

**Learning to Be Strategic: Learning *By Doing* and Learning *Before Doing* in Strategic
Decision Making**

PATRICIA VIDAL¹

MELISSA A. SCHILLING²

ABSTRACT

In this paper, we examine whether the learning curves so familiar in production settings are also manifest in strategic decision making by teams. We also analyze the degree to which learning before *doing* impacts the intercept or slope of such learning curves. We use data on 10 teams performing complex strategic problem-solving tasks over time. Results indicate that team strategic problem-solving does exhibit learning curve effects, and that investment in learning before doing (study of a process that takes place prior to commencing the task) significantly increased initial performance, but decreased the slope of the learning curves (the learning rate). These results lend support to the argument that learning before doing can jumpstart initial performance, but having "picked the low hanging fruit" (Pisano, 1997:40), less room is left for improvement through learning by doing.

¹ PhD - Boston University - School of Management - pativdal@bu.edu

² Professor at New York University - Stern School of Business - mschilli@stern.nyu.edu

Support for this research was provided by the *Systems Research Center* at Boston University.

The authors would like to gratefully acknowledge the help and support of Michael Lapre, William Starbuck, Mason Carpenter, Tom Cloherty, Alexander Marangoni, Mahesh Rajan, Tammy Madsen, and Cassandra Vasco.

Key words: Learning curves; learning by doing; learning before doing, group learning, organizational learning; strategic decision making, problem solving

INTRODUCTION

It has been repeatedly demonstrated that firms experience learning curves in production processes (e.g., e.g., Argote, 1993, 1999; Baloff, 1971; Dutton & Thomas, 1984; Hatch & Mowery, 1998; Levy, 1965; Mukherjee, Lapre & Van Wassenhove, 1998; Yelle, 1979). As a firm increases its experience with a production process increases, the knowledge about this process increases (Dutton & Thomas, 1984), the firm is able to make improvements in the process over time which translate in decreased cost per unit or increased performance (Argote, 1993, 1999; Baloff, 1971; Darr, Argote, & Epple, 1995; Hatch & Mowery, 1998; Yelle, 1979).

It is much less clear, however, that management teams experience learning curves in their strategic decision making experience. In fact, there is a considerable body of research suggesting that management teams are unable (or unwilling) to adapt their management strategies at all, nipping most opportunities for improving their strategic decision making abilities in the bud (e.g., Christensen, 1997; Leonard-Barton, 1992; Starbuck, 1992).

Though it makes sense intuitively that managers should get better at anything in which they have experience over time, the processes of improving strategic decision making are very different than for improving the production processes. First, the objectives for improving the production process are often much more clear than for strategic decision making. For example, production process improvement might target reducing waste (Mukherjee, Lapre, Van Wassenhove, 1998), reducing accidents (Greenberg, 1971), or decreasing time spent at a task (Argote, 1999). The objectives of strategic decision making tend to be much "fuzzier" (Schwenk & Cosier, 1993), such as improving the firm's competitive positioning, or leveraging its core competencies. Second, while it may be possible to change the production process incrementally and observe the outcome, changes in strategic decision making are often much more complex and systemic, and

their outcomes are dependent on a myriad of factors (e.g., economic conditions, customer preferences, competitor response), that make it difficult to attribute performance to the change. Finally, production process improvements are often embodied in changes to machinery or production procedures. The nature of the change is thus explicitly observable, and its impact tends to endure over time. By contrast, improvements in strategic decision making ability may be difficult to codify, making it difficult to replicate the improvements or pass them down to successive generations. Furthermore, their effectiveness may be highly path dependent, indicating that even if managers were able to codify the strategic decision making methods, they might not result in consistent performance improvements. Consequently it is much harder to argue (and even harder to demonstrate empirically) that management teams experience learning curves in their strategic decision making experience. Our first research question is thus, "Does team performance at strategic decision making tasks improve with experience, consistent with a learning curve hypothesis?"

Our second line of inquiry relates to the role of prior study or training in shaping learning curves in strategic decision making. Many organizations invest in research prior to commencing production (what Pisano calls "learning before doing") in order to either increase their initial efficiency, improve their rate of improvement, or both (Pisano, 1994, 1996, 1997; Argote, 1999; Carrillo and Gaimon, 2000). If management teams do experience learning curves in their strategic decision making abilities, is it possible for this process to be accelerated through study or training? Could learning before doing enable the learning curve to begin at a better starting point, essentially improving the intercept of the curve? If so, would the rate of learning (the slope of the learning curve) be unaffected? Or can learning before doing improve the learning rate, making the slope of the learning curve steeper? Our second research question is thus, "If teams experience learning curves in their strategic decision making experience, does learning before doing impact the intercept or slope of such curves?"

To address the research questions above, we utilize a carefully controlled experimental design that allows us to compare multiple learning curves on a strategic decision making task, while collecting accurate performance measures, and controlling for the learning context. Though there is abundant evidence for organizational learning curves in production processes, and some evidence for learning before doing, our research extends prior work by examining the impact of learning by doing and learning before doing on teams' performance at a strategic decision making task. Our research also extends prior work in psychology on team problem solving. Though psychologists have examined how teams perform on complex problem solving relative to individuals (e.g., Hill, 1982; Michaelsen, Watson and Black, 1989), the mechanisms by which they exchange information and resolve conflict (e.g., Pelled, Eisenhardt and Xin, 1999), and develop a group-level transactive memory (Wegner, 1987), there has not (to our knowledge) been any group-level experimental research comparing learning by doing and learning before doing. In the first section, the paper presents the underlying theoretical perspectives and develops the research hypotheses. In the second and third sections, we describe the methods of our study and the results obtained. The fourth section discusses the meaning of our results, and their implications for future research and practicing managers.

THEORETICAL PERSPECTIVES

Learning by Doing

The learning curve, or "learning by doing," refers to the process by which an individual or group increases their performance with experience in a task (Arrow, 1962). As articulated by Levitt and March (1988), organizations learn "by encoding inferences from history into routines that guide behavior" (pg. 320), and one of the purest examples of organizational learning is manifested in the effects of cumulative production on cost and productivity. Organizations experience productivity improvements as a "consequence of their growing stock of knowledge" (Dutton & Thomas, 1984:235), and the application of this knowledge to increase the effectiveness and efficiency of production technologies (Amit, 1986; Hall & Howel, 1985).

Organizational learning scholars typically model the learning curve as a function of cumulative output: performance increases, or cost decreases, with the number of units of production, usually at a decreasing rate (see Figure 1). This pattern has been found to be consistent with production data on a wide range of products and services (Argote, 1993, Baloff, 1971; Hatch & Mowery, 1998; Yelle, 1979), and for a variety of dependent variables, including total costs per unit (Darr, Argote, & Epple, 1995), accidents per unit (Greenberg, 1971), and waste per unit (Mukherjee, Lapre, & Van Wassenhove, 1998).

Insert Figure 1 about here

Though there is a considerable body of research on learning and productivity in groups (see Bettenhausen, 1991 and Williams & O'Reilly, 1998 for reviews), there are very few studies that attempt to analyze group learning curves. Those studies that have used a learning curve framework have found evidence that group learning demonstrates a learning curve pattern similar to that found in studies of individual and organizational learning (e.g. Argote, Insko, Yovetich & Romero, 1995; Guetzkow & Simon, 1955; Leavitt, 1951; Shure, Rogers, Larsen & Tassone, 1962). For example, Argote, Insko, Yovetich and Romero (1995) use a learning curve framework to examine the impact of turnover and task complexity on a group's ability to produce simple (6 steps) or complex (12 steps) products over six 12-minute experimental periods. They found that group performance conformed to learning curve patterns that have been demonstrated at the organizational levels, and that performance was negatively affected by group turnover and task complexity. Another stream of research that considers the role of group experience over time (but does not use a learning curve framework) suggests that as team members accumulate experience working together, they become better able to communicate, and utilize each other's skills, which should lead to improved performance (Amir, 1969; Harrison, Price & Bell, 1998;

Moreland, 1999; Wegner, 1987). We thus begin our study with the traditional learning curve hypothesis:

Hypothesis 1 (H1): Team performance will increase with cumulative experience ("learning by doing").

In the many studies of organizational learning curves, one finding that consistently stands out is the substantial and persistent differences in the rates at which organizations learn (Argote, 1999). Understandably, both managers and scholars are very interested in understanding why one firm reaps great improvement in a process whereas another exhibits almost no learning. Many studies have examined various reasons for this variability in learning rates, including looking at how the firm's learning rate is influenced by process-improvement projects, intentional innovation, or contact with customers and suppliers (Dutton & Thomas, 1984; Levy, 1965; Lapre, Mukherjee, & Van Wassenhove, 2000). One such factor that has begun to receive increasing attention is the degree to which the organization invests in research prior to commencement of production, or "learning *before* doing"

Learning before Doing

Learning before doing is typically considered to include all of the activities that the organization has to go through before implementing a new organizational process or routine. These activities include the training of personnel, research and development (R&D), pilot production, and other experiments that might occur prior to commencement of the production process. Such activities might influence initial performance, the rate of improvement, or both.

While most productivity studies that use a learning curve framework assume that the initial task performance is given, other research employing a problem-solving framework explicitly poses learning as a function of initial performance. In the problem-solving framework, it is argued that learning is triggered by a gap between potential and actual performance (Tyre and Orlikowski,

1993, 1996). A process has a target performance that is not usually met at the initial performance, but as the stock of knowledge about the process is acquired the performance gap (potential minus actual performance) is reduced (Cyert & March, 1963; Levy, 1965; Tyre & Orlikowski, 1993, 1996). These studies do not always utilize a learning curve framework, and have used a variety of innovative ways of measuring learning (Tyre and Orlikowski, 1996).

Levy's (1965) study was one of the first to combine both a productivity learning curve approach with a problem-solving approach. He derived a learning function based on the assumption of gap performance, and concluded that the rate of learning was proportional to the amount a process could improve (1965:B-138). He also posed that learning could be divided into three classes: 1. *Planned or induced* learning in which the organization sponsors processes or cost improvements that occur prior to the production starts (i.e., learning *before* doing); 2. *Random or exogenous* learning as knowledge received unexpectedly from the environment (e.g., spillovers); and 3. *Autonomous* learning as "the improvement due to on-the-job learning or training of employees" (i.e., learning by doing) (p. B-140). Levy also suggests that planned pre-production learning is "inversely related to the rate of learning but enhances the firm's initial efficiency" (p. B139)-- a sentiment echoed in Pisano's second scenario, which is described later in this paper. To our knowledge, the relationship between learning before doing and learning by doing was not as explicitly examined again until Pisano's (1994) study.²

² Yelle notes in his 1979 review two other early studies that examined the relationship between a learning curve's intercept and its slope. As discussed by Yelle, Cole's 1958 study of non-aircraft companies found no relationship between the learning curve intercept and slope. Also referred to by Yelle is Baloff's 1967 study which used data on manufacturing experience and experimental studies in group learning to estimate learning curve parameters. Baloff assumed a relationship between the slope and the intercept, but his results were inconclusive.

Pisano (1994, 1997) points out that pre-production activities (learning before doing) can help the organization to achieve a better initial performance, or to acquire knowledge that will be used to solve problems that may occur during the production process. Therefore, these activities might impact the learning curve intercept, the learning rate, or both. Exploring the different possible combinations of these effects, Pisano constructs four possible scenarios of how learning before doing may impact the learning curve (see Figure 2).

Insert Figure 2 About Here

The first scenario discussed by Pisano (1997) is that learning before doing has no impact on the starting point of the learning curve, but it accelerates the rate of improvement. For instance, study of the process prior to its implementation might enable a deeper understanding of the process that subsequently leads to improved problem solving on the factory floor. Pre-production efforts might also lead to the development of processes that are more amenable to improvement over time by incorporating such features as flexible machinery and quality control checkpoints. Pisano terms this "design for improvability."

Lieberman found evidence for this effect in his 1984 study. Lieberman analyzed the effect of several production process factors on the learning rate, and found that R&D expenditures had an influence on the learning curve but that "...the level of R&D expenditure does not translate directly into an average rate of cost reduction. Rather, R&D tends to accelerate the learning process and to increase the steepness of the learning curve. One interpretation is that the learning curve relation defines the overall potential for process improvement and hence the manner in which returns to R&D diminish over time... Learning-by-doing and learning-by-spending on R&D are closely linked in practice..." (p. 226-227). As learning by spending on R&D can be considered a form of learning before doing, Lieberman's conclusion indicates that learning

before doing (at least through R&D investment) might make the slope of the learning curve more steep.

The second scenario described by Pisano is one where learning before doing results in a better starting point, but decreases the rate of improvement over time. In essence, the pre-production process development results in solving problems that would normally be corrected during production. However, having picked the "low-hanging fruit," subsequent improvements are more difficult to generate (Pisano, 1997: 40). In this instance, learning before doing does not enhance the maximum performance achievable from a process, but rather allows production to begin at a point that is already closer to the maximum. One can think of this type of learning before doing effect as jumpstarting the organization to a later point in the learning curve. This scenario is consistent with Levy's (1965) arguments about the impact of planned or induced learning.

In the third scenario proposed by Pisano, efforts invested in learning before doing improve the starting point of the learning curve, but have no effect on the learning rate. In this scenario, learning before doing solves problems that cannot be solved through learning by doing. For example, Pisano notes that learning before doing may lead to the adoption of a technology that is superior to previously considered technology. Once committed to either technology, process improvements are incremental, but since one technology is fundamentally better than the other, the entire learning curve is shifted.

This scenario is captured by arguments made by Young in 1993. He notes that learning by doing alone does not explain economic development, and that a problem in the economic models of learning-by-doing is that they assume "that the potential productivity gains from learning are essentially unbounded... [there should be a bound to learning, since there is a] recurring pattern of technological improvement and stagnation apparent in pre-modern history" (p. 444). Or as Young writes: "...in the absence of further costly invention, learning by doing is fundamentally

bounded and that, in the absence of further development, most of the new technologies developed by research are broadly inferior to existing productive techniques.” (p. 465). Thus Young argues that while learning by doing is incremental and bounded, investments in invention and development may enable improvements to recommence. Unlike the previous scenario, in this scenario learning before doing has the potential to increase the maximum achievable performance.

In the fourth scenario discussed by Pisano, learning before doing both improves the starting point of the learning curve, and accelerates the learning rate. In this scenario, the technology developed or chosen through learning before doing not only has a better initial performance, but has greater potential for improvement over time. Pisano gives as an example a situation where a craft technology is replaced by a technology based on well-understood scientific principles. He notes (1997:41), "Because the science-based process can be better controlled, initial performance--in particular, yield--is likely to be better. At the same time, this superior understanding and control over the process provides a foundation for improvement."³ Learning before doing has the most potential for increasing maximum achievable performance in this scenario.

In his empirical studies of learning before doing, Pisano (1994, 1997) found that learning before doing was far more effective in chemicals-based pharmaceuticals than in biotechnology-based pharmaceuticals. He concludes that there is much more to be gained through pre-production research and other activities in fields where there is a well-developed knowledge base (as in chemicals-based pharmaceuticals), than in fields in which the underlying knowledge base is not

³ An example of such a situation offered by Pisano (1997) is the adoption of the Pilkington float glass process, which was based on deep knowledge of the molecular structure of glass, and resulted in more consistent quality and lower manufacturing costs than plate glass processes.

yet well developed (as in biotechnology-based pharmaceuticals). The repository of scientific knowledge underlying a field provides a wellspring of information that may be of use in solving production problems before production commences. In absence of a significant repository, the firm's efforts might be better spent on experiential learning (learning by doing) rather than pre-production activities.

Relating this back to strategic decision making reveals a key tension in the strategic management literature. Some management scholars would argue that there is a highly developed knowledge base in strategy that individuals can draw upon in developing their strategic problem-solving abilities. This knowledge base comprises both theoretical and empirical work, and dates back at least as far as Sun Tzu's Art of War (which was written approximately 2500 years ago). Presumably, studying this knowledge base would help managers benefit by the experiential learning of others that have gone before them, i.e., they can see what has worked -- or not worked -- before (Makridakis, 1996). Furthermore, a growing body of theory and evidence suggests that the use of a science-based knowledge repository might help an individual make better strategic decisions because it offers theories about *why* something should work or not, in addition to the experiential knowledge of *whether* something works (Gavetti & Levinthal, 2000; Nelson, 1982; Vincenti, 1990). Theories might enable managers to better extrapolate results obtained in the past to the situation at hand. The combination of prior results and theories about why particular strategies might succeed or fail, should thus help managers target strategic actions that are more likely to be successful, and avoid wasted effort on strategic actions that have been proven unsuccessful.

Other scholars might argue that prior writings on strategy have little or no bearing on any particular strategic decision making context an individual or firm faces. Most strategic decision making occurs within a dynamic context wherein both the environment and competitors are constantly changing in response to the individual or firm's actions (Terreberry, 1968). Both prior

results and theories based on those results are likely to be irrelevant beyond the context in which they were originally obtained. Too much preparation or planning could, in fact, hamper the firm by stifling its ability to adapt (Mintzberg, 1993). Consistent with this, Eisenhardt and Tabrizi (1995) found that learning by doing was significantly more useful than learning before doing in launching uncertain new products in the computer industry. They argued that in uncertain and volatile settings, improvisational approaches that combine real-time learning with experiential testing was much more effective than planning in advance. Of course, the tricky part here is that learning before doing need not be the same thing as learning by planning, despite their seeming equivalence in the literature. It might be possible, for example, for managers to invest effort in learning to think more strategically in a general sort of way. Managers could, perhaps, study methods of being strategically adaptive rather than utilizing the science of strategy to create a structured plan for the future. In so doing, managers might increase their absorptive capacity and dynamic capabilities, making them better able to respond to a wide variety of strategic situations (Cohen & Levinthal, 1990; Teece, Pisano & Shuen, 1997; Zollo & Winter, 2002). This brings us to our second set of questions: Does studying strategy improve management's ability to behave strategically? And if so, does this study impact only their initial ability, or also their rate of improvement?

Utilizing the four scenarios developed by Pisano, we form several hypotheses about the relationship between learning before doing and learning by doing. First, we hypothesize that investments in learning before doing improve the starting point of the learning curve (the intercept), with the null hypothesis being that learning before doing has no impact on the starting point of the learning curve⁴:

⁴ Note that a competing hypothesis that learning before doing negatively impacts the intercept of the learning curve could also be posed, but has no basis in the prior work on learning before doing.

Hypothesis 2 (H2): Use of learning before doing will significantly increase teams' initial performance at a strategic decision making task.

Next we pose the following two competing hypotheses about how learning before doing will impact the rate of learning (the slope of the learning curve), with the null hypothesis being that learning before doing has no impact on the rate of learning:

Hypothesis 3a (H3a): Use of learning before doing will significantly decrease the relationship between team performance and cumulative experience.

Hypothesis 3b (H3b): Use of learning before doing will significantly increase the relationship between team performance and cumulative experience.

METHODS

The data used in the analysis was collected using an experimental design in which ten teams of three subjects each performed strategic decision making tasks repeatedly for a total of ten hours each. By tracking performance carefully over time, we are able to assess each team's learning curve and compare their intercepts and slopes.

The Strategic Decision Making Task

The first challenge of constructing the experimental design was to identify a decision making task that would 1) be explicitly strategic in nature, and rely on problem-solving skills rather than motor skills, 2) enable repetition without one ideal solution, 3) enable controlling for difficulty, 4) enable accurate and consistent performance appraisal, and 5) permit extended learning. After consideration of many alternatives, we decided to have the teams play games against a computer. We evaluated more than 300 game possibilities and narrowed the list down by consulting several games review sources, including [A History of Traditional Games](#) (Masters, 1999), [A History of](#)

Card Games (Parlett, 1990), The Oxford History of Board Games (Parlett, 1999), and a very extensive index of card games by type, developed by McLeod (1998).

We finally chose the game of Go (or "Weiqi"). Go is an ancient strategic board game that is one of the most highly revered of Asian games, and is the national game of Japan. The game is played with stones on a grid, where the objective is to surround and capture territory. The game is generally considered to be far more subtle and complex than Chess. According to legend, Go was invented by emperor Yao of China (2357-2256 BC) for the purpose of developing the mind of his son Tan Chu. The game has been widely used by military leaders of China, Korea, and Japan for the development of strategic skills, indicating that the game is believed to draw on the same strategic problem solving abilities required in an actual competitive situation. The game is also played by Buddhist monks as a route to enlightenment, with some individuals dedicating their entire lives to developing a mastery of the game (Parlett, 1999), demonstrating both the game's robustness, and its ability to capture and hold the attention of participants for an extended period of time.⁵

In addition to meeting our criteria above, the game of Go also provided a number of additional benefits: a) it was unfamiliar to most of the individuals who responded to the solicitation (more information on screening is provided below); b) it is very simple to learn, yet very difficult to master; c) and a game can be completed in ten minutes or less (using a 13 X 13 board grid).

Solicitation and Screening

The participants were solicited with flyers that specified that subjects would perform basic problem-solving tasks. The flyers did not indicate that the experiment would entail playing games (to avoid creating a response bias). The flyers also offered a \$100 payment to subjects

⁵ There are numerous resources for further information on Go; a good starting point may be found at www.gobase.org

upon completion of the experiment period, with a stipulation that no partial participation would be compensated, and that no subject would be permitted to participate in more than one experiment period.

In order to ensure that the teams all began at the beginning of the learning curve, respondents to the flyer were screened to avoid inclusion of any participant with experience with Go. The respondents were asked a series of questions, including those asking for demographic information (e.g., age, gender, occupation), and about a variety of activities in which they engage (e.g., reading the newspaper, tennis, golf, etc.). A wide variety of questions were asked to avoid signaling the respondents about the particular activities of interest.

Experimental Condition

Respondents were randomly assigned to teams (three subjects per team, for a total of ten teams). Each team had their own computer, and played the game against the computer repeatedly for five hours a day, for two consecutive days. All teams were subjected to identical experimental conditions. Prior to commencing play of the games, all individuals were asked to complete an entrance survey (which collected basic demographic information, personality assessment information, and prior game experience), and were given instruction sheets for playing the games (discussed in greater detail in the next section). Teams were told that the experiment was a learning study, and that their objective was to work together to get as good at the game(s) as possible. Teams were also instructed that each game was to be a group endeavor (i.e., delegation of game playing among the team members was not allowed), however they were also instructed not to speak to any members of any *other* team, to prevent information leaks between teams.

Teams were permitted to play at their own speed, and were given detailed score sheets to track the time at which they began and ended each game, their score on each game played, and the computer's score. Three monitors observed the teams at all times to ensure that teams adhered to

the rules of the experiment and to note any unusual activity. After completion of the ten hours, the individuals were asked to complete an exit survey about how the team interacted during the experiment.

Learning Before Doing

In the beginning of the experiment, each team received a set of rules and game suggestions (three pages total). However, there was no stated requirement that they read these rules and suggestions before starting play, nor was there any time designated as "study time" for the teams to read the sheets--thus the use of instructions was left to the teams' discretion. Learning before doing was observed by a survey questionnaire administered at the end of the experiment. Two survey questions required the subjects to rate on a five-point scale their degree of study and use of the instructions provided by the examiners before playing the games. The Pearson correlation for the two items is 0.35 (alpha = .52) and significant at $p < 0.005$.

We followed the procedure proposed by James, Demaree and Wolf (1984) to assess the agreement among the responses made by a group of people on a single variable (multiple item estimator for interrater reliability):

$$r_{WG(J)} = \frac{J [1 - (\overline{s_{xj}^2} / \sigma_{EU}^2)]}{J [1 - (\overline{s_{xj}^2} / \sigma_{EU}^2)] + (\overline{s_{xj}^2} / \sigma_{EU}^2)}$$

Where $r_{WG(J)}$ is the within-group interrater reliability for judges' mean scores based on J essentially parallel items, $\overline{s_{xj}^2}$ is the mean of the observed variances on the J items, and σ_{EU}^2 is the variance of a rectangular or uniform distribution (Mood, Graybill and Boes, 1974).

All of the teams had very high inter-rater agreement reliability (average of .87), every team surpassing the recommended cutoff of .60 (Edmonson, 1996) (see Table 1). We therefore created a variable for each team with the average of the individual responses, named LBD (learning before doing). This variable was then divided into dummy variables representing high

levels of learning before doing (HLBD) and low levels of learning before doing (LLDB) using a median split (median=2.57).

Table 1 About Here

Dependent Variable

Each time the team completed a game, they recorded their score. These scores range from zero to 149, with many instances of zero scores, and no instance of a perfect (169) score. It is important to note that in production studies, it is unusual for firms to experience significant decreases (lapses in their performance) in their production for the reasons discussed in the beginning of the paper: incremental improvements are embodied in machinery, configurations, procedures, etc. so their effects endure over time. Therefore production learning curves tend to demonstrate fairly consistent improvement. However, in robust computer versions of the game of Go, the computer is a very skilled player, and even teams that have acquired good Go playing skills (and demonstrate increasing moving averages of their score) will occasionally earn very poor scores. This is particularly likely to occur when teams experiment with new strategies. Thus, as would be expected in actual strategic decision making situations, there is more variability in the learning curves for the game of Go than one would expect in production learning curves.

Overview of Analyses

The standard form of the learning curve is formulated as:

$$y = ax^{-b},$$

Where y is the number of direct labor hours required to produce the x th unit, a is the number of direct labor hours required to produce the first unit, x is the cumulative number of units produced, and b is the learning rate. By rewriting the formula in logarithmic form we obtain the following formula which enables the learning coefficient to be obtained through linear regression:

$$\log y = \log a - b \log x.$$

Our specification uses this standard formulation but with one exception; rather than modeling the outcome as a decrease in labor hours, we model the outcome as an increase in scores, resulting in a negatively accelerated increasing curve. Our learning curve form is thus:

$$y = ax^b, \text{ or } \log y = \log a + b \log x.$$

In learning rate studies, one typically regresses the dependent measure on the number of learning trials. In this study, this translates into regressing game score on the number of games played (each game played is a learning trial). Though total play time was carefully controlled in our study, there is variability in the number of games played since teams were allowed to play at their own speed. Some teams tended to play quickly, in a trial and error fashion, while others played more deliberately, discussing each move in advance. There was also variability in speed of play over time for individual teams (i.e., a team might play quickly for awhile, and then take a more measured approach).

To standardize the number of observations across teams and permit comparison across the multiple learning curves, we employ an analytical approach similar to that used by Darr, Argote and Epple (1995). They examined learning curves of individual pizza stores and the effect of belonging to a particular franchise by gathering weekly data on the pizzas made and average cost per unit. By aggregating the number of pizzas per week for each store, franchise, and across all franchises, they were able to analyze the degree to which learning occurred with experience at the store, franchise, and interfranchise level despite variation in the number of pizzas produced. In a similar fashion, we aggregated our data to the hour level, and regressed the average score a team achieves on Go games in a given hour t on the cumulative number of Go games played by the end of the previous hour (i.e., cumulative number of Go games played from the beginning of the first hour to the end of the $t-1$ hour). This allows us to control for both time and the number of games played by any particular team over time, while also yielding an equal number of

observations (ten) for each team, for a sample size of 100. As mentioned above, there was variability in the speed and deliberateness of play both across teams and over time. To control for effects due to speed of play and exploratory trial-and-error strategies, we include the number of Go games played in the hour as a control variable. We then test whether the use of learning before doing significantly influences the intercept (initial game playing ability) and the learning rate. This is accomplished by entering both a dummy variable (0,1) for high use of learning before doing, and interaction terms for high use of learning before doing X cumulative games, and low use of learning before doing X cumulative games. We use the following symbols for our variables:

s_{it} -- average score earned by team i in hour t on Go games

G_{it} -- number of Go games played by team i in hour t

Q_{it-1} -- cumulative number of Go games played by team i from the beginning of the first hour through the end of hour $t-1$

H -- dummy variable for high use of learning before doing

Our most basic model was:

$$(1) \text{Ln } s_{it} = b_0 + b_1 G_{it} + b_2 \text{Ln } Q_{it-1} + e$$

If b_2 is statistically significant and positive in the first model, then overall the teams significantly improved their performance as they played games, indicating a learning curve effect (learning by doing) consistent with H1. To control for the impact of different play strategies (e.g., fast versus slow), the model includes a control variable ($b_1 G_{it}$) for the number of Go games played in the hour. In the second model, we add the dummy variable for high use of learning before doing:

$$(2) \text{Ln } s_{it} = b_0 + b_1 G_{it} + b_2 \text{Ln } Q_{it-1} + b_3 H + e$$

If b_3 is statistically significant, it indicates that overall, high use of learning before doing significantly impacts the average of the scores achieved in a given hour, but this model does not allow us to separate the effect of high use of learning before doing on the intercept and on the learning rate. That is achieved in the fourth model, which includes both the dummy variable for high use of learning before doing, and an interaction term whereby the dummy for high use of learning before doing is multiplied by the cumulative games variable:

$$(3) \text{Ln } s_{it} = b_0 + b_1 G_{it} + b_2 \text{Ln } Q_{it-1} + b_3 H + b_4 H \text{Ln } Q_{it-1} + e$$

In this model, if b_3 is statistically significant and positive, then high use of learning before doing increases the initial performance (the intercept of the learning curve), consistent with H2. If b_4 is statistically significant, then use of learning before doing significantly impacts the rate of learning by doing (the slope of the learning curve). Specifically, if b_4 is negative, then teams that invest in high use of learning before doing have a weaker relationship between cumulative experience and performance (weaker learning by doing effect, or a less steep learning curve slope) than teams that invest in lower levels of learning before doing, consistent with H3a. By contrast, if b_4 is significantly positive, then teams that invest in high use of learning before doing have a stronger relationship between cumulative experience and performance (stronger learning by doing effect, or a steeper learning curve slope) than teams that invest in lower levels of learning before doing, consistent with H3b.

Relating these back to Pisano's scenarios: If b_3 is not significant, and b_4 is significantly positive, then learning before doing did not impact initial performance, but enhanced the learning rate, consistent with Pisano's first scenario; if b_3 is significant and positive, and b_4 is negative, then learning before doing increased initial performance, but dampened the learning rate, consistent with Pisano's second scenario; if b_3 is significant and positive, but b_4 is not significant, then high use of learning before doing has increased the initial performance, but has not significantly impacted the slope of the learning curve, consistent with Pisano's third scenario; and finally, if

both b_3 and b_4 are significant and positive, then high use of learning before doing has increased both the initial performance and the learning rate, consistent with Pisano's fourth scenario.

RESULTS

After ten hours, the teams had played an average of 142 games of Go each. Though teams could have attempted to delegate play to an individual member, the monitors noted that the team members interacted vigorously over the entire experimental period. Though some individuals became more involved with the game than others, team performance appeared to almost always be a collective effort. Within the teams, members actively discussed potential moves, evaluating what appeared to be successful or unsuccessful, and formulating strategies. Though the individual holding the computer mouse would execute the move, most move decisions were arrived at through group interaction. Many teams demonstrated a pattern whereby control of the mouse was rotated from individual to individual. Furthermore, very often the individual controlling the mouse was not the most active proponent of the next move (that is, control over the mouse did not appear to indicate decision authority over game moves). For most teams, the emotional involvement with the game appeared to escalate over time. Discussion among team members sometimes became quite heated, with individuals occasionally voicing anger or frustration. Particularly high scores often resulted in an eruption of cheers.⁶

For each team, the average of the scores for each hour, and cumulative number of games played by the end of each hour were tallied (descriptive statistics and correlations are provided in Table 2). Notably, average score is positively and significantly related to the number of games played per hour, and the cumulative experience variable (learning by doing). Learning by doing is not

⁶ It is interesting to note that during the Xihan Dynasty, Go (known as Weiqi) was widely criticized as being addictive. This may explain why teams appeared to become increasingly interested in the game over time, rather than becoming bored.

significantly correlated with learning before doing. Ordinary least squares regression was used to estimate the models described in the methods section (see Table 3). We tested for first order autocorrelation in the residuals, and found no significant autocorrelation. Scatterplots of the residuals were also examined, and they did not indicate autocorrelation or heteroscedasticity.

Table 2 and 3 around here

In Table 3, Model 1 presents the base model, in which the control variable for games per hour and the cumulative games variable are entered. The overall model explains substantial variance (Adjusted $R^2 = .447$). The control variable for games per hour is not significant, but the cumulative games variable is positive and significant (.167, $p < .001$), indicating a learning curve effect and supporting H1. Model 2 examines the impact of the learning before doing variable on overall performance. The model indicates that adding the variable improves the model significantly (change in R^2 was .03, $p < .05$). The learning before doing variable is positive and significant (.135, $p < .05$), indicating that overall, high use of learning before doing improves performance. In Model 3 we include the interaction term in order to assess separately how learning before doing affects the initial performance at start of play, and how it influences the rate of learning by doing. The inclusion of the interaction term significantly increases the explanatory power of the model (change in R^2 was .03, $p < .05$). This model shows that learning before doing increases the initial scores of the teams significantly (.409, $p < .005$). Thus learning before doing does appear to have a significant and positive impact on initial performance, supporting H2. The results also show, however, a significant and negative coefficient for the interaction term (-.09, $p < .05$), indicating that high use of learning before doing dampened the slope of the learning curve, supporting H3a.

Overall, then, we found that high use of learning before doing increased initial performance, but decreased the rate of learning by doing, consistent with Pisano's second scenario. To better

illustrate these results, the functions are plotted on scatterplots of the combined data provided in Figure 2.

Insert Figure 2 About Here

DISCUSSION

The analysis presented here indicates that both strategies of learning (learning before doing and learning by doing) had a significant impact on the performance of the teams, as demonstrated by the significant R^2 increment for the models incorporating cumulative output, use of learning before doing, and the interaction term. In addition, though the models indicate that the most substantial portion of the variance is explained by learning by doing, learning before doing does have a significant and positive impact on initial performance. However, learning before doing seems to somewhat dampen the learning rate.

This finding is consistent with both Levy's (1965) argument, and with Pisano's (1997) second scenario. In this scenario, he posits that the effect of learning before doing may be to correct most of the straightforward problems (the "low hanging fruit") before starting the task itself, thus leaving less room for improvement through learning by doing (1997: 40). Similarly, Epple, Argote and Devadas (1991) note that the reason learning curves are negatively accelerated is because learning by doing yields large gains as experience and knowledge accumulate, but the *rate* of knowledge accumulation declines as the stock of knowledge grows (1991; p. 67). Learning before doing may increase the stock of knowledge prior to start the task, conceivably pushing the learner to start at a point on the learning curve where the slope is less steep.

To the degree that the game of Go realistically tests strategic thinking abilities, our results suggest that that teams experience learning curves in their strategic decision making experience,

and that investment in prior study of strategy may enable teams to "jumpstart" to a later point on the learning curve. Future research will need to test the generalizability of this finding, and the role of individual learning, team composition factors, etc. before much can be said about implications for practice, but this research represents an important step toward understanding whether and how teams learn to be strategic.

These results also provide some evidence for the problem-solving framework in which learning is posited as being triggered by a gap between the potential and actual performance. Though no team in our sample ever achieved a perfect score (169 in a 13 X 13 game of Go), it can be argued that 169 was the potential performance of a team. By the end of the experimental period, most teams were regularly beating the computer (earning a score of 85 or greater) and many teams earned scores in excess of 100, with one team achieving a single incidence of a score of 149 (the maximum score achieved over the experimental period by any team). It can be argued, then, that our study demonstrated that teams with poorer initial performance at the game (and thus a larger gap between actual and "potential" performance) had more learning triggered through experience than teams that had greater initial performance.

Therefore, we can argue that the teams that had a lower initial performance at the game, which translate in a larger gape between actual and "potential" performance,

Limitations and Suggestions for Future Research

Though this research yields fairly robust results for learning by doing and learning before doing for our particular strategic decision making scenario, it still leaves open a number of questions. For instance, the game of Go relies on very abstract principles of spatial strategy which have applicability to a wide range of Go game scenarios. Do such abstract principles apply in the real world of business competition, and does the ability of management teams to get better at strategic decision making depend on their existence? Additionally, in our version of the game of Go the competitor (the computer) would respond based on the team's moves, but could not learn

over time. How would our results have changed had the competitor learned to anticipate or imitate the team's strategies? Furthermore, Pisano's work raises the point that the value of learning before doing is related to the state of existing knowledge that may be tapped (i.e., in fields that have relatively well developed knowledge bases that may be studied, there is more potential for learning before doing). The game of Go has an incredibly long history, but our instruction sheets provided only basic strategy suggestions and the rules of play. We did not permit our subjects to use the Internet or any other resource to research the game. A very interesting experiment might utilize different games that had knowledge bases that were clearly in different states of development, and allow teams to research games at their own discretion. The challenge, however, would be to ensure comparability of difficulty and performance outcomes.

There are a variety of other factors that may condition the importance of learning before doing, including 1) team specific factors: Do the characteristics of some teams make them better at learning before doing, while the characteristics of others make them better at learning by doing? Our study demonstrates that some teams are certainly more *inclined* to use learning before doing versus learning by doing (but the individuals were randomly assigned to the teams), but we have not examined what individual or team characteristics might make this true, and we do not have data that would permit assessing whether such teams chose the strategy that maximized their performance; 2) task specific factors: Are different learning strategies more fitting for different types of tasks?; 3) situation specific factors: Does the context of the learning environment affect the usefulness or practicality of a learning strategy? For instance, our teams were given a finite time in which to play, and were required to remain at their stations with their team at all times. If teams were given an extended amount of time to play a fixed quantity of games and were allowed to leave their stations, team members might have attempted to research the game. Such a scenario may be more analogous to many real-world strategic decision making settings.

In conclusion, this study indicates that teams do exhibit learning curves in their performance at a strategic decision making task, and that learning before doing may impact the efficacy of learning by doing. This indicates that future learning curve studies should attempt to control for this crucial source of variation in learning rates, and also indicates that future research into the factors conditioning the intricate roles of each strategy should prove as practical to managers as it is interesting to researchers.

REFERENCES

- Amir, Y. 1969. *Contact hypothesis in ethnic relations*. Psychological Bulletin, 11: 319-342.
- Anzai, Y., and Simon, H. A. 1979 "The theory of learning by doing." Psychological Review, 86 (2): 124-140.
- Argote, L. 1993 "Group and organizational learning curves: Individual, system and environmental components." British Journal of Social Psychology, 32:31-51. 1999 Organizational learning: Creating, retaining and transferring knowledge. Boston: Kluwer Academic Publishers.
- Argote, L. and Epple, D. 1990 "Learning Curves in Manufacturing." Science, 247: 920-924.
- Arrow, K. J. 1962 "The Economic Implications of Learning by Doing." Review of Economic Studies, XXIX: 5-23.
- Bahk, B. H., Gort, M. 1993 "Decomposing Learning by doing in new plants." Journal of Political Economy 101:561-583.
- Baum, J.A., Li, S.X., Usher, J. M. 2000 "Making the next move: How experiential and vicarious learning shape the locations of chains' acquisition." Administrative Science Quarterly, 45:766-792.
- Bewonder, B. and Miyake, T. 1994 "Creating and sustaining competitiveness: An analysis of the Japanese robotics industry." International Journal of Technology Management, 9:575-612.
- Carrillo, J.E. and Gaimon, C. 2000 "Improving manufacturing performance through process change and knowledge creation." Management Science, 46:265-288.
- Cyert, R.M and March, J.G. 1963. "A behavioral theory of the firm." New York: Prentice Hall.
- Darr, E. D., Argote, L. and Epple, D. 1995 "The Acquisition, Transfer, and Depreciation of Knowledge in Service Organizations: Productivity in Franchises." Management Science, 41: 1750-1762.
- Day, G.S., and Montgomery, D.B. 1983 "Diagnosing the experience curve." Journal of Marketing, 47: 44-58.
- Edmondson, A. 1996 "Learning from mistakes is easier said than done: Group and organizational influences on the detection and correction of human error." Journal of Applied Behavioral Science, 32:5-28.
- Eisenhardt, K. & Tabrizi, B. 1995. Accelerating adaptive processes: Product innovation in the global computer industry. *Administrative Science Quarterly*, 40:84-110.
- Epple, D., Argote, L. and Devadas, R. 1991 "Organizational Learning Curves: A method for investigating intra-plant transfer of knowledge acquired through learning by doing." Organizational Science, 2: 58-70.
- Gavetti, G. & Levinthal, D. 2000. Looking forward and looking backward: Cognitive and experiential search. *Administrative Science Quarterly*, 45:113-137.
- Haleblian, J. and Finkelstein, S. 1999 "The influence of organizational acquisition experience on acquisition performance: A behavioral perspective." Administrative Science Quarterly, 44:29-57.
- Hatch, N. W. and Mowery, D. C. 1998 "Process Innovation and Learning by Doing in Semiconductor Manufacturing." Management Science, 44: 1461-1477.
- Hill, G.W. 1982 "Group vs. individual performance: Are N+1 heads better than one?" Psychological Bulletin, 91:517-539.
- James, L. R., Demaree, R. G. and Wolf, G. 1984 "Estimating within-group interrater reliability with and without response bias." Journal of Applied Psychology, 69: 85-98.

Excluido: ¶
Beyond relational demography: Time and the effects of surface- and deep-level diversity on work group cohesion .
Academy of Management Journal; Mississippi State; Feb 1998; David A Harrison; Kenneth H Price; Myrtle P Bell; -
¶
Volume: ... [1]

- Lapre, M., Mukherjee, A., Van Wassenhove, L. 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 Science, 11:B136-B154.
- Lieberman, M. B. 1984 "The Learning Curve and Pricing in the Chemical Processing Industries." Rand Journal of Economics, 15: 213-228.
- Lounamaa, P. H., March, J. G. 1987 "Adaptive coordination of a learning team." Management Science, 33: 107-123.
- Masters, J. 1999 "A history of traditional games." <http://web.ukonline.co.uk/james.masters/TraditionalGames/TraditionalGamesHome.htm>
- McLeod, J. 1998 "Index of card games by type." <http://www.pagat.com/index.html>
- Michaelson, L.K., Watson, W.E, and Black, R.H.A 1989 "Realistic Test of Individual Versus Group Consensus Decision Making." Journal of Applied Psychology. 74: 834-839
- Mintzberg, Henry; The pitfalls of strategic planning; *California Management Review*, Berkeley; Fall 1993; Vol. 36, Iss. 1; pg. 32.
- Moreland, R. 1999. Transactive memory: Learning who knows what in work groups and organizations, in L.L. Thompson, J.M. Levine & D.Messick (eds) *Shared cognition in organizations*: 3-31. Mahwah, NJ: Lawrence Erlbaum Associates.
- Mood, A.M., Graybill, F.A., and Boes, D.C. 1974 *Introduction to the theory of statistics*. New York: McGraw Hill.
- Mukherjee, A., Lapre, M., Van Wassenhove, L. 1998 "Knowledge driven quality improvement." Management Science, 44: S35-S49
- Nelson, R. 1982. The role of knowledge in R&D efficiency. *The Quarterly Journal of Economics*, 97:453-570.
- Parlett, D. 1999 *The Oxford history of board games*. New York: Oxford University Press.
- Pelled, L.H., Eisenhardt, K. and Xin, K.R. 1999 "Exploring the black box: An analysis of work group diversity, conflict, and performance." Administrative Science Quarterly, 44:1-28.
- Pisano, G. P. "Knowledge, integration, and the locus of learning: an empirical analysis of process development." Strategic Management Journal, 15: 85-100. 1994
- _____. "Learning-before-doing in the development of new process technology." Research Policy, 25: 1097-1119. 1996
- _____. *The Development Factory*. Boston: Harvard Business School Press.1997
- Schwenk, Charles, and Richard Cosier 1993 "Effects of consensus and devil's advocacy on strategic decision-making." *Journal of Applied Social Psychology*, 23: 126-139.
- Senge, P. 1990 *The fifth discipline: The art and practice of the learning organization*. New York: Doubleday/Currency.
- Starbuck, W. H. (1992) 'Strategizing in the real world,' International Journal of Technology Management, Special Publication on Technological Foundations of Strategic Management, 1992, 8(1/2): 77-85.
- Teece, D. J., Pisano, G., & Shuen, A. 1997. *Dynamic capabilities and strategic management*. Strategic Management Journal, 18: 509-533.
- Terreberry, S. 1968. "The Evolution of Organizational Environments," Administration Science Quarterly, 12: 590-613.
- Thorndike, E.L. 1898 "Animal intelligence: An experimental study of the associative processes in animals." Psychological Review: Series of Monograph Supplements, 2:1-109.
- Tyre, M.J. and Orlikowski, W.J. 1993 "Exploiting opportunities for technological improvement in organizations." Sloan Management Review, 35(1):13-27. 1996

-
- "*The episodic process of learning by using.*" International Journal of Technology Management, 11:790-799.
- Tyre, M.J. and Von Hippel, E. 1997 "*The situated nature of adaptive learning in organizations.*" Organization Science, 8: 71-83.
- Vincenti, W. 1990. ***What engineers know, and how they know it.*** Baltimore, MD: John Hopkins University Press.
- Von Hippel, E. and Tyre, M. J. 1995 "*How learning by doing is done: Problem identification in novel process equipment.*" Research Policy, 24: 1-12.
- Wegner, D.M. 1987 "Transactive memory: A contemporary analysis of the group mind." In B. Mullen, G. R. Goethals (Eds) *Theories of group behavior.* New York: Springer-Verlag.
- Wright, T.P. 1936 "*Factors affecting the cost of airplanes.*" Journal of Aeronautical Science, 3 (4):122-128.
- Yelle, L. E. 1979 "*The learning curve: Historical review and comprehensive survey.*" Decision Science, 10: 302-328.
- Young, A. 1993 "*Invention and bounded learning by doing.*" Journal of Political Economy 101 (3): 443-472.

Table 1: Multiple Item Estimators for Interrater Reliability

Team	Estimator	Team	Estimator
1	0.79	6	0.91
2	0.83	7	1.00
3	0.91	8	0.89
4	0.82	9	0.91
5	0.86	10	0.74

Table 2: Descriptives and Pearson Correlation

	Mean	S.D.	N	1	2	3
1. Average score _t	56.60	19.91	100	1.00		
2. Games played in hour _t	14.25	9.32	100	.48**	1.00	
5. Cumulative games _{t-1}	50.59	53.26	100	.70**	.63**	1.00
6. High learning before doing	.49	.50	100	.14	.22*	.171

** p < .01 *p < .05

Table 3: Regression Results for the Learning Models

	Model 1	Model 2	Model 3
	B	B	B
Intercept	3.49*** (.062)	3.45*** (.063)	3.34*** (.075)
Games played in hour _t	-.001 (.004)	-.005 (.005)	-.003 (.005)
Cumulative games _{t-1} (learning by doing)	.167*** (.026)	.178*** (.026)	.206*** (.027)
High use of learning before doing		.135** (.061)	.409*** (.122)
High learning before doing X cumulative games _{t-1}			-.090** (.035)
R ²	.46	.49	.52
Adjusted R ²	.45	.47	.50
F	41.04***	30.14***	25.56***
Δ R ²		.03	.03
F increment		5.0**	6.6**

Note: standard error in parentheses

*** p < .01 **p < .05 *p < .10

Figure 1: Standard learning curve forms, (a) cost decreasing and (b) performance increasing

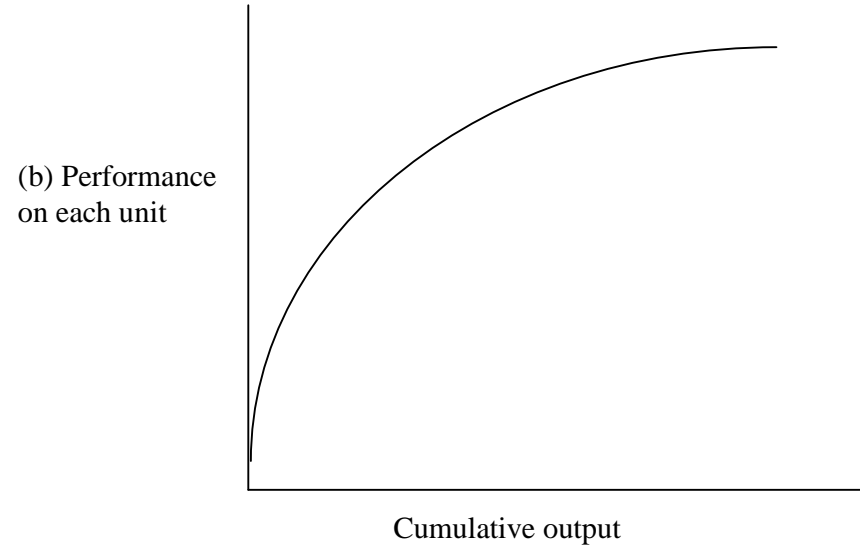




Figure 2: Scatterplots of Learning Curves Under High and Low Levels of Learning Before Doing



Beyond relational demography: Time and the effects of surface- and deep-level diversity on work group cohesion

Academy of Management Journal; Mississippi State; Feb 1998; [David A Harrison](#); [Kenneth H Price](#); [Myrtle P Bell](#);

Volume: 41
Issue: 1
Start Page: 96-107