Chapter 10

Principle 4:
Level Out the Workload (*Heijunka*)

In general, when you try to apply the TPS, the first thing you have to do is to even out or level the production. And that is the responsibility primarily of production control or production management people. Leveling the production schedule may require some front-loading of shipments or postponing of shipments and you may have to ask some customers to wait for a short period of time. Once the production level is more or less the same or constant for a month, you will be able to apply pull systems and balance the assembly line. But if production levels—the output—varies from day to day, there is no sense in trying to apply those other systems, because you simply cannot establish standardized work under such circumstances.

—Fujio Cho, President, Toyota Motor Corporation

Following the lead of Dell Computer and other successful companies, many businesses in America are rushing to a build-to-order model of production. They want to make just what the customers want when they want it—the ultimate lean solution. Unfortunately, customers are not predictable and actual orders vary significantly from week to week and month to month. If you build product as it is ordered, you may be building huge quantities one week, paying overtime, and stressing your people and equipment, but then, if orders are light the next week, your people will have little to do and your equipment will be underutilized. You will also not know how much to order from your suppliers, so you will have to stockpile the maximum possible amount of each item the customers might possibly order. It is impossible to run a lean operation in this way. A strict build-to-order model creates piles of inventory, hidden problems, and ultimately poorer quality and in the end lead times are likely to grow as the factory is disorganized and chaotic. Toyota has found it can create the leanest operation and ultimately give customers better service and better quality by leveling out the production schedule and not always building to order.

Some of the businesses I work with that try to "build to order" are in actuality asking customers to wait six to eight weeks for their "build-to-order" product. A few "special" customers may cut in line and get their orders expedited at the expense of the large majority of customers. But why ruin the pace of your operation to build an order in hand today when the customer will not get the product for six weeks? Instead, accumulate orders and level the schedule and you may be able to reduce production lead times, cut your parts inventories, and quote much shorter standard lead times to all your customers, resulting in greater overall customer satisfaction than a "hurry up, then slow down" build-to-order approach to production.

Toyota managers and employees use the Japanese term *muda* when they talk about waste and eliminating *muda* is the focus of most lean manufacturing efforts. But two other M's are just as important to making lean work, and all three M's fit together as a system. In fact, focusing exclusively on only the eight wastes of *muda* can actually hurt the productivity of people and the production system. The Toyota Way document refers to the "elimination of *Muda*, *Muri*, *Mura*" (see Figure 10-1). The three M's are:

- *Muda—Non-value-added*. The most familiar M includes the eight wastes mentioned in earlier chapters. These are wasteful activities that lengthen lead times, cause extra movement to get parts or tools, create excess inventory, or result in any type of waiting.
- *Muri—Overburdening people or equipment*. This is in some respects on the opposite end of the spectrum from muda. Muri is pushing a machine or person beyond natural limits. Overburdening people results in safety and quality problems. Overburdening equipment causes breakdowns and defects.
- *Mura—Unevenness*. You can view this as the resolution of the other two M's. In normal production systems, at times there is more work than the people or machines can handle and at other times there is a lack of work. Unevenness results from an irregular production schedule or fluctuating production volumes due to internal problems, like downtime or missing parts or defects. *Muda* will be a result of *mura*. Unevenness in production levels means it will be necessary to have on hand the equipment, materials, and people for the highest level of production—even if the average requirements are much lower than that.

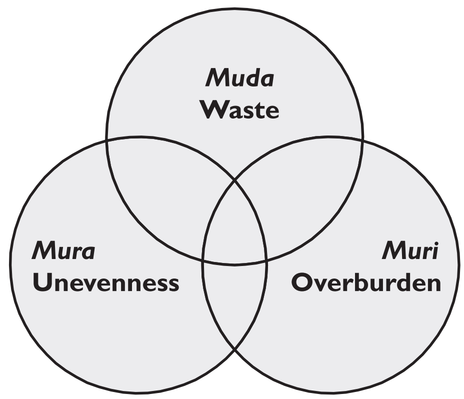

Figure 10-1. The three M's

Let's say you have a production schedule that swings wildly and a production process that is not well balanced or reliable. You've decided to start applying lean thinking and focus only on "eliminating *muda*" from your production system. You start to reduce inventory in your system. Then you look at the work balance and reduce the number of people from the system.[1] Then you organize the workplace better to eliminate wasted motion. Finally, you step back and let the system run. What you'll sadly witness is a system that will run itself into the ground due to spikes in customer demand that force people and equipment to work harder than they efficiently can! When work first begins to flow one piece at a time across work centers, without inventory, the pace and product mix of production jerk all over the place. The only thing you get is erratic one-piece flow. Workers will be overburdened. Equipment will break down even more than before. You will run out of parts. Then you'll conclude, "Lean manufacturing doesn't work here."

Interestingly, focusing on *muda* is the most common approach to "implementing lean tools," because it is easy to identify and eliminate waste. But what many companies fail to do is the more difficult process of stabilizing the system and creating "evenness"—a true balanced lean flow of work. This is the Toyota concept of *heijunka,* leveling out the work schedule. It is perhaps the Toyota Way's most counterintuitive principle. Achieving *heijunka* is fundamental to eliminating *mura*, which is fundamental to eliminating *muri* and *muda*.

Having starts and stops, overutilization then underutilization, is a problem because it does not lend itself to quality, standardization of work, productivity, or continuous improvement. As explained by Taiichi Ohno:

> *The slower but consistent tortoise causes less waste and is much more desirable than the speedy hare that races ahead and then stops occasionally to doze. The Toyota Production System can be realized only when all the workers become tortoises. (Ohno, 1988)*

I have heard this repeated from other Toyota leaders: "We would rather be slow and steady like the tortoise than fast and jerky like the rabbit." U.S. production systems force workers to be like rabbits. They tend to work really hard, wear themselves down, and then take a siesta. In many U.S. factories, workers will sometimes double up on the assembly line, one doing two jobs while the other has free time, and so long as the workers make production quotas for the day, management looks the other way.

Heijunka—Leveling Production and Schedules

Heijunka is the leveling of production by both volume and product mix. It does not build products according to the actual flow of customer orders, which can swing up and down wildly, but takes the total volume of orders in a period and levels them out so the same amount and mix are being made each day. The approach of TPS from the beginning was to keep batch sizes small and build what the customer (external or internal) wants. In a true one-piece flow, you can build Products A and B in the actual production sequence of customer orders (e.g., A, A, B, A, B, B, B, A, B…). The problem with building to an actual production sequence is that it causes you to build parts irregularly. So if orders on Monday are twice those on Tuesday, you must pay your employees overtime on Monday and then send them home early on Tuesday. To smooth this out, you take the actual customer demand, determine the pattern of volume and mix, and build a level schedule every day. For example, you know you are making five A's for every five B's. Now you can create a level production sequence of ABABAB. This is called leveled, mixed-model production, because you are mixing up production but also leveling the customer demand to a predictable sequence, which spreads out the different product types and levels volume.

Figure 10-2 gives an example of an unleveled schedule from an engine plant that makes small engines for lawn care equipment (based on an actual case).

In this case, a production line makes three sizes of engines—small, medium, and large. The medium engines are the big sellers, so these are made early in the week—Monday through part of Wednesday. Then there is a several-hour changeover of the line to make small engines that are made the rest of Wednesday through Friday morning. Finally, the large engines—in smallest demand—are made Friday afternoon. There are four things wrong with this unleveled schedule:

1. *Customers usually do not buy products predictably.* The customer is buying medium and large engines throughout the week. So if the customer unexpectedly decides to buy an unusually large number of large engines early in the week, the plant is in trouble. You can get around this by holding a lot of finished goods inventory of all engines, but this leads to a high cost of inventory, with all its related costs.

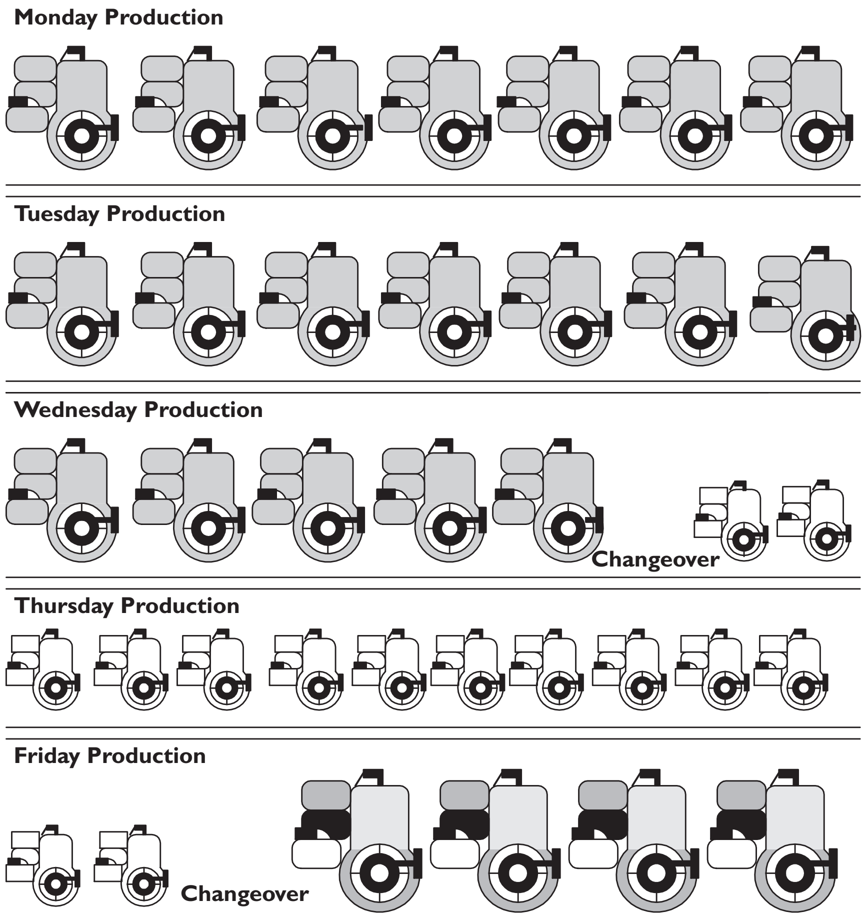

Figure 10-2. Traditional production (unleveled)

2. *There is a risk of unsold goods.* If the plant does not sell all its medium engines built up Monday to Wednesday, it must keep them in inventory.
3. *The use of resources is unbalanced.* Most likely, there are different labor requirements for these different-sized engines, with the largest engines taking the most labor time. So the plant needs a medium amount of labor early in the week, then less labor in the middle of the week, and then a lot of labor at the end of the week. There is potentially a lot of *muda* and *muri.*
4. *Placing an uneven demand on upstream processes.* This is perhaps the most serious problem. Since the plant is purchasing different parts for the three types of engines, it will be asking its suppliers to send certain parts Monday through Wednesday and different parts for the rest of the week. Experience

tells us that customer demand always changes and the engine plant will be unable to stick to the schedule anyway. Most likely there will be some big shifts in the model mix, e.g., unexpected rush order of large engines and the need to focus on making those for a whole week. The supplier will need to be prepared for the worst possible scenario and will need to keep at least one week's worth of all parts for all three engine types. And something called the "bullwhip effect" will multiply this behavior backward through the supply chain. Think of the small force in your wrist creating a huge and destructive force at the end of the whip. Similarly, a small change in the schedule of the engine assembly plant will result in ever-increasing inventory banks at each stage of the supply chain as you move backward from the end customer.

In a batch-processing mode, the goal is to achieve economies of scale for each individual piece of equipment. Changing over tools to alternate between making product A and product B seems wasteful because parts are not being produced during the changeover time. You are also paying the equipment operator while the machine is being changed over. So the logical solution is to build large batches of product A before changing over to product B. But this approach does not allow for *heijunka.*

In the case of the motors, the plant did a careful analysis and discovered the long time to changeover the line was due to moving in and out parts and tools for the larger engine and moving in and out new parts and tools for the smaller engine. There were also different-sized pallets for the different engines. The solution was to bring a small amount of all the parts on flow racks to the operator on the line. The tools needed for all three engines were mounted over the production line. It was also necessary to create a flexible pallet that could hold any size engine. This eliminated the equipment changeover completely, allowing the plant to build the engines in any order it wanted on a mixed-model assembly line. It could then make a repeating (level) sequence of all three engine sizes, so it matched the mix of parts ordered by the customer (see Figure 10-3). There were four benefits of leveling the schedule:

1. *Flexibility to make what the customer wants when they want it.* This reduced the plant's inventory and its associated problems.
2. *Reduced risk of unsold goods.* If the plant makes only what the customer orders, it doesn't have to worry about eating the costs of owning and storing inventory.
3. *Balanced use of labor and machines.* The plant can create standardized work and level out production by taking into account that some engines will require less work and others will require more work. As long as a big engine that takes extra work is not followed by another big engine, the workers can handle it. Once the plant takes this into account and keeps the schedule level, it can have a balanced and manageable workload over the day.

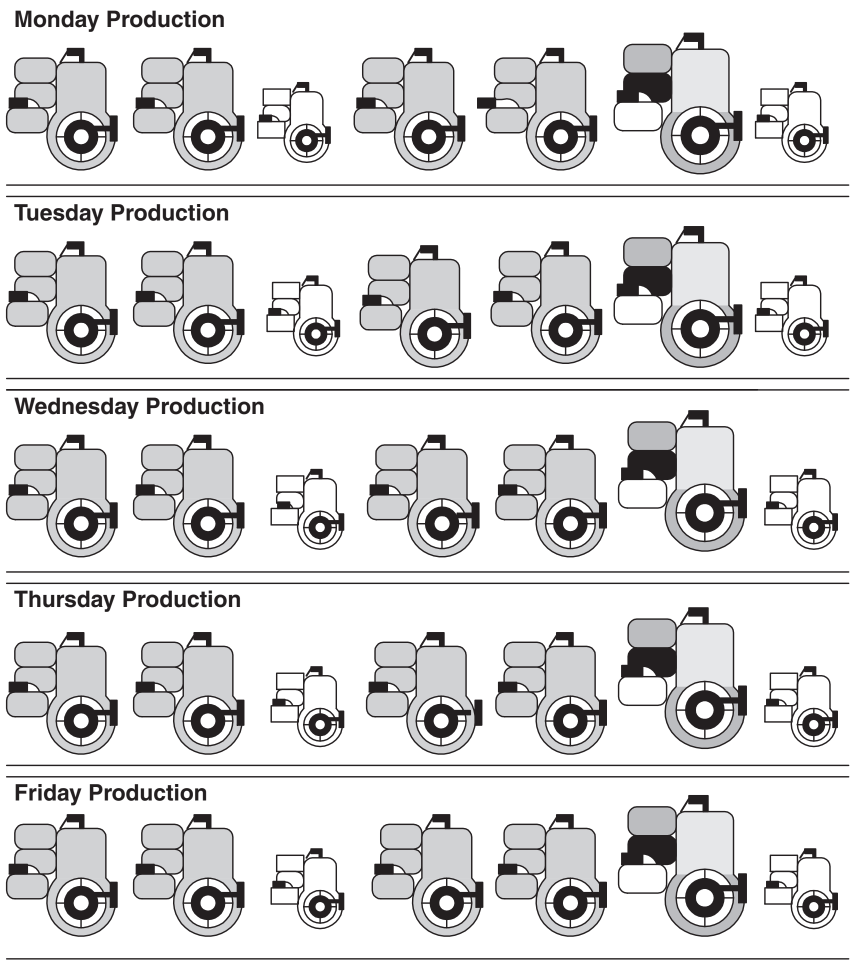

Figure 10-3. Mixed model production (leveled)

4. *Smoothed demand on upstream processes and the plant's suppliers*. If the plant uses a just-in-time system for upstream processes and the suppliers deliver multiple times in a day, the suppliers will get a stable and level set of orders. This will allow them to reduce inventory and then pass some savings on to the customer so that everyone gets the benefits of leveling.

None of this would have been possible if the plant hadn't found a way to eliminate the setup time for changeover.

Though it may seem unrealistic that you could do this in every circumstance, several decades ago Shigeo Shingo proved in his time studies that this was exactly what you had to do. Shingo was not a Toyota employee, but worked closely with Toyota. He was a meticulous industrial engineer who paid attention to every microscopic reach and grasp of the worker. In the Toyota style, he thoroughly analyzed the setup process for large stamping presses and discovered that most of the work performed fell into one of two categories: it was muda or it was something that could be done while the press was still making parts. He called the second category "external setup," as opposed to "internal setup," which was work that had to be done while the press was shut down.

In traditional mass production, the first thing the setup teams did when they performed the changeover of a production line from one model to another was to shut down the press. Shingo wondered how much of the changeover he could perform while the press was still running, so he organized an operator's workplace for that purpose and made other technical improvements until there was no more setup the operator could do while the press was running. Things like getting the next die and tools, preheating the die, and setting it in place beside the press were external and could be done while the press was making parts. When he finally shut down the press, all that was left to do was basically to swap the dies and start it up again. Amazingly, these several-hundred-ton presses that previously took many hours to change over could, it turned out, be changed over in minutes. Think of it like a racing pit crew that quickly services and gets the car back on the track, often in less than a minute.

Over the years changeover has become a kind of a sport in Japan, a manufacturing equivalent of an American rodeo. On one trip I took to Japan in the 1980s, I visited a Mazda supplier of stamped door panels whose team had recently won a prize in a national competition for changing over a several-hundred-ton press in 52 seconds.

Leveling the Schedule—Inventory's Role

Leveling the schedule has profound benefits throughout the value stream, including giving you the ability to plan every detail of production meticulously and standardizing work practices. If you visit a Toyota plant or a Toyota supplier, you will see the great pains taken to level the schedule. The best Toyota suppliers also work on the assumption that Toyota's demand for their parts will be level. This is a risk, because not keeping finished goods inventory means leaving themselves fully exposed to any wild variations in their customer's volume and mix of products. They can do this and still sleep at night because Toyota is a very reliable customer and levels out its production schedule.

For example, Trim Masters is a U.S. supplier in Georgetown, Kentucky that makes seats for the Camry and the Avalon manufactured there. Trim Masters builds and delivers seats just in time, based on a broadcast from the Toyota plant that orders one seat at a time. From the time the orders are placed, Trim Masters has three hours to build the seats, put them on the truck in sequence, and deliver them to the Toyota plant, so they appear on the assembly line in the exact order needed for production. Trim Masters orders parts just in time from its suppliers and keeps very little inventory, with inventory turns of 128 times per month. The Avalons and Camrys take different seats that require different parts, so Trim Masters has to trust Toyota will make the mix of Avalons and Camrys that it projects. If there is a sudden spike in Avalon seat production, Trim Masters will run out of parts and must pay for emergency delivery of parts. This happens routinely for U.S. auto companies, providing many truckers and helicopter pilots a good living on high-priced expedited freight. This happens from time to time with Toyota, but by and large it carefully maintains a leveled schedule and builds what it says it will build.

Most suppliers are not like Trim Masters and must satisfy customers whose demand fluctuates significantly. In these cases, TPS experts will often recommend keeping at least a small inventory of finished goods. This seems to contradict lean thinking. Theoretically, the leanest solution is to build to order and ship just what the customer wants. (If you are going to keep inventory, why keep the most expensive inventory—finished goods? Instead, build to order and store only raw material inventory.) But this reasoning doesn't consider the importance of *heijunka*. A small inventory of finished goods is often necessary to protect a supplier's level production schedule from being jerked around by sudden spikes in demand. It may seem wasteful, but by living with the waste of some finished goods inventory, you can eliminate far more waste in your entire production process and your supply chain, if you keep your production level.

This is one reason why companies that have successfully applied TPS often schedule their production with a combination of building to order and maintaining a pre-determined level of finished goods inventory. The case example at the end of this chapter shows a company that builds high-volume seasonal products to hold in inventory and then builds other products to order. This combination allows the company to level the schedule over the year, have a smooth flow, and build most of its products to order.

"Build to Order" Yet *Heijunka*

Cho's quote at the opening of this chapter suggested customers may have to wait a little longer if they want to order a vehicle specially built for them. He is not willing to sacrifice the quality and efficiency benefits of *heijunka* for the sake of "build-to-order." Yet, other car manufacturers have developed build-to-order systems, potentially giving them a competitive advantage. One of the conventional build-to-order solutions is to keep a lot of finished vehicle inventory in huge dealer lots around the country and swap vehicles among dealers that match custom orders.

So is Toyota satisfied with asking customers to wait while they may be able to get the specific car they want from a competitor? In response to this challenge, Toyota has developed a solution that will allow it to level the schedule and at the same time build to order. They are never satisfied with either/or. Alan Cabito, Group Vice President of Toyota Motor Sales, explained:

> *The Toyota system's not a build-to-order system. It is a "change to order" system. And the big difference is that we have cars moving down a line that we change specs on. We've always done that. But we're just taking it another big notch up. We pick a car on the line, any car, and change it. And obviously there are guidelines on how many changes you can do in a day, so we consistently have the parts available to do it.*

This is all done within the leveled schedule created several months in advance. Cabito explained further the realities of the mixed model production line:

> *You might have a van unibody and a truck and then you might have another truck, so that the van was every third vehicle. That isn't going to change. You can change the color, which is not simply paint, it's interior and everything else. You can have matching mirrors, etc. There's a lot of complexity to changing color— you have to change virtually all the accessories. And the way that gets managed is on the allowance of how much change can take place. There will be a limit to the number of green, leather-interior Siennas we can make in the same day.*

As usual, Toyota experimented with building to order with an actual pilot—the Solara, a sporty coupe version of a Camry—in the Canadian plant. It is relatively low volume. For Solara they achieved 100% "change to order." For the Tacoma truck there are a huge number of engine combinations, and they were able to achieve about 80% "change to order" from dealers who called in with customer requests. Cabito gave me a sales perspective on how this works:

> *We place a single month's order three times. We'll order it four months out, three months out, two months out. During that time, they set up all the components and suppliers. For July production, the final order will be placed in May. So your order's out there 60 days in advance. Then every week we can change the order in the U.S. plants. Every week we can modify anything that's unbuilt, except for the basic body type.*

The important point here is that the Toyota culture does not allow managers and engineers to conclude, "That cannot be done here." The rigid principle of *heijunka* does not stay rigid for long. On the other hand, it is not simply thrown away because of a new trend such as build-to-order. The question is: How can we accommodate the customer's desire to make choices and get the car quickly without compromising the integrity of the production system? In true Toyota Way problem-solving style, the engineers carefully studied the situation, experimented on the shop floor, and implemented a new system.

Heijunka in Service Operations

Leveling out a work schedule is easier in high-volume manufacturing than in typically lower-volume service environments. How do you level schedules in a service operation where service providers are responding to customers and the lead times on service work vary widely case by case? The solutions are similar to the solutions in manufacturing:

1. *Fit customer demand into a leveled schedule*. This is more common in service operations than you may think. Why is it that doctors and dentists schedule procedures and you need to fit into their schedule? So they can level the workload and have a constant stream of income. Time is money in service operations.
2. *Establish standard times for delivering different types of service*. Again, the medical field is instructive. Even though everyone has somewhat different medical needs, doctors and dentists have been able to establish standard times for different types of procedures. And they separate diagnosis from the procedure. You visit, they diagnose you, and then, in most cases, they can predict the time that will be required for your procedure.

Toyota has effectively been able to level the schedule for product development even though lead times are months or even years. In most cases, Toyota will make minor updates to a vehicle every two years, adding features and changing styling, and will do a redesign of the vehicle every four or five years. Toyota product development works according to a matrix, where the rows are the different Toyota vehicles—Camry, Sienna, Tundra—and the columns are years. They decide when each vehicle will be freshened and go through a major redesign. They intentionally level the schedule so a fixed percent of the vehicles are being redesigned in any one year.

Planning when vehicles are scheduled for redesign would be futile if the lead times required to actually design and develop a vehicle were unpredictable. This is where Toyota has a big edge over some of its competitors. While some auto companies let the start of production slip by months or even a year, Toyota is like clockwork. Development milestones are met with virtually 100% accuracy. So the leveled plan becomes reality.

Toyota also has found there is a cadence to the workload requirements over the life of the development project: the workload is relatively light early in the conceptual stage, then builds up as they get to detailed design, and then reduces again in launch. By offsetting different vehicle projects, they know when one is peaking and others are in the light period and can assign the numbers of engineers to products accordingly. They also can flex the number of people needed by borrowing engineers from affiliated companies (suppliers and other divisions of Toyota, such as Toyota Auto Body). Affiliates can come onto projects as needed and then go back to their home companies, allowing an extremely flexible system and requiring minimal full-time employees. This is the result of other Toyota Way principles, particularly standardization. Toyota has standardized its product development system and the product designs themselves to the point that engineers can seamlessly come in and out of design projects, because their engineers have a standardized skill set similar to the Toyota engineers' and years of experience working in Toyota's system. The principle of long-term partnering that we will discuss in Chapter 17 allows Toyota to have a trustworthy and capable set of partners who they can depend on for extra help when needed.

In short, it is possible to level the schedule in service operations. But there are some base requirements. You must follow all of the other Toyota Way process principles—flow, pull, standardization, and even visual management—to get control over lead times. Standardization is critical to controlling lead times and also to bringing people on and off the projects to address peak workloads. You must also develop stable partnerships with outside companies that are capable and that you can trust.

Putting Leveling and Flow Together—A Tough Sell

Every business would like to have a consistent volume over time so there is a consistent and predictable workload. That is an easy sell in concept. But if your sales department does not behave like Toyota Sales by cooperating to avoid spikes in demand, what can you do?

The TPS expert might suggest that a manufacturer hold some finished goods inventory and build at a leveled pace to replenish what the customer takes away in a pull system (discussed in Chapter 9). The manufacturer screams, "But we have 15,000 part numbers!" The expert says, "Look for a smaller number of part numbers that are in big demand and perhaps even seasonal, build those when you have fewer real orders, and then keep those in inventory." That is, use a combination of build-to-stock and build-to-order like the aluminum gutter company in the case study at the end of this chapter. This sounds reasonable to the manufacturer. But then comes the hard sell. The TPS expert says to changeover frequently to level the mix of products built every day. Most manufacturers balk. After all, it is so convenient building in batches, making product A for a while, then retooling, and making product B for a while, and so on. Quick retooling does not seem possible, until an expert shows them how they can do a three-hour changeover in five minutes. Even then, it is difficult for many manufacturers to maintain the discipline of quick changeover. And the real root cause of the problem may be sales promotion strategies that contribute to uneven customer demand. The most sophisticated lean enterprises begin to change their policies in sales to maintain a level customer demand. This requires a deep commitment at the very top of the company, but these organizations quickly find the enormous benefits of *heijunka* make it a worthwhile investment.

It cannot be overstated. To achieve the lean benefits of continuous flow, you need Principle 4: *Level out the workload (Heijunka)*. Eliminating *muda* is only one-third of achieving flow. Eliminating *muri* and smoothing *mura* are equally important. Principle 4 focuses on *muri* and *mura* by leveling your product volume and mix and, most importantly, leveling out the demand on your people, equipment, and suppliers. Standardized work is far easier, cheaper, and faster to manage. It becomes increasingly easy to see the wastes of missing parts or defects. Without leveling, wastes naturally increase as people and equipment are driven to work like mad and then stop and wait, like the hare. Working according to a level schedule applies to all parts of Toyota, including sales. Everyone in the organization works together to achieve it.

Case Example: Building Aluminum Gutters on a Level Production Schedule

These days aluminum gutters for houses, at least in the U.S., are mostly built to order, on-site at the house. Rolls of materials are brought to the job site, where they are cut to length, end caps are formed, and the gutters are installed. A plant in the Midwest makes much of the material that installers use—the rolls of painted aluminum. While these rolls of aluminum are not complex, there is variation in the width of the gutters, the length, and the colors. They also are packaged in different boxes, depending on the customer.

This company originally adopted the build-to-order model. Deliveries were mostly made on time, but the process of getting raw materials, scheduling operations, building the product, moving the finished goods to a warehouse, and then shipping those goods from the dozen or so shipping docks was chaotic, to say the least. There was inventory everywhere. Yet the plant regularly was short of critical materials needed for the gutters ordered. Costs of expediting shipments to large customers were getting higher and higher. People were added and laid off with regularity. A big problem was the seasonality of the business. Big box warehouse stores like Home Depot bought large quantities of gutters in the spring and early summer and then business dropped dramatically for the rest of the year. So a large number of temporary workers were added in the peak season.

The Midwestern gutter plant decided to hire a consultant who used to work for the Toyota Supplier Support Center. The consultant said the unthinkable—the overall operation would be leaner if the plant would build select products to store away in inventory! This meant selectively adding some waste. They followed the consultant's advice.

He knew there is not one type of finished goods inventory, but four types. The first is real built-to-order product that should be set in a staging lane to be put on a truck immediately. The second is seasonal product for high-volume items the plant knows it will sell, that should be built steadily throughout the year and accumulated in a seasonal inventory buffer, which then will be drawn down in the busy spring/summer season. The third is safety stock, which is inventory used to buffer against unexpectedly high demand for products that are not in the seasonal buffer; it is customer-driven variation. The fourth is buffer stock, which is held to cushion against downtime in the factory, so customers will continue to get their product even when machines are down for repair; it is plant-driven variation.

On the consultant's recommendation, each of these four types of inventory was stored in a separate area at the aluminum gutter plant, so that everyone would always be able to see how much inventory of each type was available (Toyota Way Principle #7).

The inventory was replenished by using the *kanban* system (cards instructing the production line to make a certain quantity of a certain end product) explained in Chapter 9. For example, the largest amount of inventory is the seasonal inventory buffer. It is built up during the off-season and reaches a peak just before the spring, when sales are highest. There is a pre-specified amount of seasonal buffer and, based on that forecasted amount, *kanban* is used by the production cell to make only that remaining number of packages needed. In front of the inventory is what looks like a clothesline labeled with months of the year. For example, the amount that should be completed by August, based on a constant level of production over the year, has a sign saying "August." In August, if the inventory pile is larger than should be built by that time, the pile of inventory will have moved beyond the August sign and everyone will know that there is an excess inventory problem needing to be solved.

In *kanban*, discussed in Chapter 9, the information flow begins with the customer order and works backward through the operation. In this company, a final cutting and packaging (one-piece flow) cell gets customer orders that it has to build to order. But when those orders are low, the workers do not have to sit around with nothing to do. They can build to the seasonal inventory buffer or build to replace any safety stock or buffer stock that has been used. The seasonal inventory, safety stock, and buffer stock that need to be built are represented by *kanban* cards. The cards are sorted by a planner into a visual scheduling box called a "*heijunka* box," which levels the schedule. For each product, the box says what to make at 8:00 a.m., 8:10 a.m., 8:20 a.m., etc. Cards are put in the slots and delivered to the production cell. These tell the cell what to make and pace the work of the cell. As the cell uses materials, like the painted aluminum product, a *kanban* is sent back to the prior operation asking it to make more. Pull has been established all the way back to suppliers, like the paint supplier.

At the suggestion of the TPS consultant, other improvements were made, such as standardizing work procedures, reducing changeover time, and putting in error-proofing devices (discussed in Chapters 11 and 12). The result was a very smooth flow of product through the facility, so smooth that all outbound shipping could be handled through two docks with the other 10 closed down. In addition, the plant achieved incredible performance improvements. The overall lead time for making product was reduced by 40%, changeover time was reduced by 70%, WIP of painted product was reduced by 40%, inventory obsolescence was reduced by 60%, and on-time delivery was close to 100%. A typical lean transformation!

Note

1. Toyota would never let go of or demote workers displaced by productivity enhancements. This shortsighted cost-saving move would create ill will toward the company and prevent all other workers from cooperating in future *kaizen* efforts. Toyota always seeks alternative value-added work for workers displaced by production improvements.