that R&D led to three types of process changes.

The first type of process change involved eliminating the filtering step that normally followed the reaction step in the manufacturing process. Table 2 summarizes the hours spent in reaction and filtering for the first and last batch produced of the six products for which the filtering step was eliminated. Although the time spent in reaction actually increased for four of the six products, the hourly cost of reaction was only one-fourth that of filtration, so the elimination of the filtering step more than compensated for the increased reaction time. The filtering step was eliminated by making the product with sufficient purity in the reactor. This was achieved by starting with a higher purity raw material, changing the catalyst, modifying the rate of reactor heat-up, changing the operating pressure, and reducing the time needed to analyze in-process samples.[11] All but the last of these were the result of laboratory work by chemists running experiments under different conditions to achieve higher purities. The work principally tapped the general knowledge of the lab chemists. The batch records provided a benchmark for current practice, but otherwise neither the production records nor production workers were consulted.

Table 2 Hours Spent in Process 1 (Reaction) and Process 2 (Filtering) for the Six Products Where R&D Eliminated a Process Step

	First Batch		Last Batch	
Product	Process 1	Process 2	Process 1	Process 2
1	23.7	10.2	28.5	0
2	18.0	3.0	26.0	0
3	61.0	17.0	67.5	0
4	14.3	10.5	41.5	0
5	24.4	47.4	14.8	0
6	17.0	29.5	17.0	0

The second way R&D reduced cost was through the development of auxiliary equipment to increase the throughput of primary reactors that were bottlenecks in production. For reactors used to capacity, the business unit estimated that throughput could be doubled with the addition of auxiliary equipment to preheat raw materials, premix the raw materials, and unload the reactor faster. Table 3 summarizes the hours spent in the primary reactor for charging raw materials, reacting raw materials, filtering, and cooling down or packaging for the first and last batch produced of the seven products for which auxiliary equipment was developed. It indicates that after the R&D, all the nonreaction steps that had been performed in the primary reactor were performed in auxiliary equipment at minimal cost. Much of the development work on the auxiliary equipment was subcontracted to outside vendors. They designed and installed the auxiliary equipment based on specifications provided

chemicals. The pilot plant was used primarily to insure that production processes were safe before they were scaled up.

[11] Reducing sampling time can be important because all chemical reactions are reversible and the product can become less pure if left in the reactor beyond the optimal time.

Table 3 Reduction in Reactor Cycle Time From Auxiliary Equipment

	Reactor Hours For First Batch				Reactor Hours For Last Batch
Product	Charge	React	Cool	Total	Total (React)
1	6.3	20.0	5.5	31.8	19.8
2	18.7	21.0	3.4	43.1	24.8
3 (Process 1)	6.4	40.4	0.4	46.8	36.2
3 (Process 2)*	2.4	8.4	5.8	16.6	5.0
4	4.0	18.5	0.0	22.5	15.5
5	6.0	11.5	3.0	20.5	13.5
6	7.0	24.5	0.0	31.5	24.0
7	1.0	12.0	1.7	14.7	13.0

Note. *The time spent under "React" was for filtering.

Table 4 Cycle Time in Hours Before and After Process-Related R&D

Product	First Batch	Last Batch
1	19.0	9.5
2	39.5	24.0
3	26.1	8.0
4	35.5	7.0

by engineers of the business unit. In developing these specifications, the engineers exploited their general expertise but not any particular experience in the production of the affected products.

The last type of R&D lowered cost by reducing the cycle time of the manufacturing process. This was achieved by adding computerized control systems, changing the operating conditions of the reactor, such as increasing the temperature or pressure, improving the design of the agitators, and implementing statistical process control. The effects of these changes on cycle time for the four affected products is summarized in Table 4. The changes resulted from laboratory experiments, such as studying the effects of changing temperature or pressure in different circumstances and experiments with the manufacturing process itself. The batch records again provided a benchmark for current practice, but otherwise the projects primarily relied upon knowledge of fundamental chemical relationships that were explored in the lab.

5.2. Selection of R&D Projects

Another avenue by which production experience might have affected cost is through the choice of R&D projects. Projects were first proposed and then a subset of the proposed projects were undertaken. We investigated the role played by production experience in both the proposal and selection steps.

Proposals for R&D projects came principally from three sources. In 10 of the 22 products subject to process-related R&D, the R&D was in response to the plant no longer being able to produce the product to specifications. In two instances, this was due to a raw material no longer being available. In the other eight cases, the source of the problem was never discovered.[12] When it was no longer possible to produce a product to specification, R&D was requested to develop a new production process. The next most common request for R&D, which pertained to eight of the 22 products, came from marketing. A large and growing demand for these products was forecast, and process-related R&D was needed to produce the projected quantities and lower unit costs. The requests for the R&D on the remaining four products came from sales. The sales group would sometimes find they could not be competitive on a certain product, perhaps due to higher material costs or a less efficient production process than competitors, and would challenge the business unit to reduce its price to meet or better the competition. Clearly, production experience played no role in the initiation of projects requested by marketing and sales, and it was only when production experience was not useful that the first type of R&D project, designed to bring the product back up to specification, was initiated.

Production experience also did not play a role in which of the proposed R&D projects were undertaken. The proposed projects were formally evaluated according to their technical goals, the expected sales of the affected product in the next year, and the expected increase in the per unit profit of the affected product. Because the product of the latter two factors is the

[12] Difficulties can arise from subtle changes, such as a small variation in a natural raw material or an inadvertent change in manufacturing conditions, that are hard to detect even by experienced personnel.

expected profits from the project in the next year, the projects were essentially chosen according to their technical feasibility and expected annual profits. Neither criterion involved production experience. Thus, production experience per se did not affect which of the proposed projects were undertaken.

A closer look at the selected projects, though, suggests how it might appear that production experience played a role in the choice of R&D projects and ultimately the reduction in costs. We ranked the 221 products in terms of their cumulative output before and during the sample period. We found that 19 of the 22 products (86%) subjected to R&D ranked in the top 125 (57%) in terms of cumulative output during the sample period. This is consistent with expected future sales being one of the criteria on which the products were evaluated for R&D. Perhaps more surprising is that the products subjected to R&D were also above average in their (estimated) presample cumulative output. Six of the 22 (27%) were in the top 25 (11%) and 15 of the 22 (68%) were in the top 125 (57%) in terms of presample output.[13] The primary reason for the link between presample cumulative output and the choice of R&D projects is that presample cumulative output was a good predictor of subsequent output. With R&D being the primary force behind cost reduction, the result is that products with greater cumulative output were more likely to experience cost reduction. The connection between cumulative output and cost reduction was not, however, due to production experience but to the link between cumulative output and the expected returns from R&D. Thus, it was by conditioning the *incentives* to conduct R&D that cumulative output was related to cost.

5.3. R&D, Production Experience, and the Learning Curve

Our findings suggest that for the R&D products, cost did not progressively decline with production experience, as is assumed in the learning curve model. Nonetheless, the estimates of β for the R&D products were generally negative, substantial in absolute value, and often statistically significant at conventional levels. To consider how this could occur, we examined the fit of the model for the R&D products. To minimize the noise from sampling error, we concentrated on the ten R&D products with at least 30 observations.

Figure 1 plots the batch observations and the fitted regression line for each of the ten products, numbered 1 to 10. The horizontal axis measures the log of the time since production began, which is our measure of production experience. The vertical axis measures the log of the cost of the batch plus 0.4 times the log of the size of the batch. This takes account of differences in batch costs due to differences in the size of the batch assuming the coefficient γ in Equation (1) equals -0.4, the value it was constrained to equal in the estimates of the learning curve model. The regression line is simply $a + bX$, where a and b are the estimates of α and β in Equation (1) respectively and X is the log of the time since production began. For comparison, Figure 1 presents the same information for the four sampling products with at least 30 observations, numbered 11 to 14.

In all 10 of the R&D products in Figure 1, the regression line is negatively sloped, indicating cost reduction over time. This is certainly to be expected if cost reduction is primarily due to R&D initiatives leading eventually to permanent reductions in cost. With cost eventually declining permanently and production experience rising monotonically, cost and experience must be negatively related. If R&D is the primary driver of cost reduction, though, the learning curve model is misspecified. This is supported by the scatter diagrams in Figure 1. The scatterplots reflect considerable variation in cost across batches that is not captured by the learning curve model. If the learning model generated the data, the batch observations should be randomly distributed about the regression lines. This is not the case, however. For some of the products, the initial batches lie above the regression line and the intermediate batches of the products lie below the regression line, which is consistent with R&D reducing cost as of the intermediate batches. This is best illustrated by Products 7, 8, and 9. The first five batches of Product 7, the first seven for Product 8, and 13 of the first 14 for Product 9 lie above their respective

[13] If the three products in the bottom 50 that replaced a product that ranked high in presample output were classified in the top 125, then 82% of the products would be in the top 125 in presample output.

Figure 1 **Estimated Model for 10 R&D Products (#1–10) and 4 Sampling Products (#11–14) with 30 or More Observations**

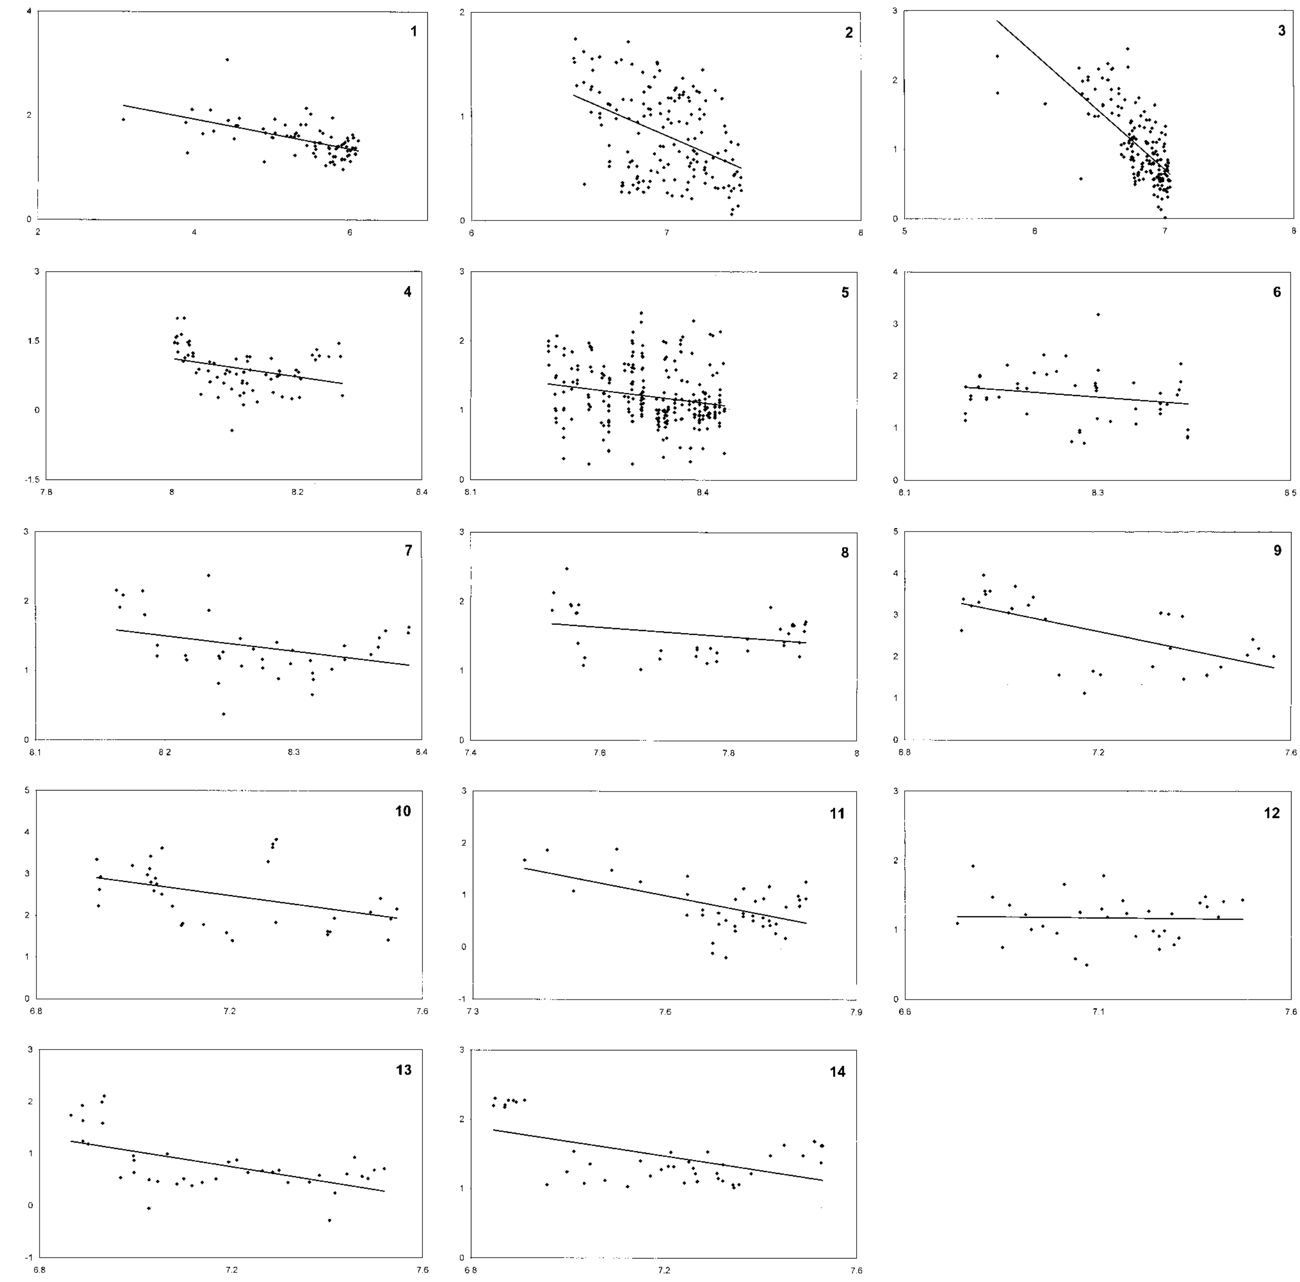

Note. Ln (time) on the Horizontal Axis, Ln (unit cost) + 0.4 Ln (batch size) on the Vertical Axis.

regression lines. In contrast, 21 of the next 25 for Product 7, the next 15 for Product 8, and the next five for Product 9 lie below their respective regression lines. A similar but less pronounced pattern holds for the other seven R&D products except 1 and 5, and it also holds for the sampling products as well. In addition, for some of the products, particularly 4, 7, 8, 9, and 14, the final batches lie above the regression line. This clustering of consecutive batches about the regression line implies that the residuals of the learning curve model will be positively serially correlated, which in turn implies a Durbin-Watson statistic below two. As expected, the Durbin-Watson statistic is below two for nine of the 10 R&D products, with seven of the nine significantly different from two at the 0.05 level.[14]

These departures from the learning curve model are consistent with R&D contributing to a permanent reduction in cost and production experience having no direct effect on cost. An extreme version of such a process would imply that the cost of batches is constant initially, then declines to a lower level once the R&D is implemented in production, and then remains constant for all subsequent batches. This is consistent with the patterns noted above in which costs initially are higher than predicted by the learning curve model, less than predicted for intermediate batches, and then greater than predicted for the final batches. Our investigation of how R&D operated suggested, though, that R&D did not generally lead to a precipitous drop in costs after a specific batch. Rather, R&D often required some experimentation across various batches before cost was consistently lowered. Furthermore, after the experimentation was completed and R&D personnel ceased monitoring production, costs could drift up as batches were left in equipment longer than necessary and production workers generally became less attentive to the products. A closer look at the patterns for the 10 R&D products supports this more nuanced view of how R&D operated. In a number of the products, after the initial high cost batches the variation in costs across batches increased, with some low cost batches mixed with some high cost ones. A good example is Product 7. The first five batches are clustered well above the regression line and then the next 11 batches span a wide range of costs, with one batch even having a higher cost than all the previous batches. After these 11 intermediate batches, costs settle down around the level attained by the lower-cost intermediate batches. This is followed by a modest rise in cost over the final six batches, but to a level still well below the initial batches, consistent with the sharp drops in cost reflected in Tables 2–4. A number of the other R&D products display similar patterns.[15] Thus, even if cost reduction is largely driven by R&D, our estimates of the learning curve indicate that it may still appear as if production experience contributes to cost reduction.

6. Discussion

In §2, we suggested that the learning curve literature leaves open a number of questions related to the role of production experience in cost reduction. One concerns the extent to which cost reduction is the result of process innovation and the role of production experience in such efforts. Relatedly, if purposeful efforts to improve the production process plays an important role in cost reduction, why do variables that measure these efforts not perform well in statistical studies of cost reduction (e.g., Adler and Clark 1991 and Mishina 1992)? Moreover, even if cost reduction is related to process innovation, why is there such a regular relationship between costs and cumulative output? In this section, we use our analyses to reflect on these and related questions.

Unlike many learning curve studies, we examined products of different vintages and produced under different conditions. A majority were produced using equipment with excess capacity. Judging from the estimates of the learning curve model for these products, they were subject to little cost reduction.[16] This

[14] The serial correlation in the 10 products was not only substantial but also considerably greater than the average product, with the mean Durbin-Watson equal to 1.4 for the 10 R&D products versus 1.8 for the entire sample of 221 products.

[15] The products that conform most closely to this pattern are Products 4, 7, and 8. These were products for which cycle time was reduced through specific R&D projects designed to modify the temperature in the reactors. While plant personnel could not recall the timing of the projects, they were confident that the sharp drop in costs reflected in the batch records was attributable to the R&D projects. In the other seven R&D products, R&D involved continuing efforts over the sample period rather than a single project, which may account for the more gradual decline in costs in the scatterplots for these products. The changes associated with the sampling initiative generally involved less uncertainty than the changes motivated by R&D. Consistently, the scatter diagrams for Products 11 to 14 show a more uniform decline in costs after the initial high cost batches than for the R&D products.

[16] Among these 122 products, Table 1 indicates that a nontrivial number, 23, had (constrained) estimates of β less than -0.6. Although we attributed this to sampling error, we cannot rule out

reflected the absence of any incentive to economize on the main element of cost, the use of equipment. Thus, among our products there was nothing automatic about cost reduction. Among the other products, we found that the conventional learning curve model fit best for the ones subject to specific initiatives such as R&D, campaigning, and the sampling program. We concentrated our analysis on the role of production experience in shaping these initiatives and in the selection of the products for the initiatives. Focusing particularly on the R&D products, we found that experience gained from production did not influence the products selected for R&D nor did it generally inform the R&D itself. R&D operated by overcoming production bottlenecks and reducing the time for production, both of which contributed to greater manufacturing throughput.[17] While the key source of knowledge underpinning the cost reductions was the business unit's R&D personnel, the initiatives behind the projects typically came from elsewhere. Management played a key role, as in the case of the sampling initiative. The quality control staff in combination with management initiated numerous projects when it was learned that products were not meeting specifications. In other instances, sales and marketing personnel suggested a strong need for cost reduction on specific products. Thus, in response to initiatives coming from a number of sources, costs were reduced through purposeful investment in process innovation, involving principally the deployment of specialized personnel and related resources.

If it was not production experience, what kind of knowledge contributed to cost reduction in the R&D products? Building on their training and broad experience with chemical manufacturing processes generally, the R&D chemists identified cost reducing strategies largely by performing experiments in the R&D lab, and the chemical engineers subsequently figured out how to implement the changes suggested by R&D. Each exploited a repertoire of solutions to past problems in performing its task. For example, engineers exploited their past experience in developing specifications for the specialized kinds of auxiliary equipment used to overcome bottlenecks. As problems were solved, repertoires of solutions expanded. Thus, to the degree that cost reduction benefited from learning from experience, the learning originated from confronting similar technical problems, which yielded know-how of a rather generic variety.[18] The know-how acquired in this learning process is not principally related to the production of any one product but to the design and modification of production processes for a broad class of products.

Our finding that considerable cost reduction resulted from R&D and related efforts contrasts sharply with the statistical learning curve studies that control for R&D and engineering efforts (e.g., Adler and Clark 1991, Mishina 1992). Our examination of the links between R&D and cost reduction may help explain the disparity in findings. For specialty chemicals, the lags between the initiation of R&D, the implementation of changes based on the R&D, and the correction of problems resulting from the changes were quite variable. These variable lags obscured the relationship between unit costs and R&D sufficiently that it was necessary to compare batches well before and after the R&D to see the effects of the R&D on cost. If R&D generally operates with a variable lag, it would be difficult for studies like Adler and Clark (1991) that have quantitative measures of R&D to find a clear statistical relationship between the timing of cost reductions and R&D. Moreover, if R&D is sufficiently disruptive when first implemented, it may be positively correlated with cost. Under such circumstances, if R&D and cumulative output are positively correlated, as we found, then controlling for R&D will actually increase the estimated effect of cumulative output on unit costs in conventional learning curve

that some of these products may have experienced conventional learning by doing. This is also true of the 19 estimates in the last column of Table 1 that are less than -0.2

[17] This accords with the case studies, which emphasize the importance of incremental technical changes to eliminate bottlenecks and speed up production, and with the findings of Bell and Scott-Kemis (1990) and Mishina (1992) concerning World War II Liberty ships and airframes.

[18] Bell and Scott-Kemis (1990) note the importance of this type of learning for World War II products. Findings by Argote et al. (1990) on the transferrence of learning across different shipyards producing Liberty ships are also consistent with this type of learning.

analyses. This would help explain the findings of Adler and Clark (1991) and Mishina (1992) that inclusion of R&D and engineering efforts in their regressions actually strengthened the relationship between unit costs and cumulative output.

If the critical knowledge underpining cost reduction comes from sources outside manufacturing operations such as R&D, why is there generally such a close relationship between cumulative output and costs? Our analysis of the R&D products demonstrated that it could arise from R&D alone. With cumulative output rising monotonically and R&D eventually leading to a permanent reduction in costs, cost and cumulative output will be negatively related. For our products, R&D involved a single initiative that generally extended over multiple batches. For products like commodity chemicals (cf. Lieberman 1984) or even the war-time products (cf. Bell and Scott-Kemis 1990, Mishina 1992) that are subject to repeated process R&D and engineering initiatives, it seems likely that the relationship between unit costs and cumulative output induced by R&D would be even stronger. Our findings suggest that if cost and cumulative output are related due to periodic decreases in cost resulting from R&D and related initiatives, there will be positive serial correlation in the residuals, a common finding in learning curve studies.[19] We also found that the incentive to perform R&D depended upon expected future output, which was in turn related to cumulative output. Cohen and Klepper (1996a, 1996b) demonstrate that many empirical regularities linking firm output and R&D can be explained by a simple theory in which past output conditions expectations of future output and in turn the incentives to perform R&D, suggesting that the relationship we found between the propensity to perform process R&D and cumulative output might hold widely. For products subject to repeated R&D, we suspect the connection between past output and R&D incentives will further strengthen the relationship between cost and cumulative output.

[19] Indeed, in his study of cost reduction during World War II in the B-17 bomber, Mishina (1992) uses serial correlation or the absence thereof as the basis for evaluating alternative models of progressive cost reduction.

Although we found that production experience did not generally inform R&D, we suspect that another reason cost and cumulative output are closely related in many studies is because in many other settings production experience does inform initiatives to lower cost. This was certainly the case in Hollander's (1965) study of DuPont's rayon plants. More recently, von Hippel and Tyre (1995) have argued that it may be a sensible strategy to *plan* on using production experience to reveal problems that need to be solved to lower cost, and Pisano (1994) finds that relying on production experience to design a new production process is more valuable when there is less a priori knowledge to draw upon to design the production process. Whatever the general role of production experience in guiding initiatives to lower cost, though, our findings demonstrate that R&D can induce a conventional relationship between cost and cumulative output even when learning from production plays little role in the process. Moreover, our findings suggest a novel way production experience may be related to cost, namely through its connection with incentives to undertake initiatives to lower cost.

7. Conclusion

For one major specialty chemical manufacturer, we developed an understanding of what stands behind the link between cumulative output and cost reduction, otherwise known as the learning curve. Using detailed cost data, information about R&D, and access to company personnel, we found that for the bulk of the products that experienced progressive cost reduction, cost reductions were largely the result of small technical changes in production that were based on R&D and related activities. While the initiative behind these changes came from a number of sources, the changes were achieved through process innovations which required the application of specialized personnel and other resources. The process changes were not generally motivated or informed by learning from experience in production. Nevertheless, we found that costs were inversely related to various measures of experience, including cumulative output and the total time in production. We also found that cumulative output was related to cost through its connection with expected future output, which conditioned the incentives to undertake process innovation. Thus, to the extent that cumulative output was related to cost at all, it was as a proxy for incentives to reduce cost through R&D and related initiatives and not as a measure of production experience.

Focusing on one type of production process, that associated with specialty chemicals, for one company offers both advantages and disadvantages. While surely unique, the batch production process we examined offered a number of distinctive advantages. First, very little changed over the production period that might have complicated the analysis. With the exception of investments in the auxiliary equipment, there were no significant changes in the capital stock, and we had detailed cost records that enabled us to quantify the effects of the auxiliary equipment. We were able to study cost reduction without having to be concerned about changes in the products themselves, as any product change resulted in the designation of a new product. Finally, unlike most products we had a natural measure of the scale of production, the size of the batch, which enabled us to distinguish scale from "learning" effects. As Lucas (1993) noted about the various studies of Liberty ships, this made for a nice clean experiment.

Surely our focus on one firm limits the generality of any claims that can be made for our study. Yet, partly because of that focus, and because we were provided detailed data and broad access to key personnel, we were able to probe more deeply than most studies of progressive cost reduction. We found a pattern that looked like learning from production experience but was really shaped by process innovation and the incentives underlying it. It would be valuable to analyze whether our findings generalize to other settings, particularly products that are produced with more labor and dedicated machinery, have a less developed knowledge base, and allow production workers more discretion to adjust production than specialty chemicals. While information from production may well inform initiatives to reduce cost for such products, we suspect there is still typically much that needs to be done with such information and a lot of other expertise that needs to be combined with it to yield process savings. If so, this would leave an important role for incentives in conditioning cost reduction, suggesting that our findings might help explain the close empirical relationship commonly found between cost and measures of production experience.[20]

[20] The authors thank Linda Argote, Kathleen Carley, Marvin Lieberman, three anonymous referees, and an associate editor for helpful comments. Klepper gratefully acknowledges support from the Economics Program of the National Science Foundation, Grant No. SBR-9600041.

References

Adler, P. S., K. B. Clark. 1991. Behind the learning cuve: A sketch of the learning process. *Management Sci.* **37** 267–281.

Alchian, A. 1963. Reliability of progress curves in air-frame production. *Econometrica* **31** 679–693.

Argote, L., S. L. Beckman, D. Epple. 1990. The persistence and transfer of learning in industrial settings. *Management Sci.* **36** 140–154.

Arrow, K. J. 1962. The economic implications of learning by doing. *Rev. Econom. Stud.* **29** 155–173.

Bell, R. M., D. Scott-Kemis. 1990. The mythology of learning-by-doing in World War II airframe and ship production. Mimeo, Science Policy Research Unit, University of Sussex.

Cohen, W. M., S. Klepper. 1996a. Firm size and the nature of innovation within industries: The case of process and product R&D. *Rev. Econom. Statist.* **78** 232–243.

——, ——. 1996b. A reprise of size and R&D. *Econom. J.* **106** 925–951.

Darr, E. D., L. Argote, D. Epple. 1995. The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises. *Management Sci.* **41** 1750–1762.

Dasgupta, P., J. Stiglitz. 1988. Learning-by-doing, market structure, and industrial and trade policies. *Oxford Econom. Papers* **40** 246–268.

Dutton, J., A. Thomas. 1984. Treating progress functions as a managerial opportunity. *Acad. Management Rev.* **9** 235–247.

Enos, J. L. 1967. *Petroleum Progress and Profits.* M.I.T. Press, Cambridge, MA.

Fudenberg, D., J. Tirole. 1983. Learning by doing and market performance. *Bell J. Econom.* **14** 522–530.

von Hippel, E., M. Tyre. 1995. How learning by doing is done: Problem identification in novel process equipment. *Res. Policy* **24** 1–12.

Hollander, S. 1965. *The Sources of Increased Efficiency: A Study of DuPont Rayon Plants.* M.I.T. Press, Cambridge, MA.

Jarmin, R. S. 1994a Learning by doing and competition in the early rayon industry. *Rand J. Econom.* **25** 441–454.

Jarmin, R. S. 1994b. Asymmetric learning in U.S. manufacturing industries. Mimeo, Center for Economic Studies, Bureau of the Census.

Krugman, P. 1987. Increasing returns and the theory of international trade. Truman F. Bewley, ed. *Advances in Economic Theory*: *Fifth World Congress*, Cambridge University Press, Cambridge, England.

Lieberman, M. B. 1984. The learning curve and pricing in the chemical processing industries. *Rand J. Econom*. **15** 213–228.

Lucas, R. E. 1993. Making a miracle. *Econometrica* **61** 251–272.

Mishina, K. 1992. Learning by new experiences. Mimeo, Harvard Business School.

Pisano, G. P. 1994. Knowledge, integration, and the locus of learning: An empirical analysis of process development. *Strategic Management J.* **15** 85–100.

Sinclair, G. 1994. Learning and progressive cost reduction in a specialty chemicals business unit. Unpublished doctoral dissertation, Carnegie Mellon University, Pittsburgh, PA.

Spence, A. M., 1981. The learning curve and competition. *Bell J. Econom*. **12** 49–70.

Stern, B. 1933. Labor productivity in the automobile tire industry. Bureau of Labor Statistics Bulletin 585, Government Printing Office, Washington, D.C.

Accepted by Hau Lee; received July 1997. *This paper has been with the authors* 10 *months for* 2 *revisions.*