

ADVERTISEMENT

Microsoft

people  ready

Online Version

Harvard Business Review

Steven Spear is an assistant professor of business administration at Harvard Business School in Boston. **H. Kent Bowen** is the Bruce Rauner Professor of Business Administration, also at Harvard Business School. Professor Bowen is the coauthor of "Regaining the Lead in Manufacturing" (HBR September–October 1994).

FEATURE

Decoding the DNA of the Toyota Production System

The Toyota story has been intensively researched and painstakingly documented, yet what really happens inside the company remains a mystery. Here's new insight into the unspoken rules that give Toyota its competitive edge.

by *Steven Spear and H. Kent Bowen*

The Toyota Production System has long been hailed as the source of Toyota's outstanding performance as a manufacturer. The system's distinctive practices—its kanban cards and quality circles, for instance—have been widely introduced elsewhere. Indeed, following their own internal efforts to benchmark the world's best manufacturing companies, GM, Ford, and Chrysler have independently created major initiatives to develop Toyota-like production systems. Companies that have tried to adopt the system can be found in fields as diverse as aerospace, consumer products, metals processing, and industrial products.

What's curious is that few manufacturers have managed to imitate Toyota successfully—even though the company has been extraordinarily open about its practices. Hundreds of thousands of executives from thousands of businesses have toured Toyota's plants in Japan and the United States. Frustrated by their inability to replicate Toyota's performance, many visitors assume that the secret of Toyota's success must lie in its cultural roots. But that's just not the case. Other Japanese companies, such as Nissan and Honda, have fallen short of Toyota's standards, and Toyota has successfully introduced its production system all around the world, including in North America, where the company is this year building over a million cars, mini-vans, and light trucks.

So why has it been so difficult to decode the Toyota Production System? The answer, we believe, is that observers confuse the tools and practices they see on their plant visits with the system itself. That makes it impossible for them to resolve an apparent paradox of the system—namely, that activities, connections, and production flows in a Toyota factory are rigidly scripted, yet at the same time Toyota's operations are enormously flexible and adaptable. Activities and processes are constantly being challenged and pushed to a higher level of performance, enabling the company to continually innovate and improve.

To understand Toyota's success, you have to unravel the paradox—you have to see that the rigid specification is the very thing that makes the flexibility and creativity possible. That's what we came to realize after an extensive, four-year study of the Toyota Production System in which we examined the inner workings of more than 40 plants in the United States, Europe, and Japan, some operating according to the system, some not. We studied both process and discrete manufacturers whose products ranged from prefabricated housing, auto parts and final auto assembly, cell phones, and computer printers to injection-molded plastics and aluminum extrusions. We studied not only routine production work but also service functions like equipment maintenance, workers' training and supervision, logistics and materials handling, and process design and redesign.

We found that, for outsiders, the key is to understand that the Toyota Production System creates a community of scientists. Whenever Toyota defines a specification, it is establishing sets of hypotheses that can then be tested. In other words, it is following the scientific method. To make any changes, Toyota uses a rigorous problem-solving process that requires a detailed assessment of the current state of affairs and a plan for improvement that is, in effect, an experimental test of the proposed changes. With anything less than such scientific rigor, change at Toyota would amount to little more than random trial and error—a blindfolded walk through life.

The fact that the scientific method is so ingrained at Toyota explains why the high degree of specification and structure at the company does not promote the command and control environment one might expect. Indeed, in watching people doing their jobs and in helping to design production processes, we learned that the system actually stimulates workers and managers to engage in the kind of experimentation that is widely recognized as the cornerstone of a learning organization. That is what distinguishes Toyota from all the other companies we studied.

The Toyota Production System and the scientific method that underpins it were not imposed on Toyota—they were not even chosen consciously. The system grew naturally out of the workings of the company over five decades. As a result, it has never been written down, and Toyota's workers often are not able to articulate it. That's why it's so hard for outsiders to grasp. In this article, we attempt to lay out how Toyota's system works. We try to make explicit what is implicit. We describe four principles—three rules of design, which show how Toyota sets up all its operations as experiments, and one rule of improvement, which describes how Toyota teaches the scientific method to workers at every level of the organization. It is these rules—and not the specific practices and tools that people observe during their plant visits—that in our opinion form the essence of Toyota's system. That is why we think of the rules as the DNA of the Toyota Production System. Let's take a closer look at those rules (for a summary, see the sidebar "The Four Rules").

The Experiments of the Toyota Production System

When organizations are managed according to the four rules, individuals are repeatedly conducting experiments, testing in operation the hypotheses built into the designs of individual work activities, customer-supplier connections, pathways, and improvement efforts. The hypotheses, the way they are tested, and the response if they are refuted are summarized below.

Rule	Hypotheses	Signs of a problem	Responses
1	The person or machine can do the activity as specified. If the activity is done as specified, the good or service will be defect free.	The activity is not done as specified. The outcome is defective.	Determine the true skill level of the person or the true capability of the machine and train or modify as appropriate. Modify the design activity.
2	Customers' requests will be for goods and services in a specific mix and volume. The supplier can respond to customers' requests.	Responses don't keep pace with requests. The supplier is idle, waiting for requests.	Determine the true mix and volume of demand and the true capability of the supplier; retrain, modify activities, or reassign customer-supplier pairs as appropriate.
3	Every supplier that is connected to the flow path is required. Any supplier not connected to the flow path is not needed.	A person or machine is not actually needed. A nonspecified supplier provides an intermediate good or service.	Determine why the supplier was unnecessary, and redesign the flow path. Learn why the nonspecified supplier was actually required, and redesign the flow path.
4	A specific change in an activity, connection, or flow path will improve cost, quality, lead time, batch size, or safety by a specific amount.	The actual result is different from the expected result.	Learn how the activity was actually performed or the connection or flow path was actually operated. Determine the true effects of the change. Redesign the change.

 The Four Rules (Located at the end of this article)

Rule 1: How People Work

Toyota's managers recognize that the devil is in the details; that's why they ensure that all work is highly specified as to content, sequence,

timing, and outcome. When a car's seat is installed, for instance, the bolts are always tightened in the same order, the time it takes to turn each bolt is specified, and so is the torque to which the bolt should be tightened. Such exactness is applied not only to the repetitive motions of production workers but also to the activities of all people regardless of their functional specialty or hierarchical role. The requirement that every activity be specified is the first unstated rule of the system. Put this baldly, the rule seems simple, something you'd expect everyone to understand and be able to follow easily. But in reality, most managers outside Toyota and its partners don't take this approach to work design and execution—even when they think they do.

Let's look at how operators at a typical U.S. auto plant install the front passenger seat into a car. They are supposed to take four bolts from a cardboard box, carry them and a torque wrench to the car, tighten the four bolts, and enter a code into a computer to indicate that the work has been done without problems. Then they wait for the next car to arrive. New operators are usually trained by experienced workers, who teach by demonstrating what to do. A seasoned colleague might be available to help a new operator with any difficulties, such as failing to tighten a bolt enough or forgetting to enter the computer code.

This sounds straightforward, so what's wrong with it? The problem is that those specifications actually allow—and even assume—considerable variation in the way employees do their work. Without anyone realizing it, there is plenty of scope for a new operator to put the seat into the vehicle differently than an experienced employee would. Some operators might put the front bolts in after the rear bolts; some might do it the other way around. Some operators might put each bolt in and then tighten them all; others might tighten as they go along. All this variation translates into poorer quality, lower productivity, and higher costs. More important, it hinders learning and improvement in the organization because the variations hide the link between how the work is done and the results.

At Toyota's plants, because operators (new and old, junior and supervisory) follow a well-defined sequence of steps for a particular job, it is instantly clear when they deviate from the specifications. Consider how workers at Toyota's Georgetown, Kentucky, plant install the right-front seat into a Camry. The work is designed as a sequence of seven tasks, all of which are expected to be completed in 55 seconds as the car moves at a fixed speed through a worker's zone. If the production worker finds himself doing task 6 (installing the rear seat-bolts) before task 4 (installing the front seat-bolts), then the job is actually being done differently than it was designed to be done, indicating that something must be wrong. Similarly, if after 40 seconds the worker is still on task 4, which should have been completed after 31 seconds, then something, too, is amiss. To make problem detection even simpler, the length of the floor for each work area is marked in tenths. So if the worker is passing the sixth of the ten floor marks (that is, if he is 33 seconds into the cycle) and is still on task 4, then he and his team leader know that he has fallen behind. Since the deviation is immediately apparent, worker and supervisor can move to correct the problem right away and then determine how to change the specifications or retrain the worker to prevent a recurrence. (See the sidebar "How Toyota's Workers Learn the Rules" for a short description of the process by which workers learn how to design work in this way.)

 How Toyota's Workers Learn the Rules (Located at the end of this article)

Even complex and infrequent activities, such as training an inexperienced workforce at a new plant, launching a new model, changing over a production line, or shifting equipment from one part of a plant to another, are designed according to this rule. At one of Toyota's suppliers in Japan, for example, equipment from one area of the plant was moved to create a new production line in response to changes in demand for certain products. Moving the machinery was broken into 14 separate activities. Each activity was then further subdivided and designed as a series of tasks. A specific person was assigned to do each task in a specified sequence. As each of the machines was moved, the way the tasks were actually done was compared with what was expected according to the original design, and discrepancies were immediately signaled.

In calling for people to do their work as a highly specified sequence of steps, rule 1 forces them to test hypotheses through action. Performing the activity tests the two hypotheses implicit in its design: first, that the person doing the activity is capable of performing it correctly and, second, that performing the activity actually creates the expected outcome. Remember the seat installer? If he can't insert the seat in the specified way within the specified amount of time, then he is clearly refuting at least one of these two hypotheses, thereby indicating that the activity needs to be redesigned or the worker needs to be trained.

Rule 2: How People Connect

Where the first rule explains how people perform their individual work activities, the second rule explains how they connect with one another. We express this rule as follows: every connection must be standardized and direct, unambiguously specifying the people involved, the form and quantity of the goods and services to be provided, the way requests are made by each customer, and the expected time in which the requests will be met. The rule creates a supplier-customer relationship between each person and the individual who is responsible for providing that person with each specific good or service. As a result, there are no gray zones in deciding who provides what to whom and when. When a worker makes a request for parts, there is no confusion about the supplier, the number of units required, or the timing of the delivery. Similarly,

when a person needs assistance, there is no confusion over who will provide it, how the help will be triggered, and what services will be delivered.

The real question that concerns us here is whether people interact differently at Toyota than they do at other companies. Let's return to our seat installer. When he needs a new container of plastic bolt covers, he gives a request to a materials handler, who is the designated bolt-cover supplier. Commonly, such a request is made with a kanban, a laminated card that specifies the part's identification number, the quantity of parts in the container, and the locations of the part supplier and of the worker (the customer) who will install it. At Toyota, kanban cards and other devices like *andon* cords set up direct links between the suppliers and the customers. The connections are as smooth as the passing of the baton in the best Olympic relay teams because they are just as carefully thought out and executed. For example, the number of parts in a container and the number of containers in circulation for any given part are determined by the physical realities of the production system—the distances, the changeover times, and so on. Likewise, the number of workers per team is determined by the types of problems expected to occur, the level of assistance the team members need, and the skills and capabilities of the team's leader.

Other companies devote substantial resources to coordinating people, but their connections generally aren't so direct and unambiguous. In most plants, requests for materials or assistance often take a convoluted route from the line worker to the supplier via an intermediary. Any supervisor can answer any call for help because a specific person has not been assigned. The disadvantage of that approach, as Toyota recognizes, is that when something is everyone's problem it becomes no one's problem.

The requirement that people respond to supply requests within a specific time frame further reduces the possibility of variance. That is especially true in service requests. A worker encountering a problem is expected to ask for assistance at once. The designated assistant is then expected to respond immediately and resolve the problem within the worker's cycle time. If the worker is installing a front seat every 55 seconds, say, then a request for help must be answered and dealt with in less than the 55 seconds. If the problem cannot be resolved in less than 55 seconds, that failure immediately challenges the hypotheses in this customer-supplier connection for assistance. Perhaps the request signal is ambiguous. Perhaps the designated assistant has too many other requests for help and is busy or is not a capable problem solver. Constantly testing the hypotheses in this way keeps the system flexible, making it possible to adjust the system continually and constructively.

The striking thing about the requirement to ask for help at once is that it is often counterintuitive to managers who are accustomed to encouraging workers to try to resolve problems on their own before calling for help. But then problems remain hidden and are neither shared nor resolved companywide. The situation is made worse if workers begin to solve problems themselves and then arbitrarily decide when the problem is big enough to warrant a call for help. Problems mount up and only get solved much later, by which time valuable information about the real causes of the problem may have been lost.

Rule 3: How the Production Line Is Constructed

All production lines at Toyota have to be set up so that every product and service flows along a simple, specified path. That path should not change unless the production line is expressly redesigned. In principle, then, there are no forks or loops to convolute the flow in any of Toyota's supply chains. That's the third rule.

To get a concrete idea of what that means, let's return to our seat installer. If he needs more plastic bolt covers, he orders them from the specific material handler responsible for providing him with bolt covers. That designated supplier makes requests to his own designated supplier at the off-line store in the factory who, in turn, makes requests directly to his designated supplier at the bolt cover factory's shipping dock. In this way, the production line links each person who contributes to the production and delivery of the product, from the Toyota factory, through the molding company, to even the plastic pellet manufacturer.

The point is that when production lines are designed in accordance with rule 3, goods and services do not flow to the next available person or machine but to a *specific* person or machine. If for some reason that person or machine is not available, Toyota will see it as a problem that might require the line to be redesigned.

The stipulation that every product follow a simple, prespecified path doesn't mean that each path is dedicated to only one particular product, however. Quite the contrary: each production line at a Toyota plant typically accommodates many more types of products than its counterparts do at other companies.

The third rule doesn't apply only to products—it applies to services, like help requests, as well. If our seat installer, for example, needs help, that too comes from a single, specified supplier. And if that supplier can't provide the necessary assistance, she, in turn, has a designated helper. In

some of Toyota's plants, this pathway for assistance is three, four, or five links long, connecting the shop floor worker to the plant manager.

The third rule runs contrary to conventional wisdom about production lines and pooling resources—even contrary to how most people think the Toyota Production System works. According to received wisdom, as a product or service is passed down the line, it should go to the next machine or person available to process it further. Similarly, most people assume that help should come from the first available person rather than from a specific person. At one auto parts supplier we studied, for example, most of the parts could be stamped on more than one press machine and welded at more than one welding station. Before the company adopted the Toyota system, its practice was to pass each part on to the first available press machine and to the first available welder. When the plant switched over, under Toyota's guidance, each type of part followed only one production path through the plant.

By requiring that every pathway be specified, the rule ensures that an experiment will occur each time the path is used. Here the hypotheses embedded in a pathway designed according to rule 3 are that every supplier connected to the pathway is necessary, and any supplier not connected is not necessary. If workers at the auto parts supplier found themselves wanting to divert production to another machine or welding station, or if they began turning for help to someone other than their designated helpers, they'd conclude that their actual demand or capacity didn't match their expectations. And there would also be no ambiguity about which press or welder was involved. Again, the workers would revisit the design of their production line. Thus rule 3, like rules 1 and 2, enables Toyota to conduct experiments and remain flexible and responsive.

Rule 4: How to Improve

Identifying problems is just the first step. For people to consistently make effective changes, they must know how to change and who is responsible for making the changes. Toyota explicitly teaches people how to improve, not expecting them to learn strictly from personal experience. That's where the rule for improvement comes in. Specifically, rule 4 stipulates that any improvement to production activities, to connections between workers or machines, or to pathways must be made in accordance with the scientific method, under the guidance of a teacher, and at the lowest possible organizational level. Let's look first at how Toyota's people learn the scientific method.

How People Learn to Improve.

In 1986, Aisin Seiki, a Toyota Group company that made complex products such as power trains for the auto industry, created a line to manufacture mattresses to absorb excess capacity in one of its plants. Since 1986, its range has grown from 200 to 850 types of mattresses, its volume has grown from 160 mattresses per day to 550, and its productivity has doubled. Here's an example of how they did it.

On one of our visits to this plant, we studied a team of mattress assembly workers who were being taught to improve their problem-solving skills by redesigning their own work. Initially, the workers had been responsible for doing only their own standardized work; they had not been responsible for solving problems. Then the workers were assigned a leader who trained them to frame problems better and to formulate and test hypotheses—in other words, he taught them how to use the scientific method to design their team's work in accordance with the first three rules. The results were impressive. One of the team's accomplishments, for instance, was to redesign the way edging tape was attached to the mattresses, thereby reducing the defect rate by 90%. (See the exhibit "On-Demand Production at the Aisin Mattress Factory.")

 On-Demand Production at the Aisin Mattress Factory (Located at the end of this article)

To make changes, people are expected to present the explicit logic of the hypotheses. Let's look at what that can involve. Hajime Ohba, general manager of the Toyota Supplier Support Center, was visiting a factory in which one of TSSC's consultants was leading a training and improvement activity (for a description of the role of the Toyota Production System promotion centers, see the sidebar "Toyota's Commitment to Learning"). The consultant was helping factory employees and their supervisor reduce the manufacturing lead time of a particular line, and Ohba was there to evaluate the group's progress.

 Toyota's Commitment to Learning (Located at the end of this article)

Group members began their presentation by describing the steps by which their product was created—delineating all the problems they identified when they had first studied the process for changing over a machine from making one part to making another, and explaining the specific changes they had made in response to each of those problems. They concluded by saying, "When we started, the changeover required 15 minutes. We were hoping to reduce that by two-thirds—to achieve a five-minute changeover—so that we could reduce batch sizes by two-thirds. Because of the modifications we made, we achieved a changeover time of seven and a half minutes—a reduction of one-half."

After their presentation, Ohba asked why the group members had not achieved the five-minute goal they had originally established. They were a bit taken aback. After all, they had reduced the changeover time by 50%, yet Ohba's question suggested he had seen opportunities for even greater improvement that they had missed. They offered explanations having to do with machine complexity, technical difficulty, and equipment upgrade costs. Ohba responded to these replies with yet more questions, each one meant to push the consultant and the factory people to articulate and challenge their most basic assumptions about what could and could not be changed—assumptions that both guided and constrained the way they had solved their problems. Were they sure four bolts were necessary? Might the changeover be accomplished with two? Were they certain that all the steps they included in the changeover were needed? Might some be combined or eliminated? In asking why they had not achieved the five-minute goal, Ohba was not suggesting that the team had failed. Rather, he was trying to get them to realize that they had not fully explored all their improvement opportunities because they had not questioned their assumptions deeply enough.

There was a second reason for Ohba's persistence. He was trying to show the group members that their improvement activity had not been carried out as a bona fide experiment. They had established a goal of five minutes based on the premise that faster changeovers and smaller batches are better than slower changeovers and larger batches. But here they were confusing goals with predictions based on hypotheses. The goal was not a prediction of what they believed they would achieve through the specific improvement steps they planned to take. As a result, they had not designed the improvement effort as an experiment with an explicit, clearly articulated, verifiable hypothesis of the form, "If we make the following specific changes, we expect to achieve this specific outcome." Although they had reduced the changeover time considerably, they had not tested the hypotheses implicit in their effort. For Ohba, it was critical that the workers and their supervisor realize that how they made changes was as important as what changes they made.

Who Does the Improvement.

Frontline workers make the improvements to their own jobs, and their supervisors provide direction and assistance as teachers. If something is wrong with the way a worker connects with a particular supplier within the immediate assembly area, the two of them make improvements, with the assistance of their common supervisor. The Aisin team we described earlier, for example, consisted of the assembly line workers and the supervisor, who was also their instructor. When changes are made on a larger scale, Toyota ensures that improvement teams are created consisting of the people who are directly affected and the person responsible for supervising the pathways involved.

Thus the process remains the same even at the highest levels. At Aisin's mattress factory, we found that the plant manager took responsibility for leading the change from three production lines back to two (the number had risen to three to cope with an increase in product types). He was involved not just because it was a big change but also because he had operational responsibility for overseeing the way work flowed from the feeder lines to the final assembly lines. In this way, Toyota ensures that problem solving and learning take place at all levels of the company. Of course, as we have already seen, Toyota will bring in external experts as necessary to ensure the quality of the learning process.

In the long term, the organizational structures of companies that follow the Toyota Production System will shift to adapt to the nature and frequency of the problems they encounter. Since the organizational changes are usually being made at a very low level, however, they can be hard for outsiders to detect. That's because it is the nature of the problems that determines who should solve them and how the organization is designed. One consequence is that different organizational structures coexist quite happily even in the same plant.

Consider Toyota's engine-machining plant in Kamigo, Japan. The plant has two machine divisions, each of which has three independent production shops. When we visited in summer 1998, the production people in the first machine division answered to shop heads, and the process engineers answered directly to the head of the division. However, in the second machine division, the engineers were distributed among the three shops and, like the production workers, answered to the various shop heads. Neither organizational structure is inherently superior. Rather, the people we interviewed explained, problems in the first division happened to create a situation that required the engineers to learn from one another and to pool engineering resources. By contrast, the problems that arose in the second division required the production and engineering people to cooperate at the level of the individual shops. Thus the organizational differences reflect the fact that the two divisions encountered different problems.

Toyota's Notion of the Ideal

By inculcating the scientific method at all levels of the workforce, Toyota ensures that people will clearly state the expectations they will be testing when they implement the changes they have planned. But beyond this, we found that people in companies following the Toyota Production System share a common goal. They have a common sense of what the ideal production system would be, and that shared vision motivates them to make improvements beyond what would be necessary merely to meet the current needs of their customers. This notion of the ideal is very pervasive, and we believe it is essential to understanding the Toyota Production System.

When they speak of the ideal, workers at Toyota do not mean something philosophically abstract. They have a concrete definition in mind, one that is remarkably consistent throughout the company. Very specifically, for Toyota's workers, the output of an ideal person, group of people, or machine:

- is defect free (that is, it has the features and performance the customer expects);
- can be delivered one request at a time (a batch size of one);
- can be supplied on demand in the version requested;
- can be delivered immediately;
- can be produced without wasting any materials, labor, energy, or other resources (such as costs associated with inventory); and
- can be produced in a work environment that is safe physically, emotionally, and professionally for every employee.

We consistently found people at plants that used the Toyota Production System making changes that pushed operations toward this ideal. At one company that produced electromechanical products, for example, we found that workers had come up with a number of ingenious error-detecting gauges that generated a simple, unambiguous yes-or-no signal to indicate whether their output was free of defects—as specified in the ideal. At yet another plant, which manufactures injection-molded parts, we found that workers had reduced the time it took to change a large molding die from an already speedy five minutes to three minutes. This allowed the company to reduce the batch sizes of each part it produced by 40%, bringing it closer to the ideal batch size of one. As Toyota moves toward the ideal, it may temporarily hold one of its dimensions to be more important than another. Sometimes this can result in practices that go against the popular view of Toyota's operations. We have seen cases where Toyota keeps higher levels of inventory or produces in batch sizes larger than observers generally expect of a just-in-time operation, as we describe in the sidebar “Countermeasures in the Toyota Production System.”

 Countermeasures in the Toyota Production System (Located at the end of this article)

Toyota's ideal state shares many features of the popular notion of mass customization—the ability to create virtually infinite variations of a product as efficiently as possible and at the lowest possible cost. In the final analysis, Toyota's ideal plant would indeed be one where a Toyota customer could drive up to a shipping dock, ask for a customized product or service, and get it at once at the lowest possible price and with no defects. To the extent that a Toyota plant—or a Toyota worker's activity—falls short of this ideal, that shortcoming is a source of creative tension for further improvement efforts.

The Organizational Impact of the Rules

If the rules make companies using the Toyota Production System a community of scientists performing continual experiments, then why aren't these organizations in a state of chaos? Why can one person make a change without adversely affecting the work of other people on the production line? How can Toyota constantly introduce changes to its operations while keeping them running at full tilt? In other words, how does Toyota improve and remain stable at the same time?

Once again, the answer is in the rules. By making people capable of and responsible for doing and improving their own work, by standardizing connections between individual customers and suppliers, and by pushing the resolution of connection and flow problems to the lowest possible level, the rules create an organization with a nested modular structure, rather like traditional Russian dolls that come one inside the other. The great benefit of nested, modular organizations is that people can implement design changes in one part without unduly affecting other parts. That's why managers at Toyota can delegate so much responsibility without creating chaos. Other companies that follow the rules will also find it possible to change without experiencing undue disruption.

Of course, the structures of other companies have features in common with those that follow the Toyota Production System, but in our research we found no company that had them all that did not follow the system. It may turn out in the end that you can build the structure only by investing the time Toyota has. But we believe that if a company dedicates itself to mastering the rules, it has a better chance of replicating Toyota's DNA—and with that, its performance.

The Four Rules

The tacit knowledge that underlies the Toyota Production System can be captured in four basic rules. These rules guide the design, operation,

and improvement of every activity, connection, and pathway for every product and service. The rules are as follows:

Rule 1: All work shall be highly specified as to content, sequence, timing, and outcome.

Rule 2: Every customer-supplier connection must be direct, and there must be an unambiguous yes-or-no way to send requests and receive responses.

Rule 3: The pathway for every product and service must be simple and direct.

Rule 4: Any improvement must be made in accordance with the scientific method, under the guidance of a teacher, at the lowest possible level in the organization.

All the rules require that activities, connections, and flow paths have built-in tests to signal problems automatically. It is the continual response to problems that makes this seemingly rigid system so flexible and adaptable to changing circumstances.

How Toyota's Workers Learn the Rules

If the rules of the Toyota Production System aren't explicit, how are they transmitted? Toyota's managers don't tell workers and supervisors specifically how to do their work. Rather, they use a teaching and learning approach that allows their workers to discover the rules as a consequence of solving problems. For example, the supervisor teaching a person the principles of the first rule will come to the work site and, while the person is doing his or her job, ask a series of questions:

- How do you do this work?
- How do you know you are doing this work correctly?
- How do you know that the outcome is free of defects?
- What do you do if you have a problem?

This continuing process gives the person increasingly deeper insights into his or her own specific work. From many experiences of this sort, the person gradually learns to generalize how to design all activities according to the principles embodied in rule 1.

All the rules are taught in a similar Socratic fashion of iterative questioning and problem solving. Although this method is particularly effective for teaching, it leads to knowledge that is implicit. Consequently, the Toyota Production System has so far been transferred successfully only when managers have been able and willing to engage in a similar process of questioning to facilitate learning by doing.

On-Demand Production at the Aisin Mattress Factory

Aisin Seiki produces 850 varieties of mattresses, distinguished by size, firmness, covering fabric, quilting pattern, and edge trim. Customers can order any one of these in a retail store and have it delivered to their homes in three days, yet Aisin maintains an inventory at the plant equal to just 1.5 days of demand. To be able to do so, Aisin has made thousands of changes in individual work activities, in the connections linking customers and suppliers of intermediate goods and services, and to the overall production lines. This table captures how dramatic the results of those changes have been.

	1986	1988	1992	1996	1997
Styles	200	325	670	750	850
Units per day	160	230	360	530	550
Units per person	8	11	13	20	26
Productivity index	100	138	175	197	208
Finished-goods inventory (days)	30	2.5	1.8	1.5	1.5
Number of assembly lines	2	2	3	3	2

Toyota's Commitment to Learning

All the organizations we studied that are managed according to the Toyota Production System share an overarching belief that people are the most significant corporate asset and that investments in their knowledge and skills are necessary to build competitiveness. That's why at these organizations all managers are expected to be able to do the jobs of everyone they supervise and also to teach their workers how to solve problems according to the scientific method. The leadership model applies as much to the first-level "team leader" supervisors as it does to those at the top of the organization. In that way, everybody at Toyota shares in the development of human resources. In effect, there is a cascading pathway for teaching, which starts with the plant manager, that delivers training to each employee.

To reinforce the learning and improvement process, each plant and major business unit in the Toyota Group employs a number of Toyota Production System consultants whose primary responsibility is to help senior managers move their organizations toward the ideal. These "learner-leader-teachers" do so by identifying ever more subtle and difficult problems and by teaching people how to solve problems scientifically.

Many of these individuals have received intensive training at Toyota's Operations Management Consulting Division. OMCD was established in Japan as an outgrowth of efforts by Taiichi Ohno—one of the original architects of the Toyota Production System—to develop and diffuse the system throughout Toyota and its suppliers. Many of Toyota's top officers—including Toyota Motor's new president, Fujio Cho—have honed their skills within OMCD. During their OMCD tenure, which can extend for a period of years, Toyota's employees are relieved of all line responsibilities and instead are charged with leading improvement and training activities in the plants of Toyota and its suppliers. By supporting all of Toyota's plant and logistical operations in this way, OMCD serves as a training center, building its consultants' expertise by giving them opportunities to solve many difficult problems and teach others to do the same.

In 1992, Toyota founded the Toyota Supplier Support Center (TSSC) in the United States to provide North American companies with training in the Toyota Production System. Modeled on OMCD, TSSC has given workshops to more than 140 companies and direct assistance to 80. Although most of these companies are auto suppliers, few are exclusively Toyota suppliers; participants come from other industries and from universities, government organizations, and industry associations. Indeed, much of the research for this paper was derived from the experience of one of the authors, who was a member of a TSSC team for five months, promoting the Toyota Production System at a plant that supplies Toyota and two other auto assembly plants.

Countermeasures in the Toyota Production System

Toyota does not consider any of the tools or practices—such as kanbans or *andon* cords, which so many outsiders have observed and copied—as fundamental to the Toyota Production System. Toyota uses them merely as temporary responses to specific problems that will serve until a better approach is found or conditions change. They're referred to as "countermeasures," rather than "solutions," because that would imply a permanent resolution to a problem. Over the years, the company has developed a robust set of tools and practices that it uses as countermeasures, but many have changed or even been eliminated as improvements are made.

So whether a company does or does not use any particular tool or practice is no indication that it is truly applying Toyota's rules of design and improvement. In particular, contrary to the impression that the concept of zero inventory is at the heart of the Toyota system, we've observed many cases in which Toyota actually built up its inventory of materials as a countermeasure. The ideal system would in fact have no need for inventory. But, in practice, certain circumstances may require it:

- **Unpredictable downtime or yields.** Sometimes a person or a machine is unable to respond on demand when a request is made because of an unexpected mechanical breakdown. For this reason, safety stock is held to protect the customer against random occurrences. The person responsible for ensuring the reliability of a machine or process owns that inventory and strives to reduce the frequency and length of downtimes so that the amount of the safety stock can be reduced.
- **Time-consuming setups.** Difficulties in switching a machine from processing one kind of product to another can prevent a supplier from responding immediately. Therefore, suppliers will produce the product in batch sizes greater than one and hold the excess as inventory so it can respond immediately to the customer. Of course, suppliers will continually try to reduce the changeover time to keep batch sizes and stores of inventory as small as possible. Here, the owners of both the problem and the countermeasure are the machine operator and the team leader, who are responsible for reducing changeover times and batch sizes.
- **Volatility in the mix and volume of customer demand.** In some cases, variations in customers' needs are so large and unpredictable that it is impossible for a plant to adjust its production to them quickly enough. In those instances, buffer stock is kept at or near the shipping point as

a countermeasure. The buffer stock also serves as a signal to production and sales managers that the person who works most directly with the customer must help that customer eliminate the underlying causes of any preventable swings in demand.

In many cases, the same type of product is held in different types of inventory. Toyota does not pool its various kinds of inventory, even though doing so would reduce its inventory needs in the short term. That might sound paradoxical for a management system so popularly known to abhor waste. But the paradox can be resolved when we recognize that Toyota's managers and workers are trying to match each countermeasure to each problem.

There's no link between the reason for keeping safety stock—process unreliability—and the reason for keeping buffer stock—fluctuations in customer demand. To pool the two would make it hard to distinguish between the separate activities and customer-supplier connections involved. The inventory would have many owners, and the reasons for its use would become ambiguous. Pooling the inventory thus muddles both the ownership and cause of the problems, making it difficult to introduce improvements.