Chapter 20: Measuring to Improve Performance and Results

Using the information that shows up in management dashboards on a screen can wreck your company. Measuring a process does not put you in control of it. A process is totally self-controlled. You always get what the process gives you. Organizational processes work well or do not because of their design.[1] A business process that cannot achieve its aims is the wrong design for its purpose. It's often a mockery to use business performance indicators to manage a company or department.

Financial performance and productivity results are interesting for a business to know, but they are point-in-time samples of what your processes were producing at that moment. The data in weekly, monthly, and annual performance charts are the results your processes let you have. When you take another sample in the following period, you will get a different result. If a performance chart contains only natural outcomes, then that is the best the process can give, and there is nothing more you can do about it. If you don't like it, the only choice is to design and build a better process.

Plotting the frequency of process results taken over a long time gives you a distribution curve of process performance. It may look like the bell curve shown in Figure 20.1, or it may take some other shape. Every time you measure a process variable, the value will be somewhere on the distribution curve. When you take a sample and show its measure in a monthly report, it is a result from some point on the curve. One month you get one value for the indicator. The next month the sample comes from elsewhere on the distribution, and you get another result. Wait a while, and the result will change again. Variation in results is how a process works. You get fluctuating values over a period—some good, some bad, and many in between. If you expect great performance results every month, you are expecting the impossible because a process produces outcomes anywhere along the curve. You would get far more management value from a monthly report if it also showed the frequency distribution of the performance measure alongside its trend plot and where on the distribution the point was located.

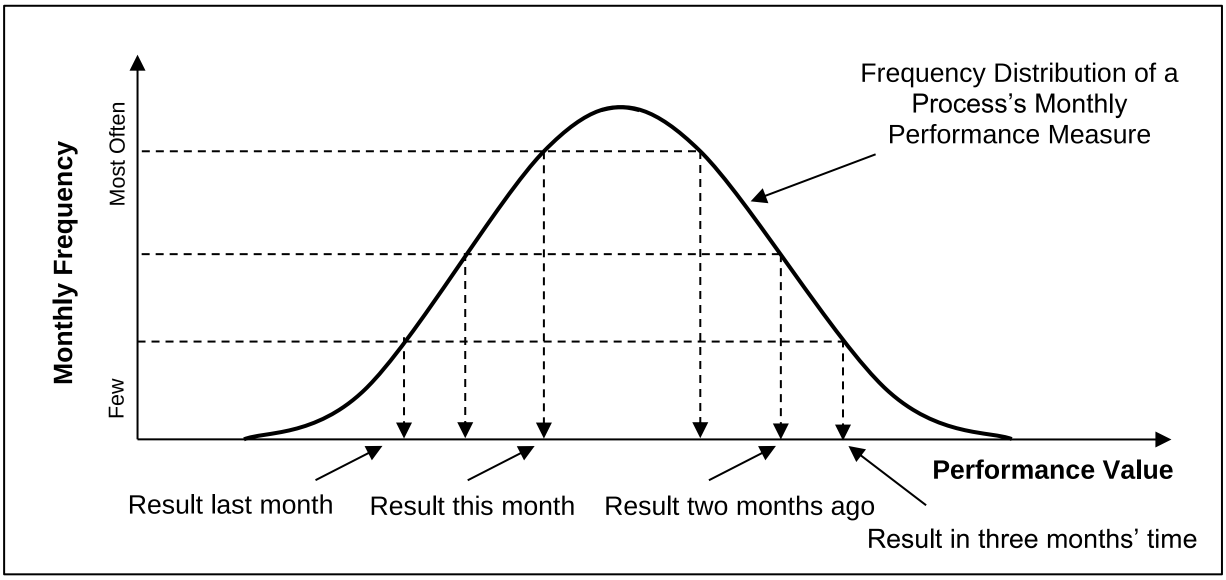

Figure 20.1—Distribution of a Process Performance Measure

Once you have a lousy process design, you will get lousy results until you get a better design—nothing else you do will bring process management success except having the right process design in the first place. You might think that collecting management indicators will aid in correct decision making. But if the process is already stable, every change made to it can worsen performance. Once managers and people get a result that is not ideal, they will tinker and adjust the process. But it is no good to make changes unless a process is out of control. First it is necessary to check whether the process is running stably, and if it is, then leave it alone. Any tweaking will only ruin it. To be of any use for process management, your performance indicators need to measure the extent of the deviation from the ideal process outcome, its target "bulls-eye", and see whether the distribution is steady. Indicators are useful when they help people make the right decisions to improve process designs or help them find out-of-control process effects and remove them. If they are used for other purposes, they will cause you to ruin your business performance.

Monitoring a Process and Its Process Steps

Key Performance Indicators (KPIs) are the measures used to monitor process outcomes—they show the total effect of all influences on the process. Performance Indicators (PIs), on the other hand, are used in the process steps—they tell you about a step's variability.

How do you monitor and measure a process such as production, maintenance, or a continuous improvement process so that you can make it world class? What about a job—how do you monitor and measure work and use the information to make it highly successful? What about the tasks in the job—how can you make them more efficient and effective? In the Plant Wellness Way, every one of these questions is answered the same way: by monitoring the trend and stability of step inputs and outputs, as well as checking the trend and stability of the whole process to see cumulative effects.

Once you have the process design captured in process maps and procedures, you have everything you need to properly monitor, measure, and improve the process. On the process map, important inputs and outputs from each step are identified. In ACE 3T procedures, all variables in the process step are identified and quality standards are set for them. The ideal results and tolerances are defined using measurable performance terms with values that can be charted. The process map and supporting documents contain all the performance measures you need to monitor the whole process and its steps' performance. Useful statistical charts and graphs can be developed with the results and used to observe and compare process and step behaviour.

Figure 20.2 is a process map of a purchasing process. The process is used to buy what is needed in as simple and straightforward a manner as possible. In the end, what you want from the buying process are the things you asked for. A measure of total process success is whether you received the items you wanted. A second measure of overall success is whether you got your order on time. A third is whether the order was in perfect condition. A fourth is whether the invoiced amount matched the price in the offer. These purchasing requirements are what the process should always deliver with 100% certainty. Less than 100% conformity on all measures means that the process design failed. Somewhere, one process step, or more, went wrong. If you monitor the step outputs, you will know which ones failed.

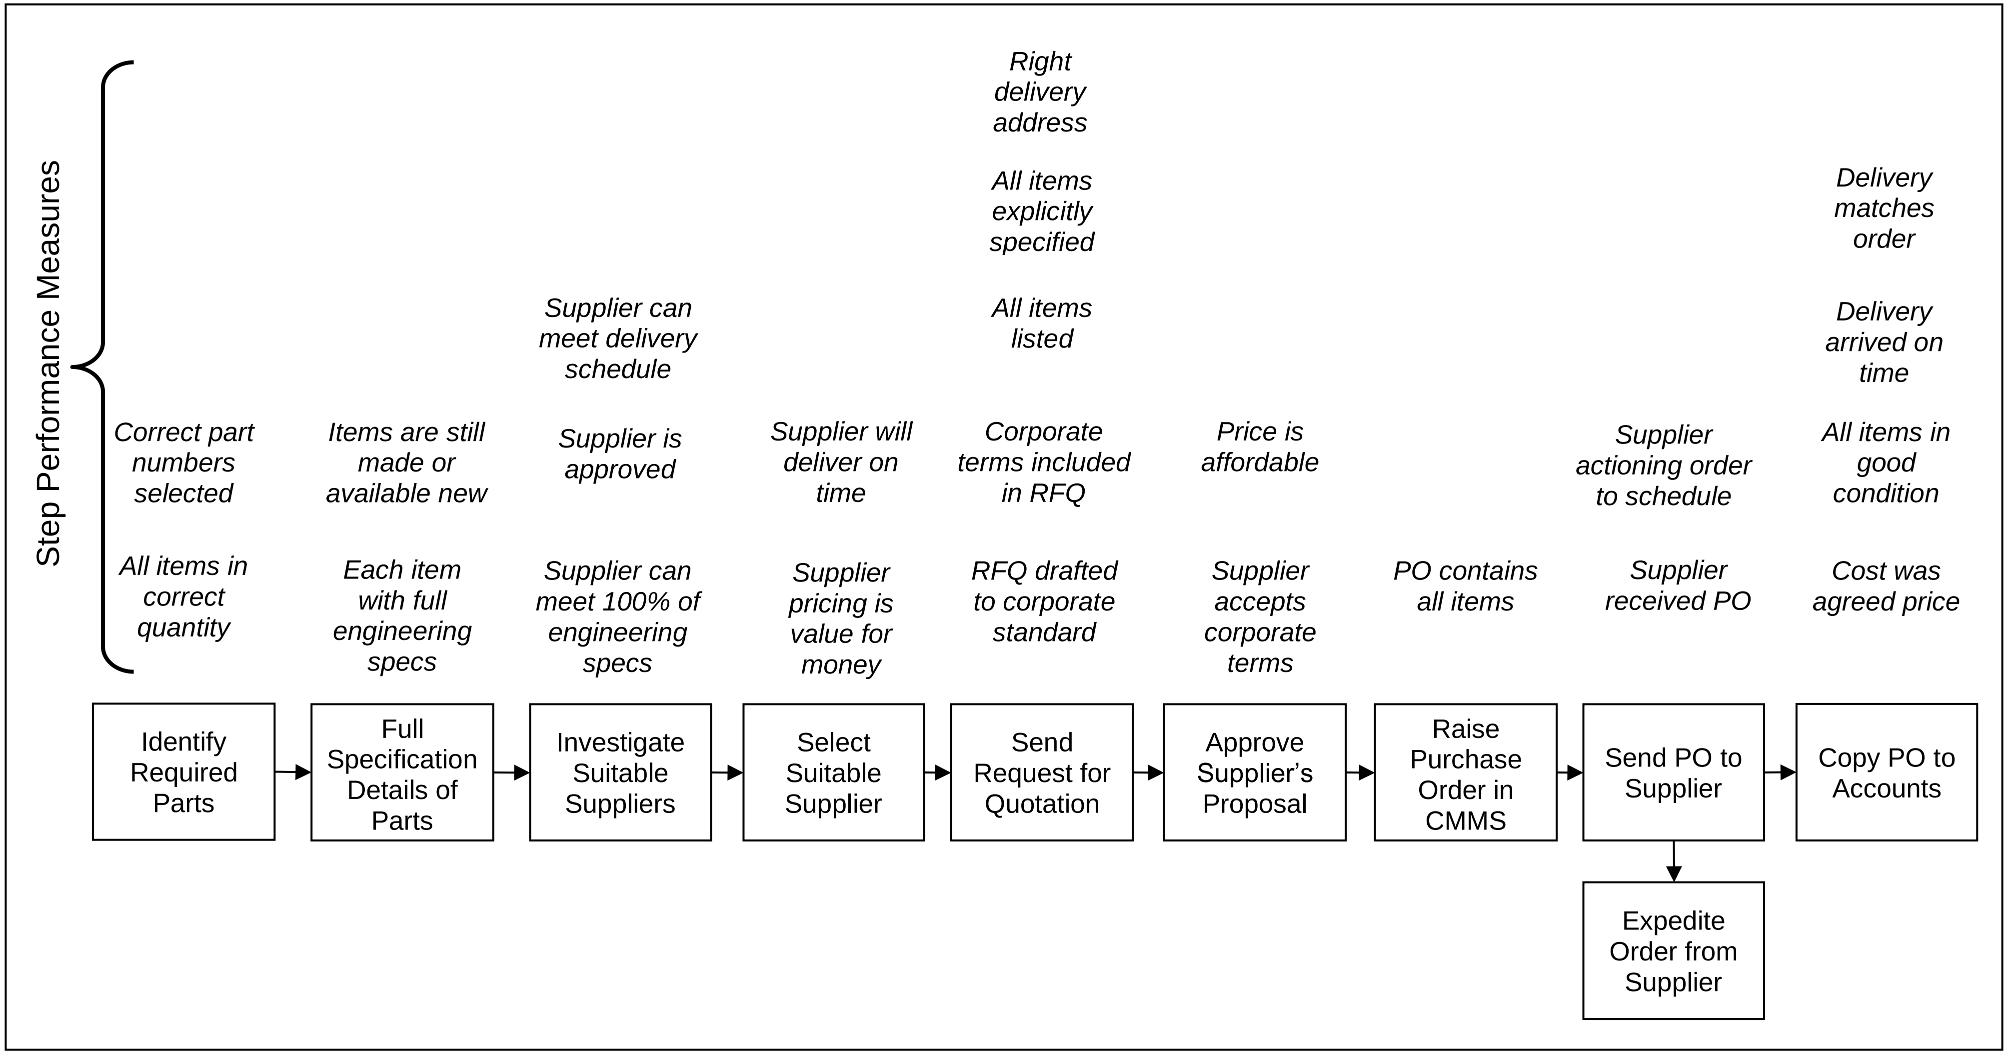

Figure 20.2—Outcome Quality Measures for Monitoring the Performance of a Procurement Process

A process fails when a process step fails. If a result from the complete process is not acceptable, you need to find and correct the steps in the process that control the result. To do that, you need indicators of each step's success. The measures you use are the required step outputs noted on the process map. The beauty of using the measures written in the process map is they are what the process design is meant to deliver. They directly result from the behaviour of the step and reflect the activities done in a step by those running the process. The performance values are available because the procedure asks for them to be generated as the work is done. If it is necessary to understand why and how a step failed internally, you observe and measure the variables in the step, even going as far as statistically charting them to look for the tell-tale signs caused by the presence of defects. When you report on the process and talk to people about its performance, it will be with facts and answers specific to the process.

Process Performance Distribution Curves

The performance indicators you take from a process are samples of the process behaviour. Each KPI and PI is a result generated as the process was being run. These indicators test the process's performance. Over time, the indicators will tell you the reliability and stability of your business processes. Look at the run charts and frequency distribution curves of your maintenance and operations processes if you want to see your company's operating future. Those graphs and plots represent future performance if all remains the same. The distribution curves of your KPIs and PIs reflect your business's outlook until you change policies, process designs, and practices to those that produce more successful outcomes.

It's interesting to watch the faces of senior managers when they first see the distribution curves for their Key Performance Indicators, and you explain what they mean. For years, they have seen individual monthly KPI results plotted on graphs. Typically, the KPI graph shows the trend line across a period. Using a run chart alone to observe process behaviour is incomplete process management. When you also show company and department managers the frequency distribution curves of the KPIs they are responsible for and explain to them that the shape indicates the process's chance of achieving the KPI, they start to understand what their business and its processes have been doing. When the curves are wide and success is uncertain, managers get a worried look. When the curve is tight but over the wrong performance value, they get a frown from wondering how they can afford to reengineer the process. When the distribution is tight and "right on the money," you get a big smile.

The run chart and the failure distribution curve developed from the chart data are two simple techniques for analysing both system and step performance. Figure 20.3 is adapted from W. Edwards Deming's book *Out of the Crisis.*[2] It is a run chart of fire events in an industrial operation over a five-year period.

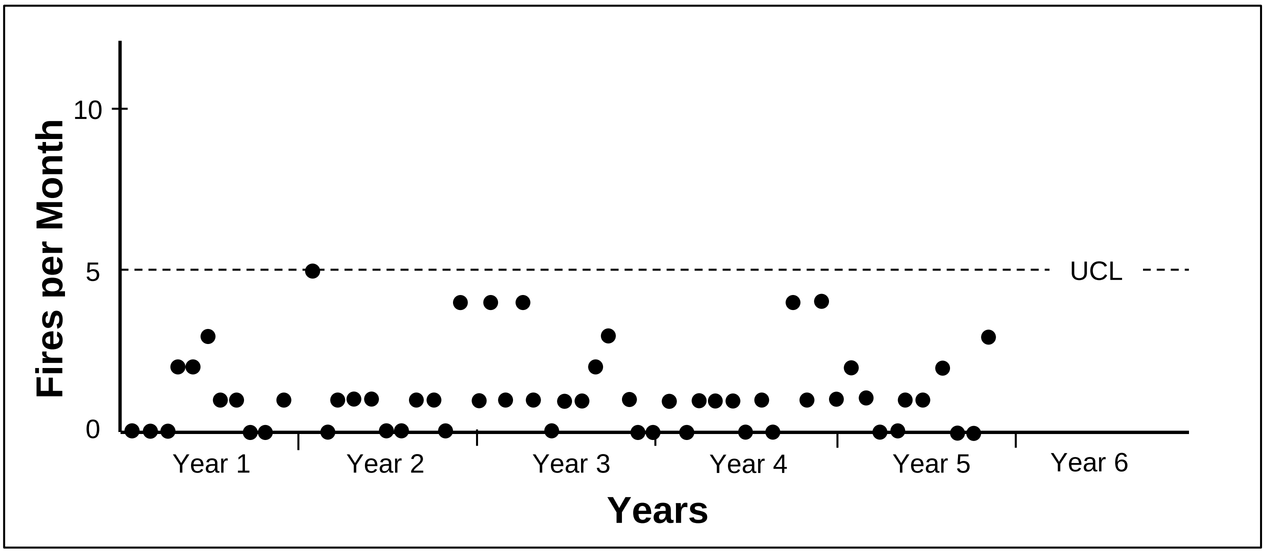

Figure 20.3—Run Chart of Fire Events in an Industrial Plant

Looking at the run chart, unless something changes, Years 6, 7 and beyond will be the same as the past five years. The run chart confirms that a persistent problem exists and will continue next year, and the year after, and again the year after that unless the process is fixed. The repetitiveness of the fires indicates that a common cause problem(s) is built into the way the company works—it gets fires by design. This place must have been endless trouble for its people and management to operate.

Within the run chart is “hidden” information about the probability of the next fire event. Over 58 months, there was one or more fires in 38 months. Therefore, the odds are 38/58 (a 0.66 chance) that there will be a fire next month. Figure 20.4 is a distribution plot showing the frequency of fires per month in the months with fires.

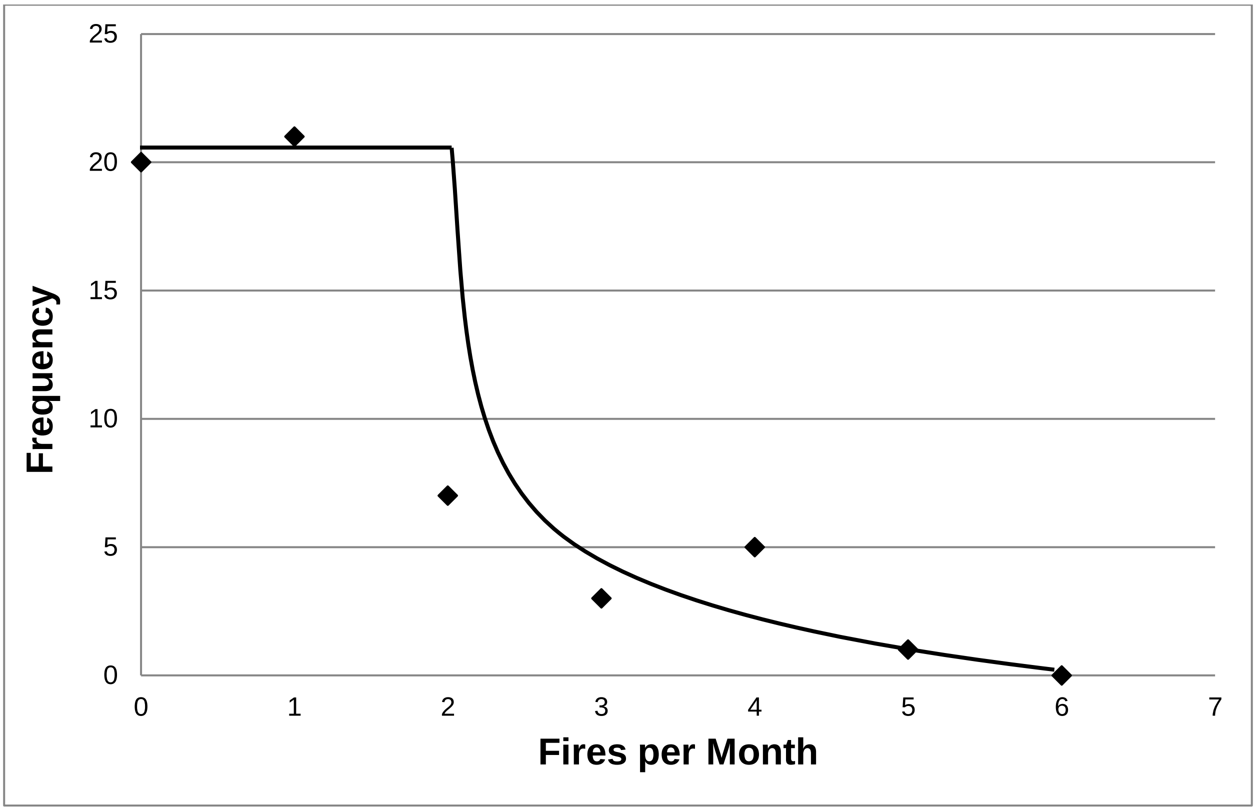

Figure 20.4—Probability Distribution Curve of Fires in Months with Fires

In the 38 months with fires, there were 68 fires. In 24 of the 38 months, there was one fire per month (24/38 = 0.63 chance of one fire in months with fires), and 5 months had two fires (5/38

= 0.13 chance of two fires). In fact, there is a 0.66 x 0.13 = 0.09 chance—odds of about 1 in 11—of having two fires in the next month. In 3 months, there were four fires (3/38 = 0.08 chance of 4 fires in months with fires), 5 months had four fires (5/38 = 0.13 chance of four fires), and 1 month had five fires (1/38 = 0.03 chance of five fires). No months had six fires, but from history, the possibility existed that it would happen one day.

There is still information in the run chart. The frequency and density of months with one fire versus the density and frequency of months with three or more fires indicates that the months when one fire happened were not the same types of months as those when three or more fires happened. Something significantly different happened in the months with many fires.

Another way of using run charts is shown Figure 20.5, which graphs the production downtime caused by equipment breakdowns each week in the industrial plant introduced in Chapter 3 (see Figure 3.12).

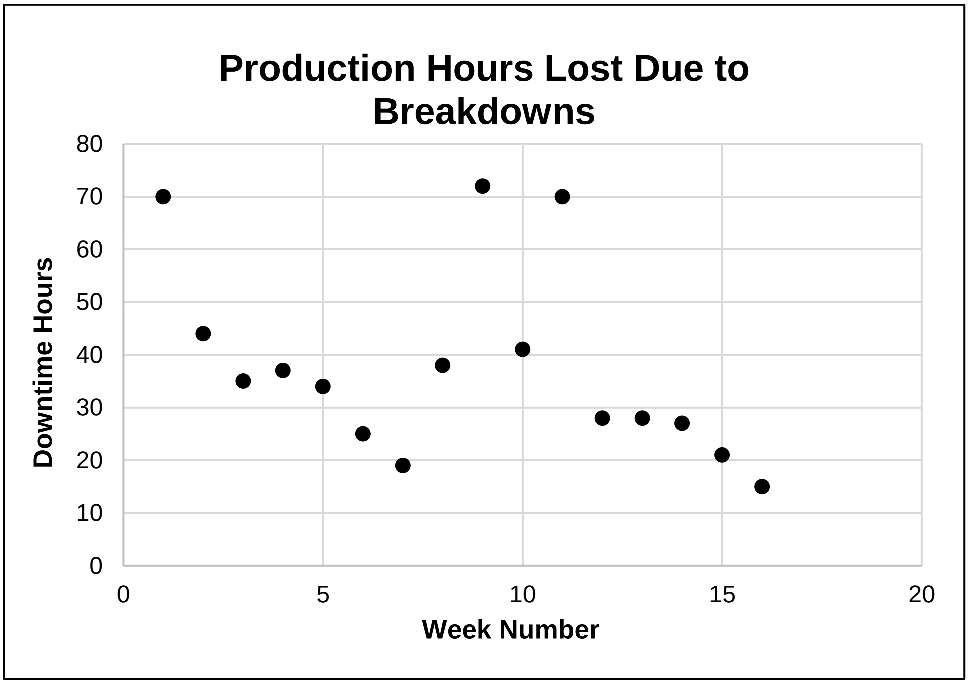

Figure 20.5—Production Hours Lost Due to Breakdowns

Figure 20.6 is the frequency distribution curve for the data. At first it looks like a normal bell-shaped distribution curve, but suddenly there is a discontinuity. This company has two types of breakdowns—usual and severe. The lost time from usual or standard breakdowns average 25 to 30 hours a week (consistently between 15 to 45 hours), but from time to time the company has breakdowns that are catastrophic to production uptime.

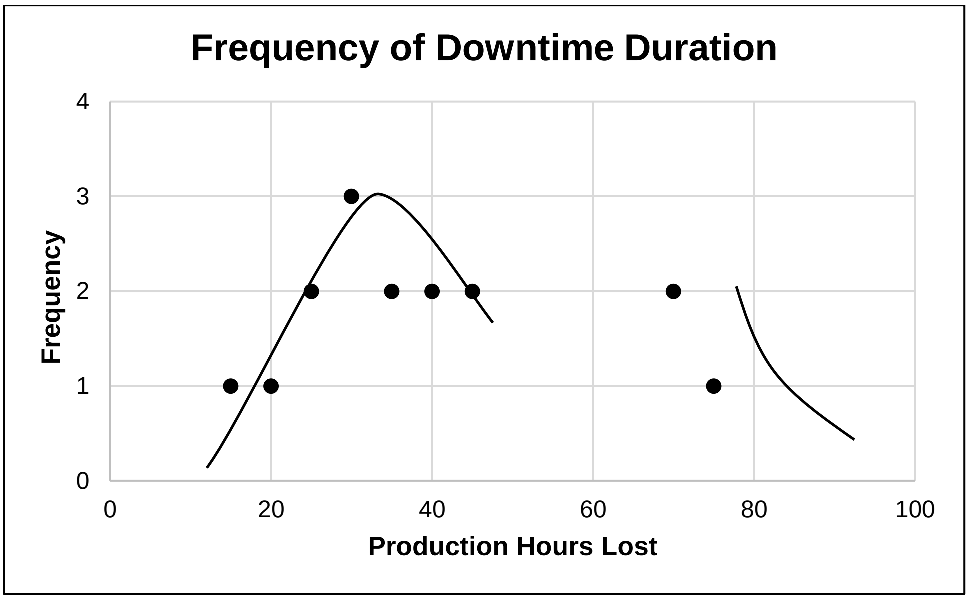

Figure 20.6—Distribution of the Production Breakdown Hours

Example 20.1: Production Plant Performance

Figures 20.7 and 20.8 are examples of getting stuck trying to fix your current business when instead you ought to throw your troubles away and build a better business system. This business is a renowned company in its home country. It is well respected and profitable enough. But it could easily be much wealthier. There are vast new fortunes sitting in the business, but they will never be seen by its owners and managers. They are totally focused each day on trying to make the existing business processes and system work properly. Their operational processes need to be redesigned to remove the problems stopping them from delivering the organization's purpose.

The two charts cover a period of 10 years of operation (the years and company are intentionally disguised). Figure 20.7 is a run chart showing the dates and durations of all outages. The periods between downtimes are the operation's uptime. Figure 20.8 is the uptime frequency distribution curve derived from the run chart. It is a "chance of success" chart because it shows how often each uptime period happened during the 10 years. The shape of the frequency distribution curve foretells what outcomes can be expected in the future.

Look at the density of outages in the run chart. There were years of frequent trouble and times of less. Things must have gone bad often to create the dense regions. The last two years show fewer minor trips and some of the best uptimes. Note the stratification of dots: many less than 50 hours of outage and far fewer greater than that. Having random outages of 10- to 48-duration is common in this operation.

The frequency distribution curve turns the splatter of events in the run chart into a clear message about this company's chance of uptime success. There are dozens of short running periods of less than 2 days and quite a few periods from 2 to 10 days. Very few uptime periods last longer than 20 days. The story hidden in the run chart becomes obvious in the frequency distribution plot: this business is missing out on a fortune in operating profits (worth several million dollars annually) because it has so many stoppages interrupting production. Yet, the same operation has several times delivered over two months of uninterrupted production. The process design can get long production runs, but things happen in the organization to prevent success.

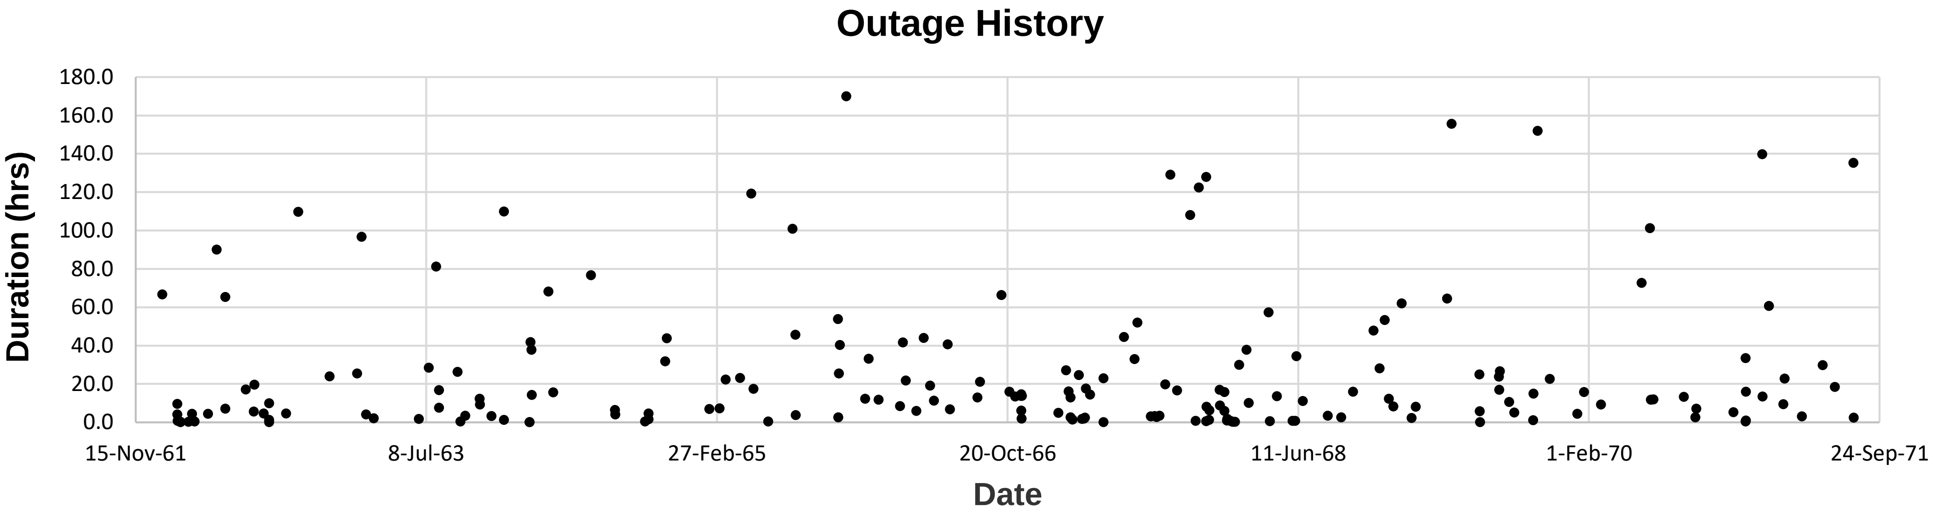

Figure 20.7—Plant Outage Duration Run Chart

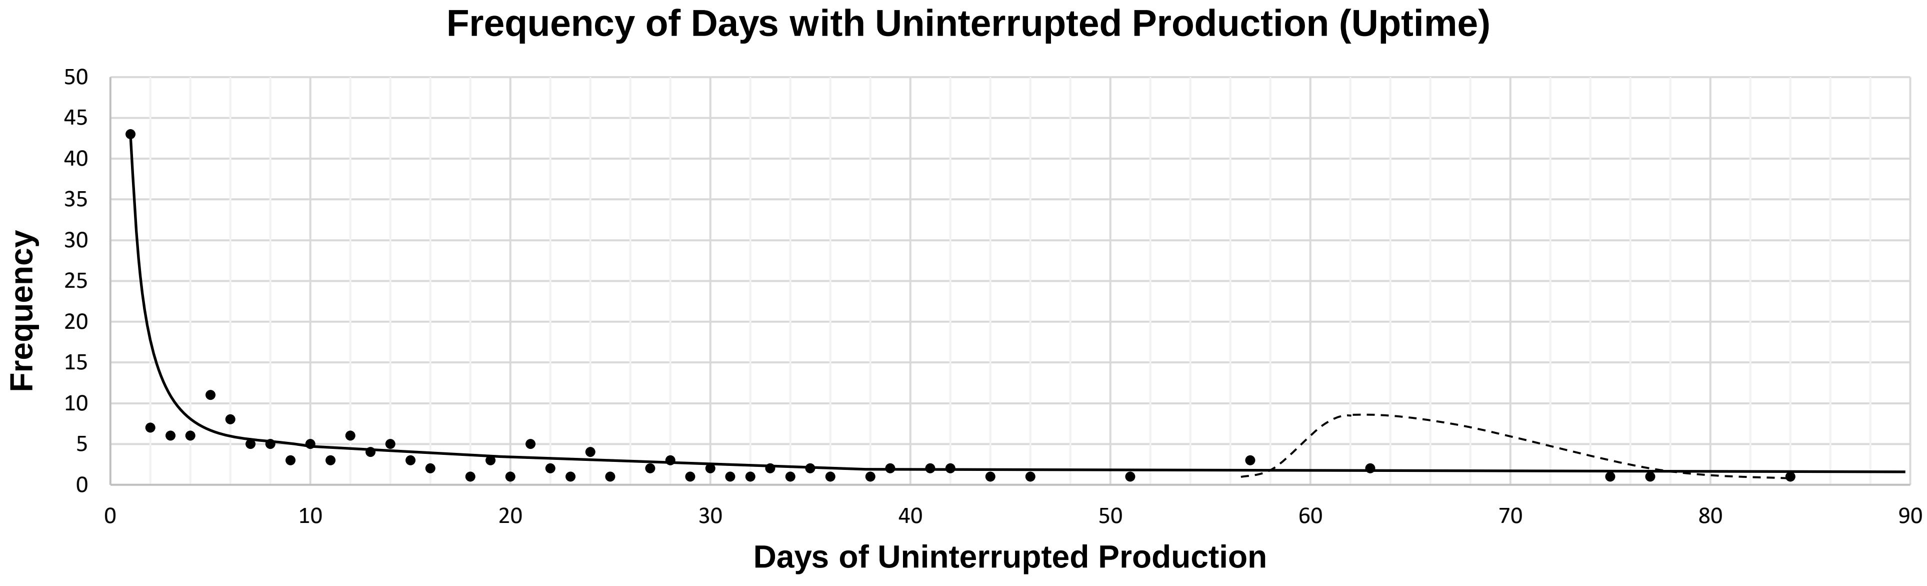

Figure 20.8—Distribution of Uninterrupted Production Days

The whole area under the curve in the frequency distribution is the probability of plant uptime. You can estimate by eye that the area between 0 and 20 days is larger than the area from 20 to 80 days by about four times. It is four times more likely that the next plant outage will be fewer than 20 days away than it will be more than 20. It looks about three times more likely that an outage will happen in fewer than 10 days rather than longer. It is clear from the “hump” in the curve between 0 to 2 days that short duration outages are a regular problem for the business. It would be very valuable to analyse the reasons for the outages in the hump to learn what causes are producing so many stoppages and eliminate them so the process can work to its best design performance.

The shape of the uptime frequency distribution curve contains information about the business process design. The clustering of results along a negative downward curve tells us there are destructive influences at work within this company. This company has an in-built outage-causing process—surely unintentionally introduced—that is producing the poor uptime results. The early failure peak is an indicator of poor business process quality control—there are a lot of defects sitting in the business waiting for the chance to go wrong. The negative slope at the back of the curve means that the company’s current designed and intended processes can never get it to the production performance it wants, which is close to 60 days or more of uninterrupted production uptime between outages (the dotted curve in Figure 20.8). That’s happened only five times in 10 years. Those five successes were all attributable to luck. We know it was lucky because the long duration uptime points form a flat line, meaning that they were accidental results without consistency. Those 50-day or longer uptimes were chance events; they are not repeatable by intention because of the self-destroying influences within the company processes. The only sure thing in this business is that uptime has a great chance of being fewer than 20 days and some chance of being up to 30 days; the company will be very lucky if it gets the 60 days it wants (about a 10% chance). Their maintenance costs must be higher and their availability lower because of the many unwanted outages that don't need to happen.

The engineering, maintenance, and operational processes this business uses can only deliver the current uptime results seen in the plots. Remember, this is 10 years of real production data—this performance is what results. Until the company adopts reliability creation processes that guarantee the success it wants, future operating performance and losses will be the same as in the past.

Plant run charts are more than simply indicators of the dates when you had problems in your plant or with your equipment. They contain knowledge of the likely behaviour of your operation in the future. That behaviour is the result of the cumulative effects of your capital projects, asset management, operational, and maintenance management policies and practices. Once you convert run chart data into a frequency distribution, you see the natural behaviour of your operation and its processes.

To do useful process improvement, it is vital to know the reason for each PI and KPI data point. A run chart of each type of event cause is then plotted, and the cause frequency distribution is graphed. Using the distribution curve of causes lets you identify the frequency of each event cause. If you then plot the cumulative TDAF costs from each event, you will understand the high-cost problems you have in your company. Repeating causes are likely to be indications of common cause problems in your operation. Single events may indicate a special cause impacting the process. If the same cause reoccurs, the business-wide losses from it will justify starting an improvement project to solve the problem. Once the problem is removed, the money once lost will turn into new operating profits. Keep full and complete records of your production and maintenance problems—they are worth solid gold to you in the future.

With distribution curves, you have a means to monitor the effectiveness of changes made to a process. Each change you introduce will produce a new performance result for the steps you alter and cumulatively for the process. Once you collect about 10 to 12 results, superimpose the curve for the new outcomes on the past distribution curve to get feedback on the impact of the change. If the new curve does not plot within the higher performance part of the old curve, you need to keep looking for better answers.

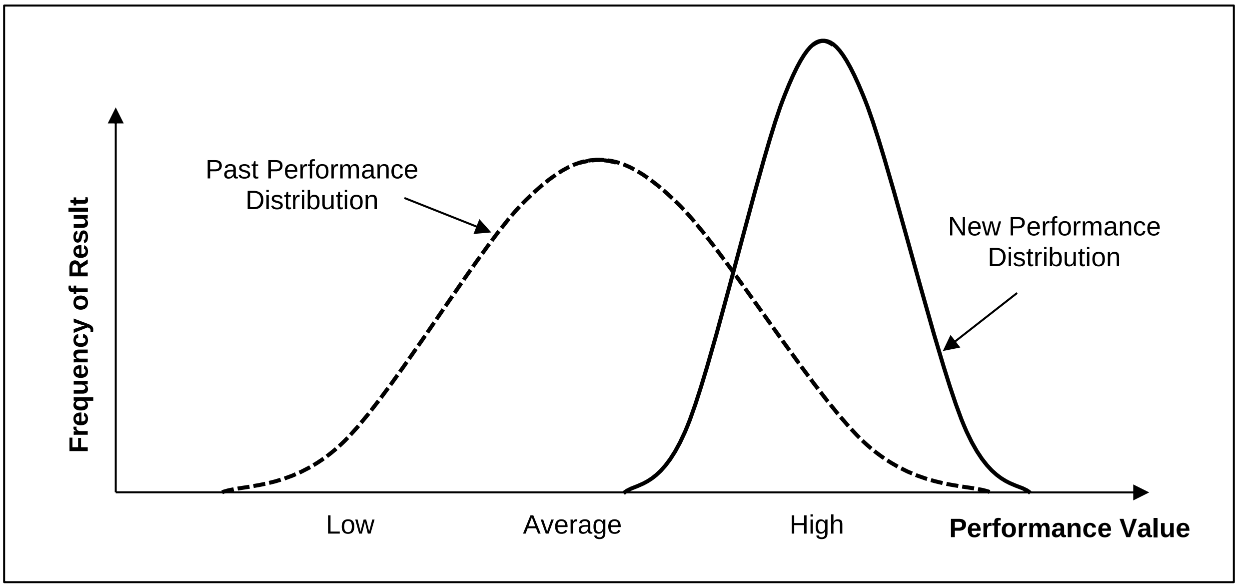

Figure 20.9—New Distribution Superimposed on the Past Distribution

Monitoring and Measuring Maintenance

The types of performance measures useful to the maintenance effort in an operation are those that do the following:

- Identify where your equipment failures start in your suite of business processes
- Track the influence of life-cycle factors on failure
- Direct efforts so that time and resources are used effectively and efficiently to remove the causes of failure, create higher reliability, and reduce operating risk
- Guide maintenance to deliver greater business benefits

Measure What Maintenance Is Doing with Its Time and Resources

Equipment fails because parts lose structural integrity from microstructure overstress or atomic attack. There are six major causes of mechanical equipment failure: lubricant degradation and contamination, out of balance, misalignment, working component distortion, incorrect fastening, and induced vibration. For electrical equipment, the six major failure causes are contamination, induced vibration, high temperature, moisture, distortion, and poor power supply stability. These failure mechanisms must be purged from your plant. This requirement goes far beyond just using condition monitoring to observe equipment health. You want maintenance to find what is creating those mechanisms in the mechanical and electrical equipment and get rid of them. You need measures to prove that maintenance is focused on successfully eliminating problems. You want the maintenance crew to be proactively creating sure reliability. To make this happen in your operation, you need to use maintenance indicators that drive reliability improvement.

It’s important to know where and why maintenance allocates time, people, and effort each month. Unfortunately, maintenance resources are easily expended doing anything and everything to keep the operation running. The best maintenance strategy delivers reliability and lower risks that give back fortunes year after year. The secret is not to focus on doing maintenance but to focus on creating reliability and removing operating risk. Maintenance has the duty to stop problems from starting, and where there are problems, they permanently remove them. If you measure maintenance using the factors listed below and a significant portion of time and resources is not used on them, then the future will be the same as it is today and highly likely to get worse.

- Maintenance work orders spent on improving equipment
- Maintenance time and effort spent removing breakdown causes
- Maintenance time and effort spent improving maintenance procedures
- Maintenance time and effort spent improving maintainer skills and knowledge
- Maintenance time and effort spent reducing operating problems
- Time spent removing wasted effort and cost from maintenance processes
- Efforts spent improving stores management processes and stored parts reliability
- Maintenance work orders spent improving safety

The size and mix of your maintenance crew reflects the effectiveness of its work processes and of the equipment reliability creation processes in use. In a world-class, reliability-driven business, maintainers spend a lot of time designing and implementing productivity improvements on equipment and removing maintenance costs.[3,4] The Plant Wellness Way solution is to use a small, highly versatile maintenance crew with the skills to do high-quality, precision work that is right the first time. When maintainers have more engineering knowledge, use precision skills, create standardized work management processes, and employ work quality assurance solutions, the maintenance crew can be used in very profitable ways.

FOOTNOTES

1. W. Edwards Deming, *Out of the Crisis* (Cambridge, MA: MIT Press, 2000).

2. W. Edwards Deming, *Out of the Crisis*, Pg. 324.

3. Jim Wardhaugh, extract from 2004 Singapore IQPC Reliability and Maintenance Congress presentation "Maintenance—The Best Practices."

4. "Maximising Operational Efficiency Presentation," E. I. Du Pont de Nemours and Co., 2004.