PROFESSOR: It's a pleasure to get the chance to talk, particularly [AUDIO OUT] because I think this, in some ways, is the first big punch line of the term, something that-- material you've been hearing about has been leading up to. I guess it depends on whether you consider yourself an engineer or a theoretician. If you're a theoretician, this is going to seem kind of pretty simplistic. If you're a practitioner, an engineer, you're going to say this is-- I hope you'll say this is what we've been waiting for.

And what it is really how we take a lot of the stuff that you've been learning about physical processes and the reasons they go bad and some basic theory on statistics, particularly normal statistics, and actually put it into some sort of use. And we'll do this-- this is really sort of SPC in-- I was going to say 90 minutes, but in less than 90 minutes. And if you follow everything that's in today's lecture and follow what's in the book, you kind of have it all, OK? And then it's just a matter of a lot of shades of gray around that. Making it more applicable, and understanding some of the subtleties.

The other part of it, if I don't get too wordy and we don't have any technical problems, is process capability. And process capability in my mind is the meeting of the two great forces in manufacturing or in a manufacturing company. And that's the relationship-- those are the organizations of design and manufacturing. And it really is just where the two are brought together.

And it's ultimately very simple. It has a lot of implications, but it's ultimately a very simple concept. And both of these things are things that historically you might say were discovered or put together and learned in the United States. They were unlearned-- it's an interesting history on this-- completely unlearned and transferred to other countries, in particular Japan, and then relearn back in this country. So there's some interesting history. If you read about Shewhart, if you read about Edwards Deming, and you read about Juran, these are the people that kind of invented and then brought it back to this country.

And it's now-- you know, again, I find myself now having worked in this area long enough to tell stories that you have to understand that your background, this has probably been in your entire adult life, this has just been standard practice in industry pretty much around the world. But I can tell you, 15 years ago, which for me wasn't-- which, for me, is yesterday, this stuff was still kind of being discovered. But now, at least, it maybe isn't dogma yet everywhere, but it's pretty close to that.

OK. So enough of my soapbox. There's an interesting paper that was written by Shewhart back when he was with Bell Laboratories. Bell Laboratories used to be the premier industrial research laboratory in the United States, maybe in the world. And of course, did a lot of work on communications and telephony and things like that and a lot of things on basic information-- on, of course, on the hardware of that. Software didn't really exist back then, but the hardware of it and other things was done there.

And they had a large theoretical group. And Shewhart, as I understand, was a statistician. And Bell Labs also, in their connection with Western Electric-- Western Electric was the manufacturing arm of the Bell System, which was the one big monopoly in the United States at that time. So they owned the phone system, and they manufactured all the phones. That was an interesting era in history, too. You could get your phones from anybody as long as it was Western Electric. And you could get any color you wanted as long as it was black. They made very solid phones, but there wasn't a lot of variety. But they always worked.

Anyway, they did a lot of work on improving telephony and other things like that, and a lot of statistics and mathematics got developed around control systems. Those if you who've taken control courses, I think-- I'm trying to think, well, I might get some of these names wrong, but some of the techniques and other things that are there came from the concept of developing better amplifiers for undersea cables and other things like that for telephony.

At the same time this development of statistics and applied statistics was being used in that area, this guy Shewhart comes along and says, wait a second. We have this manufacturing. And what this little quote here, which is the introduction to this one paper-- he wrote a series of papers. It basically says is that-- he's stating what maybe is the obvious back in, I think, 1925, which is the objective of manufacturing typically is to create a products that are as uniform as possible.

And you can set up all the conditions you want around the production to make things as uniform as possible, which is exactly what we've been talking about. Control your inputs. Control your process variables, things like that. But there will still be variation in it. He's basically saying, if you ever make things in any sort of quantity, you will always see variation.

And he basically did from that is-- and I don't know how much of this came from the idea of quantum mechanics, which was becoming popular at that time or was coming into being well known by the scientific community at least, which is that at some point, once you reduce things down, sort of explain everything that is to be explained, there's still some inherent randomness in things. And so this view of the idea that there's inherent randomness in any physical thing. Our ability to describe exactly what's going on eventually gets limited.

And so that kind of led to this idea of statistical process control, which is really based on the idea that if I do everything I can to know everything I can about a process, then the only thing that's left is that unknowable part, if you will. The purely random part.

And at one level, SPC is all about, first of all, determining, have I indeed taken care of all those things that I should know about or could know about? And those are things that are related to the physics of the process that we've been talking about so far. And once that's done, what is the underlying statistical behavior, random behavior of the process, and how well can I characterize that.

And SPC all gets down to really basically saying, have I achieved this state of pure randomness? And SPC-- another way of looking at the Shewhart hypothesis is, if the process is acting in any way except purely random, then you haven't done your job. Then you can still improve the process. OK?

OK. So I think you can probably read that on your print out. I can't read it here. So, OK. So the hypothesis in effect, or the approach, is to basically say, all processes have a certain degree of randomness. And this idea of assignable causes, which are really what-- if you're in a control systems, from a control systems background, it would say disturbances or changes. If they've all been eliminated, then it is a purely random process.

So this concept of common causes or, again, purely random effects, if that's all that's left, then you do indeed have a process that is in a state of statistical control. Now, the other problem here again, and the reason I go on about this, is I came to this whole field-- all my formal education, with the exception of one or two classes, never talked about the concept of uncertainty or randomness.

And particularly in the area of control systems, you think of control as actively watching something and doing something about it in a closed loop fashion. Statistical process control is quite different from that. And the term is used quite differently. It really means, have I eliminated or controlled everything external to the process that could cause it to vary? And indeed, is it in a state of what you'll see in a moment we'll define as purely random, stationary behavior. There's nothing active about the control at all.

Now, for those of you who have come up with a control systems or system dynamics point of view-- and I think I warned you this is the last time I talked to you. I do everything with block diagrams. So I have to do this. We have this manufacturing process. And we have certain inputs to the process. And these again, are the things that we can control, and we should be able to control. And we either set them and leave them alone, or we control them in a way that makes the process do what we want to do.

And that produces a certain output through all of the different materials and other things like that. And then in a control system sense, there are two reasons that the relationship between these two would vary. And again, the purely deterministic view of the world says, look. If I set some inputs, I'll get a certain output. There's a unique mapping between the inputs and the outputs.

But two things can happen. You can have things that are called noise. And you'll hear that term throughout the term, actually. Or disturbances. And you'll notice that's sort of external to the-- this is shown as external to the process. And from a mathematical point of view, it's shown as being additive.

So imagine that this is a purely deterministic process. I hold the input perfectly constant. The output stays perfectly constant, except that I have a noise up here that's described, let's say, by a normal distribution. So at any instant in time, this variable will be following a normal distribution. So the output will follow normal distribution.

And you might say that one way of looking at the Shewhart hypothesis is, yeah, here's this process, and I've got it completely under control. I know everything that's going on about it. Nothing inherent in the process is changing, but there's this noise process out here somehow, and it adds in. And it's always there.

Now, this turns out to be a bit of an artifice because what we all know is what's really causing most of the changes that we see is a change in the process. So we have things like a process noise or disturbance.

And so that means something inside here, like those alphas that we talked about. The basic parameters of the process. The equipment states, the equipment the parameters. The material again. Especially the material changes. Material is inherent in here. And so that's going to change.

One of the reasons we actually do this out here is that again, mathematically, having this change, having the function change, becomes really hard to handle. Additive disturbances are easy to handle. Simon.

AUDIENCE: Would you say that temperature and material and everything across [INAUDIBLE]? What examples do you have for the noise disturbance?

PROFESSOR: Yeah, good question. I was going to ask you the same thing. You can see that everything we talked about when we talked about the process modeling had to be here. OK? So a couple of things here. So if the output from here is a dimension and you think of the output of the process here being a dimension that results from the combination of temperature, pressure, force, displacement, machine parameters and all those, what would ever mess that up?

Measurement. Yeah. OK. So now, we aren't going to actually talk a whole lot about measurement in this class. But indeed, one of the good examples of this is measurement. So your perceived output is different from the real output because of some uncertainty from the measurement. So this is also a way of dealing with measurement noise.

And indeed, we deal with measurement noise exactly the same way we do with processed noise. So that's why I was saying that this is a bit of an artifice. If I assume perfect measurement, then this is really a better model of doing it. Either way, our point is going to be that, and the reason I put this in here-- let's assume we did the best job we ever could with this. In some senses, the Shewhart hypothesis doesn't matter. Always going to have something here.

And he doesn't really try to explain where it comes from. That's one of the differences here. This is not really physical origins of variation. It's just, hey. Look. Based on observation, no matter what you do, you get some of this variability. And your job is to get it down to the point where it's, in effect, unexplainable.

OK. So here's the concept of being in control. And this is really an iconic diagram. And it tries to illustrate a number of things. And I think if you understand this diagram, then you really have the essence of statistical process control.

Now, what you see here are-- again, this is, as the chart says, going in this way, you're going in time. And think of these as samples or instance of time or, from a manufacturing point of view, it's cycles of the process. So I make something at time I, and then I make something at time I plus 1 and I plus 2.

Or a better way to think of it is I measure what I've made at I. I measure something else at I plus 1. I measure another product at I plus 1 and I plus 2, and so on. Each of these, of course, is representing a probability distribution. And they're normal here because the Shewhart stuff is all based on assuming the underlying statistics are normal.

What this is illustrating is not the data itself but the distribution of the data as time moves on. And what this diagram tells you is the distribution is identical. It doesn't mean we're getting the same part every time. It just means the probability of getting a particular dimension never changes. OK?

This is the concept of statistical-- a state of statistical control. The underlying parent distribution, the true random behavior of the process, is following this curve and never changes. If you've achieved that, then by the Shewhart hypothesis, you're in a state of statistical control. And that's essentially the best you can do. The best you can do.

And so your process will never get any better than that. That's why we have the second half of the class, by the way, because what comes up after this part, after statistical process control, we say, OK. Now, you've got it in the state of statistical control. It's following this distribution for all time. That's great.

What if that's not good enough, or what if it needs to be better? What if it still has too much variability compared to a design specification? What do you do? Any ideas on what you do if it's in perfect state of control, it's doing exactly what Shewhart says it should be doing to be in a state of control, and it's not good enough?

AUDIENCE: [INAUDIBLE] parameters that narrow the distribution.

PROFESSOR: Yes, exactly. Did you guys hear that in Singapore? No? OK.

AUDIENCE: OK, I heard that.

PROFESSOR: Oh, he heard it. OK, go ahead. Oh, you heard it? [LAUGHING]

AUDIENCE: The research says about just change, optimize the parameters of the process. So try to optimize the process. So I was going to say that that is one way. But the other methods, for example, like positive controls, user feedback control of the system, and try to improve the process as well. So.

PROFESSOR: Yeah, exactly. If you go back to that, again, the somewhat iconic equation that we had in one of the first lectures on the variation equation and the three things you can do to reduce the variation, one was to eliminate all these influences, which is what we're talking about with SPC. Another is to try to reduce the sensitivity of the process to those inherent variations, and that's what Richard was talking about with changing process parameters. And then the third would be to actually use some-- well, we'll talk about using feedback control using some sort of feedback to improve the process.

There's a fourth one, which we didn't really put in there. Well, that's what he was talking about. Active control. Yeah. Can you generalize that? Yeah. [INAUDIBLE].

AUDIENCE: [INAUDIBLE]

PROFESSOR: Oh, yeah, Yeah, that's cheating. Selective assembly? Ooh.

[LAUGHTER]

None of that. No, selective assembly-- indeed, selective assembly is something that you do when this is the best you can do, and it's not good enough. You just take individual parts and say, will these fit? Or you gauge-- and that's was also my airspace example from a couple of weeks ago where you've got the people behind the curtains with the padded table and the padded hammers, and they're pounding things, and that sort of custom assembly. So you can do those sorts of things.

But the last one I was thinking of is you go out and buy a new process. OK? If this is inherent in the process and that's the best that it can do, then you sort of invent a new process, if you will, or you do things a different way that gives you a higher degree of precision.

So you can take all these steps. But in effect, this is-- as you'll see, this is the cheapest, the easiest, and the first one you always do. You never make a decision about what to do until the process is in the state of statistical control. Why is that? Why am I so adamant about that? You never go to the next step until we've done this. Any of these things we talked about. No ideas? Yeah.

AUDIENCE: [INAUDIBLE] involved a lot of efforts [INAUDIBLE] from a an effort point of view, resources. [INAUDIBLE] inherent to the process [INAUDIBLE].

PROFESSOR: From an economic point of view, you certainly should do it. It is essentially almost free, although it takes a little bit of someone's effort and other things like that. But there's maybe even a better reason.

AUDIENCE: The process starts doing something you don't like and you're going to do a feedback system or active control. You don't really know why, so you don't know what your feedback is going to do. If it can actually make it worse because I don't know the underlying cause anymore. So it could actually spiral out of control.

PROFESSOR: This is really getting at the heart of the issue. Richard?

AUDIENCE: Another point is that you have no chance to actually get an approximation of the underlying distribution when the process [INAUDIBLE].

AUDIENCE: Yeah, I was going to be building on Dan and just saying, it's understanding the root cause problem, and you can't really solve anything or understand how to improve it until you understand what's causing the [? error to ?] be on

PROFESSOR: Yeah. Yeah, exactly. The whole point with SPC is, if you remember in this-- in fact, I barely remember this, and I made up the slide. But maybe back in the first or second lecture, we had a list of things that lead up to good process control. The first thing was good housekeeping. Just have to have a sensible shop and don't play around with the process. So hold these things consonant.

The next-- so that's kind of obvious, but you'd be surprised how long it took companies to decide that things like standard operating procedures were important. The next thing after that is, OK. Let's take that one level better, which is if something's going wrong with the process by the definition that's here, then that means that there's something out there causing the process to deviate that I could fix. I could fix it pretty easily.

Bad control over material variability. Bad temperature control. Poor maintenance of the machine. That sort of thing. Shouldn't you fix all those things first before you go investing in fancy controls or something like that? And sure. If you're going to then go on to do process optimization, it's going to assume the process is in a state of statistical control. So it's like you're applying a theory to a situation that it wasn't meant to be applied to. So this is definitely where you start.

OK. I didn't mean to be so preachy today. I guess it's just the weather. OK. Why isn't this going forward? I have to do this now? OK. All right. So here's an example of a process being not in control. And this is maybe more important than the other one because this is the kind of stuff you're trying to detect.

Unfortunately, when you take data, you don't get pictures that look like this. Because, for example, at time instance I, you get one piece of data. You don't get a distribution. You get one piece of data. So this, in effect, is the challenge of working with real data.

Now, keep in mind, and this is a hard thing to keep in mind here in a classroom, but put yourself in the position of a production supervisor or an operator who's making something and products come out one at a time. And as they come out, you measure them. And of course, by the time you've measured it, the next one is ready to be measured and that sort of thing. So things are happening pretty fast.

And what you're trying to decide as you make those measurements is, is everything OK? And you don't have, with each measurement, you don't have enough data to say, well, here's the distribution. But we do the best we can. So here's how it's done.

What you're trying to do with this or what this is basically saying is, well, look what's happening here. In the first two intervals, things looked OK. And then what happened here? Well, obviously, as you can see from the diagram, the whole underlying random behavior shifted it's mean value. And it's clear, it also got greater invariance.

So the underlying parent distribution, meaning the underlying random behavior of the process, did not stay the same. It changed at integral 2. And then down here somewhat later, it got even worse. It followed a bimodal distribution. So it actually meant that the probability of getting a certain dimension had two peaks. And the least likely point was maybe somewhere near-- oh, I guess I've drawn it the other way. But yeah. Well, here's the point. It's anything but a normal distribution anymore. And it's anything but the original distribution.

So these are some examples of the process being not in a state of statistical control. Now, the other thing, of course, that could be happening here-- this is basically saying, with the exception of the bimodal thing, it's basically saying, OK. You got a normal distribution, but it's underlying statistics, the mean and the variance, are no longer constant. The other thing could be happening is that it's still random, but it's switched its distribution to something like a uniform distribution.

The real fact of the matter is that it's probably not following any distribution because it's got a lot of deterministic behavior underneath it. But here's the challenge if you're actually applying SPC. How do I look at data in real time and distinguish the difference among something that looks like this? This is actually a mean shift here now that I look at it. Something that looks like this, which is exactly what I wanted to see. Something that looks like this, where you've got a mean shift. Something that looks like this where I've got a mean shift and standard deviation shift. And something like this, where I just change the whole probability distribution.

So how could I look at individual data points one by one and decide what's going on there? So that's really the challenge. And when we do come up with these rules and say, ah, it's not in control, it's really because our hypothesis is that something like this has happened.

OK. So back to Shewhart, and back to practical. What the Shewhart chart does, the classic Shewhart chart, it's an x-bar, and often, it's done as an r-chart or a range chart. I like to use the s-chart for reasons I'll get to in just a second. But what it does is it plots not the data points themselves. Not the individual run data, but a sequential average of the data. And I'll show you this diagram in just a second.

But the basic approach here is not to take each measurement and put it on a plot but instead to take groups of measurements and plot the average of those. A simple arithmetic average of them. OK? So if I get a plot, a data point here, this is actually the average of a set of data that occurred around the interval i. And the next one here is the average at i plus 1. OK? Why would we do that? Why not just plot the run data, meaning the actual data? Why start off and say, now, we're going to look at averages?

AUDIENCE: Because the evidence or the distribution of the evidence [INAUDIBLE].

PROFESSOR: Yeah. That's one very good reason. You get the narrowing effect with the sample.

AUDIENCE: Also, you're able to calculate or estimate the standard deviation.

PROFESSOR: And that's exactly right. So for each of these, I can come up with a standard deviation as long as I have enough data. Each of those. So each of these is x-bar sub i is-- well, let me just do it this way-- is, of course, a sum over n of some set of data. And so the other part of this now is how many data points do I take? OK, that's a really good reason. There's an even more fundamental reason why I would plot this x-bar.

What is x-bar an estimate of? It's the estimate of the true theoretical mean. So, if you will, it's our estimate of the mean value. If the process is in a state of statistical control, this thing, what is the true value of the mean? Or maybe I didn't state that question as well as I should have.

How does the true value of the mean vary over time? It doesn't. Exactly. So what my ideal for in control is that if this is the theoretical mean right here of the process, then the true mean value should be a constant. So what I'm really plotting here in a sense are my sample statistics. Not the data, but my sample statistics.

And the reason is, in effect, you have to say, at this level, I don't care about the data. What I care about are the statistics of the data and what they're telling me about the underlying behavior. So that's one of the key reasons for plotting this. We're not plotting the data. We're plotting the statistics of the data.

There's another reason that's even maybe-- you'll see some other reasons why. But what other advantage do I get from calculating the average instead of-- doing this kind of averaging instead of just looking at the raw data? We already said it allows us to-- I've said here, it allows us to be actually plotting our estimate of the mean. It allows us to calculate at each interval a standard deviation. What else does it allow us to do?

Remember, I said earlier that this is all based on this being normal. What's the best way to guarantee that whatever it is that I'm looking at is-- whatever variable I'm actually plotting is following a normal distribution? Go ahead.

AUDIENCE: Central limit theorem.

PROFESSOR: Central limit theorem. Exactly. So if you go back to that, if I'm doing any sort of averaging here, I'm invoking the central limit theorem. So even if the underlying behavior of the process naturally is not purely normal, it'll tend to be normal when I do the average. So it's a sort of a self-fulfilling method by doing it that way.

OK. Now, since I can, with respect to this S chart, since I can calculate these at each intervals, you might say, I could have a second chart where I could be plotting sequentially the standard deviation. And so I can actually end up plotting-- let me put that over here. I can plot my sample standard deviation. And I could have another data point here and another data point here.

And so I'm actually looking at both. I'm plotting the sample statistics for the mean and the sample statistics for the standard deviation. And again, what should the underlying standard deviation be? It should be a constant. And these will be samples, estimates of that. And they'll be following a normal distribution as well. And so again, we'll have to estimate what that value should be, but we can get this.

Now, in your text, in most texts on this, they talk-- I know Montgomery talks about S charts. He also talks a lot about R-charts. And R-charts basically just says, for the sample of end data, what's the range? The min-max range of it, OK?

Now, range, why do you think we used range for so many years as opposed to-- we use R, and a lot of the literature is on R. Especially if, let's assume N is is a number of 5. I'm taking five samples. I'm plotting data. Making decisions on the factory floor. Why would I use range which just involves a subtraction of a minimum-maximum number versus standard deviation, which is sum of the squares and all that other stuff? Why would anybody ever want to use range?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Um, no. But I see where you're going. Yeah, it basically-- you know, it bounds the data. But again, what we're trying to get at is the underlying random behavior. So why would I ever use R?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah, think about the most basic user. We're talking about the simplicity of SPC. You got a pad and a pencil and paper. Averaging is pretty easy to do. Subtracting two numbers is pretty easy to do. You write them down. Doesn't take a calculator. Doesn't take a computer.

And I know I'm really starting to sound old now, but pocket calculators weren't around for a long time. And to something that could take a square or square root, forget it. So in the production floor basis, you can have a couple of columns of numbers and write these things down and get it very quickly.

Now, the thing is that this range is actually a pretty good estimate of the variance, and there's some nice stuff in the text about this in other books on this. So range is an estimate of the sample standard deviation. So it's not bad in their correction factors and all that. But the reason I like to do the S is, look. That's what we do now. It's easy to take the data. And very few if any organizations are doing it the other way. So that's what we do.

OK. Now, a little bit about the sampling because this is-- such as it is. That's the theory of SPC. Now, here's the practice of it. How do we actually do this? Well, here it is. Here's a comb, which is actually meant to show you instance in time or actual production instances. So each of these vertical lines is a product, an opportunity to make a measurement. That's the real product.

When we do that sequence then, what we're saying with this sequential average, which is what the Shewhart approach does, it says, well, let's take a group of those, a sample size N, and we'll call that sample interval J. So our time, if you will, is now when did this set of samples-- when was this set of samples taken?

And let's do our measurements, let's do our plotting based on that. So for this orange block, we can calculate, which has, in this case, six measurements in it. So yeah. So n is 6. So I can do an average value for that, for that set of data, and I can do a standard deviation for that. A sample standard deviation from that.

And then I wait a certain amount of time. And I do it again. And so on. So for each of these, I get-- do I have one here? So for each of these, I can get, of course, an x-bar and an S. OK? So you can see there's a trade off between how many samples I take and how far apart I take them.

All right. First question is-- I'll ask it before somebody does because somebody always does. How big does that need to be? Greater than 1. You can do it with 2, especially if you use the range. But how big does N need to be? How big should N be?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. It's exactly the confidence interval issue. But look at this relationship here. Well, yeah. You remember when we talked about the mean and then the invariance of the mean. And the invariance of the mean meaning how-- or, sorry, the variance of the estimate of the mean. This job.

So if I'm plotting my estimate of the mean value and N is a really small number like 1 or 2, then I get the maximum variance of that estimate. It's going to be all over the place. If I increase N, that gets tighter and tighter. So I'm getting a better and better estimate of what the true mean value is. So obviously, maximizing N minimizes the variance of my estimate of the mean. And so I'm doing a better job of estimating the mean.

So why not just make N a really big number? Adam.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. Yeah, I mean, it could take you all day to get a decent sample size, right? And something else-- there's a real issue here that comes up called time variability. So within, if you took a large sample, 50, 100, the whole day's worth-- you could take a whole day's worth of production and say, OK. That'll give me a great estimate of the mean, right?

But during that time, the process could have been going all over the place, changing here and there, and you just say, well, let's wait till the day is over and see how we did. OK? So it doesn't give you a lot of time to intervene. And again, I can't do it on the screen here, but imagine that I cover this up, and all you can see-- of course, you guys can't see anything now, but all you can see is what's happening right here.

You don't know what's going to happen in the future. And you're trying to estimate what's going to happen. If you wait until here to do that, think of all the product that you've made and think about trying to make your next guess. It really gets kind of ridiculous.

So there's a trade off. How big should N be? Well, big enough to get a reasonable estimate. Big enough to be able to detect changes in a timely manner. So if I take a lot of data points, it's just going to take me forever to get an estimate. And then there is the issue of how far apart they are.

AUDIENCE: [INAUDIBLE] depend on the value add of the process that we're looking at? Where if you're adding a lot of value, you might want to look at it or more often because it's costly to miss minor changes.

PROFESSOR: Yeah, sure. Sure. I think at this point, we look at this in a value-- let's say a value-neutral fashion, saying, look. At this point, there is-- your point is well taken in the following. There's a cost to every measurement we make. There's a cost to every average that we take and all that sort of thing.

In the past, those costs could be pretty high. They're tending to be much lower these days with automated inspection and things like that as a matter of course. But yes, there's a cost benefit. So if it's a part where it's not that important, you know, so forth and so on. If you look downstream and you say it's not a critical-- what's the term I'm thinking of? It's not a key characteristic of the product, it can vary, then maybe this isn't important.

But that actually gets to the next step, which is, how critical is a part? So if I hurry up, we'll get to that. But you're exactly right. So there is a cost of quality associated with this. Nowadays, we tend to think there's a cost, high cost associated with not doing this.

OK. What's the effective-- why would we actually wait and not have these boxes right next to each other? Sample 6, OK. So 1 through 6 is the first sample, 7 to 12 is the next sample, and so on. Why would I-- why am I actually showing a gap between these two? Why don't I just don't take any measurements? Ignore the data. Why would I ever do that?

AUDIENCE: [INAUDIBLE]

PROFESSOR: That's exactly right. Say a little bit more about that.

AUDIENCE: I guess if there is certain variation [INAUDIBLE], It would be [INAUDIBLE] any more than other causes of the variation that we take in our data samples right next to each other. It will have the same cause of variation affect [? two ?] of our samples, and in effect, we [INAUDIBLE] same thing twice. So [INAUDIBLE] we will still have [INAUDIBLE] causes of variation.

PROFESSOR: That's exactly right. I hope you guys heard that in Singapore. I couldn't I couldn't say it any better, but that's basically it. And we'll talk about this I think a little bit more theoretically later. But if there is an underlying-- if something happens here that has some memory, some history to it, which causes correlation from one incident to the next, you're violating the assumption of independent random behavior.

If you make this integral large enough, there's something actually called a correlation time. You can measure that for a process, and I think you'll see that near the end of the term. If you sample outside that correlation time, even though the process is not truly independent, it will appear to be independent, which is good for what we want to do. It's sort of like the other half of using the central limit theorem. We get normal, independent behavior, even when the individual data points don't follow that.

There's another reason. Which is less important today than it was in the past, and that is it takes time to do this. And you just don't have time to take all this data. So you wait. You grab six, you take them over to the bench, and you do whatever you do with them, and then you come back as soon as you can for the next one. So you have to allow for this kind of behavior.

Nowadays, again, most production for anything that's critical is actually taking the data, in many cases, on every part. Every part that's coming out has some measurement made on it. I shouldn't say most. Many. And this is all available, and you can choose to use it all. You can choose to ignore it. If you think you have a lot of correlation, you might ignore it. Question?

AUDIENCE: [INAUDIBLE]

PROFESSOR: It should be. For these reasons, yeah. It should be. Let me go on because I'm doing exactly what I didn't want to do, which is talk too much. So, OK. I think we know all that. OK. So this is what you would actually see if you took measurements, grouped them, in this case, in groups of N, and plotted the data. This is what you would see if the process were governed by purely random stationary, meaning fixed mean invariance, normal behavior, because this was generated, in fact, in the PowerPoint that's on the website.

This is-- I think you can click on this, and the Excel spreadsheet that generated it is there. So you can re-randomize it and see what happens. But it's just a normal distribution. And this is what you would get with a normal distribution. In this case, a random number generator that's following normal distribution. And you get stuff that looks like this.

So these are your estimates of the mean value. So you can look at that and say, OK. Well, you know, what does it say? Well, I can observe that-- I can eyeball that the mean value of the mean value looks like a constant. So it looks to me, in this particular case, like x-bar is telling me-- my estimate of x-bar is that the underlying meaning is a constant, and it's probably about right here.

And likewise, I don't see that my estimate of the invariance is telling me that the invariance is anything but a constant. And again, I can look at its average value. So this is what you get. You plot a new data point, and this is what you start to expect.

Now, of course, I can look at this now and get an overall-- a grand mean, which is an average of all the average values, and a grand standard deviation estimate. I should say a grand average and a grand standard deviation estimate. And say, OK. Now I know the underlying distribution of the process. Once I say that, then I know what to expect, right? Now, I can put some lines on this chart and say, this is where the data should fall if everything's in a state of control.

So this whole idea of setting chart limits is extremely important and often abused and misused. But here's the idea. I now have all this data. And this is before-- this is just taking raw data and doing the average on it. I haven't tried to quantify this other than that.

Now, I need to start to put some limits on this or put some information on. So the first thing you do, of course, is say, well, there's the-- on this set of data, there's the average in each case. Now, another thing to keep in mind-- this has already happened. I'm not making any decisions based on this about what to do next. This has already happened.

So the first thing I need to do, and this is part of the art of SPC, is I have to look at that and say, does it look like it's in a state of statistical control? And frankly, at this point, you eyeball it and you say, well, you know, I don't see anything that makes us look non-random.

And we're going to quantify that in just a second. But what else could I do looking at, say, a run of data like this to say, yeah. It kind of looks like it's in a state of statistical control. In other words, if the process is truly stationary, if it is truly following a normal distribution, what would I expect this data to look like? Yeah, it's supposed to look random. But how do I look at it and say, oh, yeah. That's random.

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. If it has a pattern. If it has a trend. Things like that. If it does improbable things. That's the one that I think-- obviously, a trend. If it starts to go like this. If it starts to move up instead of being flat, that's a trend. If it goes like this, that's a trend. Then those are common trends that we'll talk about.

There's some ones that are a little bit more subtle. Think about confidence intervals. OK? Let's put some confidence intervals on here. Let's just put some lines in here. And I'll, just for fun, put three lines above the average and three lines below the average. And I don't even know if those in the right place.

But remember, confidence intervals that said that I should get 63%-- is that right? 63% in the plus or minus 1 sigma. So most of my data should fall in there. That one interval. Even more of it should fall within the next two intervals and the plus or minus 3 sigma. And it's very unlikely that I'll get very many points in the 2 to 3 sigma. A little bit more likely that'll get them within the 1 to 2 sigma.

So if I look at this data and I start to see a lot of data points up here or down here, that's pretty unlikely. That shouldn't be happening. That's an indication that it's not following the behavior that I expected. Or what if I-- in fact, I just drew it here. But what if, in fact, what I saw-- let me get rid of those lines. They're distracting.

What if I saw instead maybe the same data, but it looked like this instead? I was getting, I don't know, two or three data points in a row all in the same band. That's pretty unlikely, too. So based on these things and actually using things like confidence intervals, you can come up with some rules on this.

I'm going to skip this. Is covered in the text very nicely, but there are ways of calculating the sample mean and the sample standard deviation. And there's some correction factors because of biases. But-- for both the x-bar and the S-chart.

So what is typically done, and you always hear about this, is the plus or minus 3 sigma limit. So here's an x-bar chart. Historical x-bar chart. I've calculated the grand mean. Arithmetic average. I've calculated plus or minus 3 sigma of the grand mean with the correction factors and all that.

And so what do these two lines tell me? These, often called control limits. What do those lines mean? They mean that I expect 99.7% of all the data to fall within those two lines. That's all it means. So a simple SPC thing and one you hear about all the time without much thought, frankly, is, oh, there's a data point up there or there's a data point up there. Impossible. Shut down the process. Everything's wrong.

So you can see that's a good first order thing to do. You know, there's only 3 chances in 1,000 that one would be out there, would be in one of these two limits. So that's a pretty good first cut. But it's a little bit simplistic because it is possible that you'd get a process that's perfectly in control, and you'd have one point that's up here. It's just not probable, but it's possible.

So we can do a little bit better than that. We can-- and again, we can do the same thing with the S-chart. We can do upper and lower control limits. But the more important thing, I think, is not waiting for that one data point that pops up here. For example, one of the reasons that might be a problem-- what if I have underlying behavior that looks like this? And you can see that this is, again, this is truly normal data. And it's a smaller sample set that it's highly improbable I'll get anything near the 3 sigma limit.

And this is what you typically get if you had a process that this was in control. But what if it was doing this? And you eyeball that. And would you say, is anything wrong with this process? Yeah. Look, it's trending up like crazy. When did that trend begin? Maybe here.

Now, if all I do is sort of have an alarm when I cross over this upper limit, can you see that I've been-- something's been wrong for a long time. And this could be hours. It could be milliseconds, but it could be hours. Something going wrong for a long time.

It would be nice to have something a little bit more subtle than just simply waiting for that one random, potentially just purely random event. So-- although we'll talk about that in just a second. So one of the things you can do, as I said, is look at the data on the basis of looking at confidence intervals on the frequency of getting any sorts of extremes and trends.

And so there are a number of different things. Now, there's a text we used to use for the class by DeVore, Chang, and Sutherland. We stopped using it because it's just got some rough edges and theoretical problems. Not problems, but lack of-- it's theoretically correct, but it's got some gaps in the theory. So you need another book anyway. Montgomery is a little bit more basic on it.

This is a little bit better on some of the process optimization and pragmatic-- these guys are engineers and less statisticians. But one of the things they did is they came up with the so-called eight rules. In fact, I'm going to-- let me go past those for just a second because-- oops. OK, never mind. I was going to talk about that, but I don't know if it's in here. OK.

So the eight rules are based on really just four tests. The probability that the data falls into one of these confidence interval bands. A measure of periodicity. You should have no periodicity in the data. No regularity in it. Linear trends. Something that's changing up or down over time. And mean shift-- a jump change in the mean.

All these things are things are going to be looking for in the data. Now, mean, you're going to hear a lot about mean shift throughout the term because that's your most-- in some ways, you can think of that is as one of the most likely things to happen. You're making product with a certain batch of material, and you change the batch. Something changes.

You have a shift change. You have PM on the machine, Preventive Maintenance on the machine. And the next time you run it, things jump up. So the underlying behavior, something in this model, just made a step change, and you saw that out there. Linear trends, obviously. Something. Tool wearing. Something not-- something changing over time. Periodicity. It could be temperature changes. It could be something that's underlying. That sort of thing.

OK. So test for out of control. I think this comes out to eight. So if you have a set of data and you look at it, first thing you might do is say is anything outside the plus or minus 3 sigma limits. That's really unlikely. If I get any extreme points like that, pull the Andon cord. Stop the production. Something's wrong. Let's go out and fix it.

Improbable points. Now, they've quantified this by saying, look. If I get 2 out of 3 points, if I look at any triad of points, if 2 or 3 of those are in the plus or minus 2 sigma band, that's pretty improbable based on confidence intervals. 4 out of 5 in the plus or minus, outside the plus or minus 1 band. That's pretty improbable. Or all the data inside the plus or minus 1 band. That shouldn't happen, either, if my process is stationary. OK? So each of these things you can see could indicate that it's not following the expected normal behavior.

Another one is in runs of 8 or more points, am I always-- is there a run of eight points? That's always above the mean value. That's improbable. It's supposed to be random on both sides. Linear trends, of course. Six points. And they pick the numbers to trade off between sensitivity and resolution.

Six points in a consistent direction. Always heading upwards or downwards. And bimodal data. You got points-- if I get points outside of the central region, that could indicate-- you go back to that picture I had. I have a distribution here. I have a distribution here. It's actually following two different mean values on either side of the mean value.

OK. So how would I apply the Shewhart charting concept? OK. Well, keep in mind, the first thing I've got to do-- well, let me back up. Not the first thing. What I'm trying to accomplish is this. I'm in production. Let's just look at the x-bar chart.

I'm in production, and this is now-- I've got some data back here. This is now-- and I took another piece of data. And I want to say, OK. Is everything all right? Time to take any action? I can't do that until I know I have at least done my grand mean and I have at least my upper control limit, my plus 3 sigma, and my lower control limit. So I can't really do SPC until I have these limits.

Where do they come from? They came from historical data. They came from the stuff that already happened. So this is a little bit of a catch-22. What do you do if the process is horribly out of control. And back here, it's all over the place. And you just say, oh. Well, here's the mean value on the standard deviation.

You can see that these will be totally bogus limits. And if it's out of control back here and you put those limits on here, guess what? It'll appear to be in control here, too. It'll stay within these limits. So it's a bit of a catch-22.

But what do you do with these-- how would I take these, I don't know, 25 to 50 points before I actually start doing control? How would I take those and say, OK. Looks good. What kind of things would you do if I were taking data to say, yep. Looks like I can now apply the chart. What am I looking for in the data?

AUDIENCE: [INAUDIBLE]

PROFESSOR: That could be very dangerous for a reason I'm going to show you in a second. But yeah, you have certain expectations for the process. Yeah, absolutely. But let's say that, again, here's the historical data. I've got the historical data. I always draw it looking periodic.

So I've got the historical data. I've got to look at that data and say, what do I need to say about it to say it's useful for charting? Number 1, it needs to be normal. Number 2, it needs to be stationary. Then I can start doing the charting. What's the test I would do for it being normal?

AUDIENCE: [INAUDIBLE]

PROFESSOR: Yeah. Exactly. You could do things like-- that's one of the reasons we looked at this. Do a histogram. What else? Yeah, yeah. You do a QQ or normal distribution plot. Really good thing to do. That's pretty fast and easy to do. Is it following normal distribution? If it's not, we got a real problem.

I just had a student of mine who plotted some measurement data, and he plotted on a QQ, plot and it was nowhere near following it. Didn't make any sense. And then I found out that, underlying, he'd been doing some very deterministic things between measurements. Changing things in a very deterministic way. And it came out exactly in the data. It should have followed the normal distribution, and it didn't because of this. What else? Hayden, what else could we do?

AUDIENCE: Well, I'm thinking you could get the upper and lower control limits there in your specification.

[INTERPOSING VOICES]

PROFESSOR: You got to careful of that, though. Yeah. Good point. That's the same issue of, you know, here's what I expect it to be. But the whole point about process capability, which I'm going to get to in the last five minutes, is process capability is what you're dreaming of getting. This is what you actually get. And they can do really far off.

The reason I ask you is you just put that derivation of the kurtosis up there. So you could do other tests on the data like kurtosis and skewness. There are all these different tests to see if the data is normal. And then that's all well and good, but then you have to look at it. Is it trending up or down? Does it have any of these things?

Before I ever put a limit on it, can I say, yeah, that kind of looks random? And even without a limit, you can say, OK. The data looks like it's not trending up. I don't see any periodicity in the data. In the case with my student's data, there was this big hump in the data where it spent a lot of time doing this, and then it came back being [INAUDIBLE]. It doesn't-- I may be wrong, but that doesn't look random to me.

OK. So once I've done that-- and this is where the art is-- then I compute the center lines and the limits and I begin plotting and do this. And then I start to apply the eight rules. Or there are some simpler rules, often called the Western Electric rules, that are actually in your text.

And again, what these rules mean is these are alarm points. These mean, to the best of our ability, your process is not following a state of statistical control, meaning it's not stationary. There is some sort of cause on the outside that's making a change, and it can be eliminated and identified.

And these are the Western Electric rules that are in your text. Any points outside the limit, 2 to 3 points outside 2 sigma and 5 points outside 1 sigma, that sort of thing. OK. So you can see that just reviewing what we said, this is in control. Here's what the data should look like if it's in control. And if you see these types of behaviors in the data, it means something like this has happened. Mean shift, trends, bimodal, whatever.

I'm going to skip over this thing on average run length. It's very interesting, and it gives you an issue-- some idea of what to do about sample sizes. And it's very carefully covered in the text. And you can play with this beautiful animation that was made up a long time ago.

And let me just summarize this by saying, keep in mind, what we're really doing here is hypothesis testing. So all the stuff you just had on hypothesis testing is exactly what we're doing. We're testing the hypothesis that the data is defined by a normal distribution.

And these charts are allowing us to test it, but it's sort of testing it as we go along. And these issues of alpha and beta probabilities all fall into this. So the power of the test is important. I think I'm getting redundant here.

OK. So again, repeating what I said earlier, remodel the process as a normal, independent, random variable. By doing sampling the way we did, we sort of invoked the central limit theorem, and by potentially by spreading the samples out, we have a better chance of it being independent, even if it's not. We say, as a result of it being normal, there are only two things I need to know about the distribution. That's the beautiful thing about a normal distribution. Two parameters, mean and variance.

So it's completely described by that. I only need to know those two things. So I estimate those with x bar and S on the data and then enforce this concept of stationary conditions. So if it's a stationary, normally distributed process, everything's fine. If it's not, everything's not fine.

Oh, yeah. What happens if it's not right? This is great. I read this in the statistics book a long time ago. If it's not right, you call an engineer. So if you're from a purely statistical point of view, you're just saying something's not right.

OK. Now, I'm rushing because I did want to mention this. I'm sure [INAUDIBLE] will do some more on it. But I really wanted to get to this because of some of the questions that came up. This is our empirical model of the process now. It's an empirical, statistical model of the process.

The real world, our estimate of the real world, our underlying model, is this theoretical parent normal distribution. That's what you get. It's not what you want. It's what you get. So the question is, how good is the process?

This is what it is. Is that good, or bad, or indifferent? And that's where you get into the concept of process capability. So we assume the process is in control. We assume it follows that normal distribution. Then we compare it. Then we compare it to things like tolerances.

Where do I want my parts to be? And there's also this concept of quality loss, which is a better way, in some ways of looking at design specifications. The simplest design specification we can think of is an upper and lower limit. A tolerance. If you can make me a part that falls between these two limits, it's a good part.

So can think of that as a nominal value, a target, and then an upper specification limit and a lower specification limit. Not control limit, but specification limit. Tolerances basically say, anything in there is a good part. Has equal quality, if you will.

Quality loss, which was introduced by Taguchi some time ago, basically said, that's baloney. And there's some really good case examples of this, but that's baloney. He said, look. If there's a target value for some specification, it's a dimension or a property of something.

And you can specify that. That's the best one. Any deviation from that is less good, and there is an actual cost associated with that deviation. And he said, not knowing anything more about it, we'll call it a quadratic loss. So the farther you are from the specified value, the worse the product is. And of course, you should always work your way down this curve.

And there's a way to calibrate this, which basically would say-- one way to think of it is-- a simple way to calibrate this. If I actually have upper and lower specification limits here, I can say, what's the cost of scrapping apart? And where this curve crosses that, you say, OK. That's my calibration constant.

But again, there's some great case examples that show that if you work against quality loss, you can actually end up with much-- once you combine together parts into the whole thing, you actually end up with a much better overall quality. But for the moment, let's just look at the tolerances. Process capability in it's simplest form is just this. I start with my desired.

I'd like the part to fall between these two limits. Once I say that, it doesn't matter what the target is, right? It's just between the two limits. I'd like it to fall between the two limits. There's my process. That's the real world. How did I do? What's the likelihood that the part that I make will fall inside those limits? That's process capability.

You can, and their numbers-- this is the chart. You can look at that and say, OK, how likely is it that I make a bad part? Well, it's the area of that little space right there. The way I've drawn this. It's that area right there, and it looks like nil here.

So I could look at that and say, that's. The probability that I make a bad part. Again, to try and give us a more common basis for this, things like this have been defined. The CP parameter. Process capability, which basically says, what's the ratio of the width of that window to my plus or minus 3 sigma width? So is the variability of my process similar to or different from the underlying, the desired?

But it compares ranges only. And if you look at this carefully, that could say that I could have the following. I could have a process where these are my specification limits. And here's my distribution. It's pretty narrow.

So the width of this, plus or minus 3 sigma, is well within this width. So this number would actually be pretty large. This window is wide compared to the window of this. But how many bad parts do I make? They're all bad, right?

So CP is one way of looking at how well you're doing with the process. But it's really only looking at variability. So we come up with this CPK measure, and there are a whole pile of these different ways of looking at stuff. CPK now penalizes deviations from the mean. OK?

So now, this would have either negative or zero CPK. And if you do the calculation, this is sort of a blend between how well it's centered in here and what the deviation is. So if it's perfectly centered, then these two numbers-- this is actually-- if it was perfectly centered and it exactly fits like this, what's CPK going to be?

So here's the center specification and here's x-bar. This is exactly plus or minus 3 sigma. This distance here will be 3 sigma. So your process capability, CPK, will be 1. So let me just end up with this and just say that there are a couple ways you can look at this.

If I have a process perfectly centered and exactly plus or minus 3 sigma limits to my upper and lower specification limit, both CP and CPK will be 1. OK. If that's the case, what's my parts per million, which many companies will use to measure their quality? Once my parts per million bad parts?

AUDIENCE: [INAUDIBLE]

PROFESSOR: 3 out of 1,000. So it's a lot. So yeah. So CP of 1's not particularly that good. Companies like to have it a lot higher than that. Here's an interesting one. Look at the difference between this. CP is 1. CPK is 1. CP is still 1 because the width of this distribution didn't change.

But the mindshift shift was such that I actually have-- half my parts will now be bad, and that, if you calculate, gives you CPK of 0. Here's another interesting one. CP went up to 2 because I narrowed the distribution, but it wasn't centered. And so I got a CPK of 1. This is really good. This is really good.

This is now plus or minus-- can you see that's plus or minus 6 sigma out to here? So what's the likelihood of making a bad part now? A couple per million, I think. this-- I don't have time to go into it, but this is sort of the origin of the 6 sigma revolution of saying, wow. If we could do this, we'd never make a bad part. There's a little bit more to that, but that's process capability. There are all sorts of different forms of it.

But the last thing-- let me just say one more thing about process capability. All these things, CP, CPK, six sigma, they're all great. Keep in mind that, again, if you can do it, if you actually have the data to do it, it's useful to do, I think. What it really says is here's the design, and somewhere here is manufacturing.

And both of these can change. OK? Both of these can change. I can do better manufacturing. I can center it. I can do worse manufacturing. I can buy new equipment. I can change this all the ways we've talked about. I can also go back to design or start with design and say, you know, this part really isn't that important. This window can actually be here.

And if I move the window out there, look at how the process capability just changed. In case 1, the process capability of this distribution against this design is terrible. The CPK is almost 0. Now, if I opened up the window, CPK now skyrockets. I'm great.

The trend usually is just the opposite, though. And this is certainly true in semiconductor manufacturing is you might say, here's the inherent underlying capability of many of the processes that we've had historically. Here were the design specifications back in 1980. Line width or something like this. Characteristic to mention.

Here they are now. Oops. Can't make this product with this process. And it was all because-- so you can look at how I have to improve my process capability. And it's interesting to talk to people who've been in the semiconductor industry their whole lives, and they'll tell you this exact story that, well, we didn't worry about this back when the specifications for this wide. But all of a sudden, it became much more important, so we had to do a lot of things, including not just SPC, but some extraordinary things. OK. Thanks for your patience. I ran over, but it was a technical difficulty at the beginning. OK. Bye bye.