1

Introduction to Robust Optimization

The automotive industry is very dynamic and the product is continuously changing. The competition is so cut-throat that it is becoming increasingly important to deliver quality products at all times. The customers are demanding the highest quality product at a cheaper price. Robust optimization is the mantra for automotive product development organizations both for original equipment manufacturers (OEMs) and their suppliers, especially in this competitive environment. Dr. Genichi Taguchi's *Robust Optimization* idea is simply revolutionary. To practice robust optimization correctly, product development and manufacturing organizations need to change the way they work, the way work is done needs to change, the way work is managed needs to change, knowledge and skills need to change, the way organizations are led needs to change. Obviously, all of these take time. Not accepting this reality will be more devastating in the future for any organization that wants to win customers' hearts by consistently delivering highest quality products.

Dr. Genichi Taguchi talked about quality as loss to society and how that loss is estimated using a "Quality Loss Function." He talked about robustness – the functional stability of products or processes in the face of ubiquitous variation in the usage conditions (noise factors). He talked about a product development process involving system, parameter and tolerance design steps.

Robust Optimization: World's Best Practices for Developing Winning Vehicles,
First Edition. Subir Chowdhury and Shin Taguchi.

Published 2016 by John Wiley & Sons, Ltd.

He suggested that engineers focus less on meeting requirements and more on discovering combinations of design variable values that (1) stabilize the function and (2) control the adjustment or "tuning" of that function. He talked about ideal functions.

Dr. Taguchi asked engineers and engineering leadership to look at technical work in an entirely different light.

What happened?

Well, since the word "quality" was part of the "Quality Loss Function," the quality experts in the organization took over that concept.

Robustness sounded like product performance in the field. So robustness was delegated to the reliability and validation engineers. Noise factors seemed similar to best case and worst case conditions, so that, too, was a good fit to reliability and validation engineering.

His recommended product development system sounded a lot like existing concurrent engineering and optimization methodologies. System engineers looked at Dr. Taguchi's comments and said, "We already do this – there's nothing new here!"

Parameter design was seen as setting design variable values at levels that met requirements in all conditions. Since parameter design borrowed orthogonal arrays from design of experiments, Taguchi's methods were often seen as a form of Design of Experiment. In most engineering organizations, Designed Experiments were organized by a quality expert when engineering had a problem. Parameter design was delegated to quality and product engineering. Often, an experiment was conducted only if a problem of sufficient magnitude presented itself. Taguchi's parameter design methods were roundly criticized by statisticians for, among many other things, a lack of statistical rigor. Even today, "Taguchi Designs" remain a subset of most statistical computer programs. A subset only "recommended" for preliminary, screening experiments.

1.1 What Is Quality as Loss?

One of our client engineers once had a car with a noisy transmission. He took it to the dealer because the noise bothered him. The dealer attached a machine to the transmission. It printed out a report.

"Your transmission is within specification," the dealer said.

There was nothing more to be done. He drove the car for a couple of years. He was glad when he could replace it with a new one. He never bought that brand of car again – even though their transmission was in specification. The dealer's machine and the printout said so.

Dr. Genichi Taguchi defines quality as "Quality may be assessed as the minimum loss imparted by the product to society from the time the product is shipped." The larger the loss, the poorer is the quality. This kind of thinking says that there is a difference among products even if they are within specification.

The "ideal" amount of noise from an automotive transmission is zero (yes, it's impossible to achieve). As the noise from the transmission increases it will bother some people more than others. But when the noise bothers someone enough, he or she will suffer a loss. They have to take the time to drive to the dealer and wait while the service technician conducts a diagnosis. There will be a dollar value for his time. The drive, diagnosis and report out will take about two hours. Two hours at that time in this person's life is probably worth about $250. Is that the total loss? What about the company's loss of a future sale? How much is that worth? What is the profit the company would make from the sale? The loss suffered by the company who made the noisy transmission is certainly more than $250.

If an automotive manufacturer makes a very, very noisy transmission, a customer might insist that it be replaced. It doesn't matter if the transmission is in or out of specification. The customer wants it replaced. The total loss to society is probably around $3500 (including customer inconvenience). It doesn't matter whether the transmission is under warranty or not. If under warranty, the manufacturer pays; if not, the customer pays. Either way "society" is out $3500 for each transmission that is so noisy it needs to be replaced.

Using this type of data, the quality in regards to audible noise of any transmission can be estimated. The actual amount of audible noise in decibels could be placed along the bottom axis. Dr. Genichi Taguchi is suggesting that every transmission that makes any noise at all contributes a slight amount of loss to society.

The redefinition of quality that you, as the technical leader of your organization, need to embrace is that producing parts within specification is absolutely necessary. However, only producing parts that meet requirements is no longer competitive.

For long-term success in the marketplace, we must focus on producing low-cost products that lower the loss to society. The average dollars lost by society due to audible transmission noise can be estimated for the transmissions made by your company versus the transmissions made by your competition. The long-term competitive position of your company correlates well with such estimates. Products with lower quality loss to society do better over time in the market. Where do your products rate?

While automobiles provide value to society such as transportation and pleasure of driving, automobiles are producing significant amounts of losses. Those losses include emissions, global warming, and automobile accidents. Dr. Taguchi always dreamt about accident-free automobiles and automobiles that clean air.

1.2 What Is Robustness?

What is robustness? You may have to dust off some of your old textbooks (or go online), but you can do it. The ideas aren't that complicated for a technically trained person like you. Let's define robustness as *the ability of a product or process to function consistently as the surrounding uncontrollable or uncontrolled factors vary*.

An example is the power window system in the driver's side door of your car. Does it perform today as well as it did the day you took delivery of it? On an extremely cold morning? On a hot summer day? When you are sitting in the car with the motor off? At 50 mph? Has the window ever stopped working entirely?

If two window systems are being compared, the more robust window system is the one that performs most consistently over a large number of cycles, at low and high temperatures, when running on battery power, or when the car is moving a high speed.

Higher robustness means that a product will last longer in the field, that is, in the hands of the customer. No matter how old the vehicle, no customer should have to awkwardly open the door of her car on a cold winter day to pay and pick up her order at the drive-through window. Only window systems with high levels of robustness can meet that requirement.

Robustness is easy to understand. We appreciate the chain of coffee stores that provides a cup of coffee with consistent taste, aroma, and temperature, regardless of whether we buy it in Seattle or Shanghai. We gravitate toward products that perform consistently over a long useful life. A carpenter needs a circular saw that will last for years of hard use after being thrown into the back of a pickup truck. The expensive two-fuel stove in our kitchen shouldn't have the control panel fail in the first month we own it.

One common misunderstanding about robustness is that more expensive products tend to be more robust. We think that we have to pay for robustness. But is a luxury brand car more robust than a small traditional sedan of one-quarter of the price? In many regards, probably not. More importantly, robust optimization provides methods by which high robustness can be achieved at low cost.

1.3 What Is Robust Assessment?

Robustness is a measurement, not a requirement to be reached. Robustness is only meaningful in comparison. Is my product more or less robust than my competitor's? By how much? Is the new design more or less robust than the old design? By how much? The measure or robustness is the signal-to-noise ratio (S/N ratio). The higher the S/N ratio, the more robust the product or process.

Use the creativity of your people to develop methods to assess (estimate) the robustness of your products in 15 minutes! Usually no more than six measurements are needed to estimate robustness. Most companies that use these ideas strategically develop special fixtures to help engineers estimate robustness quickly and efficiently.

After learning and applying Robust Assessment, an Engineering Vice President at Ricoh said, "From now on, our assessment on a paper handling system will take only two sheets of paper." At Nissan, a robust assessment technique was developed that takes only 15 minutes to assess robustness of a power window system with a high confidence level.

John Elter, a former VP of Engineering at Xerox, said that engineering labs used to be filled with prototype copy machines running continuously for life test and to estimate failure rate. After Robust Assessment, they are filled with jigs and fixtures to measure functions and robustness; functions include paper feeding, toner dispensing, toner charging, toner transfer, fusing, etc.

1.4 What Is Robust Optimization?

Robust optimization, a concept as familiar as it is misunderstood, will be clarified in this chapter. We conduct robust optimization by following the two-step process: (1) Minimize variability in the product or process, and (2) adjust the output to hit the target. In other words, first optimize performance to get the best out of the concept selected, then adjust the output to the target value to confirm whether all the requirements are met. The better the concept can perform, the greater our chances to meet all requirements. In the first step we try to kill many birds with one stone, that is, to meet many requirements by doing only one thing. How is that possible?

We start by identifying the ideal function, which will be determined by the basic physics of the system, be it a product or process. In either case, the design will be evaluated by the basic physics of the system. When evaluating a product or a manufacturing process, the ideal function is defined based on energy transformation from the input to the output. For example, for a car to go faster, the driver presses down on the gas pedal, and that energy is transformed to increased speed by sending gas through a fuel line to the engine, where it is burned, and finally to the wheels, which turn faster.

When designing a process, energy is not transformed, as in the design of a product, but information is. Take the invoicing process, for example. The supplier sends the company an invoice, and that information starts a chain of events that transforms the information into various forms of record-keeping and results, finally, in a check being sent to the supplier.

In either case, we first define what the ideal function for that particular product or process would look like; then we seek a design that will minimize the variability of the transformation of energy or information, depending on what we are trying to optimize.

We concentrate on the transformation of energy or information because all problems, including defects, failures, and poor reliability, are symptoms of variability in the transformation of energy or information. By optimizing that transformation – taking out virtually all sources of "friction" or noise along the way – we strive to meet all the requirements at once.

To understand fully this revolutionary approach, let's first review how quality control has traditionally worked. Virtually since the advent of commerce, a "good" or acceptable product or process has been defined simply as one that meets the standards set by the company. But here's the critical weakness to the old way of thinking: It has always been assumed that *any* product or process that falls *anywhere* in the acceptable range is equal to any other that falls within that range.

Picture the old conveyer belt, where the products roll along the line one by one until they get to the end, where an inspector wearing goggles and a white coat looks at each one and tosses them either into the "acceptable" bin or the "reject" bin. In that case, there are no other distinctions made among the finished products, just "okay" or "bad."

If you were to ask that old-school inspector what separates the worst "okay" specimen from the best reject – in other words, the ones very close to the cutoff line – he'd probably say something like, "It's a hair difference, but you've got to draw the line somewhere." But the inspector treats all acceptable samples the same: He just tosses them in the "okay" bin, and the same with the rejects. Even though he knows there are a million shades of gray in the output, he separates them all into black or white.

Now if you asked a typical consumer of that product if there was any difference between a sample that barely met the standards to make into the "okay" bin and one that was perfect, she'd say, "Yes, absolutely. You can easily tell the difference between these two."

The difference between the inspector's and the customer's viewpoints can be clarified further with the following analogy: If both people were playing darts, the inspector would only notice whether or not the dart hit the dartboard, not caring if it landed near the edge of the board or right on the bull's-eye. But to the customer, there would be a world of difference between the dart that landed on the board's edge and the one that pierced the bull's-eye. Although she certainly wouldn't want any dart not good enough to hit the board, she would still greatly prefer the bull's-eye to the one just an inch inside the board's edge. The point is: With the old way of inspecting products, the manufacturer or service provider made no distinctions among acceptable outputs, but the consumer almost always did, which made the company out of step with the customer's observations and desire.

This dissonance between these two perspectives demonstrates that the traditional view of quality – "good enough!" – is not good enough for remaining competitive in the modern economy. Instead of just barely meeting the lowest possible specifications, we need to hit the bull's-eye. The way to do that is to replace the oversimplified over/under bar with a more sophisticated bull's-eye design, where the goal is not merely to make acceptable products, but to reduce the spread of darts around the target.

The same is also true on the other side of the mark. In the old system, once you meet the specification that was that. No point going past it. But even if we're already doing a good job on a particular specification, we need to look into whether we can do it better and, if so, what it would cost. Would improving pay off?

Robust optimization requires you to free your employees – and your imaginations – to achieve the optimum performance by focusing on the energy/information transformation described earlier. This notion of having no ceiling is important, not just as a business concept, but psychologically as well. The IRS, of course, tells you how much to pay in taxes, and virtually no one ever pays extra. Most taxpayers do their best to pay as little as legally possible. Charities, on the other hand, never tell their donors what to pay – which might explain why Americans are by far the most generous citizens around the world in terms of charitable giving.

The point is simple: Don't give any employee, team, or project an upper limit. Let them optimize and *maximize* the design for robustness. See what's possible, and take advantage of the best performances you can produce! Let the sky be the limit and watch what your people can do! A limitless environment is a very inspiring place to work.

The next big question is: Once the energy/information transformation is optimized, is the design's performance greater than required? If so, you've got some decisions to make. Let's examine two extreme cases.

When the optimum performance exceeds the requirements, you have plenty of opportunities to reduce real cost. For example, you can use the added value in other ways, by using cheaper materials, increased tolerances, or by speeding up the process. The objective of robust optimization is to improve performance without increasing costs. Once you can achieve that, you can take advantage of the opportunities that cost reduction can create.

On the flip side, if the optimum performance comes in below the requirements, it's time to rethink the concept and come up with something better. The problem is that, in most corporate cultures, it is very difficult to abandon a concept because so many people have already spent so much time and effort on the project.

But this is where leadership comes in. Despite the heartbreak of letting an idea go, if it's not good enough, it's not good enough. So instead of spending good money on a doomed project and fighting fires later, it's best to cut your losses, reject the concept (salvaging the best ideas, if any), and move on to the next one, instead of locking yourself into a method of production that's never going to give you the results you want. Thus, it is extremely important to detect poor designs and reject them at the early stages of development.

Dr. Genichi Taguchi has built a model based on this concept that demonstrates the impact that variations from the target have on profits and costs. As the function of the product or process deviates from the target – either above or below it – the quality of the function is compromised. This in turn results in higher losses. The further from the target, the greater the monetary losses will be.

1.4.1 Noise Factors

The bugaboos that create the wiggles in the products and processes we create can be separated into the following general categories:

- manufacturing, material, and assembly vitiations;
- environmental influences (not ecological, but atmospheric);
- customer causes;
- deterioration, aging, and wear;
- neighboring subsystems.

This list will become especially important to us when we look at parameter design for robust optimization, whose stated purpose is to minimize the system's sensitivity to these sources of variation. From here on, we will lump all these sources and their categories under the title of *noise*, meaning not just unwanted sound, but anything that prevents the product or process from functioning in a smooth, seamless way. Think of noise as the friction that gets in the way of perfect performance.

When teams confront a function beset with excessive variation caused by noise, the worst possible response is to ignore the problem – the slip-it-under-the-rug response. Needless to say, this never solves the problem, although it is a surprisingly common response.

As you might expect, more proactive teams usually respond by attacking the sources of the noise, trying to buffer them, or compensating for the noise by other means. All these approaches can work to a degree, but they will almost always add to the costs.

Traditionally, companies have created new products and processes by the simple formula design-build-test, or, essentially, trial and error. This has its appeal, of course, but is ultimately time consuming, inefficient, and unimaginative. It's physically rigorous but intellectually lazy.

1.4.2 Parameter Design

Parameter design takes a different approach. Instead of using the solutions listed above, which all kick in *after* the noise is discovered, parameter design works to eliminate the effect of noise *before* it occurs by making the function immune to possible sources of variation. It's the difference between prevention and cure, the latter being one of the biggest themes of design for six sigma.

We make the function immune to noise by identifying design factors we can control and exploiting those factors to minimize or eliminate the negative effects of any possible deviations – rather like finding a natural predator for a species that's harming crops and people. Instead of battling the species directly with pesticides and the like, it's more efficient to find a natural agent. The first step toward doing this is to discard the familiar approach to quality control, which really is a focus on failure, in favor of a new approach that focuses on success.

Instead of coming up with countless ways that a system might go wrong, analyzing potential failures, and applying a countermeasure for each, in parameter design we focus on the much smaller number of ways we can make things go right! It's much faster to think that way, and much more rewarding, too. Think of it as the world of scientist versus the world of engineers. It is the goal of scientists to understand the entire universe, inside and out. A noble goal, surely, but not a very efficient one. It is the engineer's goal to understand what he needs to understand to make the product or process he's working on work well. *We need to think like engineers, looking for solutions, not like pure scientists, looking for explanations for every potential problem*.

The usual quality control systems try to determine the symptoms of poor quality, track the rate of failure in the product or process, then attempt to find out what's wrong and how to fix it. It's a backward process: beginning with failure and tracing it back to how it occurred.

In parameter design we take a different tack: one that may seem a little foreign at first, but which is ultimately much more rewarding and effective. As discussed earlier, every product or process ultimately boils down to a system whereby energy is transferred from one thing to another to create that product or process. It's how electricity becomes a cool breeze pumping out of your air conditioner. In the case of software or business processes, a system transforms information, not energy, and exactly the same optimization can be applied.

In the parameter design approach, instead of analyzing failure modes of an air-conditioning unit, we measure and optimize the variability and efficiency of the energy transformation from the socket to the cool air pumping out of the unit. In other words, we optimize the quality of energy transformation.

This forces us to define each intended function clearly so that we can reduce its variability and maximize its efficiency. In fact, that's another core issue of parameter design: the shift from focusing on what's wrong and how to fix it to focusing on what's right and how to maximize it. Mere debugging and bandaging are not effective.

To gain a deeper understanding of the distinctions between the old and new ways of thinking, it might be helpful to walk through an example. Let's look at the transfer case of a brand new four-wheel-drive truck. Now, as you probably know, the basic function of this system is as follows: The fuel system sends fuel to the engine, which turns it into active energy and sends it on to the transmission, which sends it on to the transfer case, whose job is to take that energy and distribute it to the front and rear axles for maximum traction and power. The transfer case, therefore, acts as the clearinghouse, or distribution center, for the car's energy.

Even with new transfer cases, common problems include audible noise, excessive vibration, excessive heat general, poor driving feel, premature failure or breakdown, and poor reliability. When engineers see any of these conditions, they traditionally have jumped right in to modify the transfer case's design to minimize the particular problem. The catch is, however, that often "fixing" one of these problems only makes another one worse. For example, we could reduce audible noise, only to find a dangerous increase in friction-generated heat.

It's like squeezing one end of a balloon only to see the other end expand, or quitting smoking only to see your weight increase. Using this approach, instead of eradicating the problem, we've only shifted the symptom of variability from one area to another, and have spent a lot of time, energy, and money in the process.

With parameter design, however, instead of trying to debug the transfer case bug by bug, which often results in us chasing our tails, we focus on reducing the variability of energy transformation, then minimizing the energy that goes through the transfer case cleanly. In other words, we shift our focus from defense to offense.

The theory goes like this: if we could create a perfect transfer case with zero energy loss, there would be no "wasted" energy necessary to create audible noise, heat, vibration, and so on. Sounds good, of course, but obviously building the perfect transfer case is still a pipe dream. But the thinking behind the perfect transfer case, however, can help us build a better one. Wouldn't it be better to try to achieve the perfect energy-efficient transfer case than to try to achieve perfection through endless debugging, putting out fire after fire in the hope of eliminating fires forever? As Ben Franklin said: "An ounce of prevention is worth a pound of cure." We try to build that prevention into the design. It's estimated that in a typical US company, engineers spend 80% of their time putting out fires, not preventing them. Smart companies reverse this ratio.

Usually, the single biggest source of function variation stems from how the customer uses a product or process. (Recall noise factors.) The reason is simple: Labs are sterile places where sensible scientists test the product or process under reasonable conditions, but customers can use these products in a thousand different ways and environments, adding countless variables, including aging and war. Virtually no one can anticipate the many ways that customers might be tempted to use a product or process. This is how we get warning labels on lawnmowers advising consumers not to use them on hedges.

But that's the real world. We cannot prevent customers from using their four-wheel-drive cars in just about any manner they wish. So how do we solve this problem?

Let's take a simple pair of scissors as an example. When designed well, as almost all of them are, they can cut regular paper and basic cloth well enough. But what can you do about customers who buy them to use on materials for which they were never intended, such as leather or plastic?

Most companies would do one of two things. Either they would include stern warnings in the owner's manual, and on the product itself, that the scissors are not intended for use on leather or plastic and that using them on those materials would render the warranty null and void; or companies can give up trying to educate customers, assume the worst, and bolster the design of the scissors so that they actually can cut leather and plastic.

The problem with the first approach is that such warnings only go so far; your company might still be found liable in court. In any case, even intelligent customers might be turned off by a pair of scissors that cannot cut through leather and plastic, even if they never intend to use theirs in that way. The problem with the second approach – making the scissors all but bulletproof – is that, for the vast majority of customers, the extra materials and joint strengthening would be overkill and would raise the price of the product, even for people who will never need such additional force.

With parameter design, however, you don't need to resort to either unsatisfactory solution, because the method helps you create "perfect scissors" that require virtually no effort to cut almost any material. Instead of simple bolstering the device, parameter design streamlines the product to avoid the problems that arise when the product is being used on tough materials, in much the same way that offices solved their "paper problem" not by merely building more and bigger file cabinets, but by converting their information to microfilm, microfiche, and finally to computers.

Making the scissors more efficient reduces the odds of damage and deterioration, and therefore effectively makes the scissors immune to the extremes of customer usage variation without burdening the product with undue costs.

The same concept of parameter design for robust optimization can be applied to the design of a business transactional process. Let's take the efficiency of hospital service, for example. Even for a case like this, we can look at the system as an energy transformation.

Each patient visiting an emergency room (ER) represents the input to the system. Each of them has a different level of demands. One may require a simple diagnosis and a prescription; another may require immediate surgery. The total time spent by a patient in the hospital represents the output. Therefore, we can define the ideal function as the ideal relationship between the input demanded and the actual output. Then we want to optimize the system for robustness. We want the relationship between the input and output to have the least variability at the highest efficiency.

In other words, we want the design to address the number of beds, number of nursing staff, number of health unit coordinators on staff, number of doctors on staff, pharmacy hours, in-house coverage, ER coordinator, dedicated x-ray services, private triage space, and so on. And we want the design to be the most robust against noise factors such as total number of patients visiting, time of patient visit, equipment down time, lab delay, private MD delay, absenteeism, and so on. In essence, we want the relationship between the inputs (the demands of each patient) and the outputs (the time spent on each patient) to have the smallest variability with the highest efficiently. Next, we formulate an experiment with this objective in mind which can be executed by computer simulation instead of more expensive, real-life models.

In summary, teams will learn how to apply the principles of parameter design to optimize the performance of a given system in a far more elegant fashion than just debugging or bolstering it would ever accomplish.

1.4.3 Tolerance Design

In parameter design we optimize the design for robustness by selecting design parameter values, which means defining the materials, configurations, and dimensions needed for the design. For a transfer case in a four-wheel-drive truck, for example, we define the type of gears needed, the gear material, the gear heat treatment method, the shaft diameter, and so on. For a hospital, we define the number of beds, pharmacy hours, and so on. In sum, in parameter design we define the nominal values that will determine the design.

The next step is tolerance design, in which we optimize our tolerances for maximum effect, which does not necessarily mean making them all as tight as they can be. What it does mean is making them tight where they need to be tight, and allowing looser ones where we can afford to have looser ones, thus maximizing the quality, efficiency, and thrift of our design.

For tolerance design optimization, we use the quality loss function to help us evaluate the effectiveness of changing dimensional or material tolerances. This allows us to see if our results are better or worse as we tweak a particular element up or down.

Let's start with tolerance design optimization. *Tolerancing* is a generic label often applied to any method of setting tolerances, be they tolerances for dimensions, materials, or time, in the case of a process.

Tolerance design means something more specific: a logical approach to establishing the appropriate tolerances based on their overall effect on system function (sensitivity) and what it costs to control them. As mentioned earlier, the key model employed in tolerance design is quality loss function. To say it another way, tolerance design describes a specific approach to improving tolerances by tightening up the most critical tolerances (not all of them, in other words) at the lowest possible cost through quality loss function.

This requires us first to determine which tolerances have the greatest impact on system variability, which we accomplish by designing experiments using orthogonal arrays. These experiments are done by computer simulation (occasionally, by hardware). This allows us to prioritize our tolerances – to decide which changes reap the greatest rewards – and thereby helps us make wise decisions about the status of our various options, letting us know which ones we should tighten, loosen, or leave alone.

Think of it as a baseball team's batting order, and you're the manager. Your job is to maximize run production, and you do it by trying different players in different spots in the lineup. The key is isolating who helps and who does not. Substituting various players in the lineup and changing the order will give you the results you need to determine who works best in which position.

Tolerance design will help teams meet one of the primary objectives of the initiative: developing a product or process with six sigma quality while keeping costs to a minimum. Tolerance design is intended to help you and your team work through the process of establishing optimal tolerances for optimal effect.

The goal is not simply to tighten every standard, but to make more sophisticated decisions about tolerances. To clarify what we mean, let's consider a sports analogy. Billy Martin was a good baseball player and a great manager. He had his own off-field problems, but as a field general, he had no equal. One of the reasons he was so good was because he was smart enough, first, to see what kind of team he had, then to find a way to win with them, playing to their strengths and covering their weaknesses – unlike most coaches, who have only one approach that sometimes doesn't mesh with their players.

In the 1970s, when he was managing the Detroit Tigers, a big, slow team, he emphasized power: extra base hits and home runs. When he coached the Oakland A's a decade later, however, he realized that the team could never match Detroit's home-run power, but they were fast, so he switched his emphasis from big hits to base stealing, bunting, and hitting singles. In both places, he won division crowns, but with very different teams.

It's the same with tolerance design. We do not impose on the product or process what we think should happen. We look at what we have, surmise what improvements will obtain the best results, and test our theories. In Detroit, Martin didn't bother trying to make his team faster and to steal more bases because it wouldn't have worked. He made them focus on hitting even more home runs, and they did. In Oakland, he didn't make them lift weights and try to hit more homers, because they didn't have that ability. He made them get leaner and meander and faster and steal even more bases. And that's why it worked: he played to his team's strengths.

You don't want to spend any money at all to upgrade low-contributing tolerances. You want to reduce cost by taking advantage of these tolerances. You don't want to upgrade a high-contributing tolerance if it is too expensive. If the price is right, you will upgrade those high contributors. Tolerance design is all about balancing cost against performance and quality.

One common problem is that people skip parameter design and conduct tolerance design. Be aware of the opportunities you are missing if you skip parameter design. By skipping parameter design, you are missing great opportunities for cost reduction. You may be getting the best possible performance by optimizing for robustness, but if the best is far better than required, there are plenty of opportunities to reduce cost. Further, you are missing the opportunity to find a bad concept, so that you can reject the bad concept at the early stage of product/process development. If the best concept you can come up with is not good enough, you have to change the concept.

The result of tolerance design on designs that have not been optimized is far different from the result of tolerance design after robust optimization has taken place. In other words, you end up tightening tolerances, which would have been unnecessary if the design had been optimized for robustness in the first place. Think of all firefighting activities your company is doing today. If the design were optimized, you would have fewer problems and the problems would be different. Hence, solutions would be different.