

Learning Curve? Which one?

Autoria: Paulo Prochno

Abstract

Learning curves have been studied for a long time. These studies provided strong support to the hypothesis that, as organizations produce more of a product, unit costs of production decrease at a decreasing rate (see Argote, 1999 for a comprehensive review of learning curve studies). But the organizational mechanisms that lead to these results are still underexplored. We know some drivers of learning curves (Adler and Clark, 1991; Lapre et al., 2000), but we still lack a more detailed view on the organizational processes behind those curves. Through an ethnographic study, this paper brings a comprehensive account of the first year of operations of a new automotive plant, describing what was going on in the assembly area during the most relevant shifts of the learning curve. The emphasis is then on how learning occurs in that setting. The paper suggests that the overall learning curve is in fact the result of an integration process that puts together several individual learning curves happening at different areas throughout the organization. A model to understand the evolution of these learning processes and their supporting organizational mechanisms is proposed.

Literature review

The first works on organizational learning curves date back to the 1930s, based on a simple yet powerful finding: unit costs decline with cumulative output. This effect happens beyond scale economies or increased inputs of labor and capital, and it reflects learning by doing at the organizational level. This finding has fostered research in different industries and, although the effect varies in magnitude, results give support to the learning by doing hypothesis. In management, learning curves started to be researched more systematically from the late 1980s on. Argote et al. (1990) showed that, although learning by doing does happen, the knowledge acquired through this process depreciates rapidly. They also found out that vicarious learning happens: organizations beginning production later are more productive than those with early start dates. This learning, however, happens only before production starts – after that, organizations do not benefit from learning in other organizations. Epplé et al. (1996) analysed the introduction of a second shift in an automotive plant and found out that virtually all knowledge acquired during the period on one-shift operation was carried forward to the period of two-shift operation in less than two weeks. Argote and Darr (2000) found that production knowledge in pizza stores depreciated less rapidly than service knowledge, and proposed that the difference was due to the fact that production knowledge was embedded in training materials whereas service knowledge was not codified and thus embedded primarily in individuals.

Most research in learning curves results from regression models where output (either cumulative or over time) is correlated with some measure of costs (direct labor hours, cost). In the past few years, authors have tried to open the black box of learning curves, shifting the focus from outcomes to processes. Adler (1990) suggested that shared learning (across the development/manufacturing interface and between plants) is one of the main drivers of productivity improvement; Adler and Clark (1991) studied the effects of engineering changes and workforce training (which for them represents second-order learning) on learning curves and

found out that these effects vary substantially across processes. Lapre et al. (2000) studied the effects of conceptual ('know-why') and operational ('know-how') learning on quality improvement. Pisano et al. (2001) compared learning rates in 16 hospitals performing the same surgery and suggested that learning by doing may be a firm specific capability: if an organization does not develop mechanisms for capturing knowledge and implementing learning, experience may not translate into increased performance. Another departure from traditional learning curves has been to adopt quality measures instead of cost measures (Lapre et al., 2000; Levin, 2000). Levin (2000) shows that quality learning also exhibits a 'learning curve' behavior, except that the curve is more a function of time than a function of cumulative experience, and most of improvements come when a product is first introduced (rather than at subsequent production).

Though most of recent papers stress that understanding the process behind learning curves are crucial, they fall short on bringing a detailed model of those processes. The most detailed papers so far brought at most two process dimensions in their models (e.g. Adler and Clark, 1991; Lapre et al., 2000). Pisano et al. (2001) is a step forward: although their model does not bring process dimensions, they bring qualitative evidence from case studies to explain differences in their quantitative findings. But no paper yet has brought a 'thick' view of the processes driving learning curves. Although most of the conceptual works on learning see it as an essentially social process, the social dimension has been underexplored in learning curve studies.

The setting

This paper adds the social dimension to the analysis of learning curves. The intention is not to show that a learning curve happened, but to show how it happened. In order to be able to see the organizational processes supporting the evolution of learning curves, I developed my study in the early days of a new setting. The setting chosen was a new automotive plant. This plant was the first one that the parent company was building in the country where the plant is located. Most workers, even at the managerial level, had no previous experience in automotive companies. Shop floor workers had no experience in factory work, since the plant was set in a region with no manufacturing tradition. Given that, they had to go through an eight-week training period (organized by the company in association with local institutions) before being eligible for a position in the company.

The plant launched two car models in their first year of operations; both models were already produced in other plants from the parent company. They attempted to transfer some their practices to the new setting. But, given the difficulties in transferring practices (e.g. Szulanski, 1996), the fact that most of the workers had no experience in automotive factories, and the different environmental forces, the practices developed were not a simple and immediate replication of the ones held by the parent company (Fruin, 1998). Some practices were transferred directly, some recontextualized and some developed from scratch.

Methodology

I developed an ethnographic study of the initial period in the life of an automotive plant, with the overall aim of observing the evolution of organizational capabilities. I stayed an average of 4 days a week in the plant for a period of one year. Total time spent in the field was roughly

2000 hours. Before the period spent in the specific setting, I did a set of 45 interviews in 3 companies. These interviews were aimed at 1) building relationships with the companies, 2) gathering information to select the most appropriate site and 3) having a view from managers on what would be the capabilities they wanted to develop in the new plants. I then selected one company, did some further interviews in their headquarters and started my ethnography in the new plant. During the first weeks I interviewed fifteen managers from different areas. Soon I started to follow some meetings across the organization, especially the daily production meetings in two areas: body shop and assembly. I was also doing observations and informal interviews in productive areas. After two months I started to follow one assembly line ever since the day they produced their first car. From then on, I spent around 60% of my time doing observations and the remaining 40% following managerial meetings and doing informal and formal interviews. My notes add to 4000 small notepad pages. They comprise facts (e.g. who was doing what at what time), interpretations (spontaneous or induced reactions people had at specific events), meeting transcriptions, informal and formal interviews, and my own reactions to the events I was witnessing. I also collected company documents and performance data (production and quality).

Data interpretation followed a coding process in many ways similar to the process of grounded theory building (Glaser and Strauss, 1967), but with some important conceptual differences that are specific to ethnography. The two methods do share some similarities in data interpretation, as both use comparable processes of generating understanding with iterative comparisons of data and theory. But the comparisons made at data analysis in grounded theory are focused on concepts rather than particular actors or contexts, resulting in a theory abstracted from the specificities of particular situations (Stewart, 1998). For the present paper, my main aim when coding the data was to generate categories that could represent well the organizational processes happening at specific stages during the learning curve evolution.

The learning curve

Figure one brings the learning curve for the assembly area in terms of cumulative production for the first year of operations. For confidentiality reasons, all numbers are omitted. The axis represents direct labor costs. To calculate that number, I took into account the added costs of working overtime (overtime was used quite extensively in the first months). While this does not represent the true evolution of number of hours taken to complete a vehicle (since some hours are calculated at higher costs, so the curve underestimates learning in terms of hours per vehicle), I chose to present the data in this way because it represents more accurately what was going on at the organization as a whole. The extensive use of overtime had its reasons (which I will explain below), and analyzing these reasons is important to understand organizational learning. If my focus were specifically on shop floor learning, number of hours would be a more accurate number. But as I want to describe what was happening in the organization as a whole, especially in the assembly area and the units that had to interact with it daily, labor costs are a better measure.

I divided the learning curve into three periods (see Figure 1) that will be described below. In the first one, costs are high and the process is very unstable. In period two, costs go down quite rapidly, but the process is still subject to big variations that cause a series of spikes in the learning curve. In the third period, the cost decrease rate goes down but the process is more stable. My analysis here will be purely qualitative, describing what was going on in the organization

during these periods based on my observations, informal and formal interviews and participation in production meetings. The aim is to identify drivers of the learning curve behavior at each stage, presenting a comprehensive and detailed view of the activities in the plant during its first year.

Insert Figure 1 about here

Period 1: Slack and ambiguity

The first period in the learning curve evolution (see figure 1) was characterized by a lot of slack in the shop floor. Workers were performing their activities slowly and had a good deal of free time between cars. This was not because they were not skilled enough to do it faster: as one assembler told me during that period, “don’t assume it is going to be like that always - in our training we were working at a much faster pace. I don’t know why they are going so slow now... I guess it is the planning, isn’t it? But why?” In fact, she was right. A lot of what was going in this period was the result of what had been planned. The organization, based on an assessment of its own resources and suppliers’ capabilities, planned a quite steep ramp-up – after some weeks of production, daily production would increase more than tenfold (with only twice the number of workers). Such a situation necessarily resulted in slack at shop floor during the initial weeks. But there were other problems in the organization that were disturbing the evolution of production. Even with all the slack at shop floor, objectives were not being met in the initial weeks due to problems at the production management level.

The first weeks of production (after almost three months of tryouts) were quite messy at the organizational level. As the daily production meetings started, representatives from different areas had a hard time trying to understand what was going on and what role they should play. Up to that moment, areas had worked mostly in isolation, developing their own cognitive frameworks. When production started, the different knowledge bases spread at different areas had to be integrated, but the process was not an easy one. The main initial difficulty was to select which pieces of information were needed to perform activities and solve problems, and determine where to find these pieces of information. There were no set channels of communication among areas, so relevant information usually arrived too late. Even when areas had the information they needed, procedures were still unclear and many documents were unknown.

Areas were still working under their own logic and felt they did not ‘belong’ to the assembly area. During one meeting, for example, a process engineer was referring to assembly in the 2nd person: “when you start using the new tool...” The assembly manager corrected him immediately: “when WE start using the new tool, not YOU... you’re also one of us!” This behavior was rather common; members that would be working at the assembly full time still kept a distance towards its activities. Their reference point was their own functional area.

One important issue at this stage was that these different areas were still going through their own learning curves. Although they had accumulated a good deal of declarative knowledge, the procedural knowledge concerning how to interact with other areas was still at an early stage of development. And, as they developed this knowledge linked to action, they discovered they needed more declarative knowledge as well. This process was similar to what has been conceptualized by Cook and Brown (1999) as “the generative dance between organizational

knowledge and organizational knowing”. Knowing is “to interact with and honor the world using knowledge as a tool” (Cook and Brown, 1999: 389), or the epistemological dimension of action itself. This component was missing at early interactions. Areas started with some stock of knowledge; the start of production triggered the development of knowing, which called for more knowledge to solve new problems arisen with the start of interactions. This was happening with all areas in parallel; they were all going through something of an ‘organizational learning curve’, expanding the boundaries of their local knowledge to meet other bodies of knowledge being developed elsewhere in the organization.

Since most of workers were new to the organization (and to the auto industry), they did not know their exact roles, so there was a lot of confusion on who should do what. There was a tendency to generate more and more reports to reduce uncertainty, but this information was still not properly used and shared. Under such circumstances, solving problems was a difficult task. But, given that the planned daily production was still low, many problems were still hard to notice. Daily production meetings served the purpose of information sharing, where people tried to learn more about the functioning of production and the role of each area.

These problems at the production management level had an impact on shop floor performance evolution. There were many days where production was halted because parts were missing (usually due to problems in information sharing between production and parts supply), or cars did not come from the paint shop (which was also facing many problems and going through its own learning curve). Some tools were still missing, so workers had to use alternative ways to do some of their tasks. That way, there were two major reasons behind the slack at the shop floor: the planning for the first few weeks and the high uncertainty and ambiguity at the production management level. With that, shop floor learning was quite irrelevant at this stage. They learned more when they were doing their training than on these initial periods. But the situation was going to change soon, with the planned increases in production.

Period 2: Integration through negotiation

As the planned daily production increased, the organization started to face new and difficult challenges. The performance of individual areas was still not at the expected level for regular production, and interaction patterns were still being formed. Problems were more exposed, and could not be avoided through slack anymore. Areas now had to work together and achieve ambitious common goals. This shift brought many uncertainties as organizational members had to face new problems that they were not prepared to solve. During production meetings, there were long discussions regarding sources of problems, many times with no clear conclusions. A perceived quality problem, for example, could have different sources, and it usually took a long time to discover the right cause. Feedback was still scarce, so many decisions had to be taken without enough information.

These discussions led to more fights between areas, which tried to avoid taking responsibility for problems. But there was a positive side: these conflicts triggered a response mechanism that became the main driver of integration among areas. With the pressing need for improvements, members started to negotiate many issues that were ambiguous or conflicting. These negotiations happened usually during the daily production meetings, or right after it if the topic was too specific to be treated in the meeting. Sometimes more conflicting, sometimes reflecting shared beliefs, negotiation served many purposes. It became a way to reduce

ambiguity, distribute responsibilities, structure roles and develop rules of interaction. It combined the various knowledge bases distributed in different areas. Most of all, it became the basis for routinization of production management, its outcomes getting embedded in relationships among areas and locking the organization into specific development paths.

This second phase was therefore marked by a shift from area to organizational learning. Areas had now developed more of their knowing, as well as began to improve their knowledge based on new necessities brought by knowing. The major trend at this phase was the integration of the learning curves developed by each area through the negotiation processes cited above. Sometimes the knowledge of the areas involved in a specific problem was still not adequate to solve it, and the usual result was a failed negotiation. In other words, their learning curves still had not reached levels where they could be integrated. These episodes of failed negotiations explain many of the spikes that can be observed in the learning curve during this period (see Figure 1). For example, one quality problem that happened during the first few weeks of production took a long time to be solved. As areas did not know the exact roots of the problem, nobody took responsibility to solve it. During meetings there were many attempts at negotiating this topic; all of them ended without a clear plan of action. Different possible causes of the problem were discussed, but meeting participants had no sufficient knowledge of the problem to choose any of the hypotheses, so action plans kept on being postponed. Negotiations failed because, under high ambiguity regarding sources of the problem, no agreement could satisfy all areas involved. The problem became so big that, when evidence for one of the possible causes was conclusive, top management decided to stop production until all cars produced up to that moment were checked and, if needed, reworked. All areas had to stop their work and help on that task until it was finished, regardless of who had the 'blame' for the problem.

The shop floor was now going through its crucial test. The company had made a risky decision, setting the plant in a region where potential employees had no experience in a factory setting. Their training had been carefully designed, but doubts persisted among top managers whether shop floor workers would be able to meet the objectives. During this phase, their performance was quite satisfactory. Most problems at this phase were outside the control of shop floor workers. Once again, missing parts and problems in the paint shop were the main reasons behind delays in the line. The only point where they were not performing very well regarded quality. With the big pressures to meet production objectives, coupled with frequent shortages of parts or cars to assemble, they had to work at a very fast pace when the line was running. With that, they gave more priority to quantity than quality – also because they were rewarded mostly on the basis of quantity. The result was an increasing number of quality problems that had to be corrected through re-work. In order to do rework (and also to meet production objectives in weeks where the line had stopped too much), the organization started to rely more and more on overtime.

After a first phase where individual learning at each area was the main mechanism for improvement, this second phase was marked by integration – of different learning curves and their underlying knowledge bases. With this integration, the overall assembly learning curve could progress in a much higher rate than before. Looking at the learning curve, this is the period where the most significant reductions in cost were achieved. These reductions were not due to increased skills at the shop floor, but increased organizational knowledge through integration of separate learning curves, which up to that moment were being developed at each area separately.

Period 3: Routinization

As areas developed and negotiated rules of interaction, the daily management of production became more routinized. Meetings were more objective, with less discussion and more focused problem-solving. During the previous phase, discussion of one specific topic during a meeting could last up to thirty minutes, especially when there was a lot of negotiation involved. This number went down considerably: after some months of production, each topic was discussed for less than five minutes. Meetings had now set scripts and their main aim was to share information. At this stage, members had already developed a transactive memory. Transactive memory is defined as a shared system for encoding, storing and retrieving information, and it begins when individuals learn something about one another's domain of expertise (Wegner, 1986). Developed during the previous phase, transactive memory now allowed a faster access to relevant information.

At the shop floor, they were also able to develop more stable working patterns. The constant interruptions and changes that characterized the initial months were almost over. Some work posts were still being redesigned, so there were changes in the sequence of activities, but by now most workers had the skills to work on different posts along the line. Workers showed some improvement in speed of activities, but most improvements came in quality. The percentage of cars that left the line without any problem more than tripled in a period of three weeks, reducing considerably the need for rework and overtime. The interesting feature here is that this rate of faultless cars leaving the line was low and without significant improvements for more than six months, then suddenly jumped up in a relatively short period. This suggests that only when the integrative efforts at management reached a level that allowed production to stabilize could workers develop quality-related skills. They needed practice and stability to achieve the desired quality results.

At this phase learning happened mostly through incremental processes, both induced and spontaneous. Experience brought an increased sensibility to identify and solve problems at the managerial level, and the number of improvement suggestions coming from operational levels increased – an indicator that experience was also bringing more awareness to explore opportunities at the shop floor. There were still sources of instability, especially when the organization had to face novel problems. But the effect of these problems was less disruptive because the organization had already set structures for problem solving. As it happened in the first phase, learning was happening again mostly in individual areas, but now the organization as a whole could benefit from this learning because information and knowledge sharing processes were embedded in the daily practices of the plant.

What is behind the learning curve?

As I started my field study, focused on the evolution of routines and capabilities in a new setting, I had some expectations regarding what was going to happen. I expected an incremental process where workers would gradually get better at what they did, and these improvements would result in the famous learning curve. While this is adequate as a general description of the process, it misses a lot of the intricacies that make such learning possible. I will concentrate my discussion on the most important drivers of learning in the setting that I analysed, conceptualizing on the organizational processes behind the learning curve. The intention is to

generate insights that can explain the phenomenon in more depth and help future studies in the topic by pointing important directions for study.

One first remark regards the role of planning in driving learning curve patterns. Learning curve is such a well-known phenomenon by now that it became an input to the planning process that precedes the launch of a new setting. That way, when analyzing only performance numbers, it is difficult to separate what truly represents learning from what simply represents the way the organization planned the early evolution of the process. As my study suggested, planning numbers were an important driver of the early behavior of the learning curve. Even with the ambiguity problems at the production management level, there was considerable slack. Production could be higher if planning numbers were more ambitious, which would result in a different learning curve. Planning had also a role in triggering learning at the second stage. As expected daily production increased steeply, areas had to work under a sense of urgency that accelerated the integration of their knowledge bases and the increasing stabilization of procedures. That way, learning curve may be as much a result as a cause (through planning) of the evolution path at the initial stages of a setting.

One of the most interesting findings of the study regards the role played by the different learning curves happening at different places in the organization. Understanding how learning evolves at each area and how this localized learning is shared within the organization is a key step to understand the major drivers behind the learning curve.

The learning process began at each area separately, in the pre-production phase before daily interactions were necessary. Individuals and small groups learned through cognitive and experiential processes, and developed a specific knowledge base (containing both explicit and tacit knowledge) within their areas. For example, workers at production planning learned to use the company's specific software, the rules that should be followed when programming production, the structure of their area and how to interact with their peers. People working at parts supply developed relationships with suppliers, learned about transportation options, customs rules, internal and external logistics. Some of this learning came through highly codified knowledge (company documents), some through interaction with peers (especially expatriates, which brought with them a lot of the tacit knowledge embedded in the company's practices). But this localized knowledge, though very important to the activities of the company, was not enough to bring significant performance evolution.

As interactions started, there was the need for the development of 'knowing', or an epistemology of action. This knowing was developed through experience, as members learned more about each other and formed a transactive memory. Through the development of knowing, organizational members also noticed they needed more knowledge that was required by other areas. With that, the 'generative dance' between knowledge and knowing started, as suggested by Cook and Brown (1999). In order to be able to combine their knowledge with other areas' knowledge and apply this combination to identify and solve problems in the organization, knowledge and knowing needed to reach a certain threshold level at each area. Before that, integration and application of knowledge was very difficult. After that threshold was achieved, negotiation became the most important driver for integration. Figure 2 brings a summary of how knowledge and knowing evolved within areas.

Insert Figure 2 about here

The following phase resulted in something similar to the ‘negotiated order’ concept developed by Strauss (1978). As Strauss defined it: “The negotiated order on any given day could be conceived of as the sum total of the organization’s rules and policies, along with whatever agreements, understanding, pacts, contracts, and other working arrangements currently obtained. These include agreements at every level of the organization, of every clique and coalition, and include covert as well as overt agreements.” (1978: 5-6). These agreements set the links among the different knowledge bases, resulting in a kind of synchronization of the individual area learning curves that enabled the evolution of the overall learning curve. From what I could observe in the field, the only difference from the concept proposed by Strauss is that, once a negotiated order is achieved, it gets embedded in routines and it is no longer open to further negotiations or agreements – at least for a considerable period of time (probably until routines change due to external or internal dynamics).

During the negotiation period, big increases in rates of learning were observed. This is because organizational members were learning together two critical processes, problem identifying and problem solving. But the process was not stable, since solving some problems led to identification of novel issues to be dealt with. This triggered more individual and shared learning, further negotiations and new procedures.

The negotiation process proceeded in loops until most of the major issues were resolved. Once a comprehensive negotiated order was achieved, organizational members started to act in a more routinized way. Meetings followed everyday the same structure, and much of the previous conflicts were absent. People knew already what to expect from their colleagues and from their areas. Routinization brought one important dimension for the further evolution of learning: stabilization. This stabilization facilitated learning in the shop floor, especially regarding quality. At this phase, learning was once again happening mostly within areas. But, as mechanisms for sharing were already in place, localized learning could bring improvements at the organizational level.

Figure 3 brings a summary of the organizational processes that supported the evolution of the learning curve, showing the main characteristics of learning at each phase and the related organizational mechanisms, as well as characteristics of the knowledge development process and performance evolution. Transitions between periods were not as marked as my description suggests. I described the strongest characteristics of each period to show their distinctive features, but the evolution of performance in the plant was an incremental process where characteristics of different periods could be found at the same time during some moments of the plant’s life. The transition from the first to the second period was triggered by the increase in planned daily production; from the second to the third period, by the accomplishment of a negotiated order that could be used to solve most problems.

Insert Figure 3 about here

Implications and Limitations

The article illustrated the organizational processes behind the evolution of the learning curve in the first year of a new manufacturing plant. As such, it is one step toward a shift on focus of learning curve studies: from the outcomes to the organizational processes behind it.

Understanding these processes is crucial to explain drivers of performance and to inform research on topics such as routines, knowledge development and integration, organizational learning, path-dependencies.

There are some potential contributions to management practice as well. The case suggests that managers can try to accelerate the learning curve by developing knowing before actual production starts through an early integration among areas. They should also concentrate resources on reaching knowledge threshold levels within their areas to allow for a faster integration across areas. Similar to Adler's (1990) results, shared learning seems to be crucial for performance development – in the case described here, it was shared learning across different areas at the launch phase of the productive process.

One limitation of the study is its time frame. One year may be a too short period to observe the evolution of a learning curve. But the strong competitive pressures to start producing as fast as possible in the country and the fact that the models were already produced by the company elsewhere helped to accelerate that curve. And, as it can be seen in figure 1, the process was quite stable in the last part of the learning curve. I concentrated my efforts in observing the most critical phase, where both the rate of learning and instability were higher.

There are also the usual limitations of ethnographic work. The objectives of ethnography are not the same as the ones in quantitative research. One cannot, for example, talk about 'replication' in ethnographic studies. Or claim for generalizability, at least in the typical usage of the word. But there are some evaluation criteria that can substitute for those coming from statistics-oriented research. Stewart (1998) proposes that, instead of validity, reliability and generalizability, ethnographies should aim at veracity, objectivity and perspicacity. Veracity means that the researcher has really observed what his descriptions claim. Objectivity is achieved when the study transcends the perspectives of the researcher and of the informants. Perspicacity implies that the research generates insights that are applicable to other times and other places. Although the theory generated from ethnographic data is obviously linked to the specific social setting, ethnographers can specify the contingencies for which the findings apply so they can be compared with other settings and other theories. This paper made an effort to achieve these three characteristics, generating data-grounded insights that contribute to understand an important topic in organization studies.

REFERENCES

- Adler, P. (1990), "Shared Learning", *Management Science*, 36, 938-957.
- Adler, P. S. and K. B. Clark (1991), "Behind the Learning Curve: A Sketch of the Learning Process", *Management Science*, 37(3), 267-281.
- Argote, L. 1999. *Organizational Learning. Creating, Retaining and Transferring Knowledge*. Boston: Kluwer Academic Publishers.
- Argote, L., S. L. Beckman and D. Epple (1990), "The Persistence and Transfer of Learning in Industrial Settings", *Management Science*, 36(2), 140-154.
- Argote, L. and E. Darr. 2000. Repositories of Knowledge in Franchise Organizations. In Dosi, G., Nelson, R. and Winter, S. (Eds.) *The Nature and Dynamics of Organizational Capabilities*, New York: Oxford University Press.
- Cook, S. and J. S. Brown (1999), "Bridging Epistemologies: The Generative Dance Between Organization Knowledge and Organization Knowing", *Organization Science*, 10(4), 381-400.
- Epple, D., L. Argote and K. Murphy (1996), "An Empirical Investigation of the Microstructure of Knowledge Acquisition and Transfer Through Learning by Doing", *Operations Research*, 44(1), 77-86.
- Fruin, W. M., (1998) "Site-Specific Organization Learning in International Technology Transfer" in Liker J. K., M. Fruin and P. Adler (eds.), *Remade in America: Transplanting and Transforming Japanese Management Systems*, Oxford University Press.
- Garud, R. (1997), "On the Distinction between Know-How, Know-Why, and Know-What", *Advances in Strategic Management*, 14, 81-101.
- Glaser, B. and A. Strauss (1967), *The Discovery of grounded theory*, Wiedenfeld and Nicholson, London.
- Lapre, M., A. Mukherjee and L. Wassenhove. 2000. Behind the Learning Curve: Linking Learning Activities to Waste Reduction. *Management Science*. 46: 597-611.
- Levin, D. 2000. Organizational Learning and the Transfer of Knowledge: an Investigation of Quality Improvement. *Organization Science*. 11: 630-647.
- Pisano, G., R. Bohmer and A. Edmondson. 2001. Organizational Differences in Rates of Learning: Evidence from the Adoption of Minimally Invasive Cardiac Surgery. *Management Science*. 47: 752-768.
- Stewart, A. (1998), *The Ethnographer's Method*, Sage Publications, London.
- Strauss, A. 1978. *Negotiations*. San Francisco: Jossey Bass.

- Szulanski, G. (1996), "Exploring Internal Stickiness: Impediments to the transfer of best practice within the firm", *Strategic Management Journal*, 17, 27-43.
- Wegner, D. 1986. Transactive memory: A contemporary analysis of the group mind. In B. Mullen and G. Goethals (Eds.), *Theories of Group Behavior*: 185-205. New York: Springer-Verlag

FIGURE 1

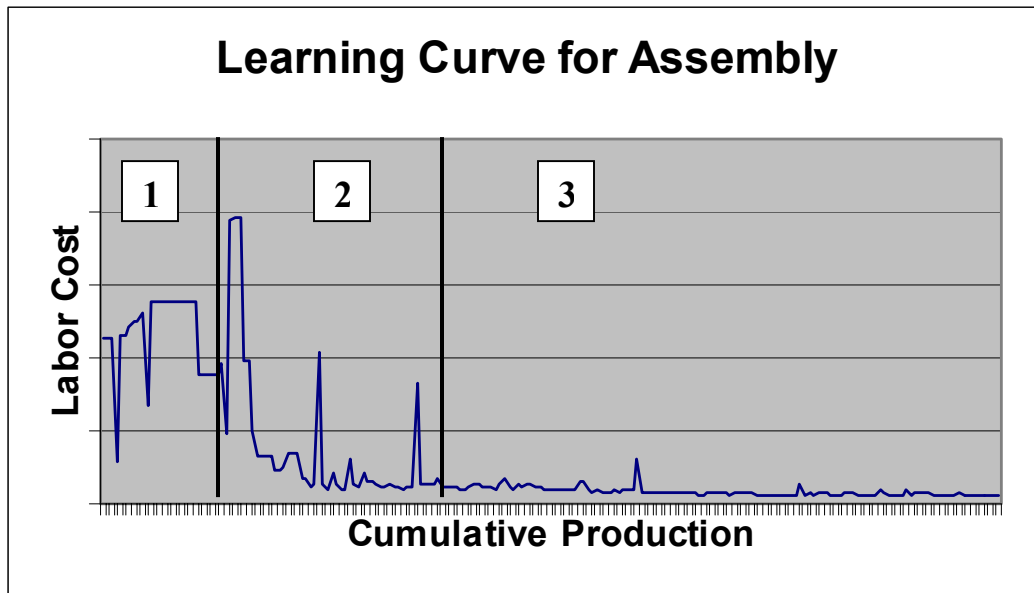


FIGURE 2

Evolution of knowledge within areas

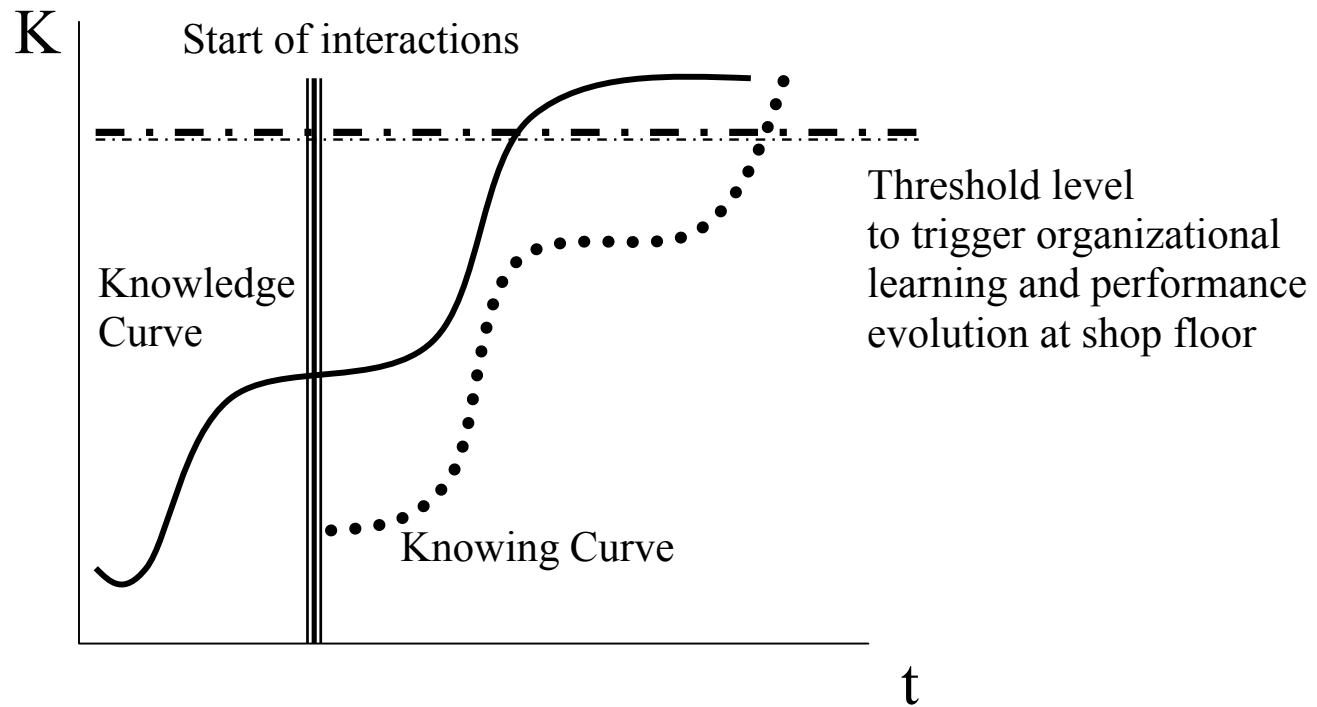


FIGURE 3

