Learning Curves in Manufacturing

LINDA ARGOTE AND DENNIS EPPLE

Large increases in productivity are typically realized as organizations gain experience in production. These "learning curves" have been found in many organizations. Organizations vary considerably in the rates at which they learn. Some organizations show remarkable productivity gains, whereas others show little or no learning. Reasons for the variation observed in organizational learning curves include organizational "forgetting," employee turnover, transfer of knowledge from other products and other organizations, and economies of scale.

As ORGANIZATIONS PRODUCE MORE OF A PRODUCT, THE unit cost of production typically decreases at a decreasing rate. This phenomenon is referred to as a learning curve, a progress curve, an experience curve, or learning by doing. A learning curve for the production of an advanced military jet built in the 1970s and 1980s (Fig. 1) illustrates the two salient properties of learning. The number of direct labor hours required to assemble an aircraft decreased significantly as experience was gained in production, and the rate of reduction of assembly hours declined with rising cumulative output.

Learning curves have been documented in many organizations, in both the manufacturing and service sectors. The unit costs of producing aircraft (*1, 2*), ships (*3*), refined petroleum products (*4*), and power plants (*5, 6*) have been shown to follow the characteristic learning-curve pattern. Learning curves have also been found to characterize outcomes as diverse as success rates of new surgical procedures (*7*), productivity in kibbutz farming (*8*), and nuclear plant operating reliability (*9*).

The productivity gains associated with organizational learning curves are often quite large. For example, during the first year of production of Liberty ships during World War II, the average number of hours of labor required to produce a ship decreased by 45%, and the average time it took to build a ship decreased by 75% (*10*). A recent study of a truck plant reported a remarkable growth in productivity of approximately 190% over the first year of the plant's operation (*11*).

Organizations vary considerably in the rates at which they learn (*12–14*). Whereas some organizations show extraordinary rates of productivity growth as cumulative output increases, others fail to show expected productivity gains from learning. Lockheed's production of the L-1011 Tri-Star in the 1970s is an example of a program with little evidence of learning (*15, 16*). Lockheed lost over $1 billion on the Tri-Star program in the 1970s (*16*).

Why did little or no productivity growth occur in production of the Lockheed Tri-Star while the truck plant mentioned earlier showed impressive growth in productivity? For U.S. manufacturing and other organizations to compete effectively, we need to understand why some organizations show rapid rates of learning and others fail to learn. Thus, we need to identify factors affecting organizational learning curves and use this knowledge to improve manufacturing performance.

Understanding factors affecting learning can enable managers to improve the performance of a firm in many areas. Applications include formulating manufacturing strategy (*17*), production scheduling (*12*), pricing and marketing (*18*), training (*19*), subcontracting production (*20*), and predicting competitors' costs (*21*). The rate and transfer of learning are also important issues for antitrust policy (*22*) and trade policy (*23*).

We examine evidence from several disciplines on organizational learning curves, particularly in manufacturing. Our focus is primarily on empirical studies that analyzed organizations or work groups. We show that organizations vary considerably in the rate at which they learn and identify factors responsible for the variation.

Research on Organizational Learning Curves

The first documentation of an organizational learning curve was published in 1936 by Wright (*1*), who reported that unit labor costs in air-frame production declined with cumulative output (*24*). Further interest in learning was stimulated by Alchian's 1948 study of learning in 22 aircraft production programs (*2*).

The conventional form of the learning curve is a power function:

$$y = ax^{-b} \tag{1}$$

where y is the number of direct labor hours required to produce the xth unit; a is the number of direct labor hours required to produce the first unit; x is the cumulative number of units produced; and b is a parameter measuring the rate labor hours are reduced as cumulative output increases.

As this expression shows, the standard measure of organizational experience in the learning-curve formulation is the cumulative number of units produced, a proxy variable for knowledge acquired through production. If unit costs decrease as a function of this knowledge, other variables being equal, organizational learning is said to occur.

Learning curves are often characterized in terms of a progress ratio, p. With the learning curve in Eq. 1, each doubling of cumulative output leads to a reduction in unit cost to a percentage, p, of its former value (*25*). Thus, an 80% progress ratio means that each doubling of cumulative output leads to a 20% reduction in unit cost.

Before the discovery of learning curves in organizations, the learning-curve pattern had been found to characterize the perform-

L. Argote is associate professor of industrial administration and D. Epple is professor of economics at the Graduate School of Industrial Administration, Carnegie Mellon University, Pittsburgh, PA 15213.

ance of individual subjects as they gained experience with a task (*26*). For example, an early study of individual learning curves focused on the number of errors individual students made as they progressed through a typing course (*27*). Organizational learning curves, by contrast, focus on the performance of entire organizations or organizational subunits (for example, manufacturing plants). Although the productivity of an organization may be affected by individuals learning how to perform their jobs better, it is also affected by many additional factors such as technological developments and improved coordination of the production process. Thus, organizational learning involves more than individuals becoming better at their particular jobs.

Much of the work on organizational learning curves has focused on specifying the functional form of the relation between unit costs and cumulative output and on studying the phenomenon in different industries (*12, 13*). Several new trends in research on organizational learning curves are apparent. The set of outcome measures has been broadened to include, for example, industrial accidents per unit of output (*28*) and defects or complaints to quality control per unit of output (*29*). The transfer of productivity gains acquired through learning by doing across organizations is also being studied (*5, 6, 30*). Increasing attention is being given to disentangling the various factors that contribute to organizational learning (*30, 31*).

Variation in Organizational Learning Rates

The frequency distribution of progress ratios found in more than 100 studies in different industries is presented in Fig. 2 (*13*). Note that the progress ratios vary a great deal, reflecting the variation in the rate productivity grows with increasing cumulative output. Also, the modal progress ratio falls at 81 to 82%—giving rise to the general assumption of an "80% learning curve" (*32*).

Understanding the reasons why learning rates vary is a major challenge for research. The different rates of learning (Fig. 2) are not simply a function of the different products studied, although differences in products are, of course, a source of variation. There is often more variation across organizations or organizational units producing the same product than within organizations producing different products. For example, productivity gains varied more within shipbuilding production programs than between production programs during World War II (*10*). Similarly, Hayes and Clark (*14*) found considerable variation in the rate of learning across plants in the same firm producing the same product with similar equipment and materials.

Different plants producing the same product that have different rates of learning are shown in Fig. 3. The data are from three truck plants producing the same product within the same company. The cumulative number of trucks produced is plotted against the number of direct labor hours required to assemble each truck. Although each plant shows the characteristic learning-curve pattern, the pattern is different for each plant. Thus, there is considerable variation in productivity among these plants that is not explained by the conventional learning-curve model (*33*).

This variation in the rate that organizations learn may be due to organizational "forgetting," employee turnover, transfer of knowledge, and the failure to control for other factors, such as economies of scale, when estimating learning curves.

Organizational forgetting. When production is resumed after an interruption such as a strike, unit cost is often higher than the level achieved before the interruption (*34*). Similarly, there is evidence that knowledge acquired through learning by doing depreciates: recent output rates may be a more important predictor of current production than cumulative output (*30*). Theoretical research and simulation results have also indicated that forgetting has implications for planning and scheduling (*35*).

Organizational forgetting may explain why Lockheed's costs for the L-1011 Tri-Star did not follow the learning-curve pattern. The production of the L-1011 Tri-Star was characterized by wide variations in the rate of output (Table 1). Lockheed estimated that its production costs would fall below price in mid-1974 (*36*). The conventional learning-curve formulation applied to the Tri-Star yielded a prediction that costs would fall below price about the time the 50th plane was built, sometime in 1973 (*37*). In November 1975, Lockheed reported that unit costs at that time were less than the price at which planes were being sold (*38*). Planes were sold for

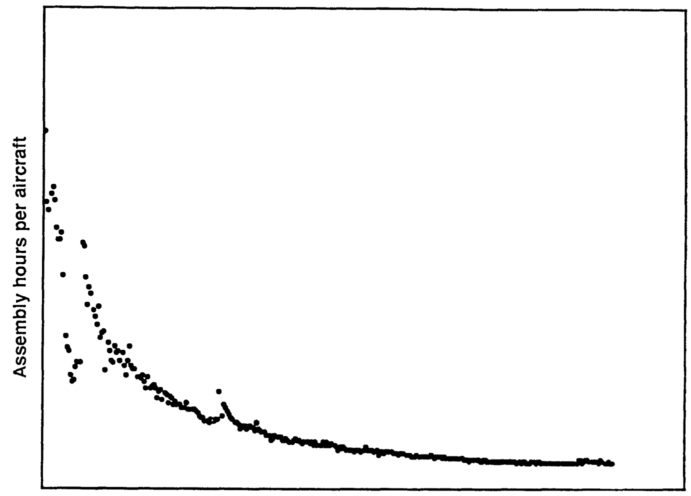

Fig. 1. Relation between assembly hours per aircraft and cumulative number produced. Units omitted.

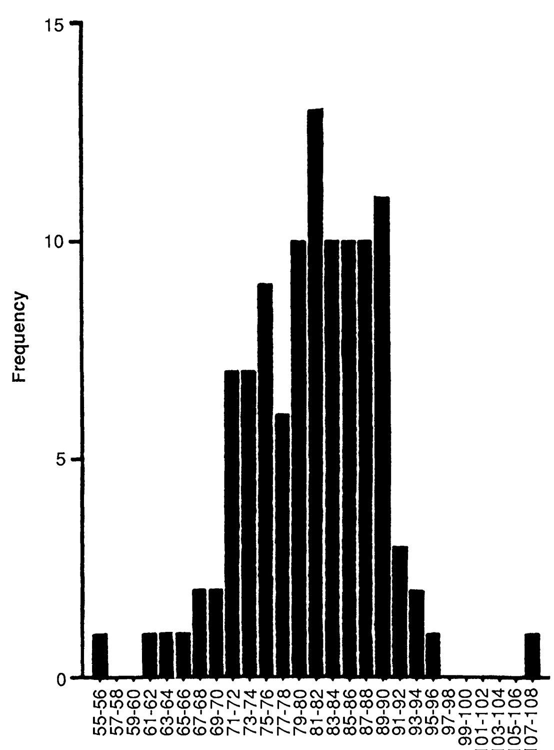

Fig. 2. Distribution of progress ratios observed in 22 field studies (n = 108) (*13*).

$20 million in 1975. Cuts in production occurred in late 1975. Costs rose to exceed price and, apparently, remained above price for the rest of the production program (*39*). In 1982, the L-1011 was sold for $50 to $60 million per plane. This corresponds to $29 to $35 million in 1975 dollars. Thus, production cost per plane was less than $20 million in real terms in 1975 but greater than $29 million in real terms in 1982.

In the conventional learning-curve model, unit costs decrease as a function of cumulative output. This model does not explain the Lockheed data, however, since costs rose as cumulative output continued to increase. A model in which knowledge depreciates and recent output is more important than cumulative output in predicting costs can explain the Lockheed data. Lockheed's costs rose when production was cut and recent output was relatively low. Even though a detailed analysis of the L-1011 data would be required to test the hypothesis that depreciation of knowledge occurred, the pattern of costs reported by Lockheed is consistent with the depreciation hypothesis.

Why might knowledge acquired through learning by doing depreciate? Knowledge could depreciate because individual employees forget how to perform their tasks or because individuals leave the organization and are replaced by others with less experience. For example, it would not be surprising if many managerial and line employees who worked on the L-1011 during its early period of high annual output (1973 to 1975) were no longer on the project when the company resumed comparatively high output levels later in the program (1979 to 1980). Depreciation could also be due to changes in products or processes that make previously acquired knowledge obsolete.

Depreciation can also result if organizational records or routines are lost or become difficult to access. An example, recently described in *Science*, is provided by the difficulty in accessing data collected by Landsat, an earth surveillance program. It is estimated that 90% of the data collected before 1979 is inaccessible because the data were recorded by equipment that no longer exists or cannot be operated and "bleeding" of magnetic images occurred over time (*40*).

Thus, forgetting or depreciation of organizational knowledge can cause organizational learning rates to vary. When depreciation occurs and the conventional learning curve is used, two organizations that have achieved the same level of cumulative output will be at different points on the learning curve if the recent output level of one is different from that of the other. Such differences in recent output levels may arise for a host of reasons including strikes, materials shortages, and fluctuations in product demand that lead to temporary shutdown of some plants but not others. A method for extending the analysis of learning to encompass depreciation is provided in Argote, Beckman, and Epple (*30*).

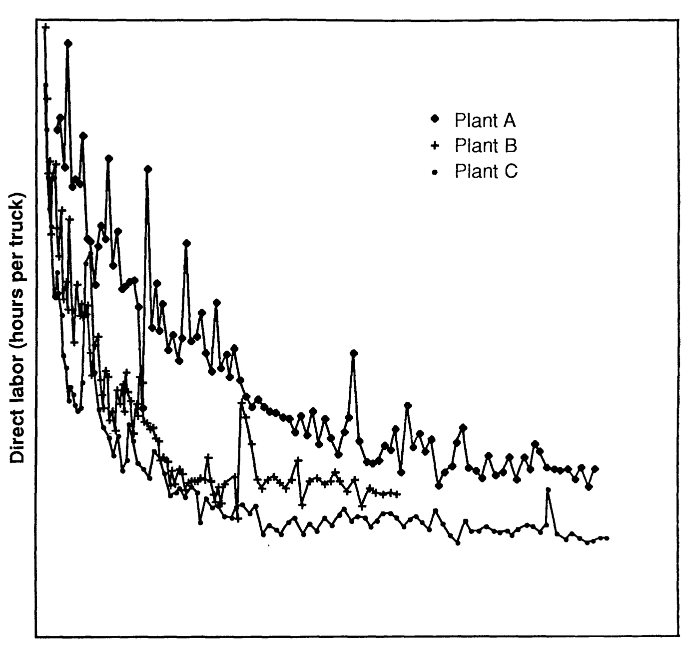

Fig. 3. Relation between direct labor hours per truck and cumulative number produced for three truck plants. Units omitted.

Table 1. Lockheed's production of the L-1011 Tri-Star (*15, 38, 53*).

Year	L-1011 production	
	Annual units	Cumulative units
1972	17	17
1973	39	56
1974	41	97
1975	25	122
1976	16	138
1977	6	144
1978	8	152
1979	24	176
1980	25	201
1981	18	219

Turnover. When organizational knowledge is possessed by individual employees, employee turnover can be expected to have an impact on learning and forgetting in organizations. Thus, differing rates of turnover across organizations could explain the differences observed in organizational learning curves.

Does turnover affect the rate of learning and forgetting in firms? Research indicates that turnover of direct production workers did not have a significant effect on the rate of learning or forgetting in World War II shipyards (*30*). This result is striking, given that turnover in these organizations averaged more than 10% per month. The result is consistent, however, with results from several laboratory studies that found increases in the performance of groups over successive trials in the face of turnover (*41*).

Why did turnover not matter in these production environments? The jobs of production workers in the shipyards were standardized and designed so that a new employee could become proficient with minimal training (*42*). Procedures existed for training and transmitting knowledge to new members.

Many production environments today also experience considerable turnover. For example, the corporate office required one plant that we studied to accept, over a 2-month period, more than 300 employees from a neighboring plant that closed. When these new employees arrived, more than 300 employees left the plant and another 150 moved to different jobs within the plant. Thus, at the end of the second month, 15% of employees at the plant were either new to the plant or at different jobs within the plant than at the beginning of this 2-month period. Plant managers at this corporation recognize that high turnover may occur and attempt to design their operations to mitigate its effects.

Although results to date do not suggest that turnover affects the rate of organizational learning, in the limit, turnover would surely affect learning and forgetting in firms (*43*). Moreover, organizations confronted with high rates of turnover may insulate themselves from its effects by routinizing jobs and procedures. The consequence may be a lower rate of learning than is achieved by organizations not confronted with such turnover. Turnover may matter more in organizations where jobs are not standardized and procedures do not exist for transmitting knowledge to new members. Turnover of managers and technical support staff, such as engineers, may also matter more than turnover of direct production workers (*44*).

Transfer of productivity gains. Another possible reason for different rates of organizational learning is the transfer of knowledge across products or across organizations. Experience gained in the production of one product can be transferred to the production of related products (*45*). For example, suppose two organizations produce the same product. The first organization produces only one product while the second organization produces a related product using some of the same operations as were used to make the first product. The second organization should benefit from the additional cumulative output generated by the second product and have lower costs on the shared operations. Thus, if the two organizations have similar cumulative output levels for the product they both produce, the second organization should have lower costs and be farther down the learning curve than the first because of the transfer of knowledge from the related product.

Similarly, an organization that produced a related product in the past may be able to transfer knowledge to the manufacture of a product currently in production. Thus, an organization with previous experience producing a related product may appear to have a faster rate of learning than an organization without prior experience, even though their cumulative output levels for the product currently in production are the same.

Transfer of knowledge across organizations might also occur (*46*). Transfer might occur through personnel movement, communication, participation in meetings and conferences, training, improved supplies, modifications in technology, or "reverse-engineering" of products. Knowledge transferred from outside the firm is difficult to measure. One approach to measuring this knowledge involves aggregating cumulative output across all firms in the industry. This measure of industry experience has been found to have a significant effect on unit costs in some industries (*5*) but not in others (*6*). Organizations coming on line later have been found to begin with higher productivity levels than their counterparts with early start dates (*30*). Once organizations began production, however, they did not benefit from knowledge acquired from other organizations.

Another approach to measuring knowledge acquired outside the firm is to assume that calendar time is an adequate proxy variable for knowledge acquired in the general environment. Several studies in which this approach was used have found that calendar time is not as good a predictor of an organization's productivity growth as is its own cumulative output (*3, 11, 31*). Kelsey *et al.* (*7*) found that calendar time was a significant predictor of surgical success rates only for the first 20 operations performed. For later operations, calendar time was not significant but experience was. The researchers suggested that surgeons were more likely to learn from the experience of others when they first begin to perform the procedure but not later. Thus, there is evidence that transfer may occur across organizations, and it seems particularly likely to occur in the early phases of production.

If transfer occurs for one organization but not another, the organizations will appear to have different rates of learning, even if their "internal" rates of learning from their own past production experience are the same. For example, consider two plants operated by the same company. One plant leads by beginning production first. The corporation invests in transferring knowledge acquired by the lead plant to the second plant. If transfer occurs, the second plant will have higher productivity than the first plant for the same level of cumulative output; the learning curve of the second plant will lie below that of the first.

Differences in learning rates across plants can also arise from incomplete transfer across shifts within plants (*11*). Managers at one plant we studied were disappointed that incomplete transfer occurred from the first shift to the second when the second shift was introduced at the plant. They speculated that the incomplete transfer was due to inadequate documentation of lessons learned from the first shift.

As an example of how incomplete transfer can cause differences in learning rates, consider two plants producing the same product. One plant operates with one shift per day while the other operates with two shifts (not an unusual occurrence). If the rate of learning per shift is the same in both plants but incomplete transfer occurs across shifts, the learning curves at the two plants will be different. When unit cost is plotted versus cumulative output from plant data, the plant operating with one shift per day will exhibit greater learning than the plant operating with two shifts per day. For example, suppose that unit costs for the two plants are compared at the point where both have produced 10,000 units. For the plant operating with two shifts per day, the cumulative output per shift will be only 5,000, that is, half the cumulative output per shift of the plant operating with only one shift. Thus, if there is no transfer across shifts, the plant operating with two shifts per day will have the productivity at a cumulative output of 10,000 units that the plant operating with one shift per day had at a cumulative output of 5,000 units.

Other factors affecting learning rates. An investigator should control for other variables that affect production because exclusion of such variables may bias the estimated rate of learning. For example, suppose economies of scale are present, so that a given increase in inputs results in a more than proportionate increase in output. If the scale of operation is gradually increased over time, productivity will rise because of increasing exploitation of economies of scale. If one estimates the rate of learning without controlling for the changing scale of operation, this increasing exploitation of scale economies will result in an overestimate of the amount of learning.

Womer (*47*) has cogently argued for the importance of integrating estimation of learning with production function estimation as a vehicle for controlling for the effects of factors other than learning. A production function is a relation specifying output per period as a function of inputs that period, the state of technical knowledge, and other variables that may affect output. Symbolically, this may be written:

$$q = F(n, k, z) \tag{2}$$

where n denotes productive inputs, k denotes measures of the state of technical knowledge, and z denotes other variables that may affect production (*48*).

In general, issues that must be addressed in estimating a production function are selection of a functional form; choice of the variables n, k, and z; specification of the properties of random factors affecting the production process; and choice of an appropriate method of estimating the parameters of interest. There is some evidence that a plateau occurs, especially in machine-intensive industries (*49*). The choice of functional form should be flexible enough to accommodate this leveling out of the learning curve. It is also important to correct for problems that may arise if data are collected on a per period basis when several periods are required to produce each unit (*50*). Other issues in choice of functional form, specification of error structures, and estimation methods are addressed by others (*51*).

The choice of variables to be included in the model varies according to the production process being studied. For example, in a single plant with unchanging physical facilities, labor hours may be the only input that varies over time. In studying multiple plants, it may be appropriate to include measures of capital investment and other inputs that differ across plants, and such measures would also be needed if the facilities in a given plant change over time. An early example of empirical work on organizational learning that controlled for additional factors in the analysis was done by Rapping (3). He found both economies of scale and learning to be present in his study of productivity gains in shipbuilding.

Although cumulative output is typically used as the measure of knowledge acquired through learning by doing, measures that place relatively greater weight on recent output than on output in the distant past are appropriate if depreciation occurs (30). When production occurs at several plants, additional variables such as cumulative output aggregated across plants may be included in addition to a plant's own cumulative output as measures of the transfer of knowledge. If the plant has the potential to benefit from improvements in technical knowledge in the larger environment, proxies for the pace of such improvements are appropriate. One such proxy is calendar time (52). Finally, it may be necessary to control for factors such as labor turnover, product mix, and adjustment costs associated with changing inputs.

Conclusion

Although learning curves have been found in many organizations, there is great variation in the rate at which organizations learn, ranging from production programs with little or no learning to those with impressive productivity growth. We identified reasons why organizational learning rates vary. These include organizational forgetting, employee turnover, transfer of knowledge across products and across organizations, incomplete transfer within organizations, and economies of scale. Learning is a powerful source of productivity growth, and better understanding of learning can enhance manufacturing performance.

REFERENCES AND NOTES

1. T. P. Wright, *J. Aeronaut. Sci.* **3**, 122 (1936).
2. A. Alchian, *Econometrica* **31**, 679 (1963).
3. L. Rapping, *Rev. Econ. Stat.* **47**, 81 (1965).
4. W. B. Hirschmann, *Harvard Bus. Rev.* **42**, 125 (1964).
5. M. B. Zimmerman, *Bell J. Econ.* **13**, 297 (1982).
6. P. L. Joskow and N. L. Rose, *Rand J. Econ.* **16**, 1 (1985).
7. S. F. Kelsey *et al.*, *Am. J. Cardiol.* **53**, 56C (1984).
8. H. Barkai and D. Levhari, *Rev. Econ. Stat.* **55**, 56 (1973).
9. P. L. Joskow and G. A. Rozanski, *ibid.* **61**, 161 (1979).
10. A. D. Searle and C. S. Gody, *Mon. Labor Rev.* **61**, 1132 (1945).
11. D. Epple, L. Argote, R. Devadas, Graduate School of Industrial Administration (GSIA), Carnegie Mellon University, Working Paper, in preparation.
12. L. E. Yelle, *Decis. Sci.* **10**, 302 (1979).
13. J. M. Dutton and A. Thomas, *Acad. Manage. Rev.* **9**, 235 (1984).
14. R. H. Hayes and K. B. Clark, *Harvard Bus. Rev.* **64**, 66 (1986).
15. "Tri-Star production costs," *Aviat. Week Space Technol.*, 15 October 1979, p. 32.
16. "Tri-Star's trail of red ink," *Bus. Week*, 28 July 1980, pp. 88 and 90.
17. R. H. Hayes and S. C. Wheelwright, *Restoring Our Competitive Edge: Competing Through Manufacturing* (Wiley, New York, 1984).
18. R. J. Dolan and A. P. Jeuland, *J. Mark.* **45**, 52 (1981).
19. F. K. Levy, *Manage. Sci.* **11**, B136 (1965).
20. J. V. Jucker, *AIIE Trans.* **9**, 321 (1977).
21. B. D. Henderson, *J. Bus. Strategy* **4**, 3 (1984).
22. A. M. Spence, *Bell J. Econ.* **12**, 49 (1981).
23. H. K. Gruenspecht, *J. Int. Econ.* **12**, 225 (1988).
24. Empirical evidence of learning inspired theoretical work to reconcile the phenomenon with economic production theory [J. Hirshleifer, *J. Bus.* **35**, 235 (1962) and W. Y. Oi, *Econ. J.* **77**, 579 (1967)]. Implications of learning by doing for both competitive and noncompetitive markets have also been investigated [K. J. Arrow, *Rev. Econ. Stud.* **29**, 155 (1962); D. Fudenberg and J. Tirole, *Bell J. Econ.* **14**, 522 (1983); S. Rosen, *Q. J. Econ.* **86**, 366 (1972); R. H. Smiley and S. A. Ravid, *ibid.* **98**, 353 (1983); A. M. Spence, *Econometrica* **52**, 101 (1984)].
25. The parameter, b, in Eq. 1 is related to the progress ratio, p, by the expression: $p = 2^{-b}$. Thus, an 80% learning curve would be indicated by a value for b in Eq. 1 of approximately 0.322.
26. Early studies of individual learning and forgetting were performed by H. Ebbinghaus [*Memory: A Contribution to Experimental Psychology* (Dover, New York, 1885)] and E. L. Thorndike [*Psychol. Rev. Ser. Monogr. Suppl.* **2**, 1 (1898)]. Recent syntheses of the literature on individual learning can be found in M. K. Johnson and L. Hasher, *Annu. Rev. Psych.* **38**, 631 (1987); W. A. Wickelgren, *ibid.* **32**, 21 (1981); Y. Anzai and H. A. Simon *Psychol. Rev.* **86**, 124 (1979); J. E. Mazur and R. Hastie, *Psychol. Bull.* **85**, 1256 (1978); and J. R. Anderson, *Cognitive Psychology and Its Implications* (Freeman, New York, 1985). Many researchers of individual learning fit their data to an exponential rather than to a power function. There is some evidence, however, that power functions may fit individual learning data better than exponential functions [A. Newell and P. S. Rosenbloom, in *Cognitive Skills and Their Acquisition*, J. R. Anderson, Ed. (Erlbaum, Hillsdale, NJ, 1981), pp. 1–55].
27. L. L. Thurstone, *Psychol. Monogr.* **26**, 1 (1919).
28. L. Greenberg, *Min. Eng.* **23**, 51 (1971).
29. L. Argote, GSIA Working Paper, in preparation.
30. L. Argote, S. L. Beckman, D. Epple, *Manage. Sci.*, in press.
31. M. B. Lieberman, *Rand J. Econ.* **15**, 213 (1984).
32. P. Conley, *IEEE Spectrum* **7**, 63 (1970).
33. Progress has been made on a theory of organizational learning [B. Levitt and J. G. March, *Annu. Rev. Sociol.* **14**, 319 (1988) and J. F. Muth, *Manage. Sci.* **32**, 948 (1986)]. The theory, however, does not explain why learning rates differ across production programs.
34. W. Z. Hirsch, *Rev. Econ. Stat.* **34**, 143 (1952); N. Baloff, *IEEE Trans.* EM-17, 132 (1970).
35. D. R. Sule, *Int. J. Prod. Res.* **21**, 771 (1983); T. L. Smunt and T. E. Morton, *IEE Trans.* **17**, 33 (1985); T. L. Smunt, *Int. J. Prod. Res.* **25**, 689 (1987).
36. T. E. Fandell, *Wall Street Journal*, 3 June 1974, p. 1.
37. U. E. Reinhardt, *J. Financ.* **28**, 821 (1973).
38. S. J. Sansweet, *Wall Street Journal*, 7 November 1975, p. 1.
39. "Lockheed loses hope," *ibid.*, 14 May 1980, p. 12; R. J. Harris, *ibid.*, 20 April 1982, p. 7.
40. E. Marshall, "Early data: Losing our memory," *Science* **244**, 1250 (1989).
41. C. A. Insko *et al.*, *J. Pers. Soc. Psychol.* **39**, 431 (1980); C. A. Insko *et al.*, *J. Exp. Soc. Psychol.* **18**, 557 (1982).
42. F. C. Lane, *Ships for Victory: A History of Shipbuilding Under the U.S. Maritime Commission in World War II* (Johns Hopkins Univ. Press, Baltimore, MD, 1951).
43. Reviews of the literature on the effects of turnover on groups and organizations can be found in R. L. Moreland and J. M. Levine [in *Psychology of Group Influence*, P. Paulus, Ed. (Erlbaum, Hillsdale, NJ, 1989), pp. 143–186] and D. R. Dalton and W. D. Tudor [*Acad. Manage. Rev.* **4**, 225 (1979)].
44. For discussions of the effect of managerial turnover on firm performance see O. Grusky, *Am. J. Soc.* **69**, 21 (1963); D. Helmich, *Acad. Manage. Rev.* **2**, 252 (1977); D. L. Worrell, W. N. Davidson, P. R. Chandy, S. L. Garrison, *Acad. Manage. J.* **29**, 674 (1986); J. Pfeffer and A. Davis-Blake, *ibid.*, p. 72.
45. G. S. Day and D. B. Montgomery, *J. Mark.* **47**, 44 (1983).
46. W. Cohen and D. Levinthal, *Econ. J.* **99**, 569 (1989).
47. N. K. Womer, *Manage. Sci.* **25**, 312 (1979).
48. The conventional learning curve in Eq. 1 can be expressed as a production function. The left-hand variable in Eq. 1 is labor hours per output ($y = h/q$). Thus, Eq. 1 can be rewritten in the form: $q = a^{-1}hx^{b}$. Thus, the production function implicit in the conventional learning curve has a single input, labor hours, and a single measure of the state of technical knowledge, x, cumulative output. This formulation also imposes the assumption of proportionality between output and hours.
49. R. Conway and A. Schultz, *J. Ind. Eng.* **10**, 39 (1959); N. Baloff, *ibid.* **17**, 25 (1966); N. Baloff, *Oper. Res. Q.* **22**, 329 (1971).
50. N. K. Womer, *Manage. Sci.* **30**, 982 (1984).
51. E. R. Bernt and M. S. Khaled, *J. Polit. Econ.* **87**, 1220 (1979); L. R. Christensen and W. H. Greene, *ibid.* **84**, 655 (1976); M. B. McElroy, *ibid.* **95**, 737 (1987); M. Fuss, D. McFadden, Y. Mundlak, in *Production Economics: A Dual Approach to Theory and Applications*, M. Fuss and D. McFadden, Eds. (North-Holland, Amsterdam, 1978), pp. 217–268; A. Zellner, J. Kmenta, J. Dreze, *Econometrica* **34**, 784 (1966); W. E. Diewert, *J. Polit. Econ.* **79**, 481 (1971).
52. R. Solow, *Rev. Econ. Stat.* **39**, 312 (1957).
53. Lockheed Aircraft Corporation, *Lockheed Annual Report* (Corporate Publications, Burbank, CA, 1972), p. 1; *ibid.* (Corporate Publications, Burbank, CA, 1974), p. 8; S. J. Sansweet, *Wall Street Journal*, 31 March 1976, p. 2; *ibid.*, 16 March 1977, p. 11; "Lockheed Corp. says losses on Tri-Star spurted in period," *ibid.*, 20 June 1980, p. 21; R. J. Harris, *ibid.*, 1 July 1981, p. 10; *ibid.*, 8 December 1981, p. 3; *ibid.*, 9 December 1981, p. 22.

54. We gratefully acknowledge the support of the National Science Foundation's Decision, Risk, and Management Science Program and Economics Program (grant number SES-8808711) and of the Center for the Management of Technology and Information in Organizations at Carnegie Mellon University. We thank E. Bailey, P. S. Goodman, R. S. Kaplan, L. Lave, T. McGuire, R. L. Moreland, T. Morton, and G. Salancik for their helpful comments.