

Creating and Transferring Knowledge for Productivity Improvement in Factories

Michael A. Lapré • Luk N. Van Wassenhove

Owen Graduate School of Management, Vanderbilt University, Nashville, Tennessee 37203

INSEAD, Boulevard de Constance, 77305 Fontainebleau Cedex, France

michael.lapre@owen.vanderbilt.edu • luk.van-wassenhove@insead.edu

Can a firm accelerate its learning curve if knowledge about the production function is incomplete? This article identifies a production line specifically set up to create technological knowledge about its production function through scientific experimentation (formal learning) as opposed to learning by doing. The organizational structure of this line was very successful in creating technological knowledge. Formal learning resulted in huge productivity improvements. Replication of this organizational structure on three production lines in other plants within the same firm fell short of expectations. Formal learning did not result in similar productivity improvements. Our research suggests two factors that may facilitate creation and transfer of technological knowledge: management buy-in and knowledge diversity to solve interdepartmental problems.

(Learning Curve; Total Factor Productivity; Learning by Doing; Formal Learning; Technological Knowledge; Knowledge Transfer; Replication)

1. Introduction

Scholars have frequently studied the learning curve phenomenon. Managers have extensively used learning curves for planning purposes (Yelle 1979). Basically, the logarithm of unit cost decreases with the logarithm of cumulative number of units produced at a uniform rate—the learning rate. Learning rates, however, show considerable variation within industries, within firms, even within plants (Levy 1965, Hayes and Clark 1985). Drawing on more than 200 learning curve studies, Dutton and Thomas (1984) conclude that a learning rate should no longer be treated as a given constant based on past performance, but as a dependent variable influenced by a firm's behavior.

A disappointing managerial implication of the larger part of the learning curve literature is that the only way to speed up cost improvement is to speed up cumulative volume, which may not always be desirable nor feasible. Yet, experts have emphasized the competitive potential of learning rates. Some

even argue that “the rate at which individuals and organizations learn may become the only sustainable competitive advantage” (Stata 1989). Consequently, we need to improve our understanding on how to manage learning curves. In particular, scholars have called for research on managerial levers that accelerate learning curves (see, e.g., Jaikumar and Bohn 1992).

Few studies have incorporated managerial variables in learning curve analyses. The ones that did focused on training and engineering activity to capture deliberate activities that aim to accelerate learning curves. Levy (1965) used direct labor training hours to explain different learning rates across workers in a cross-sectional study. Adler and Clark (1991) used longitudinal data on cumulative number of hours spent by workers on training and cumulative number of hours spent on engineering changes. Interestingly, the authors found that each managerial variable could enhance as well as disrupt total factor productivity. The learning process behind the

learning curve, therefore, was internally complex. As the authors admitted, training and engineering activity were proxies for more explicit learning efforts. The authors acknowledged that future research on these managerial learning variables was necessary. Hatch and Mowery (1998) showed that cumulative engineering significantly enhanced learning curves in a sample of 52 semiconductor processes. When new processes were introduced in manufacturing, however, cumulative engineering disrupted ongoing learning activities of existing processes. Furthermore, for some product groups, cumulative engineering did not affect learning curves.

These mixed outcomes of the impact of managerial learning variables on learning curves led to an in-depth study of the nature of knowledge created in deliberate learning activities. Mukherjee et al. (1998) analyzed 62 quality improvement projects undertaken in one factory over a decade. The authors found that learning processes in quality improvement projects exhibit considerable variation along two learning dimensions: conceptual and operational learning. *Conceptual learning* is the process of acquiring a better understanding of cause-and-effect relationships, i.e., the acquisition of know-why. *Operational learning* is the process of obtaining validation of action-outcome links, i.e., the acquisition of know-how.

In a follow-up paper, Lapré et al. (2000) used these learning dimensions to categorize quality improvement projects into different types of knowledge acquisition. For each type of knowledge acquisition, the authors constructed a cumulative number of completed projects variable. Only 25% of the projects, the ones that acquired both know-why and know-how, accelerated the factory's learning rate. Projects that produced know-why without the corresponding know-how were disruptive. Projects that failed to acquire know-why did not affect the learning rate. Case evidence suggested that it was possible to consistently produce the learning rate enhancing mix of know-why and know-how. An organizational structure called *model line*, a production line run as a learning laboratory, consistently produced know-why and know-how that was successfully transferred to the rest of the factory.

We studied subsequent learning efforts in the same company. Impressed by the success of this model line, the company decided to (i) set up model lines in other factories, and (ii) transfer the management team of the first model line to central R&D headquarters to coordinate knowledge creation and transfer on these new model lines. Several years later, management was disappointed with the performance of the new model lines.

Given the importance of the nature of knowledge created by deliberate learning activities and the apparent success of the organizational structure of the first model line in creating the "right knowledge" to accelerate the learning curve, we focus on two issues here. First, we investigate replication of the first model line: What is the impact of model line projects on a learning curve? Second, we investigate knowledge sharing between model lines: What is the impact of projects from other model lines on a learning curve? As these model lines are primarily concerned with productivity improvement, we use total factor productivity learning curves to investigate our questions.

This article is organized as follows. In the balance of this section, we discuss knowledge creation in factories when knowledge about a firm's production function is fundamentally incomplete. We draw on evolutionary economics and the dynamic approach to operations management. After addressing some of the challenges factories face in creating better knowledge about production functions, we review some empirical learning curve studies of knowledge transfer between factories. Section 2 describes the research site, §3 the data collection and analysis plan. In §4, we present and discuss the learning curve estimates. Section 5 summarizes our findings and questions for future research.

1.1. Knowledge Creation in Factories

According to evolutionary economics, knowledge about a firm's production function is typically incomplete (Nelson and Winter 1982). It cannot be fully articulated and codified in "blueprints." Instead, knowledge resides—for a large part—in organizational routines and often takes the form of tacit knowledge. Causal depth of knowledge is limited.

The dynamic approach to operations management views manufacturing environments in a very similar way (Jaikumar and Bohn 1992). Bohn (1994, p. 62) defines technological knowledge as “understanding the effects of the input variables [of a manufacturing process] on the output.” The author identifies eight stages of technological knowledge ranging from complete ignorance to complete understanding. This scale applies to every possibly relevant variable a plant needs to manage. In complex and dynamic production environments, there can be hundreds of variables, while only a few are completely understood. To proactively manage a plant’s learning curve, personnel has to undertake deliberate activities to learn about the production process and create better technological knowledge. Ultimately, a manufacturing process is governed by relatively immutable laws of nature, such as physics, chemistry, metallurgy, etc. As a plant progresses on the stages of knowledge, plant personnel need to engage in scientific experimentation. Scientific formulae—if correct—facilitate the derivation of statistical process control rules. Consequently, deliberate learning activities that augment learning by doing involve much more “formal learning.” *Formal learning* consists of theory building, formulating hypotheses, and testing hypotheses with statistical experimentation (cf. Bohn 1987, Mukherjee et al. 1998).

As Pisano (1994) shows, learning cannot always be “outsourced” to an R&D laboratory. If a production technology is closer to an art than a science, the best place to learn is the full scale manufacturing environment because R&D laboratories typically lack the complexity encountered in factories.

Can a plant manage efforts to engage in formal learning and create better technological knowledge? Lapré et al. (2000) demonstrate the importance of creating both know-why and know-how in the context of quality improvement efforts. Winter (1994) discusses several factors that may enhance quality/productivity improvement efforts. First, the “buy-in problem” requires strong top management commitment. Second, a high level of inefficiency is like “high-quality ore” for management to mine. Significant portions of the high-quality ore could well be

interdepartmental sources of waste. Tackling interdepartmental problems, though, is more costly and challenging. To do so, firms cannot rely on departmental managers. Instead, interdepartmental teams with diverse knowledge backgrounds are needed.

In conclusion, faced with incomplete knowledge about a manufacturing process, interdepartmental teams are more likely to create better technological knowledge than departmental teams are. Knowledge creation needs to take place in the plant. Top management’s buy-in is likely to enhance knowledge creation.

1.2. Knowledge Transfer Between Factories

“We know relatively more about knowledge . . . transfer than we know about knowledge creation in organizations” (Argote 1999, p. 203). In this section, we review some patterns observed in empirical learning curve studies of transfer of learning by doing. We focus on patterns relevant for intrafirm transfer between production units. For a comprehensive review, see Argote (1999).

In a cross-sectional study of 41 kitchens in a commercial food firm, Chew et al. (1990) found large differences in total factor productivity, even after controlling for structural and other characteristics. From this persistence of variation, the authors concluded that local know-how was not shared within the firm. In fact, existing incentive systems led to the widespread belief in the firm that plants and managers were unique, thus inhibiting transfer of knowledge. Hence, transfer of experience should not be taken for granted.

Recently, scholars have started to address transfer of learning by doing with longitudinal data on learning variables. Whereas cumulative production volume is the traditional measure for learning by doing within a production unit, cumulative production volume summed across all production units is used to investigate transfer of knowledge acquired via learning by doing. In a study of 16 shipyards, Argote et al. (1990) identified transfer of learning at the start-up of a yard. However, the authors did not find ongoing sharing of learning once a yard started production. Comparing intercepts and slopes of total factor productivity learning curves, Adler (1990) found similar results in a multiplant electronics firm.

Epple et al. (1991) found partial transfer of learning between two shifts in a truck plant. After adding a night shift in another automotive assembly plant, the day shift continued to learn (Epple et al. 1996). Transfer of this experience was the sole source of learning for the night shift. The latter study shows that scholars need to be careful in assuming symmetry. The day shift learned by doing, whereas the night shift did not. Sharing of experience occurred only in one direction.

So, productivity differences can persist across production units in the same firm. Transfer of learning by doing may occur at start-up and during ongoing production. But as knowledge typically does not become fully embedded in the technology, neither type of transfer is guaranteed to occur. Finally, production units in a single firm do not necessarily have identical learning rates, nor do production units share learning to the same extent.

Transfer of learning by doing concerns transfer between existing organizational units. At a "higher level" of transfer, firms can engage in replication of organizational units. In discussing the replication of a business model, Winter and Szulanski (2001) introduce the notions of a template and the Arrow core. A *template* refers to an existing business model that is the guiding example a firm wants to replicate. "There are typically competing *ideas* about the keys to success, some of them tested to some degree, some of them not" (p. 11). Conversely, the *Arrow core* is "the kernel of information that accounts for the value-creation potential of the business model when it is leveraged by replication" (p. 12). Loosely, it is everything one needs to know, nothing more, nothing less, yet, it is typically unknown. Hence, replicators are faced with the challenge to identify the Arrow core in a template.

2. The Context

The research site for this study was N. V. Bekaert, S.A., a Belgian multinational corporation. Bekaert was the world's largest independent producer of steel wire. In particular, its steel cord division, which hosted this research, produced about one-third of the world's output of steel wire (called "tire cord") used in the production of steel belted radial tires. By European standards the firm was among the first


to embark on a total quality management program, for which it earned two European Foundation for Quality Management awards in the early 1990s. This research project received the enthusiastic backing of CEO Rafaël Decaluwé, who provided us with unlimited access to plants, people, and management information systems.

A simplified version of Bekaert's basic process flow is depicted in Figure 1. Thick wire (called "wire rod") was drawn through dies that progressively reduced its diameter. At intermediate points, wires were heat treated, and a chemical process coated wires with brass. Coated wires were drawn even thinner through dies submerged in a soap solution. Very thin wires (called "filaments") were bunched around each other to form tire cord. The simplest cord had two filaments; the most complex, hundreds.

Knowledge about the production function of this process was very incomplete. No wire rod supplier could guarantee homogeneity of properties for raw materials. As a result, microscopic flaws in the wire rod could cause fractures (the industry's biggest productivity problem) at any process stage. Nobody could specify the combination of the hundreds of process settings that would prevent fractures. For example, cylindrical dies wore out nonsymmetrically. What level of ovality would cause an unacceptable level of fractures in subsequent process stages? In absence of any technological knowledge, plant personnel would just use some experience-based heuristic for die replacement. Sometimes variables were thought to be important, but could not be measured, e.g., the dryness with which wire entered a furnace. In those cases, plant personnel even ignored tacit knowledge.

In the late 1980s, Bekaert realized that central R&D laboratories lacked characteristics of the dynamic production environment encountered in factories (cf. Pisano 1994). Bekaert therefore relocated process optimization to the factory. In 1988, the company established a "model line" at its Belgian flagship plant A (MLA) for an important, representative product, and top management gave a senior R&D manager the responsibility for creating fundamental process control knowledge without sacrificing production of saleable wire.

Figure 1 Simplified, Schematic Process Flow of a Bekaert Factory



Note. In reality, several steps were repeated. Knowledge on the production function was very incomplete, e.g., nobody could specify the combination of the hundreds of process settings that would prevent the occurrence of defects. Traditionally, personnel and machines were organized by departments corresponding to the stages in the production process.

The MLA was a fundamentally new organizational structure at Bekaert. Typically, Bekaert steel cord factories organized personnel and machines by production departments corresponding to the stages in the production process depicted in Figure 1. The MLA, on the other hand, was an integrated line corresponding to a single product *cutting across all production departments*. People and machines were *dedicated to a single product*, and managed by a team led by a senior R&D manager.

The MLA team routinely defined and conducted *projects* for productivity and quality improvement. The team set improvement goals based on both R&D and production experience. The team collected data on any product/process variable deemed relevant. Its unique data capture allowed the team to link data across process stages. The MLA team typically solved problems by building scientific models from which it derived testable hypotheses. Hypotheses were tested with natural and controlled experiments.¹ The resulting regression output was used to modify limits for statistical process control (Wadsworth et al. 1986).

After two years, the MLA was perceived to be very successful. By 1991, impressed with the productivity and quality improvements achieved on the MLA, Bekaert started model lines for similar, important,

representative products at two other major European steel cord plants B and C, which we will refer to as MLB and MLC1, respectively. The first manager of the MLA was appointed as the corporate model line manager with the responsibility for coordinating these model line efforts at central R&D. In 1993, a second model line was added in plant C (MLC2).

Project knowledge was typically codified in monthly activity reports and shared with central R&D and other model line teams. Twice a year, model line teams met at central R&D to exchange such knowledge. Further, model line managers also communicated project knowledge as well as production experience directly between them and with central R&D, predominantly by phone and fax, occasionally by visits.

By 1995, the new model lines had not delivered results comparable to the MLA record. Particularly, creation of technological knowledge on the new model lines had not progressed to the same extent. Looking back in 1996, the corporate model line manager commented:

Initially, we thought that we could solve everything in a [model] line with natural data, just doing the statistical experiments. But this is not true. [Experiments] need to be guided by scientific models. This, however, is very difficult. How do you create this? ... Two things are key: responsibility, which so far has been largely driven by existing culture in a plant, and the creativity, or research skills of [model line personnel].

¹ For a description of natural and controlled experiments see Bohn (1987).

2.1. The Model Line Experiment

There were many examples where Bekaert's production personnel either ignored tacit knowledge or used routines based on tacit knowledge. Hence, knowledge of a plant's production function that translated inputs such as raw materials and process settings into good output was very incomplete (cf. Nelson and Winter 1982, Bohn 1994). Furthermore, plant complexity and plant personnel's predominantly tacit knowledge base inhibited learning.

The model line at plant A (MLA) allowed plant personnel to create technological knowledge via formal learning that individual production departments would never have been able to create. Mukherjee (1992) described two in-depth case studies of projects dealing with the same problem in the wet wire drawing (WWD) department of plant A (Figure 1). In one project, personnel from the WWD department tried to solve the problem by solely considering process variables controlled in the WWD department. The personnel were unable to solve the problem. The MLA, on the other hand, was able to demonstrate that variables determined in two upstream production departments—dry drawing and heat treatment—caused the problem. To do so, the MLA manager constructed a chemical model based on fragmentary R&D knowledge he had acquired during his 16 years in R&D. At R&D, he had worked on many projects involving all the production steps shown in Figure 1. For example, in the 1970s he had participated in an R&D project on the ability of tire cord to withstand corrosion. From this R&D project, he remembered that some copper-related variables determined in the brass coating step were relevant for the problem at hand in the WWD department. The MLA team tested the model with controlled experiments. As a result the MLA obtained a sharp improvement in productivity.

We make two observations. First, the MLA team operated with full authority over production and resources reflecting top management's buy-in. Second, this project was a typical example of the MLA manager's use of his diverse knowledge base to solve interdepartmental problems (Mukherjee 1992).

Table 1 summarizes key dates for the model lines, as well as the model line team compositions. As Epple et al. (1991, p. 69) note, it is a challenge in field

Table 1 Start-up Dates, Team Composition, and Results of the Model Lines

	MLA	MLB	MLC1	MLC2
Start production of product	1988	1990	1988	1972
Start model line concept	1988	1991	1991	1993
Team: no. Engineers	2	1	1	
no. Technicians	2	1	2	
no. Workers	4	2	2	
Indirect labor cost compared to MLA	100%	50%	66%	
Prior experience on the team	R&D, plant	plant	plant	
Year ML manager joined the company	1972	1989	1990	
Improvement in TFP compared to MLA	100%	-4%	25%	7%

research on learning "to find quasi-experiments in the field that make it possible to control for some factors while varying others." In comparing the four model line implementations we can control for firm, product type,² resource market, and production technology factors. Factors that differ across the four lines summarized in Table 1 include resources allocated, experience, and product maturity. We present additional data to illustrate how these factors led to variation across the four model lines in terms of management buy-in and interdepartmental problem solving (Winter 1994). We explore the impact of this variation in the discussion of how the four lines improved productivity. The next section describes our data.

3. Data and Method

3.1. Total Factor Productivity

Following Hayes and Clark (1985), Adler (1990), Adler and Clark (1991), and Ittner (1994), we study monthly total factor productivity (TFP) learning curves. TFP is the value of output divided by the sum

²To protect confidentiality, we cannot report the two key defining characteristics of the four model line products: the number of filaments and the diameters of the filaments. However, all four products consisted of a comparable number of filaments with comparable diameters. Moreover, Lapré et al. (2000) found evidence of transfer of MLA project knowledge to other products in plant A. The MLB, MLC1, and MLC2 products fell within the range of plant A products in terms of number of filaments and diameters of filaments. So, model lines could benefit from transfer of project knowledge. Model line managers confirmed this assessment.

of the costs of four inputs (direct and indirect labor, capital, and raw materials):³

$$TFP_t = \frac{p_{93}y_t}{\sum_j w_{j,93}x_{jt}}$$

where p_{93} is the 1993 price per ton, y_t output produced in month t , $w_{j,93}$ 1993 unit cost for input j , and x_{jt} quantity of input j used in month t . We chose 1993 as the base period for fixing prices and costs, as 1993 was the only year during which all four model lines were operating. Inflation does not affect TFP provided the same inflation rate applies to all prices and unit costs. Later in the article, we address the sensitivity of our results to changes in 1993 prices and unit costs. We collected data to compute TFP primarily from the management information systems used by model line personnel. These systems reported productivity data for each process step on a monthly basis. For the MLA, data were available from September 1989 until production of the product was transferred to a U.S. plant (August 1993). For the other lines, data were available from start-up (between July 1991 and February 1993) until April 1995. For each line, we aggregated the data for all process steps to derive TFP for the entire model line. We defined output and input quantities as follows:

Output. For plants A and B, output was directly available as good tonnage produced at the final step. For plant C, we estimated output by multiplying the amount of raw material released to the process with the product of yields at every step. Yields were easily derived from available reject rates per step. The corporate model line manager estimated 1993 prices for all four lines. We verified these estimates in the plants.

Direct Labor. Direct labor hours per ton were directly available for each step. Multiplying with the tonnage processed on the corresponding steps and

then summing across all steps gave total direct labor hours. We multiplied total direct labor hours with 1993 average direct labor cost per hour (including benefits).

Indirect Labor. A model line was essentially a plant-within-a-plant, run by a model line team. So, indirect labor was simply a headcount of process engineers, technicians, and indirect workers assigned to a model line team. Evolutions of these headcounts obtained from model line managers were confirmed by model line personnel and documents. Headcounts were weighed by annual 1993 labor costs (including benefits) for each category divided by 11. (Bekaert operated 11 months per year.)

Capital. Machine hours per ton for each step, directly available, were multiplied by output produced at the corresponding step, yielding total machine hours per step. For each machine type, we obtained 1993 machine prices (including installation costs) from R&D documents. Using steel cord plant linear depreciation schemes, we transformed machine prices to an hourly basis. Following Hayes and Clark (1985), Adler (1990), and Adler and Clark (1991), we operationalized total cost of capital as the return of capital (depreciation) plus the return on capital (opportunity cost). For the former, we obtained linear depreciation schemes from the plants' accounting departments. For the latter, we followed recent TFP studies (Adler 1990, Adler and Clark 1991, Ittner 1994) by estimating the return on capital at 7%, reflecting the long-run average inflation corrected cost of a typical mix of debt and equity (Kaplan 1986). We obtained total capital costs by multiplying the resulting machine cost per hour for each machine type with total machine hours at the corresponding process step and then summing across all steps.

Raw Materials. At plant C, wire rod tonnage released to the process was directly available. For plants A and B, we derived wire rod tonnage from final output and reject rates for the process steps. On model lines A, C1, and C2, there were repeated shifts from expensive, high quality wire rod to cheaper, lower quality wire rod. We also collected monthly percentages of each wire rod used. For each wire rod

³ There are two other inputs: energy and supplies. Data on these two inputs, however, were not consistently recorded. For one model line, available energy data showed a perfectly constant amount of energy consumption per ton. Excluding energy therefore has little impact on the variation in TFP. Supplies (mainly soap for the drawing steps) make up less than 5% of the unit cost, so exclusion does not have a big impact on TFP either.


quality, we obtained 1993 prices from plants' purchasing departments.

Following Hayes and Clark (1985), we graphed TFP and discussed resulting patterns with model line managers to understand evolution of productivity (e.g., trends, peaks, plateaus). Although the model line managers were our key informants, we have discussed these patterns with managers at all levels within Bekaert, literally ranging from foremen on factory floors to the CEO. Through these discussions we learned that we had to recalculate TFPs, accounting for different raw material qualities and changes in the composition of model line teams. Although Bekaert used many single dimensional quality and productivity measures, such as rejects and machine hours per ton at a particular step, management completely lacked an overall measure for total model line productivity performance. Managers at all levels felt that the TFP measure was particularly well suited for this purpose 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 evolution, even though they lacked such a measure before. Again following Hayes and Clark (1985), we did not relate TFP estimates to learning variables until Bekaert's management agreed on the computation of TFP and felt that the resulting pattern gave a good representation of the model lines' productivity evolution (see Figure 2). To protect Bekaert's proprietary data, we do not report the scale. It is important to note that we started collecting our data in 1996. So, none of our discussions with Bekaert personnel could have confounded performance on model lines. The variation between the TFP curves, ranging from sharp improvements to no improvements at all, demonstrates that different model line implementations gave different results. Next, we discuss factors to explain this variation.

3.2. Learning Variables

We use the conventional measure for learning by doing: Q_{it} is cumulative production volume on model line i at time t . This measure is based on the start of

Figure 2 Logarithm of Total Factory Productivity for the Four Model Lines



production reported in Table 1, i.e., it includes any production history prior to the availability of TFP data. For transfer of learning by doing to model line i , we sum output over the other model lines since the start of model line i to compute AQ_{it} . This measure excludes production histories prior to model line start dates, because we want to distinguish between transfer at start-up of model lines and ongoing transfer of learning by doing in a model line context.

Model line teams conducted improvement projects. These projects typically dealt with specific quality/productivity problems. We use cumulative number of projects completed (P_{it}) to measure formal learning within a model line. Lapré extensively interviewed the model line managers to obtain completion dates of all model line projects. In his presence, each model line manager went through all monthly activity reports from the start of a model line until 1996 to identify all projects and completion dates. In addition, this process identified for each project (i) the production departments involved and (ii) whether projects were defined by model line teams or plant management. Summing over the other model lines we compute AP_{it} to measure transfer of formal learning to model line i .

Table 2 summarizes the four learning variables in a categorization proposed by Dutton and Thomas (1984). Causal factors for learning curves differ along two dimensions. The endogenous/exogenous dimension refers to whether learning is internal or external to a production unit. The autonomous/induced dimension refers to automatic improvement due to

Table 2 Measures for Creation and Transfer of Two Types of Knowledge in Model Lines

	Learning by doing	Formal learning
Endogenous	Cumulative production Quantities $Q_{it} = \sum_{s=S_j}^t q_{is}$	Cumulative Projects $P_{it} = \sum_{s=M_j}^t p_{is}$
Exogenous	Aggregate cumulative production Quantities $AQ_{it} = \sum_{j \neq i} \sum_{s=M_j}^t q_{js}$	Aggregate cumulative Projects $AP_{it} = \sum_{j \neq i} \sum_{s=M_j}^t p_{js}$

Note. q_{it} (p_{it}) is the production volume (number of projects completed) on line i in month t . S_j (M_j) is the month production (model line concept) started for line i .

sustained production versus more deliberate investments in learning. In our case, learning by doing resembles autonomous learning, whereas formal learning is a case of induced learning.

3.3. Control Variables

Because of fixed indirect labor cost, TFP can be sensitive to volume changes. To separate economies of scale due to increased capacity utilization from productivity improvements that reflect learning, TFP studies typically include a measure for capacity utilization (see, e.g., Adler 1990). Following Hayes and Clark (1985), we define capacity utilization (CU) as the ratio of actual output to estimated capacity, where the latter was constructed by linearly interpolating between successive peaks of actual output during the scale-up phase. Once the maximum output level for the time horizon under study was reached, we used maximum output to estimate capacity. This approach is justified because capacity was never reallocated to other production lines during the time horizons studied in Figure 2.

Hayes and Clark (1985) and Ittner (1994) found that rejects can have a profound effect on TFP. To isolate any changes in quality policies that were not learning related, we include rejects as well. We used reject rates for individual process steps to compute an aggregate model line reject rate (R).

Capacity utilization and rejects have been used before as predictor variables in TFP research at process firms. Discussions of TFP patterns at plant C

revealed another one. When asked why TFP repeatedly dropped significantly, the MLC team used archival data to discover that low TFP values turned out to be associated with changes in wire rod suppliers. Plant personnel chose to switch to expensive, high quality rod (i) if a regular supplier could not deliver or (ii) to reoptimize the process if a customer modified specifications for product properties. In the latter case, use of higher quality raw material was thought to facilitate the search process, because it reduced noise compared to more heterogeneous low-quality raw material. However, regardless of whether the process switched from low to high or high to low quality, every change required readjustments in process settings, or even the introduction or deletion of an entire process step. Consequently, all changes in raw material suppliers could disrupt TFP. To isolate the effect of changes in raw material from learning (in)efficiencies, we introduce ΔRM , specifying the share of wire rod for which a different quality is used compared to the previous month. Appendix A gives the correlation matrix for all variables. We now turn to the analysis plan.

3.4. Analysis Plan

The basic TFP learning curve controlled for capacity utilization is (Hayes and Clark 1985)

$$\ln \text{TFP} = \beta_0 + \beta_1 \ln Q + \beta_2 \ln CU + u.$$

We use the same functional form to estimate TFP learning curves with the learning and control variables discussed above, specifically (let i and t denote model line and time indices, and let $D_j = 1$ if $j = i$, 0 otherwise):

$$\begin{aligned} \ln \text{TFP}_{it} = & \beta_{01} + \sum_{j>1} \beta_{0j} D_j + \beta_1 \ln Q_{it-1} \\ & + \beta_2 \ln AQ_{it-1} + \beta_3 \ln P_{it-1} \\ & + \beta_4 \ln AP_{it-1} + \sum_k \beta_k X_{kit} + u_{it}, \end{aligned} \quad (1)$$

where the X_{kit} s are the control variables $\ln CU_{it}$, R_{it} , and ΔRM_{it} .

Equation (1) makes a very restrictive assumption. It assumes identical coefficients across model lines for each independent variable, i.e., it assumes that for

each model line learning by doing was equally important, formal learning was equally important, and so on. Looking at Figure 2, it seems highly unlikely that learning rates for the four model lines are identical. Moreover, there is ample evidence in the TFP literature that even within firms learning rates differ across production units and that productivity differences can persist (see, e.g., Hayes and Clark 1985, Chew et al. 1990). Consequently, we relax the assumption of identical coefficients and allow both the intercept and all slope parameters to change:

$$\begin{aligned} \ln \text{TFP}_{it} = & \beta_{0i} + \sum_{j>1} \beta_{0j} D_j + \beta_{1i} \ln Q_{it-1} \\ & + \beta_{2i} \ln AQ_{it-1} + \beta_{3i} \ln P_{it-1} \\ & + \beta_{4i} \ln AP_{it-1} + \sum_k \beta_{ki} X_{kit} + u_{it}, \end{aligned}$$

or equivalently (remember $D_j = 1$ if $j = i$, 0 otherwise),

$$\begin{aligned} \ln \text{TFP}_{it} = & \beta_{0i} + \sum_{j>1} \beta_{0j} D_j + \sum_j D_j \\ & \times \left(\beta_{1j} \ln Q_{jt-1} + \beta_{2j} \ln AQ_{jt-1} \right. \\ & \left. + \beta_{3j} \ln P_{jt-1} + \beta_{4j} \ln AP_{jt-1} \right. \\ & \left. + \sum_k \beta_{kj} X_{kjt} \right) + u_{it}. \end{aligned} \quad (2)$$

Even though the two equations are equivalent, we will refer to Equation (2), because it emphasizes that we will estimate more parameters in a single analysis. In the next section, we estimate Equations (1) and (2) with the AUTOREG procedure in the SAS package, specifying an AR(1) process for the error term $u_{it} = \rho u_{it-1} + e_{it}$ (SAS/ETS 1993).

To assess “management buy-in” for a model line, we tabulate the number of projects defined by model line teams versus plant management. Management buy-in for the model line concept is higher if plant management leaves project definition to model line teams. To assess “interdepartmental problem solving” on a model line, we tabulate the number of projects that are (i) confined to a single department, (ii) interdepartmental between two neighboring departments, and (iii) interdepartmental spanning at least three departments.

Finally, we use case data in the discussion of our findings. Our primary sources of information for the discussion include dozens of interviews we held with the corporate model line manager and corporate model line team members at central R&D and with model line managers, model line team members, and plant managers at plants A, B, and C. At these interviews we frequently obtained copies of relevant company documents, such as memos, presentation hand-outs, and reports. For the MLA, we also used company documents and interview notes collected for prior research at Bekaert (Mukherjee 1992, Mukherjee et al. 1998, Lapré et al. 2000).

4. Empirical Results

Table 3 reports the regression results. Intercepts and coefficients for model-line-specific dummy variables are not reported to protect Bekaert’s proprietary data. Model (1) shows that model lines improved TFP weakly via learning by doing ($p = .0527$) and strongly via formal learning. Comparing Models (1) and (2), though, demonstrates that care needs to be taken in assuming identical learning rates across production units. First, the R^2 for Model (2) is significantly higher as indicated by the F statistic of 16.96 for the change in R^2 . Second, Model (1) suffers from autocorrelation—an indication of an ill-specified model—whereas Model (2) does not. Third, the Akaike information criterion (AIC) computed as $-2 \ln L + 2k$ (where $\ln L$ is the logarithm of the likelihood function evaluated at the parameter estimates and k is the number of explanatory variables) is lower for Model (2). So, even when taking the trade-off between parsimony and precision (Judge et al. 1988) into account, Model (2) is the preferred specification over (1). All coefficients for the control variables have the expected signs, although some are not significant.⁴ Only one dummy was significantly different from zero: The MLC2 dummy was negative.

⁴To appreciate the importance of the learning variables, we estimated Equation (2) omitting the control variables. The resulting R^2 was 0.9040 compared to 0.9796 in Table 3. So, the learning variables account for a very important part of the variation explained.

Table 3 TFP Learning-Curve Estimates

	Model (1) Identical Coefficients	Model (2) Individual Coefficients			
		MLA	MLB	MLC1	MLC2
Learning by doing ($\ln Q_{it-1}$)	0.073 (1.95)	-0.081 (-0.87)	-0.051 (-0.58)	0.202* (1.98)	1.887** (2.70)
Transfer of learning by doing ($\ln AQ_{it-1}$)	0.002 (0.71)	0.004 (1.35)	0.004 (0.83)	0.010* (2.01)	0.003 (0.41)
Formal learning ($\ln P_{it-1}$)	0.036** (3.47)	0.093** (3.41)	-0.034 (-1.53)	0.017 (1.09)	-0.080 (-1.75)
Transfer of formal learning ($\ln AP_{it-1}$)	-0.015 (-1.65)	-0.003 (-0.27)	0.041 (1.58)	-0.046 (-1.03)	-0.049 (-1.50)
Capacity utilization ($\ln CU_{it}$)	0.097** (9.05)	0.156** (10.97)	0.043** (2.07)	0.105** (4.30)	0.047 (1.24)
Reject rate (R_{it})	-1.41** (-12.23)	-1.619** (-9.03)	-1.150** (-2.04)	-1.206** (-5.16)	-1.615** (-2.02)
Changes in raw materials (ΔRM_{it})	-0.026 (-1.59)	na	na	-0.059** (-3.05)	-0.069 (-1.73)
ρ	0.423** (5.46)			0.028 (0.31)	
R^2	0.9239			0.9796	
F statistic for increase in R^2				16.96**	
AIC	-578.8			-605.3	
n	149			149	

Note. Dependent variable $\ln TFP_{it}$. T statistics appear in parentheses. *Significant at 0.05 in a 2-tailed test.

**Significant at 0.01 in a two-tailed test.

In Table 4, we tabulate the number of projects defined by model line teams versus plant management. Table 5 classifies projects based on the interdepartmental approach to problem solving. Next, we relate findings in Table 3 to the context of each model line using Tables 1, 4, and 5, as well as case data.

4.1. Learning in Model Lines

MLA. The first model line at plant A started before the others and stopped in 1993, because production

was transferred to a U.S. plant. Given the MLA's pioneering role, it is not surprising that the MLA did not benefit from exogenous learning. Table 3 shows that the MLA relied primarily on formal learning as opposed to learning by doing. This finding confirms prior research findings at Bekaert (Mukherjee 1992, Mukherjee et al. 1998, Lapré et al. 2000) as well as our own observations. In many MLA projects, the team pooled fragmented scientific, engineering, and experience-based knowledge to build a scientific model and to formulate testable hypotheses. Hypotheses were tested with natural and controlled

Table 4 Management Buy-in for Projects Conducted on the Model Lines

	MLA	MLB	MLC1	MLC2
Model line projects defined by model line team	10	0	4	1
Model line projects defined by plant management	0	4	10	6
Department-wide projects defined by plant management	0	12	4	4

Table 5 Interdepartmental Problem Solving on the Model Lines

	MLA	MLB	MLC1	MLC2
Intradepartmental projects	5	16	16	9
Projects between two neighboring departments	0	0	2	2
Projects spanning at least three departments	5	0	0	0

experiments. The resulting regression output gave the team formulae to derive statistical process control rules to control process variables. The team would find that some process variables needed to be controlled more tightly, whereas others were controlled too tightly. Through scientific experimentation, the MLA team repeatedly improved its understanding about how process variables affected good output. As Table 4 shows, top management buy-in allowed the MLA to define all projects. Moreover, a significant portion of the MLA projects were truly interdepartmental spanning more than two production departments (Table 5).

MLB. None of the learning variables was significant for the MLB (Table 3). As Figure 2 shows, TFP was flat for this line. Table 4 shows that plant management never allowed the MLB team to define its projects. This can be a serious impediment for productivity improvement. Winter (1994, p. 98) notes that "Management must behave so as to provide a supportive structure within which people at lower levels can act effectively to improve performance. It cannot achieve significant performance improvements by its own unilateral action, for it lacks the detailed knowledge required to do so." Furthermore, all projects were confined to a single production department (Table 5). Apparently, management of plant B did not buy in to the interdepartmental model line concept. Instead, plant B maintained Bekaert's traditional organizational structure of production departments. As a consequence, the MLB team did not get authority to implement changes in production on model line B.

MLC1. According to Table 3, the MLC1 relied on learning by doing as opposed to formal learning. Plant management chose the projects up to 1993. From 1994 onward, the model line team defined the projects. Table 4 shows that most of the projects in our database were indeed defined by plant management. Once management buy-in for the MLC1 was established, it was still hard to actually create scientific knowledge on the MLC1.⁵ The projects defined by the

MLC team continued to be incremental productivity improvement projects compared to the breakthrough projects on the MLA, which yielded fundamental scientific knowledge. Moreover, few projects were interdepartmental (see Table 5). The few projects that were interdepartmental worked on issues between two neighboring departments. The MLA, on the other hand, routinely tackled issues that related cause-and-effect spanning at least three departments. Comparing the MLC team with the MLA team in Table 1 shows that the MLA possessed a rich diverse knowledge base: The manager had been with Bekaert for 16 years, and he could draw from both R&D as well as production experience. The MLC manager, on the other hand, possessed only one year of experience with Bekaert, only in plant C, before he was appointed MLC manager. Consequently, he lacked the knowledge diversity necessary to engage in formal learning—particularly to reduce interdepartmental waste. Several years later, the corporate model line manager summarized:

What we didn't understand when we started model lines in plants B and C is that a model line manager needs to have authority over production and projects, and a young engineer is not a good choice to run a model line. Young engineers lack experience in formal problem solving. Experiments failed.

MLC2. In 1993, the model line team at plant C was given a second model line. However, the MLC team invested only 15% of its time and efforts in the MLC2; the other 85% remained focused on the MLC1. Table 3 shows that the only significant learning variable for the MLC2 was learning by doing. Even though the TFP graph (Figure 2) does not show a lot of improvement, the strong learning by doing indicates that the MLC2 was already "down its learning curve." The mature MLC2 product had been in production for over 20 years (see Table 1).

Although the MLA was perceived to be a success, model line implementations at plants B and C did not yield similar improvements. Why did formal learning not have an impact in plants B and C? Based on

⁵We redefined P_{it} and AP_{it} to include only model line projects defined by model line teams (first row in Table 4). We reestimated Equation (2) with these redefined formal learning variables. None

of the results for formal learning changed, i.e., there was neither formal learning for model lines B, C1, C2, nor transfer of formal learning.

classifications of projects and case data, we suggested the following plausible explanations. For MLB, management buy-in was missing. For MLC1, knowledge diversity to solve interdepartmental problems was underdeveloped compared to the MLA, thus inhibiting formal learning. The MLC2 was already “down its learning curve” due to the maturity of its product. We now address the robustness of the results.

4.2. Robustness of the Results

The model lines at plants B and C started later. Did these lines learn at start-up when the model line concept was adopted? We estimated Model (2) with $\ln \sum_{j \neq i} P_{jM_i}$, the amount of formal learning available at the start of model line i , to test for initial transfer at start-up. As this variable is a linear combination of the intercept and the model-line-specific dummies D_j , we had to omit model-line-specific dummies (cf. Argote et al. 1990). The coefficient for $\ln \sum_{j \neq i} P_{jM_i}$ was not significant, indicating that no initial transfer took place. We obtained a similar result with $\ln \sum_{j \neq i} Q_{jM_i}$, the amount of learning by doing available at start-up of model line i . These findings are not surprising given that none of the model-line-specific dummies were significantly greater than zero.

The technology depicted in Figure 1 did not fundamentally change over the time horizon studied. We included calendar time in Model (2) as a proxy for any technological development. Calendar time was not significant, indicating that our assessment of no fundamental change in technology appears to be correct.

If each model line had quality problems with a different production step, specific model-line knowledge may not have been useful for the others to improve. As one reviewer suggested, we examined the distribution of reject rates for the process steps. For all four model lines, rejects were of the same order of magnitude on a particular step. This was true for all steps. Between steps, however, reject rates could differ an order of magnitude. So, reject rates for the process steps followed a similar distribution for all four model lines. Consequently, it is unlikely that different quality problems impeded knowledge transfer.

The MLB did not learn at all in the time horizon studied in Figure 2. Was there simply nothing left to learn? After this time horizon, the steel cord division started to put more emphasis on cost. The MLB changed to a cheaper raw material that resulted in a 20% improvement in TFP. So, clearly, there was significant room for learning, underscoring our explanation that management of plant B did not buy in to the model line concept.

In Appendix B, we discuss sensitivity analyses for alternative ways of calculating TFP. These analyses show that the estimates reported in Table 3 are robust for large changes in 1993 prices and unit costs used to calculate TFP.

5. Conclusion

We still know relatively more about knowledge transfer than we know about knowledge creation (Argote 1999). Further research is needed to better understand organizational structures and problem-solving routines that allow firms to accelerate learning curves.

In our research, we identified an organizational structure—the first model line at plant A—that was extremely successful at proactively creating the “right” knowledge to accelerate its learning curve, i.e., to engage in formal learning. Replication efforts of this model line within the organization fell short of expectations. Model lines in plants B and C did not improve productivity through formal learning. We used classifications of projects to suggest that lack of management buy-in and lack of interdepartmental problem solving may have impeded formal learning in plants B and C. Future research should address the generalizability of our single-firm study. How can we measure these constructs in other settings? What other factors might enhance formal learning? For example, what incentive issues need to be addressed?

A comparison of MLA and MLC1 suggests that knowledge diversity is an important factor. Future research should address the role of knowledge diversity in creating technological knowledge, especially when interdepartmental sources of waste make for a high quality ore to mine. How can

firms build knowledge diversity? Is it necessary for organizational members to accumulate 16 years of diverse experience in R&D, or can interdepartmental teams be formed with fewer years of “overlapping” experience?

Model line replications at plants B and C did not result in transfer of formal learning. Several factors may have contributed to this. First, projects on model lines B, C1, and C2 did not result in productivity improvements. So, other lines might have simply ignored projects from these lines. Second, up to 1993 there was no management buy-in for the model line concept at plants B and C—possibly prohibiting transfer of MLA projects. Third, if a model line team has difficulty engaging in formal learning, it may be equally difficult to incorporate formal learning from others. Cohen and Levinthal (1990, p. 134) note that “Learning by doing does not contribute to the diversity that is critical to learning about or creating something that is relatively new.” For example, the MLC1 focus on learning by doing might have blocked the build up of diverse formal knowledge and consequently impeded innovative performance. Future research in other settings is needed to sort out these possible explanations.

This learning curve study explicitly distinguished transfer of improvement projects from transfer of other experience. We hope that similar studies will increase our understanding of the conditions that facilitate transfer of deliberate learning activities. Addressing these issues will help organizations to better measure and manage their learning curve processes.

Acknowledgments

Roger Bohn, Robin Plager, Lars Hendrik Röller, Gabriel Szulanski, Ludo Van der Heyden, Patricia Vidal, and participants at Wharton’s Technology Miniconference, in particular discussant Sidney Winter, provided helpful comments. Departmental Editor Linda Argote and three anonymous referees provided very thoughtful comments. The authors thank the management and employees of N. V. Bekaert S.A. for their unstinting cooperation. Michael Lapré gratefully acknowledges a summer research grant from the Dean’s Fund for Faculty Research from the Owen School at Vanderbilt, as well as financial support by Arthur D. Little, N. V. Bekaert S. A., Boston University School of Management, INSEAD, and the Sasakawa foundation.

Appendix A: Correlations

	ln Q	ln AQ	ln P	ln AP	ln CU	R	ΔRM
ln TFP	0.62	0.68	0.16	0.55	0.10	-0.76	0.19
ln Q		0.47	0.39	0.47	-0.18	-0.40	0.04
ln AQ			0.44	0.76	-0.27	-0.66	0.17
ln P				0.54	-0.38	-0.13	-0.07
ln AP					-0.20	-0.42	0.06
ln CU						0.22	0.07
R							-0.13

All correlations ≥ 0.16 or ≤ -0.16 are significant at 0.05.

Appendix B: TFP Sensitivity Analyses

We calculated total factor productivity as $TFP_t = p_{93}y_t / \sum_j w_{j,93}x_{jt}$, where p_{93} is the 1993 price per ton, y_t output produced in month t , $w_{j,93}$ 1993 unit cost for input j , and x_{jt} quantity of input j used in month t . We chose 1993 as the base period for fixing prices and costs, as 1993 was the only year during which all four model lines were operating. We investigated how sensitive the results in Table 3 are to changes in p_{93} and $w_{j,93}$.

We increased and decreased price for one model line by 10%, keeping prices for the other lines constant. We did this for all four model lines, giving us eight alternative calculations of TFP. We also increased and decreased $w_{j,93}$ by 10% for each of the four inputs (across all lines), keeping the other $w_{k,93}$ ($k \neq j$) constant. This gave us eight more alternative calculations for TFP. We reestimated Models (1) and (2) for all sixteen alternative calculations of TFP.

It should be noted that 10% represents a very large margin of error. Our prices are much more accurate. Our $w_{j,93}$ are even more accurate. Results available from the authors show that the original estimates in Table 3 are rather insensitive to 10% changes in p_{93} or $w_{j,93}$. Comparison of Models (1) and (2) yields the same results. The significant increase in R^2 is at least 3.86%. Model (1) suffers from autocorrelation, whereas Model (2) does not. AIC is lower for Model (2). Moreover, the MLA shows strong formal learning, the MLC2 shows strong learning by doing, and p -values for learning by doing and transfer of learning by doing on the MLC1 are close to 0.05. The only change concerns the MLC1 dummy. In one scenario (decrease p_{93} by 10% for the MLC1), the MLC1 dummy becomes significant at 0.05. But as it is negative, none of the results for transfer at start-up change.

References

- Adler, P. S. 1990. Shared learning. *Management Sci.* **36** 938–957.
- , K. B. Clark. 1991. Behind the learning curve: A sketch of the learning process. *Management Sci.* **37** 267–281.
- Argote, L. 1999. *Organizational Learning: Creating, Retaining and Transferring Knowledge*. Kluwer Academic Publishers, Norwell, MA.
- , S. L. Beckman, D. Eppler. 1990. The persistence and transfer of learning in industrial settings. *Management Sci.* **36** 140–154.

- Bohn, R. E. 1987. Learning by experimentation in manufacturing. Working paper 88-001, Harvard Business School, Boston, MA.
- . 1994. Measuring and managing technological knowledge. *Sloan Management Rev.* **36**(1) 61–73.
- Chew, W. B., T. F. Bresnahan, K. B. Clark. 1990. Measurement, coordination, and learning in a multiplant network. R. S. Kaplan, ed. *Measures for Manufacturing Excellence*. Harvard Business School Press, Boston, MA.
- Cohen, W. M., D. A. Levinthal. 1990. Absorptive capacity: A new perspective on learning and innovation. *Admin. Sci. Quart.* **35** 128–152.
- Dutton, J. M., A. Thomas. 1984. Treating progress functions as a managerial opportunity. *Acad. Management Rev.* **9** 235–247.
- Epple, D., L. Argote, R. Devadas. 1991. Organizational learning curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing. *Organ. Sci.* **2** 58–70.
- , ———, K. Murphy. 1996. An empirical investigation of the microstructure of knowledge acquisition and transfer through learning by doing. *Oper. Res.* **44** 77–86.
- Hatch, N. W., D. C. Mowery. 1998. Process innovation and learning by doing in semiconductor manufacturing. *Management Sci.* **44** 1461–1477.
- Hayes, R. H., K. B. Clark. 1985. Exploring the sources of productivity differences at the factory level. K. B. Clark, R. H. Hayes, C. Lorenz, eds. *The Uneasy Alliance: Managing the Productivity-Technology Dilemma*. Harvard Business School Press, Boston, MA.
- Ittner, C. D. 1994. An examination of the indirect productivity gains from quality improvement. *Production Oper. Management* **3** 153–170.
- Jaikumar, R., R. E. Bohn. 1992. A dynamic approach to operations management: An alternative to static optimization. *Internat. J. Production Econom.* **27** 265–282.
- Judge, G. G., R. C. Hill, W. E. Griffiths, H. Lütkepohl, T. Lee. 1988. *Introduction to the Theory and Practice of Econometrics*. 2nd ed. Wiley, New York.
- Kaplan, R. S. 1986. Must CIM be justified by faith alone? *Harvard Bus. Rev.* **64**(2) 87–95.
- Lapré, M. A., A. S. Mukherjee, L. N. Van Wassenhove. 2000. Behind the learning curve: Linking learning activities to waste reduction. *Management Sci.* **46** 597–611.
- Levy, F. K. 1965. Adaptation in the production process. *Management Sci.* **11** B136–B154.
- Mukherjee, A. S. 1992. The effective management of organizational learning and process control in the factory. Unpublished doctoral dissertation, Harvard Business School, Boston, MA.
- , M. A. Lapré, L. N. Van Wassenhove. 1998. Knowledge driven quality improvement. *Management Sci.* **44** S35–S49.
- Nelson, R. R., S. G. Winter. 1982. *An Evolutionary Theory of Economic Change*. Belknap Press, Cambridge, MA.
- Pisano, G. P. 1994. Knowledge, integration, and the locus of learning: An empirical analysis of process development. *Strategic Management J.* **15** 85–100.
- SAS/ETS. 1993. *User's Guide Version 6, 2nd ed.* SAS Institute Inc., Cary, NC.
- Stata, R. 1989. Organizational learning—The key to management innovation. *Sloan Management Rev.* **30**(3) 63–74.
- Wadsworth, H. M., K. S. Stephens, A. B. Godfrey. 1986. *Modern Methods for Quality Control and Improvement*. Wiley, New York.
- Winter, S. G. 1994. Organizing for continuous improvement: Evolutionary theory meets the quality revolution. J. A. C. Baum, J. K. Singh, eds. *Evolutionary Dynamics of Organizations*. Oxford University Press, New York.
- , G. Szulanski. 2001. Replication as strategy. *Organ. Sci.* Forthcoming.
- Yelle, L. E. 1979. The learning curve: Historical review and comprehensive survey. *Decision Sci.* **10** 302–328.

Accepted by Linda Argote; received January 30, 1998. This paper has been with the authors 25 months for 3 revisions.