QP
QUALITY PROGRESS
THE OFFICIAL PUBLICATION OF ASQ

FILL IN THE GAPS

DEC 2022

ASQ's annual salary survey gives you the tools to paint yourself into the future you want in the quality profession

ASQ
Excellence Through Quality®

QUALITYPROGRESS.COM

CONTENTS

December 2022 VOLUME 55 NUMBER 12

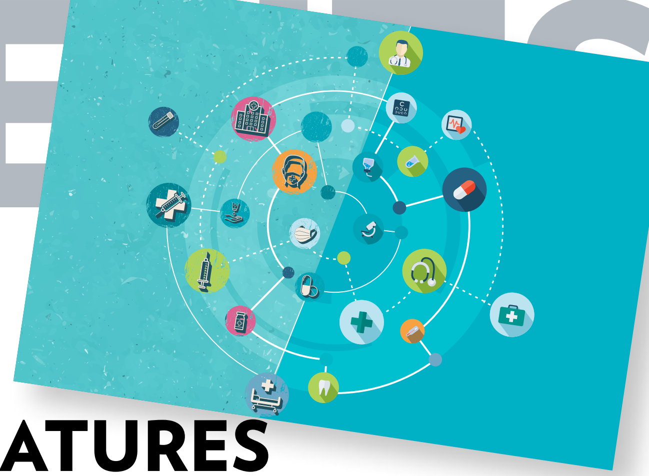

COVER STORY

FEATURES

DIGITAL EDITION

QP is now available in a convenient digital format to make it easier to access the quality information you need on the go. Make sure to log in to view the link.

digital.asq.org/qualityprogress/library

A BETTER FRAMEWORK

JUST THE FACTS

Many managers make projections based on single-point estimates, such as average revenue, average cost or average defect rate. But complex decisions with broad potential consequences require more in-depth estimates that consider all available data.

The Markov Chain Monte Carlo (MCMC) method is a more thorough and accurate technique used to analyze data that allows you to extract more insights than with a single-point estimate technique.

The authors present a theoretical case study in which a manufacturer uses MCMC to analyze the monthly incoming quality cost data for two suppliers to determine whether the manufacturer can save money by single sourcing its axles.

A technique for analyzing random variables by Mark Fiedeldey and Ray Harkins

Business today is more competitive than ever, so successful business leaders often must make quick decisions with incomplete data. The wrong decision could result in significant losses, layoffs or worse. This is the point at which quality professionals and other data-savvy specialists can offer assistance—by making the best analysis possible given the available data.

Without a deeper understanding of analytics, most managers make their projections based on single-point estimates, such as average revenue, average cost or average defect rate, for example. After all, the arithmetic average of any of these values is generally considered the best guess to use for projections. But complex decisions with broad potential consequences deserve more than "back-of-the-napkin" estimates. They require the use of all available data, not just their averages.

Consider the following hypothetical case study.

Monster Machines

Monster Machines ("Monster"), a manufacturer of custom off-road vehicles, began purchasing axles six months ago from Axles "R" Us. Axles "R" Us is struggling to meet its production demand, resulting in Monster sourcing additional axles from A1 Axle four months ago. However, the quality coming from A1 Axle is lower than expected. As a result, Monster's quality manager, Mike Rometer, wonders whether an investment in additional workstations at Axles "R" Us, which would allow it to meet the full production demand, is a good decision for Monster.

Given the various costs and savings associated with sourcing axles from a single supplier, Monster's CFO, Penelope Pincher, estimated that a monthly savings of $750 on incoming quality is needed to justify eliminating A1 Axle. Rometer must determine whether Monster could save $750 per month by sole sourcing the axles from Axles "R" Us.

He tallied the monthly incoming quality cost records for both suppliers. The results are shown in Table 1 (p. 76).

Because historical records showed no seasonality with the production of the vehicles, the simplest comparison Rometer can perform is to examine the average monthly quality costs. The average monthly cost for Axles "R" Us is $2,206.95, and the average monthly cost for A1 Axle is $3,231.74.

With Monster's proposed investment, Axles "R" Us would be able to meet the production demand without the need for a second supplier. Axles "R" Us also has guaranteed its quality costs will not increase even with the added production. Thus, Monster's proposed investment coupled with Axles "R" Us's quality commitment results in an average quality cost savings of $1,024.79 per month. This is well above Penelope's threshold of $750 per month, and therefore, looks like the obvious choice.

Each supplier's average quality cost, however, is a mere single-point estimate of its incoming quality. While a valid estimate, the arithmetic average isn't complete enough for use in critical decisions because it ignores any data variation.

Rometer decided to perform a more in-depth analysis of the quality data beginning with determining its underlying probability distribution. While random variables often can be modeled as normally distributed, more detailed analyses require verifying the underlying probability distribution before proceeding.

To better understand the spread of the cost data over time, Rometer started his new level of analysis by first calculating the standard deviation of the quality cost data for both suppliers. The standard deviations are $166.99 and $117.44 for Axles "R" Us and A1 Axle, respectively. He used the Shapiro-Wilk[1] statistical test on the data, which allowed him to infer that, in fact, the quality costs could be modeled with a normal distribution.

With the arithmetic average and standard deviation for each, Rometer subtracted Axles "R" Us's distribution of cost data from A1 Axle's distribution of cost data, resulting in a new normal distribution of monthly savings data.

Subtracting normal distributions involves two steps for noncorrelated variables such as these:

1. Subtract the means to derive the mean of the resulting normal distribution.
2. Take the square root of the sum of the squared standard deviations to derive the standard deviation of the resulting normal distribution.

When Rometer stepped through this process on his cost data, he arrived at the results shown in Table 2.

Monster's monthly savings distribution has an average of $1,024.79 and a standard deviation of $204.15. Entering these new distribution parameters into Excel's norm.dist formula,[2] Rometer estimated that the probability of saving at least $750 is about 0.91. He also generated the chart in Figure 1 showing the cumulative estimated monthly savings distribution.

Rometer's analysis showed the staff at Monster that single sourcing its axles from Axles "R" Us could save the company more than $750 per month, but it's not the sure bet suggested by his analysis using the averages only. Understanding that Monster's risk of failing to exceed the $750 threshold is about 9% may influence its decision.

Rometer knew that his second analysis, while intricate, used more of the available data and offered an estimation of risk unattainable by simply examining averages. He also knew that the second analysis depended on the accuracy of each supplier's distribution parameters. After all, repeated sampling to estimate these same parameters would generate slightly different results each time.

In fact, random variables can be modeled with an infinite number of normal distributions, each with a slightly different set of parameters. Some models would be better fitting than others, but there is no way to determine which pair of parameters, if any, describe the "true" distribution.

These various pairs of possible parameters form their own distribution—a joint probability distribution—that analysts can use to account for uncertainty in the parameter estimates.

TABLE 1

Monthly incoming quality cost

Month	Axles "R" Us	A1 Axle
1	$2,285.09	NA
2	$2,136.09	NA
3	$1,927.29	$3,315.28
4	$2,167.75	$3,121.23
5	$2,364.41	$3,349.67
6	$2,361.09	$3,140.77

TABLE 2

Distribution of savings

	Mean	Standard deviation
Axles "R" Us	$2,206.95	$166.99
A1 Axle	$3,231.74	$117.44
Difference formula	$\bar{X}_{A1\ Axle} - \bar{X}_{Axles\ R\ Us}$	$\sqrt{\sigma^2_{A1\ Axles} + \sigma^2_{Axles\ R\ Us}}$
Monthly savings	$1,024.79	$204.15

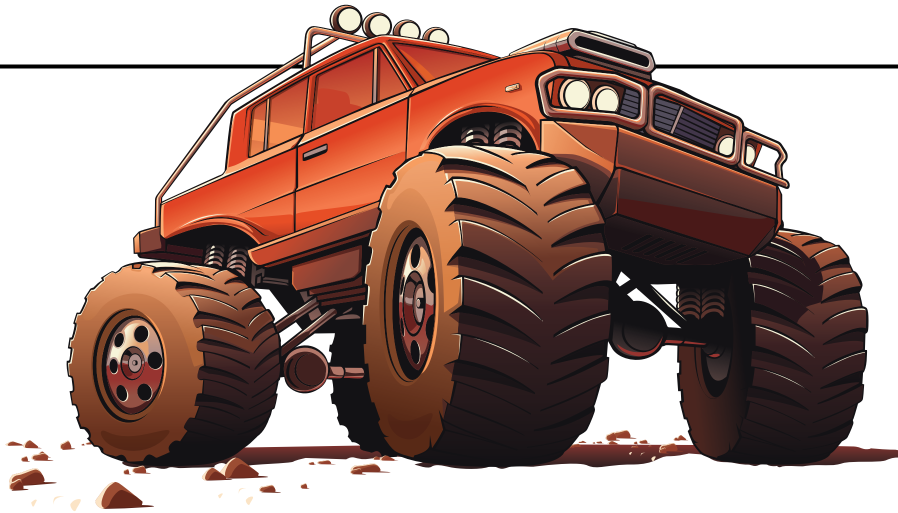

FIGURE 1

Estimated monthly savings

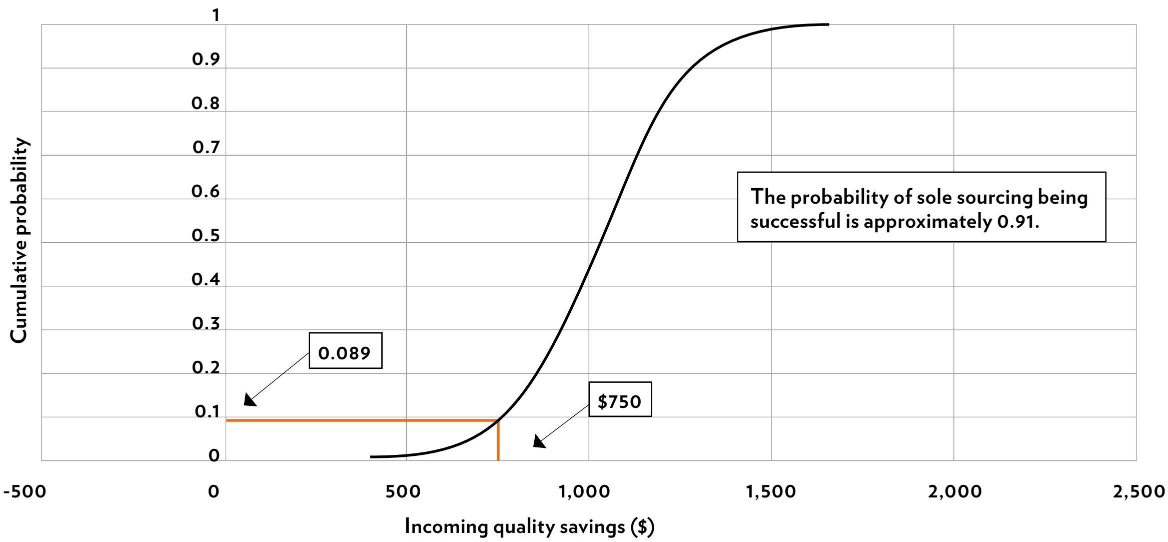

Unlike common probability distributions—such as normal and exponential—that have equations to describe them, there is no equation to describe this joint probability distribution of parameters. The only option is to use numerical techniques to generate a sample of the distribution. With sufficient samples in a numerically generated estimate, the uncertainty in the parameters can be adequately estimated.

One class of algorithms called the Markov Chain Monte Carlo (MCMC) method is commonly used for this purpose. The specific algorithm used in this example is called the Metropolis-Hastings.[3]

When Rometer applied the MCMC method to the incoming quality cost data for Axles "R" Us and A1 Axle, he generated the joint mean/standard deviation distributions in Figures 2 and 3 (p. 78), respectively. Each distribution was based on 750,000 samples.

The first detail to note in these estimated joint probability distributions is that the single-point estimates of the parameters, namely mean and standard deviation, are right in the heart of the MCMC sampled distribution. This should make Rometer feel confident his earlier analysis was valid.

The second interesting detail is the range of possible pairs of parameters that could model the cost data. For example, normal distributions with means between $2,000 and $2,400 and standard deviations between $100 and $400 appear from the graphics to be useful model contenders for the Axles "R" Us data. Perhaps the wide range is primarily a consequence of having only six data points. But the fact remains that the uncertainty in these parameters is quite large. To reduce this range of values, Rometer may choose to postpone his decision to single source axles until he can collect more data. Nonetheless, this is the best summary he can provide given the current data.

FIGURE 2

Axles "R" Us joint distribution of normal parameter pairs

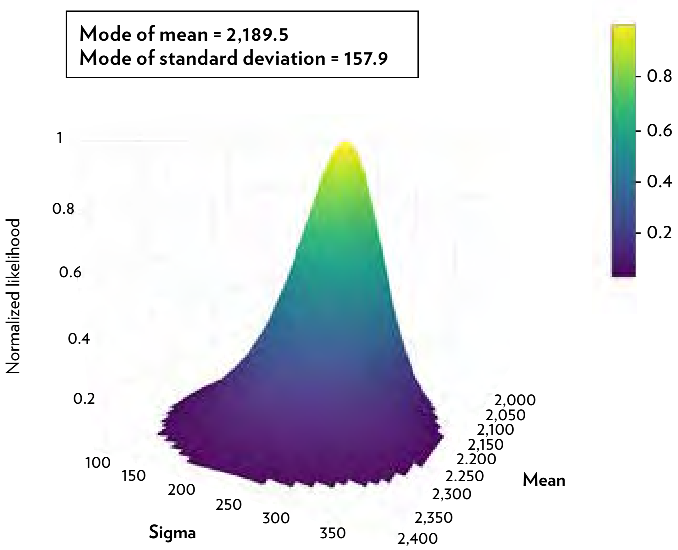

FIGURE 3

A1 Axle joint distribution of normal parameter pairs

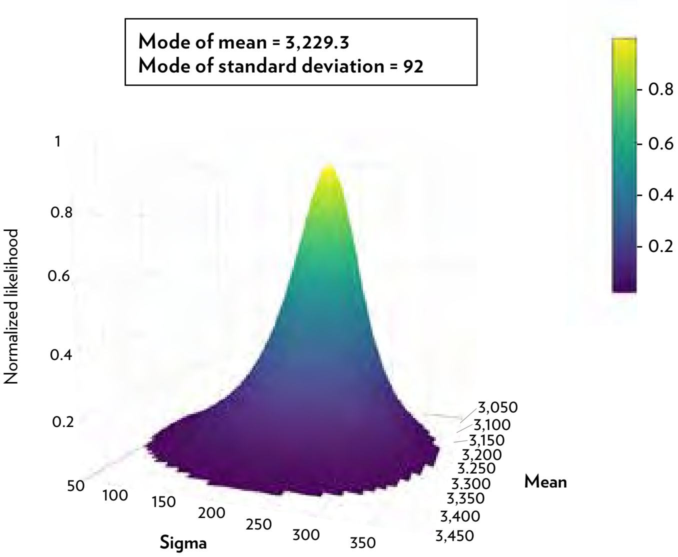

Based on this joint probability distribution for each supplier, Monster can estimate the future monthly incoming quality cost distribution for each supplier in a manner that includes the uncertainty in the distribution parameters. The monthly cost difference between the two suppliers can be simulated and the savings estimated. The predictive distribution of simulated monthly savings, based on 25,000 simulated monthly costs, is shown in Figure 4.

Rometer's newest analysis, accounting for the uncertainty in each supplier's cost distribution parameters, raised the risk of not meeting the threshold of $750 monthly savings from 9% to 17%—a significant enough shift to potentially influence Monster's decision.

This case study illustrates that basing decisions on analyses using point parameter estimates ignores significant sources of uncertainty, potentially resulting in bad decisions. The first analysis, which was a point estimate average accounting for no uncertainty, suggested sole sourcing to Axles "R" Us was an absolute win. The second analysis, which accounted for uncertainty in monthly incoming quality costs but used point estimate parameter values, suggested a risk of failure of about 9%. The final analysis, which accounted for the uncertainty in monthly costs as well as the uncertainty in the distributional parameter values, raised the risk of failure to about 17%.

Forging process

In an actual quality-related case familiar to the authors, engineers were asked to examine capability data from a forging process using a similar analytical approach. A manufacturer was hot forging the rough shape used to make the inner rings for an auto wheel bearing. The forged rings were CNC turned to form their final shape used in the bearing assembly. Therefore, the dimensions of the forged components were critical to the subsequent turning process. Too much material on the forged components resulted in premature tool wear in the turning process. Too little material resulted in "nonclean-up" surface defects on the final part.

Thirty samples of the forged rings were randomly drawn from an initial production run of 500 pieces. The overall width of these samples was measured on a coordinate measuring machine and the data were analyzed. The results are shown in Figure 5.

FIGURE 4

Estimated predictive monthly savings accounting for parameter uncertainty

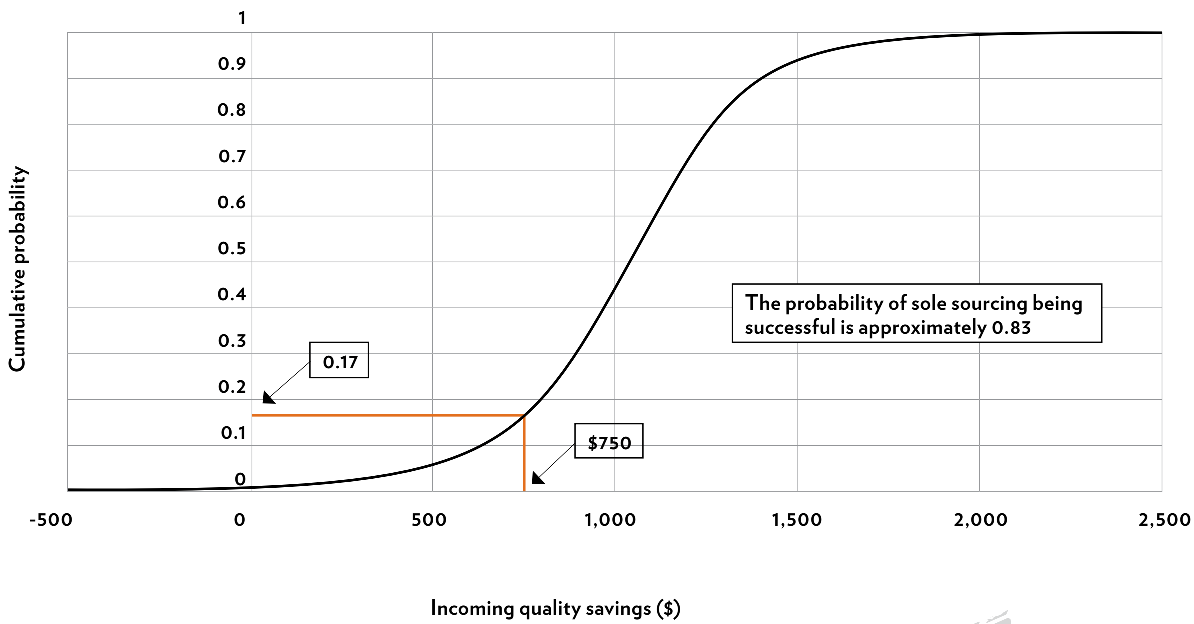

FIGURE 5

Data analysis of 30 samples

32.136	32.044	32.161	32.069	32.176
32.138	32.206	32.068	31.941	32.059
32.134	32.108	32.039	32.032	32.057
32.167	32.078	32.135	32.06	32.149
32.138	31.965	32.065	32.1	32.073
32.12	32.14	32.139	32.12	32.058

Count	30
Max	32.206
Min	31.941
Mean	32.0958
StDev	0.0599
P_{pk}	1.65

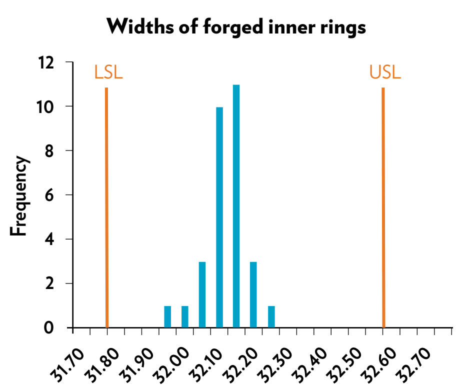

LSL = lower specification limit
USL = upper specification limit

FIGURE 6

Estimate of individual and joint distributions of the parameters

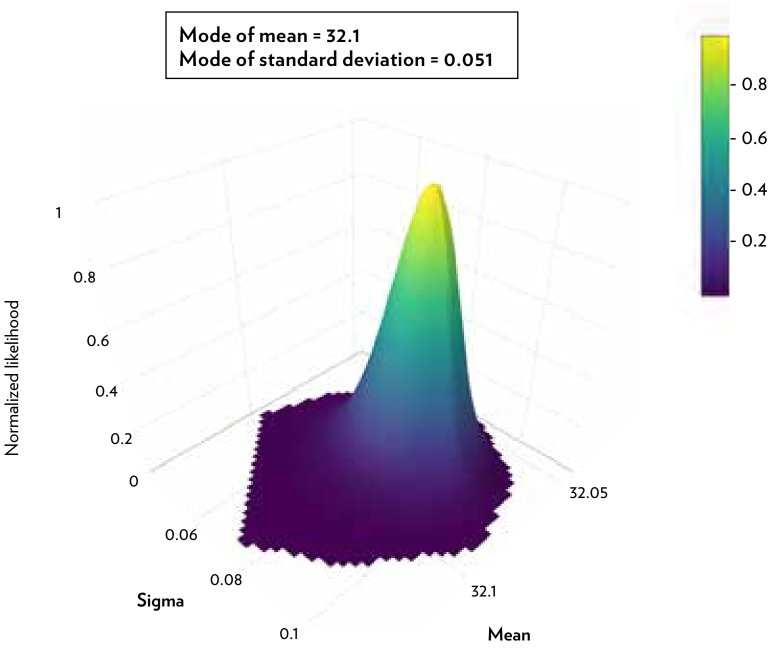

The engineering specification for the width was 32.2 +/- 0.4 mm. Using the standard capability index formula results in a P_{pk} of 1.65. Observing the histogram, the data are slightly shifted toward the lower specification limit (LSL). Using the Excel formula norm.dist (x, mean, standard deviation, cumulative), the expected parts per million (PPM) below the LSL equals 0.39.

These results, like the earlier supplier quality cost example, are based on single-point estimates of the process distribution and the assumption that future production will produce parts with the exact same distribution parameters. Ignoring the inherent uncertainty of these parameters and assuming the future is as certain as the past builds into our conclusions an optimistic bias.

Given that historical data have shown that the manufacturing equipment used to produce parts such as these has variation that is well modeled with a normal distribution, the same MCMC technique employed for the Monster Machines example can be applied to estimate the individual and joint distributions of the parameters. The results are shown in Figure 6.

Note that the mean of the 30-piece sample measurements is nearly identical to the mode of the MCMC distribution of means. Likewise, the standard deviation of the 30-piece sample data is nearly equal to the mode of the MCMC distribution of standard deviations.

The parameter distributions shown earlier represent the parameter values for a normal distribution from which the sample data are likely to have been derived. So, the single point estimates found in Excel are not necessarily wrong—they are just one of many potential sets of parameters for the underlying distribution. Because there is no way to know the true parameter values, all potential parameter sets that are likely to be true should be used to generate the results.

Using these parameter distributions, you can generate an estimate for the distribution of future production. The cumulative distribution function (CDF) of that predictive distribution is shown in Figure 7.

From the predictive CDF, a large sample of measurements is drawn and used to calculate the mean, standard deviation, resulting future capability index and PPM defective (Online Table 1, found on this article's webpage at qualityprogress.com). The P_{pk} found in Excel was 1.65. The P_{pk} for future production found by considering model parameter uncertainty was 1.56. The simulated PPM defective is about 17 compared to the Excel result of 0.39. The risk of producing 17 defective PPM may be sufficient to investigate the process and implement actions to reduce this risk.

Monster can estimate the future monthly incoming quality cost distribution for each supplier in a manner that includes the uncertainty in the distribution parameters.

FIGURE 7

Predicted inner ring width

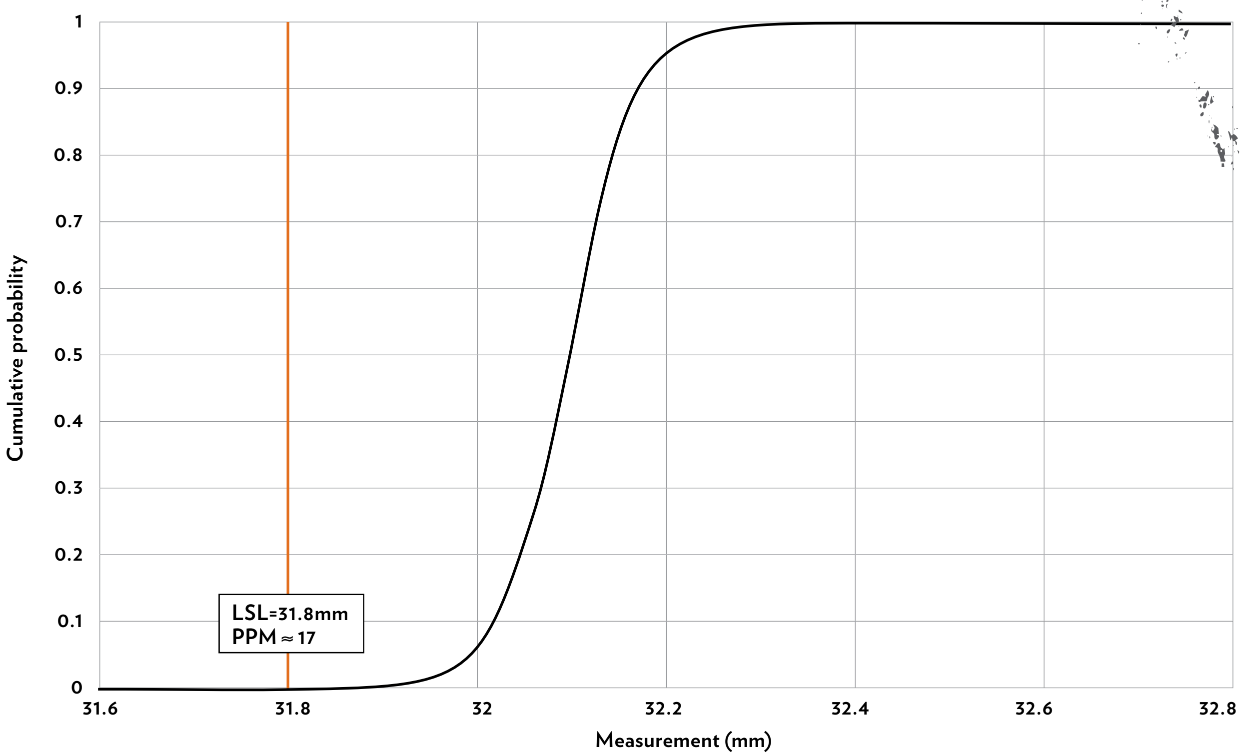

LSL = lower specification limit **PPM** = parts per million

Whether the manufacturing process is deemed qualified is not the analyst's decision. The analyst's role is to provide results using the best techniques available, and—as we have seen—single-point estimates ignore the inherent uncertainty in distributional parameters.

As quality engineers and data analysts, decision makers often look to us to capture and crunch the data available to make major decisions. Analytical tools, such as MCMC, allow us to extract far more insights than we could with the simpler, single-point estimate techniques. **QP**

NOTES

1. The Shapiro-Wilk test (swtest) for normality used was an Excel add-in function swtest (data, original, linear) downloaded for free from Charles Zaiontz, "Shapiro-Wilk Original Test," Real Statistics Using Excel, https://tinyurl.com/3rwrv2d7.
2. Excel's norm.dist formula is =1-norm.dist (threshold, average, standard deviation, cumulative yes or no).
3. The R code for the Metropolis-Hastings algorithm was downloaded for free from Mark Powell at www.attwaterconsulting.com and modified for each case study. The R programming language (also free to download) has a package called MHadaptive, which includes a function called Metro-Hastings.

Mark Fiedeldey is a senior reliability engineer. A member of ASQ, Fiedeldey is an ASQ-certified reliability engineer and quality engineer.

Ray Harkins is the quality and technical manager of Ohio Star Forge Co. in Warren, OH, and teaches quality and business-related courses online at www.themanufacturingacademy.com. He received a master's degree in manufacturing leadership and project management from Rochester Institute of Technology in New York. A member of ASQ, Harkins is an ASQ-certified quality engineer, quality auditor and calibration technician.