PERGAMON

Intl. Trans. in Op. Res. 7 (2000) 595–607

INTERNATIONAL TRANSACTIONS IN OPERATIONAL RESEARCH

www.elsevier.com/locate/orms

The learning curve: a new perspective

Willard I. Zangwill[a,*], Paul B. Kantor[b]

[a]*Graduate School of Business, University of Chicago, 1101 E. 58th Street, Chicago, IL, USA*
[b]*Rutgers Center for Operations Research, Rutgers University, 4 Huntington Street, New Brunswick, NJ 08901, USA*

Received 21 February 2000; accepted 24 February 2000

Abstract

Despite the popularity of the ''learning'' curve, it should be called the ''forecasting'' curve. Given an industrial process, the traditional learning curve forecasts how costs are expected to decline (or other performance measures improve) in the future. Forecasting is certainly very useful. Nevertheless, the previous literature is notably mute on how the improvement or learning occurs. Without that information about how the process operates, the traditional learning curve cannot improve the rate of learning. This paper transforms the curve so that it not merely forecasts but rather actively creates the learning and knowledge. By extending work of the authors (Zangwill and Kantor, 1998), this paper constructs a new theoretical framework for learning and making improvements based upon learning cycles. This approach allows what is learned in one period to be intelligently applied to the next, improving the rate of learning right as production is occurring. Previous approaches could not do this. Because it should quickly increase the rate of learning, this approach might have important industrial applications.

Keywords: Learning curve; Progress curve; Experience curve; Continuous improvement; Forecasting; Learning; Improvement; Learning cycle; Lotka; Volterra; Cost reduction; Learning speed; Productivity improvement

1. Background

What has the learning curve traditionally done? Has it actually helped increase the rate of learning? For instance, say a firm manufactures a product. Each time the item is made,

* Corresponding author. Tel.: +1-773-702-7289; fax: +1-773-702-0458.
E-mail addresses: willard.zangwill@gsb.uchicago.edu (W.I. Zangwill), kantor@scils.rutgers.edu (P.B. Kantor).

0969-6016/00/$20.00
PII: S0969-6016(00)00022-8

the employees presumably discover how to make it better and less expensively. What the learning curve does is to forecast the expected amount of learning (cost reduction) that should take place. For instance, the learning curve (also called experience curve, progress curve) will predict that when the plant makes the 100th of an item, or the 1000th or 10,000th, what its cost should be. Indeed, since its discovery over half a century ago (Wright, 1936), the learning curve has been widely employed to forecast cost and other performance measures[1].

1.1. Forecasting is not learning

However, forecasting is a far cry from learning. In fact, a more accurate name for the traditional curve would be the "forecasting" curve. The learning curve should help a firm not just to forecast costs but to more rapidly cut costs (or more rapidly improve other important performance measures). Companies that learn faster are more likely to be successful. In this paper, we want to transform what up to now has primarily been a "forecasting" curve into an authentic learning curve. We want the curve to identify what specific techniques and actions produce improvement, and to do this quickly right in the midst of production. That allows management to apply these successful techniques immediately. Then the learning curve will genuinely increase the rate of learning, and will live up to its name.

To provide the background for this, let us first explore how the learning curve is able to forecast. Once that is done, it will be clearer how to make the learning curve help improve and accelerate learning.

1.2. Why the learning curve forecasts

What makes forecasts possible was the discovery that the cost reduction (learning) did not occur in a random, haphazard way. On the contrary, it occurred so regularly that the expected cost reduction could be described by a simple curve, the power law. (The power law is quite common, although there are other forms such as the exponential, finite and a general formulation, as discussed in Zangwill and Kantor (1998)). In particular, the amount of cost reduction was seen to depend upon the cumulative output, q. Explicitly, let $C(q)$ be the expected cost (or other metric) of making the qth item in a series and $C(1)$ be the cost of producing the first one. Then for p, the parameter of learning, the classical power law becomes

$$C(q) = C(1)q^{-p} \tag{1.1}$$

In many circumstances, time (t) is substituted for the quantity producred (q).

Given a particular operation in a company, the parameter p typically holds for a class of

[1] The learning curve can be employed for almost any performance measure such as cycle-time, defects, customer satisfaction, software bugs and so on (Zangwill and Kantor, 1998)

items, say aircraft engines, and is empirically determined. In this way, it then becomes possible to forecast ahead of time what the costs should be. The learning curve was born.

1.3. Problems

Yet, a variety of problems exist with the learning curve. The overriding one, as mentioned, is that the learning curve does not improve learning. In Eq. (1.1), the learning parameter p is fixed. If the learning curve were actually to enhance learning, then p would increase as more items were made, reflecting the faster learning rate. No longer would the learning curve passively forecast, as it does now, but it would actively accelerate the learning as production is occurring.

To resolve this most basic problem, however, we first require understanding of some other problems with the learning curve, because the problems are interconnected and all require resolution.

2. Problem two: empirical basis

On the heels of the basic problem is a second problem, namely, that the learning curve is empirical. We have already noted that the parameter p was obtained empirically. But, even more fundamentally, the learning curve itself is a totally empirical observation.

With its complete basis empirical, there is no underlying theory and no real understanding of how it operates. And without that knowledge and theory it is very difficult to improve the rate of learning. Mishina (1987) underlines this point by noting:

> "Once one starts questioning what is going on, however, he or she is likely to encounter more and more questions. This is mainly due to the fact that the concept of the learning curve has been established as a strictly empirical phenomena. Although the learning curve literature has accumulated significant amount of evidence, little has been done as to the explanation of the phenomena. As a natural result, we are left with numerous confusing arguments and irreconcilable recommendations."

This paper provides a theoretical construct designed to help explain the learning curve better, pinpoint the improvements, and eliminate many of the issues Mishina observed.

3. Problem three: after the fact observation

Part and parcel with the empirical issue is a third problem, specifically, that previous analyses were longitudinal and focused on the curve's behavior over a long time. Of course, this macro entire curve type of analysis did produce some interesting results.[2] Nevertheless, and this is the point here, detecting these results required examination of the curve over an extended period, perhaps years. By the time any conclusions could be drawn, it was too late, the production was over. To improve a production process specific actions must be taken, say a machine adjusted or software program changed, as production occurs. The traditional macro analysis, however, does not detect these individual actions, because it looks longitudinally and not at individual events.

Identifying what happens in the micro realm, the approach of this paper, requires investigation on an incident-by-incident basis. This is done, as will be discussed later, by use of "learning cycles". This approach considers individual improvement efforts (the learning cycles) as mini-experiments, and are generally done quickly, in one or two periods. The experiment's outcome exposes which improvement techniques are effective and which are not. The most successful techniques then can be applied immediately, thereby improving the overall learning rate at once. Traditional approaches, being after the fact, cannot accomplish this.

4. Problem four: statistical jumble

What is worse, with the traditional approach it is virtually impossible to determine the effect of any specific advancement or improvement. This is another problem with the traditional approach and is quite serious. Suppose management wants to determine whether the strategies it uses in a period really help to reduce the cost (see Zangwill and Kantor, 1998). The "obvious" way to determine this information is to subtract the total cost of making the item at the end of the period from the total cost at the end of the previous period. Using t for time, the amount of improvement is thus

$$\Delta C(t) = C(t) - C(t-1) \tag{4.1}$$

[2] For instance, Abernathy and Wayne (1974) noted that for the production of Ford's model T, during the period from 1909 to 1923 the cost of making the model dropped from $5000 to less than $900. Similarly, studies identified the learning curve in almost every industry including motorcycles, chemical, paper, steel, electronic, knit production, mechanical goods, and even productivity at an Israeli Kibbutz (Day and Montgomery, 1983; Dutton and Thomas, 1984).

More recently, the learning curve has been proven to be very important in high technology, where the production process is often very complex and often not very stable. (Ghemawat, 1985; Terwiesch and Bohn, 1998). In this regard the famous Moore's Law of computer technology predicts that the cost of computer capability will drop by a factor of 2 in 18 months. Although this is not a power law directly, if a power law applies to the whole industry, and the industry grows exponentially, such a result could be predicted (See Zangwill and Kantor, 1998).

Likewise, looking at the curve as a whole suggested that various factors might contribute to learning. Terwiesch and Bohn (1998), Day and Montgomery (1983) and Uzumeri and Nembhad (1998) summarize a number of studies which suggest that the learning depends on such factors as the training of labor, capital improvements, technology advancements, organizational structure and so on. Adler and Clark (1991), for instance, examine engineering changes and workforce training. A recent reserch path has examined issues such as forgetting due to delay or distributions in the production. (Argote et al., 1990; Argote, 1993, and see discussions in Eden et al., 1998; Li and Rajagopalan, 1998).

Although calculating Eq. (4.1) appears straightforward, it is not. To demonstrate, suppose the cost in a period is 100. Let the cost in the next period be 95. Using basic subtraction, the resulting cost improvement is 5. That result, however, is generally not correct. Strange as it seems, subtraction is wrong[3].

The issue, as Bohn (1991) stresses, is that the usual cost data are extremely noisy. In reality, the cost is not exactly 100, but is 100 with an error of (say) ± 10 (standard deviation of 10). Similarly, let the cost in the next period be 95 with an error of ± 10. Suppose the common case where the errors are independent and normally distributed. Then after subtraction and considering the even distributions, the improvement is seen to be 5 ± 14.

Here, although the initial error was about 10% of its base value, after the subtraction to calculate the improvement, the error has ballooned to 280% of the calculated difference. The real value of the improvement is almost anyone's guess.

This example is not an anomaly. An extremely large error will nearly always result even if the data are reasonably correlated because the subtraction involves two large numbers that are similar in value. Due to this inherent statistical problem, the traditional learning curve approach is so inaccurate that, for any practical application, it fails and cannot determine what the real improvement is.

4.1. Analysis of errors

To understand and overcome this problem, note that the total cost of producing an item is actually the sum of the costs of numerous small activities, say $a_1, \dots, a_n$, each with its own random variation. These small activities might be: entering sales data, billing, contacting clients, training, telephoning, heat etc. Assuming $n = 100$ activities, cost can be written as

$$C(t-1) = a_1 + a_2 + a_3 + \cdots + a_{100}$$

Suppose during period t, management attempts to reduce the cost of some activities. Let b_i denote the cost of those particular activities at the end of the period. If management works on two activities during the period, the cost of activity $i = 1, 2$ is cut from a_i to b_i. Since management does not work on any of the other activities, their costs should be unchanged except for random disturbances. Let those costs at the end of period t be designated $a_3^*, \dots, a_{100}^*$. The total cost in period t is thus

$$C(t) = b_1 + b_2 + a_3^* + \cdots + a_{100}^*$$

Calculating the improvement, ΔC, from Eq. (4.1),

$$\Delta C(t) = C(t) - C(t-1) = [a_1 + a_2 + a_3 + \cdots + a_{100}] - \left[b_1 + b_2 + a_3^* + \cdots + a_{100}^*\right] \tag{4.2}$$

Using the approach of subtracting total costs in Eq. (4.2), the improvement calculated includes the errors of two hundred cost estimates, which is why the error is so large. Some observers

[3] Of course, the substraction itself is not incorrect. But it can be misapplied. A similar misapplication is the calculation of profits by substracting costs from revenues without considering uncertainity.

have suggested that the error might be reduced by taking a long series of observations. But that does not help because it takes too much time to obtain the data, perhaps many periods. Our goal is to quickly obtain information so that we can rapidly learn how to make better improvements. To overcome some of these error difficulties, we propose to measure the amount of cost improvement by looking only at the activities targeted for improvement. In the above example, since only two of the 100 activities were changed, we can calculate the improvement using the equation

$$\Delta C(t) = (a_1 - b_1) + (a_2 - b_2) \tag{4.3}$$

Since this equation has the errors of only four terms to contend with, the error is substantially reduced. The sizable random variation caused by the other 98 components is eliminated.

4.2. Example of obtaining data

But how can we obtain the data needed for Eq. (4.3)? To illustrate, suppose management moves machine A next to machine B, eliminating the need to forklift the material from A to B, which produces as cost saving per unit of \$2. Consequently, the quantity $a_1 - b_1 = \$2$.

As a second action during the period, suppose management revises the software on their enterprise resource planning system. Suppose that reduces the scrap and defect rates, thereby cutting from the cost of making the item by \$6. Hence $a_2 - b_2 = \$6$. Thus,

$$\Delta C(t) = a_1 - b_1 + a_2 - b_2 = \$8. \tag{4.4}$$

for a cost reduction of \$8.

While this \$8 value undoubtedly has some error, it is likely to be a small fraction of the error produced by subtracting the total costs as in Eq. (4.1) or (4.2). This is because the errors of only a very few terms accumulate in Eq. (4.3), but Eq. (4.2) bears the error of 200 terms (assuming 100 small activities).

This approach of looking directly at the actions improved assumes that the improvements made do not interact and that other parts of the process are not affected. In practice, these assumptions should be checked, and when necessary some or all of the additional terms in Eq. (4.2) might have to be included. In most situations, however, only a few of the terms in Eq. (4.2) will be needed.

The traditional macro approach of looking at the learning curve over a long period of time is clearly inadequate. The approach herein suggests a period-by-period perspective, and by focusing on only the costs actually changed, is more likely to obtain the data to examine what happened.

4.3. Learning cycle

The key to this paper's micro approach is that each period is considered an opportunity to conduct a learning cycle (actually a learning experiment). In each period an action is taken, say a change is made, machine is adjusted or software program altered. Then, at the end of the period, the data are examined to determine if an improvement occurred. This comprises the learning cycle. By repeatedly doing learning cycles, we produce knowledge about which actions work, which do not, and how to improve the process.

4.4. Difference equation

But how do the results of the learning cycles get incorporated into the learning curve? We will describe the learning curve using a difference equation. The difference indicates the change in learning that occurs in a learning cycle. As will be provided mathematically below, integrating the differences (differentials) then yields the entire learning curve.

The difference equation permits the knowledge gained during the learning cycles to be directly incorporated into the learning. If the learning cycle produced a sizable improvement, the difference will be large for that period. If there is no improvement, then the difference will be zero for that period.

This approach provides management the information to learn more rapidly. To illustrate, during one learning cycle management might conduct a pilot project to train some workers. If the training helps, then that training can be done with all relevant employees. For another learning cycle, management might adjust the temperatures and pressures in certain equipment. If that helps, the adjustment could be installed permanently and employed on other similar machines. It is by conducting such learning cycles that the learning rate is improved. The traditional learning curve model does not do this, but the approach of this paper does.

5. Problem five: overestimation of improvements

Another issue with the traditional learning curve approach, as Lapre et al. (1998) discovered, is that most attempts at improvement fail to produce improvement.[4] Management often assumes that virtually all improvement projects work successfully and make the process operate more effectively. But, when closely examined, that turns out to be false. Most attempts to make improvement do not succeed. That corresponds to what Zangwill (1999) has determined, as well.

This suggests that many efforts at employee training, to make technical improvements or to enhance processes do not, in fact, produce improvement. The traditional learning curve approach does not have a way of validating if the improvement really worked, and so can lead to difficulties. This fact spotlights the importance of the learning cycle since its testing and experiment type of orientation is more likely to distinguish what works from what does not.

5.1. Viewpoint

Overall, there are a number of reasons why the traditional approach to the learning curve

[4] The numbers they obtained are dramtic. Only about 25% of the improvement projects helped. About 50% had little affect, while 25% harmed performance. The study was done in a plant that was considered very well managed, and had received awwards for that.

does not enhance the rate of learning. The approach of this paper, because it considers individual attempts to improve as learning cycles done period-by-period, provides a framework that is directly focused on isolating the specific causes of learning and on increasing the tempo of learning.

6. Mathematical development

We now provide the mathematical structure. (Although focused on a different context, Zangwill and Kantor (1998) contains more details.)

Since learning can be associated with metrics (measures) other than cost, we employ a more general metric that could also represent cycle-time, defects, performance inadequacies, flaws, bugs, customer dissatisfaction and so on. In particular, after q items are produced, let $N(q)$ be the amount of metric remaining to eliminate. Intuitively $N(q)$ represents non-value added work.

Next let $E(q)$ be the fraction of $N(q)$ eliminated in the period after q items are produced. $E(q)$ measures "management effectiveness", or how effectively management is improving the process.

6.1. Lotka–Volterra equations

The differential equation relating $E(q)$ and $N(q)$ is of the Lotka–Volterra genre. More precisely, in the 1920s Lotka and, separately, Volterra (see Murray, 1989) suggested an equation describing populations of predators and prey. As a simple instance, wolves prey on moose. Over a wide range of conditions, the rate of elimination of moose over time ($-\mathrm{d}S/\mathrm{d}t$) will be proportional to:

(a) the number of moose $S(t)$, as the opportunities for each wolf are proportional to the number of moose,
(b) the number of wolves $W(t)$, as the moose population will decline in proportion to how many wolves are preying upon them.

This analysis led Lotka and Volterra to the multiplicative relationship:

$$\frac{\mathrm{d}S}{\mathrm{d}t} = -aW(t)S(t).$$

In the present context, the "prey" is the metric left to eliminate $N(q)$. The "predator" is the management, or more precisely the effectiveness of management's effort to improve the process. Also, $E(q)$ denotes the effectiveness of management.

Consequently, we may rewrite the Lotka and Volterra observations as follows. The rate of improvement in the metric is proportional to

(i) the effectiveness of management, $E(q)$, and
(ii) the amount of metric left to eliminate, $N(q)$.

6.2. The learning curve differential equation

Assuming that low values of the metric are "better",

$$\frac{-\mathrm{d}N}{\mathrm{d}q} = \text{ rate of improvement in the metric} \tag{6.1}$$

Then from the Lotka–Volterra theory, for c, a coefficient of proportionality:

$$\frac{\mathrm{d}N}{\mathrm{d}q} = -cE(q)N(q) \tag{6.2}$$

To measure improvement in any practical application, express Eq. (6.2) in finite differences:

$$\frac{\Delta N}{\Delta q} = -cE(q)N(q) \tag{6.3}$$

These forms, (6.2) and also (6.3), are the "learning curve differential (or finite difference) equation" (LCDE).

Intuitively, the term $\Delta N/\Delta q$ represents how quickly learning is taking place at any particular instant in time. It shows how much the metric improves, ΔN, over the interval when Δq items are made. That term depends on two primary factors. One, $E(q)$, is the "effort" that management is devoting. The effort term subsumes not just time and resources, but also ingenuity and skill. The greater the "effort", the faster the learning. The second term $N(q)$, the non-value added, depicts the waste, duplication, poor practices, defects, lack of training and so on that is making the process less effective and efficient than it should be. Early in the production cycle $N(q)$ is likely to be large since the non-value added work is likely to be large and there is much to learn. After many improvements are made, $N(q)$ should become much smaller.

7. Application

In applying the LCDE to actual operations, $N(q)$ should be known, since that is the value of the metric, say cost, after q items are processed. Further, c would be set to the scale of the data, so without loss of generality we can set $c = 1$. To apply the LCDE, thus, we need only determine the rate at which the metric improves, $\Delta N/\Delta q$, as that will provide the effectiveness of management: $E(q)$.

Specifically,

$$E(q) = \frac{-\Delta N}{\left(\Delta q N(q)\right)} \tag{7.1}$$

During any period of interest, Δq is the number of items processed. (If we are examining how the metric improves from the time the 36th item is made until the 39th item is made, then $\Delta q = 3$.)

Next we must estimate ΔN, or how much the metric changed during that interval.

But this is simply the outcome of the improvement effort. For instance, if the experiment cut the cost to produce an item by \$2, then $\Delta N = 2$. (And similarly for other metrics such as yield, cycle-time and so on.)

With that done, $E(q)$ is then known. The improvement efforts that produce a high $E(q)$ would be implemented and also applied more widely to other applications and areas.

In actual practice, suppose that during different intervals of time a firm attempts various techniques in order to obtain improvements, such as SPC, moving the machine, improving maintenance, training. Then for any specific technique attempted, calculating its $E(q)$ determines its ability to reduce the metric. That information lets us identify *which improvement techniques are more powerful and therefore should be implemented.*

8. Power law

But, what about the power law, Eq. (1.1)? How can that be derived from the LCDE? For a parameter, $\kappa > 0$, and coefficient K, let

$$E(q) = KN(q)^{\kappa} \tag{8.1}$$

Eq. (8.1) represents how management changes its level of effort (effectiveness) as improvement is accomplished. In particular, $N(q)$, being the non-valued added work, decreases as improvements are being made. The parameter $\kappa > 0$ reflects that the effectiveness of management then decreases also. That typically occurs because after the large initial improvements are obtained and the improvements start to get smaller, management's attention generally begins to shift to other projects. Eq. (8.1) thus models the common phenomenon that management decreases its attention and effort over time as the severity of the problems diminish.

Plugging Eq. (8.1) into the LCDE Eq. (6.2) we obtain

$$\frac{\mathrm{d}N}{\mathrm{d}q} = -Kc\big(N(q)\big)^{\kappa+1}. \tag{8.2}$$

For $\kappa > 0$, the solution is a power law, which, as in Eq. (1.1), is the classical form of production learning.

$$N(q) = N(1)q^{-p} \quad q > 1 \tag{8.3}$$

The methodology of this paper thus reproduces the traditional power law of learning. The LCDE can also generate the other forms of the learning curve (see Zangwill and Kantor, 1998). The LCDE thus yields the traditional forms of the learning curve, but is more general. By varying the terms of the LCDE, almost any type of industrial learning can be tracked. But, of course, the LCDE goes much farther since also it helps identity how to accelerate the learning.

9. Overview of approach

This paper attempts to reorient the viewpoint and application of the learning curve. Conceptually, the difference (differential) equation and the learning cycle are symbiotic. The learning cycle is an experiment that evaluates if a given improvement effort actually produces an improvement. The difference equation then incorporates the result of the experiment into the mathematical framework of the learning curve. For instance, suppose in a given learning cycle that a specific improvement technique works and decreases cost. Then the amount of that improvement will show up as the value of the difference equation for that period.

In this way we quickly identify the improvement and can employ the successful technique in other similar contexts forthwith. By applying this approach consistently, we should be able to identify many techniques and approaches that help, deploy them on a wider scale, and thereby lift the overall rate of learning.[5]

The traditional learning curve cannot accomplish this since it lacks a difference equation. Because of that, it cannot identify which techniques are working and which are not working, period-by-period.[6] This inability to easily detect whether an approach helps or not severely impedes its ability to enhance and accelerate the learning. Indeed, this explains why, as mentioned above, the traditional learning curve has been used for prediction instead of for accelerating the learning.

Moreover, summing the difference equation (integrating the differential equation) produces the learning curve as a whole. Thus, the approach in this paper generates the traditional curve as well. The micro approach thereby simultaneously provides the macro.

10. Problems revisited and resolved

This paper's approach also resolves many of the problems posed previously, as follows:

Resolution of problem 1 — forecasting versus learning: When the difference equation is used with learning cycles, we obtain a means to identify which improvement techniques are successful. Since we can do this each and every learning cycle, we can employ the information gained immediately to boost the rate of improvement. The approach of this paper thus helps change the role of the learning curve from prediction to that of accelerating the learning.

Resolution of problem 2 — empirical basis: The foundation of the curve is no longer purely empirical. The differential/difference Eqs. (6.2) and (6.3) provide a conceptual framework. Moreover, both the efficiency of management term, $E(q)$, and the non-value added term, $N(q)$, can be further broken down into their underlying components. That enables them to be examined at greater depth. It also permits the theory to be extended because the more we understand the factors that comprise $E(q)$ and $N(q)$, the more we can extend the theory.

[5] See Lapre et al. (1998) who discuss how and when an improvement technique is found to help one process, it is spread to other processes throughout the firm.

[6] Recall as mentioned above, most improvement techniques do not help.

Resolution of problem 3 — after the fact observation: Traditional approaches analyze data over an entire production run, which might take years, so results are after the fact and too late. This paper's approach employs the learning cycle or period as the unit of analysis. This permits the examination of individual incidents as they occur and enables improvements to be made right away. This increases the rate of learning in the midst of production, something that is very important in today's highly competitive environment.

Resolution of problem 4 — statistical jumble: To reduce the statistical error, discussed above, the approach herein looks only at the costs (or other metric) changed. This becomes possible using the conceptualization of this paper because examination of changes fits the difference equation framework. (The difference is the change.) The approach of examining changes is almost impossible under the traditional framework because it considers a long period of time. Over a long period of time costs change for a variety of reasons, making it difficult to isolate what is happening.

Resolution of problem 5 — overestimation of improvements: As noted, many, if not most, improvement efforts fail. In the traditional framework, a failed improvement is difficult to discern because such a long period of time is examined that individual improvement efforts get lost in the statistical noise. Using learning cycles, however, such failures can generally be determined quickly, virtually as they occur.

11. Conclusions

By building on an earlier paper (Zangwill and Kantor, 1998), this paper attempts to articulate a conceptual framework for the learning curve. Previously the curve was employed almost exclusively for prediction and was a "forecasting" curve. The approach of this paper, by employing the learning cycle concept, provides information on which improvement efforts fail and which succeed. Since this is done on a cycle-by-cycle or period-by-period basis, we can employ that information at once to enhance the tempo of learning. And this can be done while production takes place so that learning improvement becomes dynamic. In this way, the curve can begin to live up to its name.

Acknowledgements

The research by Paul B. Kantor was supported in part by the National Science Foundation under Grant SEES 88-21096 to Tantalus Inc.

References

Abernathy, William J., Wayne, Kenneth, 1974. Limits of the learning curve. Harvard Business Review No. 74501, (Sept.–Oct.) 109–118.

Adler, Paul S., Clark, Kim B., 1991. Behind the learning curve: a sketch of the learning process, Management Science 37(3), 267–282.

Argote, L., 1993. Group and organizational learning curves: individual, system and environmental components. British Journal of Social Psychology 32, 31–51.

Argote, L., Beckman, S., Epple, D., 1990. The persistence and transfer of learning in industrial settings. Management Science 32, 140–154.

Bohn, Roger, 1991. Noise and learning in semiconductor manufacturing. Report, University of California, San Diego.

Day, George S., Montgomery, David B., 1983. Diagnosing the experience curve. Journal of Marketing 47, 44–56.

Dutton, John M., Thomas, Annie, 1984. Treating progress functions as a managerial opportunity. Academy of Management Review 4 (2), 235–247.

Eden, Colin, Williams, Terry, Ackerman, Fran, 1998. Dismantling the learning curve: the role of disruptions on the planning of development projects. The International Journal of Project Management 16 (3), 13–138.

Ghemawat, Pankahj, 1985. Building strategy on the experience curve. Harvard Business Review No. 85206, (March–April) 2–8.

Lapre, Michael A., Mukherjee, Amit Shankar, van Wassenhove, Luk N., 1998. Behind the learning curve: linking learning activities to waste reduction, INSEAD. Working paper 98/66/TM (CIMSO 5).

Li, George, Rajagopalan, S., 1998. A learning curve model with knowledge depreciation. The European Journal of Operational Research 105, 143–154.

Murray, J.D., 1989. Mathematical Biology. Springer–Verlag, Berlin, pp. 63–72.

Mishina, Kazuhiro, 1987. Behind the flying fortress learning curve. Harvard Business School. See Learning by new experiences: revisiting the flying fortress learning curve. In Lamoreaux, Naomi, Raff, Daniel, Temin, Peter (Eds.), Learning by Doing in Markets, Firms, and Countries. University of Chicago Press, Chicago.

Terwiesch, Christian, Bohn Roger E., 1998. Learning and process improvement during production ramp-up. Report 98-01, The Information Storage Industry Center, University of California, San Diego.

Uzumeri, Mustafa, Nembhad, David, 1998. A population of learners: a new way to measure organizational learning. Journal of Operations Management 16 (5), 515–528.

Wright, T.P., 1936. Factors affecting the cost of airplanes. Journal of Aeronautical Sciences 3, 122–128.

Zangwill, Willard I., 1999. Business 368, operations strategy and management: class notes. Graduate School of Business, University of Chicago, Chicago, IL.

Zangwill, W.I., Kantor, Paul B., 1998. Toward a theory of continuous improvement and the learning curve. Management Science 44 (7), 910–920.