Dynamics of Learning by Doing under Constraints: Analysis of the Tipping Point

(Extended Abstract)

J. Bradley Morrison
MIT Sloan School of Management

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. 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 purpose of this paper is to extend learning curve theory by embedding the learner in a context in which he or she must achieve certain output objectives. Building on the basic notion that accumulating experience leads to improved productivity, I formulate and analyze a two-loop system dynamics model that incorporates a constraint to achieve a specified level of output as well as the forgetting or deterioration of knowledge. I use 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. I characterize the tipping point that distinguishes these two modes. The discussion highlights key implications for managing implementation efforts and organizational change.

A Model of Learning by Doing under Constraints

The learner 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. However, the learner has mastered the old way, so productivity using the old, proven method is high. 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 – so all else is not equal. Second, experience accumulated using the new way has a limited useful life. If the new skill is not used, the experience will atrophy.

Figure 1 shows a stock and flow diagram of a model of learning by doing. The time spent with the new skill leads to learning that accumulates in a stock of experience. The stock is increased by a flow of learning and decreased by forgetting. Accumulating experience increases the relative productivity with the new skill which in turn leads to more time spent with the new skill and thus more learning, forming the Learning by Doing Loop, reinforcing loop R. Accumulated experience with the new skill also atrophies, as shown in the Forgetting Loop, balancing loop B.

Figure 1: Feedback Structure of Learning by Doing

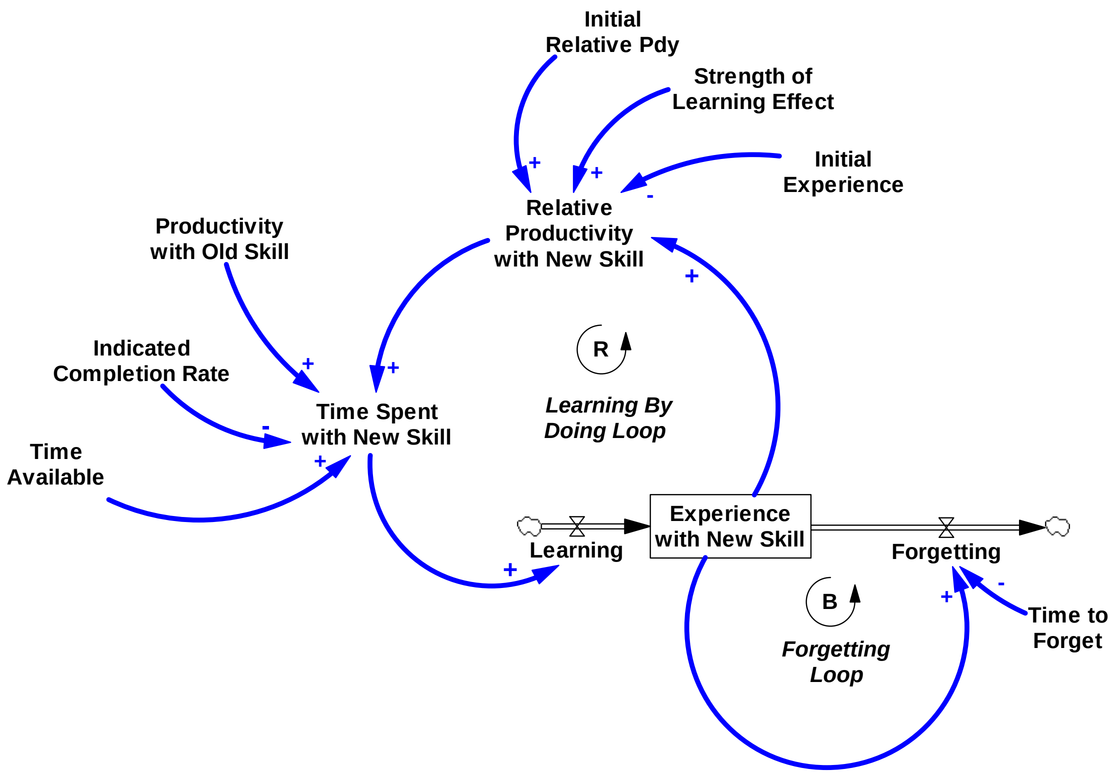

The complete model structure is described here.

$$Exp_t = \int_o^t (L_t - F_t)dt + Exp_0$$

$$L_t = TSNS_t$$

$$F_t = Exp_t / TF$$

where

Exp = Experience with New Skill (hours)
L = Learning (hours/week)
F = Forgetting (hours/week)
TSNS = Time Spent with New Skill (hours/week)
TF = Time to Forget (weeks)

Experience with New Skill is the accumulation of Learning less Forgetting. Learning is based on the Time Spent with 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. Forgetting is modeled as a loss from the stock of experience at a constant fractional rate, as given by the Time to Forget.

$$TSNS_t = max(0, min(TA, ITSNS_t))$$

$$ITSNS_t = \frac{TA * PdyOld - ICR}{PdyOld * (1 - RPdyNew_t)}$$

where

ITSNS = Indicated Time Spent with New Skill (hours/week)
TA = Time Available (hours/week)
PdyOld = Productivity with Old Skill (widgets/hour)
ICR = Indicated Completion Rate (widgets/week)
RPdyNew = Relative Productivity with New skill (dimensionless)

The Time Spent with the New Skill is the Indicated Time Spent with the New Skill, constrained to be not less than zero or more than the total time available. The Indicated Time Spent with the New Skill is based on the learner's resource allocation policy. The learner chooses to use all available time for learning the new skill consistent with the need to achieve a set output objective, given by the Indicated Completion Rate. The learner's allocation of time to the new skill must satisfy two equations:

$$TSOS*PdyOld + ITSNS*PdyOld*RPdyNew = ICR$$

$$TSOS + TSNS = TA$$

The first equation forces the total output to equal the Indicated 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 (TSOS) and the productivity of that time (PdyOld). Similarly, output from the new way is the product of time spent with the new skill (ITSNS) and the productivity of that time (PdyOld*RPdyNew). 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 ITSNS yields the allocation policy for the learner's time, as shown above.

$$RPdyNew_t = max(0, min(0.99, IRPdyNew_t))$$

$$IRPdyNew_t = RelPdy_0 + SLE*(Exp_t / Exp_0 - 1)$$

where

IRPdyNew = Indicated Relative Productivity with New skill (dimensionless)
$RelPdy_0$ = Initial Relative Productivity with new skill (dimensionless)
SLE = Strength of Learning Effect (dimensionless)
Exp_0 = Initial Experience with New Skill (hours)

The Relative Productivity with the New Skill is the ratio of the productivity with the new skill to the productivity with the old skill and 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 Indicated Relative Productivity with New skill is a straight line with a slope equal to SLE/Exp_0 and an intercept equal to ($RelPdy_0$ – SLE). This linear formulation 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.

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

TF = 12 weeks
PdyOld = 1 widget/hour
ICR = 30 widgets/week
TA = (ICR/PdyOld)/(1+IFNS*(IRPdy-1))(hours/week)
$RelPdy_0$ = 0.5
SLE = 0.25
$Exp_0 = TF * TA * IF$ (hours)
IF = Initial Fraction of time spent with new skill = 0.3 (dimensionless)

Model Behavior:

In this section, I use simulation analysis to explore 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. To give the learner an opportunity to learn more, the first test introduces extra time available. The hope is that the additional time available will set in motion the reinforcing Learning by Doing Loop (R in Figure 1). If the learner is endowed with additional time available with 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 2 shows the results of a test in which a pulse of extra time (120 hours) is added in week 10. The amount of experience (and thus also productivity) increases immediately

Figure 2: Simulating the Learner's Response

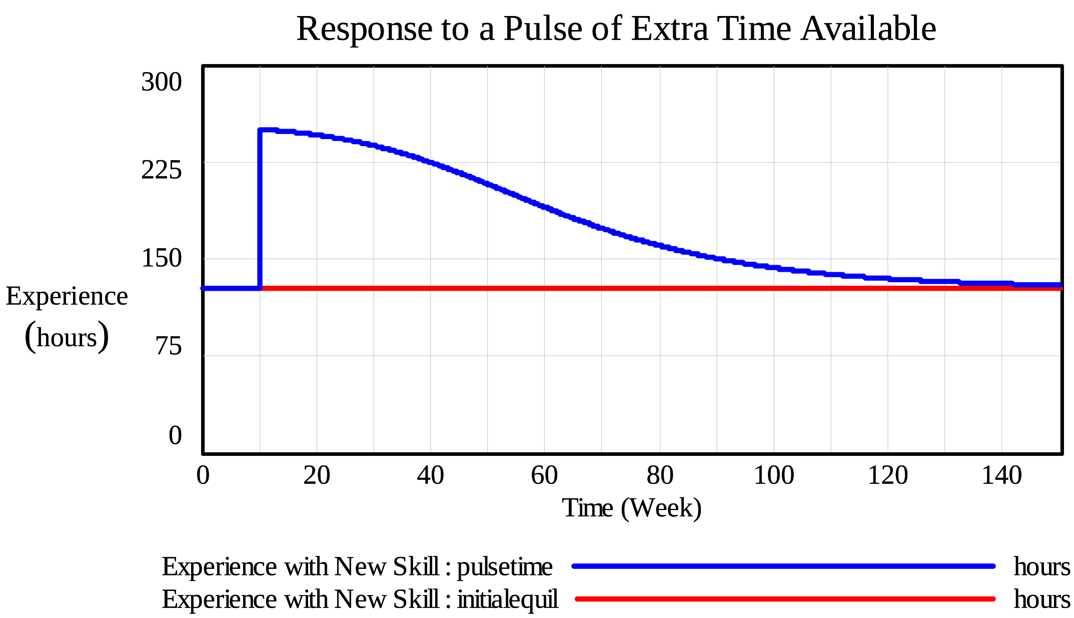

as expected. However, after this initial increase, experience degrades slowly from the peak achieved at the time of the pulse back to the original levels. The addition of extra time had led to an improvement, but this improvement is only temporary.
As seen in Figure 2, experience with the new skill begins a decline from the original peak at an increasing rate from week 10 until approximately week 46. During this period, the reinforcing loop R is dominating the behavior, working as a vicious cycle against the goal of increasing learning. 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 1, is dominating the behavior, guiding the system back to its original equilibrium conditions.

Figure 3 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 three different terminals values for the level of experience. First, for some pulse sizes such as the lowest four simulation runs in Figure 3, experience temporarily improves but then slowly decreases over time back to the original level, a pattern identical to the test done in Figure 2. Second, for one unique pulse size (the blue line in Figure 3), the pulse immediately sets experience to a new level that stays constant thereafter. Third, for some higher pulse sizes such as the highest five shown in Figure 3, a third outcome is reached in which the learner accumulates experience to higher levels, eventually reaching a level which is maintained thereafter. In these 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. That threshold is known as a tipping point. The one unique pulse (in the blue line) brings the stock of experience just exactly to the tipping point. With any greater amount of experience, the feedback structure brings the system to a new, higher

Figure 3: Response to Various Pulses of Experience

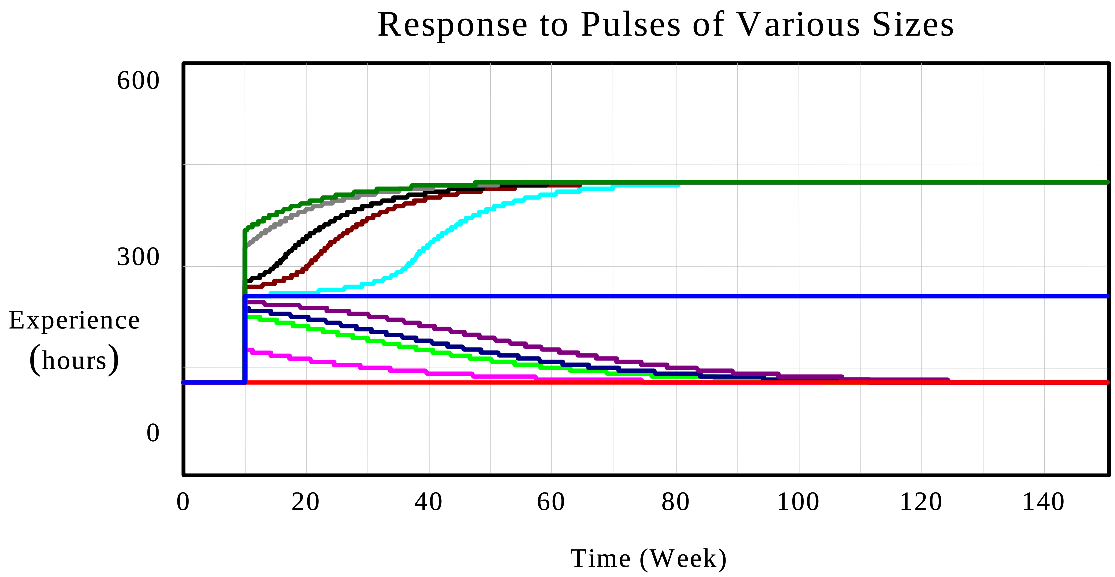

equilibrium level of experience. With any smaller amount of experience, the system returns to its original level of experience.

Model Analysis:

Under what conditions will the learner 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.

Figure 4 shows a rate-level plot constructed by examining how the rates of flow in the model depend on the stock of Experience with the New Skill. The plot shows the flow of Learning as a function of the level of 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 4 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

Figure 4: Rate-Level Plot of Learning by Doing

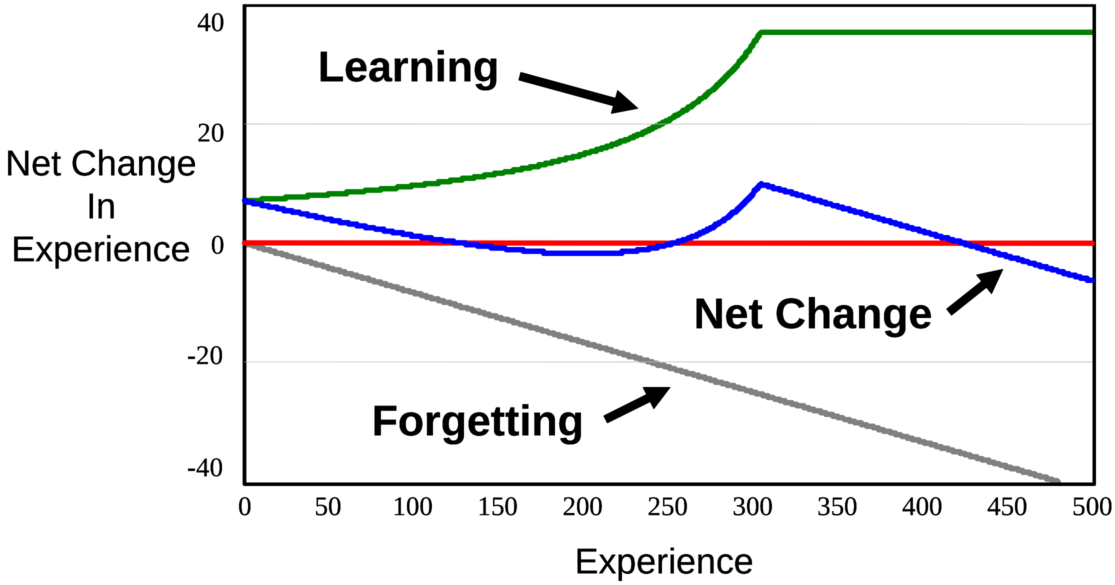

straight line with a downward slope equal to the reciprocal of the time constant, Time to Forget.

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 4, the Net Change in Experience curve is the sum of the other two curves.

Figure 5 reproduces the Net Change in Experience curve from Figure 4 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 5 is the starting point for the simulations. This equilibrium 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 so 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 5 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

Figure 5: The Tipping Point on the Rate-Level Plot of Learning by Doing

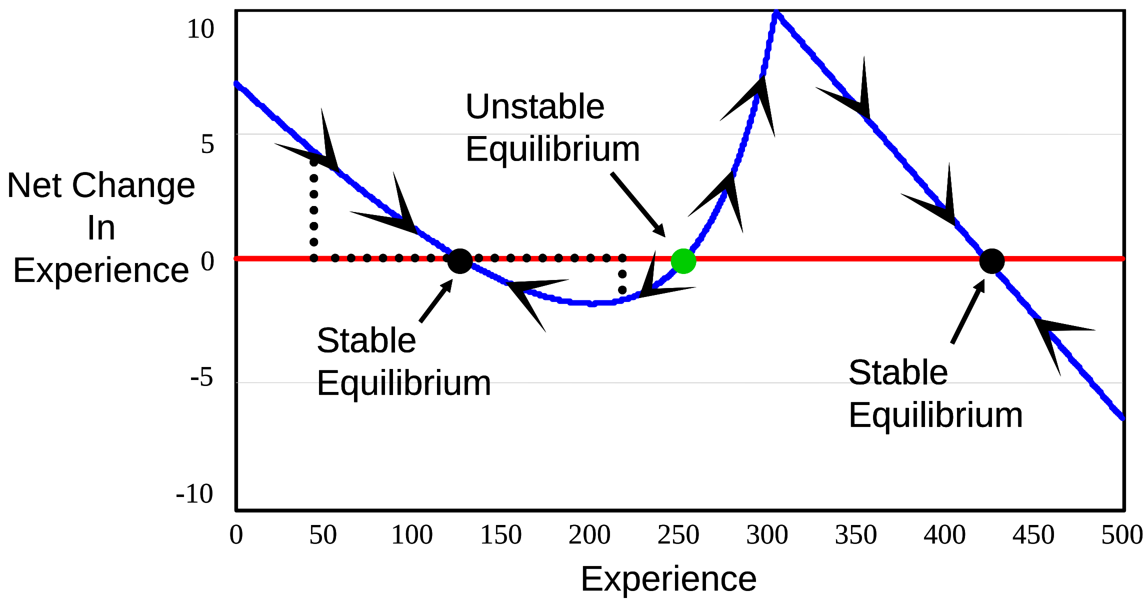

to the right.)

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 point here is also the 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.

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 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 Experience with New Skill yields a quadratic equation. The two roots of this equation are the leftmost stable equilibrium in Figure 5 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:

$$Exp^{tippt} = Exp_0 * \left\{ \frac{-1 + \sqrt{1 - \frac{4 * \left(TA - \frac{ICR}{PdyOld} \right)}{SLE * TA * IF}}}{2} \right\}$$

For the parameter values used in the simulations shown here, the quantity in brackets is equal to two. Thus, to reach the tipping point, 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 6 displays the rate-level plots obtained using five different values for the Strength of Learning Effect (the slope of the effect of experience on productivity). 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 6: Sensitivity Analysis: Rate-Level Plots by Strength of Learning Effect

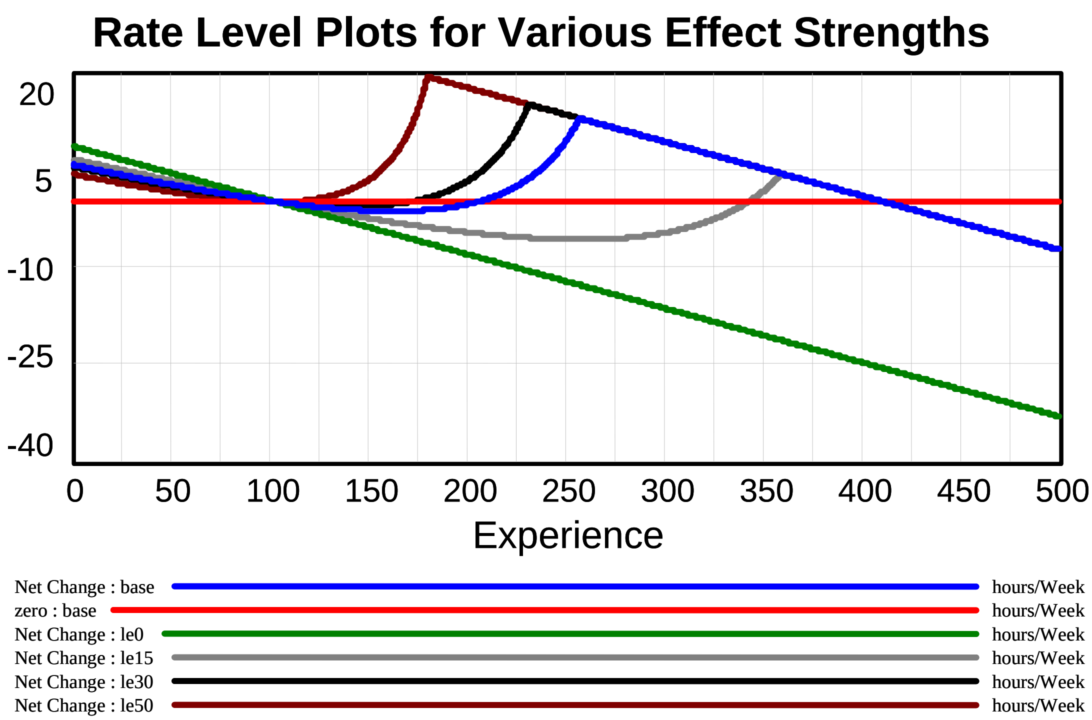

Discussion:

The paper has 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. 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 these two modes of behavior, the preceding analysis has added to our understanding of the dynamics of change.

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.

Back to the Top