Implementation as Learning: An Extension of Learning Curve Theory

J. Bradley Morrison
Brandeis University International Business School

ABSTRACT

Firms often attempt to imitate successful practices of other firms. When implementing new practices, individuals in organizations learn new ways of doing things, develop new skills, and adopt new organizational routines. In the paper, we view implementation as a learning process and apply learning curve theory to the understanding of implementation dynamics. We extend classic learning curve to include a required output level for an individual who must choose between an old and a new way to achieve the output. Doing work the new way builds experience, increasing productivity and thus favoring continued use of the new skill, but this reinforcing process works to favor the new skill only at relatively high levels of productivity. Otherwise, the same process is a vicious cycle, driving out the new skill. We use a system dynamics model to demonstrate a mode of behavior in which learning begins and then stalls and another mode in which the new skill becomes the preferred one. We identify the tipping point between these two modes and characterize the transition problem: Learning by doing is a dynamic process, a transition from use of an old way to a new way that requires accumulating experience beyond a threshold.

Implementation as Learning: An Extension of Learning Curve Theory

J. Bradley Morrison
Brandeis University International Business School

Implementation is a learning process. During the implementation of an organizational innovation, people learn new ways of doing things. What is learned ranges from new skills in data entry with an information technology system, improved methods for machining a manufactured part, and novel routines for problem solving in process improvement programs to evolving norms for group functioning, updated guidelines for execution of a change in strategic direction, and nuances of interrelating with newfound colleagues from process or organizational redesign. In many instances of implementation, the degree to which people learn these new ways is a critical factor in determining the effect on organizational capability and performance. In this paper, we view implementation as learning and draw on learning curve theory to develop key insights into the dynamic nature of learning by doing.

The central notion in learning curve theory is that accumulating experience leads to improved performance, or "learning by doing." The concept occupies a central role in many strands of strategy and organization theory and forms the basis for such ideas as the specialization of labor, organizational learning, knowledge transfer, and core competences of the firm(Argote, 1996, 1999). Despite the apparent importance of learning as a key factor in the success of implementation, there are surprisingly few applications of learning curve theory to studying the challenges of implementation. We are unaware of any published work that directly applies learning curve theory to rigorously analyze the dynamic patterns of behavior observed in implementation processes. One reason for this gap in the literature may be that studies in the learning curve literature have generally viewed the productive activity of interest in isolation from other organizational activity, whereas implementation researchers are often concerned with richly contextualized settings in which implementation processes unfold. Another reason may be that if organizations always learn according to learning curves – e.g., monotonically increasing productivity as experience accumulates – many of the questions motivating the implementation literature – such as what are the reasons for implementation failure or why do some improvement initiatives fail while others succeed – are moot.

The purpose of this paper is to examine the dynamic character of learning by doing in the context of implementation. To apply learning curve theory in this context, we extend the theory in two critical ways. First, we model a context in which the learner must achieve a targeted rate of production while learning. Second, we incorporate a choice between a current way of working and a new way, both of which are means to accomplish the desired production. In our analysis of the feedback structure that characterizes this learning by doing under constraints, we find a mode of behavior in which learning begins and then stalls and another mode in which learning dominates so that the new skill becomes the preferred manner of doing. The paper contributes to the learning curve literature by applying learning curve theory to understanding possible reasons for the failure to learn. In particular, we show how learning by doing is prone to bifurcation dynamics, in which the distinction between reversion to an old way of doing and conversion to a new preferred way is whether the learner accumulates enough experience to cross a critical threshold, or tipping point.

The paper is organized as follows. The next section presents some highlights of the learning curve literature as background for the reader. The following section begins by representing the learning curve in a system dynamics model that captures the basic notion that accumulating experience leads to improved productivity and includes the forgetting or deterioration of knowledge. Then, we depart from traditional theory by formulating a model of learning under constraints. The model explicitly incorporates a constraint that requires the learner to achieve a specified level of output. Model analysis begins in the next section with equilibrium analysis and the construction of a rate-level plot. The next section uses simulation analysis to demonstrate a mode of behavior in which learning begins and then stalls and another mode in which learning dominates so that the new skill becomes the preferred manner of doing. We then turn to mathematical analysis to characterize the tipping point that distinguishes these two modes and present some sensitivity analyses. Finally, the discussion highlights key implications for managing implementation efforts and organizational change.

Background

The seminal work of Wright in 1936 first established an empirical relationship between cumulative production and the quantity of input needed to achieve that production (Wright, 1936). Wright observed that as the quantity of units an organization has manufactured doubles, the number of direct labor hours needed to produce the next individual unit decreases at a uniform rate. This phenomenon has been called the learning curve, the experience curve the progress curve, and learning by doing. Learning curves have been demonstrated in a wide range of settings at the individual, group and organizational levels. Several excellent reviews of the literature can be found in Yelle (1979) and Argote (1999).

Learning curves have become important in the practice of management (Dutton & Thomas, 1984). For example, the anticipation of future cost reduction that accrues as production experience is gained suggests setting prices aggressively, even below cost of manufacturing, early in a product life cycle in order to build market share. In addition to a large empirical literature, some authors have proposed theories of micro level activity that examine the underlying learning processes (Adler & Clark, 1991; Argote, 1999; Zangwill & Kantor, 1998).

The Basic Learning Curve Model

In this section we begin by developing a model of the learning curve based entirely on the current literature to serve as our departure point. We will then build on the basic model in the next section. The complete model is a set of equations that forms a first-order non-linear system of differential equations that we use for simulation analysis and then solve analytically.

The core notion in learning curve theory is the increase in productivity (or decrease in direct labor hours required for a unit of production) observed as experience accumulates. The theory is based on an empirical regularity rather than any explicit causal mechanism to explain learning. Our model will consider the learning of a new skill by repeated use over time. We represent the accumulation of experience as the integral of time spent using the new skill, as shown in Figure 1.

Figure 1: A Model of the Learning Curve with Forgetting

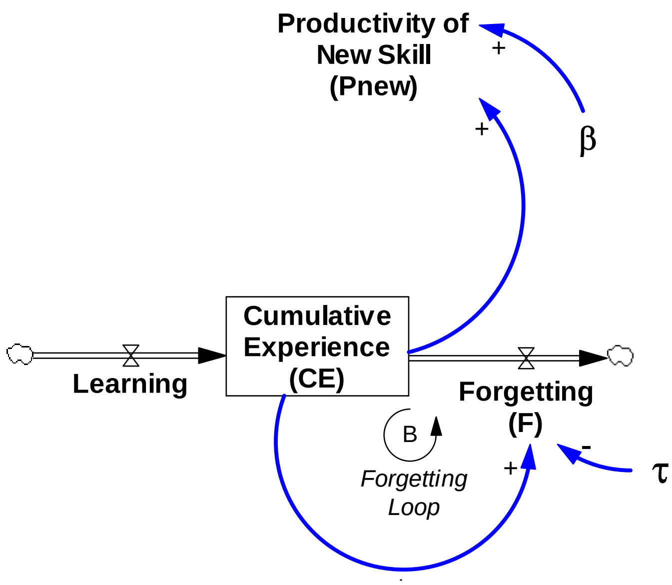

Cumulative Experience is the accumulation of Learning less Forgetting. Learning is based on the time spent using the new skill and is represented as a separate variable here only for convenience in the diagram. Experience accumulates as hours of time spent with the new skill. The stock is increased by learning and decreased by forgetting. Forgetting is modeled as a loss from the stock of experience at a constant fractional rate, as given by the Time to Forget (τ).

Representing the phenomenon of forgetting, or the depreciation of knowledge, as directly proportional to the stock of experience is typical of other models of forgetting found in the learning curve literature (Anderson & Parker, 2002; Epple, Argote, & Devadas, 1991). Thus, the mathematical representation of the accumulation and loss of experience is:

$$CE(t) = \int_{o}^{t} (L_t - F_t)dt + CE_0$$

$$L(t) = T_{new}(t)$$
$$F(t) = CE(t) / \tau$$

where

$CE(t)$ = Cumulative Experience with New Skill (hours)
$L(t)$ = Learning (hours/week)
$F(t)$ = Forgetting (hours/week)
$T_{new}(t)$ = Time Spent with New Skill (hours/week)
τ = Time to Forget (weeks)

According to learning curve theory, the accumulation of experience increases productivity, or alternatively reduces costs. The model captures this notion in the link from the stock of Cumulative Experience to the Productivity with New Skill. The Productivity with the New Skill increases with the learning effect based on accumulated experience. Here it is assumed that productivity of the new skill will be greater than zero and less than (99% of) that of the old skill during the learning phase of interest here. The Productivity with New skill as a function of Cumulative Experience is modeled as a straight line characterized by a parameter for the initial relative productivity (IRP) compared to the productivity of the old skill and one for the strength of the learning effect (β). By normalizing Cumulative Experience with its initial value, we have a line with a slope equal to β/CE_0 and an intercept equal to $(IRP_{new0} - \beta)$. Various functional forms of the learning curve appear in the literature, with the power form and the exponential form most common (Yelle, 1979). The linear formulation used here may be interpreted as a first-order approximation of a small portion of a learning curve formulated as a power function or an exponential function.

$$P_{new}(t) = P_{old} * \{IRP_{new0} + \beta * (CE(t) / CE_0 - 1)\}$$

$$0 \leq P_{new}(t) \leq 0.99 * P_{old}$$

where

P_{new} = Productivity with New Skill (tasks/hour)
P_{old} = Productivity with Old Skill (tasks/hour)
IRP_{new0} = Initial Relative Productivity with New Skill (dimensionless)
β = Strength of Learning Effect (dimensionless)
CE_0 = Initial Experience with New Skill (hours)

Thus far, the model has drawn on learning curve literature and formalizes three key notions. First, learning accumulates as experience or a stock of knowledge. Second, this accumulation increases productivity. Third, this accumulation depreciates over time. We now build on this basic formulation from the previous literature by explicitly incorporating a constraint on the learner's time and the need to achieve a given level of output.

A Model of Learning by Doing under Constraints

Traditional learning curve theory considers the productive activity of interest, such as the manufacture of airframes, in isolation from other demands for critical resource inputs, such as direct labor hours. In contrast, many learning situations are characterized by a competition for the learner's time between a new skill to be learned and an old, proven means of accomplishing tasks. The learner's time is a limited resource. The learner faces the challenge of allocating this resource to meet the demand for certain output objectives while simultaneously trying to learn how to do things a new and possibly better way.

The learner in our stylized model has two choices for how to do the work that will achieve the productive output of interest – an old way and a new way. As the learner uses the new method, the experience accumulated leads to increases in the productivity of this new method. The learner has mastered the old way, so productivity using the old, proven method is high. For simplicity, we also assume that productivity using the old skill is constant – that is forgetting of the old skill is too slow to warrant including in the model. At the outset, working in the new way requires considerably more time to accomplish the same quantity of output. (The present analysis assumes that both methods lead to similar or acceptable quality.) Since the learner wants to learn the new way, all else equal, he or she will choose the new method in order to gain experience. There are two complications. First, the learner must achieve a set rate of output in the given. Second, the learner has a fixed amount of working hours per week – so all else is not equal.

Figure 2 displays a model of learning by doing that shows the feedback structure characterizing the learner described above. The time spent with the new skill leads to learning that accumulates in a stock of experience. Accumulating experience increases the productivity with the new skill which in turn leads to more time spent with the new skill and thus more learning, forming a reinforcing feedback loop, the Learning by Doing Loop, labeled loop R. Accumulated experience with the new skill also atrophies, as shown in the Forgetting Loop, balancing loop B.

Figure 2: A Model of Learning by Doing under Constraints

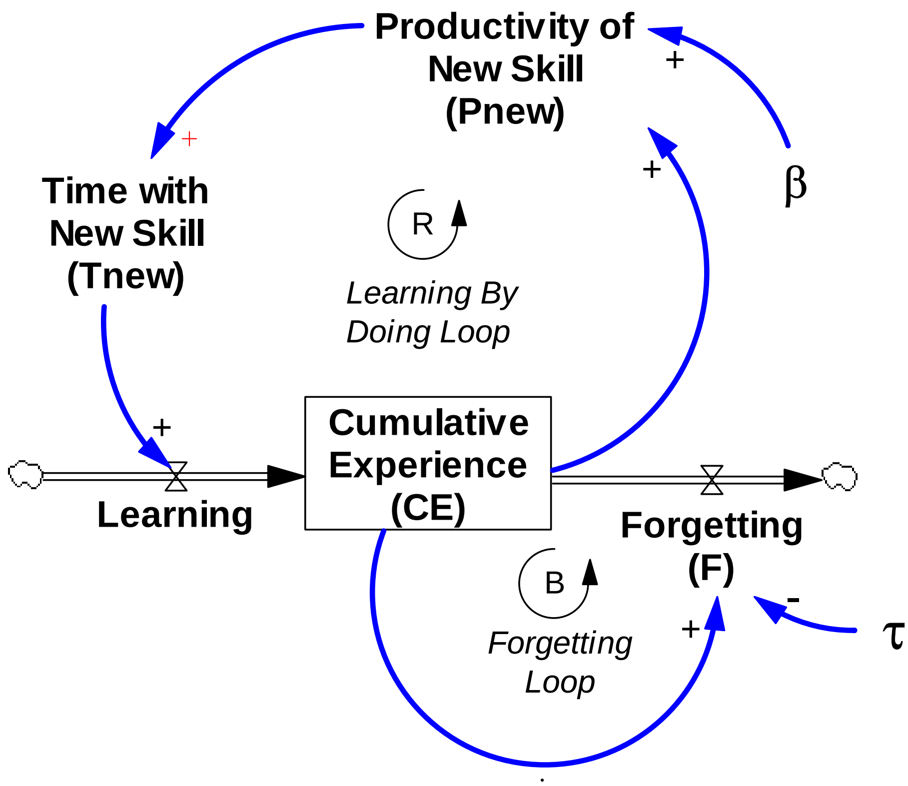

The Time Spent with the New Skill is based on the learner's resource allocation policy. The learner chooses the amount of time for learning the new skill consistent with the need to achieve a set output objective, given by the Required Completion Rate. The learner's allocation of time to the new skill must satisfy two equations:

$$T_{old}*P_{old} + T_{new}*P_{new} = Q*$$

$$T_{old} + T_{new} = T_{tot}$$

where

T_{tot} = Time Available (hours/week)
T_{old} = Time Spent with Old Skill (hours/week)
P_{old} = Productivity with Old Skill (widgets/hour)
$Q*$ = Required Completion Rate (widgets/week)

The first equation forces the total output to equal the Required Completion Rate. Total output is achieved as the sum of output from the old way and output from the new way. Output from the old way is the product of time spent with the old skill (T_{old}) and the productivity of that time (P_{old}). Similarly, output from the new way is the product of time spent with the new skill (T_{new}) and the productivity of that time (P_{old}). The second equation assures that time spent with the old skill and time spent with the new skill sum to the total time available. Solving these two equations for T_{new} yields the allocation policy for the learner's time. The Time Spent with the New Skill is constrained to be not less than zero and not more than the total time available. The allocation policy is thus:

$$T_{new} = \frac{T_{tot} * P_{old} - Q*}{(P_{old} - P_{new})}$$

$$0 \leq T_{new}(t) \leq T_{tot}$$

To fully specify the model, the following parameter values are used:

τ = 12 weeks
P_{old} = 1 widget/hour
$Q*$ = 30 widgets/week
$T_{tot} = (Q*/ P_{old})/(1+\mathrm{IF}*(IRP_{new0} -1))$(hours/week)
IRP_{new0} = 0.5
β = 0.25
$CE_0 = \tau * T_{tot} *IF$ (hours)
IF = Initial Fraction of time spent with new skill = 0.3 (dimensionless)

Model Analysis:

With the model now fully specified, we turn to model analysis. In this section, we do equilibrium analysis and construct a rate-level plot. In the following sections, we conduct simulation analyses, present an analytical solution for the model, and perform sensitivity analysis to understand the effect of parameters on the model behavior.

The model of learning by doing under constraints is a first-order non-linear system. The behavior of the system thus can be summarized by its one state variable, Cumulative Experience, and we can construct a rate-level plot that shows how the rates of flow in the model depend on the stock of Cumulative Experience. We use Figure 3 to derive the rate-level plot. We begin with the flow of Learning as a function of the level of Cumulative Experience. Even at very low levels of experience, learning is positive (because the learner's throughput goals are low enough that he can afford to spend some time learning even if productivity is very low). The plot for learning shows that as experience increases, learning increases. Learning increases at an increasing rate because there are two benefits to increasing experience. First, increasing experience increases the productivity of doing the work with the new skill to achieve the required throughput. Second, as productivity with the new skill increases, the opportunity cost of using the new skill decreases. Learning eventually reaches a maximum when all of the allowed time is allocated to using the new skill, as seen in the flat portion of the curve. Figure 3 also shows the flow of Forgetting. There is no Forgetting when there is no experience, so the curve starts at the origin. For any positive quantity of experience, forgetting occurs at a constant fractional rate. The curve is thus a straight line with a downward slope equal to the reciprocal of the time constant, Time to Forget.

Figure 3: Rate-Level Plot of Learning by Doing

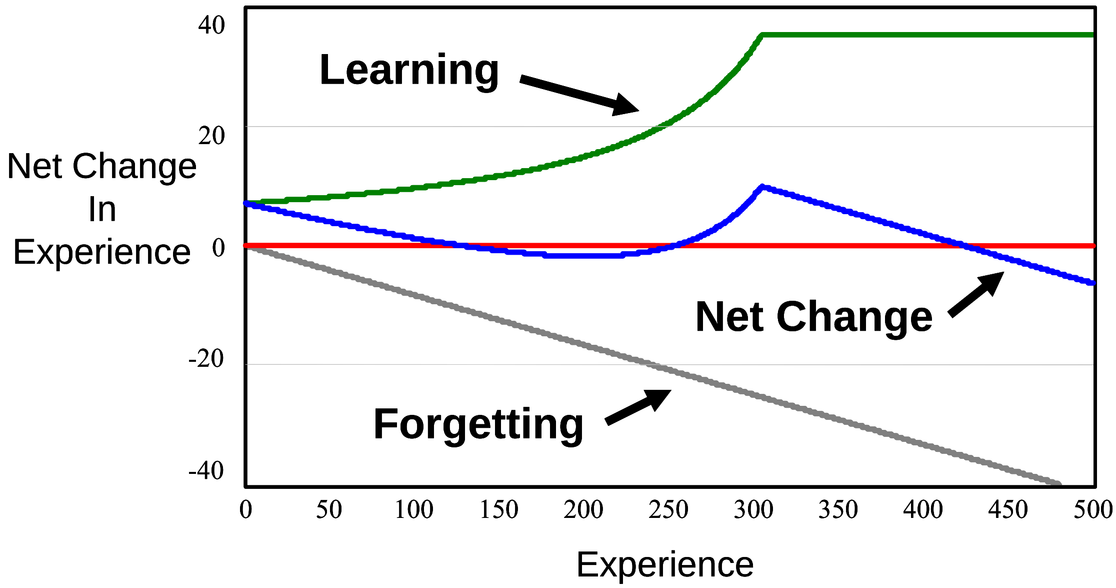

The net change in the level of experience as a function of the level of experience is the difference between the inflow Learning and the outflow Forgetting. As seen in Figure 3, the Net Change in Experience curve is the sum of the other two curves.

Figure 4 reproduces the Net Change in Experience curve from Figure 3 on an expanded scale. The Net Change curve crosses the zero line at three different points. Each of these points represents a level of experience at which the inflow from learning is exactly equal to the outflow from forgetting and are thus levels at which the stock of experience is in equilibrium. The arrows show the trajectory of the system in disequilibrium conditions, at all other levels of experience. The leftmost equilibrium in Figure 4 is formally denoted a stable equilibrium, meaning that small perturbations from it are compensated for by the system's dynamics. The balancing forgetting loop dominates behavior in this region, bringing stability to the system. If experience drops below the equilibrium level (a shift to the left in the figure), forgetting slows a bit more than learning does. The net change is positive, and the equilibrium level is restored. If experience increases a little (moving right in the figure), forgetting speeds up so net change is negative, again returning the level of experience to its original equilibrium value. This equilibrium occurs where the rate-level plot crosses the zero axis with a downward slope, thus characterizing a stable equilibrium. Similarly, the rightmost black dot in Figure 4 labels another stable equilibrium. (The downward sloping region in the right portion of the curve arises from the fixed constraint on the maximum amount of time available to learn. In this region, the learner is spending all allowable time using the new skill. If this constraint were removed, the curve would continue unbounded up and to the right.)

Figure 4: The Tipping Point on the Rate-Level Plot of Learning by Doing

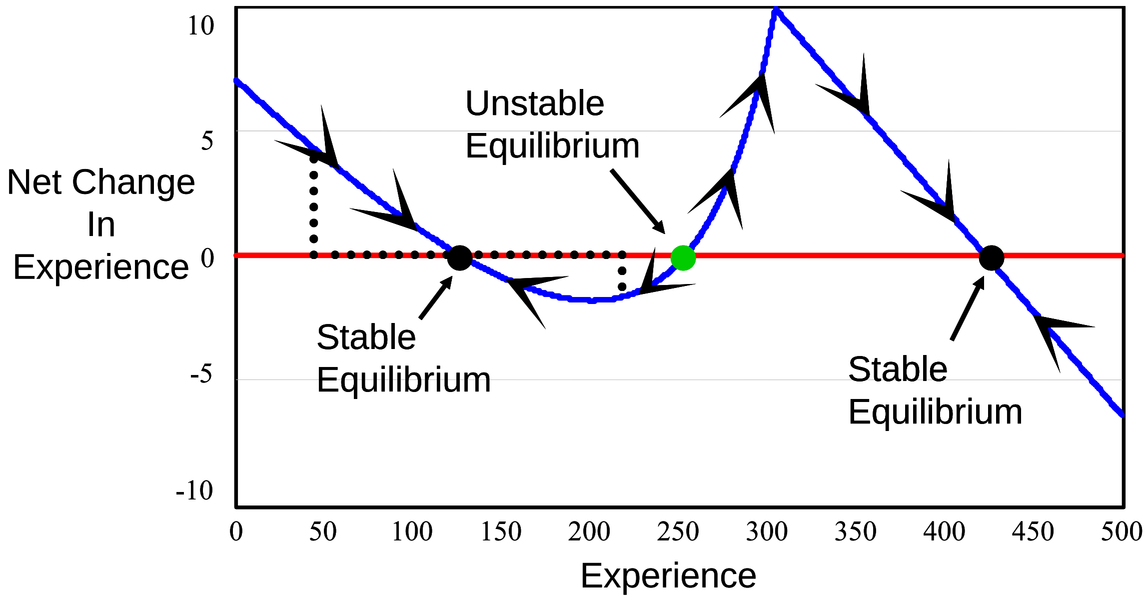

In contrast, the third equilibrium, designated with a green dot, is in a region where the rate-level plot crosses the zero axis in an upward sloping direction. The reinforcing loop dominates behavior in this region, so the equilibrium is unstable. As the arrows show, small perturbations away from the equilibrium are amplified, sending the system off towards one of the other two equilibria. The unstable equilibrium is critically important in understanding the dynamic behavior of the system. Consider a learner whose experience is at the level of the leftmost equilibrium, a stable equilibrium, and gains experience through learning by doing the new skill. From the rate-level plot, we can see that when learning increases experience to a level that is still to the left of the unstable equilibrium, the system will compensate and bring the learner back to the original equilibrium. Yet, if learning increases experience to a level that is just to the right of the unstable equilibrium, the reinforcing loop takes the system off towards the rightmost equilibrium. A small quantitative difference in the amount of learning from to the left of and to the right of the unstable equilibrium results in qualitatively different behavior. An unstable equilibrium is also known as a tipping point, and the rate-level plot explains why. At all levels of experience to the left of the tipping point, the system will be drawn to the leftmost stable equilibrium. Once experience accumulates enough to move just rightward of the tipping point, the system transitions onto a path towards the rightmost equilibrium.

The rate-level plot can now be used to restate the learner's goal. For the learner wishing to develop a sustained proficiency with new skill, the goal is to build enough experience to get just past the tipping point. Although stability is often a desirable characteristic of systems, it is the inherent instability in a critical region of this system's state space that creates the opportunity for enduring change. In the learner's case, reaching the unstable equilibrium is the key to the transition to a sustained level of higher proficiency with the new skill.

Model Behavior:

In this section, we use simulation analysis to explore model behavior and understand what it takes for the learner to develop a lasting proficiency with the new skill.

The simulations begin with a learner that has a small amount of experience with the new skill. The learner is exactly accomplishing the indicated completion rate and is learning at exactly the rate necessary to offset the forgetting that is occurring. That is, the learner is spending exactly the amount of time with the new skill needed to maintain a constant level of experience, and thus constant productivity. The first test introduces extra experience to the stock of Cumulative Experience (as might happen if a learner were endowed with some extra time to dedicate solely to learning – by doing – the new skill). The hope is that adding experience will set in motion the reinforcing Learning by Doing Loop (R in Figure 2). At a higher level of experience, the productivity of the new skill will be higher. If the learner is endowed with this higher productivity and no increase in the throughput objective, he will have slack capacity. He should allocate more time to working with the new skill, thus stimulating learning, which in turn will build experience and increase productivity, leading to further slack capacity and thus still more time allocated to working with the new skill. This virtuous cycle of learning will fuel the increase in experience needed to master the new skill.

Figure 5 shows the results of a test in which a pulse of extra experience (120 hours) is added in

Figure 3: Simulating the Learner's Response

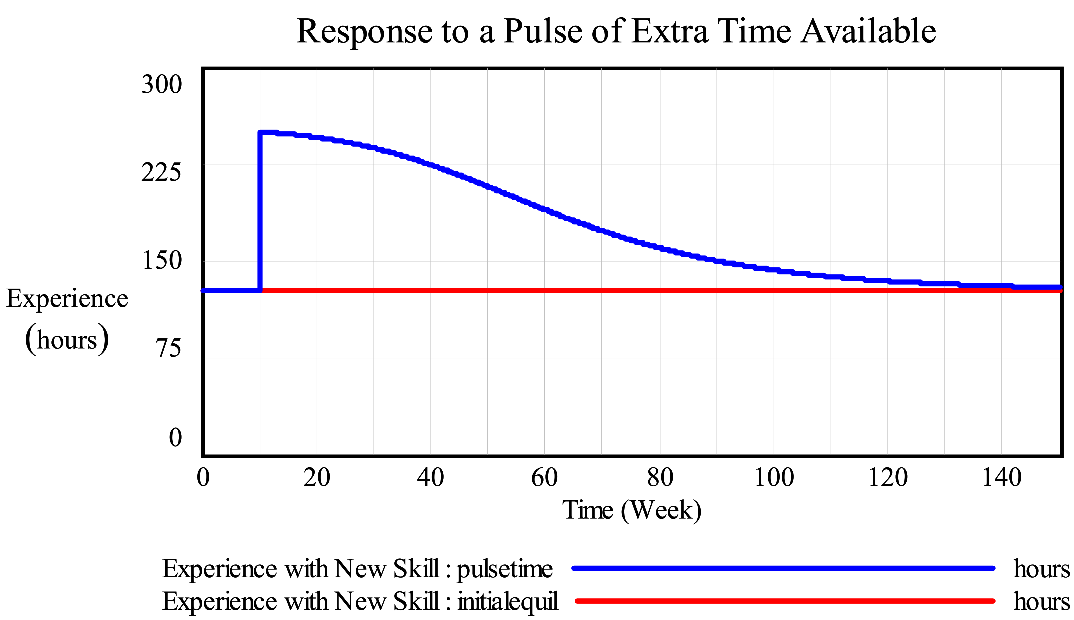

week 10. The amount of experience (and thus also productivity) increases immediately as expected. However, after this initial increase, experience degrades slowly from the peak achieved at the time of the pulse back to the original levels. The addition of extra time had led to an improvement, but this improvement is only temporary.

As seen in Figure 5, cumulative experience with the new skill begins to decline from the original peak in week 10. By closely examining the simulation output (and looking also at a tabular form of the output) we can see that cumulative experience declines at an increasing rate from week 10 until approximately week 46. The increasing rate of decline signals the dominance of a reinforcing loop. During this period, the reinforcing loop R from Figure 2 is dominating the behavior, working as a vicious cycle against the goal of increasing learning. Recall that the rate-level plot (see Figure 4) is upward sloping in the region just to the left of the unstable equilibrium. Then, from week 46 onwards, experience continues to decline toward the original level but now at a decreasing rate of decline. During this period, the Forgetting Loop, balancing loop B in Figure 2, is dominating the behavior, guiding the system back to its original equilibrium conditions.

Figure 6 shows the results from introducing pulses of various sizes to the stock of experience. This set of simulations shows that pulses of various sizes lead to two different terminals values for the level of experience. First, for some pulse sizes such as the lowest four simulation runs in Figure 6, experience temporarily improves but then slowly decreases over time back to the original level, a pattern identical to the test done in Figure 5. (The blue line in Figure 6 results from the same simulation conditions as for the blue line in Figure 5.) Second, for some higher pulse sizes such as the highest five shown in Figure 6, a second outcome is reached in which the learner accumulates experience to higher levels, eventually reaching a level which is maintained thereafter. In these latter simulation runs, the system has passed a critical threshold and entered into a regime in which the new skill is sustained at a permanently higher level. The threshold distinguishing the two behavior patterns is the tipping point we identified in the rate-level plot.

With any greater amount of experience, the feedback structure brings the system to a new, higher equilibrium level of experience. With any smaller amount of experience, the system returns to its original level of experience.

Figure 4: Response to Various Pulses of Experience

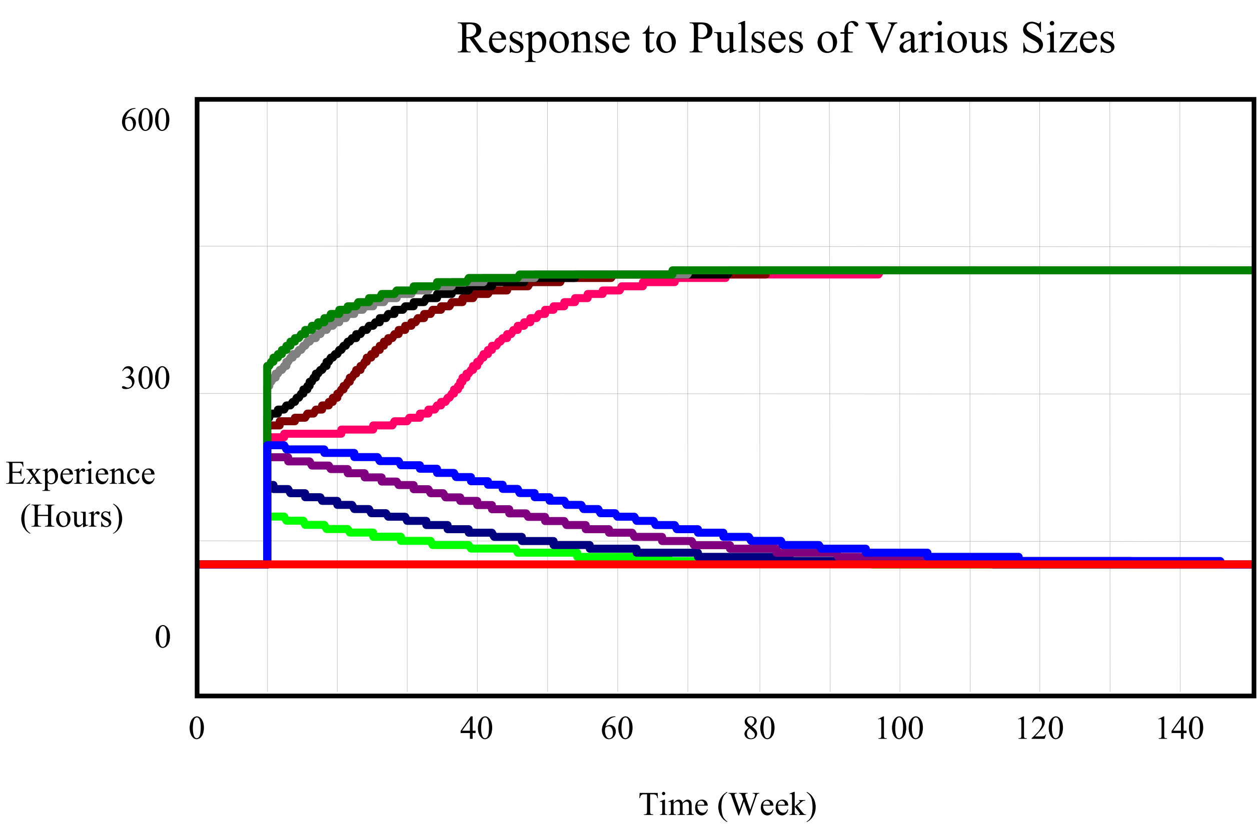

The bifurcation shown in Figure 6 has important implications. In the simple model we use here, the two steady-state equilibria to which the system can be attracted are fundamentally different in character. The lower of the two equilibria is the same equilibrium at which the simulation starts, and the dynamic patterns that lead to this endpoint are examples of fleeting improvement. Performance improves, or perhaps more precisely adherence to the use of a new way of doing increases, but the increases are only temporary. Despite an intervention that did indeed increase learning, over time performance reverts to the original levels. There is no lasting effect resulting from what may have at first appeared to be a successful intervention. This pattern is reminiscent of a common failure mode in implementation settings in which the use of a new approach starts out successfully but eventually fades away. On the other hand, the higher of the two equilibria represents a change in preferences such that the new skill has become locked in. The learner has reached a higher degree of proficiency and consequently the much lower opportunity cost of using the new skill no longer works to squeeze out its use. This pattern is characteristic of a successful transformation. Moreover, although the two equilibria appear dramatically different with regard to the mastery of the new work practices, the stimuli that led to them differ by only a tiny amount – and even more importantly, that tiny difference took place during the intervention in week 10.

Mathematical Analysis

How much experience is enough to cross the tipping point? In the stylized model presented here where all parameter values are known, an analytical solution is possible. The tipping point is an equilibrium, so it must satisfy the condition that the inflows to the stock of Cumulative Experience equal the outflows. Setting the flow equation for learning equal to the flow equation for forgetting, substituting to remove the auxiliary variables, and solving for Cumulative Experience yields a quadratic equation. The two roots of this equation are the leftmost stable equilibrium in Figure 4 and the unstable equilibrium. The first is the equilibrium for the initial conditions of the simulations, and the second is the tipping point. The tipping point occurs when:

$$E^{tippt} = E_0 * \left\{ \frac{1 - IRP}{\beta} \right\}$$

The equation for the tipping point is expressed as the product of the Cumulative Experience at the initial time and a multiplier. The multiplier (the quantity in brackets) is the quotient of two terms. Recall that IRP is the Initial Relative Productivity of using the new skill and ranges from 0 to 1. The numerator, 1 – IRP, is the proportional opportunity cost (in terms of productivity loss) that the learner incurs from spending time using the new skill. The denominator is the strength of the learning curve effect. Thus, we see that the tipping point is close to the initial experience level when IRP is large, which is in situations where there is only a small difference between the productivity of the old skill and the productivity of the new skill. In such cases, the opportunity cost of using the new skill is low, and the tipping point should be relatively easier to reach. Similarly, we see that the tipping point is close to the initial experience level when the strength of the learning effect is large. Small gains in experience will lead to relatively large gains in productivity, so again the tipping point should be easier to reach. For the parameter values used in the simulations shown here, the multiplier is equal to two. Thus, to reach the tipping point under these parameters, the learner must accumulate enough time spent with the new skill to double his experience.

The equation for the tipping point is useful to develop an analytical understanding of the system's properties, but its usefulness in practice will be limited by the degree to which the parameter values are known or can be estimated with reasonable accuracy. Sensitivity analyses that examining the effect of changes in various parameters can help build understanding of system performance. For example, Figure 7 displays the rate-level plots obtained using five different values for the Strength of Learning Effect (β). With no learning (Strength = 0, as in the green line), forgetting dominates, the rate-level curve is always downward sloping, and there is no tipping point. For the strong learning curve (Strength = 0.5) shown in the brown line, any increase from the initial equilibrium sends the system towards the higher equilibrium point. For moderate learning effects as in the other lines, the tipping point moves to the left as the Strength of Learning Effect increases. Efforts to increase the strength of this learning effect will bring the tipping point closer to the initial level of experience.

Figure 7: Sensitivity Analysis: Rate-Level Plots by Strength of Learning Effect

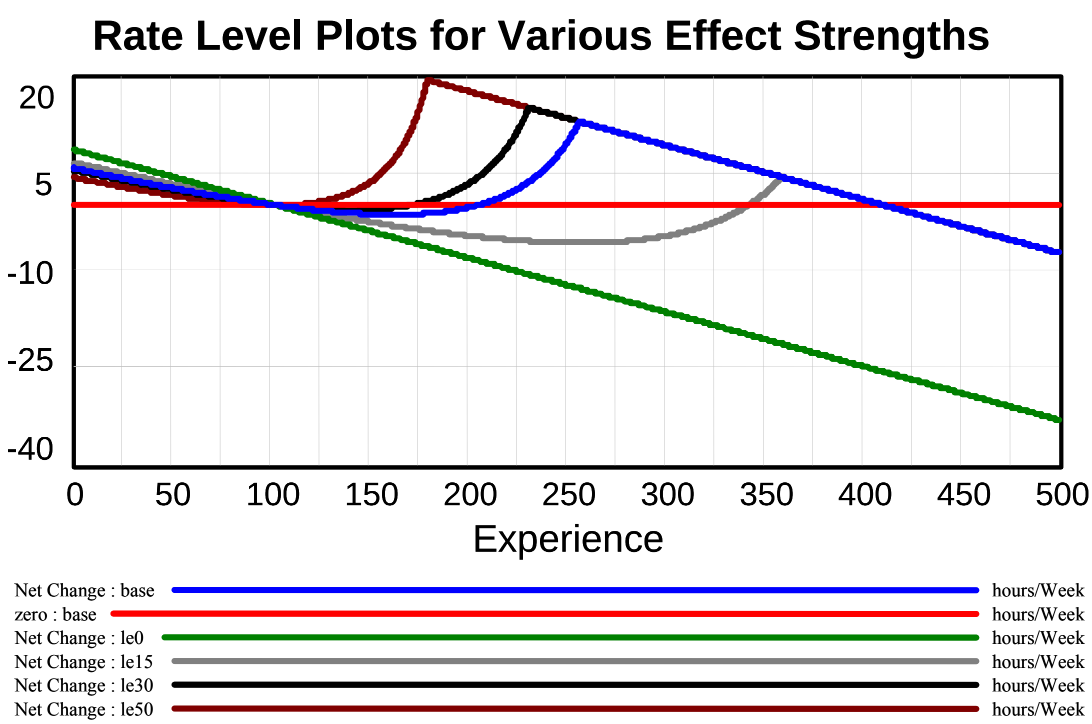

Discussion:

The paper has adopted a view of implementation as learning and drawn on learning curve theory to ask not just about learning but also about the failure to learn. To apply learning curve theory in the context of implementation, we developed a system dynamics model based on extensions of learning curve theory that incorporates a learner's need to accomplish ongoing work while also meeting the challenge of learning new skills. Formal analysis of the feedback structure showed that learning by doing generates tipping dynamics. Simulations demonstrated that some attempts to learn will be short-lived, while others will move the learner into a regime of sustained higher proficiency. By characterizing the tipping point that distinguishes the two modes of behavior, the preceding analysis has added to our understanding of the dynamics of change.

There are important implications of this new contextualized view of the learning curve. One implication is that extending learning curve theory by incorporating a throughput constraint sheds light on possible causes for failure to learn. Learning does not take place in a vacuum. The learner's time is often a constrained resource. More attention to constraints on time and competition among means to achieve performance objectives may prove fruitful in the study of implementation failure. Empirical study is also needed of both the dynamics of forgetting and the character of learning curves at low levels of experience.

A second implication for scholars relates to the shape of a learning curve under constraints. In traditional learning curve analyses, productivity increases with cumulative experience, but the rate of learning (the change in experience) decreases. With the addition of throughput constraints, we see that in some regions, learning as a function of experience increases at an increasing rate. This is because as experience increases, productivity with the new skill increases, and thus the opportunity cost of the new skill decreases, encouraging more use of the new skill and consequently a higher rate of learning.

A third implication is that transitions to new ways of doing must be understood as fundamentally dynamic phenomena. Similar feedback structure underlies many approaches to building organizational capabilities. Scholars may benefit from a reconceptualization of "managing change" as "moving past the tipping point." The tipping point occurs when the reinforcing (positive) loop is dominant, and more attention to the dynamics of reinforcing loops during organizational transitions will likely yield more insights.

Managers of implementation efforts where learning is necessary for success can apply the lessons of this analysis. A key challenge is to understand what conditions will allow the learner to transition to and sustain the higher level of experience with the new skill. The learner will do so when the level of experience becomes great enough so that the reinforcing Learning by Doing Loop is working in a favorable direction and dominating the Forgetting Loop. Policy analysis should search for means of bringing the tipping point within reach. Policies might aim to strengthen the positive loop, increase learning from experience, reduce forgetting, improve retention, recognize and moderate the pressure to achieve output, and monitor how work is done not just output rates. To enhance the success of innovation implementation, managers need to manage transitions, recognizing not just the need for new skills like the masters, but for skill building. The masters are already past the tipping point. Studying the masters – e.g., benchmarking and best practices – overlooks the journey and underestimates the resource needs.

A more effective approach to managing performance in systems with such tipping potential may be to identify symptoms of system behavior that signal being near or past the tipping point. The nature of unstable equilibria is such that systems rarely, if ever, operate at or even very close to them. Learning about how to effectively manage performance will thus be quite challenging. Learners may come quite close to a tipping point and never realize how close they were to "getting over the hump." More insidiously, in the region of experience just below the tipping point, the system behavior is dominated by a reinforcing loop acting as a vicious cycle to keep the learner away from the tipping point. Operating in this region is likely to feel much like swimming upstream. The learner who perseveres to get past the critical threshold will then gain the benefit of a reinforcing loop acting in his favor, as though the current were almost pulling him a long.

References

Adler, P. S. & Clark, K. B. 1991. Behind the Learning Curve: A Sketch of the Learning Process. Management Science, 37(3): 267-281.

Anderson, E. G. & Parker, G. 2002. The Effect of Learning on the Make/Buy Decision. Production and Operations Management, 11(3): 313-339.

Argote, L. 1996. Organizational Learning Curves: Persistence, Transfer and Turnover. IJTM, Special Publication on Unlearning and Learning, 11(7/8): 759-769.

Argote, L. 1999. Organizational Learning: Creating, Retaining, and Transferring Knowledge. Boston: Kluwer Academic Publishers.

Dutton, J. M. & Thomas, A. 1984. Treating Progress Functions as a Managerial Opportunity. Academy of Management Review, 9(2): 235-247.

Epple, D., Argote, L., & Devadas, R. 1991. Organizational Learning Curves: A Method for Investigating Intra-plant Transfer of Knowledge Acquired through Learning by Doing. Organization Science, 2(1): 58-70.

Wright, T. P. 1936. Factors Affecting the Cost of Airplanes. Journal of Aeronautical Science, 3: 122-128.

Yelle, L. E. 1979. The Learning Curve: Historical Review and Comprehensive Survey. Decision Sciences, 10: 302-328.

Zangwill, W. I. & Kantor, P. B. 1998. Toward a Theory of Continuous Improvement and the Learning Curve. Management Science, 44(7): 910-920.