

Today's Drive Toward Better Quality Control

Quality Control is an ever-changing field designed to ensure patient safety, but what does the future of QC look like? In this informative webinar, Sten Westgard, Director of Client Services and Technology at Westgard QC, looks at the past for insight into how the next generation of QC will look. Westgard is a leader in the field, an adjunct faculty member and honorary visiting professor at schools around the world who has authored multiple books, essays and reports on the subject.

This webinar, *Today's Drive Toward Better Quality Control, and What Will the Next Generation of QC Look Like?* features Sten Westgard, and was hosted by Werfen.

The objectives of the webinar include: 1) to identify strategies to address the risks of variability in sample collection for enhanced patient safety, 2) to review quality by design, and 3) to discover how quality by design can help achieve higher quality in blood gas testing.

Sten Westgard is the Director of Client Services and Technology at Westgard QC and an adjunct faculty member at the Mayo Clinic School of Health Sciences in Minnesota, the University of Alexandria in Egypt, and the Kasturba Medical College of Manipal in India. He is an honorary visiting professor at Jiao Tong University in Shanghai, China. Sten has managed the Westgard QC website course portal and blog for well over 20 years. He has written and edited hundreds of reports, essays and applications on quality control, method validation, and Six Sigma among others. And finally, he has contributed numerous books on quality, including the basic QC practices, basic quality management systems, Six Sigma QC Design and Control, CLIA Final Rules, and the Poor Lab's Guide to the Regulations among others.

Sten Westgard: I'm very happy to have the opportunity to present this information and to have the support of Werfen in presenting these educational topics.

I need to make sure you understand which generation of Westgard I am. When the name Westgard is invoked, most often it's my father who people think of. He is the original, the classic, the reference Westgard. He's James O. Westgard. When we talk about Westgard Rules, he developed those. I am the next generation Westgard. I follow the Rules, I teach the Rules, but I did not come up with them, so you cannot blame me whenever they are frustrating you. But that puts me in a unique position to discuss what the next generation of QC will look like.

As we look ahead and take into account all of the changes in diagnostics and in the engineering of instruments, it's the same QC we used 40, 50 years ago, and what we're going to

be using 10, 20 years from now.

Let's start by going back to the original Westgard Rules that were introduced in 1981. What was the first thing they were trying to address? Before the Westgard Rules, if you think about the dawn of the laboratory era, people were using the 1_{2S} rule. That's two standard deviations from the mean. And even in the early days of the laboratory, it was known that the $2SD$, while sensitive, was very good at detecting errors, but was atrocious at generating false rejections. It created a lot of false alarms. Specifically, with two control measurements it was giving a 9% false rejection rate. And with three control measurements, it was giving a 14% false rejection rate.

Maybe it was that the laboratory of the past had less to do, fewer tests on their plate, fewer patients they were addressing, so they had more time to deal with false rejections. But if we talk about today's laboratory, we have a huge menu of tests. We have a massive volume of samples to address. And, we don't have time to add 9% or 14% more effort, particularly when it does not contribute anything useful to patient care. When you walk into work, I'm sure you don't look at your list of things to do and say to your supervisor or your administrator, "I know I have this very long list of things to do. I would like 9% more work, or I would like 14% more work today. And that 9-14% more work, please make it pure waste. Please make it an activity that doesn't help any patient and just drains the resources of the laboratory."

Back in the '70s, as more instruments were becoming more automated with more menu, and higher volumes, there was a need to address and try to eliminate the 1_{2S} false rejection problem. And that's what the first Westgard Rules were about, trying to eliminate all of the false rejection problems from $2SD$. And it took that 1_{2S} rule and instead of making it a rejection, it said, "Okay, if that 1_{2S} alarm goes off, that's a warning. It's just a warning. What you need to do is check five other rules, and only if one of those is violated, will you reject the run. If you look at those five other rules and none of them are violated, you're in control." And that takes a 9 to 14% false rejection rate, and it reduces it to 2, 3 or 4%. It dramatically reduces the number of false rejections.

And at the same time, these other rejection rules, the five other rules after the warning rule, when combined, still provide a relatively high level of error detection. And even better, the

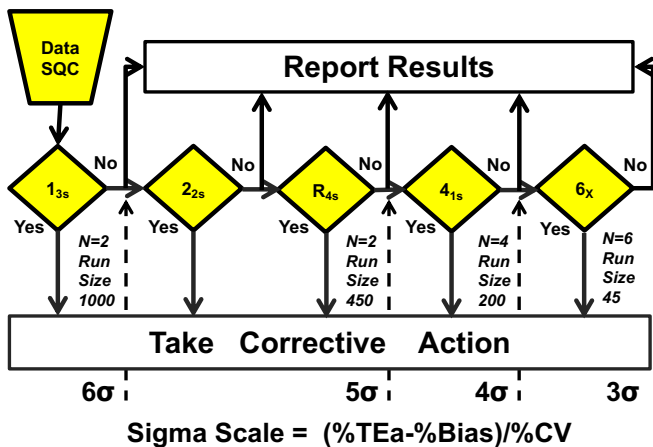
If you would like to participate in this feature, as a company or healthcare provider, please contact Steve Goldstein at s.gold4@verizon.net.

Westgard Rules allow you to use a few measurements and still get relatively high error detection.

Outside the laboratory, in other applications of quality control, if two controls aren't providing enough error detection, you run 10 controls. And if 10 controls aren't giving you enough error detection, you run 20, 30, 40, 50, 100 controls. Because in other applications, and especially in manufacturing, controls are inexpensive. Whereas in the laboratory, we can only afford two or three levels because controls are expensive. So we're trying to deal with our false rejection problem and still maximize our error detection problem with a relatively small number of measurements. This is a classic laboratory workaround. We're trying to do more with less.

That was the original Westgard Rules. They were quickly adopted worldwide because it became clear, "Oh, this big false rejection problem that laboratories are experiencing, we can dramatically reduce it if we switch to Westgard Rules." But that was 1981. Now we are far into the 21st century. Do we still need to use the same formulation of Westgard Rules that we were using four decades ago? Well, the Westgard Rules have not stayed in place. They haven't been carved in stone in 1981. The Westgard Rules have been evolving the Westgard Rules to try to keep up with the changing and the improving quality of performance in today's instrumentation.

The latest version of Westgard Rules is called Westgard Sigma Rules, and that's where we take Six Sigma and combine it with the Westgard Rules. And with Six Sigma, we're trying to benchmark the level of quality. So how good is the method at achieving the level of quality needed by patients? The better it is, the closer that number is to Six Sigma. As you get closer to Six Sigma, you don't need as much QC. With the Westgard Sigma Rules, the flowchart looks similar. The 1_{2s} rule is gone because we no longer recommend it, even as a warning rule.



Westgard Sigma Rules with Run Sizes for the numbers of patient samples between SQC events. Note sigma-scale at the bottom of the diagram. To apply, determine sigma-metric, locate on the sigma scale, identify control rules, total number of control measurements (N), and frequency of SQC events specified as Run Size.

Figure 1. Westgard Sigma Rules

But the biggest change comes in at the bottom. See the bottom of that flowchart (Figure 1). It says Six Sigma, Five Sigma, Four Sigma, Three Sigma, and below that is a Sigma metric calculation. The equation tells you how good you are at achieving the required level of quality. There are three ingredients: 1) the allowable total error (TEa), 2) the bias, and 3) the CV.

The TEa is the quality goal you need to hit. It's often provided by regulations. In the US, CLIA, for most analytes defines how good you have to be. Then bias and CV, are things you operationally calculate already. It's not a difficult calculation because all of the ingredients, all of the variables are already at your fingertips. Once you know that Sigma metric, you can look at the Westgard Sigma Rules and determine how many actual rules are needed.

Six Sigma, which is world-class quality, represents less than four defects out of a million reported results. It's as close to perfect as most processes can get. But if you're at that level where you're almost completely defect-free, that's when you only need one rule. You don't need all the Westgard Rules, you just need one rule, and those limits are even wider.

They go to 3SD. Instead of 9-14% false rejection, now you're less than 1%. There's a benefit to having the best quality test because it allows you to perform the least amount of QC.

And then as your Sigma metric goes lower, which indicates lesser performance, Five Sigma being excellent, Four Sigma being good, and Three Sigma considered the minimum acceptable quality, you need more and more Westgard Rules. At Three Sigma, for instance, you need all the Westgard Rules. The minimum level of quality requires the maximum number of Westgard Rules. So now there's a flexible approach to how many QC rules are needed, depending on how well your method is performing.

But this can still be more advanced, once we start to think about closed box operations. To recap, if we can get to Six Sigma, we won't need all the Westgard Rules. And at Three Sigma we'll need all of the Westgard Rules. That's the minimum quality, requiring the maximum Westgard Rules. And that obviously means when we're in open system processes—systems that are more exposed to human error—we're going to need more Westgard Rules. As the Sigma metrics go down, the more pre-analytical and analytical errors are possible, are occurring, and we'll need more Westgard Rules. But if we have a closed system that can operate at a high level of quality, we may not need Westgard Rules at all.

Let's switch now from the technical application of Westgard Rules and QC to the operator of QC. Who is performing this QC? It would be nice if not only could we customize the QC, but also if we could automate it. A couple of years ago, we conducted a survey of point-of-care users and how they did QC, and more importantly, what they would like out of their point-of-care devices. Some of them said, "Well, we'd like more training on the device." Others said, "We need more directions from the manufacturer on how to run QC."

But the most common thing asked by users is, "Can you put more of the QC system inside the device?" Take more out of the hands of the user and automate it inside the device itself. That parallels many processes we see today in other industries. For instance, cars and car safety, we've seen a huge number of innovations in the area of helping cars not crash and helping cars minimize the impact of a crash. This is part of their QC system. They have lane and distance detectors.

I'm old enough to remember when automatic brake systems were new. The original Westgard was old enough to remember when seat belts didn't exist, much less airbags. But over the

decades, new systems keep being added. Many are automated and they take the safety part out of the hands of the driver. Because if you look at who causes accidents, few of them are caused by the car. And few of them are caused by the environment. Over 90% are caused by the driver. That's why there's this drive to make drivers driverless, to make the next car, completely automated and take the driving out of the hands of the human.

A year or two ago, we thought that we were going to see the end of driving. Now we're a more skeptical, but there is more automated driving. And we're at the point where you might look at yourself and if you have children, you might look at them and think, "Are my grandchildren going to be the first generation who doesn't drive or are they going to be the last generation who drives?" Basically, we're going to turn driving into a closed system that doesn't suffer from human error.

Let's pivot again and talk about quality in an abstract level. What are we trying to manage? Why do we run QC? Why don't we accept the device and say, "Well, they said it wasn't going to have errors, so I'll accept that and I don't need to do QC." The first thing you'll encounter is the manufacturer, people who sell you the device as well as the regulators of your operations, they will tell you, "No, no, you have to run QC. You cannot assume everything is going to be fine." And then as you go through your training, of course, you place at the center of your focus the patient's health. "I need to make sure every time I'm delivering the right answer, I'm not going to harm the patient."

Now, luckily, we're in a position where we can afford some QC, and we can see a benefit for ourselves when we do it correctly. There's also not an insignificant reason, that when we do QC, we get useful information. So if we have our Levey-Jennings charts, if we're plotting values, we not only can catch immediate errors, we can even see errors as they develop. We may even be able to anticipate problems and prevent them from occurring.

That's in the ideal application of QC. There are still laboratories all over the world that do QC as theater. They look at the regulations and say, "What's the least amount of effort I can make? Do you want me to run some controls? Well, I'll run some controls. Do I need a chart? Okay, I'll draw some lines. I'll throw some dots on the chart, but I'm not really paying attention to them." But we want to do something useful with these dots. And most relevant at the point of care is, we have an immediate answer and an immediate application. If that's a single result, it needs to be reliable. We need to know if it's reliable. We need to know, can we make a decision? Can we change a medical diagnosis or a therapy at that very moment?

To support that, if that is our goal, we need a QC that will work immediately.

There's technical reasoning that goes behind the powering of immediate QC. One of the things we have to do is define all kinds of sources of errors. What could go wrong? Then we need to say, "Well, for every kind of error, we also need a way to monitor it." And not only do we monitor it, but we have to define, "How much error is acceptable and how much is unacceptable?" Because some little amount of variation is going to happen. There is no method, no instrument, no process that is perfect. There is variation and it exists.

Some variation can be tolerated. And if it's small enough, it doesn't change the final result and we can look at that and say, "Yes, there's some variation, but we're not going to raise an alarm over it." That's part of the Six Sigma theory. What we need to do is define the tolerance specification.

Let's figure out, "This is the size error that is acceptable and isn't going to change the diagnosis, isn't going to hurt the patient." But we're going to add another line and say, "Look, any error that goes over here or over there, that's too much error. That could flip a diagnosis. That could give a false positive, or a false negative. The patient isn't going to get the right diagnosis or the right treatment." Those are the things we need to prevent. So there's a technical way to do that. We can determine the total allowable error. And then we can model what happens to a method when an error occurs. What's going to be our ability to detect that error?

On the left of the screen below (Figure 2), is the way we want things to occur. This is the setup. There's no error. The distribution of results is centered where it should be and the limits are wide, so there's little false rejection. We're going to catch any error that occurs. And then on the right you see, "If an error does occur, if the error is of this size, we can look at the area under the curve between these two lines." We can say, "Oh, we've got 97%, that's a 97% chance we're going to detect that error with very little false rejection." There's a mathematical way to figure out how good we need to be, and when things go wrong, how likely we are to be able to detect that error.




Figure 2. Defining Quality and Methods for Control: The Gritty Details.

We can take that error detection capability and turn it into an average run length, which tells us how quickly we will detect the error. You'd like your average run length to be 1 or less. If your average run length is 1, then you detect the error when it occurs. The same run, same time that the error occurs, you're going to detect it. But if your average run length goes up and up and up, it gets larger. If it's 2, then it says, "Ah, well, on average it's going to take me two runs before I detect the error. That means one run might go out the door and be acted upon and be used for diagnosis before I'm able to detect the error."

We want that to be less than one, but then we come to our definition of what a run is. Is a run a single result or is it 24 hours? If a run is 24 hours long, then a day could pass before we catch an error. And for some of those results, you don't hold them for a day, you act on them quickly.

Even having an average run length of 1, depending on the circumstances of how you're running QC and how you're defining the run, that could be too late. That's too long.

		pH	pCO ₂ (mmHg)	pO ₂ (mmHg)	Sodium (mmol/L)	Potassium (mmol/L)	Calcium (mmol/L)	Chloride (mmol/L)	Glucose (mg/dL)	Lactate (mmol/L)	tHb (gr/dL)	O ₂ Hb (%)	COHb (%)	MetHb (%)	HHb (%)	tBili (mg/dL)
PCS A (N>45000)	Mean	6.90	64	113	107	7.1	1.84	46	144	3.3	14.6	89.4	2.4	1.6	6.6	20.1
	SD (CV%)	0.003 (0.0%)	0.5 (0.8%)	1.1 (1.0%)	0.4 (0.4%)	0.03 (0.4%)	0.015 (0.8%)	0.2 (0.5%)	2.1 (1.5%)	0.06 (1.7%)	0.03 (0.2%)	0.01 (0.0%)	0.01 (0.4%)	0.00 (0.3%)	0.00 (0.1%)	0.03 (0.2%)
PCS B (N>53000)	Mean	7.41	33	181	156	2.0	0.79	85	0	0.0	0.0	N/A	N/A	N/A	N/A	0.0
	SD (CV%)	0.004 (0.1%)	0.5 (1.4%)	2.2 (1.2%)	0.7 (0.5%)	0.01 (0.5%)	0.007 (0.9%)	0.4 (0.4%)	1.6 (N/A)	0.03 (N/A)	0.04 (N/A)					0.04 (N/A)
PCS C (N>6000)	Mean	8.05	33	3.1	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A	N/A
	SD (CV%)	0.007 (0.1%)	0.3 (1.0%)	1.0 (32.5%)												
PCS D (N>13000)	Mean	7.35	24	48	165	7.3	1.22	142	349	8.0	7.4	80.0	4.1	4.0	11.9	10.4
	SD (CV%)	0.004 (0.1%)	0.4 (1.5%)	3.2 (6.6%)	0.6 (0.4%)	0.04 (0.6%)	0.011 (0.9%)	1.6 (1.1%)	3.3 (1.0%)	0.24 (3.0%)	0.03 (0.4%)	0.14 (0.2%)	0.09 (2.1%)	0.10 (2.6%)	0.33 (2.8%)	0.03 (0.3%)
PCS E (N>13000)	Mean	7.22	69	92	129	4.5	0.56	101	71	1.6	16.5	50.0	10.1	8.1	31.9	20.0
	SD (CV%)	0.005 (0.1%)	1.0 (1.5%)	2.2 (2.4%)	0.5 (0.4%)	0.03 (0.6%)	0.008 (1.4%)	0.8 (0.8%)	1.0 (1.5%)	0.06 (4.1%)	0.07 (0.4%)	0.11 (0.2%)	0.05 (0.5%)	0.07 (0.9%)	0.23 (0.7%)	0.07 (0.4%)
Sigma Average		9.0	11.2	6.9	8.3	25.2	12.4	10.5	7.0	8.5	18.4	122.2	70.1	128.3	209.8	43.0
Overall Pfr		0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.000	0.152	0.152	0.152	0.152	0.152	0.167
Overall Ped		0.974	0.953	0.938	0.412	0.979	0.990	0.950	0.097	0.933	0.997	1.000	0.957	0.957	0.998	0.923
Overall ADT (min)		2	2	2	5	2	2	2	21	2	2	2	2	2	2	2

Table 1. Method Sigma.


Now we're going to look at the actual performance that we see in the GEM® Premier™ 5000 blood gas system with Intelligent Quality Management 2 (iQM® 2). The table is dense (Table 1). Here are some points to anchor on. First, the columns represent different parameters that are going to be measured, from CO₂ to sodium, chloride, glucose, and then you have the five different process control solutions (PCS), that are run automatically inside the system at various times.

The last five rows at the bottom (Table 1) have the most important details. What's the Sigma that's being achieved by these methods? And if you look at the Sigma average across that row, they are all higher than six. We have better than Six Sigma performance on the whole battery of tests. What does that mean when it comes to average time of detection (ADT), which is essentially our run length, our average run length, how quickly can we detect errors? For most, it's two minutes. It's going to take us two minutes to pick up an error instead of the 24 hours traditional laboratory might be doing.

Another row in the bottom to take a look at is overall false rejection rate (Pfr) What is the false rejection rate that these QC techniques inside the box are going to generate? For most of them, it's zero. Or I should say it's 0.000. If we added more digits, we might see a little more false rejection.

But it's compared to what we see outside the box, in traditional laboratories. There are laboratories still doing 2SD and they've got 9% or 14% false rejection. And there are labs that might be using Westgard Rules and they've got it down to 2-4%. And then there are Six Sigma assays that might use the Westgard Sigma Rules, and they might use 3SD, which gets down to less than 1%. And then you have the GEM Premier 5000, which is running five different levels of QC at different intervals, and false rejections down to essentially zero. They aren't going to be bothered by false rejections.

To summarize Table 1, detection time is as small as two minutes for most analytes. And the average of the Sigma metric is greater than six for all analytes. You get world-class quality with very rapid error detection.



decision paper



free chart tool

Westgard JO. A method evaluation decision chart (MEDx chart judging method performance. Clin Lab Sci. 1995;8(5):277-283.

Figure 3. Visualization of Sigma Performance: Method Decision Charts. Westgard JO. A method evaluation decision chart (MEDx chart judging method performance. Clin Lab Sci. 1995;8(5):277-283.

Another way to look at the Sigma metric, other than calculating it, is to use a Method Decision Chart (Figure 3). The QR codes will take you to the original papers as well as a free tool that will generate these charts for you. Method Decision Charts provide a visualization of the Sigma metric, and it takes the same information that we used in the equation and displays it as a plot. Now we take the imprecision, on the X-axis, and the bias on the Y-axis. If you know the CV and bias of your methods, those are the X-Y coordinates. You now know where to plot your performance. And on this chart, are these diagonal lines, carving up that space into different Sigma zones. At the graph's origin near the zero point, that's world-class or Six Sigma performance. What is that? That's less than four defects per million reported results.

And then you have Five Sigma, which is excellent. That's about 233 defects per million. Four Sigma is around 6,200 defects per million. Three Sigma, the marginal quality or the minimum acceptable quality, is around 67,000 defects per million. Two Sigma is poor. Less than Two Sigma, we hope you'll consider unacceptable. And that's where we get into hundreds of thousands of defects per million reportable results.

If you want to simplify your thinking about this chart, those diagonal lines are the rings of a target and you're trying to hit the bull's-eye. The closer you get to the bull's-eye, the fewer defects you have, the more reliable your result is, the more confidence you and the laboratory and the clinician and the patient can have in that result and any decision being made, based on that result. As you have more imprecision, more bias, you're generating more defects, you're adding more noise around all of the signals. And in the upper right quadrant, you're probably confusing the clinician. You're not helping to confirm the diagnosis.




Figure 4. Cut to the Chase: Method Decision Charts Make it Clear.

We put that Method Decision Chart to use in hundreds of studies and laboratories around the world. One study that's particularly relevant to our topic today is from an institution in Belgium. They performed a study looking at three blood gas instruments. It's rare to have a head-to-head, apples-to-apples comparison and method performance. Here was an institution that had access and was able to perform a quality assessment of three different blood gas instruments. And when you look at the RapidPoint, you see seven tests where they had three levels of control (Figure 4). A handful of them are green or in the bull's-eye, and that's Six Sigma. Two are minimum quality, one that's Two Sigma. And two that are less than Two Sigma.




Figure 5. Cut to the Chase: Method Decision Charts Make it Clear.

When you look at the Radiometer (Figure 5), the performance gets worse. Yes, there are nine tests, 27 levels. You see a number of assays in the bull's-eye, but it's concerning that you have six levels below Two Sigma. Those could be generating errors. And if you're only doing your QC once every 24 hours, you probably have impacted results going out to clinicians and patients before the QC detects the error and you're able to address them.




Figure 6. Cut to the Chase: Method Decision Charts Make it Clear.

And then finally, we have the GEM Premier 5000 (Figure 6). Again, another nine tests, three different levels. It's not a perfect box. There is no perfect box, but you have the vast majority of performance in the bull's eye. Of the three, the most number of levels in the bull's eye happens with the GEM Premier 5000.

How is that performance achieved? How are you able to do better QC on better performance? It comes down to the closed analytical system. You're taking things out of the hands of the users, which they appreciate. Less work for the end user, but by controlling it within a closed system, you also make it perform better. And in the figure (Figure 6), you can see all the different process control solutions. Remember, the system uses five different QC solutions, and the user does not have to prepare and insert five different levels. That would create an opportunity for more pre-analytical mishandling. And on top of the closed system, there's a pattern recognition software, which compares the current observations to essentially a database of all previous experiences with this instrument and this method and matches them up. "Oh, the way your performance looks right now, it matches historically to an error condition we've seen in the past." It can help catch errors that no other system is able to detect.

Then there's the frequency of QC. When we talk about 24 hours as the traditional QC frequency, that's an arbitrary choice. Why is 24 hours so magical? It's where the regulations and even the manufacturers of most devices came down. We don't know how often you should perform QC, so we'll say once a day.

If you venture into the United States, the CLIA regulations, there is a more extreme option called IQCP, the Individualized Quality Control Plan. If you do an IQCP, you no longer have to base your QC frequency on the earth. You can base it on the moon. That is, you can reduce your QC frequency to once every lunar cycle, once every month. A once a month QC frequency is literal lunacy. If you do that, it's crazy. If you have an error and you catch it one month later, how do you go back to patients and say, "Yeah, 20 days ago, three weeks ago, we ran a test on you, and now we just realized that it's wrong?"

So those are two ways that QC frequency occurs. There is another one that is driven once every eight hours, that's almost driven by the shifts. You have three shifts? You have the morning, you have the noon, you have the overnight. Every eight hours, you have every shift run QC. Interesting approach.

You're basing it on your personnel, but what should be the heart of our focus for QC? It should be the patient. We should have patient-based, performance-driven QC frequency.

What do we need to do to make sure the result we give the patient is not wrong? How do we eliminate or reduce to the number of erroneous results we're delivering to patients? That theory has been calculated and developed. One approach that has become more popular, at least to talk about, is patient-based real-time QC, or PBRT-QC. This gets offered as well. We don't need to worry about traditional QC because we'll do patient-based QC. And patient-based QC has the benefit of using patient samples. You don't have to spend extra money like with traditional controls.

But even PBRTQC requires that you accumulate a certain number of typically normal patient samples. It might be 10, it might be 20. If you work out the math, sometimes you need several hundred patient samples before you can calculate an average of normals and then be able to track the performance of the method. Your PBRTQC might give you signals faster than every 24 hours, but it's still behind the clock for an individual patient's sample.

The most patient-based QC would be to look at a single patient specimen and only focus on measurements on that specimen. This is one area where the GEM Premier 5000 excels. It takes 15 measurements in 15 seconds on a single patient specimen. That's the most patient-based QC of all. And then it will compare those measurements against, again, like a database of all well-performing tests, as well as the erroneous ones. And it'll match up and say, "The trajectory of that test in those 15 measurements in 15 seconds, looks like an error we've encountered in the past. We have to be careful about this particular test result." Which gives you a big contrast, because now you have a QC system that is executing on every patient sample. Not only is it doing pattern recognition, called IntraSpect, but it also has an interference detection. It detects micro-clots. This is in contrast to our traditional QC system, which may be happening only every eight hours, or only every 24 hours.

There might be a few additional electronic checks that a manufacturer promotes. "Well, we do traditional QC and we have some electronic system checks." But they're not looking at individual patient samples. This additional approach that the GEM Premier 5000 has taken expands the number of errors it's able to cover and monitor. Our traditional QC has been focused on the analytical step. But with the distribution of errors in the total testing process, many are happening at the pre-analytical level. And then a lot of them also occur in the post-analytical era phase of the total testing process. To provide a better QC system, we would like to use something like the GEM Premier 5000, that will encompass the pre-analytical, post-analytical, and the analytical phases of the total testing process.

One example is bubbles. In one paper, bubbles were introduced into samples and then different devices were tested to see if they could detect these micro-bubbles in the samples themselves. And the GEM Premier 5000 rejected all of them. It was able to effectively respond, "There are micro-bubbles in here and we can't use these samples."

Another device, the ABL 90, had different error messages. The system didn't recognize what was going on. It flagged results

with strange errors and reported results that were impacted by those bubbles. Here's an area where that pre-analytical error, introduction of bubbles into the specimen, is only addressed and detected by the GEM Premier 5000, and its predecessor, the GEM Premier 4000.

Now let me wrap up the discussion. We've tried to show you that the next generation of QC is here. It's present in the GEM Premier 5000, with iQM[®] and IntraSpect[™], and the additional systems Werfen has brought to the instrument. It has a wider reach of monitoring.

It looks beyond the analytical phase. It has novel ways to monitor quality. It has an unmatched ability to detect and to perform QC on each patient sample. And it will catch errors that traditional QC cannot. Approximately one out of every 100 patient samples has an error that only iQM2 can detect, and other devices will miss those errors.

And with that, I thank you.

Disclosure. Sten Westgard received a consultant fee and a speaker fee for this presentation from Werfen.