The Dynamics of Organizational Learning

Bernardo A. Huberman

Dynamics of Computation Group
Xerox Palo Alto Research Center
Palo Alto, CA 94304

Abstract

This paper presents a dynamical theory of organizational learning that explicitly accounts for both the introduction of new routines into the manufacturing process *and* improvements in the selection of which procedures to follow. Besides producing a power law of organizational learning with a rate that depends on the effectiveness of the decision procedure, the theory also accounts for observed anomalies characterized by price increases with cumulative output.

May 15, 1997

1. Introduction

The phenomenon of organizational learning is one of the few strong empirical regularities observed in the life of institutions. Simply stated, organizational learning refers to the increases in productivity that are observed as firms gain experience in production[1]. [4, 5]. What makes this phenomenon interesting is that the rate at which these increases occur is nearly universal, in the sense that regardless of the product, the unit cost of production usually decreases as a power law of the number of units produced, with a functional form, $y = y(1)x^{-b}$, where $y(1)$ is the number of hours required to produce the first unit, x is the xth unit, and b is an exponent that tends to vary across industries but that it is typically observed to have a value near 1/3 [6]. It is this value that underlies the common observation that each doubling of cumulative output roughly leads to a 20% reduction in unit cost.

The power law behavior of the learning curve has been observed in a wide class of industries, such as egg production, aircraft manufacturing, surgical procedures, nuclear plants, the production of umbrellas and the operation of fast food franchises, to name a few [7–9]. In addition, a number of thorough empirical studies have highlighted some of the factors that are important in determining some of the mechanisms underlying organizational learning and the rate at which it takes place [10]. It is known that contrary to what would seem obvious, in many cases employee turnover does not strongly affect learning rates, whereas in others it can hurt organizational productivity. More specifically, movement of new employees at moderate levels enhances productivity while the departure of high performing individuals tends to lower it [11]. Other studies have shown that low turnover, large enterprise size and high productivity are highly correlated [12]. Also, the ability to transfer knowledge from one part of the organization to the other, or among different organizations delivering the same product, has an influence on the rate of learning [10]. Equally important, it has been shown that organizational learning can depreciate in the sense that recent output is at times a more reliable predictor of current production than output from the distant past [11].

[1] The term organizational learning is sometimes used to describe the mechanisms of learning by individuals embedded in an organizational setting [1]. While individual learning also exhibits a power law of practice [2, 3], and affects the performance of a firm, it is not the sole determinant of organizational learning and remains distinct from the latter.

The ubiquitous nature of the learning curve points to the existence of general mechanisms responsible for its existence, mechanisms that while reflecting the procedures of an organization do not strongly depend on the particular nature of the product or service being created. Thus it is of interest to produce a theory of the learning curve that considers production as a dynamical entity and that also accounts for the factors that determine the rates at which learning takes place.

Several attempts have been made at providing a theoretical understanding of the learning curve, mostly from a trial and error approach to the process of manufacturing [13, 14]. An often quoted search theory of production by Muth [14] generated a power law by considering a random search for lower cost methods among a fixed population of technological improvements. One problem that one encounters with these trial and error theories is that they do not account explicitly for the different observed rates of learning in terms of parameters that can be associated with the manufacturing processes. Another objection that has been raised [15] concerns the fact that these approaches assume that the routines underlying organizational behavior are fixed, whereas it is known that procedures do change as the organization learns which new ones to pursue [16, 17].

These shortcomings point to the need for a theoretical framework that can account for both the varying rates that are observed in learning curves, and the fact that routines or procedures are dynamical entities in the life of organizations. This paper presents such a theory. Its main assumption is that organizational learning stems both from the introduction of new shortcuts and procedures into the manufacturing process *and* the improvement in the selection of which routines to follow. By taking into account the dynamics of procedural changes brought about by the discovery of shortcuts while searching for better and faster ways of producing a given product, we show that a power law of learning follows quite naturally.

Besides producing a power law of learning with a rate that depends on the effectiveness of the decision procedure, this theory also accounts for anomalies that have been reported in the literature. These anomalies are characterized by situations where the price of a product actually increases as more units are produced [18]. They arise from a dynamical effect whereby the effectiveness of the organizational search cannot compensate for the high rate at which new routines enter the organization.

2. Learning by Doing

Consider a general manufacturing process, which can be described as the conversion of a number of inputs into a finished product through several intermediate stages. These stages do not only involve the assembly of particular parts into components but also other processes and organizational routines associated with them, such as keeping inventories, writing and propagating documentation through the firm, and the many other activities involved in producing a physical product or service.

The production process can be mathematically represented by a connected graph whose nodes represent stages in the manufacturing process and the links the procedures that connect them. The goal node represents the finished product. There are many ways in which parts can be assembled together in order to come up with manufactured components, and this translates into the fact that the goal node representing the finished item can be reached from many intermediate states. Since each step requires a certain amount of time required to be completed, the total cost of manufacturing a given product will be proportional to number of steps (or links traversed in the graph) that it takes to reach the goal node. Within this representation, the learning process consists in finding decreasingly shorter paths from the initial node (the inputs) to the goal node (the output).

In this model, learning takes place in two ways. First, a new procedure or routine can be discovered for going from the initial to final states, in analogy to a shortcut. Depending on the generality of this new procedure, this corresponds to adding one or more links to the graph. This new procedure can be general enough so as to be transferable to other parts of the production process or even to other organizations if it is made public. Notice that if the number of old processes that a new shortcut replaces is large, this new procedure will lead to an effective organizational "forgetting", making the new routines more important in determining the time to production than the old ones.

Second, the ability to select an appropriate procedure among many possible ones can improve, as is the case when search takes place. This means that out of the very many possible links emanating from each node to reach the following one, the organization can search for those that lead to a shorter path.

After either one or both aspects of this learning process take place, a new and hopefully shorter path is followed. As this process repeats the associated performance changes accordingly. Thus formulated, the problem consists in calculating the expected number of steps required to reach the final node as a function of the number of trials when learning takes place.

In order to do so, we need to point out that we are interested in the learning behavior of typical organizations, as opposed to specific ones. This means that all possible procedure spaces should be considered and that there is no particular basis to choose one set over the other. This translates into the consideration of a space that has the structure of a random graph, i.e. a collection of nodes with an initial set of links between them that are effectively distributed at random. As learning by the organization takes place new links are added independently, shortening the path to the goal. Similarly, the decision procedure consists in the selection of which link to follow when there are multiple ones out of any given node.

The process of adding links to the graph is mathematically equivalent to initially having them exist between each pair of states with independent probability, p, and having p increase as new procedures are learned during the manufacturing process [19]. Also, the effectiveness of the search part of the process can be described by eliminating at any node each unproductive link with probability r. Thus, an improvement in the decision procedure corresponds to an increase in the value of r. Note that r=*1* corresponds to a perfect decision procedure in which search and backtracking are never required, while r=*0* corresponds to a random walk on the graph.

To obtain an explicit expression for the average length of the manufacturing process, y, in terms of p and r, we will assume that all nodes of the graph have the average number of links $\mu = (n-1)p$. If the cycles in the graph are typically long, at any node there will be one link to a node one step closer to the goal while the others will make the path to the goal longer. This makes the procedure space into a regular graph consisting of n nodes with uniform branching ratio μ . Because of the initially exponential growth in the number of nodes with distance, the initial node and the final one will be typically separated by a distance given by the diameter, D, of the graph, which is given by $D = \frac{\ln n}{\ln \mu}$, or in terms of the number of shortcuts, by

$$D = \frac{\ln n}{\ln (n-1)p} \tag{1}$$

From most nodes, there is only one short path to the goal state so that there will be one correct link to choose (which will move the path one step closer to the goal) and μ—1 incorrect ones (which will move the path one step farther from the goal). Since however, the decision procedure eliminates each incorrect choice with probability r, there are effectively only (μ-1)(1–r) incorrect choices at each node. Thus the whole process is equivalent to a biased random walk in which one starts at a distance D from the goal and moves randomly until the goal is reached. This motion is constrained to remain within distance D of the goal and, at each step, to move towards the goal with probability p_-, given by

$$p_- = \frac{1}{1 + (\mu - 1)(1 - r)} \tag{2}$$

and away from the goal with probability $p_+ = 1 - p_-$. Within this model the average time required to reach the goal as a function of the number of shortcuts learned and the effectiveness of the decision procedure can be readily computed, as it is basically a first passage time in the random graph that we described above [20]. It is given by [3]

$$y = \frac{2}{(2p_- - 1)^2} \left[p_- + p_-(D-1) - (D/2) + \frac{(1-p_-)^{D+1}}{p_-^{D-1}} \right] \tag{3}$$

Eq. 3 describes the dynamics of organizational learning within the model described above. As we shall see, it produces a number of interesting behaviors, including a power law decrease in the number of steps to finish a product, both as a function of the shortcuts discovered and of the effectiveness of the decision procedure. This power law is characterized by an exponent whose value depends on the effectiveness of the decision procedure, r.

3. Shortcuts, Learning Rates, and Anomalies

The full dependence of the time to reach the goal, or the cost, y, as a function of both the number of shortcuts, p, and the effectiveness of the decision procedure, r, as given by Eq. 3 is shown in Fig. 1 As can bee seen, for a wide range

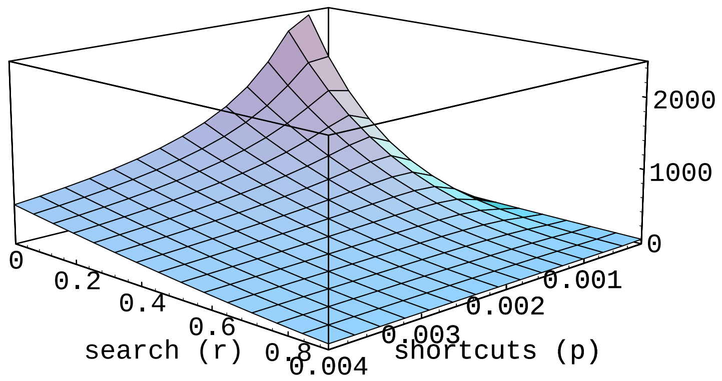

Fig. 1. The time to reach the goal, y, displayed in the vertical axis, as a function of the search efficacy and the number of shortcuts (or new procedures) introduced into the production of a good or service. The graph was computed for a random graph with 10,000 nodes.

of parameter values one observes a learning curve both along the search and shortcuts dimensions.

In order to analyze in detail the solutions exhibited by Eq.3, we first notice that since most of the time organizations engage in local improvements rather than exploration of novel possibilities, it makes sense to start by looking at the dependence of y on the number of shortcuts introduced as learning takes place, for a fixed value of r. In most organizations the latter changes slowly due to improvements in search procedures.

For the case where the decision procedure associated with search is good (r nearly one), the dependence of y on the number of shortcuts is shown in Fig2 for the case of a graph with 10,000 nodes.

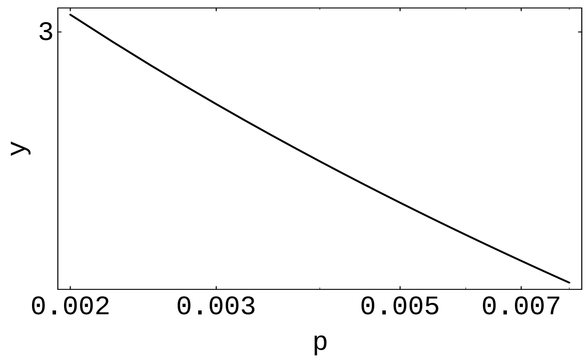

Fig. 2. Log-Log plot of the time to produce reach the goal, *y*, versus the number of new shortcuts for a nearly optimal value of the decision procedure (*r*=0.8) Notice that although the random graph contains 10,000 nodes, the path to the goal is extremely short..

As can be seen, it exhibits a linear behavior on a log-log scale, corresponding to a learning curve with a power law. Analytically, this follows from the fact that in the limit of optimal decision procedures, (*r*=1), $p_- = 1$ and *y* scales like the diameter of the graph, *D*, which in terms of the number of new shortcuts is equivalent to

$$y \approx \frac{\ln n}{\ln (n-1) + \ln p} \tag{4}$$

Moreover, away from the limit of $r = 1$, the slope of the curve, which measures the rate at which productivity grows, exhibits a strong dependence on the effectiveness of the decision procedure. This is observed in Fig.3 where we plot the slope of the learning curve at a particular point in the learning process (i.e. *p=0.003*) as a function of *r*.

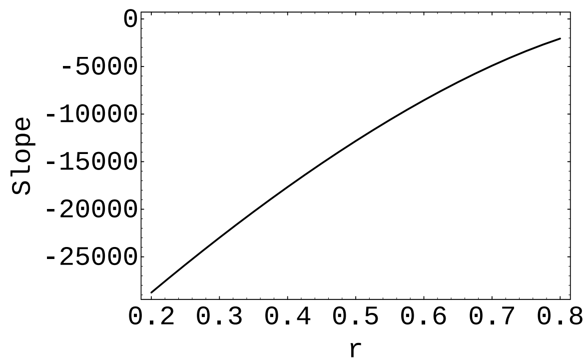

Fig. 3. Slope of the learning curve for a graph of 10,000 nodes as a function of the effectiveness of the decision procedure, at the value of the number of shortcuts, *p=0.003*.

As this clearly shows, very small changes in the effectiveness of the decision procedure bring about large variations in the learning rates.

In the case of a poor decision procedure (i.e. r very small) one also obtains a power law of learning, with a straight line in a Log-Log plot, but it is characterized by extremely long times to reach the goals. For example, if r=0.1 the number of steps to reach the goal in a random graph of 10,000 nodes varies between 1500 and 200, as p ranges from 0.001 to 0.091. This is to be contrasted with the few that were observed in the limit of a very good decision procedure.

A different behavior of the cost a as a function of the number of new shortcuts appears in the case of a mediocre decision procedure. For values of r and p such that

$$|\ln(1-r)|\ln n > [\ln(n-1)p]^2 \qquad (5)$$

Eq. 3 produces a curve where the number of steps towards the goal *increases* as new shortcuts are discovered. This is illustrated in Fig. 4, which was computed for the same graph with 10,000 nodes and a value of $r = 0.2$

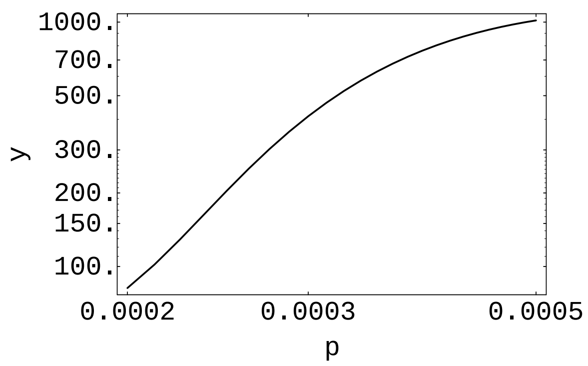

Fig. 4. log-log plot of the time to reach the goal as a function of new shortcuts when the decision procedure is mediocre.

This phenomenon, reminiscent of the production data for the Lockeed Tri-Star L-1011 [18], is due to a dynamical effect whereby the effectiveness of the decision procedure cannot compensate for the new possible ways of getting lost implied by the addition of new links or routines to the procedure space.

Finally, in the case where the number of shortcuts or procedures remains fixed and the only learning taking place is due to improvements in the decision mechanisms, organizational learning still exhibits a power law curve, but the dynamics tends to be much slower than before. In pointing out this result one should be reminded of the fact that it is seldom the case that one encounters a situation like this in real manufacturing.

4. Discussion

The existence of a learning curve for organizations has been thoroughly documented and studied since 1936, but the underlying mechanisms that cause it to be a power law are still not clearly understood. This paper presented a theory of organizational learning that explicitly deals with the introduction of shortcuts in the production process, as well as the effectiveness of the decision procedure involved in the search for better and shorter ways to reach the goal. By considering a general procedure space in which nodes represent stages in manufacturing and links correspond to the routines that connect them, learning can be simply described by the addition of shortcuts in the form of new links, and by improvements in the decision procedure when choosing among many possible ways of going from one stage to the other.

This approach yields a power law of learning with an associated learning rate that depends on the effectiveness of the decision procedure. It also produces anomalies that have been reported in the literature, and which are still lacking a general explanation. These anomalies, characterized by increasing costs with cumulative output, are shown here to arise out of the introduction of new shortcuts while not improving the decision procedures fast enough to compensate for this.

At a normative level, this suggests that as new routines and shortcuts are introduced into an organization, care must be exerted so that the associated learning of new procedures compensates for the increase in novel ways of doing work. This might entail a phase of organizational search for new routines or documentation that will ensure that the firm moves away from an effectively mediocre decision procedure into a more optimal one.

As to the empirical evidence supporting this theory, there is documented evidence of how routines evolve in laboratory settings [21, 17]. In particular, the experiments of Agre and Shrager showed how the introduction of shortcuts into the production of multiple copies of a document led to a power law of the time it took to complete the task as a function of units produced. At the organizational level, it would be interesting to see if the introduction of new routines as advocated in business reengineering practices [22] will support the predictions of this paper in a quantitative manner.

One issue that remains unsolved is the connection between individual and organizational learning. While obviously individuals and their learning[2] have an effect on the efficiency of organizational procedures, other variables such as technology and communications affect the firm in ways that go beyond individual learning. While this theory operates at an aggregate level that does not require such a connection, it still leaves open the task of building the important bridge between those two levels.

[2] Since individuals also display a power law of practice, the mathematical mechanism underlying learning at both the singular and aggregate level can be similar in structure, for it was shown earlier that cognitive chunking and good heuristics lead to power laws of learning. [2, 3].

References

[1] John S. Brown and Paul Duguid. Organizational learning and communities and communities of practice. In *Organizational Learning*, pages 58–82. Sage Publications, 1996.

[2] J. R. Anderson. *The Architecture of Cognition*. Harvard University Press, 1983.

[3] J. Shrager, T. Hogg, and B. A. Huberman. A graph-dynamic model of the power law of practice and the problem-solving fan-effect. *Science*, 242(4877):414–416, 1988.

[4] T. P. Wright. Factors affecting the cost of airplanes. *Journal of Aeronautical Science*, 3:122–128, 1936.

[5] W. B. Hirschmann. Profit from the learning curve. *Harvard Business Review*, 42:125–139, 1964.

[6] J. M. Dutton and A. Thomas. Treating progress functions as a managerial opportunity. *Academy of Management Review*, 9:235–247, 1984.

[7] L. E. Yelle. The learning curve: Historical review and comprehensive survey. *Decision Sciences*, 10:302–328, 1979.

[8] P. L. Joskow and G. A. Rozanksi. The effects of learning by doing on nuclear plant operating reliability. *Review of Economics and Statistics*, 61:161–168, 1979.

[9] Michael L. Rothschild. *Bionomics: economy as ecosystem*. Henry Holt and Company, 1990.

[10] Linda Argote and Dennis Epple. Learning curves in manufacturing. *Science*, 247(4945):920–924, Feb 23 1990.

[11] Ruhmini D. Rao Linda Argote, Dennis Epple and Kenneth Murphy. Learning and turnover in a manufacturing organization. *INFORMS*, 1995.

[12] Tad Hogg Natalie S. Glance and Bernardo A. Huberman. Training and turnover in the evolution of organizations. *Organization Science*, 1996.

[13] Peter C. Roberts. A theory of the learning process. *Journal of Operations Research*, 34:71–79, 1983.

[14] John F. Muth. Search theory and the manufacturing progress function. *Management Science*, 32:948–962, 1986.

[15] Barbara Levitt and James G. March. Organizational learning. *Annual Reviews of Sociology*, 14:319–340, 1988.

[16] R. M Cyert and J. G. March. *A Behavioral Theory of the Firm*. Prentice Hall, 1963.

[17] Michael D. Cohen and Paul Bacdayan. Organizational routines are stored as procedural memory: Evidence from a laboratory survey. In *Organizational Learning*, pages 58–82. Sage Publications, 1996.

[18] Tri-star production costs. *Aviation Week and Space Technology*, page 32, October 1979.

[19] B. Bollobas. *Random Graphs*. Academic Press, 1982.

[20] N. G. van Kampen. *Stochastic Processes in Physics and Chemistry*. North-Holland, Amsterdam, 1981.

[21] P. E. Agre and J. Shrager. Routine evolution as the microgenetic basis of skill acquisition. In *Proceedings of the Annual Conference of the Cognitive Science Society*, pages 694–701. Lawrence Erlbaum, 1990.

[22] Michael Hammer and James Champy. *Reengineering the Corporation*. HarperBusiness, 1993.