

# Inside the Organizational Learning Curve: Understanding the Organizational Learning Process

By Michael A. Lapré and Ingrid M. Nembhard

## Contents

---

<b>1</b>	<b>Introduction</b>	<b>3</b>
1.1	Organizational Learning: The Defining Elements	7
1.2	Levels of Learning: Individual, Team, and Organization	10
<b>2</b>	<b>Organizational Learning Curves</b>	<b>12</b>
2.1	Learning Curve Models: The Link between Experience and Performance	12
2.2	Evidence Regarding Learning Curves	23
<b>3</b>	<b>Behind the Learning Curve: Understanding Variation in Learning Rates</b>	<b>27</b>
3.1	Frameworks for Understanding the Variation in Learning Curves	27
3.2	Variation Derived from Experience	31
3.3	Variation Derived from Deliberate Learning	41
3.4	Other Sources of Variation in Learning Rates	48

<b>4</b>	<b>Relative Effectiveness of Experience vs. Deliberate Learning as Sources of Learning</b>	<b>54</b>
4.1	The Path to Optimal Learning: Experience or Deliberate Learning?	55
4.2	Depends on Stage of Production	58
4.3	Depends on Stage of Knowledge	59
4.4	Depends on Task Characteristics	62
<b>5</b>	<b>Moving from Learning to Performance: Steps Inside the Learning Curve</b>	<b>66</b>
5.1	From Learning to Better Organizational Knowledge	67
5.2	From Better Organizational Knowledge to Changed Behavior	69
5.3	From Changed Behavior to Organizational Performance	70
5.4	Challenges to Advancing	71
<b>6</b>	<b>The Next Frontiers in Organizational Learning Curve Research</b>	<b>79</b>
6.1	Knowledge Creation	80
6.2	Development of the Learning Organization	83
6.3	Learning Curves for Other Measures of Organizational Performance	85
6.4	Learning to Improve Multiple Measures of Performance	88
6.5	Conclusion	90
	<b>Acknowledgments</b>	<b>91</b>
	<b>References</b>	<b>92</b>

## Inside the Organizational Learning Curve: Understanding the Organizational Learning Process

Michael A. Lapré<sup>1</sup> and Ingrid M. Nembhard<sup>2</sup>

<sup>1</sup> *Vanderbilt University, Owen Graduate School of Management, 401, 21st Avenue South, Nashville, TN 37203, USA, michael.lapre@owen.vanderbilt.edu*

<sup>2</sup> *Yale University, School of Public Health & School of Management, 60 College Street, P.O. Box 208034, New Haven, CT 06520-8034, USA, ingrid.nembhard@yale.edu*

### Abstract

In this work, we aim to provide an in-depth understanding of the organizational learning curve and why significant differences in the rate of learning exist across organizations. We review what is known about organizational learning curves as well as what is unknown. In sum, much is known and much remains unknown. Few studies have “stepped inside the learning curve” to provide greater understanding of the organizational learning process underlying the learning curve. We contend that this understanding is essential for helping organizations learn better and faster, and thus, operate more effectively and efficiently in a dynamic world. Therefore, not only do we examine what is known about organizational learning curves, but also what is known about the organizational learning process. Much of the former research has been conducted by operations scholars, while much of the latter has been

conducted by organizational behavior scholars. By integrating research from both (of our) disciplines, we hope to provide a more comprehensive understanding of organizational learning and the venerable organizational learning curve.

# 1

---

## Introduction

---

The learning-curve phenomenon is widely known. As individuals gain experience with a task, they get better at performing that task. This observation is reflected in the oft-repeated adage, “practice makes perfect.” The phenomenon of practice-makes-perfect has been observed not just for individuals, but also for organizations. As organizations gain operating experience, organizational performance improves, albeit at a decreasing rate. Wright [182] was the first to document this “organizational learning curve.” He found that with every doubling of airframes manufactured, the amount of direct labor hours necessary to produce a single airframe decreased at a uniform rate. Since his study, in the vast majority of the literature, organizational learning has been inferred whenever organizational performance improved as a function of operating experience. Learning curves have been observed for several measures of performance in many different contexts. For example, Figure 1.1 shows an organizational learning curve for an airline learning to reduce customer dissatisfaction.

Interestingly, organizational learning curves show tremendous variation, even when organizations perform the same task. Some organizations learn fast, some learn slowly, and some do not learn

#### 4 Introduction




Fig. 1.1 The organizational learning curve: Customer dissatisfaction with United Airlines. *Note.* Customer complaints filed by passengers with the U.S. Department of Transportation.

at all. The extent to which organizations differ in performance of the same task is amazing. Research shows that productivity for the best performer in the insurance industry is three times that of the worst performer [168]. Similarly, a comparison of regional Bell telephone companies showed that the best performer had 50% lower unit costs compared to the worst performer. Furthermore, although most of the telephone companies learned to reduce unit cost over time, some increased unit cost [168], indicating not only a slow rate of beneficial learning for some, but also that harmful learning occurs. Chew et al. [32] studied over 40 plants in a commercial food operation, and found productivity differences on the order of 3:1. Even after controlling for characteristics such as age, size, technology, and location, productivity differences of 2:1 remained. The authors noted that, “discussions with managers and our experience with plant networks studied over longer periods of time suggest that plant-to-plant variation is not a transient phenomenon and in fact, has persisted for a number of years” [32, p. 16].

However, poor learning and performance need not persist. A study by Pisano et al. [137] demonstrated the positive potential of organizational learning. The authors investigated 16 hospitals that implemented a new

technology for minimally invasive cardiac surgery, and found considerable variation in learning rates, as measured by improvement in operative procedure time. The best hospital completed the surgery in 143 minutes after 40 cases, while the worst hospital required 305 minutes after the same number of cases. Strikingly, one hospital (Hospital M) started out slowly — almost 60% slower than the sample average. However, it caught up and surpassed the sample average, attaining procedure times that were 40% faster than the sample average, after 50 cases. The authors attributed the dramatic improvement to the hospital’s use of deliberate learning activities and how they were performed.

Experts such as CEO Ray Stata of Analog Devices have argued that “the rate at which individuals and organizations learn may become the only sustainable competitive advantage, especially in knowledge-intensive industries” [157]. The rationale for the competitive advantage of learning rates lies in several trends. First, the rate of knowledge growth in many industries is astonishing. Consider the medical industry, in which over 10,000 studies are published annually about strategies to improve the clinical and operational effectiveness of health care delivery (Institute of Medicine, [81]). With such knowledge growth comes an imperative for organizations to quickly implement an abundance of new practices in order to better serve their customers. Second, organizational learning rates are important because of shorter product life cycles; the lead time for getting new products and services to market is decreasing, requiring organizations to learn to innovate faster. Third, many new ideas and technologies are complex; organizations must learn to apply them efficiently and effectively. Finally, the tremendous variation in performance across organizations creates an imperative for organization learning. To catch up with the highest performing organization, laggards have to learn faster. Likewise, if the highest performing organization wishes to stay ahead of the competition, it must improve at rate that is faster than the competition. Thus, every organization arguably has an incentive to learn as fast as possible i.e., to accelerate its organizational learning curve.

In this work, we aim to provide an in-depth understanding of the organizational learning curve and why significant differences in the rate of learning exist across organizations. We review what is known

about organizational learning curves as well as what is unknown. In sum, much is known and much remains unknown. Few studies have “stepped inside the learning curve” to provide greater understanding of the organizational learning process underlying the learning curve. We contend that this understanding is essential for helping organizations learn better and faster, and thus, operate more effectively and efficiently in a dynamic world. Therefore, not only do we examine what is known about organizational learning curves, but also what is known about the organizational learning process. Much of the former research has been conducted by operations scholars, while much of the latter has been conducted by organizational behavior scholars. By integrating research from both (of our) disciplines, we hope to provide a more comprehensive understanding of organizational learning and the venerable organizational learning curve.

We organize the remainder of this text as follows. To provide a foundation for our discussion, we begin by reviewing the definition of organizational learning (Section 1.1.) and where it occurs in organizations (Section 1.2). In Section 2, we shift attention to our primary focus — the organizational learning curve. We review various learning curve models, describing the measures of organizational experience and organizational performance that have been used to develop these models as well as the mathematical functions used to construct these models. We then summarize the evidence from these models; the evidence shows tremendous variation in organizational learning rates. Section 3 reviews frameworks for understanding this variation in learning rates and discusses variation that arises from differences in experience, deliberate learning activities, and other key sources. Section 4 examines the relative effectiveness of experience versus deliberate learning activities as sources of learning. We contend that these sources of learning affect performance through a process. Section 5 describes the steps that characterize the learning process inside the learning curve: From learning to better organizational knowledge to changed behavior to organizational performance. We discuss the significant challenges organizations need to overcome in order to advance along these steps.

Two decades ago, scholars called for organizations to become “learning organizations” [73, 149, 157]. Empirical evidence suggests that



many organizations have struggled to attain this goal [60]. We believe that this indicates a need for more research that aims to provide a better understanding of the organizational learning process and insights to guide organizations toward achievement of their learning goals. Thus, we conclude our discussion in Section 6 by outlining areas for future research that build on the admirable research that has been conducted. We believe that these areas are the next frontiers in organizational learning research. This research is needed because the imperative of organization learning has not diminished [50]. Instead, all trends indicate that the imperative continues to grow.

## 1.1 Organizational Learning: The Defining Elements

What does it mean for an organization to learn? There are several comprehensive reviews of the organizational learning process, for example Hedberg [74], Fiol and Lyles [58], Levitt and March [107] and Huber [76]. It seems that with every review, a new definition of organizational learning is offered. Table 1.1 gives only a sample of the definitions of organizational learning that scholars have offered.

Most definitions have three elements in common. The first element is a focus on the *organization*. A member of an organization can learn something, but if that learning is not captured at the organizational level, organizational learning has not occurred. Thus, organizational learning is different from individuals learning within organizations. The second common element of organizational learning across definitions is *better knowledge*. Organizations tend to have limited knowledge about why and how their actions produce organizational outcomes [84]. A critical part of organizational learning is enhancing the knowledge and understanding inside the organization. The third element is *improving actions*. The purpose of organizational learning is to facilitate changes in actions to produce better organizational performance. Implicit in most views of organizational learning is a fourth element: *ongoing effort*. Organizational learning is not a one-shot game. Instead, it is an ongoing process that should occur throughout the lifetime of an organization. Thus, integrating the common elements of organizational learning across definitions, organizational

Table 1.1. Some definitions of organizational learning.

Argyris [10]	Organizational learning is a process of detecting and correcting error (any feature of knowledge or knowing that inhibits learning) (p. 116).
Duncan and Weiss [46]	Organizational learning is defined as the process within the organization by which knowledge about action–outcome relationships and the effect of the environment on these relationships is developed (p. 84).
Hedberg [74]	Learning takes place when organizations interact with their environments: Organizations increase their understanding of reality by observing the results of their acts (p. 3).
Fiol and Lyles [58]	Organizational learning means the process of improving actions through better knowledge and understanding (p. 803).
Levitt and March [107]	Organizations are seen as learning by encoding inferences from history into routines that guide behavior (p. 319).
Stata [157]	Organizational learning occurs through shared insights, knowledge, and mental models . . . [and] builds on past knowledge and experience — that is, on memory (p. 64).
Huber [76]	An entity learns if, through processing of information, the range of its potential behaviors is changed (p. 89).
Garvin [59]	A learning organization is an organization skilled at creating, acquiring, and transferring knowledge, and at modifying its behavior to reflect new knowledge and insights (p. 80).
Kim [92]	Organizational learning is defined as increasing an organization’s capacity to take effective action (p. 43).




Fig. 1.2 The organizational learning cycle.  
*Note:* Adapted from March and Olsen[116].

learning can be defined as the organization’s ongoing effort to use better knowledge to improve its actions.

To better understand how organizational learning occurs, it is useful to review classic models by March and Olsen [116] and Kim [92]. According to March and Olsen’s model (see Figure 1.2):

At a certain point in time some participants see a discrepancy between what they think the world ought to be (given present possibilities and constraints) and

what the world actually is. This discrepancy produces individual behavior, which is aggregated into collective (organizational) action or choices. The outside world then “responds” to this choice in some way that affects individual assessments both of the state of the world and of the efficacy of the actions (p. 149).

All four elements identified above in definitions of organizational learning are evident in March and Olsen’s description of how learning occurs. While individual beliefs and actions play a key role, organizational action is different from individual action (the organizational element). Updating of beliefs — especially about action-response relationships — represents better understanding (better knowledge). By modifying behavior, more favorable environmental responses should be obtained (improving actions). Lastly, the cycle keeps repeating itself, hopefully yielding improvements over time (ongoing effort).

Kim [92] argued that there are two additional sub-processes within the learning cycle — conceptual and operational learning — that shape the first step in the learning process (i.e., the formation of individual beliefs). *Conceptual learning* consists of *assessing* cause and effect relationships that govern experienced events, and *designing* an abstract concept — a theory — to explain this experience. Conceptual learning is in essence trying to understand why events occur; it facilitates the acquisition of know-why. In contrast, *operational learning* consists of *implementing* changes and *observing* the results of these changes. Operational learning is basically developing a skill of how to deal with experienced events; it facilitates the acquisition of know-how. This cycle of observe–assess–design–implement, depicted in Figure 1.3, has several names in the literature. For example, Deming [44] called it the “plan–do–study–act (PDSA) cycle”. As the following quote by Stata illustrates, it is challenging to obtain the right balance between conceptual and operational learning:

I think to some extent, we jump back and forth between these two extremes of over-conceptualization and pure pragmatism because we don’t have the tools to connect




Fig. 1.3 Conceptual and operational learning Kim [92].

them. The core challenge faced by the aspiring learning organization is to develop tools and processes for conceptualizing the big picture and testing ideas in practice. All in the organization must master the cycle of thinking, doing, evaluating, and reflecting. Without, there is no valid learning. (Stata quoted in Ref. [149, p. 351]).

## 1.2 Levels of Learning: Individual, Team, and Organization

As noted above, many of the definitions and models of organizational learning in the literature focus on the actions of organizations and the individuals working within them (e.g., [116]). However, there is growing belief that these conceptualizations miss a critical group of actors: Teams or workgroups. Teams consists of a group of individuals that exist within a larger organization, have a clearly defined membership, and are responsible for a shared product or service [65].

Some scholars have argued that teams and team learning are the primary vehicles of organizational learning for two reasons [49, 149]. First, an increasing amount of organizational work is performed by teams. Second, teams frequently serve as the context for organizational learning as most organizational actions are complex and require coordination among team members with different expertise [134]. As team members work together, they are able to engage in team learning. Team learning describes the activities through which members acquire, share, or combine their knowledge with the goal of adapting and improving

their work processes [4]. While there are many behaviors that may serve this purpose, three behaviors are consistently associated with team learning: Speaking up, collaboration, and experimentation [48, 130].

While the understanding of individual, team, and organizational learning has largely developed through separate streams of research, there is a growing appreciation that these three levels of learning, though distinct, are interrelated [35, 40, 92]. Moreover, the levels facilitate and depend on one another. Individual learning influences team and organizational learning. Likewise, institutionalized norms, procedures and routines at the team and organizational levels influence individuals' attention, thinking, capability, motivation, and actions [40].

In this integrated process, individual learning occurs as individuals make inferences about the relationship between their actions and outcomes based on their experiences. When the individual shares his or her lessons learned with other members of the organization, individual learning combines with the learning and interpretation of other group members to influence learning at the team level. As team members share their learning, they may develop a shared understanding of each other's experience, expertise, and perspective. This understanding can lead to the modification of current practice; effective sub-practices may be incorporated, while ineffective sub-practices are refined or replaced. Effective practice changes are likely to diffuse throughout the organization. As this happens, the organization learns and practices become institutionalized. The institutionalized practices then become the basis for new individual learning. Thus, learning is an iterative, multi-level process in organizations. Knowledge and practices move from the individual to group to organizational level. Learning at the organizational level then shapes how individuals and groups act and what they learn going forward [40, 92].

# 2

---

## Organizational Learning Curves

---

### 2.1 Learning Curve Models: The Link between Experience and Performance

As discussed above, the organizational learning curve captures the notion that practice makes perfect. The learning-curve phenomenon has also been described as “learning by doing” [13]. Models of organizational learning curves generally include two core components: A measure of experience and a measure of performance. In this section, we will review measures of organizational experience and organizational performance that have been used, and models relating organizational experience and organizational performance.

#### 2.1.1 Measuring Experience

The measure most commonly used to measure organizational experience is cumulative volume. In a factory, cumulative volume is the total amount of units produced since the factory started production. In a service organization, cumulative volume is the total amount of service encounters since the organization started serving customers. Wright [182] operationalized experience with the cumulative number of

airframes produced. Since Wright’s study, many learning-curve studies have used cumulative volume (see [4] and [184] for reviews). Cumulative volume captures the notion of “learning by doing”. Task repetition to increase volume allows an organization to fine-tune operations.

An alternative measure of organizational experience is calendar time elapsed since the start of operation. Calendar time elapsed captures the notion of “learning by thinking”. The “learning by thinking” view considers time to reflect as the most critical input for learning. According to this view, it does not matter how much is produced each period. Instead, the amount of elapsed time matters because every period allows an organization to learn [56, 106]. In a study of chemical processing industries, Lieberman [111] found that cumulative volume is a better experience measure than calendar time. Calendar time can act as a proxy for variables that change over time, yet are unrelated to actual learning. For example, calendar time can capture technological progress occurring outside the organization. Because of this possibility, learning-curve studies tend to use cumulative volume rather than calendar time as the measure of experience. Learning-curve studies that use cumulative volume should include calendar time as well to rule out alternative explanations.

A third measure for organizational experience is maximum volume — the largest amount of units produced per period since the start of operation. Maximum volume therefore represents the maximum proven capacity to date. Maximum volume captures the notion of “learning by new experiences” or “learning by stretching” [99, 123]. When an organization is scaling up production, the organization needs to figure out how to execute tasks in a changed environment. The following quote by Chew [31, p. 16] captures why even adding a single machine to a factory requires learning by new experiences:

The manufacturing plant is not simply a collection of separate pieces of machinery. It is a complex network of equipment tied together by the demands of the product and the plant’s operating systems ... It is this interdependence of processes that cause adjustment costs to be so large. There is rarely such a thing as an isolated

problem in a manufacturing plant. A change at one point of the process can impact the entire web of equipment interrelationships. A change in the way step two is performed may change the input to steps ten and eleven in such a way that those process steps must be modified. Alternatively a change in step twenty-three may require slightly different inputs thereby necessitating changes in steps two and seven . . . In this way a single new machine can impact an entire \$100 million plant.

So, any time a plant attempts to add or use more capacity, factory personnel need to figure out how to solve new challenges and produce good output in new situations.

Only two studies have used maximum volume as the key organizational experience variable. Mishina [123] estimated learning curves for the production of bomber airplanes. The author found that learning-curve estimations with cumulative volume or calendar time suffered from autocorrelation, whereas learning-curve estimation with maximum volume did not. Lapré et al. [99] estimated learning curves for the production of tire-cord (steel wire used to reinforce radial tires). The learning-curve estimation with maximum volume explained a much higher percentage of variation than with either cumulative volume or calendar time. Hence, both Mishina [123] and Lapré et al. [99] found that maximum volume is a better measure for organizational experience than cumulative volume or calendar time. Future research should investigate whether these findings hold in settings beyond airplane production and tire-cord manufacturing. The scale-up argument need not be limited to manufacturing contexts. For example, expanding capacity in a call center might lead to challenges of training and coordinating a higher number of employees. Likewise, adding a wing of beds to a hospital emergency department can add to the challenge of coordinating care across departments for a larger volume of patients with urgent care needs. How can management ensure higher productivity and quality in such high and expanding volume situations? More research is needed to answer that question.



### 2.1.2 Measuring Performance

Early learning-curve studies focused on time to produce a unit and cost to produce a unit as indicators of performance. Over the past two decades scholars have significantly expanded the set of organizational performance measures.

A learning-curve effect for unit time — the time to process one unit — indicates that an organization is getting faster. Wright [182] used the number of direct labor hours to produce an airplane. Many learning-curve studies have used direct labor hours to measure organizational performance in manufacturing contexts [184]. More recently, learning-curve studies have investigated unit time in service settings as well, particularly in hospitals. Pisano et al. [137] and Edmondson et al. [52] studied learning curves for times to complete heart surgery; Reagans et al. [139] studied times to complete joint replacement. A faster organization can effectively do more with the same amount of resources. Hence, as an organization reduces unit time as a function of operating experience, one would expect unit cost to fall as well. Therefore, learning-curve scholars have extended the learning curve for unit time to unit cost — the cost to produce a single unit. Learning curves for unit cost have been observed in contexts varying from coal-burning steam-electric generating units [88] to pizza franchises [43].

Unit time is a good proxy for unit cost when (i) quality of the units produced is not affected by faster production, and (ii) direct labor constitutes the bulk of unit cost. If quality changes, it is important to at least control for quality. When changes in quality are significant, it would be preferable to study quality as a measure of performance. One measure of quality that has been used repeatedly in learning-curve studies is yield — the percentage of final production output conforming to final product specification. Related measures of quality are rejects or defects defined as 100% minus yield, and waste or scrap — the percentage of products that are scrapped because of irreparable defects. When organizations are ramping up production of new products, yields can be quite low. To illustrate the importance of learning to improve yields, Bohn and Terwiesch [25, p. 42] cite Richard Downing, a senior

VP of manufacturing at Seagate:

It is how you can improve your yield that will get your productivity up. We are not in a business where you have a 99% yield. In many cases, there are initial yields on high-end products that are in the 50% range. So a 5% or 10% improvement in these yields is significant. (Quoted in *Electronic Business Asia*, Feb. 1997, p. 35.)

Yields which are significantly less than 100% are found in many manufacturing processes for products such as semiconductors [25], microprocessors [163], pharmaceuticals [136], color picture tubes [91] and steel wire [99]. Learning to improve yields is particularly important in “immature” processes. In semiconductor manufacturing, for example, a percentage point of yield improvement is worth millions of dollars per month [175].

Unit cost is a key measure of performance. As data on unit cost can be hard to obtain, early learning-curve studies focused on direct labor required to produce one unit. Labor productivity is the number of units produced divided by direct labor hours used. For labor-intensive processes, labor productivity is a good proxy of the reciprocal of unit cost. Hence, for labor-intensive processes, the evolution of labor productivity will closely resemble the evolution of unit cost. However, many processes have multiple inputs: Labor hours, machine hours, raw material, energy, etc. Each input — or type of resource — will have its own productivity evolution. A sound measure to combine all inputs and output in a single measure is Total Factor Productivity (TFP). TFP is the value of output divided by the value of the inputs:

$$\text{TFP}_t = \frac{\text{price}_0 \times \text{output}_t}{\sum_j \text{wage}_{j,0} \times \text{input}_{jt}},$$

where  $\text{price}_0$  is the price of the product in base year 0,  $\text{output}_t$  is good output produced in period  $t$ ,  $\text{wage}_{j,0}$  is the unit cost of input  $j$  in base year 0, and  $\text{input}_{jt}$  is the amount of input  $j$  used in period  $t$ . In essence, TFP measures how effectively multiple inputs are used to produce output. TFP is easily extended to capture the production of

multiple products as well. A desirable feature of TFP is that it ignores inflation. Factory management cannot control inflation, so leaving inflation out of TFP is appropriate. Rather, TFP focuses on what factory management can control — the conversion of inputs into outputs. Learning-curve scholars have studied TFP for manufacturing of products such as computer peripheral devices [1, 2] and steel-cord [102]. Lapré and Van Wassenhove [102] studied TFP evolutions for four production lines at Bekaert, the world’s largest independent producer of tire cord. Bekaert routinely collected many single dimension quality and productivity measures, such as labor hours per ton and process interruptions at a specific step. However, management lacked an overall measure of performance. The researchers introduced a TFP measure to management and found that:

Managers at all levels felt that the TFP measure was particularly well suited [to assess overall productivity] as TFP gave them (i) a dynamic perspective, (ii) a measure not distorted by inflation effects, and (iii) a measure that aggregated trade-offs between partial measures management typically focused on. Final TFP patterns confirmed ideas managers intuitively had about productivity evolutions, even though they lacked such a measure before [102, p. 1318].

Thus, tapping into manager’s intuition can be a useful starting point for understanding and developing productivity performance measures. Scholars have recommended that, before relating TFP to learning variables, researchers graph the evolution of TFP and discuss the evolution with management to improve measurement as well identify possibly relevant control variables, including policy changes [72, 102]. Relevant policy changes may include policies regarding equipment, quality, inventory, work force, as well as policies causing confusion (such as fluctuations in production volume and number of engineering change orders) [72].

Recently, learning-curve scholars have begun to expand the set of performance measures beyond productivity, cost and conformance

quality. At the same time, scholars have started to conduct learning-curve studies outside of manufacturing. In a study of pizza franchises, Argote and Darr [7] examined service timeliness by measuring the frequency of ‘late’ pizzas. Ingram and Simons [80] studied profitability in kibbutzim (cooperative agricultural settlements). Baum and Ingram [19] and Ingram and Baum [79] analyzed organizational survival rates in the hotel industry. In the U.S. airline industry, Lapré and Tsikriktsis [101] examined customer dissatisfaction, measured by the rate of complaints against airlines filed by passengers with the government. Haunschild and Sullivan [70] studied accidents and incidents experienced by commercial airlines.

Even with an expanded set of performance measures, productivity and quality remain important dimensions of organizational performance. At the same time, it is important to continue to expand the set of performance measures beyond productivity and quality assessed by actors inside the focal organization as those outside of the organization (e.g., customers) determine organizational success as well. Future research that investigates organizational performance based on externally-driven metrics such satisfaction, repeat purchase, loyalty, and performance relative to competition should be particularly fruitful.

### **2.1.3 The Traditional Model for Relating Experience and Performance: The Power Form**

The functional form most commonly used to model a learning curve relating experience to performance is the power form:

$$c_q = c_1 q^{-b},$$

where  $c_q$  is the unit cost to produce the  $q$ th unit,  $c_1$  is the unit cost to produce the first unit, and  $b$  is the learning rate. Wright [182] introduced the power form to relate cumulative number of airframes to direct labor hours required to produce an airframe. Since his study, many scholars have used the power form as the most common learning-curve specification (see reviews by Yelle [184], Dutton and Thomas [47] and Argote [4]).

Despite its frequent use, the power form has some fundamental drawbacks. Muth [126] listed the following often observed patterns

which are not accommodated by the power form. First, initial downward concavity represents the initial difficulty in gaining learning momentum. Second, a plateau effect means that at some point additional volume does not lead to further improvement. Third, after a significant time period without any improvement, suddenly renewed improvements can occur. Similarly, Hax and Majluf [71] observed that investments can result in shifts to steeper learning curves. As we will discuss in more detail in Section 2.2, there is huge variation in learning rates across industries, across organizations, across organizational units, and across individual workers. Furthermore, a learning rate can also show variation over time. As Dutton and Thomas [47] convincingly argue, the learning rate should really be treated as a dependent variable as opposed to a given constant. Management should be explicitly responsible for managing learning rates. Arguably the most important reason why the power form does not accommodate the empirically observed patterns mentioned above is that it lacks an underlying theory. In particular, it does not provide any insight into the actual learning process that leads to performance improvement. Lastly, there is a practical problem with econometric estimation of a power form. Typically, an econometrician has performance and output data for a given time period. If there is any unknown prior production history before the given time period, then omission of prior production history will bias the estimate for the learning rate [101].

#### 2.1.4 Alternative Model for Relating Experience and Performance: The Exponential Form

The most commonly used alternative to the power form is the exponential form, introduced by Levy [108] as the “adaptation curve.” Levy assumed that a firm has a maximum output  $P$  it would like to achieve for a new process. The rate of output after producing  $q$  units is  $Q(q) < P$ . Levy’s critical assumption is that “the rate of increase in the rate of production as the firm gains experience is proportional to the amount that the process can improve” [108, p. B-238]:

$$dQ(q)/dq = \mu[P - Q(q)],$$

where  $\mu$  represents the process's rate of adaptation. Solving this differential equation yields the adaptation curve:

$$Q(q) = P[1 - e^{-(a+\mu q)}],$$

where  $a$  represent the initial efficiency for the process. Finally, the adaptation rate  $\mu$  is modeled as a function of variables  $y_1, \dots, y_n$  that the firm can use to get closer to  $P$ . Examples include prior training and experience:

$$\mu = \beta_0 \sum_{i=1}^n \beta_i y_i.$$

Levy's adaptation curve does explicitly recognize that the learning rate is a function of variables that a firm can influence. The exponential form also has the practical advantage that omission of unknown prior production history does not bias learning-rate estimates [101]. The adaptation curve does raise further questions though. First, how should one determine the maximum target  $P$ ? Even the most ambitious targets can be overtaken, thereby violating the assumption  $Q(q) < P$ . Boeing, for example, more than quadrupled its target set for a plant producing B-17 heavy bombers in World War II [123]. Similarly, Chaparral Steel typically exceeds ambitious goals set "considerably beyond current production capabilities" [105]. Second, the adaptation curve does not accommodate the empirical observations listed by Muth [126]. Third, what is the theoretical underpinning for the assumption that the rate of improvement is proportional to the amount the process can improve?

Lapr e et al. [99] built on Levy's adaptation curve to address these concerns in the context of quality. Let  $E$  be a measure for operating experience,  $W(E)$  the waste rate after the organization has accumulated  $E$  experience,  $\mu$  the learning rate, and  $P = 0$  the optimal target level for the waste rate. The researchers made several adjustments to the assumptions in Levy's model. First, in the literature on Total Quality Management, Deming [44] advocated zero defects. For quality metrics such as defects, waste, and complaints with third parties, zero is a natural target for  $P$  that can never be overtaken. Therefore, the assumption  $W(E) - P > 0$  can never be violated. Second, variables that affect the rate of learning do not have to be constant.

Allowing explanatory variables for the learning rate to vary over time ( $y_{it}$ ), makes it possible to model and observe any of the patterns listed by Muth [126]. Third, Lapré et al. [99] provide a theoretical foundation for Levy's assumption grounded in the organizational learning literature on performance gaps [27, 46, 117]. According to Lapré et al. [99, p. 600]:

“The performance gap [ $W(E) - P$ ] induces the organization to search for alternatives to reduce this gap. A larger discrepancy spurs the organization to exert more effort in searching for better knowledge (e.g., [117]). The effectiveness of acquiring new knowledge is determined by the learning rate  $\mu$ . Consequently, we can model the rate of improvement as the product of the learning rate and the performance gap.”

$$dW(E)/dE = \mu[W(E) - P]$$

With  $P = 0$ , the solution for this differential equation is the exponential form:

$$W(E) = \exp(a + \mu E),$$

and the learning rate can be modeled dynamically as:

$$\mu = \beta_0 \sum_{i=1}^n \beta_i y_{it}.$$

The exponential form with a static learning rate has been used for yield variation in electromechanical motor assembly [56], complaints against airlines [101], and service failures by airlines [98]. If calendar time is used for  $E$ , the exponential form with a static learning rate is Schneiderman's [147] half-life curve. The half-live curve has been used by companies such as Analog Devices to measure the rate of organizational learning [157].

The exponential form overcomes several shortcomings of the power form, especially in the context of quality improvement where zero defects provide a natural, objective target. Future research is needed to

objectively determine the target level  $P$  for performance measures that lack a clear, natural target level such as cost, productivity, and lead-times. Consequently, for quality metrics the exponential form would be the preferred specification, whereas for many other measures of performance such as cost it is still an open question as to what constitutes the preferred learning-curve specification.

### 2.1.5 Other Models for Relating Experience and Performance

Kantor and Zangwill [90] note that many systems consist of subcomponents. Complicated systems such as aircraft are manufactured with a variety of advanced technologies (e.g., composites and titanium). Kantor and Zangwill [90] show mathematically that if a learning system consists of multiple learning subcomponents, the power form for the total system cannot be obtained by summing the learning curves for the subcomponents. Instead, the authors advocate the use of the exponential form to represent total system learning as a proper summation of subcomponent learning. Production data suggests that the learning curve could be a sum of a few exponentials.

Although Levy [108], Kantor and Zangwill [90], and Lapré et al. [99] succeeded in doing so, most theoretically derived learning-curve models are difficult to estimate with real production data because very specific information is required. Muth's [126] search theory, for example, requires a specification of the distribution function of possible improvements. Application of Mody's [124] work demands knowledge about the productivity of engineers and engineering costs. See Kantor and Zangwill [90] for similar estimation difficulties of other learning-curve models.

In the management science literature, scholars have proposed models to link quality and learning; see e.g., Refs. [42, 57, 160].<sup>1</sup> These models yield optimal quality control policies concerning the timing and amount of inspecting the production process. Inspection (or maintenance) provides an opportunity to learn about the process,

---

<sup>1</sup>For a review of empirical and normative literatures on process change, see Carrillo and Gaimon [29].



thus reducing the probability that the process produces defective units in the future. These models assume an exogenously given effectiveness (be it deterministic or stochastic) for all learning activities. They offer no insight why some learning activities are more effective in producing process control knowledge than others.

## 2.2 Evidence Regarding Learning Curves

There is tremendous evidence regarding learning curves. Reviewing the entire learning-curve literature falls well beyond the scope of this text. Instead we want to use this subsection to (i) demonstrate that there is ample evidence for learning curves both in manufacturing and service settings (i.e., the learning curve is a robust phenomenon), and (ii) there is tremendous variation in learning rates.

### 2.2.1 Evidence from Manufacturing

The learning rate  $b$  in the power form can be expressed as a progress ratio  $p = 2^{-b}$  [4]. For each doubling of cumulative volume, the new unit cost is  $p$  % of the old unit cost. So, with an 80% progress ratio, each doubling of cumulative volume reduces cost by  $100\% - 80\% = 20\%$ . Dutton and Thomas [47] compared progress ratios for over 100 learning-curve studies in manufacturing processes for industries such as electronics, machine tools, EDP system components, papermaking, aircraft, steel, apparel, and automobiles. The authors found tremendous variation in progress ratios not only across industries, products, and processes, but also for subsequent runs of the same product within the same plant. The steepest learning curve had a progress ratio of 55% (indicating that cost more than halved with every doubling of cumulative volume), whereas one learning curve had a progress ratio greater than 100% (indicating that costs actually increased). For 107 out of 108 studies, cost decreased as a function of cumulative volume. The modal progress ratio was around 81–82%, but Dutton and Thomas found very significant variation in progress ratios around the mode.

The tremendous variation in learning rates has been found for other performance measures as well. Adler [1] studied TFP learning curves for eight departments in a single manufacturing firm producing computer

peripheral devices. For all eight departments, cumulative volume had a positively significant effect on TFP indicating that all eight departments learned to improve TFP. The fastest learning rate was more than twice as large compared to the slowest learning rate. Likewise, in a single steel-cord manufacturer, Lapré and Van Wassenhove [102] studied TFP learning curves for four production lines. TFP learning curves ranged from steep learning to no learning at all. Epple et al. [53] studied labor productivity in an automotive assembly plant that switched from operating one shift to two shifts. The authors found that after the switch the day shift continued to learn at roughly the same rate as the plant did before the switch, whereas almost no learning occurred on the night shift. Hatch and Dyer [67] plot defect density trends for 30 semiconductor manufacturing processes in 16 firms. The plot shows that starting points for defect learning curves differ by a factor of eight. Furthermore, the slopes of defect reduction show tremendous variation.

### **2.2.2 Evidence from Services**

The study that jump-started the field to consider learning curves in services is Darr et al. [43]. The authors studied learning curves for unit cost in pizza franchises. They found significant evidence of learning albeit at a slower progress ratio of 93% compared to the 80% modal progress ratio in manufacturing. Like manufacturing however, absolute measures of organizational performance for services (e.g., efficiency) are generally believed to improve over time [19].

In contrast, relative measures of performance in services, i.e., measures that capture performance relative to competitive rivals, do not appear to improve indefinitely. Recent learning-curve studies in services that have used relative measures of performance have found a U-shaped learning-curve effect: Initially, organizational performance improves as a function of operating experience, but eventually deteriorates. U-shaped learning-curve effects have been documented for failure rates in the hotel industry [19, 79], and customer dissatisfaction against airlines [101, 98]. The explanation for a U-shaped effect is rooted in the trade-off between exploitation and exploration [115]. As organizations focus on exploitation, repeated practice with an existing set of routines

makes perfect. However, focus on exploitation at the expense of exploration can place organizations in “competency traps” [107]. The set of routines perfected by the organization can become inadequate in a changing environment. The criteria for organizational success change after the organization has learned [19].

Research shows that learning-curve heterogeneity exists in service settings just like it does in manufacturing settings. In a study of the implementation of minimally invasive cardiac surgery in 16 hospitals, Pisano et al. [137] found significant variation in the rate at which hospitals learned to perform the surgery. Lapré and Tsikriktsis [101] also found significant learning-curve heterogeneity for customer dissatisfaction with airlines, i.e., differences in learning from customer dissatisfaction across airlines. Figure 2.1 shows just a small sample of the heterogeneous organization learning curves from Lapré and Tsikriktsis [101]. United Airlines started out worse than America West, but eventually caught up with and surpassed America West. Neither airline was able to catch up with the benchmark in the industry — Southwest. Both United and America West struggled more than Southwest did




Fig. 2.1 Variation in organizational learning curves: Customer dissatisfaction with three airlines.

*Note:* Customer complaints filed by passengers with the U.S. Department of Transportation. AW, America West; SW, Southwest Airlines; UA, United Airlines.

to consolidate initial improvements. In a follow-up study, Lapré [98] decomposed customer dissatisfaction into two components: Service failure and customers' propensity to complain given the occurrence of a service failure. He found that learning-curve heterogeneity originated in the propensity to complain as opposed to service failure. In a study of the Royal Dutch Mail, Wiersma [178] found significant learning-curve heterogeneity across 27 geographically dispersed regions in the Netherlands. So, just like in manufacturing, there is tremendous variation in learning rates in services.

# 3

---

## Behind the Learning Curve: Understanding Variation in Learning Rates

---

So far, we have discussed the importance of organizational learning and evidence of organizational learning curves. In sum, the evidence consistently documents the organizational learning curve phenomenon. It also shows that there is tremendous variation in learning rates. Thus, a key goal of learning-curve research is to further our understanding of what causes the tremendous variation.

### 3.1 Frameworks for Understanding the Variation in Learning Curves

#### 3.1.1 Levy (1965): Origins of Autonomous vs. Induced Learning

As discussed in the previous section, Levy's [108] adaptation curve modeled the learning rate as an explicit function of variables that a firm could use to improve performance faster. To categorize such variables, Levy identified three types of firm learning: Planned or induced learning, random or exogenous learning, and autonomous learning. *Planned or induced learning* is defined as "learning that results from the firm's applying techniques that are designed to increase the rate of output

or, equivalently, to reduce costs in the production process” (p. B-139). Planned learning could involve pre-planning such as building prototypes and testing raw materials. Planned learning can also take place after a production process has started by engaging in “industrial engineering” activities such as time and motion studies and redesigning raw material specifications.

*Random or exogenous learning* consists of “improvements in production processes that can result when a firm acquires information unexpectedly from its environment” (p. B-139). Sources of external information that could speed up production include suppliers, government and trade publications, competitors and customers.

*Autonomous learning* is “improvement due to on-the-job learning or training of employees” (p. B-140). As employees gain experience with a production process, they become better at trouble shooting. Employees can benefit from mistakes or problems that occurred in the past, and correct issues faster. Likewise, employees working together in a group get used to each other’s actions and responses allowing the group to work faster.

### **3.1.2 Dutton and Thomas (1984): Autonomous vs. Induced and Endogenous vs. Exogenous**

In the previous section, we described Dutton and Thomas [47] impressive study of variation in learning rates. Based on the tremendous variation in learning rates, the authors concluded that the learning rate is neither fixed, nor automatic. Hence, the learning rate should be viewed as a dependent variable influenced by a number of potential policy variables. “Once it is realized and accepted that the rate of improvement is not a given, the question immediately shifts to the issue of how the [learning rate] may be managed” (p. 240). Revisiting Levy’s [108] classification, Dutton and Thomas realized that induced vs. autonomous learning make up a “learning-type” dimension, whereas “origin” (exogenous vs. endogenous) is a dimension that is orthogonal to the learning-type dimension. So, both induced and autonomous learning could have endogenous and exogenous origins. See Table 3.1.

Table 3.1. The framework of Dutton and Thomas [47] with some examples.

	Autonomous learning	Induced learning
Exogenous origin	Productivity improvement from periodic equipment replacement	Copying process settings derived elsewhere (R&D, other plants)
Endogenous origin	Learning by doing in a plant (“practice makes perfect”)	Quality/productivity improvement projects

Dutton and Thomas [47] characterize *induced learning* as requiring “investment, induction, or resources made available that are not present in the current operating situation” (p. 241). *Autonomous learning*, on the other hand, “involves automatic improvements that result from sustained production over long periods” (p. 241). *Exogenous learning* “results from information and benefits acquired from external sources such as suppliers, customers, competitors, and government”, whereas *endogenous learning* “is attributable to employee learning within a firm as manifested by technical changes, direct-labor learning, and smoothing production flows” (p. 241).

Until Dutton and Thomas [47], Levy [108] was the only study to empirically incorporate induced-learning variables in a learning-curve analysis. However, Levy’s variables were fixed, i.e., his variables of prior training and prior experience did not evolve over time. Dutton and Thomas [47] advocated using longitudinal variables from different cells in Table 3.1 to study questions such as (p. 244):

1. “Does the cumulative effect of regular short run adaptations or inducements yield significant progress relative to a long run inducement?”
2. “When does the cumulative progress due to short run inducements asymptote?”
3. “When and how does the system have to be regenerated in order that progress may continue?”

The authors note the absences of longitudinal studies that control for different factors, thus isolating relative effects. Most notably, there were no studies incorporating organizational behavior variables in learning-curve analyses. Even to date, these questions advanced by Dutton

and Thomas [47] remain understudied. Below, we will review the very few studies that have incorporated both longitudinal autonomous and induced learning variables [2, 14, 68, 67, 83, 99]. The only study that incorporated a longitudinal variable for each cell in Table 3.1 is [102]. So, even now, the research opportunity identified by Dutton and Thomas [47] remains under-explored. A significant challenge is measuring the actual learning process that underlies the learning-curve effect. We will discuss the importance of studying this process next.

### 3.1.3 Bohn (1994): Inside the Learning Curve

Traditionally, the learning-curve literature has treated ‘learning’ as a black box as in Figure 3.1, Part A. Learning is inferred if cost or quality improves as a function of operating experience. Bohn [21], however, pointed out that there is an actual learning process inside the learning curve as depicted in Figure 3.1, Part B. Learning can result from experience (autonomous learning) or deliberate activities (induced learning). Furthermore, he posited that learning can yield better organizational knowledge. Better organizational knowledge can persuade organizational members to modify their behavior. Changed behavior, in turn, can improve organizational performance. None of these steps are

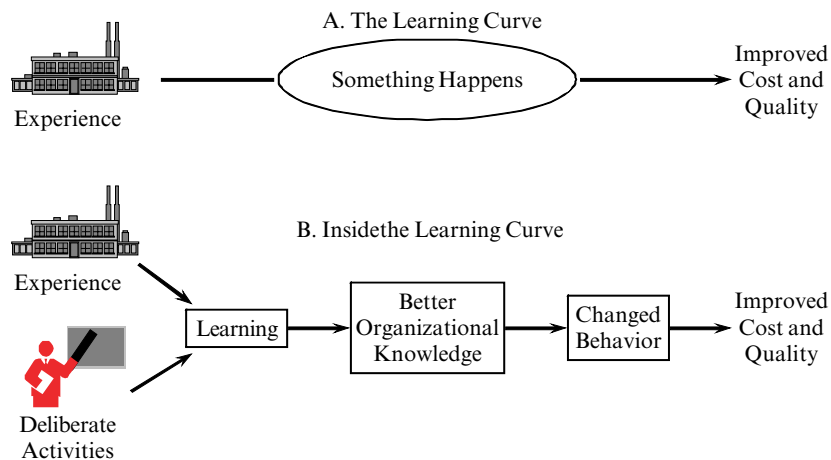


Fig. 3.1 Two views of the learning curve.  
*Note:* Adapted from [21].



trivial. Scholars have merely scratched the surface in terms of studying these steps. No single empirical study has incorporated all steps [97]. Only two studies have incorporated a step from Figure 3.1, Part B with longitudinal variables in learning-curve estimations [14, 99].

Figure 3.1 provides a compelling picture why there is so much variation in learning rates. Organizations can differ in terms of:

- the amount of experience,
- the nature of experience,
- the ability to learn from experience,
- the amount of deliberate learning activities,
- the nature of deliberate learning activities,
- the ability to learn from deliberate activities,
- the ability to translate learning into better organizational knowledge,
- the ability to change behavior in response to better organizational knowledge, and
- the ability to obtain better organizational performance as a result of changed behavior.

Next, we review some of the findings in the literature regarding variation along these factors.

### **3.2 Variation Derived from Experience**

Across frameworks, there is agreement that experience is a core mechanism for facilitating organizational learning. Additionally, there is growing appreciation that all experience is not the same. Some organizations accumulate their experience in an industry by focusing on a limited number of products or services, while others accumulate their experience by producing a wide variety of products or services over the same period. Some organizations focus on learning from their successful experiences, while others focus on learning from failures. For some, effort is made to acquire experience at the individual level, for others at the team level, and still for others, at the organizational level. Thus, the same quantity of experience can be qualitatively different on a number of dimensions. The extent to which these differences in

the nature of experience affect learning and performance has been of growing interest to researchers [78, 148]. To date though, scholars have largely focused on the three attributes of experience captured in the examples above as explanations for the variation in the rate and effectiveness of organizational learning. That is, they have focused on level of specialization (versus diversification), outcome (success versus failure), and locus of learning (individual, team or organization). In this section, we review current knowledge about the impact of each of these variables on learning rates.

### **3.2.1 Specialized vs. Diversified Experience**

A critical decision for many organizations is what level of specialization to pursue. On one hand, a high degree of specialization (i.e., a narrow market, product or operational focus) is believed to be preferred. Ever since Adam Smith's (1776) seminal work on the efficiency of division of labor, researchers and practitioners have believed that specialization improves learning and performance. This belief was reinforced 200 years later by Wickham Skinner's [155] observation about the "focused factory." Skinner studied 50 firms across six industries and observed that those firms that specialized in a limited number of activities outperformed those that maintained a more diversified portfolio of activities. Drawing on this observation, scholars (e.g., [145]) have recently questioned whether organizational learning is also maximized through specialization, or whether some variation or diversification of experience improves learning rates.

A series of studies conducted in the last decade in a variety of industries have sought to answer this question. The earliest studies provided supporting evidence for the superiority of specialized experience over generalized experience, the broadest form of diversified experience. For example, Barnett et al.'s [17] study of the banking industry showed that specialist banks had higher returns on average assets (ROAA) as a function of experience, while generalist banks had no increase in ROAA as a function of experience. Similarly, in the U.S. hotel industry, Ingram and Baum [79] 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). The former failed less frequently than the latter. Likewise, in the airline industry, Haunschild and Sullivan [70] found that specialist airlines that analyzed the heterogeneous causes of their prior accidents subsequently had lower accident and incident rates than generalist airlines that did the same. Lapré and Tsikriktsis [101] also studied the learning curves of specialist and generalist airlines and found that, although the average specialist airline learned at the same rate as the average generalist airline, the best specialist airline learned faster than the best generalist airline.

However, not all research suggests that specialized experience is optimal for learning. An emerging body of work suggests that a good balance between specialized experience and widely diversified or generalized experience maximizes learning. In this work, the researchers have differentiated between specialized experience (acquiring the same experience repeatedly), related experience (acquiring similar, but different types of experiences), and unrelated experience (acquiring different experiences, akin to generalized experience in the aforementioned studies). Using this approach in settings ranging from laboratory experiments [145] to software development [20] to mail services [178], researchers have found that the learning rate for those organizations that acquire related experience is significantly greater than the learning rate for those that acquire either specialized or unrelated experience. Additionally, they have found no significant difference between learning rates for specialized and unrelated experience. This set of studies suggests an inverted U-shape relationship between exposure to variety and performance. That is, more variety is better, but only up to a point. Thereafter, more variety limits the organizations ability to learn. Studies on individual learning also find this relationship (e.g., [127]).

Scholars have theorized that organizational learning is greatest when organizations achieve a good balance between specialized and highly diversified (unrelated) experience, rather than emphasize one or the other, because each has its advantages and disadvantages. Specialization has the advantages that come with focus — a deeper understanding of a defined area and easier transferability of knowledge

from the last experience to the current experience due to similarity. However, specialization can also lead to stagnation rather than growth due to repetition. In contrast, diverse experiences have the advantages of stimulating new ideas, prompting consideration of more associations and synthesis of ideas, and fostering more complex understanding that pertains to a wider scope of experiences. However, diversity can make it difficult to integrate and apply knowledge across experiences. Schilling et al. [145] concluded that this is why related experience is superior to specialized and unrelated/generalized experience for learning in the long-term. It has the advantages of both experiences, without having their disadvantages. More research is needed to understand the relative importance of the advantages, and what problem-specific, learner-specific, and context-specific factors mediate the effect of the advantages on learning. The need for such research has been articulated by Schilling et al. [145]. We agree that this is the next logical and important step in understanding how specialized and diversified experience alters learning. Future research should also investigate the existence of managerial practices that further minimize the disadvantages of these experiences.

### **3.2.2 Success vs. Failure Experiences**

Regardless of the organization's level of specialization, inevitably, some of the organization's experiences are successes and some are failures. What constitutes success versus failure is determined by reference to an aspirational level [41, 63, 117]. When performance achieves or exceeds the aspirational level, success is said to occur. In contrast, when performance falls below the aspirational level, a failure is said to occur. The behavioral theory of the firm argues that organizations respond differently to the experience of success and failure [41, 117]. Success prompts local search within the scope of current actions to refine and reinforce lessons from past experiences, while failure prompts non-local search for new possibilities to correct or enhance performance. The differences in search strategy and therefore potential learning have led researchers to ask whether the relevant experience required for organizational learning is success or failure.

Most learning-curve studies implicitly assume learning from success [4], however a growing group of scholars have argued that learning from failure plays a central role in organization learning [28, 122, 154], an argument that has been affirmed by several studies. For example, qualitative studies of high reliability organizations (HROs, e.g., nuclear power plants) show that these organizations' stellar performance stems in part from their active learning from failures [141, 176]. Quantitative studies have also documented the value of learning from failure. In a series of studies, Haunschild and colleagues showed that organizations in the airline and automobile industries that previously experienced failures (i.e., accidents and product recalls, respectively) had reduced rates of future failures, an indication that they learned from their past failures [69, 70]. The benefits of learning from success have been documented as well [154] in settings ranging from commercial banks [93] to railroads [18]. Organizations with more past successes are found to have more future successes and less future failures.

Thus, both success and failure experiences can facilitate organizational learning and performance. However, Kim et al. [93] recently showed that the positive relationship to performance for both experiences takes time to emerge. They studied over 2600 banks during a 15-year period, and found that both success and recovery (a form of failure) experiences were initially associated with increased bank failure. Success and failure then favorably reduced bank failure, leading to an inverted U-shaped curve between experience (success or failure) and poor performance. This relationship led Kim et al. [93] to conclude that organizations must accumulate a certain amount of the same experience — success or failure — before organizational performance will improve as a result of learning from that experience.

While the evidence is fairly robust that both success and failure experiences can facilitate performance-enhancing organizational learning, results are mixed about which experience is more advantageous. In a learning-curve study of U.S. railroads, Baum and Dahlin [18] found that accident costs declined (i.e., improved) as a function of cumulative operating experience (an indicator of success experience), but not as a function of accident experience (an indicator of failure experience), suggesting greater benefits from learning from success than

from failure. However, other studies suggest the contrary i.e., learning from failure is more beneficial than learning from success. For example, Li and Rajagopalan's [109] study showed that cumulative number of defective units (failures) contributed more to learning-curve effects in manufacturing firms than cumulative number of good units (successes). Likewise, Madsen and Desai's [113] study of firms in the orbital vehicle launch industry showed that the number of prior failed launches better explained variation in launch failure likelihood than the number of prior successful experiences. In their study, the more failed launches a firm had experienced, the lower the firm's likelihood of a subsequent failed launch. The same relationship held for past success, but to a lesser degree, and past successes did not significantly lower the likelihood of failure once models adjusted for past failures.

Failure experience is theorized to be a particularly effective stimulant for learning because it is highly salient, directly challenges the notion that current practices are adequate, and thereby provokes interest in identifying and developing alternative approaches [62, 122, 154]. Failures are the triggers in the aspiration-performance feedback relationship; they create an urgency to reflect, challenge old assumptions, and innovate to achieve aspirations. In contrast, success encourages preservation of the status quo, complacency about experimenting with new ideas, and risk aversion. Thus, success inspires a narrower scope of learning and change than failure. Several studies have shown that organizations do not initiate change when their performance is satisfactory or successful (e.g., [62]), but do embrace change when their performance is poor (e.g., [96]). Thus, part of the variation in learning curves likely stems from the distribution of success versus failure experienced by organizations.

Research indicates that whether and how an organization responds to and learns from its success and failure experiences depends on a variety of factors. First, it depends on the *nature of the success or failure*. In their aforementioned study of specialist and generalist airlines, Haunschild and Sullivan [70] showed that airlines learned more from accidents and incidents (failures) with heterogeneous causes than those with homogeneous causes. They theorized that heterogeneous causes (due to their complexity) promoted broader search for causality, and

better analysis and solutions in turn. Second, research suggests that the rate of learning depends on *the level of each experience, and the presence of the other experience*. Kim et al. [93] showed that a threshold level of success or failure experience is required for an organization to learn from the same experience. Furthermore, learning from success or failure experience is enhanced by the organization's experience with the other type of experience, a finding consistent with theories that emphasize the importance of contrasting experience in providing useful knowledge. Third, *level of aspiration* appears to influence rate of learning from experience. Baum and Dahlin's [18] study of railroads showed that when a railroad's accident rate deviates significantly from aspirational levels, the railroad benefits less from its own success and failure experience (and more from other railroad's experience). This finding supports the notion that performance away from aspirations (failure) stimulates non-local search, while performance near aspiration (success) fosters local search. Finally, learning from success or failure experience depends on *context*. Chuang and Baum's [34] study of Canadian nursing homes showed that when other organizations in the industry experience the same failure, the focal organization is less likely to learn from its own failure, potentially because it holds less regard for these failures. Additionally, if the organization has an historical investment in a strategy, it is less likely to learn from new information, which Chuang and Baum attribute to the organization succumbing to competency traps that prevent it from accumulating sufficient experience with alternative strategies to realize their value.

We believe that the four categories of factors we identify as determinants of the effectiveness of success and failure experience — nature of the success or failure experience, the level of each experience and the presence of the other experience, level of aspiration, and context — are a useful lens for continued examinations of the variation in learning caused by these experiences. As our review indicates, relatively few studies (one or two) have examined variables within these categories. More research is needed to understand the relative magnitude of influence of the variables that have been studied, as well as to identify the full scope of relevant variables within each category. At the same time, we encourage research to identify other categories of factors that

influence learning from success and failure. In addition to the four we cite, others such as the nature of the task may be highly influential.

### **3.2.3 Individual vs. Team vs. Organizational Experience**

In addition to investigating how the attributes of experience alter the organizational learning curve, researchers have focused on a series of questions about how the level at which learning in the organization occurs shapes the curve. As mentioned earlier, there are three levels at which learning is said to occur in organizations: The individual, team, and organization [3, 35, 40]. However, the question remains: At what level must experience be accumulated for organizations to learn most effectively? Must learning occur at the individual, team, and organizational level or some combination of these? Can experience at one level substitute for experience at another?

Each level of experience has been theorized to provide learning and performance benefits [139]. With increased cumulative individual experience comes individual proficiency through knowledge and skill development. With cumulative team experience (i.e., experience working together) comes better coordination and teamwork as individuals learn who knows what, who is best at performing each task, and how to trust each other. With cumulative organizational experience, staff gain the opportunity to learn from the knowledge accumulated by others. Reagans et al. [139] studied the contribution of each of these levels of experience to performance improvement in the orthopedic department of a large teaching hospital over five years, and found that the accumulation of surgical experience at each level impacted performance (i.e., procedure time), while controlling for the impact of the other two levels. As hypothesized, both team and organizational experience had a consistently positive effect on performance. The same was not true for individual experience. It had a more complicated, U-shaped relationship to performance such that at low levels of individual experience, increasing individual experience hurt procedure completion time; but at higher levels of individual experience, more individual experience contributed to better procedure completion times. Reagans et al. [139] hypothesized that the change in relationship over time resulted from



improved application of acquired knowledge: Initially individuals inappropriately apply what they learn from past encounters, and therefore harm performance. However, with more experience, they learn what knowledge to apply when, which enhances collective performance.

In the very different context of Indian software development, Huckman et al. [77] further investigated the relationship between team experience and performance. In addition to examining the experience that team members had working with each other (i.e., team familiarity) as Reagan et al.'s had done, they also examined the overall role experience of individuals in the team (i.e., average years in a given role within a team). They found that both of these dimensions of team experience were positively associated with the quality and efficiency of software teams' products. Overall, the more experience team members had working together and in their role within the team, the better the team performed on both measures of performance. However, the importance of role experience with respect to product quality varied by role, with manager role experience having no significant relationship to product quality while team member role experience was positively and significantly related. Huckman et al.'s follow-up discussions with the firm's staff led them to conclude that the difference in effects reflects role differences in the ability to monitor progress and perform mid-course corrections. When roles readily allow individuals to monitor their progress and learn from their actions in real-time, more likely with staff roles than managerial roles, collective performance improves. While their data did not allow them to test this hypothesis, their results nevertheless provide evidence that team experience, and specifically role experience, is an important contributor to organizational learning. Additionally, their results provide a possible explanation for Reagan et al.'s finding of a U-shaped relationship between individual experience and performance. The relationship may reflect the individual's role experience. Early in a new role (e.g., working with a new team), individuals may inappropriately apply knowledge gained from past experience to their new role, leading to poorer performance. However, as the individual develops better understanding of and more experience in the new role, better application of experience likely occurs. In turn, performance is likely to improve. Of course, research is needed to determine

whether these hypothesized relationships hold true. It would also be informative for future studies to evaluate whether the effect of individual role experience is contingent on team and organizational experience. It is likely that individual role experience is less detrimental to learning when team and organizational experience are high. Even if this is true, the stream of research on level of experience provides convincing evidence that learning-curve models would be enhanced by incorporating measures of experience at all levels, and decomposing experience at each level (e.g., into tenure, role, and task).

### **3.2.4 The Combined Effect of Different Experiences**

To date, the majority of studies that have examined how different attributes of experiences alter the rate and effectiveness of organizational learning have investigated one attribute. However, researchers have begun to consider how the interaction of different attributes alters the learning curve. Boh et al. [20], for example, investigated whether the benefit of specialized versus diversified experience varies with the level at which experience is acquired (individual versus team versus organization). They showed that it does: Specialized experience most enhances individual learning, while related experience most enhances team and organizational learning. In the only other study we found that examines the interaction between experience types, Haunschild and Sullivan [70] further showed that the benefit of failure experience is contingent on the level of specialized experience. In their study, specialist airlines learned more from failure than generalist airlines, leading them to conclude that learning from failure is challenging for organizations with complex designs like generalists. Complex organizations tend to be more political, more hierarchical, and more compartmentalized. These characteristics hinder learning from experience, particularly failures, because they motivate hiding or recasting experiences to present the best impression rather than effort to better understand experiences [118]. It remains unknown how other experiences combine to shape organizational learning curves. This is a critical, missing link for our understanding these curves.

### 3.2.5 Summary: Variation Caused by Different Experiences

While there remains much to be learned about the variation in organizational learning curves caused by different kinds of experiences, existing research offers several insights. First, the accumulation of experience at all levels of the organization — individual, team and organization — enhances the rate of learning. Second, the rate of learning is enhanced by acquiring multiple forms of experience, rather than emphasizing one. Too much of any particular experience stifles organizational learning as indicated by the inverted U-shaped relationship to performance found in studies of both specialized/diversified experiences and success/failure experiences. It is better for organizations to acquire related experience (a combination of specialized and diversified experience) and a combination of success and failure experiences. A third insight is that not all experiences are equally beneficial as both theory and empirical research suggest that failures accelerate learning more than success. Finally, research shows that the benefit of any experience depends on other factors (e.g., the nature of the experience, the level of other experiences, the aspirational level, and the context). More research is needed to identify the full range of experiences that influence learning, which experiences most facilitate learning, the range of conditions that alter the effectiveness of each experience, and how different experiences interact to affect learning.

## 3.3 Variation Derived from Deliberate Learning

Relative to experience, much less research has investigated the impact of deliberate learning on the rate of organizational learning, even though prominent frameworks for understanding variation in learning curves (see Section 3.1) identify deliberate learning as a determinant of organizational learning. As discussed above, deliberate learning is distinct from autonomous learning through cumulative experience in that it results from the planned activities of managers and staff conducted with the explicit intent of acquiring, creating, and implementing new knowledge. Almost by definition, variation in deliberate learning can be expected because organizations vary in both their choice of activities

and in the level of investment made by managers and staff in the chosen activities due to contextual differences. In this section, we review the research on variation in learning rates caused by choice of deliberate learning activities and caused by contextual differences that alter the investment that managers and staff make in these activities.

### **3.3.1 Types of Deliberate Learning Activity**

Research on deliberate learning has identified a variety of activities that organizations use for this purpose. Examples include training sessions, engineering programs, experiments, dry-runs of new practices, quality management programs, quality circles, and employee suggestion programs [2, 14, 99, 166]. The effectiveness of these deliberate learning activities (DLAs) has been examined in several studies. However, Adler and Clark [2] were the first to include any DLAs in a learning curve, and examine the possibility that different DLAs might have different effects on the learning curve. Using data from an electronic equipment company, they constructed a model of productivity improvement as a function of cumulative experience as well as two DLAs: Cumulative training activity and cumulative engineering activity, which included conducting experiments and learning new specifications. Much to their surprise, it was not the case that the two activities simply varied in the magnitude of their effect on the learning curve, but rather, they varied in the directionality of their effect between departments in the company. In one department, training activity helped productivity, while engineering activity harmed productivity. In a second department, the opposite was true. The researchers hypothesized that differences in capital-intensity and motivation explained the contrasting effects of the two DLAs as they qualitatively observed that training was more beneficial in the more labor-intensive, performance-oriented department. While they were unable to formally test their hypothesis, their study was historic because it provided the first evidence that DLAs can alter learning curves, and that different DLAs can have dramatically different effects.

Hatch and colleagues took the next step by assessing whether the influence of DLAs extends beyond productivity learning curves and

whether types of DLA other than cumulative engineering activity and training activity are influential [68, 67]. Their studies showed that yield learning curves were also altered by DLAs and by more DLAs than previously studied. Specifically, they found that the use of human capital activities focused on human resource selection (e.g., screening tests), deployment (e.g., statistical process control training) and use (e.g., problem-solving teams) improved organizational learning and performance. They reasoned that the learning and performance benefits derived from having “better human resources and from better practices to develop firm-specific human capital and deploy it to learning activities” in an effective and efficient way [67, p. 1173].

Three additional studies provide further evidence that variation in use of DLAs should be included alongside cumulative experience as explanation for the differences in organizational learning curves. Ittner et al.’s [83] study of 12 manufacturing plants over 9 years showed that organizations’ deliberate investment in preventive engineering activities predicted their current quality levels. Arthur and Huntley’s [14] study also showed that the cumulative number of implemented employee suggestions (a form of DLA) was significantly associated with reduced production costs in an auto parts manufacturer with a gainsharing program. Nembhard and Tucker’s [131] study further showed that the influence of DLAs extends to organizations in dynamic service settings like hospitals, although the benefit of DLAs take time to emerge in this setting. Initially, their use increases the odds of poor performance in this setting, much like success and failure experiences do in the context of service organizations like banks (Kim et al. [93]; Section 3.2.2). After a period, the use of DLAs then fosters high performance.

In the majority of studies, the use of DLAs ultimately contributes positively to organization’s learning curve. However, as Adler and Clark’s [2] early work showed, that is not always the case. Different DLAs can have different, sometimes contradictory, effects in different organizations. Lapré, Mukherjee, and Van Wassenhove provided insight into why that might be [99, 102, 125]. They studied five learning activities (and 3 functions) used in quality improvement projects at a tire cord manufacturing plant, and found through factor analysis that the DLAs mapped onto two distinct factors. The first factor — use

of statistical experiments, use of scientific models or engineering staff, and levels of Ishikawa diagrams — they labeled as activities for generating conceptual learning because these DLAs fostered the acquisition of know-why. The second factor — modification of action variables and follow-up of experimental results — they described as activities for generating operational learning because these activities provided know-how. Lapré et al. [99] observed that projects with different distributions of these two factors of DLAs affected the plant's learning curve for waste reduction differently. Projects with DLAs that generated both know-why and know-how accelerated waste reduction, whereas projects that generated know-why without know-how reduced waste reduction. Thus, their data suggested that DLA-related variation in learning curves stems from the nature of learning that different types of DLA generate.

Tucker, Nembhard and Edmondson [166] provide support for this explanation. In the very different context of hospital intensive care units, they too investigated the effects of 12 different DLAs, and found through factor analysis that DLAs in use by hospital improvement teams mapped onto two qualitative groups. One group of activities, termed learn-what, helped teams identify best practices. The second group, termed learn-how, consisted of activities that helped teams discover the underlying science of a new practice and operationalize the practice in the organization. In helping organizations achieve both of these goals, learn-how arguably facilitates both conceptual and operational learning, as described above. Tucker et al. found that use of learn-how activities were positively associated with the implementation of new practices, but that learn-what activities were not. Learn-how activities included solicitation of staff ideas, opportunities for staff to provide feedback before full implementation, education sessions with staff, pilot runs, dry runs, project team meetings and problem-solving cycles (Plan–Do–Study–Act). Learn-what activities included distribution of articles to staff, conference calls with other organizations, literature reviews, site visits to other organizations, and use of workbooks about new practices. The authors proposed that learn-how was a more effective facilitator of learning because improvement in complex services organizations like hospitals requires attention to understanding

work processes in order to create positive change. These findings, in conjunction with Lapré et al.'s, suggest that the variation in organizational learning curves not only reflects the amount of DLAs that organizations use, but also the type of DLA used. The faster learners in both manufacturing and services settings use more DLAs that generate know-how and know-why. They focus on learn-how.

While research is clear about the nature of the DLAs that are best for learning, it provides little information about the relative effectiveness of specific activities within the broad categories of effective DLA. For example, learn-how consists of a bundle of seven activities [131, 166]. We do not know whether all seven activities contribute equally to the effectiveness of learn-how, or whether the activities are best used in some sequence. We advocate for longitudinal studies to answer the latter question.

### 3.3.2 Contextual Differences

Selecting the “right” DLAs is necessary, but not sufficient to accelerate the learning curve. Organizations must also provide the proper context for deliberate learning. Context refers to the “situational opportunities and constraints that affect the occurrence and meaning of organizational behaviors” such as engaging in deliberate learning [87, p. 386]. Contexts are known to vary widely across organizations, and also across units within organizations [49]. Johns [87] has proposed that contexts vary along important dimensions, often reported in high-quality journalistic stories and typically referred to as the five W's: *who*, *what*, *where*, *when*, and *why*. Research on organizational learning suggest that the effectiveness of what — in this case, using DLAs — is affected by who is involved (i.e., leaders and staff) and their level of investment, where it occurs (e.g., in one or multiple departments), when it occurs (e.g., in conjunction with reflection) and why it has been pursued (e.g., to fulfill efficiency or reliability and quality goals).

The who in organizations generally falls into three broad categories: Senior managers, middle managers or team leaders, and staff. Each must take certain actions for the use of DLAs to translate into better performance. Senior management, for example, “must behave so as

to provide a supportive structure within which people at lower levels can act effectively to improve performance” [181, p. 98]. Evidence of the importance of senior management allowing staff to manage deliberate learning, thereby signaling management support and buy-in for staff effort, was documented in a learning curve study in factories [102]. The researchers observed that use of a “model line” (a production line run as a learning laboratory in which experiments are conducted for improvement) led to substantial productivity improvements in the first plant. Replication of the model line concept on three production lines in two other plants within the same firm, however, failed to meet expectations. A comparison of the model lines showed that senior management, not staff, defined the majority of the projects in the replicate model lines, an indication that senior management did not buy into the model line concept and lacked confidence in this staff-guided learning process. When staff sense a lack of support and interest in their knowledge, they share less of their knowledge and commit less of their effort, which undermines organizational learning and performance [21, 56, 150].

Team leaders, as the individuals most proximate to staff, play a central role in shaping staff responsiveness to deliberate learning challenges. In a follow-up study to the aforementioned study (in the Introduction) of the learning curves for 16 hospitals’ implementation of a new surgical technology, Edmondson et al. [51] found that differences in team leader behavior correlated with teams’ success in implementing the technology. Teams that had leaders that provided a rationale for engaging in learning activities, insisted on staff use of multiple and more intensive DLAs (e.g., dry runs), asked for input, and showcased their fallibility to encourage staff to speak up about problems succeeded, while teams with leaders that did not display these behaviors were less successful. Thus, variance in the behavior of who is involved contributes to variation in the effectiveness of deliberate learning.

Where deliberate learning takes place also dramatically influences the degree of learning. In the aforementioned study of learning via model lines, the researchers found that the three replicate model lines that performed relatively poorly had few projects that were inter-departmental, whereas the successful model line routinely pursued projects which spanned more than three departments. By taking this



approach, the successful model line drew upon a richer, more diverse knowledge base, allowing it to solve more systemic, interdepartmental problems that affected the organization as a whole [102]. The importance of drawing upon multiple locations was also demonstrated in a study of engineers charged with solving problems with new factory machines [167]. In order to solve the problems, engineers left their labs and went to the plant at some point for 78% of the problems studied. For almost all of these problems, the engineers then moved back to the lab for more problem solving. And, for 40% of the problems, they went between the plant and the lab three or more times. Situating themselves in multiple locations enabled the engineers to notice different clues, gather different kinds of data, use different tools and activities, and experience different pressures relevant to the problem. Each of these acts helped them learn about the problem and develop better solutions. Thus, the location(s) of learning and the willingness to span boundaries is a critical determinant of DLA effectiveness.

Even if the location, the individuals, and their behavior is ideal and equal across organizations, learning rates can vary because of the differences when deliberate learning occurs. Studies have shown that deliberate learning is optimized when it occurs in conjunction with reflection. In the course of deliberate learning, staff gather new knowledge. Haas [64] showed that whether the abundant knowledge gathered translates into improved project performance depends on the allocation of time in the workplace. When the project teams she studied had more “slack time” to process and interpret the knowledge they acquired i.e., to reflect, they were better able to translate their new knowledge into improved project performance. She reasoned that this was because staff had the time and attention to focus on developing the understanding to arrive at the best solution as opposed to rushing to locate a satisfactory solution, as is typically the case when staff face high workload or overload [117]. Supporting this logic, Ye et al.’s [183] study of deliberate learning in a health care organization showed that frontline staff advanced in their processing of knowledge from knowledge generation to knowledge articulation when their workload was moderate, allowing them time to reflect without being overwhelmed. Edmondson [49] showed this time for reflection is a critical complement to learning from

deliberate acts. In her study, whether teams focused on strategic planning or product manufacturing, they performed worse if they did not iterate between action and reflection. Neither taking action without reflection (e.g., discussing errors) nor reflecting without taking action led to better performance [49].

Finally, why deliberate learning is being pursued is a factor that appears to influence the impact of deliberate learning. Levin [106] raised the question of whether learning to improve efficiency varies from learning to improve quality. Using data on automobile reliability from *Consumer Reports*, he found that learning to improve quality occurred as a function of “offline” activities to debut new car models not as a function of cumulative experience producing the car. The former led to what he called “exceptional learning,” while the latter led to what he called incremental, everyday learning. Exceptional learning via deliberate offline activities significantly elevated the quality of cars such that debuting car models had lower repair rates measured at three and six years of ownership. Based on these findings, it appears that deliberate learning is particularly well-suited for quality improvement efforts.

In sum, existing research suggests that organizational learning rates vary as a function of the level of deliberate learning. Furthermore, differences in the type of DLA used as well as the context in which it is used alter its impact on organizational learning curves. The greatest positive impact occurs when all organizational participants (senior management, team leaders and staff) support deliberate learning by their actions and when the use of DLAs occurs across multiple locations with time for reflection and with the purpose of quality improvement. Of course, organizations vary in the degree to which their use of DLA aligns with the conditions for positive impact. Hence, the differences we observe in learning curves.

### **3.4 Other Sources of Variation in Learning Rates**

Up to this point, we have discussed what is known about how variation in the two forms of learning believed to underpin the learning curve — autonomous learning via cumulative experience and induced learning via DLAs — affects learning rates. In essence, we have reviewed

what has been learned from taking a micro-approach to understanding the variation in organizational learning. The alternative, macro-approach examines the conditions that shape the learning curve as a whole rather than conditions that influence the sub-components of the learning curve. An unlimited number of macro-factors potentially cause variation in the learning rate. However, research attention has centered on two: Task and organizational characteristics. In this section, we review research findings with respect to both sets of characteristics.

### 3.4.1 Task Characteristics

In the learning literature, tasks have been primarily characterized by the knowledge required for their completion. Scholars have observed that knowledge varies along several dimensions, including for example, tacitness, complexity, observability, provenness, causal ambiguity, and system dependence [142, 159, 180, 185]. However, much of the theoretical and empirical research has focused on how one dimension, tacitness, influences the learning of new practices.

Tacit knowledge has both technical and cognitive elements [133]. Nonaka [133] observed that, the technical elements consist of “the kind of informal, hard-to-pin down skills captured in the term know-how.” It is reflected in a master craftsman who after years of experience develops a wealth of expertise “at his finger tips,” but he is often unable to articulate the scientific or technical principles behind what he knows (p. 98). The cognitive elements consist of “mental models, beliefs, and perspectives so ingrained that we take them for granted and therefore cannot easily articulate them” (p. 98). Thus, both the technical and cognitive elements of tacit knowledge contribute to an inability to easily communicate what is known. This knowledge is often rooted in action or in a specific context, and is often contrasted with explicit or codified knowledge, which refers to knowledge that is transmittable in formal, systematic language [133]. In reality, tacit and explicit knowledge are the extremes of a communicability spectrum, not mutually exclusive categories [138]. According to Nonaka [132, 133], organizational knowledge evolves by shifting between these two extremes. In many instances, what is now explicit was once tacit. Certainly, this is true in

fields such as medicine, where many practices that are now formalized in published clinical guidelines for all to follow were once performed almost unconsciously without a script by just a few [173]. Nevertheless, across industries, there are many tasks that remain largely tacit (e.g., coordinating work across disciplines) or require a combination of tacit and explicit knowledge for task completion (e.g., implementing a new technology).

Organizational theory suggests that the proportion of tacit-to-explicit knowledge in a task explains a significant percentage of the variation in improvement rates for organizations learning to perform the same task. In a direct test of this hypothesis, the researchers who studied the learning curves for 16 hospitals' implementation of a new surgical technology compared the rates of improvement for two tasks related to the implementation [52]. The first task aimed to improve efficiency and required teams to learn new ways of working together. Knowledge about how to work together was not codified in any external sources, thus their task relied on tacit knowledge. The second task aimed to expand the breadth of use of the new technology by having it applied in more instances. As this second task required the application of a well-documented practice in more settings, it relied more on explicit knowledge. When the researchers compared organizational learning rates for the two tasks, they found different patterns: Learning rates were heterogeneous for the efficiency task, which relied on tacit knowledge, and homogeneous for the expanded-use task that relied on explicit knowledge.

The heterogeneity for tacit tasks likely reflects the difficulty of learning these tasks without any guidance. Each organization must learn for itself how to perform these tasks, which inevitably introduces variability as organizations try different approaches. The variability likely also stems from differences in how they contend with other challenging dimensions of knowledge underlying their tasks. Research shows that organizations experience great difficulty transferring best practices (a task that requires learning from another entity) when the practices' underlying knowledge is complex, unproven, or causally ambiguous (i.e., cause-and-effect relationships within the task are unclear) [159, 185].

Too few learning curve studies have incorporated task or knowledge characteristics in their models. This is unfortunate as “better knowledge” and “improving actions” are core elements of organizational learning (as we discussed in Section 1.2) and both elements are likely to depend on task and knowledge characteristics. To the extent models fail to integrate these characteristics, they provide an incomplete picture of the learning curve. Thus, we advocate for studies that explicitly integrate task characteristics, and that consider a broader range of task characteristics rather than limit their focus to tacitness.

### 3.4.2 Organizational Characteristics

A number of empirical studies have identified organizational characteristics that explain differences in organizational learning rates. Sorenson [156], for example, examined the impact of internal structure of the organization. Using time series data on 175 computer firms, he found that the learning rate was lower for vertically integrated firms in stable environments. In these firms, interdependence between business units was high, making it difficult to observe the results of actions, and therefore identify effective and ineffective routines. As a result, learning occurred at a slower pace relative to less interdependent firms, as evidenced by lower sales growth and high exit rates as a function of experience. However, Sorenson’s [156] results do not universally imply that organizations that wish to learn effectively should avoid vertical integration. Firms with this structure in volatile environments actually learned more effectively than their more independent rivals. He attributed this result to the control that vertically integrated firms have over their environment as suppliers of their own inputs. Control provides a buffer against volatility, and enables firms to stabilize their inputs and outputs so they can learn the relationship between them. The effect of organizational design on learning has been supported by other studies as well (e.g., [15, 54]).

Research also highlights organizational capacity and staffing as drivers of learning rates. The effect of capacity on the overall learning curve is the same as its effect on deliberate learning discussed above. More resource-based capacity, e.g., more slack time, leads to higher

learning rates [178] because employees capitalize on available time and resources to find smarter ways to work [114]. Cohen and Levinthal [36] suggest that more absorptive capacity, another component of organizational capacity, increases learning rates as well. Absorptive capacity refers to the “ability to recognize the value of new information, assimilate it, and apply it to commercial ends” (p. 128). Absorptive capacity is greatest when the subject has a broad knowledge base as the ability to absorb new information depends on new information having a relationship to what is already known. Organizations vary in their existing knowledge, and therefore in their absorptive capacity. As a result, they differ in their ability to appreciate new knowledge from experiences and to respond appropriately. According to Cohen and Levinthal [36], heterogeneity in organizational learning arises from this difference.

Staffing configurations can add to the variation. Wiersma [178] studied the impact of temporary employees on 27 organizational units within the Royal Dutch Mail system for three years, and found that units with a higher percentage of temporary employees had higher learning rates. This finding is consistent with the notion that a moderate amount of heterogeneity and novelty in workgroups spurs learning [115]. Temporary employees introduce diversity in workgroups as they tend to approach tasks in different ways based on their experiences working in other firms. Seeing these differences can motivate permanent employees to adopt novel ways of performing their work, including the more effective ways of their temporary coworkers. Several studies have shown that diverse teams tend to be more innovative than homogeneous teams [179]. Their members bring different ways of thinking and doing to the group, providing opportunity for other members to learn new approaches and stimulating the creation of more new approaches.

Differences in expectations and incentives also cause variation in learning curves, as evidenced by Sinclair et al.’s [152] research. In their study, they estimated learning curves linking cumulative past output to unit costs for 221 specialty chemicals produced by a large Fortune 500 company. They found that cumulative past output was related to unit costs through its role in conditioning expectations of future output, which drove incentives to perform R&D. Those products with the highest expected returns became the focus of R&D projects. Those projects

Learning process		Task Characteristics	Organizational Characteristics
Autonomous learning via Cumulative experience:	Induced learning via Deliberate learning activities:		
Specialized v. diversified Success v. failure Individual v. team v. organization	Types of activities used — Conceptual v. operational — Learn-what v. learn-how Conditions under which activities used (i.e., who, what, where, when, and why)	Knowledge properties — Tacitness — Complexity — Provenness — Observability — System dependent	Org. design/structure Staffing configuration Capacity — resource-based (e.g., time) — cognitive-based (e.g., absorptive) Incentives/expectations

Fig. 3.2 Sources of variation in organizational learning curves.

typically addressed bottlenecks for the whole production process, contributing to large performance gains for the organizations. Based on these observations, learning rates differ because organizations have different expectations about their future performance, and in turn create different incentives for staff to engage in learning efforts like R&D. The difference in incentives alters staff behavior and the organizational learning curve in turn.

While the sources of variation in organizational learning curves discussed may seem disparate, Argote et al. [9] argued that the factors that shape learning can be classified into three categories: Those that affect the motivation to learn, those that affect the ability to learn, and those that affect the opportunity to learn. The factors we discussed can all be classified using this framework. Organizational characteristics such as the incentive or reward structure derived from expectations affect motivation to learn. Task characteristics, absorptive capacity, and staffing configurations (e.g., moderate use of temporary employees) affect the ability to learn. Organization design and slack time affect the opportunity to learn. Future research should identify more influential factors in all three categories, and may identify additional categories of influence. It will also be important to examine the interactions among factors in each category. We expect that learning is greatest when motivation, ability, and opportunity for learning are highest. Figure 3.2 summarizes the identified sources of variation in organizational learning curves.

# 4

---

## Relative Effectiveness of Experience vs. Deliberate Learning as Sources of Learning

---

As evidenced above, the variation in learning can be great for many reasons. However, the evidence is convincing that both experience and deliberate learning accelerate organizational learning under favorable conditions. In theory then, organizations have two pathways to better learning and performance at their disposal. The question naturally arises: are the two paths equally effective? If not, which is the more effective path to learning and performance improvement, and when? The answer to these questions interest organizations because there are significant differences in level of resources (time, effort, money, etc.) they must invest for each path [45, 188]. While learning via experience requires little to no investment as it occurs automatically as staff repeat their tasks, deliberate learning requires managers and staff to make a significant investment of resources so staff can have the opportunity to engage in DLAs, including experiments in which failure is likely to occur. If resources were abundant or infinite, we might imagine that organizations would choose to engage in deliberate learning, particularly because a growing number of scholars have argued that organizational learning conveys a sustainable competitive advantage [50, 60]. However, resources are not infinite, and so the financial and



non-financial costs associated with deliberate learning means this choice can only be justified if the benefits are greater than the cost, and the net benefits exceed those associated with experience. In this section, we review research that has examined the relative benefit of experience versus deliberate learning as well as the conditions that govern their relative effectiveness.

#### **4.1 The Path to Optimal Learning: Experience or Deliberate Learning?**

The first study to consider the relative impact of cumulative experience and deliberate learning on learning rates was Adler and Clark's [2] aforementioned study of learning in two manufacturing departments of an electronic equipment company. In addition to showing that the use of different DLAs has different, sometimes contrasting effects on learning, their study showed that the effectiveness of one can depend on the other. Adler and Clark had expected that deliberate learning would diminish the effect of experience. However, when they included DLAs — engineering activity and training activity — alongside experience in learning curve models, much to their surprise, they discovered that these learning activities strengthened the relationship between experience and productivity. They reasoned that because experience was a significant determinant of using DLAs in the plants they studied and because one of the DLAs in turn negatively affected productivity, experience had the stronger effect on performance. It affected performance via two pathways: (1) a positive direct effect and (2) a negative indirect effect on productivity via deliberate learning. Mishina [123] also found this enhanced effect of experience on performance in the presence of deliberate learning, suggesting that experience and deliberate learning are interrelated, however, experience is the primary driver of learning. This hypothesis has yet to be confirmed by quantitative studies that explicitly examine the interaction of experience and deliberate learning in multiple settings, using traditional statistical methods such as interaction terms in regression models. Instead, research has focused on whether one is more beneficial than the other, and therefore should substitute for the other.

Hatch and Mowery [68] compared the two using cumulative production volume and allocation of engineering resources as proxies for experience and deliberate learning, respectively. Their analysis showed that improvement in yield rates in the semiconductor industry was a function of managers' systematic investment of engineering resources to solve problems, rather than cumulative production volume. However, the influence of both factors equalized as processes matured. Ittner et al.'s [83] study of electronic component manufacturers also found no significant difference in benefits derived from each approach to learning. Although past cumulative expenditure on prevention activities and past cumulative volume reduced annual product defects by 13% and 7%, respectively, suggesting that deliberate learning via prevention activities contributed more to firm's quality performance than experience, a *t*-test showed that the difference in effects was not statistically significant. Each approach contributed similarly to firms' current quality levels. Prevention activities included quality planning, developing and maintaining the quality planning and control systems, quality improvement activities, and internal quality improvement facilitation and consulting. Based on these two studies, deliberate learning enhances performance more than experience in manufacturing settings, but not significantly more in the long-term.

Some studies conducted of service organizations suggest the opposite: That deliberate learning can provide substantial benefits beyond experience. Zollo and Singh's [187] study of U.S. bank mergers showed that banks with greater deliberate knowledge codification about bank acquisitions performed significantly better post-merger, while banks' level of acquisition experience had no effect on their performance. Pisano et al.'s [137] study of the learning curves for 16 hospitals' implementation of a new surgical technology also suggested this difference in relative effectiveness. The researchers' qualitative case study analyses of the high- and low-performing hospitals indicated that the difference in learning rates stemmed from differential use of DLAs (i.e., practice session and early trials). However, the researchers were unable to test this hypothesis formally as they did not measure the use of DLA.

Following up on this work, Nembhard and Tucker [131] did measure the use of DLA in a 3-year study of 23 hospital neonatal intensive care units. They found no significant difference between the level of improved workgroup performance associated with a one standard deviation increase in the use of DLA versus cumulative experience (i.e., 18% compared to 20% reduction in patient mortality, respectively) after two years of use. Additionally, they found that, in the short-term, use of DLA increased the odds of worse performance. The researchers drew three conclusions from this set of results. First, organizational learning models for dynamic service settings should include deliberate learning alongside experience and consider how their effects change over time; failure to include both risks providing an incomplete depiction of the organizational learning process. Second, their finding of a worse-before-better effect suggests that using DLAs initially presents a challenge for many workgroups, which often must juggle deliberate learning while fulfilling current production demands. During the initial phases of deliberate learning, they must therefore be vigilant to minimize adverse effects on performance. Finally, the researchers concluded that the similarity in effectiveness of each process means managers have potentially substitutable learning tools at their disposal, which is beneficial in settings where the “paucity of experience problem” (cf. [107]) is prevalent (e.g., dynamic settings). In such settings, managers can circumvent this problem without sacrificing long-term performance by facilitating deliberate learning.

Together, the studies on the relative effectiveness of experience versus deliberate learning provide little consensus on which learning process is more effective. Some studies suggest that experience is more effective, particularly when deliberate learning occurs simultaneously [2]. Others suggest that deliberate learning is more effective [187]. Still others suggest that there is no significant difference in the long-term effectiveness of experience and deliberate learning in either manufacturing [83] or service [131] settings. The difference in relative effectiveness across studies has led to speculation about the conditions under which each is effective. In the next three sections, we discuss the findings related to the conditions that have garnered the most attention: The stage of production, nature of knowledge to be learned, and characteristics of the task.

## 4.2 Depends on Stage of Production

As most of the studies of organizational learning curves have occurred in manufacturing firms, it is natural that scholars have questioned whether the stage of production in the manufacturing process might explain the variability in findings about the effectiveness of experience and deliberate learning. In sequence, the stages of production identified in the learning and process improvement literature are three: Product development, production ramp-up, and full capacity production [162]. However, in research, scholars tend to simply distinguish between early/new stages of production versus late/mature stages of production [45, 68].

Results from the few studies that examine the effectiveness of experience and deliberate learning in relation to production stage consistently indicate that stage of production is a key determinant of relative effectiveness [45, 68, 162]. When organizations are early in production, deliberate learning benefits the organization more than experience. However, the relative benefit changes as production processes mature: The benefit of deliberate learning declines, while the benefit of experience increases. Dorroh et al. [45] demonstrated that this shift in relative effectiveness is robust to changes in a range of parameters values underlying learning curve estimations. When they varied the parameters (e.g., discount rate, production rate, initial knowledge endowment, and resource consumption) in mathematical models of a firm producing made-to-order products, they repeatedly found a shifting emphasis from deliberate learning to experience over the production cycle was optimal for the firm. Their analysis suggested that the shift reflects a change in the value of knowledge generated from each process. Early in production, deliberate learning adds considerably to the knowledge stock; later in production, the same knowledge is less valuable and not as valuable as that gained from experience.

Terwiesch and Bohn [162] further showed mathematically that, under extreme conditions (e.g., high experimentation capability), the optimal strategy during production ramp-up is: 100% deliberate experimentation and 0% production experience to start, and 0% deliberate experimentation and 100% production experience in the end. This

strategy brings to the forefront an inter-temporal paradox that organizations face. Achieving their best performance means they must forego production in the short-term when prices are highest to deliberately learn, so that they may reliably produce at high capacity in the long-term when prices are lower. Paradoxically, their performance is higher if they acquire less experience in the short-term.

Hatch and Mowery [68] verified the shifting effectiveness suggested by their predecessors' mathematical models in an empirical study of yield learning curves for semiconductor manufacturing. Their data showed that yield learning curves for processes in the early stages of manufacturing were driven by cumulative engineering not cumulative volume. For more mature processes, cumulative engineering and cumulative volume were equivalent sources for learning to improve yields. However, the introduction of new processes disrupted the learning of existing (maturing) processes because engineering resources were diverted away to resolve problems with new processes.

### 4.3 Depends on Stage of Knowledge

Both theoretical and empirical research suggests that the optimal learning strategy depends not just on the stage of production, but also on the stage of knowledge characterizing that which is to be learned. Bohn [21, p. 64] 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 [84] and Bohn [21] developed the notion that there are “stages of knowledge.” The stages correlate with how much an organization knows about  $Y = f(x)$ . The more an organization knows, the higher its stage of knowledge.

Bohn [22, 23] observed that there are two dimensions of knowledge that advance through stages — causal knowledge and control knowledge, and that both dimensions have six stages. Table 4.1 depicts the stages of causal knowledge and control knowledge. Causal knowledge refers to knowledge about the relationship between an input and

Table 4.1. Stages of causal knowledge and control knowledge.

Causal knowledge Know how “ $x_i$ affects $y$ ”		Control knowledge Know how to control $x_i$	
Stage	Description	Stage	Description
1. Ignorance	the organization is unaware that $x_i$ might affect $y$	1. Ignorance	the organization is unaware of $x_i$
2. Awareness	the organization is aware that $x_i$ and $y$ are related, but the direction of causality is unknown	2. Awareness	the organization is aware of the existence of $x_i$
3. Direction	the organization knows that $x_i$ affects $y$	3. Measure	the organization is able to measure $x_i$ routinely
4. Magnitude	the organization can quantify the impact of a small change in $x_i$ on $y$	4. Control of the mean	the organization can control $x_i$ at the mean level, but there is significant variation in the level of $x_i$
5. Scientific model	the organization has a functional specification with parameters describing the relationship between $x_i$ and $y$	5. Control of the variance	the organization can control the variance of $x_i$
6. Interactions	the organization has extended stage 5 knowledge to include interactions with all other input variables	6. Reliability	the organization can always keep $x_i$ at its target level

*Note:* Table constructed from notes taken during a presentation by [22].

output, while control knowledge refers to knowledge about how to keep an input variable at its desired level. Regardless of whether the knowledge is causally or control-related the first and last stages of development are conceptually the same. The first stage is “ignorance”; the organization is unaware of the variable or its relationship to output, respectively. In the last stage, the organization has complete knowledge. It completely understands the variable’s relationship to other variables and can reliably control the variable, respectively. In reality, the final stage is never achieved. Most tasks to be learned lie somewhere along the continuum between having a very strong knowledge base and having a very weak one.

In theory, when the underlying knowledge for a task is strong, the organization knows enough about important contributors to task completion that it can design deliberate learning activities outside of the actual operating environment (e.g., laboratory experiments) and be confident that what is learned from them is relevant to the operating environment. Few surprises will occur when they apply lessons learned from their “offline” deliberate learning activities in the final setting. In contrast, when knowledge related to the task is weak, there is a great risk that many lessons learned from “offline” deliberate learning activities are irrelevant or misleading. Yet, the organization is unable to identify these lessons that should not be applied, as it has no basis for elimination due to its limited baseline knowledge. In such circumstances, some (e.g., [135]) argue that the organization may most efficiently and effectively learn through direct experience in the actual operating environment.

Empirical research by Pisano [135] supports the idea that learning from experience is the more effective strategy when knowledge is under-developed, while deliberate learning is the more effective strategy when knowledge is well-developed. In his study, he compared how the use of laboratory experiments, a deliberate learning activity, affected process development lead times in two pharmaceutical environments with different stages of process technology knowledge: (1) Traditional chemical-based pharmaceuticals, an environment characterized by a strong theoretical and practical knowledge base, and (2) biotechnology-based pharmaceuticals, an environment characterized by a more “artistic” than scientific understanding. His analysis showed that greater use of laboratory experiments was associated with better lead times in the strong knowledge environment of chemical-based pharmaceutical. However, their use had no effect on lead times for biotechnology-based pharmaceuticals, leading to the conclusion that:

There is no one best approach to learning (learning-by-doing [via experience] vs. learning-before-doing [via experiments]), but that it depends on the nature of the firm’s knowledge environment. Deep knowledge of the effect of specific variables and their interactions

increases the leverage of research and other forms of learning-before-doing. Learning-by-doing is required when organizations lack the underlying knowledge needed to simulate and predict effects 'off-line'. [135, p. 98]

Despite Pisano's observation and findings that knowledge influences the effectiveness of different learning strategies, subsequent learning-curve research has not extended the study of the relationship between the stages of knowledge and learning. No empirical research has moved from the dichotomous view of stages in Pisano's research (early/under-developed vs. mature/well-developed knowledge) to investigating the implications of the more nuanced, six stages of knowledge in Bohn's [22, 23] framework. The absence of this research likely reflects the difficulty of identifying the six stages in practice. Additionally, resources may be a limiting factor as researchers may not have the resources to monitor the organization over the extended time required to see the shift between the six stages. Nevertheless, we encourage further study of the relative effectiveness of experience and deliberate learning at different stages of knowledge. A starting point could be the study of specific stages. Even for those without interest in this topic, the importance of the stage of knowledge suggests that researchers should be mindful of the stage in their studies and report this information to help others understand the context of their research.

#### **4.4 Depends on Task Characteristics**

Zollo and Winter [188] have theorized that the relative effectiveness of experience and deliberate learning depends on the features of the task to be learned. They argue that the returns to investment in deliberate learning relative to experience are obvious for some task characteristics. For example, tasks with high economic importance should benefit from a relatively higher investment in deliberate learning because the investment builds competence performing these tasks, which diminishes future failures and increases economic success. Likewise, tasks with a larger scope, involving multiple groups, departments, or even all segments of the organization to be performed effectively, should benefit



more from deliberate learning because DLAs provide opportunity for individuals to understand the full scope of the task, not only their fraction of the task. When individuals understand how their task contributes to the larger task, they tend to be more committed to organizational goals and more savvy in assessing opportunities for improving their actions in service of the larger task, which ultimately improves organizational performance. Thus, organizations clearly benefit from relatively higher investment in deliberate learning for certain tasks.

However, the benefit of deliberate learning relative to experience (or experience relative to deliberate learning) is less obvious for other task characteristics, such as task frequency, task heterogeneity, and causal ambiguity within the task [188]. Task frequency refers to how often the task must be performed within a specific period of time. Task heterogeneity refers to how new the task appears each time to the individual or group that has to perform it. Causal ambiguity refers to how easy it is to derive an understanding about what should or should not be done to perform the task. Instinctively and because, in practice, deliberate learning is often lower when task frequency, homogeneity (as opposed to heterogeneity), and causal clarity (as opposed to causal ambiguity) exist, it may appear that experience is relatively more effective in these situations. However, Zollo and Winter [188] argue that the opposite is true: deliberate learning is more effective relative to experience for tasks characterized by low frequency, high heterogeneity, and high causal ambiguity.

As they explain, for tasks with low frequency, at least three factors contribute to the superiority of deliberate learning over experience: Limited memory, coordination costs, and opportunity costs. Unlike tasks with high frequency, for which individuals naturally develop memories of effective and ineffective practices, tasks with low frequency provide limited occasion for the development of individual memory, which is the basis of learning from experience. The absence of this basis hinders learning from experience. Memory, however, is not a basis for the effectiveness of deliberate learning; deliberate learning enables knowledge acquisition in the absence of memory, making it more effective for infrequent tasks. It should also be more effective, given the coordination costs of low frequency tasks. As frequency decreases, deliberate

learning becomes less costly to coordinate. Furthermore, the opportunity costs decline as rarely must resources be diverted from routine operations to deliberate learning.

Zollo and Winter [188] note that if one compares studies of tasks occurring with high, moderate, and low frequency, there is evidence to support their hypothesis. In studies of quality improvement projects, which occur frequently, the effect of experience and deliberate learning is strongly positive and comparable [125]. Studies of acquisitions [186] and alliances [89], which occur with moderate frequency, show that deliberate learning contributes strongly to performance, while experience contributes only weakly. In studies of reengineering, a process which occurs with low frequency, deliberate learning has a positive effect on current performance, while experience has no effect [174]. Thus, the evidence suggests that as task frequency decreases, moving from high (e.g., quality improvement projects) to moderate (e.g., acquisitions) to low (e.g., reengineering processes), the relative effectiveness of deliberate learning to experience appears to increase, as experience diminishes in effectiveness along this spectrum.

According to Zollo and Winter, the relative effectiveness of deliberate learning is also greater when there is high task heterogeneity. They reason that:

“the hazards of inappropriate generalization can only be attenuated via an explicit cognitive effort aimed at uncovering the interdependence between the dimension(s) of heterogeneity and the action–performance relationships. For example, a firm that has made several acquisitions in a wide variety of sectors will probably find it more difficult to extrapolate rules of conduct in managing acquisition processes, compared to another one that has consistently acquired in its own domain. The former might find it comparatively more useful to invest in debriefing sessions and in detailed postmortem analyses as opposed to simply relying on its group of M&A experts. The need to understand what works and what doesn’t in the different contexts experienced

requires an explicit investment in retrospective sense-making” [i.e., deliberate learning] (p. 348)

Empirical studies of the effect of specialized versus diversified experiences on learning rates (Section 3.2.1) supports their reasoning. Recall that these studies document an inverted U-shape relationship between exposure to variety and performance. The declining portion of the curve may reflect the difficulty of learning from experience when tasks are so diverse. Lessons may be inappropriately drawn and erroneously applied, harming performance — unless effort is made to deliberately learn. Indeed, Zollo and Singh [187] showed that knowledge codification processes are strongly related to performance under these conditions, and more effective than experience.

Finally, tasks with a high degree of causal ambiguity are also believed to benefit more from deliberate learning than experience. In these tasks, it is difficult to detect cause-and-effect relationships due to the number and degree of interdependence among sub-tasks. The high degree of cognitive effort inherent in deliberate learning should facilitate identification of the relationships and application of lessons learned once the relationships are understood. The tacit process of experience is unlikely to remove ambiguity as effectively and efficiently. Empirical research is still needed to test this hypothesis. In general, the field would benefit from greater study of the interaction of task characteristics and learning, particularly in service settings, which have been under-studied relatively to manufacturing settings. It would be interesting to know how the levels of task interdependence between employee and customer, a key characteristic of service settings, alters the effectiveness of different learning strategies, for example. Additionally, we encourage researchers to examine other conditions beyond the three identified in research thus far — stage of production, stage of knowledge, and task characteristics — that might influence the relative effectiveness of experience and deliberate learning. Environmental conditions could be another determinant.

# 5

---

## Moving from Learning to Performance: Steps Inside the Learning Curve

---

Although the debate about the relative effectiveness of different learning strategies — experience and deliberate learning — continues, there is consensus that learning can alter performance. Most of the studies we have discussed up until this point, and in fact, most of the studies of organizational learning curves have focused on documenting the statistical link between learning and performance. Few studies have delved into the process by which learning results in improved performance, despite calls for such research and Bohn’s [21] provision of a framework for thinking about this process (see Section 3.1.3). In Bohn’s framework, learning, which occurs as a result of experience and deliberate activities, is only a starting point on the path to improved performance. To improve performance, learning must first result in the development of better organizational knowledge (step 1), which motivates changes in behavior (step 2), which ultimately contributes to improved cost and quality performance (step 3). Thus, there are potentially three steps “inside the learning curve” (Figure 3.1).

In this section, we discuss the research that has sought to provide better understanding of what happens inside the learning curve. This research supports Bohn’s conceptualization of the learning-performance

relationship as a multi-step process mediated by better organizational knowledge (i.e., cognitive change) and behavioral change. Therefore, we organize our discussion using Bohn's framework.

## 5.1 From Learning to Better Organizational Knowledge

Like Bohn, several scholars have theorized that effective learning behaviors create better and relevant organizational knowledge, which positively impacts performance (e.g., [58, 170]). However, only two studies have empirically considered the relationship between learning behaviors and knowledge creation. In the first, mentioned earlier (Section 3.3.1), Lapré et al. [99] built on Mukherjee et al.'s [125] finding that the use of DLAs can lead to two types of learning — conceptual learning and operational learning. As discussed, conceptual learning refers to the process of acquiring a better understanding of cause-and-effect relationships, while operational learning refers to the process of obtaining validation of action-outcome links. Using this distinction, Lapré et al. [99] classified quality improvement projects into four different types of knowledge creation projects based on whether they were high or low on each type of learning given their choice of DLAs. They then incorporated (cumulative number of) projects by knowledge type in their learning curve model for waste reduction in a tire cord plant. When they did that, they found that only one of the four types of the learning-derived knowledge projects contributed positively to performance: operationally validated theories. These projects were high on both conceptual and operational learning, enabling project teams to acquire both know-why and know-how. This bundle of knowledge overturned myths (i.e., erroneous knowledge) and validated new theories of operations. The researchers observed that “the conceptual learning enhanced the build-up of deep process knowledge and codification of the results obtained with operational learning” (p. 607), which facilitated the transfer of solutions across settings and the replication of positive results. The other three types of knowledge projects had no or poor effect on performance, as the underlying learning activities failed to develop better knowledge. The solutions they produced were ill-understood or non-validated when implemented. This Lapré et al.

[99] study is the first (and only to date) learning curve study to link learning and knowledge (via the assignment of knowledge types based on learning types), and to utilize longitudinal variables for learning and knowledge.

Choo et al. [33] explicitly examined the link between learning behaviors and knowledge creation in a cross-sectional study of organizational learning in a Fortune 500 manufacturing firm. In their study, they measured the learning behaviors used in 188 Six Sigma projects as well as the level of knowledge created in the course of each project, with level of knowledge creation measured by the degree of solution uniqueness, idea generation, and improved understanding and capability of the team members after the project was completed. As they hypothesized, the use of learning behaviors increased the level of knowledge created by teams. Furthermore, the results of their structural equation model showed that the knowledge created mediated the relationship between learning behaviors and project performance. In other words, their models showed that the use of learning behaviors enhanced knowledge creation, which in turn improved project performance. Additionally, they found that adherence to a structured method of problem-solving (i.e., Six Sigma's Define–Measure–Analyze–Improve–Control (DMAIC) method) served as an antecedent to knowledge-creating learning behaviors. Psychological safety — the belief that it is safe to take interpersonal risks such as asking questions, seeking feedback, reporting a mistake, or proposing a new idea [48] — also directly contributed to knowledge-creation. Choo et al. theorized that adherence to a structured method altered how information was acquired and therefore directly affected learning behavior, while psychological safety allowed the team to freely explore opportunities for improvement, leading to the creation of new knowledge.

Both Lapré et al. [99] and Choo et al. [33] provide solid evidence that learning is associated with knowledge creation. Lapré et al. [99] also indicates that not all learning leads to better organizational knowledge and performance. In three out of four instances, the knowledge acquired from learning led to the same or lower performance rates. More studies are needed to understand the situations in which learning hurts knowledge development, and the moderators of that effect.

These studies, like Lapré et al.'s [99] are advised to include longitudinal measures of learning and knowledge. Longitudinal studies may show opportunities to shift from learning that fosters poor knowledge to learning that fosters better knowledge.

## 5.2 From Better Organizational Knowledge to Changed Behavior

Neither of the two studies that examined step 1 inside the learning curve examined step 2: The transition from better organizational knowledge (a cognitive change) to behavioral change. However, related and follow-up studies have suggested that this transition occurs and is important. The Mukherjee et al. [125] study that preceded the Lapré et al. [99] study showed that both conceptual and operational learning, which implicitly led to better organizational knowledge, altered improvement project teams' ability to change behavior. The more conceptual and operational learning that occurred during a project, the more able the team was to change attention to the organization's rules, as evidenced by modifications in standard operating procedures and statistical process control rules. The researchers observed that conceptual learning helped to overturn the myths held by staff, which set the stage for changes in practices. Operational learning then provided staff with further evidence of the need for change as it allowed staff to witness changes in variables of interest as a direct result of practices. Thus, both types of learning, and the knowledge they generated, provided the foundation for behavioral change.

Tucker et al. [166] showed that not only does the ability to spur behavioral change grow as a result of learning (and better knowledge), but also that actual behavioral change occurs as well. In their aforementioned study of the use of DLAs in hospital neonatal intensive care units, they found that improvement teams' use of learn-how — “activities that provide an understanding of why a practice works, as well as how to carry it out” (p. 898) — was positively associated with implementation success of new practices in the unit. Higher levels of published evidence supporting new practices, an indicator of the quality of knowledge that had been learned previously, also led to

greater implementation success. Implementation success was defined as “employee commitment to and consistent use of new practices” (needs ending quote mark) [94]. Such use is a significant behavioral change in health care, an industry in which implementation failure is common [128]. More research is needed on the behavioral changes that occur in other industries as a result of better organizational knowledge. We suspect that the behavioral changes occur in stages as learning progresses. Therefore, we believe that much would be gained from research that examines the evolution of behavioral change in learning organizations. Qualitative studies may be necessary to uncover the patterns of change. The psychological literature on behavioral change may be a useful starting point for theorizing as well.

### **5.3 From Changed Behavior to Organizational Performance**

In a follow-up study to Tucker et al. [166], Nembhard and Tucker [131] investigated the relationship between use of DLAs, changed behavior, and organizational performance in the same neonatal intensive care units. As noted earlier (Section 4.1), their three-year study showed a worse-before-effect of using learn-how on organizational performance, as measured by NICUs’ risk-adjusted mortality rates for 2159 infant-patients. More notably, with respect to step 3 inside the learning curve — the shift from changed behavior to organizational performance — their study showed that a change in a critical workgroup interaction mediated the relationship between learn-how and performance. Specifically, their data showed that the use of learn-how fostered interdisciplinary collaboration, which in turn improved NICUs’ risk-adjusted mortality rates. In health care, improved interdisciplinary collaboration is a notable behavioral change as several seminal reports have documented the frequent absence of interdisciplinary collaboration in workgroup interactions, resulting in quality problems (Institute of Medicine [82]). The need for greater interdisciplinary collaboration is present in other industries as well. A McKinsey [119] survey showed that 75% of executives did not believe that effective collaboration occurred, although it is important for organizational performance. Nembhard and Tucker’s [131]



findings indicate that over time, learning activities can facilitate the interdisciplinary collaboration needed for performance improvement.

Based on the authors' theorizing, there are three explanations for why the behavioral change they studied, increased interdisciplinary collaboration, could translate into improved organizational performance. First, they note that prior research (e.g., [151]) shows that interdisciplinary collaborators make better decisions because they openly share their pertinent expertise, raise relevant questions, consider alternatives more fully, and integrate ideas across disciplines to enrich the decision choice set. These actions typically lead to higher quality decisions that ultimately improve performance [86, 112]. A second, implied explanation for improved performance is that interdisciplinary collaboration improves coordination, i.e., the integration of different pieces of the task to accomplish the collective goal [169]. Through collaboration, staff develop transactive memory about "who knows what" [110, 139], allowing them to coordinate their tasks more efficiently and effectively, which improves overall organizational performance. Third, collaborators are skilled at detecting and learning from errors, enabling them to respond sooner to minimize adverse effects on performance [177].

The idea that performance improves as a direct result of changed behavior has also been supported by a four-year study of a unionized auto parts manufacturing plant with approximately 1300 workers. In this learning-curve study, Arthur and Huntley [14] found that the cumulative number of implemented employee suggestions significantly contributed to lower production costs. The implemented suggestions originated from cost reduction ideas submitted by staff involved in the plant's gainsharing program. Implementing the ideas required changes in staff behavior that ultimately led to better cost performance.

#### **5.4 Challenges to Advancing**

While the steps inside the organizational learning curve may seem so logical as to occur naturally, ample studies suggest that organizations experience difficulty progressing from learning to improved performance. A review of these studies suggests that their efforts are thwarted by at least four sets of factors: Psychological and sociological factors

that govern the interactions among those that must learn together, cognitive factors that shape learning capacity, complexity, and the nature of learning itself as a multi-level process. Additionally, the context in which the organization exists can challenge learning, as we discussed in Section 3.3.2. In the remainder of this section, we highlight the factors that challenge the organizational learning process within the first four categories mentioned.

#### **5.4.1 Psychological and Sociological Factors**

Despite learning's potential to improve performance, many individuals fail to engage in activities that would help their organization learn, even when they are well-aware of the need for organizational improvement [165]. Their hesitation reflects an observation about organizational learning. It requires engaging in behaviors that can be interpersonally risky. In particular, it requires speaking up about problems and opportunities for improvement, collaborating with individuals with diverse backgrounds and expertise, and participating in experiments, which are likely to fail [130]. These behaviors most occur when individuals feel psychologically safe, i.e., when they believe that negative consequences (e.g., criticism, punishment, embarrassment) will not follow from their interpersonal risk-taking [48]. Unfortunately, psychological safety in the workplace is often illusive due to the nature of work and workgroups. Status and power hierarchies, the temporariness of work teams, and the virtuality of work, for example, all contribute to lack of psychological safety [130]. Lack of psychological safety, in turn, is associated with lower staff participation in learning behaviors [48, 129, 166]. Thus, a key reason that organizations do not progress along the learning curve is that low psychological safety stifles the willingness to engage in step 1 of the process, learning.

Even in psychologically safe workplaces, other psychological factors can limit learning. Humans have a deep psychological aversion to change and failure [144]. Thus, they inherently avoid activities that might lead to failure, such as experiments and other deliberate learning activities. Additionally when failures occur, they often dismiss, deny, or distort them mentally in order to preserve self-image and self-esteem,

rather than treat them as learning opportunities [144, 161]. In general, individuals and groups have the tendency to stifle broader awareness of problems [176] and avoid the kind of problem-solving that results in organizational learning. Tucker et al. [165] found that, even in hospitals where one might think that a learning orientation is the norm, only 8% of nurses' responses to problems they encountered aimed to start the process of organizational learning to arrive at a systematic solution (called "second-order problem-solving"). Instead, most nurses worked around problems and did not alert supervisors (called "first-order problem-solving"). This behavior occurs outside of health care as well. Over 85% of the managers and staff in a general survey admitted to remaining silent about a concern at work [121], denying the organization a learning opportunity.

Organizations also miss opportunities to learn due to social pressure, the use of traditional conflict management strategies, and competency traps. Social pressure for conformity creates a bias toward the leader's perspective rather than engaging in collaborative inquiry to arrive at a solution [85, 140]. Furthermore, leaders' can inadvertently stifle dissent, a learning opportunity, by simply expressing their opinion. They can also stifle interest in learning by being uninviting and unresponsive to suggestions [129]. Individuals that disagree rarely ask each other the questions necessary for them to learn from each other [11]. Instead, they tend to force their view on other party. Furthermore, when the stakes are high or the situation is ambiguous, individuals and groups tend to under-react and fall into competency traps [107], making the flawed assumption that current organizational routines are preferable to alternatives, as we discussed in Section 2.2.2. Consequently, there is no need to learn new routines.

While the psychological and sociological obstacles to learning are many, they are not insurmountable. Prior research shows that leaders through their actions can help organizations overcome the obstacles. When leaders are supportive of staff — demonstrating inclusiveness, seeking and appreciating others' input, making themselves available, displaying fallibility, and providing needed resources — they create an organizational climate in which staff members feel psychologically safer and more willing to engage in organizational learning [129, 130]. The

same is true when the organization as a whole sanctions and provides opportunities for learning, offers support and encouragement to overcome the fear and shame associated with making errors, rewards efforts in the right direction, cultivates norms that legitimize the making of errors, and promotes norms that reward innovative thinking and experimentation [12, 144].

### 5.4.2 Learning Capacity

Even when an organization provides a fertile environment, learning can be hindered by the organization's innate capacity for learning and capacity to retain what has been learned [9]. As suggested in Section 3.4.2, an organization's capacity for learning is a function of its resource and absorptive capacities. In turn, resource capacity is a function of the time, human capital, technology, and monetary resources available to the organization. Absorptive capacity is a function of the organization's existing knowledge base. When the organization has a broad knowledge base, its absorptive capacity is greater because it is able to appreciate a large scope of new knowledge by drawing associations to what it already knows [36].

For those organizations with limited resource and/or absorptive capacity, less is likely to be learned and learning likely occurs at a slower pace. It is possible that some might overcome this disadvantage by importing capacity (e.g., hiring individuals with expanded or different knowledge bases), organizing staff into interdisciplinary work teams to promote creative problem-solving, or fostering organizational identification in order to cultivate organizational citizenship behavior, i.e., staff's positive, extra-role behavior in which they draw upon their own resources to perform "additional tasks" such as helping the organization to learn. However, the effectiveness of these proposed strategies have yet to be examined empirically.

For those organizations that do have the capacity to learn and do so, another challenge can limit their ability to translate learning into improved performance: Knowledge depreciation or forgetting. Several studies have shown that knowledge acquired through experiential learning depreciates rapidly, meaning that the value of knowledge

learned declines as time passes. For example, Darr et al. [43] found that knowledge depreciates at a rate of 17% per week for pizza franchises, implying that “roughly one half of the stock of knowledge at the beginning of month would remain at the end of the month.” Argote et al. [6] found a depreciation rate of 25% per month for construction of Liberty cargo vessels during wartime. A follow-up study of the same vessels by Thompson [164] confirmed that significant knowledge depreciation occurred, but found a slower rate (3.6–5.7%).

While rates vary across settings and depending on methods of calculation, the evidence is robust that the value of knowledge learned from experience depreciates. Arthur and Huntley [14] showed that knowledge acquired through deliberate learning depreciates as well. Several explanations have been offered for this phenomenon, including the occurrence of technological advancements that make past knowledge less relevant, the tendency for individual forgetting, failure to codify knowledge in organizational memory systems, ineffective knowledge management systems, and employee turnover [16, 164]. Only the correlation to employee turnover has been studied, and the supporting evidence is mixed [6, 8]. Regardless of the explanation, the fact remains that knowledge depreciation exists. For comprehensive reviews of the research on this topic, please see Argote [4] and Argote et al. [9].

Given that learning depends on absorptive capacity and absorptive capacity depends on the stock of knowledge, knowledge depreciation presents a challenge for organizations as it limits the stock of knowledge for learning. The implication is that organizations must continually learn in order to build the stock of knowledge for learning. In other words, they must become learning organizations, as will be discussed in Section 6.1.3.

### 5.4.3 Complexity

Several complexity barriers further impede organizational learning. *Detail complexity* arises when the presence of too many variables makes it difficult to comprehend a problem in its entirety. *Dynamic complexity* arises when distance and time make cause-and-effect difficult to establish. Senge [149] argues that the real leverage in most

management situations lies in understanding dynamic complexity, not detail complexity:

When the same action has dramatically different effects in the short run and the long, there is dynamic complexity. When an action has one set of consequences locally and a very different set of consequences in another part of the system, there is dynamic complexity. When obvious interventions produce nonobvious consequences, there is dynamic complexity... [In an enterprise] it takes days to produce something, weeks to develop a new marketing promotion, months to hire and train new people, and years to develop new products, nurture management talent, and build a reputation for quality—and all of these processes interact continually (pp. 71–72).

*Incomplete technological knowledge* is the lack of a full understanding of the effects of input variables of a process on the output [21]. Incomplete technological knowledge is quite common [31, 84, 91, 103]. Faced with detail and dynamic complexity, ambiguity, and incomplete technological knowledge, individuals in an organization can create their own potentially inaccurate beliefs, or “myths,” based on subjective interpretations of events [74]. Misperception of feedback along with poor inquiry and scientific skills can make it very challenging to overturn false myths [158]. The challenge is further aggravated when the myths are held by powerful people [46].

#### **5.4.4 Multilevel Process**

As we discussed in Section 1.3, there are three levels at which learning occurs in organizations: Individual, team, and organization. Given that learning occurs at multiple levels in organization, its effectiveness is susceptible to factors at multiple levels. At the individual level, for example, individuals’ knowledge and experiences can facilitate or hinder learning. At the group level, interpersonal dynamics among group members and group norms, for example, can have an effect. At the

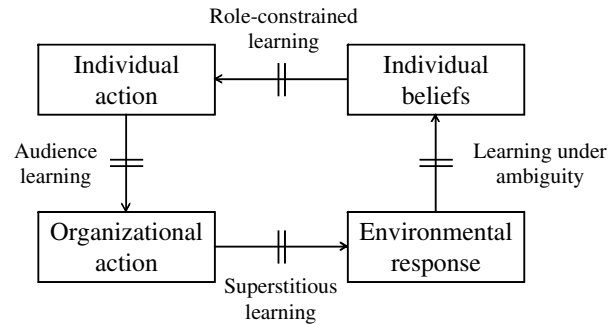


Fig. 5.1 Incomplete learning cycles.  
*Note.* Adapted from March and Olsen [116].

organizational level, characteristics such as those we discussed in Sections 3.3.2 and 3.4.2 (e.g., organizational structure and design) can alter the progression from learning to improved performance inside the learning curve.

March and Olsen [116] labeled four of the core challenges to moving from one level in the process to the next. Figure 5.1 shows the four challenges. First, *role-constrained learning* can occur, which occurs when individual learning has no effect on individual action. “One of the conspicuous things about complex organizations (or any complex social structure) is their ability to inhibit the modification of individual behavior on the basis of individual learning” [116, p. 158]. Constraining role definitions and standard operating procedures can contribute to role-constrained learning [74]. A second challenge, *audience learning*, occurs when individual action does not affect organization action. Organizational politics and general inertia can lead to audience learning. A third challenge, *superstitious learning*, happens when the link between organizational actions and environmental responses are ambiguous: “Organizational behavior is modified as a result of an interpretation of the consequences, but the behavior does not affect the consequences significantly” [116, p. 159].<sup>1</sup> In the fourth challenge, *learning under ambiguity*, it is not clear what happened or

<sup>1</sup>See also Levitt and March [107].

why it happened. The ambiguity stems from the simultaneous existence of equally plausible but mutually contradictory explanations of a situation. Thus, organizations face the significant challenge of simultaneously managing the many factors that shape organizational learning at multiple levels.

Vera and Crossan [171] proposed that strategic leadership is needed to manage the multi-level organizational learning process. Specifically, they identify the need for a leadership that combines transformational and transactional leadership behavior. Transformational leadership is charismatic, inspirational, intellectually stimulating, and individually considerate. This form of leadership inspires and excites individuals to participate in the accomplishment of the organizations' goals. In contrast, transactional leadership relies on contingent-reward exchanges and active management-by-exception. Vera and Crossan [171] propose that transformational leadership has a positive effect on the progression of learning from individual to organizational (and vice versa), when learning challenges institutionalized practices. In contrast, transactional leadership has a positive effect on the progression of learning when learning reinforces institutionalized practices. Their propositions have yet to be tested. Nevertheless, they highlight the difficulty of learning due to its multi-level nature. Managers must select the appropriate leadership style to facilitate the progression of learning across levels. It can be difficult to identify which style is needed.



# 6

---

## The Next Frontiers in Organizational Learning Curve Research

---

As discussed in the previous section, there are many challenges to advancing from learning to performance. However, successful management of those challenges and the process inside the learning curve more generally can yield a compelling competitive advantage. Thus, as we said in the Introduction, there is an imperative to advance our understanding of the organizational learning curve. In this section, we discuss some of the many opportunities for future research. We introduce research streams whose integration with organizational learning curve research is likely to yield significant insight. In Section 6.1, we propose research that extends the work on knowledge creation in organizations, as the creation of better knowledge is the first step inside the learning curve. We highlight the need for greater study of (i) how the stages of knowledge impact the effect of learning on performance; and (ii) learning by experimentation, as this is a primary vehicle for generating better knowledge. In Section 6.2, we suggest that another frontier for exploration is the development of the “learning organization,” i.e., the context that allows the steps in Figure 3.1 to occur effectively. Section 6.3 suggests extending the performance measures used in learning curve research beyond cost and quality to capture

additional dimensions of performance relevant to organizations. Finally, Section 6.4 concludes by encouraging research that looks at the likely situation of organizations needing to learn to improve more than one measure of performance.

## **6.1 Knowledge Creation**

We have observed that there is huge variation in learning rates. A significant part of the variation stems from organizations having incomplete knowledge about their operating systems. Unfortunately, organizations face many obstacles as they try to create better organizational knowledge. A useful way for organizations to gauge progress in terms of creating better knowledge is the concept of “stages knowledge” discussed in Section 4.3 [21, 22, 23, 84].

### **6.1.1 Stages of Knowledge**

Table 4.1 shows the stages of causal knowledge and control knowledge. The two dimensions of causal and control knowledge closely mirror the dimensions of the learning process identified by Mukherjee et al. [125]. 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. [99] demonstrated the significance of incorporating the two learning dimensions in learning-curve estimation. Moreover, the authors showed the importance of balancing both conceptual and operational learning. It would be a major contribution to include longitudinal progress on the stages of knowledge in learning-curve estimation. Furthermore, it would be beneficial to determine: 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, i.e., should causal knowledge and control knowledge progress at the same pace? What is the impact of climbing the stages of knowledge for primary variables versus secondary variables? Primary variables are variables that directly impact output. Secondary variables are variables that directly impact primary 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 their progress along stages. Ittner et al. [83], for example, found a research site that measured four stages of quality-based learning: “Aware of need”, “process characterized and sources of variation identified”, “critical process parameters understood”, and “knowledge institutionalized”. Additionally, in a longitudinal field study of an electromechanical motor assembly plant, Field and Sinha [56] found that actions to control the mean (control knowledge stage 4) preceded actions to control the variance (control knowledge stage 5), as predicted by the stages model. Nevertheless, scholars have yet to quantitatively measure progress on the stages of knowledge and include such measures in learning-curve estimations. To do so, would be to take the next logical step following on the sequence of studies by Levy [108], Adler and Clark [2], Pisano [135], and Lapré et al. [99].

### 6.1.2 Learning by Experimentation

Deliberate learning activities, or induced learning, can significantly enhance the rate of learning. As results by Adler and Clark [2], Lapré et al. [99], Lapré and Van Wassenhove [102], and Nembhard and Tucker [131] have shown, deliberate learning activities need careful management in order to have the desired effect of accelerated learning. Failure to carefully manage deliberate learning activities can actually harm the learning rate. Typical examples of deliberate learning activities are quality and productivity improvement projects. Such projects often rely on a series of experiments. Bohn and Lapré [24, p. 2] define an experiment as “a deliberate comparison of outcomes from a varied but repeatable set of conditions, with an effort to explain different outcomes by differences in conditions.” The authors identify four types of experiments (pp. 8–10):

- *Controlled experiments* make deliberate changes to treatments for several groups of subjects, and compare their properties.
- *Natural experiments* use normal operation as the data source. Box et al. [26] called natural experiments “happenstance data.”

- *Ad hoc experiments*, like controlled experiments, use deliberate changes. However, the changes are made without a careful control group or experimental design.
- *Evolutionary operation* experiments are a hybrid between controlled and natural experiments. Slight changes in the process are made deliberately, and the resulting changes in results are measured and statistically associated with the process shifts. The changes are small enough that the process still works and all the output is still good. Subsequent changes can move farther in whichever directions give improved results — this is the “evolution.”

There is no established theory specifying which types of experiments should be used in any given environment. Bohn and Lapré [24, pp. 11–13] identify five high-level criteria to evaluate different types of experiments:

- *Speed* is the inverse of the information cycle time, the time that elapses from beginning to end of each Plan–Do–Study–Act cycle.
- *Signal-to-noise ratio* is the ratio of the true (unknown) effect of the experimental change to the standard deviation of measured outcomes.
- *Cost* per cycle. Costs can be financial (for materials used) or non-financial (opportunity cost for labor and equipment involved).
- *Value and variety of the underlying ideas* being tested. Better ideas improve the signal-to-noise ratio and increase the benefit from the new knowledge.
- *Fidelity* of the experiment is the degree to which the experimental conditions emulate the world in which the results will be applied.

As the authors discuss, experimentation typically involves trade-offs between these criteria. For example, compare experimentation in a controlled laboratory setting versus the full-scale manufacturing environment of a factory. Speed and signal-to-noise ratio are typically higher in a laboratory, whereas fidelity is higher in a factory. Lastly,

experimental designs vary depending on the current stages of causal knowledge. Given the central role of experiments in deliberate learning activities, future research is needed to figure out optimal experimentation strategies. Mindful of the five aforementioned criteria and the existing knowledge base, what kind of experimental design needs to be performed where? What should organizations do to increase the value and variety of the underlying ideas being tested? What experimentation strategies allow an organization to climb the stages of knowledge faster? We have seen mixed results for DLAs. Poor experimentation strategies have likely contributed to disappointing results from DLAs. Better understanding of sound experimentation strategies could well be the single most important component for accelerating organizational learning curves.

## 6.2 Development of the Learning Organization

Garvin [59, p. 80] defined a *learning organization* as “an organization skilled at creating, acquiring, and transferring knowledge, and at modifying its behavior to reflect new knowledge and insights.” The idea of a learning organization is appealing because it portends the existence of an organization able to sustain high performance in changing conditions, a characterization that has come to describe most organizational environments. Despite this merit, the ideal of the learning organization has not been realized [60]. The absence of learning organizations has led scholars like Hackman and Wageman [66] to question “why does a learning orientation not blossom even in TQM organizations?” Why is it so difficult to create a learning organization? Is it even possible to create a learning organization?

Stata provides a clue that it is in fact possible to create a learning organization despite all of the challenges to learning we discussed in Section 5.4.3. Reflecting on Analog Devices’ efforts to become a learning organization, he noted that:

We have found that the best way to introduce knowledge and modify behavior is by working with small teams that have the power and resources to enact change. For example, quality improvement training

starts with the division manager and his or her direct reports. The group not only develops a common understanding of new concepts and language, but peer pressure can also help to bring along skeptics who might otherwise block progress. Moreover, the new knowledge can be immediately transformed into action as an integral part of training [157, p. 70].

Inherent in his observation are at least two principles for changing organizational behavior to emphasize organizational learning: (1) Start with empowered teams to capitalize on their shared knowledge and peer pressure and (2) focus on training. Still, the question remains: how does an organization become “skilled at creating, acquiring, and transferring knowledge, and at modifying its behavior to reflect new knowledge and insights” [59, p. 80]? Are teams and training sufficient to transform an organization into a learning organization? Garvin et al. [60] suggest that there are more building blocks for the learning organization. In particular, they identify three: (1) A supportive learning environment (which includes psychological safety, appreciation of differences, openness to new ideas, and time for reflection), (2) the presence of learning processes and practices (including experimentation, information collection, analysis, education and training, and information transfer), and (3) leader behavior that provides reinforcement of learning. The Learning Organization Survey ([los.hbs.edu](http://los.hbs.edu)) is a survey tool intended to help organizations assess their performance on these building blocks. The architects of the survey note that:

Organizations do not perform consistently across the three blocks, nor across the various subcategories and subcomponents. That fact suggests that different mechanisms are at work in each building-block area and that improving performance in each is likely to require distinct supporting activities [60, p. 110].

Little is known about the mechanisms at work in each building-block area. Therefore, an important scholarly and practical contribution would be the identification of the different underlying mechanisms

that govern the development of these building blocks. With greater understanding of the mechanisms, we should gain a better understanding of which behaviors, and in what sequence, result in the development of the learning organization. Additionally, the field would benefit from greater research on the behaviors that enable organizations to overcome the formidable challenges to advancing from step to step inside the learning curve that we discussed in Section 5.4.3.

Providing insight into the behavioral changes necessary to become a learning organization is a critical next step in the evolution of organizational learning research. The next frontier is to provide greater insight into sustaining the learning organization. Part of sustaining this organization is sustaining the knowledge it acquires, as knowledge is the basis for better current practice and the basis for better future practice. Thus, another important area for future research is how organizations manage and store their knowledge. Thus far, scholars have largely focused on routines and people as knowledge reservoirs [5]. However, many more knowledge reservoirs likely exist, including artifacts, relationships, organizational information space, culture and structure [172]. The effectiveness of these reservoirs as the supporting knowledge infrastructure for the learning organizations have yet to be investigated, yet these reservoirs are critical to understanding the development and sustainability of the learning organization. We look forward to research that provides greater understanding of how to effectively use these reservoirs for organizational learning, and to research on how to sustain learning within organizations, more generally.

### **6.3 Learning Curves for Other Measures of Organizational Performance**

Most of the measures of organizational performance in the learning-curve literature map onto cost, quality, or time. Corbett and Van Wassenhove [37] identified cost, quality, and time as the dimensions of competence in the field of operations strategy. Dimensions of competence represent the internal competence of an operation. They indicate how well an operation can perform. Dimensions of competence are also referred to as operations priorities. Competitiveness, on the other hand,

captures an organization's ability to meet market's desires. Competence can aid competitiveness, but competence is by no means a guarantee for competitiveness. We anticipate future research on organizational learning curves for both dimensions of competence and competitiveness.

### **6.3.1 Dimensions of Competence: Measuring Operational Performance**

Without a doubt, operations priorities will remain important for organizations. As organizations continue to develop and introduce new products, organizations need to learn to improve cost, quality, and time for the latest products. Thus, we expect learning-curve scholars to continue to study cost, quality and lead-time. Several developments, however, suggest that the field should study additional dimensions of operational performance. The field of Operations Management has spent considerable effort studying supply chain management. Indeed, as Lee [104, p. 105] indicates, "Supply chain management has emerged as one of the major areas for companies to gain a competitive edge." Furthermore, "Because of shorter and shorter product life cycles, the pressure for dynamically adjusting and adapting a company's supply chain strategy is mounting" (pp. 118–119). Nevertheless, there is very little overlap in studying organizational learning curves and supply chain management. An observation by Lapré and Van Wassenhove [103], however, illustrates the importance of sound supply chain management for organizational learning curves (and we certainly expect the reverse to hold as well). In a study of Total Factor Productivity for a production line "MLC1," Lapré and Van Wassenhove [103, p. 63] found that:

The biggest disruptions in TFP were all associated with changes in raw material suppliers. While MLC1 had one preferred, principal supplier ("S"), shortages and/or quality problems at supplier S forced the MLC1 to sometimes use back-up suppliers. Two types of events were particularly disruptive: Having at least 50% of raw materials provided by back-up suppliers, or having to use at least 15% of raw materials provided by a different back-up supplier compared to the previous



month. Changes in raw materials required readjustments in process settings, and sometimes the introduction or deletion of an entire process step.

We anticipate a bigger role for supply chain management practice than a “mere control variable.” As organizations in supply chains move from adversarial relationships to supplier relationships [38], organizational learning plays an important role. How should supply chain partners learn to coordinate processes, learn to reduce inventory levels in the entire supply chain, learn to be more reliable?

Another area that stands to benefit from an organizational learning curve approach is sustainable operations management. Kleindorfer et al. [95, p. 482] use the term sustainability “to include environmental management, close-loop supply chains, and a broad perspective on triple-bottom-line thinking, integrating profit, people, and the planet into the culture, strategy, and operations of companies.” Corbett and Klassen [39] have argued that adopting an environmental perspective can actually lead to unexpected benefits for supply chain performance. The field would benefit from organizational learning curve research investigating issues such as reverse flows, component recovery, energy savings, and end-of-life product disposal.

### **6.3.2 Dimensions of Competitiveness: Measuring Performance in the Market Place**

Few studies have investigated organizational learning curves for dimensions of competitiveness. Notable exceptions include failure rates [19, 79], profitability [80], and customer dissatisfaction [98, 101]. Interestingly, all five studies found that in the early stages of an organization, increasing experience improves performance, but eventually the relationship reverses and increasing experience beyond that threshold reduces performance. All five studies investigated relative measures of performance. This begs the question: Is too much experience eventually going to be detrimental, or is there a way to avoid the competency trap of focusing on exploitation at the expense of exploration? What organizational learning efforts can re-ignite improvement after experiencing

a reversal in performance? Can the reversal effect be avoided by accumulating related experience — which Schilling et al. [145] identified as a richer balance between specialized and unrelated, highly diversified experience?

Another fruitful area for future research concerns significant expansion of the set of competitive dependent variables. The Service Profit Chain framework [75] provides some clues. According to the Service Profit Chain, internal quality drives employee satisfaction; employee satisfaction drives employee productivity and loyalty; which in turn enhance service quality; service quality enhances customer satisfaction; customer satisfaction fosters customer loyalty; customer loyalty in turn increases revenues and profitability. So, there is an internal portion of the Service Profit Chain focused around employees who enhance competence. This internal portion of the Service Profit Chain enhances the external portion of the Service Profit Chain that focuses on customers who determine competitiveness. An important notion of Service Profit Chain thinking is that organizations should not only think about individual service encounters, but instead focus on the lifetime a customer spends with the organization. Loyal customers provide the bulk of the revenues for an organization. Consequently, it should be worthwhile for learning-curve scholars to study dependent variables such as customer satisfaction, customer retention, repeat purchase, customer loyalty, and lifetime value of the customer.

## **6.4 Learning to Improve Multiple Measures of Performance**

### **6.4.1 Trade-offs?**

Up to the 1980s, dominant thinking in the field of operations strategy implied inherent trade-offs between operating priorities. “Higher quality could only be obtained at the expense of higher cost.” “Higher flexibility meant longer lead-times.” Ferdows and De Meyer [55] challenged this fundamental trade-off model with their “sand cone model.” Trade-offs could be avoided provided that companies invested in operating priorities in a certain sequence: Quality first, followed by flexibility and delivery, ending with cost (see also [37]). Schmenner and Swink [146]

introduced the notion of an “asset frontier,” which is formed by structural choices made by a company (investments in plant and equipment). Schmenner and Swink proposed that close to the asset frontier firms operate under the law of trade-offs, whereas further away from the asset frontier firms operate under the law of cumulative capabilities.

Evidence has been found to support either view (trade-offs versus sand cone model), but most research has used cross-sectional data, whereas longitudinal research designs are required to truly test the notion of trade-offs versus cumulative progression [143]. The only longitudinal study of performance improvement paths to date is Lapré and Scudder [100]. The authors studied the evolution of cost and quality (measured by customer dissatisfaction) over 11 years for the major U.S. airlines. For high performing airlines the analysis confirmed the predictions from the sand cone model when operating further away from their asset frontiers, whereas trade-offs do occur when operating close to asset frontiers. Studying the learning efforts behind performance improvement paths provides a tremendous opportunity for future research. The only learning-curve study in this space is Ref. [61]. The authors advance the notion that learning to improve one dimension may come at the expense of learning on another dimension. In a sample of 16 hospitals, the authors find evidence of a learning trade-off between efficiency and application innovation. Much more research is needed to better understand organizational learning curves on multiple dimensions. Can operating experience and deliberate learning activities simultaneously drive improvement for multiple measures of organizational performance, or do different performance measures require different learning variables? How would these learning effects differ across different pairs of performance measures? What learning activities can prevent trade-offs?

#### **6.4.2 Longitudinal Data Envelopment Analysis: Competitive Learning**

Data envelopment analysis (DEA) is particularly well suited to combine indicators that are measured in different dimensions [30, 120]. Briefly, DEA combines multiple inputs and multiple outputs for several decision

making units. DEA assumes that each unit has its own value system, so no a priori weighting of outputs and inputs are assumed. Each unit is analyzed one at a time to find that unit's best specific weights. Using these weights, each unit obtains an efficiency rating from 0 (worst) to 1 (best). Units with a rating of 1 are said to be on the efficient frontier, whereas units with a rating less than 1 are off the efficient frontier. Longitudinal DEA could significantly add to the organizational learning curve literature. Future research can investigate questions such as which firms consistently perform on the efficient frontier? Which firms do not? Why do firms move on/off the efficient frontier over time [153]? The questions we raised in the trade-offs section apply here as well. What types of experiences and deliberate learning activities foster competitive learning — staying on the efficient frontier or moving on the efficient frontier? What learning activities can prevent falling off the efficient frontier?

## **6.5 Conclusion**

Two decades later, Stata's [157, p. 64] observation continues to hold: "the rate at which individuals and organizations learn may be the only sustainable advantage, especially in knowledge-intensive industries." The variation in learning rates has been the subject of much research. Scholars have furthered our understanding of the organizational learning curve, and delved into the learning process behind this curve. The evidence suggests that this process is a key contributor to the tremendous variation in organizational learning rates. Much work remains to better understand this learning process. Hopefully the avenues for future research that we have proposed will lead to better guidance for organizations to better manage their organizational learning curves.

## Acknowledgments

---

We thank Charles Corbett and Zac Rolnik for their support, guidance, and encouragement. An anonymous reviewer provided thoughtful and constructive comments. The work by Michael Lapré was supported by the Dean's Fund for Faculty Research, Owen Graduate School of Management at Vanderbilt University. The work by Ingrid Nembhard was supported by a grant from the Agency for Healthcare Research and Quality.

## References

---

- [1] P. S. Adler, "Shared learning," *Management Science*, vol. 36, no. 8, pp. 938–957, 1990.
- [2] P. S. Adler and K. B. Clark, "Behind the learning curve: A sketch of the learning process," *Management Science*, vol. 37, no. 3, pp. 267–281, 1991.
- [3] L. Argote, "Group and organizational learning curves: Individual, system and environmental components," *British Journal of Social Psychology*, vol. 32, pp. 31–52, 1993.
- [4] L. Argote, *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Norwell, MA: Kluwer Academic Publishers, 1999.
- [5] L. Argote, "Group and Organizational Learning: Current Themes and Future Directions," Presented at Texas Conference, 2009.
- [6] L. Argote, S. L. Beckman, and D. Epple, "The persistence and transfer of learning in industrial settings," *Management Science*, vol. 36, no. 2, pp. 140–154, 1990.
- [7] L. Argote and E. D. Darr, "Repositories of knowledge in franchise organizations: Individual, structural and technological," in *The Nature and Dynamics of Organizational Capabilities*, (G. Dosi, R. R. Nelson, and S. G. Winter, eds.), pp. 51–68, New York, NY: Oxford University Press, 2000.
- [8] L. Argote, D. Epple, R. D. Rao, and K. Murphy, "The acquisition and depreciation of knowledge in a manufacturing organization: Turnover and plant productivity," Working paper, Carnegie Mellon University, Pittsburgh, Pennsylvania, 1997.
- [9] L. Argote, B. McEvily, and R. Reagans, "Managing knowledge in organizations: An integrative framework and review of emerging themes," *Management Science*, vol. 49, no. 4, pp. 571–582, 2003.

- [10] C. Argyris, "Double loop learning in organizations," *Harvard Business Review*, vol. 55, no. 5, pp. 115–126, 1977.
- [11] C. Argyris, *Strategy, Change, and Defensive Routines*. Boston, MA: Pitman, 1985.
- [12] C. Argyris, *Overcoming Organizational Defenses: Facilitating Organizational Learning*. Boston, MA: Allyn and Bacon, 1990.
- [13] K. J. Arrow, "The economic implications of learning by doing," *Review of Economic Studies*, vol. 29, pp. 155–173, 1962.
- [14] J. B. Arthur and C. L. Huntley, "Ramping up the organizational learning curve: Assessing the impact of deliberate learning on organizational performance under gainsharing," *Academy of Management Journal*, vol. 48, no. 6, pp. 1159–1170, 2005.
- [15] P. G. Audia, O. Sorenson, and J. Hage, "Tradeoffs in the organization of production: Multiunit firms, geographic dispersion and organizational learning," *Advances in Strategic Management*, vol. 18, pp. 75–105, 2001.
- [16] C. D. Bailey, "Forgetting and the learning curve: A laboratory study," *Management Science*, vol. 35, no. 3, pp. 340–352, 1989.
- [17] W. P. Barnett, H. R. Greve, and D. Y. Park, "An evolutionary model of organizational performance," *Strategic Management Journal*, vol. 15, no. S1, pp. 11–28, 1994.
- [18] J. A. C. Baum and K. B. Dahlin, "Aspiration performance and railroads' patterns of learning from train wrecks and crashes," *Organization Science*, vol. 18, no. 3, pp. 368–385, 2007.
- [19] J. A. C. Baum and P. Ingram, "Survival-enhancing learning in the Manhattan hotel industry, 1898–1980," *Management Science*, vol. 44, no. 7, pp. 996–1016, 1998.
- [20] W. F. Boh, S. A. Slaughter, and J. A. Espinosa, "Learning from experience in software development: A multilevel analysis," *Management Science*, vol. 53, no. 8, pp. 1315–1331, 2007.
- [21] R. E. Bohn, "Measuring and managing technological knowledge," *Sloan Management Review*, vol. 36, no. 1, pp. 61–73, 1994.
- [22] R. E. Bohn, "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, 1995.
- [23] R. E. Bohn, "From art to science in manufacturing: The evolution of technological knowledge," *Foundations and Trends in Technology, Information and Operations Management*, vol. 1, no. 2, pp. 129–212, 2005.
- [24] R. E. Bohn and M. A. Lapré, "Accelerated Learning by Experimentation," in *Learning Curves: Theory, Models, and Applications*, (M. Y. Jaber, ed.), forthcoming, Boca Raton, FL: CRC Press, Taylor and Francis, 2011.
- [25] R. E. Bohn and C. Terwiesch, "The economics of yield-driven processes," *Journal of Operations Management*, vol. 18, no. 1, pp. 41–59, 1999.
- [26] G. E. P. Box, J. S. Hunter, and W. G. Hunter, *Statistics for Experimenters*. New York, NY: Wiley, 1978.
- [27] V. E. Cangelosi and W. R. Dill, "Organizational learning: Observations toward a theory," *Administrative Science Quarterly*, vol. 10, no. 2, pp. 175–203, 1965.

- [28] M. D. Cannon and A. C. Edmondson, "Failing to learn and learning to fail (intelligently): How great organizations put failure to work to innovate and improve," *Long Range Planning*, vol. 38, pp. 299–319, 2005.
- [29] J. E. Carrillo and C. Gaimon, "A framework for process change," *IEEE Transactions on Engineering Management*, vol. 49, no. 4, pp. 409–427, 2002.
- [30] A. Charnes, W. Cooper, A. Y. Lewin, and L. M. Seiford, *Data Envelopment Analysis — Theory, Methodology and Applications*. Norwell, MA: Kluwer Academic Publishers, 1994.
- [31] W. B. Chew, *Productivity and Change: The Short-term Effects of Investments on Factory Level Productivity*. Unpublished doctoral dissertation, Harvard Business School, 1985.
- [32] W. B. Chew, T. F. Bresnahan, and K. B. Clark, "Measurement, coordination, and learning in a multiplant network," in *Measures for Manufacturing Excellence*, (R. S. Kaplan, ed.), pp. 129–162, Boston, MA: Harvard Business School Press, 1990.
- [33] A. S. Choo, K. W. Linderman, and R. G. Schroeder, "Method and psychological effects on learning behaviors and knowledge creation in quality improvement projects," *Management Science*, vol. 53, no. 3, pp. 437–450, 2007.
- [34] Y.-T. Chuang and J. A. C. Baum, "It's all in the name: Failure-induced learning by multiunit chains," *Administrative Science Quarterly*, vol. 48, no. 1, pp. 33–59, 2003.
- [35] Y.-T. Chuang, L. S. Ginsburg, and W. B. Berta, "Learning from preventable adverse events in health care organizations: Development of a multilevel model of learning and propositions," *Health Care Management Review*, vol. 32, no. 4, pp. 330–340, 2007.
- [36] W. M. Cohen and D. A. Levinthal, "Absorptive capacity: A new perspective on learning and innovation," *Administrative Science Quarterly*, vol. 35, no. 1, pp. 128–152, 1990.
- [37] C. Corbett and L. Van Wassenhove, "Trade-offs? What trade-offs? Competence and competitiveness in manufacturing strategy," *California Management Review*, vol. 35, no. 4, pp. 107–122, 1993.
- [38] C. J. Corbett, J. D. Blackburn, and L. N. Van Wassenhove, "Partnerships to improve supply chains," *Sloan Management Review*, vol. 40, no. 4, pp. 71–82, 1999.
- [39] C. J. Corbett and R. D. Klassen, "Extending the horizons: Environmental excellence as key to improving operations," *Manufacturing and Service Operations Management*, vol. 8, no. 1, pp. 5–22, 2006.
- [40] M. M. Crossan, H. W. Lane, and R. E. White, "An organizational learning framework: From intuition to institution," *Academy of Management Review*, vol. 24, no. 3, pp. 522–537, 1999.
- [41] R. M. Cyert and J. G. March, *A Behavioral Theory of the Firm*. Englewood Cliffs, NJ: Prentice-Hall, 1963.
- [42] M. Dada and R. Marcellus, "Process control with learning," *Operations Research*, vol. 42, no. 2, pp. 323–336, 1994.
- [43] E. D. Darr, L. Argote, and D. Eppler, "The acquisition, transfer, and depreciation of knowledge in service organizations: Productivity in franchises," *Management Science*, vol. 41, no. 11, pp. 1750–1762, 1995.



- [44] W. E. Deming, *Quality, Productivity, and Competitive Position*. Cambridge, MA: MIT Center for Advanced Engineering Study, 1982.
- [45] J. R. Dorroh, T. R. Gulledge, and N. K. Womer, "Investment in knowledge: A generalization of learning by experience," *Management Science*, vol. 40, no. 8, pp. 947–958, 1994.
- [46] R. Duncan and A. Weiss, "Organizational learning: Implications for organizational design," *Research in Organizational Behavior*, vol. 1, pp. 75–123, 1979.
- [47] J. M. Dutton and A. Thomas, "Treating progress functions as a managerial opportunity," *Academy of Management Review*, vol. 9, no. 2, pp. 235–247, 1984.
- [48] A. C. Edmondson, "Psychological safety and learning behavior in work teams," *Administrative Science Quarterly*, vol. 44, no. 2, pp. 350–383, 1999.
- [49] A. C. Edmondson, "The local and variegated nature of learning in organizations: A group-level perspective," *Organization Science*, vol. 13, no. 2, pp. 128–146, 2002.
- [50] A. C. Edmondson, "The competitive imperative of learning," *Harvard Business Review*, vol. 86, nos. 7–8, pp. 60–67, 2008.
- [51] A. C. Edmondson, R. M. Bohmer, and G. P. Pisano, "Disrupted routines: Team learning and new technology implementation hospitals," *Administrative Science Quarterly*, vol. 46, no. 4, pp. 685–716, 2001.
- [52] A. C. Edmondson, A. B. Winslow, R. M. J. Bohmer, and G. P. Pisano, "Learning how and learning what: Effects of tacit and codified knowledge on performance improvement following technology adoption," *Decision Sciences*, vol. 34, no. 2, pp. 197–223, 2003.
- [53] D. Epple, L. Argote, and K. Murphy, "An empirical investigation of the microstructure of knowledge acquisition and transfer through learning by doing," *Operations Research*, vol. 44, no. 1, pp. 77–86, 1996.
- [54] C. Fang, J. Lee, and M. A. Schilling, "Balancing exploration and exploitation through structural design: The isolation of subgroups and organizational learning," *Organization Science*, vol. 21, no. 3, pp. 625–642, 2010.
- [55] K. Ferdows and A. D. Meyer, "Lasting improvements in manufacturing performance: In search of a new theory," *Journal of Operations Management*, vol. 9, no. 2, pp. 168–184, 1990.
- [56] J. M. Field and K. K. Sinha, "Applying process knowledge for yield variation reduction: A longitudinal field study," *Decision Sciences*, vol. 36, no. 1, pp. 159–186, 2005.
- [57] C. H. Fine, "Quality improvement and learning in productive systems," *Management Science*, vol. 32, no. 10, pp. 1301–1315, 1986.
- [58] C. M. Fiol and M. A. Lyles, "Organizational learning," *Academy of Management Review*, vol. 10, no. 4, pp. 803–813, 1985.
- [59] D. Garvin, "Building a learning organization," *Harvard Business Review*, vol. 71, no. 4, pp. 78–91, 1993.
- [60] D. A. Garvin, A. C. Edmondson, and F. Gino, "Is yours a learning organization?," *Harvard Business Review*, vol. 86, no. 3, pp. 109–116, 2008.
- [61] F. Gino, R. M. J. Bohmer, A. C. Edmondson, G. P. Pisano, and A. B. Winslow, "Learning tradeoffs in organizations: Measuring multiple dimensions

- of improvement to investigate learning-curve heterogeneity,” Working Paper 05-047, Harvard Business School, 2006.
- [62] H. R. Greve, “Performance, aspirations, and risky organizational change,” *Administrative Science Quarterly*, vol. 43, no. 1, pp. 58–86, 1998.
- [63] H. R. Greve, *Organizational Learning from Performance Feedback: A Behavioral Perspective on Innovation and Change*. Cambridge, UK: Cambridge University Press, 2003.
- [64] M. R. Haas, “Knowledge gathering, team capabilities, and project performance in challenging work environments,” *Management Science*, vol. 52, no. 8, pp. 1170–1184, 2006.
- [65] J. R. Hackman, “The design of work teams,” in *Handbook of Organizational Behavior*, (J. W. Lorsch, ed.), pp. 315–342, Englewood Cliffs, NJ: Prentice-Hall, 1987.
- [66] J. R. Hackman and R. Wageman, “Total quality management: Empirical, conceptual, and practical issues,” *Administrative Science Quarterly*, vol. 40, no. 2, pp. 309–342, 1995.
- [67] N. W. Hatch and J. H. Dyer, “Human capital and learning as a source of sustainable competitive advantage,” *Strategic Management Journal*, vol. 25, no. 12, pp. 1155–1178, 2004.
- [68] N. W. Hatch and D. C. Mowery, “Process innovation and learning by doing in semiconductor manufacturing,” *Management Science*, vol. 44, no. 11, pp. 1461–1477, 1998.
- [69] P. R. Haunschild and M. Rhee, “The role of volition in organizational learning: The case of automotive product recalls,” *Management Science*, vol. 50, no. 11, pp. 1545–1560, 2004.
- [70] P. R. Haunschild and B. N. Sullivan, “Learning from complexity: Effects of prior accidents and incidents on airlines’ learning,” *Administrative Science Quarterly*, vol. 47, no. 4, pp. 609–643, 2002.
- [71] A. C. Hax and N. S. Majluf, “Competitive cost dynamics: The experience curve,” *Interfaces*, vol. 12, no. 5, pp. 50–61, 1982.
- [72] R. H. Hayes and K. B. Clark, “Exploring the sources of productivity differences at the factory level,” in *The Uneasy Alliance: Managing the Productivity-Technology Dilemma*, (K. B. Clark, R. H. Hayes, and C. Lorenz, eds.), pp. 151–188, Boston, MA: Harvard Business School Press, 1985.
- [73] R. H. Hayes, S. Wheelwright, and K. B. Clark, *Dynamic Manufacturing — Creating the Learning Organization*. New York, NY: The Free Press, 1988.
- [74] B. Hedberg, “How organizations learn and unlearn,” in *Handbook of Organizational Design*, (P. C. Nystrom and W. H. Starbuck, eds.), pp. 3–27, New York, NY: Oxford University Press, 1981.
- [75] J. L. Heskett, W. E. Sasser, Jr., and L. A. Schlesinger, *The Service Profit Chain*. New York, NY: The Free Press, 1997.
- [76] G. P. Huber, “Organizational learning: The contributing processes and literatures,” *Organization Science*, vol. 2, no. 1, pp. 880–115, 1991.
- [77] R. S. Huckman, B. R. Staats, and D. M. Upton, “Team familiarity, role experience, and performance: Evidence from Indian Software Services,” *Management Science*, vol. 55, no. 1, pp. 85–100, 2009.

- [78] P. Ingram, "Interorganizational learning," in *The Blackwell Companion to Organizations*, (J. A. C. Baum, ed.), pp. 642–663, Malden, MA: Blackwell Publishers Ltd., 2002.
- [79] P. Ingram and J. A. C. Baum, "Opportunity and constraint: Organizational learning from operating and competitive experience," *Strategic Management Journal*, vol. 18, pp. 75–98, Summer Special Issue 1997.
- [80] P. Ingram and T. Simons, "The transfer of experience in groups of organizations: Implications for performance and competition," *Management Science*, vol. 48, no. 12, pp. 1517–1533, 2002.
- [81] Institute of Medicine, *To Err is Human: Building a Safer Health System*. Washington, D.C.: National Academies Press, 1999.
- [82] Institute of Medicine, *Keeping Patients Safe: Transforming the Work Environment of Nurses*. Washington, D.C.: National Academy Press, 2004.
- [83] C. D. Ittner, V. Nagar, and M. V. Rajan, "An empirical examination of dynamic quality-based learning models," *Management Science*, vol. 47, no. 4, pp. 563–578, 2001.
- [84] R. Jaikumar and R. E. Bohn, "A dynamic approach to operations management: An alternative to static optimization," *International Journal of Production Economics*, vol. 27, no. 3, pp. 265–282, 1992.
- [85] I. L. Janis, *Groupthink*. New York: Houghton-Mifflin, 1982.
- [86] K. A. Jehn and C. Bendersky, "Intragroup conflict in organizations: A contingency perspective on the conflict-outcome relationship," in *Research in Organizational Behavior*, (B. Staw and L. L. Cummings, eds.), pp. 189–244, Greenwich, CT: JAI Press, 2003.
- [87] G. Johns, "The essential impact of context on organizational behavior," *Academy of Management Review*, vol. 31, no. 2, pp. 386–409, 2006.
- [88] P. L. Joskow and N. L. Rose, "The effects of technological change, experience, and environmental regulation on the construction cost of coal-burning generating units," *Rand Journal of Economics*, vol. 16, no. 1, pp. 1–27, 1985.
- [89] P. Kale, *Building an alliance capability: A knowledge-based approach*. Unpublished dissertation. Philadelphia, PA: University of Pennsylvania, 1999.
- [90] P. B. Kantor and W. I. Zangwill, "Theoretical foundation for a learning rate budget," *Management Science*, vol. 37, no. 3, pp. 315–330, 1991.
- [91] A. Khurana, "Managing complex production processes," *Sloan Management Review*, vol. 40, no. 2, pp. 85–97, 1999.
- [92] D. H. Kim, "The link between individual and organizational learning," *Sloan Management Review*, vol. 35, no. 1, pp. 37–50, 1993.
- [93] J.-Y. Kim, J.-Y. Kim, and A. S. Miner, "Organizational learning from extreme performance experience: The impact of success and recovery experience," *Organization Science*, vol. 20, no. 6, pp. 958–978, 2009.
- [94] K. J. Klein and J. S. Sorra, "The challenge of innovation implementation," *Academy of Management Review*, vol. 21, no. 4, pp. 1055–1080, 1996.
- [95] P. R. Kleindorfer, K. Singal, and L. N. Van Wassenhove, "Sustainable operations management," *Production and Operations Management*, vol. 14, no. 4, pp. 482–492, 2005.

- [96] T. K. Lant, F. J. Milliken, and B. Batra, "The role of managerial learning and interpretation in strategic persistence and reorientation: An empirical exploration," *Strategic Management Journal*, vol. 13, no. 8, pp. 585–608, 1992.
- [97] M. A. Lapré, "Inside the learning curve: Opening the black box of the learning curve," in *Learning Curves: Theory, Models, and Applications*, (M. Y. Jaber, ed.), forthcoming, Boca Raton, FL: CRC Press, Taylor and Francis, 2011.
- [98] M. A. Lapré, "Reducing customer dissatisfaction: How important is learning to reduce service failure?," *Production and Operations Management*, forthcoming, 2011.
- [99] M. A. Lapré, A. S. Mukherjee, and L. N. Van Wassenhove, "Behind the learning curve: Linking learning activities to waste reduction," *Management Science*, vol. 46, no. 5, pp. 597–611, 2000.
- [100] M. A. Lapré and G. D. Scudder, "Performance improvement paths in the U.S. airline industry: Linking trade-offs to asset frontiers," *Production and Operations Management*, vol. 13, no. 2, pp. 123–134, 2004.
- [101] M. A. Lapré and N. Tsikriktsis, "Organizational learning curves for customer dissatisfaction: Heterogeneity across airlines," *Management Science*, vol. 52, no. 3, pp. 352–366, 2006.
- [102] M. A. Lapré and L. N. Van Wassenhove, "Creating and transferring knowledge for productivity improvement in factories," *Management Science*, vol. 47, no. 10, pp. 1311–1325, 2001.
- [103] M. A. Lapré and L. N. Van Wassenhove, "Managing learning curves in factories by creating and transferring knowledge," *California Management Review*, vol. 46, no. 1, pp. 53–71, 2003.
- [104] H. L. Lee, "Aligning supply chain strategies with product uncertainties," *California Management Review*, vol. 44, no. 3, pp. 105–119, 2002.
- [105] D. Leonard-Barton, "The factory as a learning laboratory," *Sloan Management Review*, vol. 34, no. 1, pp. 23–38, 1992.
- [106] D. Z. Levin, "Organizational learning and the transfer of knowledge: An investigation of quality improvement," *Organization Science*, vol. 11, no. 6, pp. 630–647, 2000.
- [107] B. Levitt and J. G. March, "Organizational learning," *Annual Review of Sociology*, vol. 14, pp. 319–340, 1988.
- [108] F. Levy, "Adaptation in the production process," *Management Science*, vol. 11, no. 6, pp. B136–B154, 1965.
- [109] G. Li and S. Rajagopalan, "The impact of quality on learning," *Journal of Operations Management*, vol. 15, no. 3, pp. 181–191, 1997.
- [110] D. W. Liang, R. L. Moreland, and L. Argote, "Group versus individual training and group performance: The mediating role of transactive memory," *Personality and Social Psychology Bulletin*, vol. 21, no. 4, pp. 384–393, 1995.
- [111] M. B. Lieberman, "The learning curve and pricing in the chemical processing industries," *Rand Journal of Economics*, vol. 15, no. 2, pp. 213–228, 1984.
- [112] K. Lovelace, D. L. Shapiro, and L. R. Weingart, "Maximizing cross-functional new product teams' innovativeness and constraint adherence: A conflict communication perspective," *Academy of Management Journal*, vol. 44, no. 4, pp. 779–793, 2001.

- [113] P. M. Madsen and V. Desai, "Failing to learn? The effects of failure and success on organizational learning in the global orbital launch vehicle industry," *Academy of Management Journal*, vol. 53, no. 3, pp. 451–476, 2010.
- [114] J. G. March, "Footnotes to organizational change," *Administrative Science Quarterly*, vol. 26, no. 4, pp. 563–577, 1981.
- [115] J. G. March, "Exploration and exploitation in organizational learning," *Organization Science*, vol. 2, no. 1, pp. 71–87, 1991.
- [116] J. G. March and J. P. Olsen, "The uncertainty of the past: Organizational learning under ambiguity," *European Journal of Political Research*, vol. 3, pp. 147–171, 1975.
- [117] J. G. March and H. A. Simon, *Organizations*. New York: Wiley, 1958.
- [118] J. G. March, L. S. Sproull, and M. Tamuz, "Learning from samples of one or fewer," *Organization Science*, vol. 2, no. 1, pp. 1–13, 1991.
- [119] McKinsey Global Survey of Business Executives, "The McKinsey Quarterly," 2005.
- [120] R. D. Metters, F. X. Frei, and V. A. Vargas, "Measurement of multiple sites in service firms with data envelopment analysis," *Production and Operations Management*, vol. 8, no. 3, pp. 264–281, 1999.
- [121] F. J. Milliken, E. W. Morrison, and P. F. Hewlin, "An exploratory study of employee silence: Issues that employees don't communicate upward and why," *Journal of Management Studies*, vol. 40, no. 3, pp. 1453–1476, 2003.
- [122] A. S. Miner, J.-Y. J. Kim, I. W. Holzinger, and P. R. Haunschild, "Fruits of failure: Organizational failure and population-level learning," in *Advances in Strategic Management*, (A. S. Miner and P. Anderson, eds.), pp. 187–220, Stamford, CT: JAI Press, 1999.
- [123] K. Mishina, "Learning by new experiences: Revisiting the flying fortress learning curve," in *Learning by Doing in Markets, Firms, and Countries*, (N. R. Lamoreaux, D. M. G. Raff, and P. Temin, eds.), pp. 145–179, Chicago, IL: The University of Chicago Press, 1999.
- [124] A. Mody, "Firm strategies for costly engineering learning," *Management Science*, vol. 35, no. 4, pp. 496–512, 1989.
- [125] A. S. Mukherjee, M. A. Lapré, and L. N. Van Wassenhove, "Knowledge driven quality improvement," *Management Science*, vol. 44, no. 11 (part 2 of 2), pp. S35–S49, 1998.
- [126] J. F. Muth, "Search theory and the manufacturing progress function," *Management Science*, vol. 32, no. 8, pp. 938–962, 1986.
- [127] S. Narayanan, S. Balasubramanian, and J. M. Swaminathan, "Individual learning and productivity in a software maintenance environment: An empirical analysis," Working Paper, University of North Carolina, Chapel Hill, 2006.
- [128] I. M. Nembhard, J. Alexander, T. Hoff, and R. Ramanujam, "Why does the quality of health care continue to lag? Insights from management research," *Academy of Management Perspectives*, vol. 23, no. 1, pp. 24–42, 2009.
- [129] I. M. Nembhard and A. C. Edmondson, "Making it safe: The effects of leader inclusiveness and professional status on psychological safety and improvement efforts in health care teams," *Journal of Organizational Behavior*, vol. 27, no. 7, pp. 941–966, 2006.

- [130] I. M. Nembhard and A. C. Edmondson, "Psychological safety: A foundation for speaking up, collaboration and experimentation," in *The Oxford Handbook of Positive Organizational Scholarship*, (K. Cameron and G. Spreitzer, eds.), New York: Oxford University Press, 2011.
- [131] I. M. Nembhard and A. L. Tucker, "Deliberate learning to improve performance in dynamic service settings: Evidence from hospital intensive care units," in *Organization Science*, pp. 1–16, Articles in advance (September 15, 2010), 2010.
- [132] I. Nonaka, "The knowledge-creating company," *Harvard Business Review*, vol. 69, no. 6, pp. 77–91, 1991.
- [133] I. Nonaka, "A dynamic theory of organizational knowledge creation," *Organization Science*, vol. 5, no. 1, pp. 14–37, 1994.
- [134] S. K. Parker and T. D. Wall, "Job design and modern manufacturing," in *Psychology at Work*, (P. B. Warr, ed.), pp. 333–359, London: Penguin Group, 1996.
- [135] G. P. Pisano, "Knowledge, integration, and the locus of learning: An empirical analysis of process development," *Strategic Management Journal*, vol. 15, pp. 85–100, 1994.
- [136] G. P. Pisano, *The Development Factory*. Boston, MA: Harvard Business School Press, 1997.
- [137] G. P. Pisano, R. M. J. Bohmer, and A. C. Edmondson, "Organizational differences in rates of learning: Evidence from the adoption of minimally invasive cardiac surgery," *Management Science*, vol. 47, no. 6, pp. 752–768, 2001.
- [138] M. Polanyi, *The Tacit Dimension*. Garden City, NY: Doubleday, 1966.
- [139] R. Reagans, L. Argote, and D. Brooks, "Individual experience and experience working together: Predicting learning rates from knowing who knows what and knowing how to work together," *Management Science*, vol. 51, no. 6, pp. 869–881, 2005.
- [140] M. A. Roberto, "Lessons from Everest: The interaction of cognitive bias, psychological safety, and system complexity," *California Management Review*, vol. 45, no. 1, pp. 136–158, 2002.
- [141] K. H. Roberts and D. M. Rousseau, "Research in nearly failure-free, high reliability systems: Having the bubble," *IEEE Transactions on Engineering Management*, vol. 36, no. 2, pp. 132–139, 1989.
- [142] E. M. Rogers, *Diffusion of Innovations*. New York: Free Press, 1983.
- [143] E. D. Rosenzweig and G. S. Easton, "Tradeoffs in manufacturing? A meta-analysis and critique of the literature," *Production and Operations Management*, vol. 19, no. 2, pp. 127–141, 2010.
- [144] E. H. Schein, "How can organizations learn faster? The challenge of entering the green room," *Sloan Management Review*, vol. 34, no. 2, pp. 85–92, 1993.
- [145] M. A. Schilling, P. Vidal, R. E. Ployhart, and A. Marangoni, "Learning by doing something else: Variation, relatedness, and the learning curve," *Management Science*, vol. 49, no. 1, pp. 39–56, 2003.
- [146] R. W. Schmenner and M. L. Swink, "On theory in operations management," *Journal of Operations Management*, vol. 17, no. 1, pp. 97–113, 1998.
- [147] A. M. Schneiderman, "Setting quality goals," *Quality Progress*, vol. 21, no. 4, pp. 51–57, 1988.

- [148] M. Schulz, "Organizational learning," in *The Blackwell Companion to Organizations*, (J. A. C. Baum, ed.), pp. 415–441, Malden, MA: Blackwell Publishers Ltd., 2002.
- [149] P. M. Senge, *The Fifth Discipline: The art And Practice Of The Learning Organization*. New York, NY: Currency Doubleday, 1990.
- [150] E. Siemsen, S. Balasubramanian, and A. V. Roth, "Incentives that induce task-related effort, helping and knowledge sharing in workgroups," *Management Science*, vol. 53, no. 10, pp. 1533–1550, 2007.
- [151] T. Simons, L. H. Pelled, and K. A. Smith, "Making use of difference: Diversity, debate, and decision comprehensiveness in top management teams," *Academy of Management Journal*, vol. 42, no. 6, pp. 62–73, 1999.
- [152] G. Sinclair, S. Klepper, and W. Cohen, "What's experience got to do with it? Sources of cost reduction in a large specialty chemicals producer," *Management Science*, vol. 46, no. 1, pp. 28–45, 2000.
- [153] K. K. Sinha and A. H. Van de Ven, "Designing work within and between organizations," *Organization Science*, vol. 16, no. 4, pp. 389–408, 2005.
- [154] S. B. Sitkin, "Learning through failure: The strategy of small losses," *Research in Organizational Behavior*, vol. 14, pp. 231–266, 1992.
- [155] W. Skinner, "The focused factory," *Harvard Business Review*, vol. 52, no. 3, pp. 113–121, 1974.
- [156] O. Sorenson, "Interdependence and adaptability: Organizational learning and the long-term effect of integration," *Management Science*, vol. 49, no. 4, pp. 446–463, 2003.
- [157] R. Stata, "Organizational learning — The key to management innovation," *Sloan Management Review*, vol. 30, no. 3, pp. 63–74, 1989.
- [158] J. D. Sterman, "Learning in and about complex systems," *System Dynamics Review*, vol. 10, no. 2–3, pp. 291–330, 1994.
- [159] G. Szulanski, "Exploring internal stickiness: Impediments to the transfer of best practice within the firm," *Strategic Management Journal*, vol. 17, no. (Special Winter Issue), pp. 27–43, 1996.
- [160] C. S. Tapiero, "Production learning and quality control," *IIE Transactions*, vol. 19, no. 4, pp. 362–370, 1987.
- [161] S. E. Taylor and J. D. Brown, "Illusion and well-being: A social psychological perspective on mental health," *Psychological Bulletin*, vol. 103, no. 2, pp. 193–210, 1988.
- [162] C. Terwiesch and R. E. Bohn, "Learning and process improvement during production ramp-up," *International Journal of Production Economics*, vol. 70, no. 1, pp. 1–19, 2001.
- [163] C. Terwiesch and Y. Xu, "The copy-exactly ramp-up strategy: Trading-off learning with process change," *IEEE Transactions on Engineering Management*, vol. 51, no. 1, pp. 70–84, 2004.
- [164] P. Thompson, "How much did the liberty shipbuilders forget?," *Management Science*, vol. 53, no. 6, pp. 908–918, 2007.
- [165] A. L. Tucker, A. C. Edmondson, and S. Spear, "When problem solving prevents organizational learning," *Journal of Organizational Change Management*, vol. 15, no. 2, pp. 122–137, 2002.

- [166] A. L. Tucker, I. M. Nembhard, and A. C. Edmondson, "Implementing new practices: An empirical study of organizational learning in hospital intensive care units," *Management Science*, vol. 53, no. 6, pp. 894–907, 2007.
- [167] M. J. Tyre and E. von Hippel, "The situated nature of adaptive learning in organizations," *Organization Science*, vol. 8, no. 1, pp. 71–83, 1997.
- [168] M. van Biema and B. Greenwald, "Managing our way to higher service-sector productivity," *Harvard Business Review*, vol. 75, no. 4, pp. 87–95, 1997.
- [169] A. H. Van de Ven, A. L. Delbecq, and R. Koenig, Jr., "Determinants of coordination modes within organizations," *American Sociological Review*, vol. 41, no. 2, pp. 322–338, 1976.
- [170] D. Vera and M. Crossan, "Organizational learning and knowledge management: Toward an integrative framework," in *Handbook of Organizational Learning*, (M. Easterby-Smith and M. Lyles, eds.), pp. 122–141, Oxford, UK: Blackwell, 2003.
- [171] D. Vera and M. Crossan, "Strategic leadership and organizational learning," *Academy of Management Review*, vol. 29, no. 2, pp. 222–240, 2004.
- [172] T. L. Virani, Lemieux-Charles, D. A. Davis, and W. Berta, "Sustaining change: Once evidence-based practices are transferred, what then?," *Longwoods Review*, vol. 6, no. 4, pp. 89–96, 2009.
- [173] K. G. Vosburgh and R. S. Newbower, "Moore's law, disruptive technologies, and the clinician," in *Medicine Meets Virtual Reality*, (J. D. Westwood, H. M. Hoffman, G. T. Mogel, R. Robb, and D. Stredney, eds.), pp. 8–13, Amsterdam: IOS Press, 2000.
- [174] S. Walston, "Does re-engineering really work?," Presented at the EGOS Conference, Maastricht, The Netherlands July, 1999.
- [175] C. Weber, "Yield learning and the sources of profitability in semiconductor manufacturing and process development," *IEEE Transactions on Semiconductor Manufacturing*, vol. 17, no. 4, pp. 590–596, 2004.
- [176] K. E. Weick and K. M. Sutcliffe, *Managing the Unexpected*. San Francisco, CA: Jossey-Bass, 2001.
- [177] K. E. Weick, K. M. Sutcliffe, and D. Obstfeld, "Organizing for high reliability: Processes of collective mindfulness," in *Research in Organizational Behavior*, (B. M. Staw and R. Sutton, eds.), pp. 81–123, Greenwich, CT: JAI Press, 1999.
- [178] E. Wiersma, "Conditions that shape the learning curve: Factors that increase the ability and opportunity to learn," *Management Science*, vol. 53, no. 12, pp. 1903–1915, 2007.
- [179] K. Y. Williams and C. A. O'Reilly, "Demography and diversity in organizations: A review of the 40 years of research," in *Research in Organizational Behavior*, 20, (B. M. Staw and L. L. Cummings, eds.), pp. 77–140, Greenwich, CT: JAI Press, 1998.
- [180] S. Winter, "Knowledge and competence as strategic assets," in *The Competitive Challenge — Strategies for Industrial Innovation and Renewal*, (D. Teece, ed.), pp. 159–184, Cambridge, MA: Ballinger, 1987.
- [181] S. G. Winter, "Organizing for continuous improvement: Evolutionary theory meets the quality revolution," in *Evolutionary Dynamics of Organizations*,



- (J. A. C. Baum and J. K. Singh, eds.), New York, NY: Oxford University Press, 1994.
- [182] T. P. Wright, "Factors affecting the costs of airplanes," *Journal of Aeronautical Science*, vol. 3, pp. 122–128, 1936.
  - [183] J. Ye, D. Marinova, and J. Singh, "Deliberate learning in the frontlines of service organizations," *Academy of Management Proceedings*, pp. 1–6, 2008.
  - [184] L. E. Yelle, "The learning curve: Historical review and comprehensive survey," *Decision Sciences*, vol. 10, no. 2, pp. 302–328, 1979.
  - [185] U. Zander and B. Kogut, "Knowledge and the speed of the transfer and imitation of organizational capabilities: An empirical test," *Organization Science*, vol. 6, no. 1, pp. 76–91, 1995.
  - [186] M. Zollo, "Knowledge codification, process routinization and the development of organizational capabilities: Post-acquisition integration in the U.S. banking industry," Unpublished doctoral dissertation, University of Pennsylvania, Philadelphia, PA, 1998.
  - [187] M. Zollo and H. Singh, "Deliberate learning in corporate acquisitions: Post-acquisition strategies and integration capability in U.S. bank mergers," *Strategic Management Journal*, vol. 25, no. 13, pp. 1233–1256, 2004.
  - [188] M. Zollo and S. G. Winter, "Deliberate learning and the evolution of dynamic capabilities," *Organization Science*, vol. 13, no. 3, pp. 339–351, 2002.