2 Inside the Learning Curve: Opening the Black Box of the Learning Curve

Michael A. Lapré

CONTENTS

INTRODUCTION

The learning curve phenomenon is well known. As organizations gain operating experience, organizational performance improves, although at a decreasing rate. Scholars have frequently used the power curve to model this relationship in manufacturing contexts. In these models, the logarithm of unit cost decreases linearly as a function of the logarithm of cumulative number of units produced (Yelle 1979). The decrease in cost (i.e., improvement) is attributed to organizational learning, hence the name "learning curve." Scholars have extended the power curve by incorporating forgetting (recent experience matters more than older experience) and learning from others (transfer of experience). For an overview of forgetting and learning from others, see Argote (1999). This chapter focuses on learning from own experience.

The disappointing implication of the typical use of the power curve is that management can only accelerate learning from own experience by producing more. There are several limitations to this traditional view of the learning curve. First, the

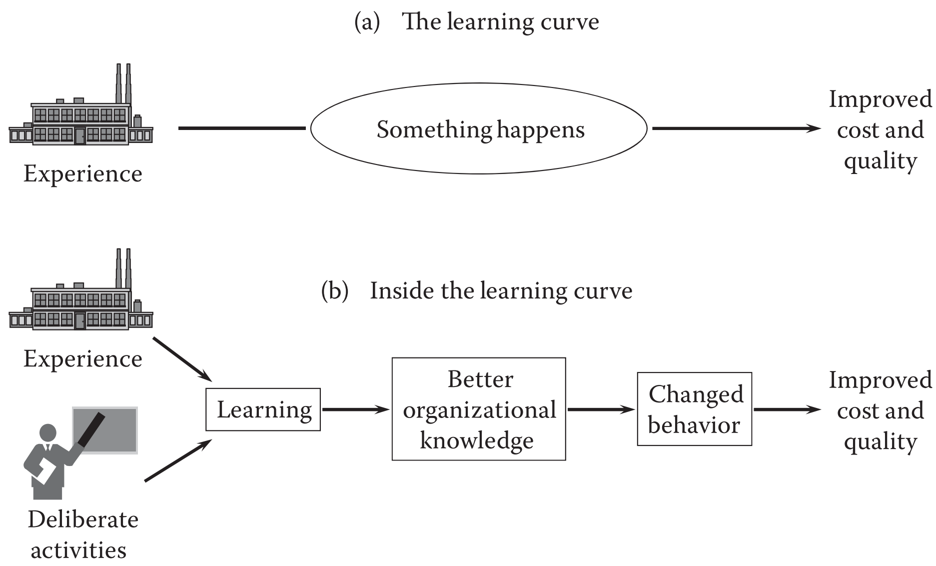

FIGURE 2.1 Two views of the learning curve. (Adapted from Bohn, R.E., *Sloan Management Review* 36(1), 1994.)

rate of improvement—the learning rate—is typically treated as some exogenously given constant. However, there is ample evidence that learning rates vary widely across industries, within industries across organizations, and within organizations across organizational units (Dutton and Thomas 1984; Hayes and Clark 1985; Lapré and Van Wassenhove 2001). Hence, the learning rate should be treated as an endogenous variable. In other words, management is actually responsible for managing the rate of improvement. Second, experience—typically measured by cumulative production volume—is not the only source for learning. Organizations can engage in deliberate learning activities such as quality improvement projects. Third, and most importantly, the traditional view treats the learning in the learning curve as some "black box" (see Figure 2.1a). Yet, there is an actual learning process inside the learning curve. Learning results from experience and deliberate activities. It can yield better organizational knowledge, and better organizational knowledge can persuade organizational members to modify behavior. Changed behavior, in turn, can improve organizational performance (Bohn 1994) (see Figure 2.1b). None of these steps are trivial. Scholars have merely scratched the surface in terms of studying these steps. No single empirical study has incorporated all of the steps.

This chapter reviews empirical findings in the literature in terms of (i) different sources for learning, and (ii) partial assessments of the steps that make up the actual learning process inside the learning curve. The chapter concludes by identifying opportunities for future research that should provide insights for organizations to better manage their learning curves.

EXPERIENCE AS A SOURCE FOR LEARNING

As far as learning from own experience goes, scholars have started to investigate three interesting themes: First, what constitutes own experience and how should it be measured? Is cumulative volume the most relevant proxy, or would a variant of cumulative volume be a more accurate measure of experience? Second, to what extent can organizations gain a competitive learning curve advantage from specialization? Third, what factors contribute to the rather significant variation in learning rates from own experience?

Nature of Experience

The learning curve literature has focused on three common experience variables: (1) cumulative volume, (2) calendar time, and (3) maximum volume. As mentioned before, cumulative production volume is the typical experience variable (Yelle 1979; Argote 1999). Repetition allows organizations to gain experience and to fine tune operations. Cumulative volume captures the notion of "learning by doing." Some scholars have used calendar time elapsed since the start of operation (Levin 2000; Field and Sinha 2005). When time for reflection is more important for learning, calendar time captures the notion of "learning by thinking." Mishina (1999) proposed the third experience variable: maximum output produced to date, or maximum proven capacity to date. When a plant is scaling up production, the production system faces significant challenges and new situations. Factory personnel need to figure out how to solve such challenges during scale-up. Mishina's experience variable captures the notion of "learning by new experiences" or "learning by stretching" (Mishina 1999; Lapré et al. 2000). Only two studies have compared learning curve estimations with all three measures of experience. Interestingly, both studies concluded that maximum volume was the best measure. Mishina (1999) found that the learning curve estimation for bomber airplanes suffered from autocorrelation with cumulative volume and calendar time, but not with maximum volume. In tire-cord manufacturing, Lapré et al. (2000) found that maximum volume explained the learning curve much better than cumulative volume and calendar time. Future research should assess whether these findings generalize beyond airplane and tire-cord manufacturing.

Not all experience is necessarily equally effective in improving organizational performance. First, organizations could learn significantly from their own failures (Cannon and Edmondson 2005). Whenever a production unit produces defective units, such defects provide opportunities to learn from and improve the production system. Li and Rajagopalan (1997) found that the cumulative number of defective units is statistically more significant than the cumulative number of good units in explaining learning curve effects. For manufacturing contexts that require defective units to be reworked, Jaber and Guiffrida (2004) propose a model that incorporates rework time. Depending on the evolution of rework time, a learning curve may continue to improve, plateau, or deteriorate. Future learning curve research is needed to empirically investigate the role of defects and rework.

Second, experience can accumulate at the individual, team, and organizational levels. Recently, scholars have investigated the impact of team experience. In addition to organizational experience measured with the usual cumulative volume variable, organizations with stable teams could potentially reach higher performance levels. In stable teams, team members learn how to better coordinate work with one another, because team members learn (1) who is best at performing which role, and (2) to trust one another. Scholars have found that team experience is a significant driver for learning curves in health care (Reagans et al. 2005) and software development (Huckman et al. 2009).

Sinclair et al. (2000) have questioned the role that experience plays in achieving better organizational performance. Their study suggests that cumulative volume provides an indication of future volume. Future expected volume conditioned future expected returns from research and development (R&D) and, by extension, the choice of R&D projects. Research and development projects—not cumulative volume—were the real source of cost reduction. Future research should continue to investigate the question of "under what conditions does a certain type of experience trigger actual learning?"

Experience and Specialization

In 1974, Skinner introduced the notion of a "focused factory." Focused factories that specialize in executing fewer tasks outperform factories that perform a wider set of tasks. Task homogeneity, coupled with a higher frequency of repetition, allows a factory to learn more quickly from its experience. Several learning curve studies have investigated the benefits of specialization. In the U.S. hotel industry, Ingram and Baum (1997) found that hotel chains operating in a limited geographic region (geographic specialists) benefited more from their own experience than hotel chains operating nationwide (geographic generalists). In a study of incidents and accidents in the U.S. airline industry, Haunschild and Sullivan (2002) found that specialist airlines benefited more from analyzing heterogeneous causes than generalist airlines. Heterogeneous causes allow for deeper analysis. The authors concluded that focus helped specialist airlines to analyze heterogeneous causes. Interestingly, however, it might not be optimal to focus as much as possible. An experimental study showed that some degree of variation yields faster learning rates. Schilling et al. (2003) found that learning rates for related experience are greater than for specialized or unrelated experience. An analysis of an offshore software services operation confirmed the potential benefits of related variation. According to Narayanan et al. (2006), exposure to a greater variety of tasks improves long-term productivity. However, investments in learning new tasks can impede short-term productivity. A learning curve study on customer dissatisfaction compared specialist and generalist airlines (Lapré and Tsikriktsis 2006). Average specialist airlines did not learn faster than average generalist airlines. However, the best specialist airline did learn faster than the best generalist airline. So, focus provides an opportunity for faster learning, but there are no guarantees for superior performance. A promising area for future work would be to investigate under what conditions does a particular level of specialization result in faster learning from a certain type of experience?

Variation in Learning Rates

It has been well documented that organizations show tremendous variation in learning rates. Dutton and Thomas (1984), for example, graphed a distribution of learning rates from a sample of over 100 studies. Dutton and Thomas (2004, 238) expressed the learning rate as a progress ratio: "When cumulative volume doubles, the cost per unit declines to p% of original cost." "P" is called the progress ratio. Progress ratios ranged from 55% to 108%. Understanding the dynamics that cause learning curve heterogeneity has been an important area for learning curve research. In a study of adoption of minimally invasive cardiac surgery in sixteen hospitals, Pisano et al. (2001) found significant learning curve heterogeneity. The authors used case data from two hospitals to explore differences that might have contributed to variation in learning rates. The two hospitals differed markedly in terms of: (i) their use of formal procedures for new technology adoption, (ii) cross-functional communication, (iii) team and process stability, (iv) team debrief activities, and (v) surgeon coaching behavior. In a follow-up study, Edmondson et al. (2003) found that learning curve heterogeneity is greater for aspects of performance that rely on tacit knowledge (as opposed to codified knowledge). Wiersma (2007) investigated how four conditions affected learning rates across twenty-seven regions in the Royal Dutch Mail Company. A higher degree of temporary employees, a higher level of free capacity, a higher degree of product heterogeneity, and less conflicting concerns about other performance measures all had a favorable impact on the learning rate. It will be worthwhile for future research to further quantify conditions that vary across organizations and include such quantitative data in learning curve analyses.

DELIBERATE ACTIVITIES AS A SOURCE FOR LEARNING

Organizations do not have to limit themselves to learning from experience. They can also engage in a more pro-active approach to managing learning curves. Levy (1965) introduced the distinction between "autonomous learning" from experience and "induced learning" from deliberate activities designed to improve production processes. Examples of deliberate activities include both pre-production planning before a process starts, as well as industrial engineering after a process starts. Levy found that prior experience and training explain differences in the estimated learning rates for individual workers. This was a landmark study even though the explanatory variables prior experience and training did not evolve over time. Adler and Clark (1991) made the next step by incorporating longitudinal variables for deliberate activities in productivity learning curves: cumulative engineering activity and cumulative training activity. In one production department, engineering activity enhanced productivity while training activity disrupted productivity. In a second production department, the exact opposite occurred. Thus, deliberate learning activities can both help and hurt. The authors provided some case-based explanations for these surprising findings. For example, if producibility concerns trigger engineering activity, engineering activity enhances productivity. On the other hand, if product performance concerns trigger engineering changes, such changes could be disruptive. Hatch and Mowery (1998) studied the impact of cumulative engineering in yield learning curves in semiconductor manufacturing. Yield learning curves for processes in the early stages of manufacturing were driven by cumulative engineering as opposed to cumulative volume. In more mature processes, cumulative engineering and cumulative volume were both sources for learning to improve yields.

However, the introduction of new processes disrupted the ongoing learning activities of existing processes. Hatch and Dyer (2004) further examined yield learning curves and showed that, in addition to cumulative engineering, human capital variables such as screening tests and statistical process control training significantly enhanced yield improvements. Human capital variables differed across processes, but were constant over time within a process.

Lapré et al. (2000) provide a systemic explanation for the seemingly unpredictable effect of deliberate learning activities observed by Adler and Clark (1991). In a tire-cord manufacturing plant, the authors studied quality improvement projects as deliberate learning activities. They found that the cumulative number of quality improvement projects that generated both know-why and know-how accelerated waste reduction, whereas the cumulative number of quality improvement projects that generated know-why without know-how slowed down waste reduction. A production line run as a learning laboratory, called a "model line," consistently produced the learning-rate-enhancing mix of know-why and know-how. Replication of this model line concept in other plants in the same firm fell short of expectations (Lapré and Van Wassenhove 2001). These replications neglected conditions of management buy-in and knowledge diversity to solve interdepartmental problems.

Two studies provide further evidence that deliberate activities should be included in learning curves in addition to autonomous learning-by-doing variables. Ittner et al. (2001) showed that both cumulative quality engineering and cumulative design engineering significantly reduced defect rates. Quality engineering includes prevention activities such as quality planning, developing and maintaining the quality planning and control systems, quality improvement activities, and internal quality improvement facilitation and consulting. Design engineering covers product design engineering expenses incurred for prevention activities. Arthur and Huntley (2005) studied cost reduction ideas submitted by staff in an auto-parts manufacturing plant's gainsharing program. The authors found that the cumulative number of implemented employee suggestions (their measure of deliberate learning) significantly reduced costs. More research is needed to understand under what conditions does a certain type of deliberate activities enhance knowledge creation, adoption, and organizational performance.

STEPS INSIDE THE LEARNING CURVE

Very few studies have addressed the steps inside the learning curve as indicated in Figure 2.1. Only two studies have incorporated a step from Figure 2.1 with longitudinal variables in learning curve estimations (Lapré et al. 2000; Arthur and Huntley 2005). This section also reviews related research—cross-sectional studies without a link to longitudinal organizational performance.

Mukherjee et al. (1998) studied 62 quality-improvement projects undertaken in a tire-cord plant over a decade. A cross-sectional analysis of these projects showed significant variation on two dimensions of the learning process: conceptual learning and operational learning. Conceptual learning consists of using science and statistical experimentation to develop a deeper understanding of cause-and-effect relationships—in other words, the development of know-why. Operational learning consists of modifying action variables and obtaining follow-up of experiments—in other words, the development of know-how. Both conceptual and operational learning enhanced changed behavior measured by modifications in standard operating procedures and statistical process control rules. Lapré et al. (2000) used the dimensions of conceptual and operational learning to split a sample of quality improvement projects into four categories according to high or low conceptual learning and high or low operational learning. For each of the four categories, the authors constructed longitudinal variables capturing the cumulative number of projects completed to date. The four cumulative project variables were incorporated in a learning curve estimation for the factory's waste rate—the percentage of products that had to be scrapped because of irreparable defects. Only two cumulative project variables had a statistically significant impact on waste evolution. Projects with high conceptual learning and low operational learning were disruptive. These "non-validated theories" were often advanced by experts from central R&D, yet the insights obtained at central R&D were not developed sufficiently for a full-scale manufacturing environment. Projects with high conceptual learning and high operational learning, on the other hand, accelerated waste reduction. These "operationally validated theories" provided solutions that worked, backed by scientific principles explaining why these solutions worked. This study explicitly incorporates the "better organizational knowledge" step in Figure 2.1 in a longitudinal learning curve estimation. The research site lacked historical data on all modifications to standard operating procedures and statistical process control rules. If such data had been available, the "changed behavior" variable could also have been included. The Arthur and Huntley (2005) study mentioned in the previous section used a cumulative number of implemented employee suggestions, which does capture "changed behavior" in Figure 2.1.

None of the steps depicted in Figure 2.1 are self-evident. Tucker et al. (2002), for example, investigated problem-solving behavior by front-line workers. Faced with a problem, nurses typically engage in "first-order problem solving"—fixing a problem without doing anything to prevent a similar problem from occurring in the future. Rarely do nurses engage in "second-order problem solving"—conducting root-cause analysis to change underlying causes. In 92% of the 120 problems observed by the authors, nurses ignored possible root causes. Such a focus on first-order problem solving prevents actual learning. Despite the opportunity for root-cause analysis, no attempt has been made to create better organizational knowledge. Tucker et al. (2002) identified several factors contributing to a continued emphasis on first-order problem solving at the expense of learning. First, front-line personnel feel good about themselves by patching a problem, demonstrating independence and competence. Second, a high workload focuses front-line workers on completing pressing tasks now, rather than thinking about improvement for the future. Third, the lower status of nurses compared to physicians might prevent nurses "to intrude upon a physician's time." The authors conclude that it is necessary to create an organizational environment that is psychologically safe (Edmondson 1999). Tucker et al. (2007) investigated the importance of psychological safety (a supportive organizational context) in deliberate learning activities. The authors conducted a cross-sectional study of organizational learning in 23 neonatal intensive care units (NICUs). The authors studied learn-how, which concerns "understanding why a practice works, as well as how to carry it out" Tucker et al. (2007, 898). Like Mukherjee et al. (1998), the authors link "learning" to "changed behavior" in Figure 2.1. Learn-how in improvement projects enhanced project implementation success. Tucker et al. (2007) found that psychological safety was an antecedent of learn-how. Furthermore, the level of published evidence for a practice (a measure of knowledge) also enhanced implementation success. In a follow-up study, Nembhard et al. (2009) investigated the impact of learn-how on organizational performance. The authors used data on 1061 infant patients from the same 23 NICUs. Learn-how significantly reduced patient mortality rates (measured at the organizational level). Moreover, interdisciplinary collaboration was found to mediate the relationship between learn-how and organizational performance.

A survey of 188 six sigma projects in a manufacturing firm provides further cross-sectional evidence for Figure 2.1. Choo et al. (2007) found that "learning behaviors" enhanced "knowledge created," which in turn enhanced "project performance." Moreover, the authors found that "use of a structured method" was an antecedent of "learning behaviors," whereas "psychological safety" was an antecedent of "knowledge created."

AVENUES FOR FUTURE RESEARCH

So far, this chapter has identified the following avenues for future research:

- What are the conditions that determine which measure of experience best captures the learning curve phenomenon?
- When does what type of experience trigger actual learning?
- Quantify conditions that vary across organizations and include such quantitative data in learning curve analyses to explain learning curve heterogeneity.
- When does how much specialization result in faster learning from what type of experience?
- Under what conditions do which deliberate activities enhance knowledge creation, adoption, and organizational performance?

A major reason for learning curve heterogeneity is the nature of the existing knowledge base in organizations. As mentioned earlier, Edmondson et al. (2003) found more heterogeneity for aspects of performance that rely on tacit knowledge. Many organizations have incomplete knowledge of their production systems (Jaikumar and Bohn 1992). Examples of settings characterized by incomplete knowledge include kitchens in commercial food firms (Chew et al. 1990), semiconductor manufacturing (Hatch and Mowery 1998), tire-cord manufacturing (Lapré and Van Wassenhove 2001), cardiac surgery (Edmondson et al. 2003), and electromechanical plants (Field and Sinha 2005). In the face of incomplete knowledge, organizations rely more on art as opposed to science. As a result, organizations depend on workers to figure out how to control processes, create better knowledge, and improve performance. "Transforming operators into quasi-engineers requires investments in human capital but pays big dividends in learning performance" (Hatch and Dyer 2004, 1173) How can we assess whether an organization is making progress in moving from an art to a science? Jaikumar and Bohn (1992), and Bohn (1994, 1995, 2005) introduced the "stages of knowledge" to gage progress of knowledge creation.

Stages of Knowledge

Bohn (1994, 62) defined technological knowledge as "understanding the effects of the input variables on the output. Mathematically, the process output, Y, is an unknown function f of the inputs, x: $Y = f(x)$; x is always a vector (of indeterminate dimension)." Inputs include raw materials, control variables, and environmental variables. Jaikumar and Bohn (1992) and Bohn (1994) developed "stages of knowledge" detailing how much an organization knows about $Y = f(x)$. In 1995, Bohn refined the concept, recognizing that stages of knowledge need to be measured along two separate dimensions: causal knowledge and control knowledge. Table 2.1 depicts the stages of causal knowledge and control knowledge. (See also Bohn 2005.)

Causal knowledge assesses how much an organization knows about the relationship between an input x_i and output y. At stage 1 "ignorance," the organization is unaware that x_i might affect y. At stage 2 "awareness," the organization is aware that x_i and y are related, but the direction of causality is unknown. At stage 3 "direction," the organization knows that x_i affects y. At stage 4 "magnitude," the organization can quantify the impact of a small change in x_i on y. At stage 5 "scientific model," the organization has a functional specification with parameters describing the relationship between x_i and y. At stage 6 "interactions," the organization has extended stage 5 knowledge to include interactions with all other input variables (never obtained in practice).

Control knowledge, on the other hand, assesses an organization's ability to keep an input variable x_i at its desired target level. At stage 1 "ignorance," the organization is unaware of x_i. At stage 2 "awareness," the organization is aware of the existence of x_i. At stage 3 "measure," the organization is able to measure x_i routinely. At stage 4 "control of the mean," the organization can control x_i at the mean level, but there is significant variation in the level of x_i. At stage 5 "control of the variance," the

TABLE 2.1
Stages of Causal Knowledge and Control Knowledge

Causal Knowledge **Know How x_i affects y**	**Control Knowledge** **Know How to Control x_i**
1. Ignorance	1. Ignorance
2. Awareness	2. Awareness
3. Direction	3. Measure
4. Magnitude	4. Control of the mean
5. Scientific model	5. Control of the variance
6. Interactions	6. Reliability

Note: Author's notes taken during a presentation by Bohn (1995).

organization can control the variance of x_i. At stage 6 "reliability," the organization can always keep x_i at its target level (never obtained in practice).

The two dimensions of causal and control knowledge closely mirror the dimensions of the learning process identified by Mukherjee et al. (1998). Conceptual learning should allow an organization to climb the stages of causal knowledge; whereas operational learning should allow an organization to climb the stages of control knowledge. Lapré et al. (2000) demonstrated the significance of incorporating the two learning dimensions into a learning curve estimation. It would be a major contribution to include longitudinal progress on the stages of knowledge in a learning curve estimation. At what stages of causal knowledge and control knowledge can an organization expect to make more than merely incremental improvements? Do breakthrough improvements require balanced climbing of the stages knowledge; that is, should causal knowledge and control knowledge progress at the same pace? Primary variables are variables that directly impact output. Secondary variables are variables that directly impact primary variables. What is the impact of climbing the stages of knowledge for primary variables versus secondary variables? One challenge in addressing such questions is in finding a research site where progress along the stages of knowledge can be captured. However, some organizations are aware of progress along stages of knowledge. Ittner et al. (2001), for example, used a research site that measured four stages of quality-based learning: (1) "aware of need," (2) "process characterized and sources of variation identified," (3) "critical process parameters understood," and (4) "knowledge institutionalized." In a longitudinal field study of an electromechanical motor assembly plant, Field and Sinha (2005) found that actions to control the mean (control knowledge stage 4) do indeed precede actions to control the variance (control knowledge stage 5). Scholars have yet to quantitatively measure progress on the stages of knowledge and include such measures in learning curve estimations.

Learning to Improve Multiple Dimensions of Organizational Performance

Learning curve scholars have focused their attention on single dimensions of organizational performance. However, organizations might have to perform on more than one dimension. The operations strategy literature has advanced cost, quality, delivery, and flexibility as typical candidates for competitive priorities. Typically, not all competitive priorities are equally important at all times. For example, an entrepreneurial firm might successfully compete on quality, whereas a mature firm might find it necessary to have low cost. Hence, the importance of different competitive priorities might change over time (Corbett and Van Wassenhove 1993).

Initially, competitive priorities were thought of as fundamental trade-offs. Higher quality implies higher cost, while cost reductions imply worse quality. Ferdows and De Meyer (1990) challenged this inherent trade-off view. Their sand-cone model proposes that capabilities are built cumulatively: first invest in improving quality, then delivery and flexibility, ending with cost reduction. Lapré and Scudder (2004) found empirical evidence for the sand-cone model in the U.S. airline industry. Airlines that ended up in a sustainable superior quality-cost position made larger initial improvements in quality compared with cost, although trade-offs do occur when operating close to asset frontiers. Gino et al. (2006) propose that learning to improve one dimension may come at the expense of learning on another dimension. In a sample of 16 hospitals, the authors found evidence of a learning trade-off between efficiency and application innovation.

Future learning curve research should address learning along multiple performance measures. Do different types of experiences enhance different dimensions of performance? How do organizations learn to improve internal performance (such as cost) and external performance (such as customer loyalty) at the same time? What is the relationship between internally oriented operating experience and externally oriented competitive experience?

In the past two decades, learning curve scholars have made important contributions to understanding what processes are "behind the learning curve." Much work, however, remains to be done. Hopefully, this chapter inspires others to further our understanding of organizational learning curves.

ACKNOWLEDGMENTS

Ingrid Nembhard and Bradley Staats provided very helpful comments. Support for this chapter was provided by the Dean's Fund for Faculty Research, Owen Graduate School of Management at Vanderbilt University.

REFERENCES

Adler, P.S., and Clark, K.B., 1991. Behind the learning curve: A sketch of the learning process. *Management Science* 37(3): 267–281.

Argote, L., 1999. *Organizational learning: Creating, retaining and transferring knowledge*. Norwell: Kluwer Academic.

Arthur, J.B., and Huntley, C.L., 2005. Ramping up the organizational learning curve: Assessing the impact of deliberate learning on organizational performance under gainsharing. *Academy of Management Journal* 48(6): 1159–1170.

Bohn, R.E., 1994. Measuring and managing technological knowledge. *Sloan Management Review* 36(1): 61–73.

Bohn, R.E., 1995. Measuring technological knowledge in the hard disk-drive industry. Presentation at the Fall meeting of the Institute of Management Sciences and Operations Research (INFORMS), New Orleans.

Bohn, R.E., 2005. From art to science in manufacturing: The evolution of technological knowledge. *Foundations and Trends in Technology, Information and Operations Management* 1(2): 129–212.

Cannon, M.D., and Edmondson, A.C., 2005. Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve. *Long Range Planning* 38(3): 299–319.

Chew, W.B., Bresnahan, T.F., and Clark, K.B., 1990. Measurement, coordination, and learning in a multiplant network. In *Measures for Manufacturing Excellence*, ed. R.S. Kaplan, 129–162. Boston: Harvard Business School Press.

Choo, A.S., Linderman, K.W., and Schroeder, R.G., 2007. Method and psychological effects on learning behaviors and knowledge creation in quality improvement projects. *Management Science* 53(3): 437–450.

Corbett, C., and Van Wassenhove, L., 1993. Trade-offs? What trade-offs? Competence and competitiveness in manufacturing strategy. *California Management Review* 35(4): 107–122.

Dutton, J.M., and Thomas, A., 1984. Treating progress functions as a managerial opportunity. *Academy of Management Review* 9(2): 235–247.

Edmondson, A.C., 1999. Psychological safety and learning behavior in work teams. *Administrative Science Quarterly* 44(2): 350–383.

Edmondson, A.C., Winslow, A.B., Bohmer, R.M.J., and Pisano, G.P., 2003. Learning how and learning what: Effects of tacit and codified knowledge on performance improvement following technology adoption. *Decision Sciences* 34(2): 197–223.

Ferdows, K., and De Meyer, A., 1990. Lasting improvements in manufacturing performance: In search of a new theory. *Journal of Operations Management* 9(2): 168–184.

Field, J.M., and Sinha, K.K., 2005. Applying process knowledge for yield variation reduction: A longitudinal field study. *Decision Sciences* 36(1): 159–186.

Gino, F., Bohmer, R.M.J., Edmondson, A.C., Pisano, G.P., and Winslow, A.B., 2006. *Learning tradeoffs in organizations: Measuring multiple dimensions of improvement to investigate learning-curve heterogeneity*. Working Paper 05–047. Boston: Harvard Business School.

Hatch, N.W., and Dyer, J.H., 2004. Human capital and learning as a source of sustainable competitive advantage. *Strategic Management Journal* 25(12): 1155–1178.

Hatch, N.W., and Mowery, D.C., 1998. Process innovation and learning by doing in semiconductor manufacturing. *Management Science* 44(11): 1461–1477.

Haunschild, P.R., and Sullivan, B.N., 2002. Learning from complexity: Effects of prior accidents and incidents on airlines' learning. *Administrative Science Quarterly* 47(4): 609–643.

Hayes, R.H., and Clark, K.B., 1985. Exploring the sources of productivity differences at the factory level. In *The uneasy alliance: Managing the productivity-technology dilemma*, eds. K.B. Clark, R.H. Hayes, and C. Lorenz, 151–188. Boston: Harvard Business School Press.

Huckman, R.S., Staats, B.R., and Upton, D.M., 2009. Team familiarity, role experience, and performance: Evidence form Indian Software Services. *Management Science* 55(1): 85–100.

Ingram, P., and Baum, J.A.C., 1997. Opportunity and constraint: Organizations' learning from the operating and competitive experience of industries. *Strategic Management Journal* 18(S1): 75–98.

Ittner, C.D., Nagar, V., and Rajan, M.V., 2001. An empirical examination of dynamic quality-based learning models. *Management Science* 47(4): 563–578.

Jaber, M.Y., and Guiffrida, A.L., 2004. Learning curves for processes generating defects requiring reworks. *European Journal of Operational Research* 159(3): 663–672.

Jaikumar, R., and Bohn, R.E., 1992. A dynamic approach to operations management: An alternative to static optimization. *International Journal of Production Economics* 27(3): 265–282.

Lapré, M.A., Mukherjee, A.S., and Van Wassenhove, L.N., 2000. Behind the learning curve: Linking learning activities to waste reduction. *Management Science* 46(5): 597–611.

Lapré, M.A., and Scudder, G.D., 2004. Performance improvement paths in the U.S. airline industry: Linking trade-offs to asset frontiers. *Production and Operations Management* 13(2): 123–134.

Lapré, M.A., and Tsikriktsis, N., 2006. Organizational learning curves for customer dissatisfaction: Heterogeneity across airlines. *Management Science* 52(3): 352–366.

Lapré, M.A., and Van Wassenhove, L.N., 2001. Creating and transferring knowledge for productivity improvement in factories. *Management Science* 47(10): 1311–1325.

Levin, D.Z., 2000. Organizational learning and transfer of knowledge: An investigation of quality improvement. *Organization Science* 11(6): 630–647.

Levy, F., 1965. Adaptation in the production process. *Management Science* 11(6): B136–B154.

Li, G., and Rajagopalan, S., 1997. The impact of quality on learning. *Journal of Operations Management* 15(3): 181–191.

Mukherjee, A.S., Lapré, M.A., and Van Wassenhove, L.N., 1998. Knowledge driven quality improvement. *Management Science* 44(11): S35–S49.

Mishina, K., 1999. Learning by new experiences: Revisiting the flying fortress learning curve. In *Learning by doing in markets, firms, and countries*, eds. N.R. Lamoreaux, D.M.G. Raff, and P. Temin, 145–179. Chicago: The University of Chicago Press.

Narayanan, S., Balasubramanian, S., and Swaminathan, J.M., 2006. *Individual learning and productivity in a software maintenance environment: An empirical analysis*. Working Paper, University of North Carolina at Chapel Hill. http://isb.edu/faculty/upload/Doc4120072316.pdf (accessed February 9, 2009).

Nembhard, I.M., Tucker, A.L., Bohmer, R.M.J., Carpenter, J.H., and Horbar, J.D., 2009. *Learn-how to improve collaboration and performance*. Working Paper 08–002. Boston: Harvard Business School.

Pisano, G.P., Bohmer, R.M.J., and Edmondson, A.C., 2001. Organizational differences in rates of learning: Evidence from the adoption of minimally invasive cardiac surgery. *Management Science* 47(6): 752–768.

Reagans, R., Argote, L., and Brooks, D., 2005. Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together. *Management Science* 51(6): 869–881.

Schilling, M.A., Vidal, P., Ployhart, R.E., and Marangoni, A., 2003. Learning by doing something else: Variation, relatedness, and the learning curve. *Management Science* 49(1): 39–56.

Sinclair, G., Klepper, S., and Cohen, W., 2000. What's experience got to do with it? Sources of cost reduction in a large specialty chemicals producer. *Management Science* 46(1): 28–45.

Skinner, W., 1974. The focused factory. *Harvard Business Review* 52(3): 113–121.

Tucker, A.L., Edmondson, A.C., and Spear, S., 2002. When problem solving prevents organizational learning. *Journal of Organizational Change Management* 15(2): 122–137.

Tucker, A.L., Nembhard, I.M., and Edmondson, A.C., 2007. Implementing new practices: An empirical study of organizational learning in hospital intensive care units. *Management Science* 53(6): 894–907.

Wiersma, E., 2007. Conditions that shape the learning curve: Factors that increase the ability and opportunity to learn. *Management Science* 53(12): 1903–1915.

Yelle, L.E., 1979. The learning curve: Historical review and comprehensive survey. *Decision Sciences* 10(2): 302–328.