Chapter 3: Variability in Outcomes

Probability, likelihood, chance; the more we learn about them, the more we realize how much they impact our lives, our businesses, and our machines.[1] All around us, things happen. People make choices and act. We only see the effects of those choices in the future. Often, we can't differentiate one effect from another because past choices interact to cause mysterious and unknowable events. Operators, maintainers, manufacturers, engineers, managers, purchasing officers, suppliers, and many others make choices all the time that impact the lives and reliability of plant and equipment.

All unknowns and vagaries introduce variability: the cause of most of our operating and business problems. Variability is "the range of possible outcomes." A business with an aim of providing a product or service with consistent specifications does not want its processes behaving randomly and producing off-spec results. Out-of-specification products are a waste of money, time, and effort. Large amounts of modern organizations' resources are devoted to controlling variability within their business and operating processes. The people involved in this duty carry the titles executive, manager, supervisor, superintendent, or the like. Their role is to ensure that outputs are within prescribed limits. Anything outside those limits is urgently controlled. A business process with high variability means that outcomes range from good to mediocre to disastrous. Things are uncontrolled and volatile. This randomness is the exact opposite of what is required. It is much more profitable to get the right result every time.

Observing Variability

A simple tabletop game can help you understand why variability is a problem. It is a great introduction to controlling variability of processes. In Figure 3.1, two lines cross at 90° with a 2 mm (millimeter) diameter circle drawn at their intersection. The next-larger circle is 10 mm diameter, and the outer circle is 50 mm diameter. Sit at a table, and with the nib at the height of a 300 mm ruler, try to drop a pen by hand into the smallest circle. A hit within the 2 mm circle is the "process" outcome that you require. Repeat the targeting and dropping process at least 30 times. After each drop, measure the Cartesian position of the new mark to an accuracy of half a millimeter. Record the horizontal distance from the vertical line (the X distance) and the vertical distance from the horizontal line (the Y distance) in a table like that shown in Table 3.1.

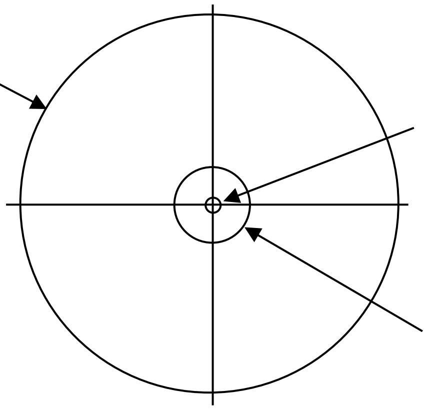

Figure 3.1—The Crosshair Game

Note the average and spread of the X and Y results in Table 3.1. No hits are within the 2 mm circle; some are on the edge, or near, but most are far away. Though great effort was made, the "process," outcomes covered a wide range—there is no repeatability. That is variability.

Hit No.	Distance X	Distance Y	Hit No.	Distance X	Distance Y	Hit No.	Distance X	Distance Y
1	8.5	16	11	1.5	5	21	1.5	5.5
2	7	9	12	1.5	20	22	3	3
3	4	16	13	3.5	3.5	23	3.5	0
4	3.5	2.5	14	2.5	12	24	2.5	6
5	5	24.5	15	3	24.5	25	0.5	2
6	5	16	16	4.5	6	26	1	2
7	7	10.5	17	4	12.5	27	3.5	10.5
8	5.5	9.5	18	5.5	5	28	1	9
9	2	3.5	19	1	9	29	4	14
10	3	2	20	6	4.5	30	0.5	3.5
			Average	X = 3.48	Y = 8.90			
			Spread	0.5–8.5	0–24.5			

Table 3.1—Record of Crosshair Game Hits

If the aim of the game is to have the pen drop inside the 2 mm circle every time, then we have a very poor process for doing that. To get better results, you must change the process. To be successful at the game, you must invent a different process that puts the pen inside the 2 mm circle every time. The results in Table 3.2 were obtained from a process in which the pen was dropped after aiming at the circle from above, much like using a targeting sight to drop a bomb from an airplane.

Hit No.	Distance X	Distance Y	Hit No.	Distance X	Distance Y	Hit No.	Distance X	Distance Y
1	8	10	11	5.5	6	21	3.5	0
2	5	6	12	2	4.5	22	2	5
3	4	3.5	13	0	1	23	0.5	1
4	3	4	14	5	2	24	6.5	0
5	2.5	1	15	4	7	25	3.5	3
6	2	0.5	16	3	1	26	0	8.5
7	13.5	7.5	17	3.5	5	27	6	1.5
8	10.5	9.5	18	4	0	28	0	4
9	1.5	7	19	4	1	29	2	1.5
10	7.5	6.5	20	2	2.2	30	0	6.5
			Average	3.82	3.87			
			Spread	0–10.5	0–10			

Table 3.2—Record of Crosshair Game Hits Using a Sighting Process

The second attempt to play the Crosshair Game using a modified process got better results; the X and Y values are virtually the same. The averages indicate that the hits were closer to the intersection than they were for the first process. There is less spread. But the second process is still not suitable for meeting the requirements. It is unlikely that any process using human hands to drop a pen within a 2 mm circle from a height of 300 mm has sufficiently accurate control. Using human hands cannot meet the required accuracy. You could tell the person dropping the pen to “try harder” or to “improve the quality of their efforts,” but it would be pointless because it is the process that cannot do what is required, not the person. To get the pen consistently within the circle requires a better process that removes the variability caused by the human hand.

Since the game was invented, it has been changed to make the target zone a 10 mm diameter circle, and the number of drops to land on target reduced to 10. The frequency distribution of results after playing the Crosshair Game some 400 times with the bigger target is plotted in Figure 3.2. Many players get one or two drops inside the 10 mm circle. A lot of players get none. Getting three or more drops within the 10 mm diameter circle is difficult. The most successful player got nine drops on target. On his second attempt, he got five drops on target. Some people are gifted in a particular way that the rest of us are not. That is what makes gold-medal-winning Olympians so rare. To win Olympic gold, you must work at improving your ability, but first you must have a capability in your specialty that is well above average—you already must be exceptional to start training to win Olympic gold. But in business, you can't wait for extraordinary people to walk through your door—there just aren't that many of them. In business, you must work with people like the rest of us.

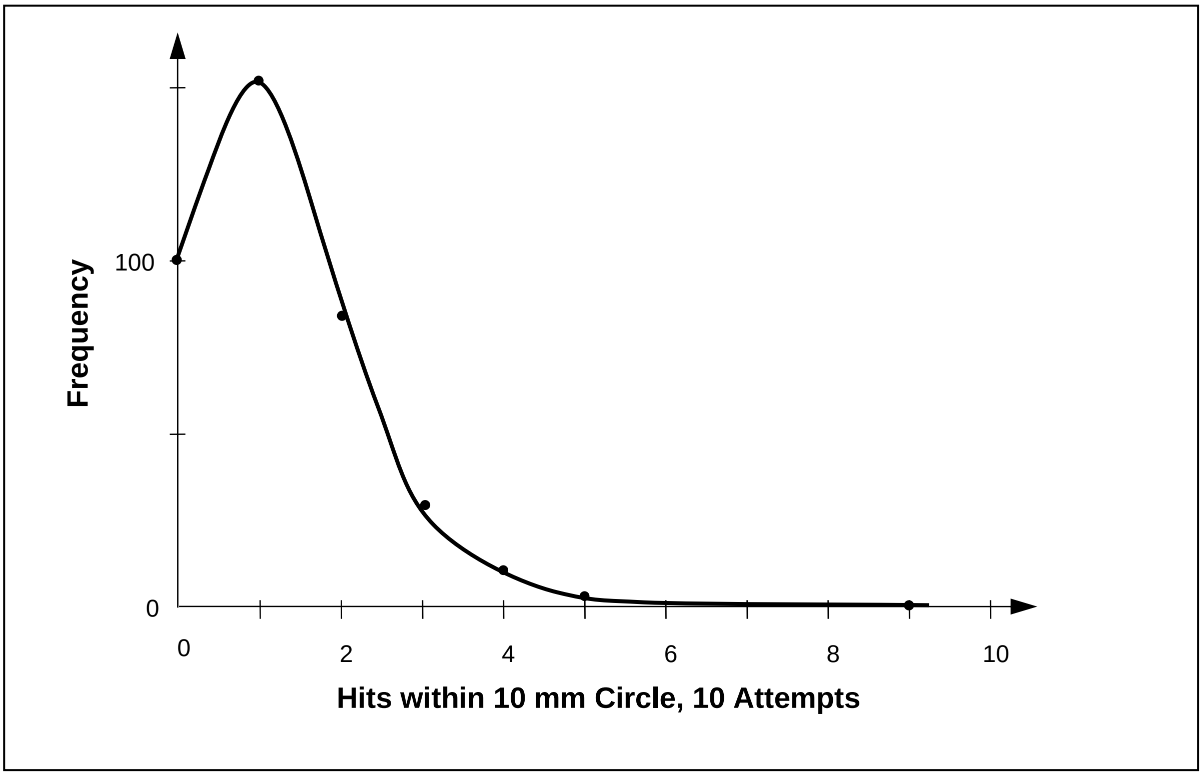

Figure 3.2—Frequency Distribution Plot of Crosshair Game Results

Several Crosshair Game solutions have been proposed by past players. These include using a long, tapered funnel to guide the pen onto the target; using a tube in which the pen slides; using a V-shaped slide to direct the pen into the circle; using a guide rod with the pen fixed in a slider that moves up and down the guide; and using a robot with a steady manipulator to drop the pen. As good as these solutions are, they still involve human interaction in locating guides and maintaining equipment. When people are involved in a process, mistakes will be made at some point. The "human factor" causes variation and inconsistency. But if the solution were mistake-proofed, it would not matter where the pen drops—it would always end up inside the circle.

There is one mistake-proof answer known to the author. It requires using the target in a different way. My thanks and respect go to the tradesman boilermaker who suggested it. Figure 3.3 depicts his solution: make the paper into a funnel with the 2 mm circle at the bottom. No matter where the pen is dropped, it is always on target. Human error has no effect on the outcome. This mistake-proofed solution turns an impossible activity into one that is always done perfectly.

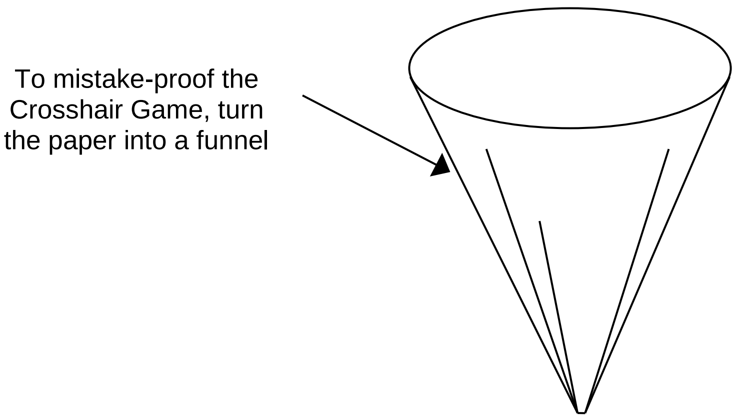

Figure 3.3—Mistake-Proofing the Crosshair Game

An answer jokingly suggested from time to time is to extend the acceptable distance up to 50 mm diameter, and then everything will be on target. This suggestion defeats the purpose of having a process that delivers accurate results. Unfortunately, many businesses unwittingly select it as the solution to their problems. They choose to “widen the target” and accept any result, good, mediocre, or disastrous, rather than set high quality standards and improve their processes to match them. A business that does not pursue excellence in the execution of its activities will not last.[2]

Examples of processes with inherent high variability are those that at some point:

- Require decisions
- Require choices
- Permit human error
- Are done without exacting training
- Have no standards
- Have no or inadequate procedures
- Lack correct information
- Require a guess to be made
- Are ill defined
- Are based on opinion
- Involve emotion
- Can be done in multiple ways
- Are not measured
- Have high rates of equipment failure
- Involve interpretation of data
- Alter settings based on historical results

In these situations, randomness and uncertainty abound. This is particularly the case in sales and marketing, finance, human resources, administration, engineering, design, customer service, production, manufacturing, procurement, dispatch, after-sales service, and maintenance. In other words, it is the case in every process in a business.

The late quality guru W. Edwards Deming advised graphing the process variables and the process outputs across time on a run chart (a time-series plot) to identify uncertainty and variability.[3] When the run charts are used together, they identify the times and causes of poor results. If you want immediate control over a process, then track the process variables—those factors that influence the result—so that they are observable as they change. If the change is bad, you have time to react and correct it before it does too much damage. If you want preemptive control of a process, then trend the variables of the process inputs before they enter the process. By ensuring that the inputs into a process are correct, you can be more certain that the process they feed will behave right.

If you only want to know how well a process performed, then monitor its final outputs after completion. Unfortunately, monitoring the final output puts you in the position of asking "what happened?" when something goes wrong, just like the company in Example 3.1, which had no idea what had changed to cause a spate of raw material stock-outs. By tracing the replenishment process on two run charts, the company exposed process fluctuations, and one phone call later, the underlying causes were revealed.

Example 3.1: Inventory Replenishment Mayhem

The stock replenishment process of a national company involved the ocean shipment of raw material from a manufacturer on one side of the country to the company's factory, located five shipping days away on the other side of the country. For some months prior to the investigation, the company had been running out of stock across a range of shipped materials. As a result, the company was unable to supply products on time to clients because the warehouse replenishment process could not maintain adequate raw material stocks. The company was using up safety stock and not getting resupply quickly enough to meet clients' orders. Annoyed clients reported the problems that this was causing in strongly worded correspondence and angry telephone calls. The company did not know why it was having the stock-outs.

The investigation began by collecting data on raw material stocked-out in the past. Table 3.3 shows the frequency of supplied material that had suffered stock-outs during the prior two years. The company was suffering increased numbers of stock-outs and the table confirmed the seriousness of the situation. The next step was to find what was causing the lack of supply.

Item	Total	Jun	May	Apr	Mar	Feb	Jan	Dec	Nov	Oct	Sep	Aug	Jul	Jun	May	Apr	Mar	Feb	Jan	Dec	Nov	Oct	Sep
T166	21	1	1	2	2	3	1	1				1	1		1		1	1		1	1	1	1
T129	14			2	1	2	1			1	1			1	1							1	
T209	13	1		2			1											1		2	2	1	2
T201	10	1	1	1		1	1		1											1	1		1
T281	10	2	1		2																1	2	
T126	9	1	1	1										1			1			1	1	1	
T169	8	1	2	1				1				1								1	1		
T241	5	1															2			1			
T321	4																					1	
T161	5			2									2	1									
T361	3			1														1					
160N	11	1			1	1	2	2	2	1	1												
120N	9	1	1	1		2											1			3			

Table 3.3—Frequency of Raw Material Stock-Out

It was necessary to look at the history of deliveries from the manufacturer. Historical records of delivery dates are plotted in Figure 3.4, which is a run chart graph of the ship departure dates. It shows a great deal of variability in the deliveries over the most recent months. Lately, deliveries were up to two weeks overdue when they should have been arriving weekly.

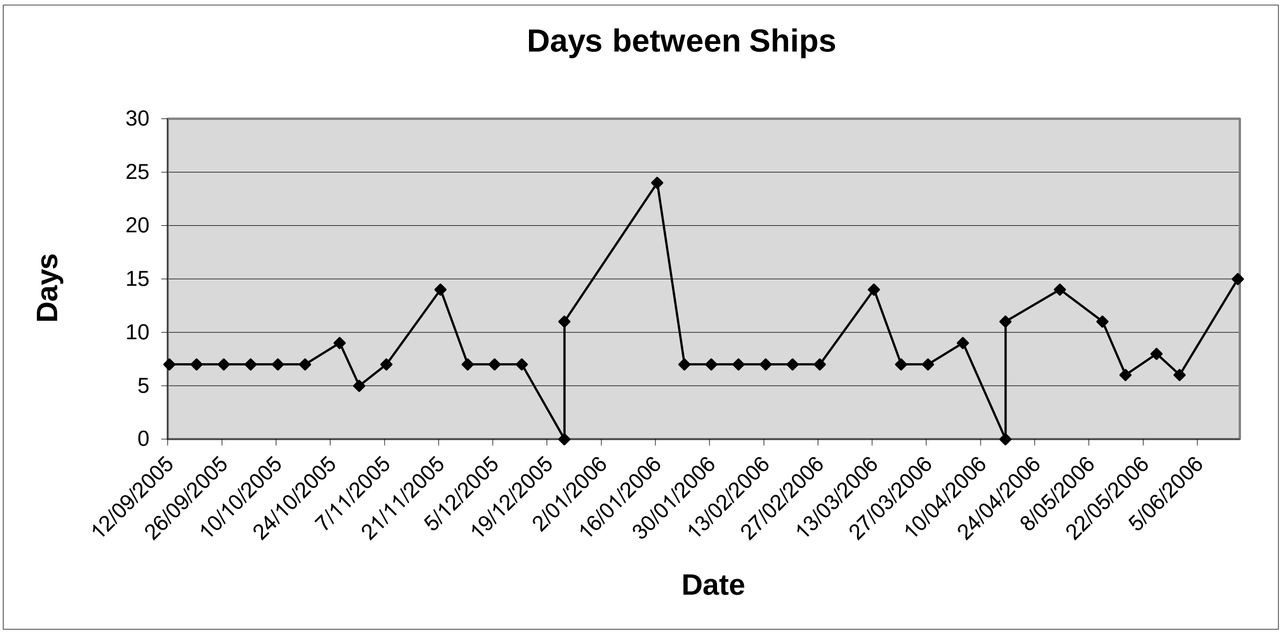

Figure 3.4—Ship Departure Dates

Figure 3.5 is a graph of the number of sea containers in each delivery. It shows variability in the amount of product sent in each shipment. Instead of having normal deliveries of 10 to 11 sea containers, the company was now receiving shipments ranging from 4 to 27 containers.

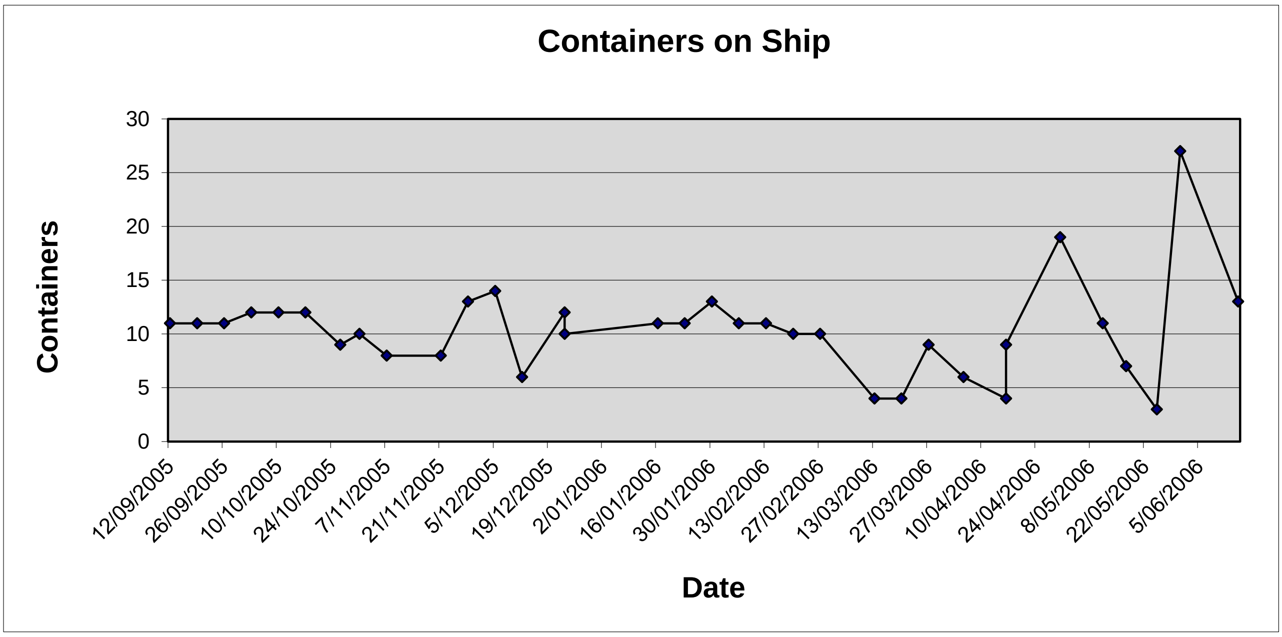

Figure 3.5—Numbers of Containers on Each Ship

Inquiries with the freight forwarder revealed that the regular national shipping line used for deliveries had a two-month maintenance outage for one of its two ships. Whereas once there had been regular weekly shipments, now the only ship left was running fortnightly. To get ordered material to the factory during the maintenance outage, the manufacturer started booking transport with international shipping companies. These ships had irregular departure schedules and only took the number of sea containers needed to fill the empty bays that were left after meeting prior commitments. Sometimes they took a few containers, and other times they took many. The irregular departure of the international carriers with either small or large amounts of raw material resulted in the stock-out mayhem.

The company's customers suffered because of the erratic supply of raw materials to the manufacturer. The irregularity was attributable to the high variability of international ocean shipping, which was further complicated by the feast-or-famine quantities of product on each ship. There were two possible responses to the temporary shipping problems until the return of the regular national carrier's serviced vessel. The company could increase the order size shipped fortnightly, which would effectively raise its inventory levels in transit, or it could book transnational rail delivery. Most importantly, to prevent future stock-outs, the manufacturer needed to change its replenishment process to regularly monitor the ocean shipper's plans and to check for possible delays in shipments.

The disruption of regular delivery to the company in the inventory replenishment example was the result of a "special cause" event—the ship repairs. A special cause event is an extraordinary occurrence in a process that is not attributable to the process itself. Had there been no ship repairs, the weekly deliveries most likely would have been normal. The ship repair was outside the control of the replenishment process, but it had a negative impact on it.

Fluctuation that is attributable to the natural variability of a process is called "common cause" variation. The Crosshair Game is an example of the effects of common cause variation. Where the pen lands will depend on the behavior of the process variables affecting the drop, such as steadiness of the hand, accuracy over the target, evenness of the release, and so on. A ±25 mm spread of hit locations is normal for the Crosshair Game. To have a pen fall into a 2 mm circle, or even land within a 10 mm circle, using a process with ±25 mm variation is mostly the result of luck. Dropping a pen by human hand from a height of 300 mm and having it fall inside a 10 mm circle every time is impossible—the common cause variability of that process is too great for the accuracy required. For the pen to fall inside the circle every time, a process without the element of luck in needed, not an increase in the knowledge and skills of the person doing the job. You could spend a fortune teaching people how to drop pens into circles, and still, everything would be scrap.

An example of a classic misunderstanding of variability that causes equipment to break down is the tightening of fasteners. This confusion is the root cause of many flange leaks, fastener looseness, and machine vibration problems. Figure 3.6 shows the variation in the typical methods used to tighten fasteners.[4] The method that produces the greatest variation, ranging ±35%, is "Feel—Operator Judgment," in which muscle tension is used to gauge fastener tightness. Even using a torque wrench has a variation of ±25% unless special practices are followed that reduce it to ±15%. Industrial experience is that fastener tensioning problems disappear when the final tension is within ±10% of the correct value for the application. There are only three tensioning technologies that, by design, can stay within that accuracy: load-indicating washers, measuring fastener elongation (for example, with hydraulic tensioning, in which a set pressure corresponds to a known shank extension, or ultrasonic measurement of the change in shank length), and mounting strain gauges to measure microstructure elongation. Other torquing techniques have too much natural variation.

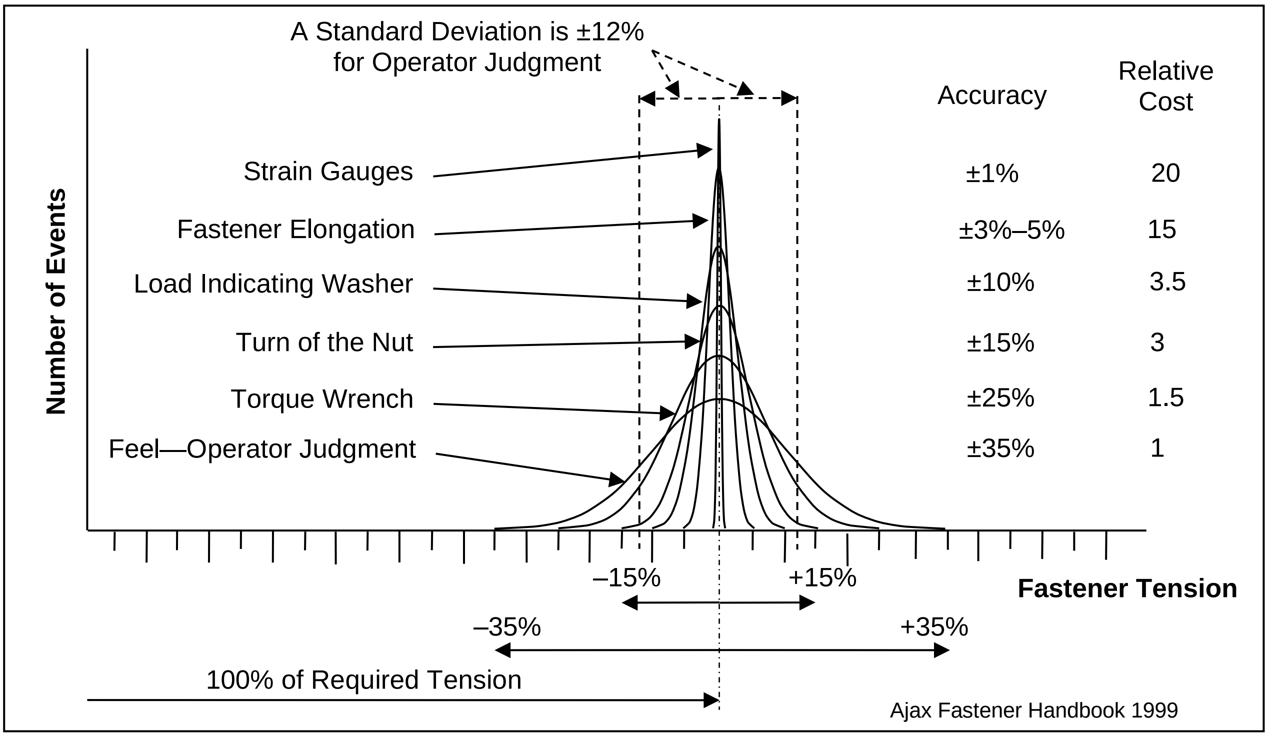

Figure 3.6—Variability in Methods of Providing Correct Torque for Fasteners

It is impossible to guarantee tension accuracy when tightening fasteners by hand. Using a process that ranges ±35% to get within ±10% of a required value is like playing the Crosshair Game—it requires much luck. Joint and connection failure are inherent in a hand tensioning process. Companies that approve the tensioning of fasteners by operator judgment must also accept that there will be many cases of loose or broken fasteners. It cannot be otherwise because processes that use hand torquing to cause a required growth of fastener length have guesswork built into them. It would be a foolish manager or engineer who demanded that his or her people stop fastened joint failures but only allowed them to use operator feel or tension wrenches to control the accuracy of their work. Such a manager or engineer would come to believe that he or she has poorly skilled and error-prone people employees, when, it is the method that the manager specified and approved that is causing the failures. In this case, the manager totally misunderstands that it is the process itself that is not accurate enough to ensure correct fastener tension, not the people in the process.

Stopping fasteners from failing necessitates a method that surely delivers a required fastener shank extension. The fastening process must be changed so that it guarantees the necessary shank stretch. Only after that management decision is made and followed through by purchasing the necessary technology, quality controlling the new method to limit variation, and training the workforce until competent in the correct practice can the intended outcome always be expected.

There is one other phenomenon about variation that you need to grasp. Notice that one standard deviation about the target tension value is ±12% for "Feel—Operator Judgment." It is very close to the ±10% variation about the required tension after which fastener problems disappear. One standard deviation to either side of the mean on a normal distribution bell curve captures 68% of all results. This implies that about two-thirds of all tensioning done using "Feel—Operator Judgment" will not be problematic—it will be close to or within ±10% of the required tension. Because nearly two-thirds of all hand fastening done is "good enough," there will be many workforces in which people on the crews have hardly ever had a fastener failure. It seems that these people have got hand tensioning down to a fine art. That might be true for some of them, but for many people, they have simply been lucky. The situation is much like golf—professional players are highly tuned to make par or better, and they consistently get those results. When nonprofessionals make par, it is luck, not their golfing expertise, that produces success.

Any operation using peoples' muscles to control fastener tension has failure built into the design—having some loose and broken fasteners is the nature of the process. It is why Deming gave his famous warning to managers: "Your system is perfectly designed to give you the results that you get." Poor equipment reliability is the result of using business and engineering processes with inherently wide variation. These processes are statistically incapable of getting the required performance with certainty, and so equipment failure from their use must be expected. Occasional failure is a design outcome of these processes, and luck helps keep their users in business.

Another process that designs failure into equipment is the common maintenance practice of changing oil after it is dirty. When managers decide to replace lubricant only after it is chemically degraded or black with solid contaminants, they unwittingly agree to let their equipment fail. Table 3.4 lists some ISO 4406 solids-in-oil contamination ranges numbers.[5] Each value has twice the count of particles in a milliliter of lubricant as the previous range (a volume equal to about 20 drops of distilled water). The range number 21 (dirty lubricant) contains 128 times the number of particles in each milliliter than 14 (clean lubricant). The size and number of solids in lubricant directly impact the likelihood of roller bearing failure.[6]

Range Number	Number of Particles per Millilitre		Increase in Particle Count from 10 Range	Visual Colour
25	160,000	320,000	32,000	
24	80,000	160,000	16,000	Dark
23	40,000	80,000	8,000	
22	20,000	40,000	4,000	
21	10,000	20,000	2,000	Dirty
20	5,000	10,000	1,000	
19	2,500	5,000	500	
18	1,300	2,500	250	From new drum
17	640	1,300	130	
16	320	640	64	Clear
15	160	320	32	
14	80	160	16	Clean
13	40	80	8	
12	20	40	4	
11	10	20	2	
10	5	10		

Table 3.4—ISO 4406 Particle Count for Lubricant

Depending on the lubricant regime (hydrodynamic, elasto-hydrodynamic), viscosity, shaft speed, and contact pressure, roller bearing elements are separated from their raceways in the load zone by a lubricant thickness of 0.025 μm (micron) to 5 μm.[7] Eighty percent of lubricant contamination is composed of particles less than 5 μm in size.[8] This means that in the location of highest stress—the load zone—tiny solid particles can be jammed against the surfaces of the roller and the race. The bottom diagram in Figure 3.7 shows particle contamination in the load zone of a bearing. A solid particle carried in the lubricant film is squashed between the outer raceway and a rolling element. Like a punch shearing a hole through sheet steel, the contaminant particle causes a high load concentration in the small contact areas on the race and roller. Depending on the size of the stress that develops, the surfaces may or may not be damaged by the particle. Low and average stresses are accommodated by the plastic deformation of the material of construction. However, an exceptionally high stress will punch into the microstructure, creating submicroscopic cracks in the surface and subsurface.[9] Once a crack is created, it becomes a stress raiser and grows under much lower stress levels than what is needed to initiate the crack.

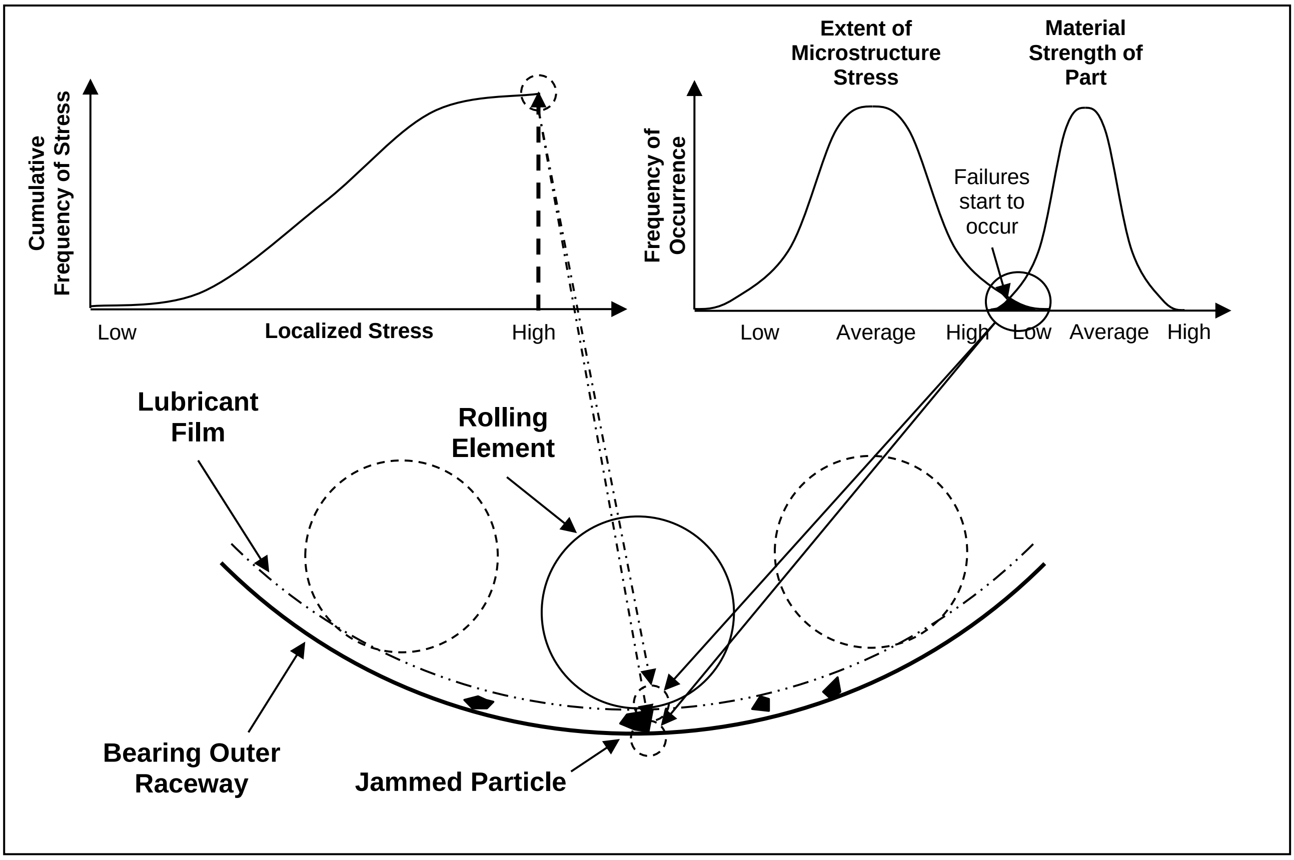

Figure 3.7—Solid Contaminant between Roller and Raceway Overloads the Microstructure]

High surface stresses also result when cumulative forces of loads, each of which is individually below the threshold that damages the atomic structure, act together. Such circumstances arise when a light load supported on a jammed particle combines with additional loads from other stress-raising incidents. Like impact loads from misaligned shafts, tightened clearances from overheated bearings, forces from out-of-balance masses, and sudden operator-induced overload. All these stress events are random. They might happen at the same time and place that a contaminant particle is jammed into the surface of a raceway, or they may not. Whether they combine to produce a sufficiently high stress to create new cracks, or they combine on already damaged locations where lesser loads will continue the damage, is a matter of probability.

When a roller bearing is in use there is relative motion between the raceways and the rotating rolling elements. The odds that a damaged area on a roller will be repeatedly stressed is low because the roller moves to a different spot. However, a damaged area on a race remains exposed to all the rolling elements that pass over it. The passage of each roller is an opportunity for solid contaminants to be squashed against the race surface. The size and frequency of stress caused to a bearing from solid particulate is random. You could have very clean lubricant, and although the odds are extremely low, you may be unlucky enough to jam the only solid particle in the neighborhood between roller and raceway while a rotating misalignment force spike passes through it. But as contamination by solid particles rises more solids are in the lubricant. With each rolling element that turns over a raceway the growing number of particles provides ever-increasing opportunity for one to be punched into the surface. As lubricant gets more contaminated, the odds of damaging bearings, blocking oil flow paths, or jamming sliding surfaces climbs.

To significantly reduce bearing failures, gear failures, and sticking hydraulic valve problems, the ISO 4406 particle count must be kept at clear levels, or below, so that the oil has hardly any solid contaminants in it. Companies mistakenly allow their gearboxes, drives, bearing housings, and hydraulic system oils to get dirty and blacken from wear particles before they replace them. Often, they wait for an oil analysis to indicate that contamination is too high, or they replace dirty oil on time-based preventive maintenance schedules. Unfortunately, by the time the lubricant becomes dark from particle contamination, the probability of jamming a particle between two contact surfaces has increased markedly, and failure sites have likely already been initiated. The risk of failure to a company's plant and equipment from wear particle oil contamination is the direct result of the management processes applied (or not applied) to decide how much solid contamination will be sanctioned in the oil. When oil is changed after it is already darkened by particulates, it is far too late to greatly reduce the probability of failure. Greases and oils must never be blackened by particle contamination in the first place if you want to stop the chance of having lubricated and hydraulic equipment breakdowns.

The failure of machinery is directly related to the corporate maintenance strategy and the processes chosen to maintain and operate the plant. You could be destroying your equipment and think you have great operating and maintenance strategies! The managers and engineers in these companies are fervent that they do the right maintenance practices and have excellent preventive maintenance processes in place. They are wrong, of course, because the processes they use cannot guarantee the results they want. Many organizations try to achieve results using business, engineering, and operating processes with common cause variation that cannot reliably produce the performance they need—they are playing the Crosshair Game in what they do. Such businesses employ processes containing inherent volatility that naturally produce outcomes outside requirements. Trying to manage an organization with systems and processes that produce highly variable results is an exercise in futility that causes waste, distress for all involved, and emotional burnout for managers, engineers, and supervisors.

Controlling Process Variation

Controlling common cause problems require design improvements in the way a process operates. In contrast, special cause variability is controlled by stopping the influence of the extraordinary event. Preventing the effect of the supply ship's repair that caused late raw material deliveries in Example 3.1, "Inventory Replenishment Mayhem," is done by using other reliable modes of transport to replace the unavailable ship. As soon as on-time delivery by ship is not possible, rail freight needs to be booked. You address special cause issues by stopping them from happening or by preventing them from impacting your business. But common cause issues are inherent in the process, and their elimination requires changing the process.

It is the nature of every process to produce variation. The challenge for business and operations processes is twofold. First, it is to have only "natural" variation and no special cause variation. Second, it is to select or develop processes with natural variation well within the required performance. This allows the organization to focus mainly on stopping special cause problems, sure in the knowledge that the process is inherently stable and produces good results. When a business or operating process no longer performs within its normal limits, look first for a special cause of the change. Only after all special causes are eliminated can you be sure that only natural common cause variation remains. If the common cause variations are still too volatile, you have justification for improving or changing the process. By following that sequence, you confirm whether any special cause variations are masking the natural process variability with effects that confuse the analysis. If a special cause is mistaken for a common cause, you will make the wrong decisions to address the problem.

So far, we have seen examples of variability in a game, in the supply chain of an organization, and in maintenance strategies used on plant and equipment. Being able to get a picture of variability with run charts and tables brings a clearer appreciation of what is happening within the process. It allows us to ask relevant questions that lead to a more profound understanding of a situation's causes and resolutions. Great value is gained when an organization observes the irregularity of its business processes. Once a picture of process behavior is available, you can make focused efforts to control unacceptable variability. The next case study is of an ore-processing plant where the consensus was to invest $250 million to expand production by 50%, when in fact that investment may have been unnecessary if production variability had first been addressed.

Example 3.2: The Hidden Factory

This is an example of the value of identifying the causes of variability in a business and removing them before spending new capital. The production from an ore-processing plant is trended in the simple bar graph of Figure 3.8. It shows the hourly production rates in a milling operation running 24 hours a day, seven days a week during eight consecutive weeks. It contains a lot of valuable information about the operation's capacity as well as a clear indication that the business is suffering from wild fluctuations in its production throughput. Examination of the graph provides insight into the facility's dilemmas.

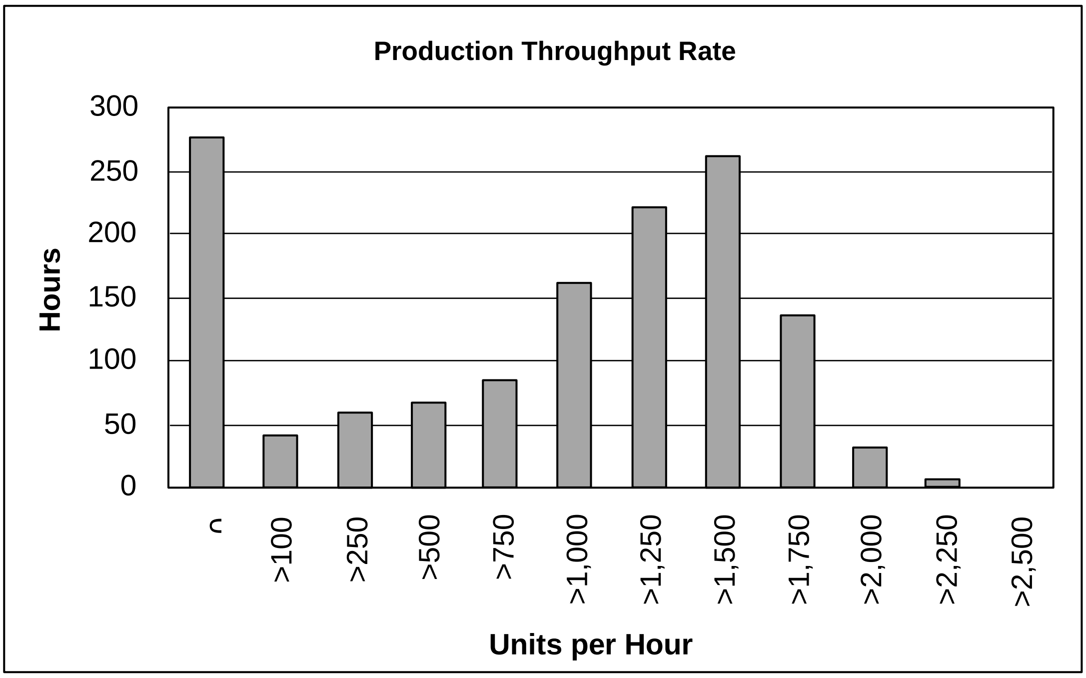

Figure 3.8—Production Rates Chart

The eight weeks of production shown on the graph represent 1,344 production hours. For 275 hours, there was no production, so for 20% of possible production time, the plant was standing still. The plant design capacity is 1,500 units per hour. For 615 of the remaining hours, 57% of the time it was running, it delivered substantially less than the designed production rate. The actual average throughput for the eight weeks was 1,000 units per hour, which is two-thirds of design duty. This facility is suffering from severe production problems and needs to investigate why it is not producing consistently at design capacity.

There is additional information in the graph. For a significant number of hours, the plant ran above its design rate. There are two possibilities here. One is that in trying to make up for lost production, the plant was overloaded, which then led to even more equipment failures and added downtime. The second is that the plant could be run at more than its design duty. Confirming each possibility requires an engineering design investigation. There is a good chance that with minimal engineering changes, the plant could be run consistently at 2,000 units per hour, which is one-third greater than the current design capacity and twice the current average production. The overstressing of parts would be a major concern at the increased production rate and would need to be addressed by a full design review. An operating risk analysis based on Physics of Failure consequences would be conducted, and problems designed out as part of the decision to increase production to a higher rate than the original design.

There are obvious questions to ask of a plant with this extent of variability in performance. What is causing the stoppages and below-design throughput so often? If the plant can produce at higher rates by accidents of circumstance, then what could be consistently produced if those circumstances were deliberate? It would be sensible to identify the causes of the disastrous production losses and solve them while making the fortuitous events of the past intentional. The total “lost” throughput represented by the work stoppage and slow running, plus the higher production rates available from reengineered capacity, means that this operation has plenty of opportunity to deliver a large production increase without significant capital investment.

This company’s decision to spend $250 million on a major capital upgrade to boost production 50% may not have been necessary. By recovering the downtimes and low production rates and reengineering bottlenecks for higher throughput, the extra capacity was probably achievable with the old plant. It was only necessary to conduct root cause investigations on why the production losses were occurring and how to solve them. The financial return on such an investment would be unbelievable. All these options become clear simply by measuring production variability.

Constructing a graph like that in Figure 3.8 requires collecting the hourly production figures for a sufficiently long time to observe the full range of variability affecting the process. The figures show a range of performance around a mode value (the quantity in a data set that occurs most often). The extent of the spread below the mode indicates whether there are production problems hampering throughput. The range of spread above the mode indicates whether there is spare capacity available. If the spread is tight about the modal production rate, then the throughput is stable, though not necessarily optimal. But if the spread below the mode is wide, as in Figure 3.8, then the plant has "hidden" efficiency opportunities to improve its production performance.

When production throughput graphs have a wide spread of production rates, there is potential to increase plant capacity by removing the causes of operating losses with minor engineering upgrades or removing the variability by adopting improved procedures and useful training. Before you invest more capital to expand plant capacity, investigate the variability of current production, as there may already be a "hidden factory" within your plant.

Controlling Business Process Performance

When any process is run, it produces a range of outputs that are its characteristic signature—you get what the process does. Business process variability produces a frequency distribution of results. The repetitiveness of outcomes can form a normal bell distribution, as in Figure 3.9, although processes often have non-normal distributions with skewed shapes.[10]

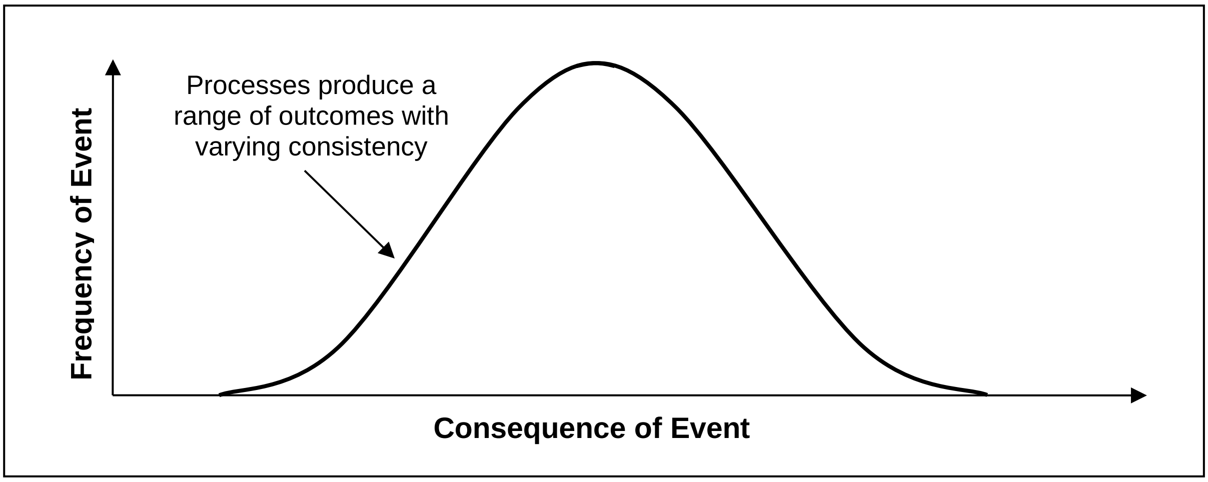

Figure 3.9—Processes Produce Their Own Results

The way to tackle variability is to put a limit on the acceptable range of variation and then build, or change, business processes to ensure only those outcomes can occur. This approach drives process innovation until the goal is reached. Figure 3.10 shows a minimum specification of performance for a process producing wide variation. The acceptable range is categorized by zones of accuracy: "good" being tolerable, "better" being a finer performance, and "best" being world-class results. Only outcomes that are within the "good" standard are acceptable. Any lesser results are defects and rejects. Changing to a process that produces a performance curve that is always inside the good range requires designing and installing better methods that remove performance fluctuations, reduce volatility, and stabilize process variation inside the tolerable range.

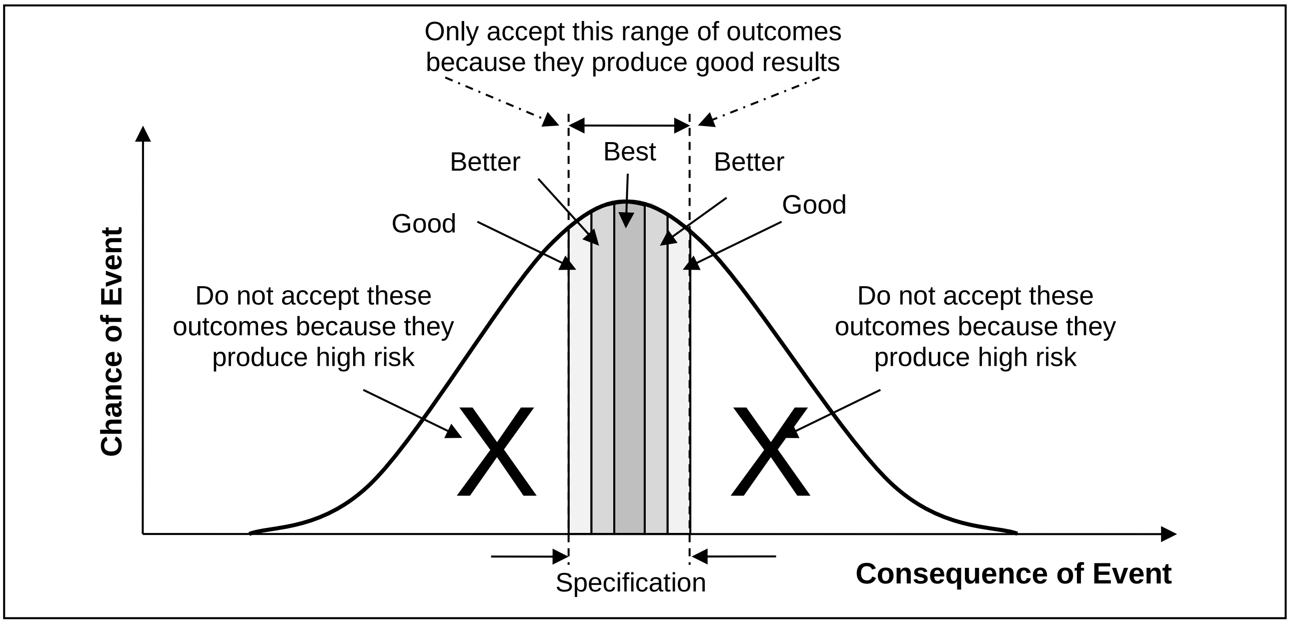

Figure 3.10—Controlling the Chance of a Process Event

With volatility controlled, the results tighten around a consistent mean, as shown in Figure 3.11. Variation still exists, but it is within the desired limits. A process that always produces stable, repeatable outcomes within its specification limits is in control and capable. It becomes highly predictable, and the results can be guaranteed. Business process performance is mostly in our control. You improve your processes by choosing policies and practices that minimize the chance of bad events happening and maximize the chance of good outcomes occurring.

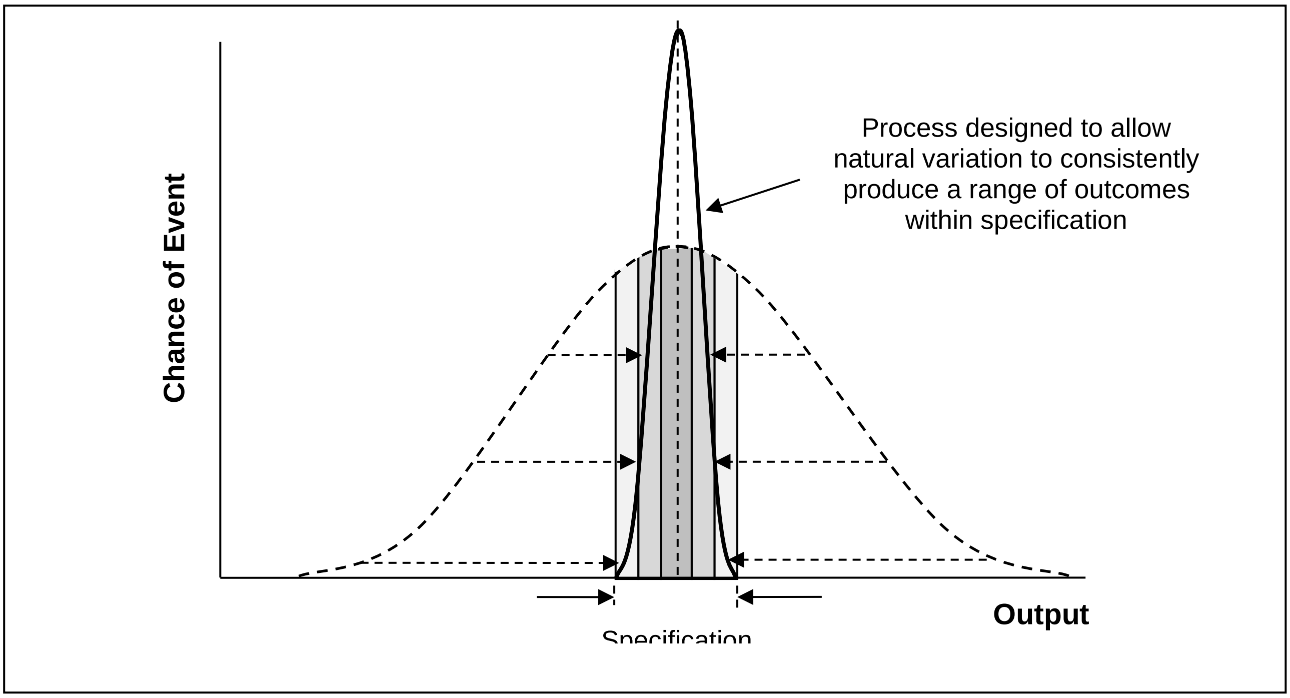

Figure 3.11—The Effect of Removing Volatility from Business Processes

The purpose of controlling variability is to provide certainty of performance. Once variability is identified, it becomes necessary to make the decision to leave the situation alone and accept the range of outcomes that happen, or to address the underlying problems causing the fluctuations. When you make process improvements, it requires finding the causes of the problems and then identifying ways to design them out of the process. Example 3.3 tells of a company struggling with constant maintenance problems despite doing the very best it knows how do.

<u>Example 3.3: The Factory That Made Its Machines Break Down</u>

Most industrial businesses make their equipment fail. Misunderstanding risk, variation and probability leads managers and engineers to use processes that cause trouble. An analysis of a real business illustrates the effects of this all-too-common management error. Figure 3.12 is a run chart of a company's total breakdown hours per week over 16 weeks. If the graph is representative of normal operation, the time series can be taken as a sample of the company's typical business performance. Important information about the company's operations is exposed by using two basic statistical analyses on the data: a frequency distribution, and a control chart.

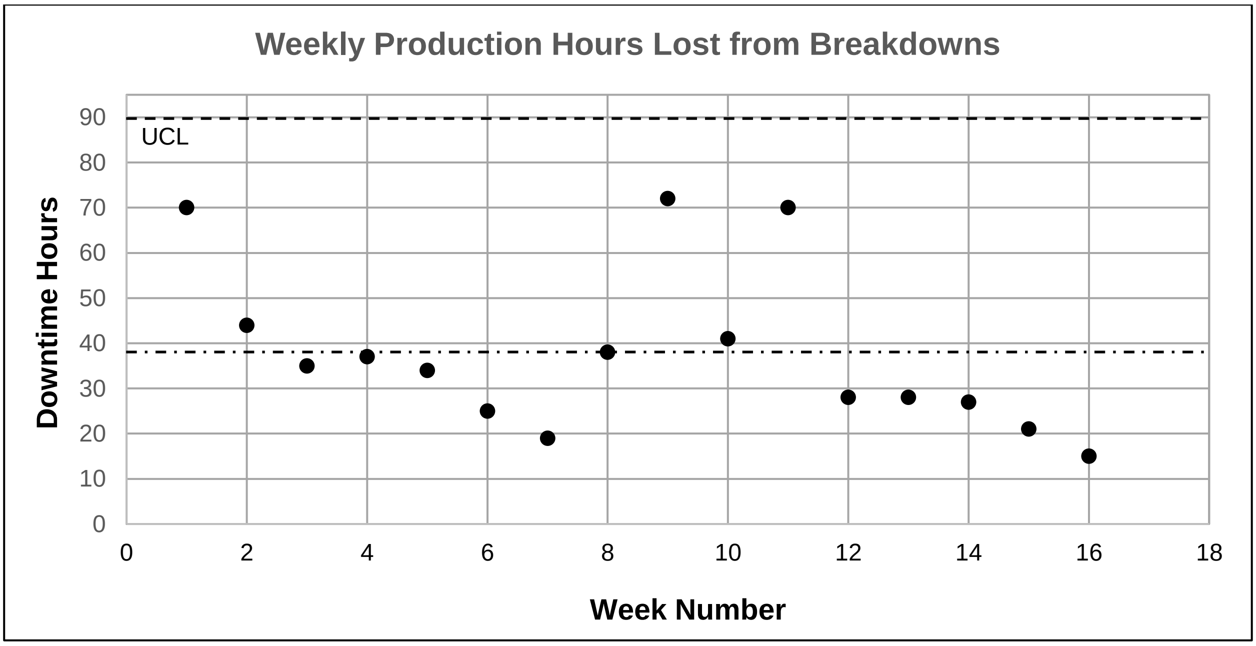

Figure 3.12—Breakdown Hours per Week

The frequency of total downtime hours each week due to breakdowns is plotted in Figure 3.13. The distribution is twin peaked, with the lower portion looking like a normal bell curve but suddenly there is a discontinuity between it and the higher breakdown hour weeks. This company has two types of breakdowns—the normal, and the catastrophic.

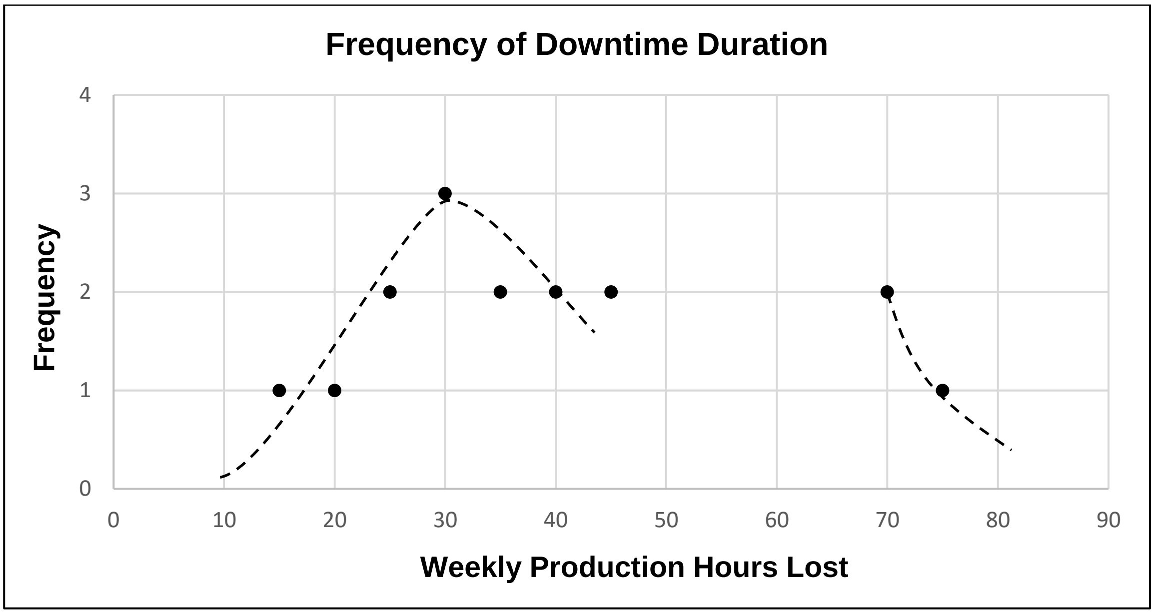

Figure 3.13—Frequency of Weekly Breakdown Hours

For an x-bar control chart, the average breakdown hours per week are 38 hours. The standard deviation is 18 hours. The upper control limit is 92 hours. The lower control limit is 0. These control lines are shown in Figure 3.12. Because all data points are within the statistical boundaries, the analysis indicates that breakdowns are common to the business processes and not caused by outside influences. This company has a statistically stable system for making its equipment break down. Breakdowns are one of the "products" made every week.

Because the breakdown creation process is stable, the future generation of breakdowns is predictable and certain. If this time series is a true sample of normal operation, it can confidently be said that there will always be an average of 38 hours lost to breakdowns every week in this business. In the three weeks following the 16-week period represented, the "normal" weekly breakdown hours were 25, 8, and 25 hours, respectively. This business has built breakdowns into the way it operates because the process of breakdown manufacture is part of the way the company works. The only way to stop the breakdowns is to change to processes that prevent breakdowns.

What Quality Is

In his book, *Out of the Crisis,* Deming advised that "quality must be built in."[11] Quality, Deming tells us, is installed at the source. It is designed in and made part of the product or service; it is delivered by the business process design itself. Quality is a definite and "hard" measure that can be clearly identified. It is quantified with engineering values—the "numbers" that, when achieved, deliver customer satisfaction. In Deming's view, a product or service has the right quality when customers are so satisfied that they boast about it to the people they meet. The quality of the product or service is designed to satisfy the customer. Word of mouth markets it.

The same certainty over quality, but applied to equipment parts, is necessary to deliver the outstanding equipment reliability and plant availability that produces world-class production performance. What is important to know about quality is that it must be measurable. Quality is not left up to people to interpret—it is management's responsibility to define it. It needs to be quantifiable—a length, a thickness, a resultant force or pressure, a color, a smell, a viscosity, a period, a rate of change. You require a specific engineering value, even a collection of values, that defines a level of performance. Once the values are attained, the performance is certain, and the required quality is achieved.

To have quality, you need a target and a range of acceptable outcomes. It is impossible to know how to control quality until standards of allowable variability are set. Once a standard is specified, measures are made to identify whether the processes used to achieve it are statistically capable of meeting the standard. For the business in Example 3.3, "The Factory That Made Its Machines Break Down," the processes will never be able to deliver long periods of breakdown-free operation. The company's current production asset management system does not work. It is not designed to produce a breakdown-free week. In fact, the practices, methods, and processes used cause the failures. It is nearly impossible in this operation to expect more than a couple of days without breakdowns. This company needs to change its fundamental life-cycle asset management processes if it wants to improve equipment reliability. If the company set a target average of, say, 10 breakdown hours a week, a search for better methods and strategies to reach 10 hours breakdown per week would start. The great challenge for this company is to replace years of destructive practices in operations and maintenance with processes and methods that produce high reliability. This change can only occur when the company decides to create business processes that deliver more uptime.

It is necessary to design a new game plan if existing processes do not produce required results. Figure 3.14 represents the strategic aim when changing processes to be capable. Deming said that it is the responsibility of management to improve a process—no one else can do it.

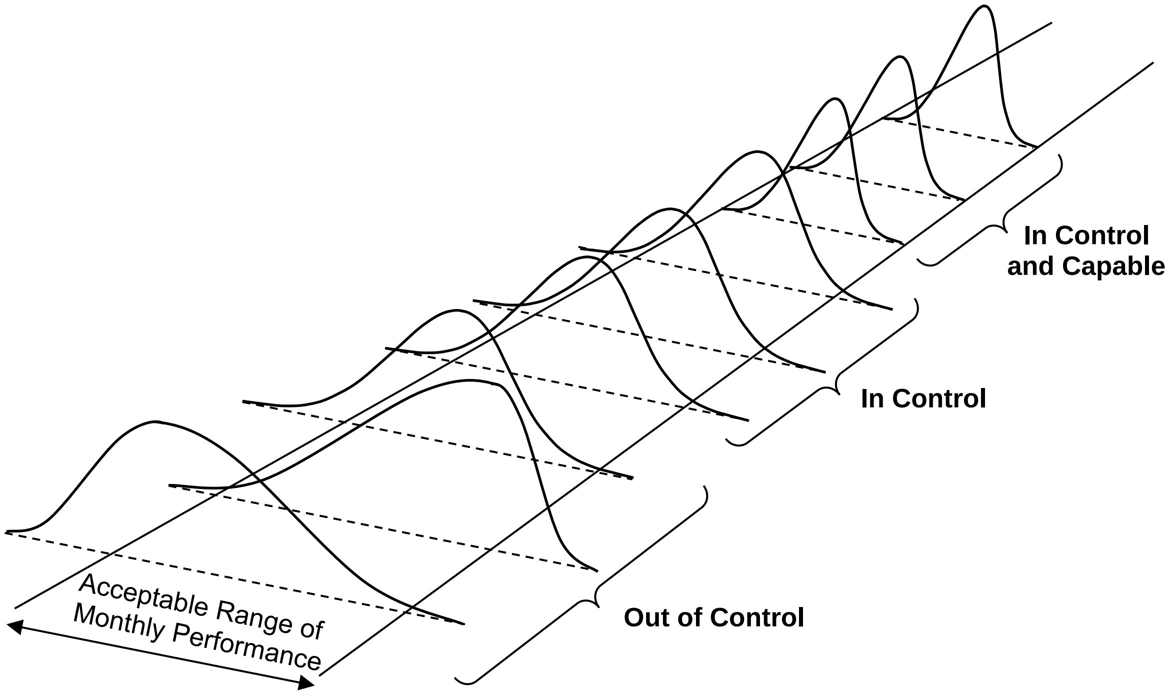

Figure 3.14—Making a Process in Control and Capable

Need for Setting Engineering and Maintenance Quality Control Standards

Reducing the influence of chance and luck on equipment parts starts by deciding what engineering and maintenance quality standards you will specify and achieve in your operation. Probabilistic plant reliability outcomes are controlled by delivering the conditions that produce excellent equipment dependability and performance.

The degree of shaft misalignment you tolerate directly impacts the likelihood of roller bearing failure.[12] If shaft misalignment is present on equipment, that does not mean that a bearing will be failed. Depending on the extent of misalignment, the operational abuse, clearance reduction from high temperatures, out-of-balance forces from unbalanced masses, and myriad other stress-raising possibilities, the size of the resulting stresses may still be lower than the strength of the bearing's material of construction. But it does mean that shaft misalignment increases the chances that its loads will combine with others and add up to produce a catastrophic failure. As more stress-causing opportunities become present in equipment, the chance of part failure grows ever greater.

If shaft misalignment between equipment is so critical to production uptime, then what is the shaft misalignment tolerance in your company? What number of contaminating particles do you permit in your lubricant? The lower the quantity of solids, the surer the likelihood that you will not have a breakdown. The frequency and scale of machine abuse permitted during operation directly affects the likelihood of machinery failure. What do your operating procedures say about running your plant and equipment for outstanding reliability? The rotating equipment balance quality that you achieve directly influences the likelihood of failing roller bearings.[13] The lower the residual out-of-balance forces, the smaller the possibility that out-of-balance loads will combine with other loads to initiate or propagate failures. What balance standard have you set for your rotors? How accurately will you specify fastener tension to prevent fasteners loosening or breaking? The more precisely the extension meets the needs of the working load, the less likely it is that a fastener will come loose or fail from overload. The temperatures at which bearings operate change their internal clearances, which directly influence the likelihood of bearing failure.[14] How well do you control bearing operating temperatures during machinery use so that clearances are always correct? Similar statements about the dependency of failure on the probability of failure-causing incidents can be said of every equipment part.

Chance and luck affect the lifetime reliability of all parts and, consequently, of all your plant items, machines, and rotating equipment. But the chance and luck that affect your equipment parts can be altered by your choices. Change your lubricant cleanliness limit so that far fewer wear particles are present, and you'll greatly reduce the number of contaminant particles in the lubricant film and, thereby, reduce chance and risk of failure. Combine that with ensuring that shafts are precisely aligned at operating temperature, that rotors are highly balanced, that bearing clearances are correctly set and sustained, and that operational abuse is banded and replaced with good operating practices to keep loads below design maximums, and you will greatly improve your "luck" with equipment reliability. You can have any equipment reliability you want by turning luck and chance in your favor through using and meeting ever-finer quality standards in your plant and machines.

Defect-Elimination Strategy

Because variability exists in all processes, a range of outcomes will always result. The Crosshair Game and the operating problems discussed earlier in this chapter highlight some of the bad effects that process variability causes for organizations. When variability becomes excessive, defects occur and failures result. A defect is a "nonconformance to requirements or function." It is a deficiency. It means that bad quality went into service. Defects that escape correction lay hidden and may not become apparent until they cause a failure. A failure is "an event or circumstance which prevents the accomplishment of an intended purpose." A failure occurs any time a thing does not do its job. A failure happens whenever a system or component is unable to perform its designed role. Figure 3.15 is a modified version of the DuPont Chemicals defect and failure model.[15] It highlights some of the many processes by which failure-causing defects and errors come into an operation.

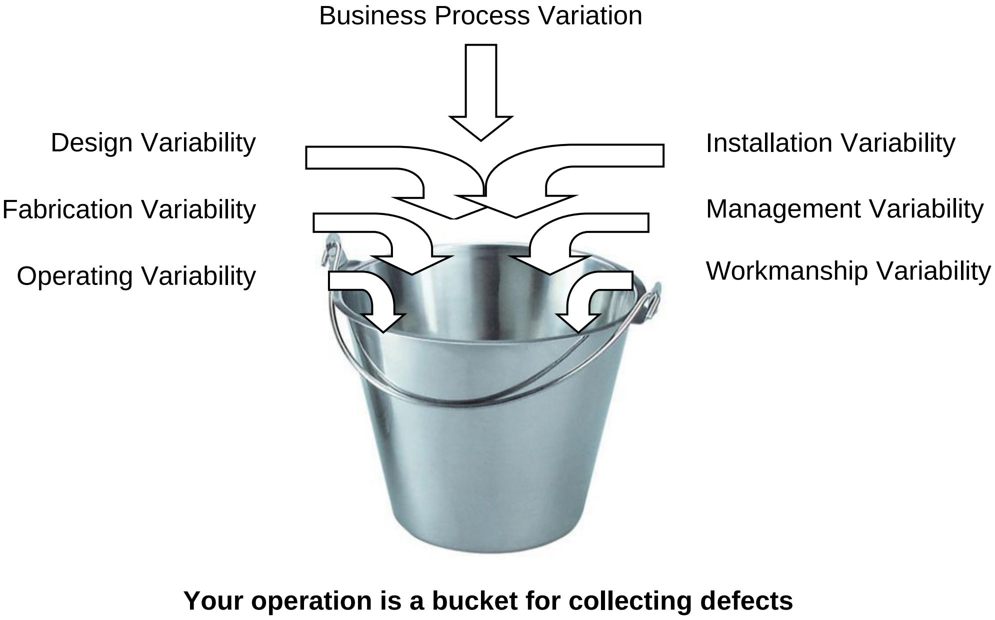

Figure 3.15—Defect Creation

Most businesses react to failure as shown in Figure 3.16. They introduce maintenance and repair systems to manage the presence of defects. They accept failure as normal. Consequently, they suffer production downtime and high maintenance costs by trying to limit the effects of the introduced problems, and fixing failures caused by the problems they could not stop.

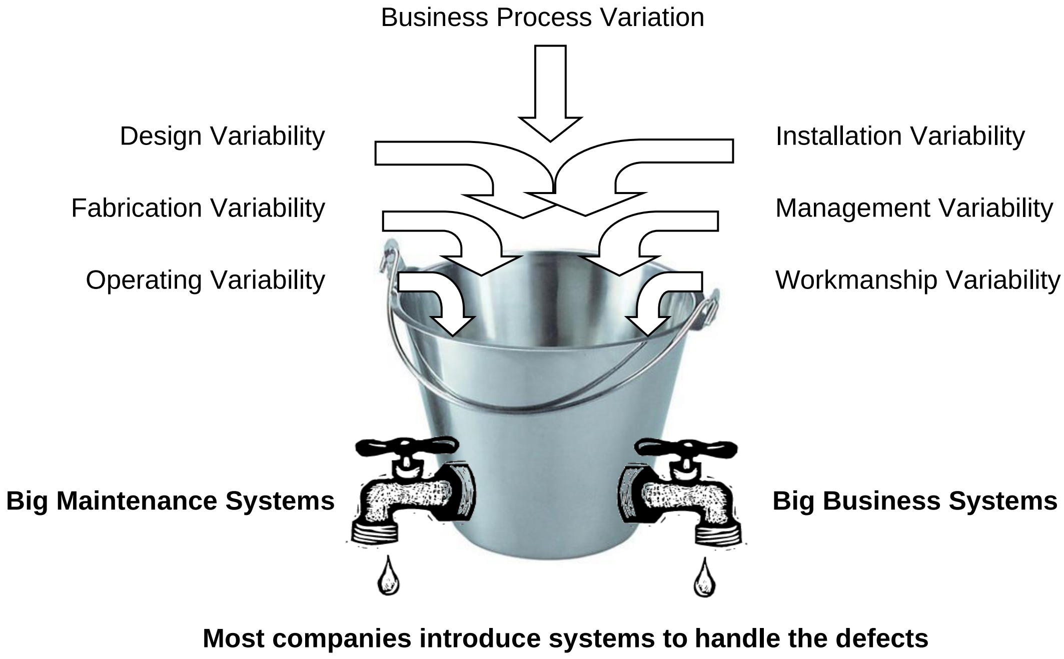

Figure 3.16—Defect Management

Figure 3.17 shows the best strategy: to stop defects from entering your business in the first place. This forces your quality standards to improve, and, as a result, maintenance costs are reduced, and production uptime is lifted.

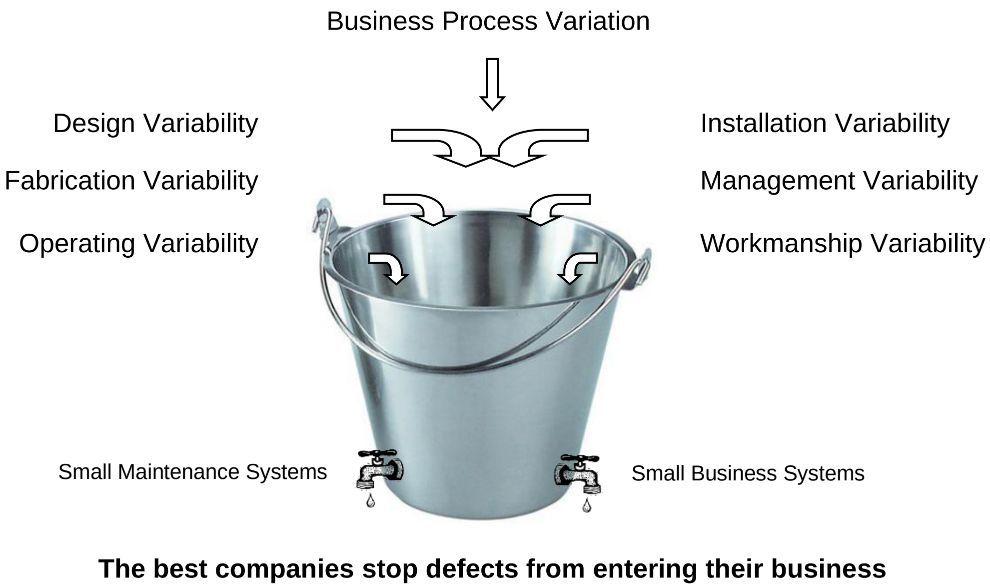

Figure 3.17—Defect Elimination

Because there are fewer failures, your equipment reliability, plant availability, and productivity rise. All the money not spent on failure correction and repairs, plus the extra income from the added throughput in the recovered production time, can be banked as new operating profit.

Variability acts across distance and time. Variations in one process can reduce the effects of variability in an interacting process. Much like an ocean wave rebounding off a cliff, variability between interconnected processes may act to calm the waters, but usually the opposite happens, and variations combine to produce problems of greater magnitude—instead of calm, a surging wave of trouble is created. This was the case in Example 3.1, in which the international shipping line's failure to adhere to a fixed schedule and to provide regular container slots compounded the replenishment problems of its users. Variability that worsens problems requires identification and redesign of the offending processes to remove the negative impacts.

Because every process in a business produces variable results, as more processes are used, there is greater opportunity for defects and failures. Those organizations that try to do many things have many processes to manage and control. Each process step introduces opportunities for variation. A product or service is exposed to risks from the full range of variability in each activity employed during its creation—concept, feasibility, design, procurement, production, assembly, and delivery.

A common supply chain philosophy that compounds operational problems is when company purchasing policy requires the same item to be bought from a range of suppliers in the questionable hope of keeping costs low through competition. When external providers are used, it is necessary to have protection against the worst excesses of their processes and ensure compliance to precise and agreed quality specifications. Companies with multiple suppliers end up suffering more problems and costs than a company using only one supplier. Working with a range of suppliers for the same item requires a great deal of extra effort and time in procurement, accounting, and warehousing. Compounding the risk, each supplier brings its own process variability. When a specified item is bought from many different suppliers, you increase the workload, and the range of variability within your business. This requires corrective measures to be added to your processes to fix the problems caused by slight differences between the goods from each supplier. Suddenly, the small amount of money saved at purchase is dwarfed by that wasted rectifying the troubles. But by staying with one supplier, you have smaller, simpler processes within your organization, you adapt your systems to the supplier's process variability, and the supplier will modify its processes to provide the product quality you want. Those companies that think having supplier competition reduces their costs increase variability problems throughout their business.

Variability of engineering component design and fabrication introduces two failure scenarios into plant and equipment. One is when the parts come from poorly controlled metallurgical or manufacturing processes and are at the weak end of their material capability. These outliers contain imperfections, defects, or flaws of one nature or another—cast, 3D printed, and welded products can have internal faults. When these parts are put into machines and equipment, they suffer operational and environmental stresses. If the capacity of the part is not up to the difficulties of the situation, any inherent defectiveness can cause premature and unexpected failure. The second scenario is when part structure is well controlled in manufacture, but the part is wrong for the duty—it cannot take the stresses and degradation of service. In such circumstances, there is nothing wrong with the item, but it was selected for a situation that is beyond its ability, and unexpected failure again occurs. Both scenarios are the responsibility of the design, reliability, procurement, and maintenance groups to prevent.

Accepting process variability as inevitable is sensible but accepting the accompanying failure consequences as inevitable is disastrous. Organizational processes need to be designed and built to hit all their quality targets right first time. Proactive defect elimination and failure prevention removes process variation. The best way to atop a problem is not to have it at all. If you want to reduce the number of failures that happen in your business, build processes that always deliver the quality standards of excellence you want.

FOOTNOTES

1. Leonard Mlodinow, *The Drunkard's Walk: How Randomness Rules Our Lives* (New York: Pantheon Books, 2009).

2. Chris Denove and James D. Power IV, *Satisfaction: How Every Great Company Listens to the Voice of the Customer* (New York: Portfolio, 2006).

3. W. Edwards Deming, *Out of the Crisis* (Cambridge, MA: MIT Press, 2000), 49.

4. *Fastener Handbook—Bolt Products* (Victoria, Australia: Ajax Fasteners, 1999), 48.

5. "ISO 4406:1999 Hydraulic Fluid Power—Fluids—Method for Coding the Level of Contamination by Solid Particles," accessed at http://www.iso.org/iso/catalogue_detail.htm?csnumber=21463, July 1, 2015.

6. "Contamination in Lubrication Systems for Bearings in Industrial Gearboxes," SKF Ball Bearing Journal, no. 242 (1993).

7. William R. Jones, Jr., and Mark J. Jansen, *Lubrication for Space Applications* (NASA, 2005) Publication NASA/CR—2005-213424.

8. Wayne Bisset, "Management of Particulate Contamination in Lubrication Systems" (presentation at the IMRt Lubrication and Condition Monitoring Forum, Melbourne, Australia, October 2008).

9. FAG OEM and A. G. Handel, "Rolling Bearing Damage—Recognition of Damage and Bearing Inspection," Publication WL82102/2EA/96/6/96.

10. Many real-world process outputs are normally distributed, but distributions can also be skewed or multipeaked.

11. Deming, *Out of the Crisis,* Page 49.

12. John Piotrowski, *Shaft Alignment Handbook,* 3rd ed. (Boca Raton, FL: CRC Press, 2007).

13. "ISO 1940-1:2003 Mechanical Vibration—Balance Quality Requirements for Rotors in a Constant (Rigid) State—Part 1: Specification and Verification of Balance Tolerances," accessed at http://www.iso.org/iso/catalogue_detail.htm?csnumber=27092.

14. FAG OEM and A. G. Handel, "Rolling Bearing Damage."

15. Winston J. Ledet, "Engaging the Entire Organization in Improving Reliability," accessed at http://manufacturinggame.com/resources/engagingtheentireorganization.pdf, July 1, 2015.