Chapter 1: Reliability of Work, Processes, and Machines

If you want a company where great results are natural and excellence abounds, you need to ensure that its processes, jobs, and machines are designed and built to deliver excellence. Every step in every process, every task in every job, and every part in the machines need to go right all the time. It can only happen in the real world when your processes, work tasks, and equipment deliver the right outcomes every time they are used.

A business, a job, a machine must work right by design before it can work right in reality. A business produces products and services from a collection of interacting processes. Build a business of world-class processes, and you'll get a world-class business. Follow a well-built, exact work procedure with properly organized and planned tasks and activities and the job gets done right. Do work without using a designed procedure to control and coordinate the job and you don't know what you'll get. Inside a machine its parts work in a prescribed arrangement carrying their loads, stresses, and strains. When the design is poorly engineered or poorly built then poor performance is what you get from the machine. If the design is robustly engineered and well-built you get a reliable machine that handsomely returns the investment.

Creating a more successful business means designing, then building, more successful processes. A successful process comprises correct inputs, effective tasks, knowledgeable people, and reliable machines working in concert. With the activities, equipment and processes in your company performing at world-class quality, world-class business results become natural.[1] Measuring the chance of business process or work success requires statistics and probability math. Such math can be difficult, but you need only simple multiplication to see what chance you have of getting work and process success in your organization.

Job and Work Process Reliability

Every job is a link in a work process chain. The results from the process depend on how well each job and its activities are done. A wrongly done activity introduces errors and defects that jeopardize job and process success—each process failure damages company performance. Figure 1.1 is a process map depicting a five-task job. From such flowcharts you can gauge how successful the work, the job and the process will be.[2]

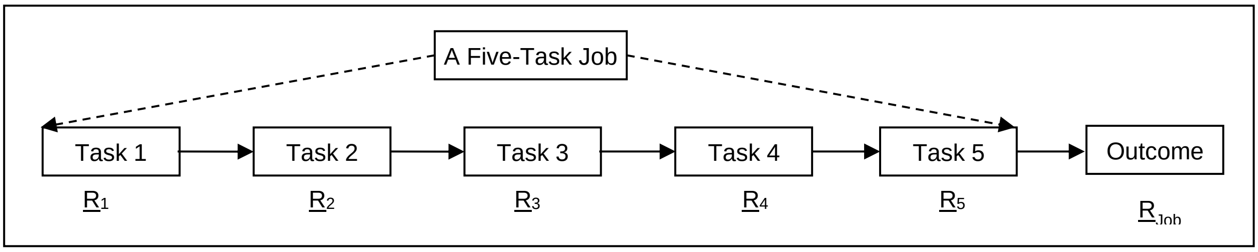

Figure 1.1—A Series of Tasks in a Work Process

To determine work task success rates, you collect data on work task failures. This lets you determine the likelihood of doing each task right, after which you can calculate the chance of doing the whole job right. If Task 1 has a 100% chance of perfect work, its probability of success is 1. If it is done right 50% of the time, it has a 0.5 probability of success. Formula 1.1 is used to calculate job reliability, or the chance of doing the whole job successfully. The underscore distinguishes work task reliability ($\underline{R}$) from system reliability (R), which does not use the underscore.

Formula 1.1 - Reliability of Work Tasks and Activities

$$\underline{R}_{Job} = \underline{R}_1 \times \underline{R}_2 \times \underline{R}_3 \times \ldots \underline{R}_n$$

We can use this formula to see the effect of mistakes on the chance of success in our five-task job. A short list of human error rates applicable to industrial plant operating and maintenance functions is given in Table 1.1.[3] Routine simple inspection and observation tasks incur 100 times fewer errors than complicated work done non-routinely. Equipment and machinery repair tasks belong to the "complicated, nonroutine" category. Usually, repairs are done irregularly on complex machinery, and human error rates during maintenance of 1 in 10, and more, are common (which means 9 out of 10 times, a task will be done right—a 0.9 chance of success).

Situation and Task	**Error Rate** (Per Task)	**No Error Rate** (Success Rate)
Routine simple tasks		
Read checklist or digital display wrongly	0.001	0.999
Check for wrong indicator in an array	0.003	0.997
Fail to correctly replace printed circuit board (PCB)	0.004	0.996
Wrongly carry out visual inspection for a defined criterion (e.g., a leak)	0.003	0.997
Select wrong switch among similar	0.005	0.995
Read 10-digit number wrongly	0.006	0.994
Routine tasks with care needed		
Wrongly replace a detailed part	0.02	0.98
Put 10 digits into a calculator wrongly	0.05	0.95
Do simple arithmetic wrong	0.01–0.03	0.99–0.97
Read five-letter word with poor resolution wrongly	0.03	0.97
Dial 10 digits wrongly	0.06	0.94
Punch or type character wrongly	0.01	0.99
Complicated, nonroutine tasks		
Fail to notice incorrect status in roving inspection	0.1	0.9
New work shift—fail to check hardware, unless specified	0.1	0.9
High stress, nonroutine work	0.25	0.75
Fail to notice wrong position of valves	0.5	0.5
Fail to act correctly after one minute into an emergency situation	0.9	0.1

Table 1.1—Selected Human Error Rates

If every task in Figure 1.1 has a 0.9 chance of success, the whole job reliability is calculated as follows:

$$\underline{R}_{Job} = 0.9 \times 0.9 \times 0.9 \times 0.9 \times 0.9 = 0.59 \ (59\%)$$

Even at 90% certainty for each of the five tasks, the chance that the whole job will be done without error is a poor 59%. In other words, the job will be wrong 41 times for every 100 times it is done. To get a 90% success rate for the whole job, the calculation below warns us that each task will need a 98% chance of success—two errors in every 100 times it is done.

$\underline{R}_{Job}$ = 0.98 x 0.98 x 0.98 x 0.98 x 0.98 = 0. 9 (90%)

As a job gets longer, each activity in it is another opportunity for mistakes. The more activities done in a job, the greater the opportunities to make errors and leave defects, and the fewer times the job will be done right. For a job of 12 tasks in length, with each task having 90% chance of success, its reliability is calculated below to be 0.28—it will contain defects and errors 72 times out of every 100 times it is done. To get the job success rate up to 90 out of a 100, every task will need to be 0.99 perfect—no more than 1 error in every 100 times it is executed.

$\underline{R}_{Job}$ = 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 x 0.9 = 0.28 (28%)

If every task in our five-task job is done right except Task 3, which is done correctly 60% of the time, the reliability of the job is as follows:

$\underline{R}_{Job}$ = 1 x 1 x 0.6 x 1 x 1 = 0.6 (60%)

The chance of the whole job ending right is just 60%. All operating and maintenance work consists of tasks done in series, all of which have far more than the five steps in our simple example. Maintenance jobs of 40 to 50 tasks long, and often longer, are common. Unless every task is done right, the job will leave behind defects and mistakes. The high human error rates for repair work make breakdown maintenance and overhaul repairs very risky if you want maximum equipment reliability and utmost production uptime. Depending on the industry, early life failure of plant and machinery represent 50%-70% of all equipment failures and is most often caused by bad work quality during a maintenance rebuild.[4] Is it any wonder that many companies suffer from poor-performing operations when their managers, engineers, maintenance crews, and operators use failure-prone series processes?

To do a job perfectly, every task must be 100% right. In a series process, such as doing a repair job, operating a production line, using a supply chain, or running a business, when there is a mistake in one step a defect is made or a problem is created, and the final outcome is wrong. This makes for a simple work reliability rule: the chance of job success is never greater than the chance of success for the worst-performed task. It’s the same with every series arrangement: “One poor, all poor; one bad, all bad” is a reliability mantra to remember. It explains why you can have constant production quality problems—make one error anywhere in a series work process, and the finished item will be defective.

Today’s aircraft industry has been outstandingly successful in controlling the outcomes of maintenance processes. It has developed highly reliable work techniques to maintain aircraft in extremely safe flying conditions. It is instructive and insightful to know what these companies do.

When you buy an airplane from a manufacturer, you also get a large set of maintenance manuals explaining in great detail exactly how to maintain the aircraft. The manuals are written by the designers. Every aircraft part is specified by a set of engineering parameters, right down to the formulation of its materials of construction. The designers define and explain the details of the ideal way to install and care for each component in the aircraft. Every maintenance activity is prescribed, including the drawings to use, the job procedures to follow, the technique to apply, any special tools required, the parts to be replaced, and all work record forms. When independent double checks are needed, the procedure specifies where and how the checks are to be done. The industry is highly regulated worldwide, and it is a universal requirement when doing any aircraft maintenance to precisely follow the manufacturer's manuals.

The first question that aircraft mechanics ask before starting a job is, "Where is the manufacturer's maintenance procedure?" They know they can only do their work right if they follow the aircraft's designers approved manuals. Aircraft maintenance technicians are trained, tested, and certified competent on a model of plane before they can get their license to work. They can only work on the specific aircraft models they are licensed for and no others. Throughout their career, aircraft technicians' work is regularly monitored for consistency of quality and accuracy. When new and improved methods are introduced by the aircraft maker the technicians are retrained and recertified. No matter where an airplane is maintained in the world, everyone working on it must be licensed for the currently approved maintenance procedures. If they are not up to the standard, they must stop working on aircraft until their competency is restored.

These are some of the processes the global airline industry uses to maintain planes and make air travel as safe as it is today. The industry has found, from many decades of experience and continual improvement, that faultless aircraft maintenance requires processes to ensure that every job and all tasks are exactly specified and perfectly achieved every time they are done.

Transferred Defect Inheritance and Quality Inheritance

Every defect in a process step has the potential to impact numerous future steps. A defect in an item or work done in a prior step that causes trouble in a later step is termed an "inherited" defect.

It is an error or fault that travels along with the item or job and becomes a future problem in the process or in another process. One defect may only become a minor irritation, while another could turn into a severe business-destroying disaster. Transferred defect inheritance is involved in many business and operational problems, and industrial equipment failures.

A common example of defect inheritance found in machinery is the adverse impact on parts from bad machining practices during manufacture.[5] Three groups of alloy coated steel parts were machined with differing surface roughness, Group 1 was coarsely rough machined, having a surface roughness of 80 micron between topographic peaks, another group was rough machined with 20-micron roughness, and the final group were given 0.32-micron roughness by grinding. All groups were heat treated to harden the surface coating and ground to a finish surface roughness of 0.16 micron, then put into wear trials to find their resistance to abrasion. The coating of the Group 1 specimens wore out the quickest and suffered the greatest number of surface cracks. Group 2 specimens had less wear and fewer cracks than Group 1, and Group 3 had little wear with no cracks at all. Under the microscope a difference in the coating microstructure was observed. The Group 1 rough machining had generated greater heat and produced high internal stresses that had caused many crevices, defects and microcracks in the coating, but these were not present in the Group 3 specimens. A quality characteristic of a prior process step had changed the behavior of a subsequent process step. Surface hardness is important for machine parts that wear during service. If a machine had Group 1 rough machined parts installed, its maintenance costs and production downtime would be far more than had Group 3 parts been fitted. The quality characteristics of a manufacturer’s machining process have dire consequences for the businesses using their machines.

Another example of defect inheritance is a shaft journal machined out of round in a rough turning step that is later turned or ground to the finished size in a fine machining step will have retained its initial oval trait. The ovality is inherited for the life of the journal. If the oval journal is within the design tolerance for its size and shape it will pass dimensional inspection and be used in service, but the ovality produces higher localized stress in the rolling bearing mounted on the journal. During operation the higher local stress combines with other stresses to increase the probability of early bearing failure. To prevent the fine turning step making oval shapes in journals it is necessary to go back into the prior manufacturing process steps to find the faults that caused the oval shape. The problems uncovered in the previous manufacturing steps would themselves have come from earlier failures in the process. Those early failures would have still earlier defects. You would find that there are ever repeating steps of transferred defects followed by the troubles they cause.

Defect inheritance occurs in all processes. Any time an error, a misjudgment, a bad decision, a fault, a deficiency, or any other possible adverse outcome that can occur in a process step happens, it will create the opportunity for problem after problem to arise later. The problems cannot be stopped when they arise, they can only be fixed, replaced or lived with. Problems stop when there is no defect present in the first place to cause the problem. The same data and examples above of defect inheritance apply equally to the exact opposite—quality inheritance. Top quality results achieved earlier in a process also transfer to future process steps. Doing fine quality work brings its own satisfaction and success, but also it brings more success later in the process because quality items perform far better than poor quality items when used in service. High quality results always contribute to the production of good results later, but poor-quality work will only harm future success. The better the quality you produce in each process step and job task, the higher is the chance of success in all the subsequent steps of the future processes that use that quality characteristic.

Business Process Reliability

Figure 1.2 shows a simple production process. Within each process step, there are subprocesses. The Raw Material step has numerous processes within it and impacting it, the Preparation step has its own processes, as will the Manufacture step, and so on for all of them.

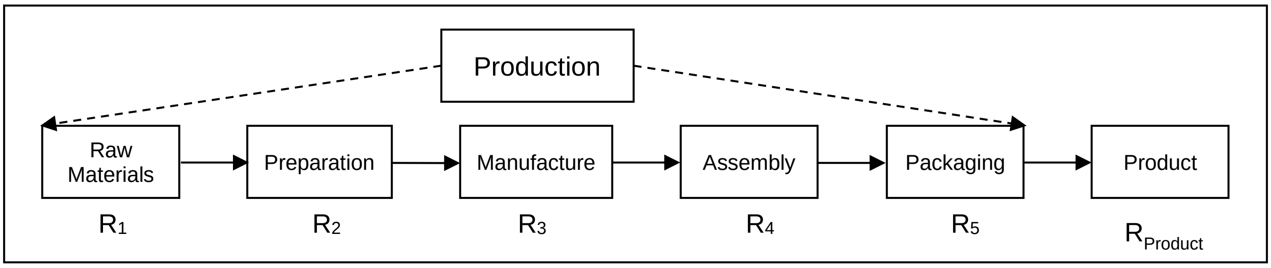

Figure 1.2—A Series of Steps in a Production Process

Figure 1.3 shows some of the processes in the Manufacture step for making a machine part. There are hundreds of activities in dozens of processes impacting an industrial operation.

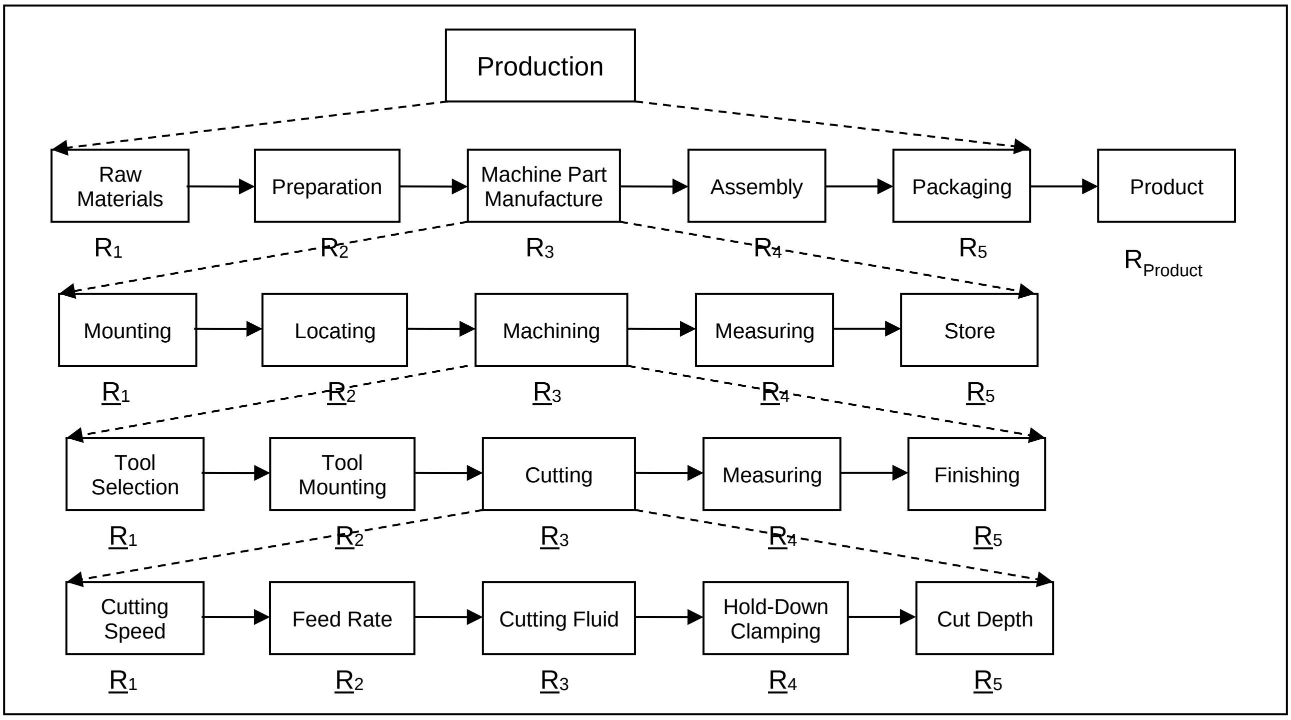

Figure 1.3—There Are Many Work Subprocesses in Every Production Process

Production plants experience many processes in their lifetime.[6] The design, manufacture, supply chain, warehousing, installation, operation, and maintenance processes comprise numerous tasks that must be done right. From time to time, mistakes and poor choices are made in all of them. Those defects eventually lead to equipment or production failures. To understand how business and work processes impact equipment performance, you need to see the interconnectivity of all processes used across the life cycle to engineer, buy, make, and run the equipment.

Figure 1.4 is a representation of the many supply-chain and operational processes involved in making a product. Process after process connects with others in a tangled web of interaction across time and space. There are dozens and dozens of processes containing task upon task. There are hundreds of tasks in most businesses; many companies have thousands of them. Companies with highly complex operations, such as constructing large industrial process plants, constructing big power stations, building spaceships or airplanes, have tens of thousands of activities to control. Each one presents an opportunity for things to go wrong. Because each process feeds many other processes, any error in one has a ripple effect that harms those downstream. A process that goes wrong that is not corrected can impact numerous others in the future. For example, a poor maintenance repair will cause a future production failure; an operator error that overloads a machine will lead to a future breakdown; the wrong choice of materials of construction by the designer of a gas-processing plant will contribute to a future explosion and possibly the death of people. That is why it is important for every step in a series process to go right every time—the future consequences are unforeseeable and can be devastating.

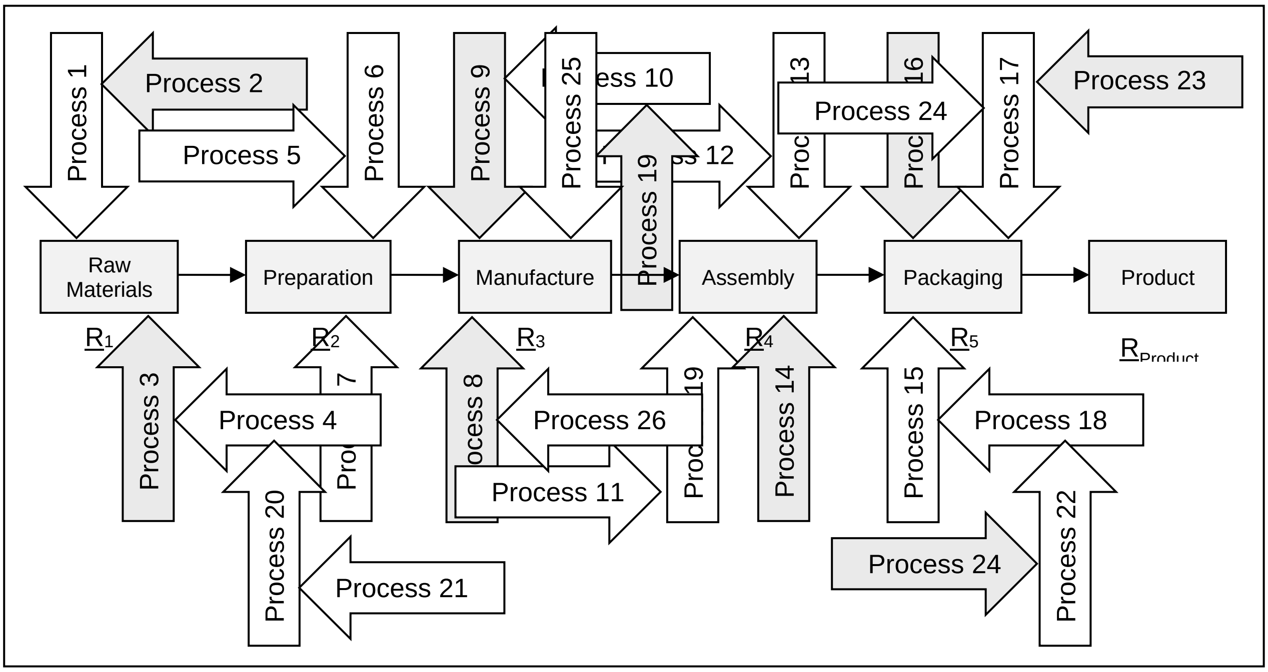

Figure 1.4—Numerous Processes Interact across Every Process Chain

Getting the individual tasks in every process 100% right the first time is a seemingly impossible challenge in running a business. Guaranteeing that every activity is done correctly cannot be left to chance. Doing dozens of processes and thousands of activities perfectly requires a standardized system of excellence. Without ensuring excellence in every process step, you cannot get excellent products or services. World-class operations recognize the interconnectedness and holistic nature of their business and work hard to ensure that everything is right at every stage in every process across the entire business life cycle.

Industrial Equipment Reliability

Figure 1.5 shows how series processes are used in operating plants. It highlights that series processes abound throughout the lifetime of every piece of equipment. During design, manufacture, assembly, operation, and maintenance, multitudes of risks exist that can adversely impact equipment and business performance.

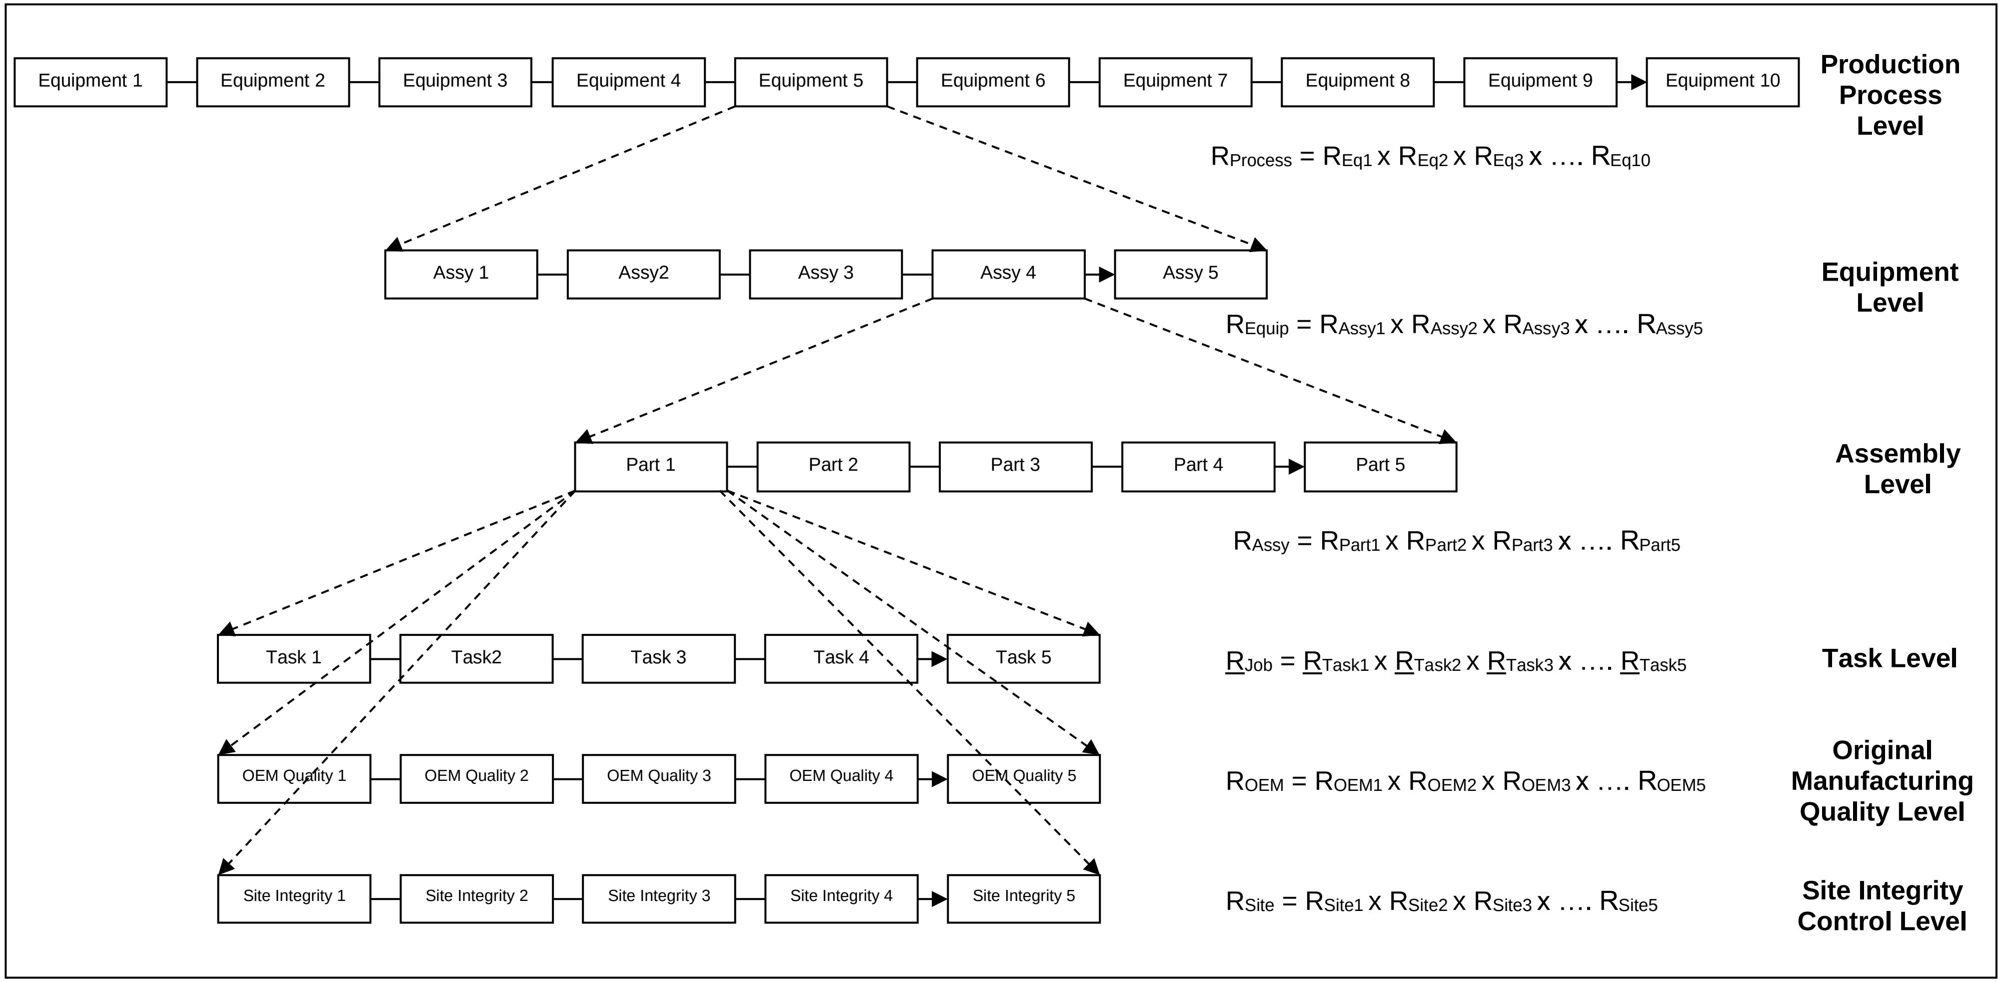

Figure 1.5—Impacts on Reliability during an Operating Equipment’s Lifetime

A machine is a series of parts configured to move and act in an organized sequence. Each part functions on another, which then causes the next part to act, and so on. The parts that suffer operating stresses during use are known as *working parts*. If when a working part fails it causes the equipment to stop or breakdown, then it is designated a *critical working part*. That is why production plants and industrial operations can have many breakdowns—it only takes one failure of one critical part in one machine to stop the whole plant. In sites and facilities with thousands of equipment items, there are millions of opportunities for plant and equipment failures.

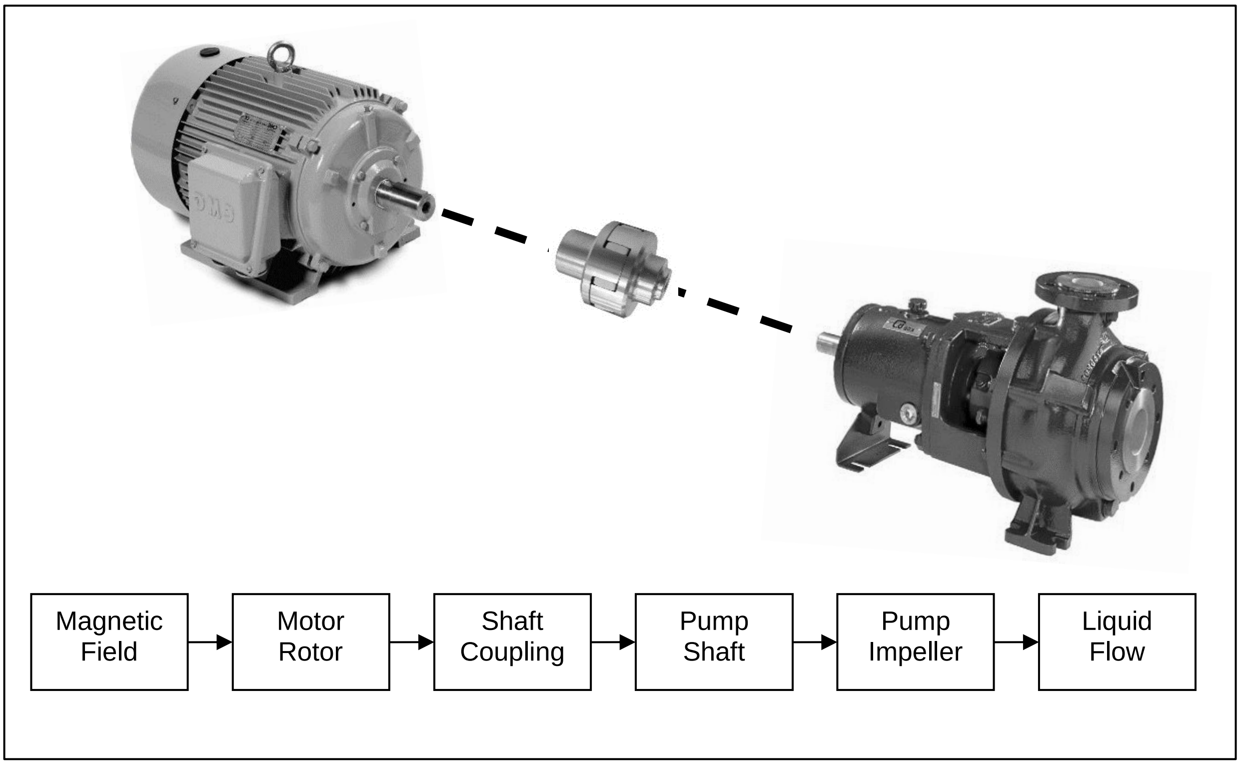

Figure 1.6—Series Arrangement of Assemblies in a Centrifugal Pump-Set

The segmented centrifugal pump-set assembly shown in Figure 1.6 is used as an example to help explain and understand equipment reliability. The electric motor turns a rotor that is connected by a drive coupling to the pump shaft, on which is mounted an impeller. For the pump impeller to spin and pump liquid, the pump shaft must rotate, as must the coupling, as must the motor rotor, as must the magnetic field in the motor. All these requirements for the impeller to turn form a series arrangement. If the diagram displayed every piece of equipment needed to make liquid flow from the impeller, the whole process would start at the power provider's generator and show dozens of process steps. If any process step in the chain fails, the impeller will not turn, and no liquid will flow.

The reliability of a series configuration is calculated by multiplying the reliability of each item in the arrangement, using the following formula:

Formula 1.2 - Reliability of Equipment and Machines

$R_{Series} = R_1 \text{ x } R_2 \text{ x } R_3 \text{ x} \ldots R_n$

As soon as the reliability of any item in the series drops to zero, the whole series goes to zero, and the entire system stops working. If the shaft coupling of the pump-set fails, its reliability becomes zero. The impeller mounted on the pump shaft cannot turn, and the pump-set fails. If the electric motor cannot rotate, the pump-set is again failed. An Internet search by the author for causes of centrifugal pump-set failures found 228 ways for the wet-end components to fail, 189 ways for a mechanical seal to fail, 33 ways for the shaft drive coupling to fail, and 103 ways for the electric motor to fail. This totals 553 ways for one common item in a plant to stop functioning. In those operations with many equipment items, there is a constant struggle against mountainous odds to keep them working. Improving the reliability of your series-constructed equipment is critically important for reducing operating plant failures.

A series arrangement has three series reliability properties.

1. **The reliability of a series system is no more reliable than its least reliable component.**

   The reliability of a series of parts (a machine is a series of parts working together) cannot be higher than the reliability of its least reliable part. If the reliability of each part in a two-component system is 0.9 and 0.8, the series reliability is 0.9 x 0.8 = 0.72, which is less than the reliability of the least reliable item. Even if work is done to lift the 0.8 reliability up to 0.9, the best the system reliability can be is 0.81.

   Series Reliability Property 1 means that anyone who wants high reliability from a series process must ensure that every step in the series is even more highly reliable.

2. **Add "k" items into a series system of items, and the probability of failure of all items in the series must fall by an equal proportion to maintain the original system reliability.**

   Say one item is added to a system of two. Each part has 0.9 reliability. The reliability with two components is originally 0.9 x 0.9 = 0.81, and with three it is 0.9 x 0.9 x 0.9 = 0.729. To return the new series to 0.81 reliability, all three items must have a higher reliability, for example, 0.932 x 0.932 x 0.932 = 0.81. In this case, each item's reliability must rise 3.6% for the system to be as reliable as it was with only two components.

   Series Reliability Property 2 means that if you want highly reliable series processes, you must remove as many steps from the process as possible so your opportunities for failure decrease—simplify, simplify, simplify!

3. **An equal rise in the reliability of all items in a series causes a much larger proportionate rise in system reliability.**

Say a system-wide change is made to a three-item system, such that the reliability of each item rises from 0.932 to 0.95. This is a 1.9% individual improvement. The system reliability goes from 0.932 x 0.932 x 0.932 = 0.81 to 0.95 x 0.95 x 0.95 = 0.86, which is a 5.8% improvement. For a 1.9% effort, there is a gain of 5.8% from the system. This is a 300% return on investment.

Series Reliability Property 3 seemingly gives big system reliability growth for free. Series Reliability Property 3 means that system-wide reliability improvements deliver far more payoff than making individual step improvements. It is the principle that delivers most operating profit most quickly.

These three reliability properties are key to great enterprise asset management and Operational Excellence success.

Reliability, Safety, and Risk

The correlation between safety and reliability is striking. Safety data from industrial sites show a clear inverse relationship between equipment reliability and injuries—plants with highly reliable equipment have fewer injuries.[7] The logical connection between higher reliability and fewer safety incidents is easy to explain: the greater your plant uptime, the fewer repairs you must do, and that means fewer opportunities for injury to both operators and maintainers. Reliable machines and equipment are safer machines and equipment. With perfect equipment reliability, there would be no failures and no risks arising from equipment failure events. This clear relationship between risk, safety, and reliability gives you two highly beneficial risk reduction strategies guaranteed to deliver a safer workplace.

1. Lift your equipment reliability so fewer breakdowns happen, thereby removing opportunities for injury during repair while delivering more plant uptime and lower maintenance costs.
2. Intentionally keep your equipment healthy and in good condition thus preventing failures and creating a safer workplace that gets more throughput of on-quality production.

Whenever the chance of a plant and equipment failure is reduced, your organization is guaranteed to get improved safety results because you create a lower-risk operation. By intentionally making your equipment more reliable, your workforce benefits from having fewer known and unknown workplace risks, and the company profits from safer and more productive plant. You'll get noticeably better workplace safety performance when you proactively identify the hazards inside your operating assets and put in place effective strategies to control them.

The Control of Series Process Reliability

Reliability engineering principles also give us the answer to series process problems—the parallel arrangement. Figure 1.7 shows a parallel layout. The second and higher-numbered items form a redundant configuration with the first item. Should the first item fail, the second item continues in operation, and the outcome from the system is maintained.

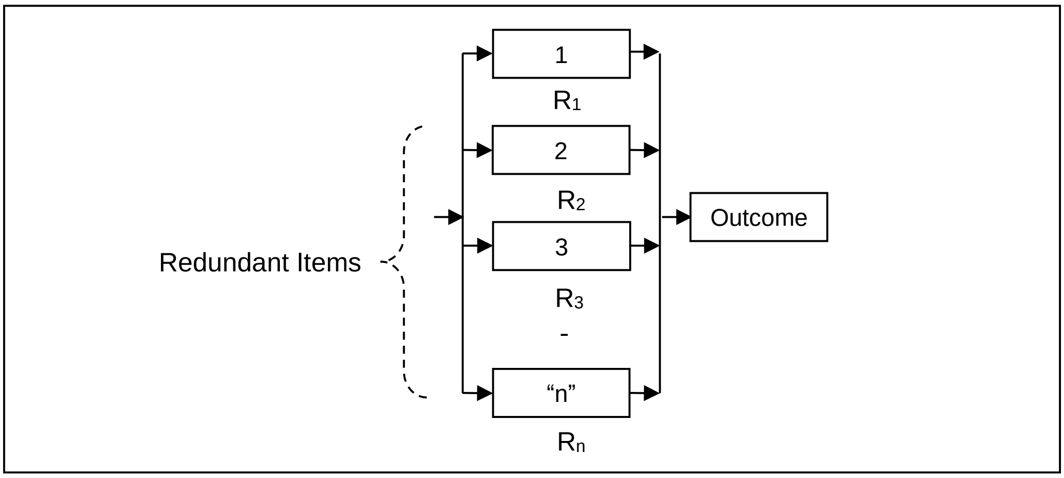

Figure 1.7—A Parallel Process

Reliability behavior in parallel arrangements is very different from that in series arrangements. Formula 1.3 is used to calculate the reliability of a parallel arrangement in which each element is in use and any one of them can do the full duty (known as fully active redundancy).

Formula 1.3 - Reliability of Parallel Arrangements

$R_{Para} = 1 – [(1 – R_1) \text{ x } (1 – R_2) \text{ x } \ldots . (1 – R_n)]$

Other system configurations of redundancy are common, such as a unit on duty and one on standby, two out of three, or three out of four, where one unit is a standby for the other concurrently operating units in the system. Each type of parallel configuration has its own reliability formula that applies to the specific arrangement.

In a fully active parallel arrangement of four items, each with a terrible 0.6 reliability (a 40% chance of failure, which is exceptionally bad odds), the whole system reliability is represented as follows:

$\underline{R} = 1 - [(1 - 0.6) \times (1 - 0.6) \times (1 - 0.6) \times (1 - 0.6)]$

$= 1 - [(0.4) \times (0.4) \times (0.4) \times (0.4)] = 1 - [0.0256] = 0.9744$

This arrangement gives a 97% chance of system success even though each item has a 40% chance of failure. We can use this fact to redesign series processes to get high reliability from them. Putting things in parallel gives you a way to lift production uptime, and it is also a powerful strategy used to get greater job reliability, and to build robust, anti-fragile business processes.

There is a natural economic limit to how many redundant items you can justify in a parallel arrangement. Each extra item requires money to acquire, install, and support. Each item needs regular maintenance and incurs ongoing operating expenditure by its presence. You want as few redundancies as possible in a process, but you can justify a redundancy when the risk of not having it is too high to accept.

Risk is the deciding factor when choosing plant, equipment, or work process redundancy. When the consequence of failure for an item in a series arrangement is excessive, it becomes practical to install parallel redundancy whenever the savings resulting from the redundant item more than pay for its cost, future upkeep, and eventual disposal. Adding a redundancy does not mean you can dismiss the risk. Providing a standby unit does not give you the right to allow anything to go wrong with the working equipment because you have covered its failure with the backup item. Once the duty equipment fails and you start the standby, you lose the benefit of redundancy. Without the standby item, the operating risk instantly jumps to total production loss. When a duty unit in a redundant arrangement stops and the standby is used, it is important to get the failed item fixed in an organized and timely manner—but do it immediately.

Parallel Tasks and the Carpenters' Creed

An example of high-reliability work is the Carpenters' Creed: measure twice, cut once. Carpenters have known for millennia that a double check will save problems and trouble later. We can turn the adage into the parallel task shown in the reliability block diagram of Figure 1.8, where a second measurement is done to confirm the first. By using a proof test activity to verify that the original task has been done right, we create a highly reliable "task system." Though the measurements are sequential, the logical purpose of the proof test measurement is to check the first one. This forms the parallel task arrangement shown in the block diagram.

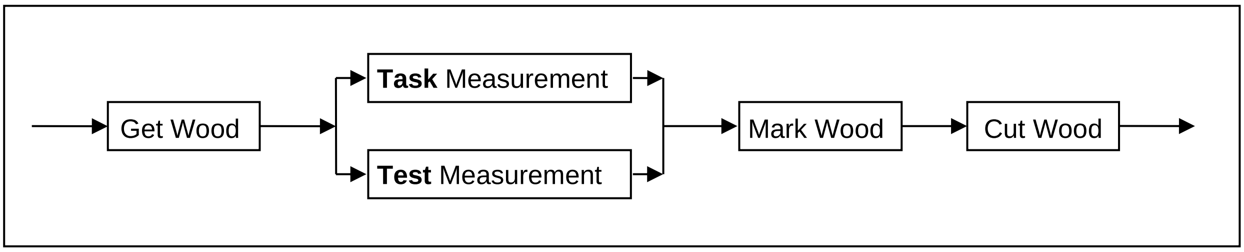

Figure 1.8—Carpenters' Creed: "Measure Twice, Cut Once" Is a Parallel Redundant Activity

The effectiveness of the Carpenters' Creed can be shown mathematically. A typical error rate in reading a tape measure is 0.005—that is, 5 times in every 1,000 it is misread, or 995 times out of 1,000 it is right (a task reliability of 0.995). This means the average carpenter will mark the wood in the wrong spot about 1 time in every 200 measurements. It is not hard to imagine a carpenter averaging 40 to 50 cuts a day. So about once a working week, the carpenter will mark and cut the wood in the wrong place and must throw the job away. When he adds the proof test required by the Carpenters' Creed, he creates a parallel arrangement in which both tasks must fail before the system of two measurements together is failed. He would have to measure incorrectly twice in a row. With the chance of making one measurement wrong being 0.005, the reliability of the two measurements combined into a “measuring system” is found using Formula 1.3.

$$\underline{R} = 1 - [(1 - 0.995) \text{ x } (1 - 0.995] = 1 - [0.000025] = 0.99998$$

With the proof test added, the chance of getting the cut position right rises to 0.99998, which is an error rate of 2 in every 100,000 times. At 50 cuts a day a measurement error is made once every 200 working days or about every 40 working weeks. Doing a check test means 40 times fewer scrapped jobs. That is the advantage of adding parallel proof test activities to work tasks: to ensure that each activity is done right before the next step is started. Note that it is the proof test alone that protects against error. It is only by doing the check test that human error is prevented, and high task reliability is achieved. Without the test, you have no error prevention.

There is one vital requirement for any proof test to reduce the chance of a common cause error. Common cause error is a shared error in which the same mistake is done in both the original and the test tasks. For the proof test, you must use a different measuring device than was used to make the original measurement. It is unlikely to have two measuring devices out of calibration at the same time unless there are systematic calibration problems within the organization. You should also have a different person do the proof test. The person and the measuring equipment form a system. Changing only the measuring device for the proof test and not the person doing the test leaves your business exposed to common cause problems from shared misunderstandings and wrong beliefs existing among your people. Having two totally independent measuring devices greatly reduces the chance of a common cause error. Similarly, by using two competent people to perform independent proof tests, you protect against common misunderstandings, incorrect information, and wrong training. It is unlikely for two knowledgeable, competent people to share the same mistaken education and bad work practices unless they were both wrongly educated and trained.

Figure 1.9 shows the five-task job depicted in Figure 1.1, with each task having a parallel inspect-and-measure proof test to confirm that it is correct. By adding test activities to all tasks in the five-step maintenance job, you create a high-reliability work process.

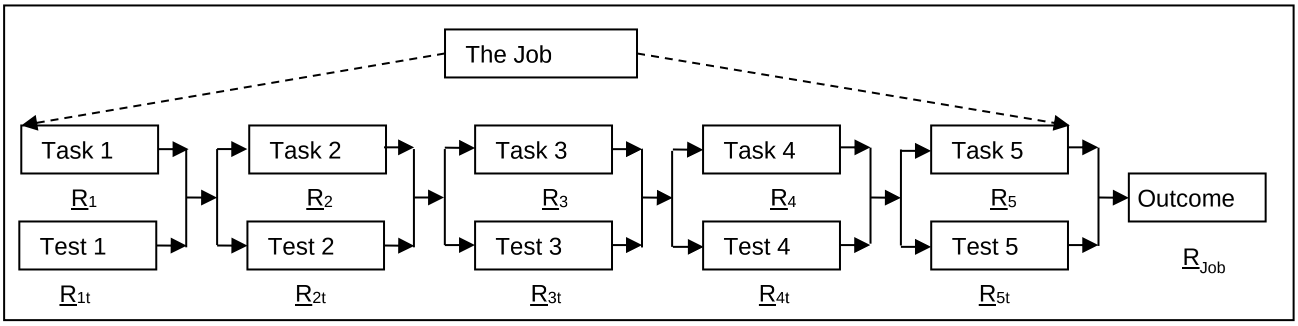

Figure 1.9—A Job with Parallel Test Tasks

If the test has 0.99 reliability—and testing is carefully performed using high-quality tools and procedures—then each parallel-tested step reliability is as follows:

$$\underline{R}_{Task} = 1 - [(1 - R_1) \text{ x } (1 - R_{1t})]$$

$$= 1 - [(1 - 0.9) \text{ x } (1 - 0.99)] = 1 - [(0.1) \text{ x } (0.01)] = 1 - [0.001] = 0.999 \ (99.9\%)$$

The reliability of the whole job is represented by the following equation:

$$\underline{R}_{Job} = 0.999 \text{ x } 0.999 \text{ x } 0.999 \text{ x } 0.999 \text{ x } 0.999 = 0.995 \ (99.5\%)$$

A job that began at 0.59 reliability without any proof tests rises to 0.995 probability of success with proof-tested tasks. But even 0.995 reliability means that 5 times out of every 1,000 opportunities, the job will be wrong. In a large, busy operation with many people, 1,000 opportunities for error accrue rapidly. Similarly, when numerous processes are used to make a product, there are hundreds, even thousands, of opportunities a day for error to happen along the process chains. We need job and process reliabilities of great certainty if we want excellence in our businesses. You can achieve this by adding another parallel activity to each "task system." Figure 1.10 is an example. The test, which involves careful inspection and/or measurement, takes a reliability of 0.99, while 0.9 reliability is used for each of the other parallel activities because "human factors" causing human errors are present when they are performed.

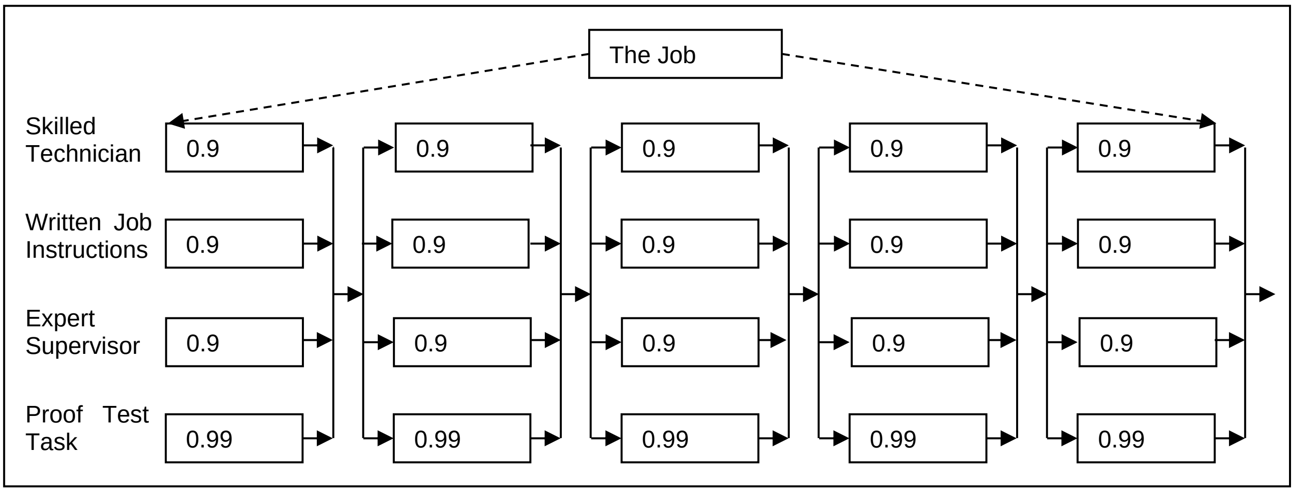

Figure 1.10—A Multi-Paralleled Task Work Process

The reliability equation for each of the multi-paralleled work tasks is as follows:

$$\underline{R}_{Task} = 1 - [(1 - 0.9) \times (1 - 0.9) \times (1 - 0.9) \times (1 - 0.99)]$$

$$= 1 - [(0.1) \times (0.1) \times (0.1) \times (0.01)] = 0.99999 \text{ (1 error per 100,000 opportunities)}$$

The reliability of the entire job of five tasks with each task paralleled in this error-preventing configuration is as follows:

$$\underline{R}_{\text{Job}} = 0.99999 \times 0.99999 \times 0.99999 \times 0.99999 \times 0.99999 = 0.99995\ (99.995\%)$$

The error rate for the whole job is very low: 5 errors per 100,000 opportunities. This is the way to drastically reduce human error and get highly reliable work. To have high-reliability work processes, build parallel inspection activities into the job tasks.

My brother-in-law, who used to work for Japan Airlines (JAL), tells a story of watching Japanese aircraft maintenance technicians overhaul a JAL airplane jet engine. He tells this story because it is so unusual. During his visit to the maintenance hangar, he was enthralled by the extraordinary maintenance procedure that the JAL technicians followed. He watched as a man on a podium in front of a jet engine being worked on read from a manual. Once he’d finished speaking, the technicians at the engine began working on the equipment. The man on the podium went and looked carefully at the work being done. When the technicians finished, they stepped away from their work, and the man, who seemed to be the supervisor, tested and checked their workmanship. As he went through the double-checking process, he would, from time to time, note comments on a form that he carried. Once his inspection was completed and the technicians had also signed off on their work, he returned to the podium and read the next instruction from the manual. The whole process was repeated while my brother-in-law watched in astonishment.

What he saw was JAL’s stringent policy of rebuilding its jet engines by following standard operating procedures paralleled to verbal instruction and supervisory monitoring. The expert supervisor read each task step, explained it, and then monitored the fully qualified and experienced aircraft technicians as they did the task. As the technicians performed the work, the supervisor watched and checked their workmanship. The task was completed only when the technicians and the supervisor confirmed that the work met the required standard, and a record of proof was made of its successful completion. Then the next task step of the job was performed in the same way. By this method, JAL absolutely ensured that its jet engines were correctly rebuilt and fully meet specification. If you fly Japan Airlines, it is reassuring to know the rigors their aircraft mechanics go through to ensure that their jet engines and planes are in top order. Figure 1.11 shows how adding proof tests as done by JAL to create five-level parallel tasks gets amazing reliability.

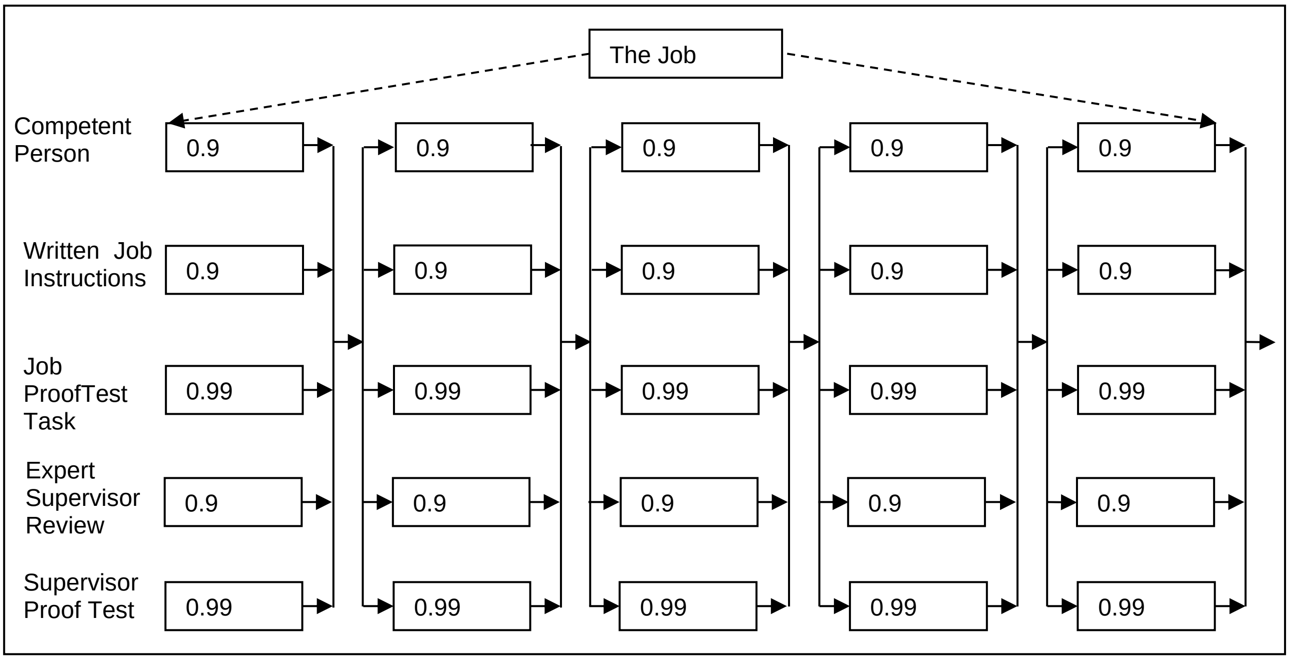

Figure 1.11—A Super-Sure Error-Prevention Work Process

The reliability of each five-level paralleled error-preventing step is as follows:

$$\underline{R}_{Task} = 1 - [(1 - 0.9) \times (1 - 0.9) \times (1 - 0.99) \times (1 - 0.9) \times (1 - 0.99)]$$

$$= 0.999999 \text{ (99.9999\%, or 1 error per 1,000,000 opportunities)}$$

The reliability of the entire job of five super-sure tasks is as follows:

$\underline{R}_{Job}$ = 0.999999 x 0.999999 x 0.999999 x 0.999999 x 0.999999 = 0.999995

(99.9995%, or 5 errors per 1,000,000 opportunities)

Performing each task independently of the other parallel tasks is a vital condition to meet to get these levels of work reliability. In the “Supervisor Proof Test,” the supervisor must use different test equipment from that used by the technician in the “Job Proof Test.”

Getting the maximum reliability from processes should drive all production management thinking and business risk decision making. The design of the work ought to ensure that high reliability is the natural outcome. You want your people to know for themselves when a thing is “done right.” With the use of parallel-tested tasks, human error is detectable and controllable to any level of risk by setting commensurate task quality standards to achieve and by independently double-checking that they are met. Make proof-testing a standard practice in your system of work—make double checks and proof tests “the way we do things around here.” When the reliability is insufficient for a situation or the risk is too high, add a parallel test activity to guarantee a higher chance of success. At least parallel critical tasks with very specific and certain error-preventing tests so you can be very sure that a work process is able to deliver the needed results and quality.

There is a sure way to fail the rigorous security of multiple-level parallel test activities—don’t do the proof tests! An example comes from the 1960s, when an international airline adopted five-level parallel inspections as part of doing critical maintenance work. The joystick and flaps on aircraft wings of the era were connected by wires and pulleys within the wing. The arrangement let the pilots control the position of the wing flaps. After a maintenance task involving the wires, the job record form was returned to the office with five signatures on it, signifying proper work completion and full five-level inspection. The aircraft was needed immediately, and it was pulled out of the hangar and handed over to the pilot ready to fly. It is a necessity that aircraft pilots conduct their own independent tests to confirm that an airplane is in a safe state to use. When the pilot worked the joystick, the flaps jammed and would not move. The urgently needed plane was pulled back into the hangar, and the wires were reinspected. A wire was found to be off its pulley, yet the completion form advised five times that the work had been done properly.

The technician who did the work had signed the maintenance record form, as had four other people. The five signatures on the maintenance record indicated that each one had personally seen the job and agreed it was right. This situation, where people do not check a thing for themselves, is not uncommon. They are mistakenly confident in the capability of the person before them doing their work well, and, seeing that it has already been approved and passed, they don't do their own cross-checks. They misunderstand their role in the work process, which is to be a proof test to protect against human error. Thank goodness the pilot was not one of the five. The aircraft pilot's inspection and test of the maintenance department's work is an error-proof activity that is intentionally designed into the handover process because the risk from aircraft maintenance failure is too great to accept.

When a company culture values expediency over accuracy, or the organization's management practices bullying to rush work to completion, people will take shortcuts and tell lies. This is how you can fail the best security and safety plans even with four extra parallel tests stipulated. Such ethics live in a company for decades, regularly causing unintended, systematic failures throughout an organization’s processes.

The Best Answer Is to Mistake-Proof by Design

Human error cannot be prevented. It is human nature to make mistakes. We will always make errors because our brains and bodies have physical limits.[8] But that does not mean a mistake must lead to a failure. There is a better way to control failure than paralleling test activities. It is to mistake-proof an equipment’s design to ensure that human error cannot cause failure. “Mistake-proof” means changing the design of a thing so the design itself ensures that mistakes have no effect on the outcome. A simple example is replacing stairs with long, inclined ramps. A fall on a ramp is unlikely to happen because there is no place to catch one’s toes and trip. Even if you fell, the injury would be only a bruise, whereas if you tumbled down a set of stairs, you could break a bone.

Figure 1.12 shows our five-task job designed so that each task is mistake-proofed. The job is always completed with perfect reliability.

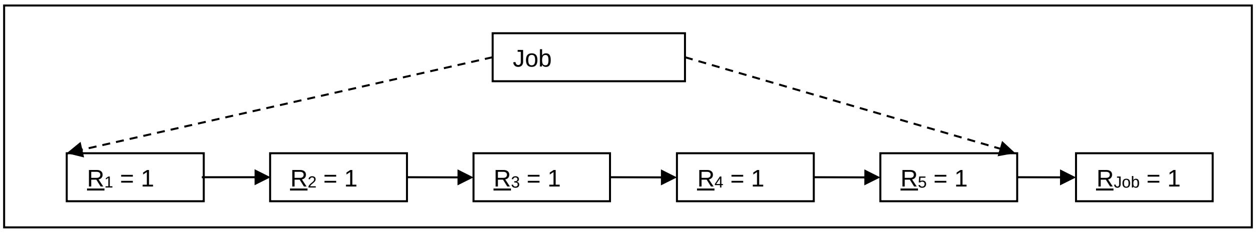

Figure 1.12—A Series Work Process with Each Task Mistake-Proofed

You get 100% reliability in a mistake-proof process. In such situations and circumstances, no human error leads to failure. Mistake-proofing does not mean errors are not made—they are inevitable. Rather, when mistakes happen, they do not fail the job or the machine. Examples of mistake-proofing equipment include changing designs of parts, so they assemble in only one way and providing parts with a telltale indication of correct positioning. We cannot stop mistakes, but we can stop human error from causing failure. Use mistake-proof designs, and the right outcomes result the first time and every time. When you can design the effects of human error away with mistake-proofing, there is no better way to guarantee utmost reliability.

The Operating Asset Life Cycle

Figure 1.13 shows the typical life cycle of physical assets in a facility. It, too, is a series—concept, feasibility, detailed design, procurement, installation, commissioning, operation, and decommissioning. There are multitudes of interconnected work processes in every phase, providing innumerable opportunities for error. By now you should not be surprised to learn that a great number of them become latent problems that play out over time to cause future equipment failures. People can make mistakes and errors anywhere, at any time. Investigations into safety incidents confirm that the root causes of failure occur at all stages of a facility's life cycle.[9] This is why you will regularly hear plant and equipment maintainers cursing equipment and plant designers for their hidden design "traps." The reliability of the operating phase is totally dependent on the reliability of all the numerous human-dependent activities performed in the prior phases.

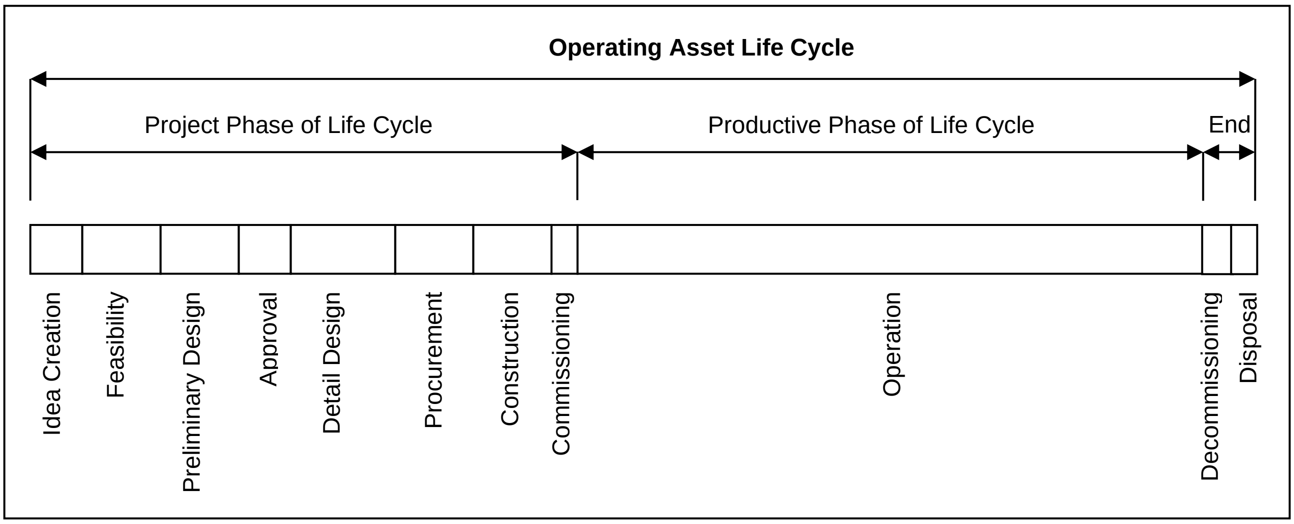

Figure 1.13—The Life Cycle of an Industrial Facility Involves Multitudes of Series Process

Getting high reliability from any series process, whether it is a business life cycle, a job, or a machine, is a decision you make, and then you put into place the necessary quality practices, error-prevention methods, and mistake-proofing techniques to deliver it with certainty.

FOOTNOTES

1. A. V. Feigenbaum, *Total Quality Control*, 3rd ed. (New York: McGraw-Hill, 1993).

2. Mike Sondalini, "Total Control over Human Error" (paper presented at the ICOMS Asset Management Conference, Fremantle, Australia, May 26–30, 2008).

3. David J. Smith, *Reliability, Maintainability, and Risk: Practical Methods for Engineers*, 7th ed. (Boston: Butterworth-Heinemann/Elsevier, 2005), appendix 6.

4. Nolan, Stanley F., Heap, Howard F., *Reliability Centred Maintenance*, Dolby Access Press, 1978

5. John Osarenren, *Integrated Reliability Condition Monitoring and Maintenance of Equipment*, Chapter 4, Section 4.1, CRC Press, 2015.

6. Benjamin S. Blanchard, *Design and Management to Life Cycle Cost* (Forest Grove, OR: M/A Press, 1978).

7. Ron Moore, RM Group, Inc., "Correlation of Injuries with PM & PdM Maintenance Work Orders and with Corrective & Reactive Work Orders," presentation, Knoxville, TN. 2007

8. Malcolm Gladwell, *Blink: The Power of Thinking without Thinking* (New York: Little, Brown, 2005).

9. A. G. Foord and G. Gulland, "Can Technology Eliminate Human Error?" Trans IChemE, Part B, *Process Safety and Environmental Protection*, 2006 84(B3): 171–173.