LAB MANUAL (INSTRUCTOR VERSION)
An Interactive VBA Tool for teaching Statistical Process Control (SPC) issues

Jaydeep Balakrishnan and Sherry Oh

Note: This exercise is to be used with the open source files posted in : Balakrishnan, J. and Oh, S.L., "An interactive VBA tool for teaching Statistical Process Control (SPC) and process management issues", *INFORMS Transactions on Education*, 5, 3, May 2005. http://dx.doi.org/10.1287/ited.5.3.19

The following are the concepts covered in this exercise:

1. Process capability
2. The role of reduced variation in ensuring better process capability
3. Upper Specification Limit (USL) and Lower Specification Limit (LSL) and differentiating these from Upper Control Limit (UCL) and Lower Control Limit (LCL)
4. False out of control and false in control indications
5. The role of z in the Type I and II error
6. The role of sample size in Type I and II error
7. The role of reduced variability on improved process control.
8. Six Sigma
9. Managerial trade-offs
10. Detecting process shifts

Note to Instructors:
The lab manual contains the correct responses to various questions asked of the students as they follow the exercise. The responses are formatted using 'hidden text' that can be displayed (or not) when printed by using the appropriate setting under:

Tools → Options → Print → Include with document [x] Hidden Text

The hidden text can be viewed on the screen by:

Tools → Options → View→ Formatting marks [x] Hidden Text

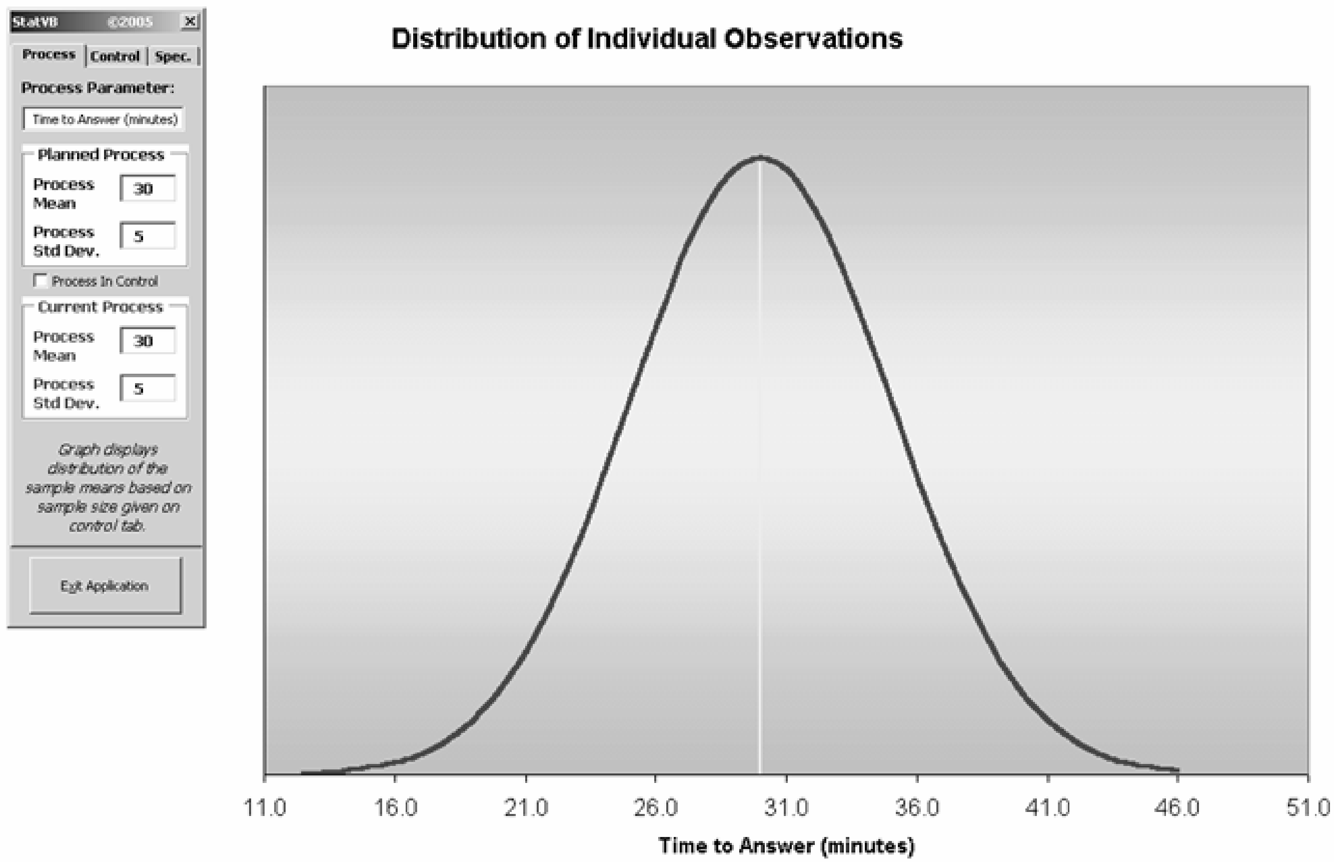

Exhibit 1: VBA Sheet (Process)

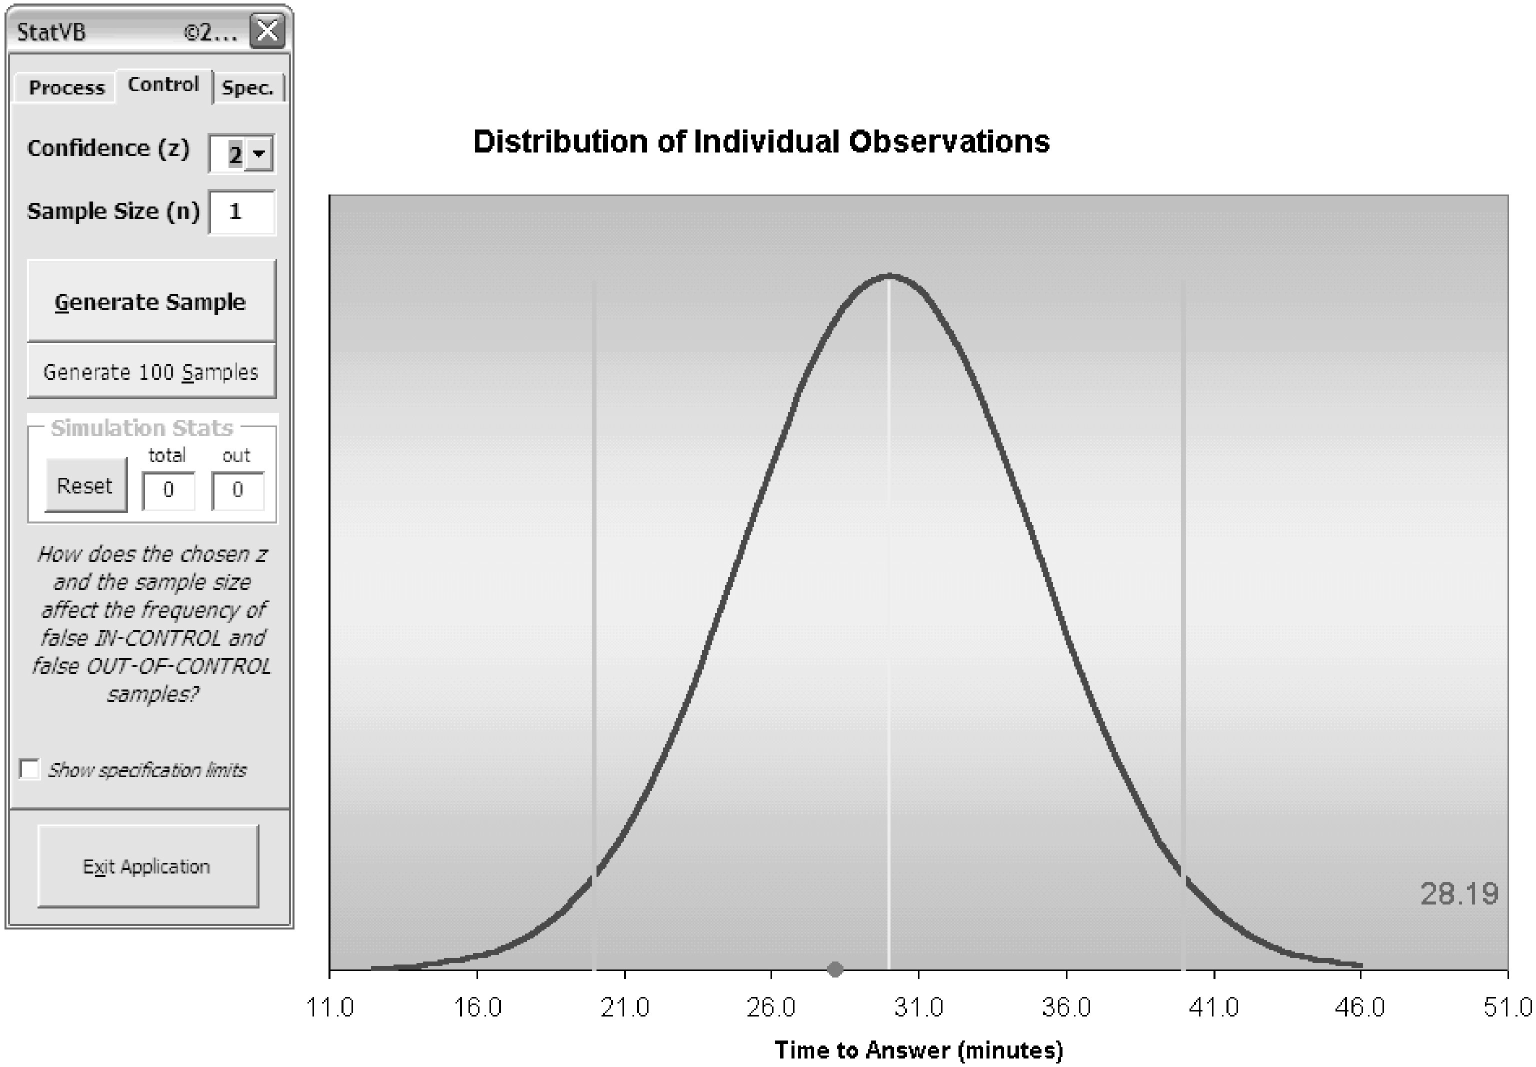

Exhibit 2: Process Control Sheet

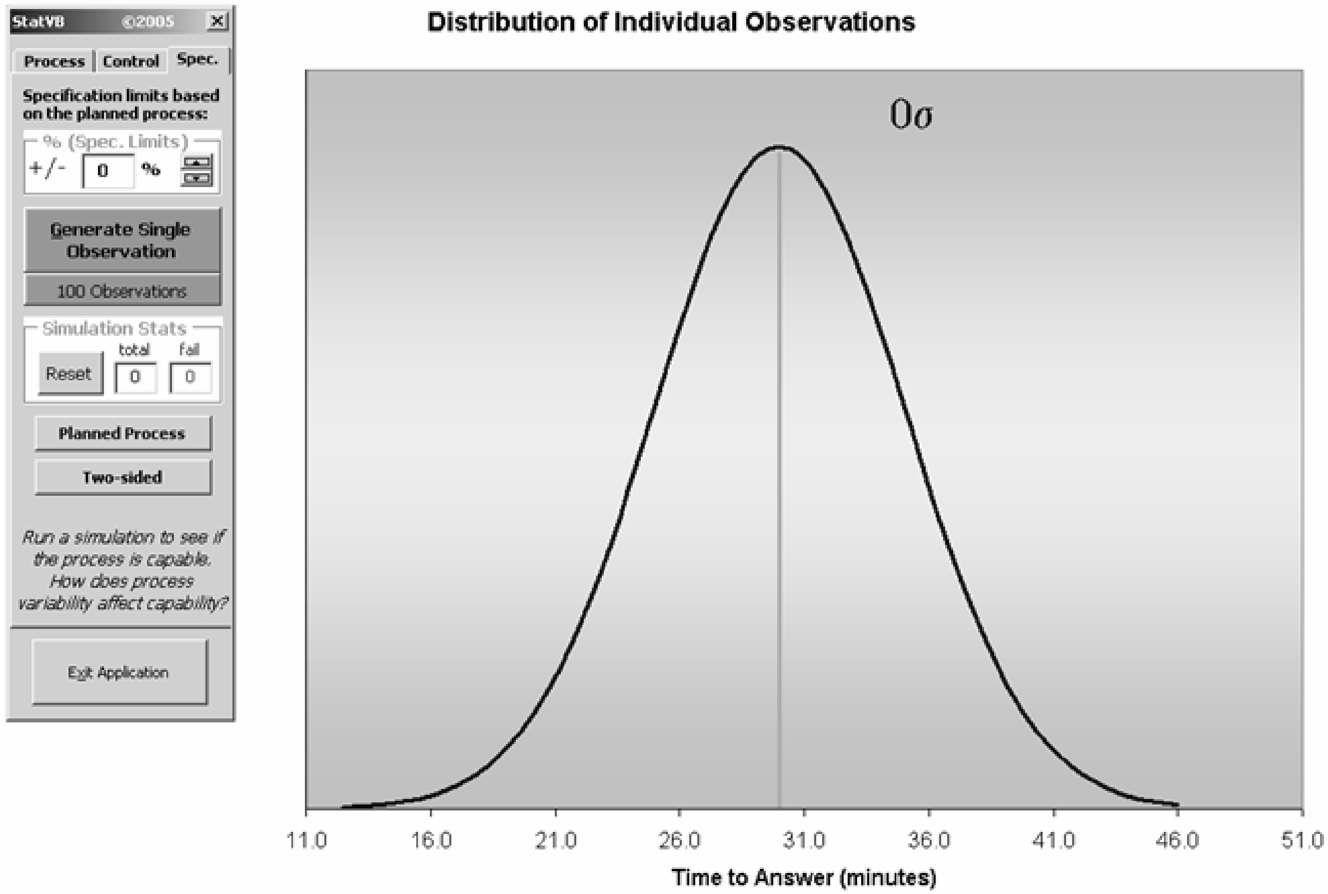

Exhibit 3: Specification Sheet

1. Software Set-up

Exhibit 1 is a screen print of what the user sees upon opening the VBA spreadsheet. The example shown is from a call center where the average time a customer is put on hold is measured, though this can be changed to suit the situation in the *Process* sheet (Exhibit 1). The *Process* sheet screen also allows one to specify two processes (both normally distributed for simplicity) through the μ and σ. The software graphs the distributions automatically. The process, as it was designed and set up, is called the *Planned Process*. The second distribution, called the *Current Process*, represents the process as it is actually working. If the process is working as planned (in control) the *Planned Process* and the *Current Process* are shown to be identical (same μ and σ). The TAB key is used to move between buttons and as an ENTER key.

There are two other sheets you can use: *Control* and *Spec* (Specification). Exhibit 2 shows the *Control* sheet where you can specify the SPC parameters such as the z value (1,2, or 3), and sample size (n). The UCL/LCL will always be shown. If $n = 1$, the chart will plot 'individual' observations as indicated in the title. If $n>1$, then sample averages are plotted. The sample mean will be displayed at the bottom right of the screen, as shown in Exhibit 2. You can superimpose the USL/LSL in the Control sheet by selecting the *Show Specification Limit* button, but ONLY if the sample size (n) is set to 1.

Exhibit 3 shows the *Spec* sheet. In the Spec sheet you can specify (by entering it directly or by using the increase/decrease buttons) the specification limit up to +/- 50%. This represents the maximum acceptable deviation from the planned process mean. You can also toggle the specification limits to be either one or two sided, as well as determine whether the current or planned process is displayed by clicking on the toggle buttons below the *Simulation Stats* display. In the situation of a call center, time on hold would generally have just a single-sided USL. In our example, the *Generate Single Observation* button generates a normally distributed ‘call on hold time’ and a corresponding marker along the x-axis of the graph

In the *Spec* sheet, the *Simulation Stats* display gives the number of total calls and the number of defective calls (in red). In the *Control* sheet, the *Simulation Stat* buttons gives the number of total samples taken and the number of sample means (in red) outside the LCL and UCL. On either sheet, you also have the option of generating 100 samples or observations to see the statistics displayed instantly. The title in Exhibit 1 and Exhibit 2 will state ‘Distribution of Individual Observations’ or ‘Distribution of Sample Means’ depending on whether the sample size is 1 or greater.

Note:

1. Hit Tab or Enter after any changes in values to ensure that it is recognized
2. After finishing a sampling experiment, use the *Simulation Stat Reset* button before proceeding.

2. Process Capability

Consider the following example. You are in charge of a call center for the Internet division of a major bank. Historically, the mean time (μ) that a caller has been placed on hold is 10 minutes and the standard deviation (σ) has been 5 minutes. Top management dictates that all calls have to be fielded within 15 minutes, (i.e., cannot be put on hold for longer that that). This is your USL. Your call center would be considered 'capable' if there was very little chance of a call exceeding this time. Now consider whether the following cases are capable:

Case I
Is the call center currently capable? To answer this, in the *Process* sheet set the *Planned Process* μ to10 minutes and σ to 5. Select the *Process in Control* button to ensure that the process is in control. Only one normal distribution will be seen. As the variability is fairly high relative to the mean, the distribution will be 'truncated' on the left-hand side, and any simulated calls with negative call times will be counted as having a wait time equal to zero.

In the *Spec* sheet, increase the specification limit to +/- 50%. For a μ of 10, it means that the USL is 50% greater (15 minutes). Since management considers only long waits as defective, click the two-sided toggle button to switch to a one-sided test. This will make it correspond to the call center example. Now click the *Generate Single Observation* button to generate a normally distributed 'call on hold time' along the x-axis of the graph. If the marker is a green dot, it means that that the caller was on hold for a time less than the USL of 15 minutes. If it is a red square, it indicates that the caller was put on hold for more than the USL of 15 minutes, i.e. a violation of the call center guidelines - a defective service!

Click *100 Observations* to generate 100 calls. Based on the results displayed in the *Simulation Stat* box, **how many calls took more than 15 minutes to answer?**

This can also be calculated mathematically using the properties of the normal distribution. For a given μ and σ, the *z* value corresponding to the probability that a value greater than X will be generated is given by Equation 1.

$$z = \frac{(X - \mu)}{\sigma} \tag{1}$$

Therefore, if X = 15, μ = 10 minutes and σ = 5 minutes, what is the value of z?

Given this *z*, using a normal distribution table, what is the probability that X > 15?

Tables of the Normal Distribution

Probability Content from -oo to Z

Z	0.00	0.01	0.02	0.03	0.04	0.05	0.06	0.07	0.08	0.09
0.0	0.5000	0.5040	0.5080	0.5120	0.5160	0.5199	0.5239	0.5279	0.5319	0.5359
0.1	0.5398	0.5438	0.5478	0.5517	0.5557	0.5596	0.5636	0.5675	0.5714	0.5753
0.2	0.5793	0.5832	0.5871	0.5910	0.5948	0.5987	0.6026	0.6064	0.6103	0.6141
0.3	0.6179	0.6217	0.6255	0.6293	0.6331	0.6368	0.6406	0.6443	0.6480	0.6517
0.4	0.6554	0.6591	0.6628	0.6664	0.6700	0.6736	0.6772	0.6808	0.6844	0.6879
0.5	0.6915	0.6950	0.6985	0.7019	0.7054	0.7088	0.7123	0.7157	0.7190	0.7224
0.6	0.7257	0.7291	0.7324	0.7357	0.7389	0.7422	0.7454	0.7486	0.7517	0.7549
0.7	0.7580	0.7611	0.7642	0.7673	0.7704	0.7734	0.7764	0.7794	0.7823	0.7852
0.8	0.7881	0.7910	0.7939	0.7967	0.7995	0.8023	0.8051	0.8078	0.8106	0.8133
0.9	0.8159	0.8186	0.8212	0.8238	0.8264	0.8289	0.8315	0.8340	0.8365	0.8389
1.0	0.8413	0.8438	0.8461	0.8485	0.8508	0.8531	0.8554	0.8577	0.8599	0.8621
1.1	0.8643	0.8665	0.8686	0.8708	0.8729	0.8749	0.8770	0.8790	0.8810	0.8830
1.2	0.8849	0.8869	0.8888	0.8907	0.8925	0.8944	0.8962	0.8980	0.8997	0.9015
1.3	0.9032	0.9049	0.9066	0.9082	0.9099	0.9115	0.9131	0.9147	0.9162	0.9177
1.4	0.9192	0.9207	0.9222	0.9236	0.9251	0.9265	0.9279	0.9292	0.9306	0.9319

Source: http://www.math.unb.ca/~knight/utility/NormTble.htm

Is this likely acceptable (i.e. is the call centre capable)?

The value of z indicates how many σ's the specification limit (maximum allowed wait) is from the process mean, μ. We refer to this as the σ *level* of the process and it is shown near the USL on the chart displayed with the Spec sheet (Exhibit 3). The higher the σ *level*, the less likely that we will get an observation outside the limit.

What is the σ *level* of this process?

Case II – Reduced variability

Assume that the mean 'on hold' time is still 10 minutes, but by installing more user friendly software and better employee training, the variation in the time required to help customers is reduced. Thus the on hold standard deviation σ has also been reduced. The new σ is 1.5 minutes. With these improvements, the process is now more consistent, but is the process now more capable?

To answer this, set the σ of the *Planned Process* in the *Process* sheet to 1.5 minutes and switching to the *Spec* sheet, observe the USL. Is there any change in the area of the curve that is *above* the specification limit compared to Case I? (You may need to switch back to σ = 5 to verify) **What does this imply for the likelihood of a defective product?**

Click *100 Observations* in the *Spec* sheet to generate 100 calls. Based on the results displayed in the *Simulation Stat* box, **how many calls took more than 15 minutes to answer?**

Calculate the probability of defects mathematically using Equation (1) and the Normal Distribution Table given below. **What is the *σ level*? Is the process more capable?**

Far Right
Tail Probabilities

Z	P{Z to oo}	Z	P{Z to oo}	Z	P{Z to oo}	Z	P{Z to oo}
2.0	0.02275	3.0	0.001350	4.0	0.00003167	5.0	2.867 E-7
2.1	0.01786	3.1	0.0009676	4.1	0.00002066	5.5	1.899 E-8
2.2	0.01390	3.2	0.0006871	4.2	0.00001335	6.0	9.866 E-10
2.3	0.01072	3.3	0.0004834	4.3	0.00000854	6.5	4.016 E-11
2.4	0.00820	3.4	0.0003369	4.4	0.000005413	7.0	1.280 E-12
2.5	0.00621	3.5	0.0002326	4.5	0.000003398	7.5	3.191 E-14
2.6	0.004661	3.6	0.0001591	4.6	0.000002112	8.0	6.221 E-16
2.7	0.003467	3.7	0.0001078	4.7	0.000001300	8.5	9.480 E-18
2.8	0.002555	3.8	0.00007235	4.8	7.933 E-7	9.0	1.129 E-19
2.9	0.001866	3.9	0.00004810	4.9	4.792 E-7	9.5	1.049 E-21

Source: http://www.math.unb.ca/~knight/utility/NormTble.htm

Case III – Shift in the mean

For some reason, the mean 'on hold' time has recently risen to 13 minutes. The σ is still 1.5 minutes. Is the process still capable? To answer this, in the *Process* sheet, unselect the *Process in Control* button and change μ of *Current Process* to 13. This means that the process has shifted and people on average are put on hold for 13 minutes (though as a manager, you may not detect this until you take a sample).

Observe in the *Spec* sheet the new distribution with μ = 13. Make sure that you are viewing the 'current' process by clicking on the process toggle button so that it displays the current process. Would you expect more defects when compared to Case II? Why?

Click *100 Observations* to generate 100 calls from the *Current process* distribution. **What is the defective rate?**

Calculate the probability of defects mathematically using Equation (1). **What is the σ *level*? Is the process capable?**

Case IV

You have instituted process improvement measures such that the mean 'on hold' time is back to 10 minutes and the σ has been reduced to 0.8 minutes. Is the planned process capable? To answer this, in the *Current Process* in the *Process* sheet, change the μ back to 10 and σ to 0.8, and select the *Process in Control* button. This implies that the process is very consistent. Click *Process in Control* and return to the *Spec* Sheet to view the *planned process* distribution. **Would you expect more or less defects when compared to Cases I, II and III? Why?**

Click *100 Observations* to generate 100 calls from the *Planned process* distribution. **What is the defective rate?**

Calculate the probability of defects mathematically using Equation (1). **What is the σ *level*? Is the process capable?**

What does 6σ imply for process accuracy and capability?

Note:
6σ is similar in concept but not the same as 'Six Sigma'**.** The term Six Sigma as originally coined by Motorola has a slightly different statistical interpretation as it allows for 1.5σ shifts in the mean before calculating the probability of an error which results in 3.4 defects per million. Secondly, it is important to emphasize that Six Sigma in organizations is a process improvement philosophy of which statistical methods are just one aspect. To quote former GE CEO Jack Welch (2001), "The big myth is that Six Sigma is about quality control and statistics. It is that – but its much more. Ultimately, it drives leadership to be better by providing tools to think through tough issues."

3. Difference between USL/LSL and UCL/LCL

In the *Process* sheet, set μ = 10 minutes and σ = 1.5 minutes for the *Planned Process* and select the *Process in Control* button. This means that process is working as it should. In the *Spec* sheet, ensure that USL/LSL is at 50%. In the *Control* sheet, set z to 1, sample size (n) = 1 and select the *Show Specification Limits* button. Note that the USL and LSL (in blue) are seen and are different from the UCL and LCL (in orange). The software is calculating the UCL/LCL from the SPC formulae for variables (shown below in Equations 2 and 3) while the USL is company policy.

$$UCL_{\mu} = \mu + z\left(\sigma / \sqrt{n}\right) \tag{2}$$

$$LCL_{\mu} = \mu - z\left(\sigma / \sqrt{n}\right) \tag{3}$$

Click *z* = 2. What happens to UCL/LCL in terms of spread?

What happens to USL/LSL?

What is the significance?

Set z back to 1 and increase sample size (n) to 5.

What happens to the spread of UCL/LCL as the sample size increases?

What happens to the specification USL/LSL? Why?

Set *n* = 1, *z* = 1. Click the *Generate Sample* button. This generates a sample and calculates the sample mean. It also plots this mean along the x-axis of the graph. It is represented as a green dot if within the UCL and LCL and as a red square if outside the UCL and LCL. The calculated mean value is shown on the bottom right of the chart (Exhibit 2). Since the sample size is 1, the mean is just the individual value. (*Note: A sample size of one would usually not be used for SPC in practice, but required for the purpose of this demonstration).*

Click it a few more times till you get a red square. In practice, if you were managing this process **what would you do when you get a red square**?

Does the call on hold violate company guidelines if the red square was above the UCL? In other words is the generated on hold time a defect?

Now, in the *Control* sheet, set *z* = 3, *n* = 1. In the *Spec* sheet, adjust the specification limits so that they are approximately 1σ from the mean (corresponding to 16% in the *Spec Limits*). After returning to the *Control* sheet, click the *Generate Sample* button till you get a green dot that is <u>above</u> the USL but <u>below</u> the UCL.

What does this indicate? Is this a desirable situation?

On the other hand, it is desirable is to have a system where even when the process is out of control, defectives will not be produced before the SPC mechanism detects it. To illustrate the desired system, in the *Process* sheet, unselect the *Process in Control* button, set μ of the *Planned Process* to 10 minutes and of the *Current Process* to 11 minutes. Set σ of both processes to 0.6 minute. In the *Spec* sheet, change USL to 50%. In the *Spec* Sheet, what *σ level* is this?

In the *Control* sheet, select the *Show Specification Limits* button. Click the *Generate Sample* button until you get a red square. You would naturally stop the process to investigate. **How many clicks did you need to get a red square? Is sample defective?**

What is the probability that calls will be defective?

Why is this a desirable situation?

How does this demonstrate differences between Control and Specification Limits?

4. Managerial issues in Process Control

The Effect of z on Type I /Type II error

A Type I error occurs when the process is in control but the sample indicates that the process is out of control. A Type II error occurs when the sample indicates that the process is still in control while it fact it is no longer in control. For this example, we want to start with a process that is operating as planned. In the *Process* sheet, set μ = 10 minutes and σ = 5 minutes for the *Planned Process* and select the *Process in Control* box. In practice, SPC uses the statistics from samples (n>1) gathered to evaluated the process and in this example, we will assume a sample size of four calls. In the *Control* sheet, set z = 1, n = 4 and unselect the *Show Specification Limits*, if necessary. Remember that the process is working as it should and the correct managerial decision is to NOT stop the process. When samples are generated, green dots are the correct indicators while the red squares incorrectly indicate that the process is out of control. You can try this out by generating a few samples using the *Generate Sample* button.

Now generate 100 samples. The number in red in the *Simulation Stats* box indicates the number of times the sample mean would have fallen outside the LCL/UCL. Each time this happens we would have stopped the process in ERROR (Type I error)

What is the approximate error rate (Type I) and is it acceptable?

Now set z = 2 and generate 100 samples. Note the Type I error compared to z = 1.

Set z = 3 and generate another 100 samples. **What is the effect of *z* on Type I error?**

Next we examine the Type II error. In the *Process* sheet, unselect the *Process in Control* button and change μ of the *Current Process* to 11. Note that the process is now out of control and now you have two distributions: one is the planned process and the other the current, shifted process. Assume that we do not know it has shifted and will be relying on samples to decide whether it is in or out of control. This is the situation that process managers face when determining whether or not the process is in control. You can try this out by generating a few samples using the *Generate Sample* button.

In the *Control* sheet, with $z = 3$ and $n = 4$, generate 100 samples. Note that now a red square is the correct indicator, i.e, we should be stopping the process. The green dot is actually a Type II error. **How often (%) does a Type II error occur?**

Repeat for $z = 2$ and $z = 1$. **Compare the Type II errors. Which is better? Why?**

What is the dilemma that you observe with respect to Type I and Type II error and the value of *z*?

Given that many companies choose to use $z = 3$, there are generally two different methods to reduce the possibility of a Type II error and ensure the accuracy of the monitoring:

<u>Method A</u>: Use a larger sample size by changing $n = 100$. Set $z = 3$ and generate 100 samples. **What happens to the sampling distribution?**

Has the error probability decreased? How do larger samples help?

What is the disadvantage of larger samples?

<u>Method B:</u> Set n back to 4. In the *Process* sheet, for both processes set $\sigma = 0.5$ minute. This indicates less variability through better equipment, training of employees and so on.

In the *Control* sheet, generate 100 samples. **What is the Type II error probability? What is the overlap between distributions?**

How do both of these methods help to reduce Type II error?

Both of these methods to reduce Type II error will cost money – which is better?

No matter how accurate the process, there is always a trade-off between Type I and Type II error. The issue of how low the error rates should be is a managerial decision based on the costs of allowing the error to continue versus the costs to detect it.

Other information from control charts

In the *Process* sheet, set μ of the *Planned Process* to 10 minutes, set μ of the *Current Process* to 11 minutes, σ of both to 1.5 minutes. In the *Control* sheet, set n = 4, and z = 3. Reset the Simulation Statistics and generate twelve samples from the current process. After generating each sample, roughly plot the sample mean (as displayed in the bottom right corner of graph in Exhibit 2) in Exhibit 4 below.

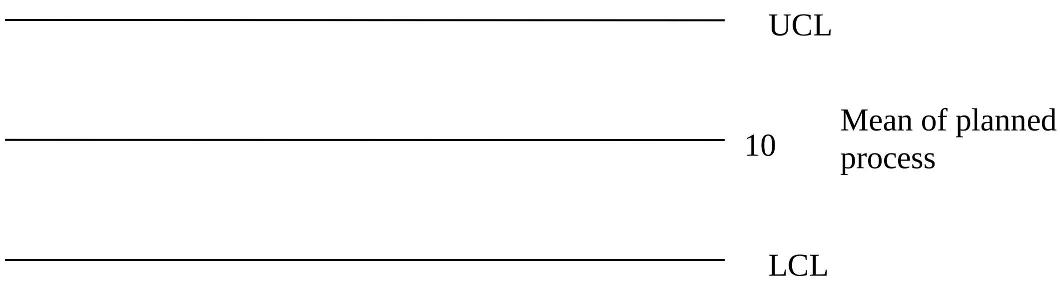

Exhibit 4

Are the plotted points randomly scattered above and below the planned process mean of 10?

What does this imply for managers looking for an indication of an out of control process?

<u>The magnitude of the shift.</u>

Ensure that in the *Process* sheet, μ of the *Planned Process* is 10 minutes, μ of the *Current Process* is 11 minutes, and σ of both processes is 1.5 minutes. In the *Control* sheet, ensure that n = 4 and z = 3. Now generate 100 samples. Recall that since the process is out of control, a red square is the correct indicator. **What is the resulting probability of a Type II error?**

In the Process sheet change μ of the *Current Process* to 15. In the *Control* sheet generate 100 samples. **What is the Type II error compared to the previous case?**

What does this imply for managers that are concerned about Type II error?

Once again, it is a managerial decision as to whether or not resources need to be deployed to reduce the probability of a Type II error for small shifts in the process. It may be adequate to allocate resources only for large shifts in process mean. Managers will need to understand the cost of process control versus the cost of not detecting process shifts.

Acknowledgements

The authors are thankful to their colleague Janice Bodnarchuk Eliasson, the anonymous reviewers and the Associate Editor for their valuable comments. Financial support was provided by the Natural Sciences and Engineering Research Council (NSERC) of Canada.