

Input and Outcome Charts and EFAST Studies

The fifth of six uses for process behavior charts

Donald J. Wheeler

While we may tweak things in production, we rarely get permission to conduct formal experiments with an operating production line. Production's job is to make product while experiments are what they do in R&D. So how can we learn about an existing production process without rocking the boat. This is where input and outcome charts have a role to play.

In a production environment things are generally assumed to be okay until the fraction of nonconforming product takes a turn for the worse. When this happens the production line is suddenly "in trouble" and process engineers are called in to put out the fire. They go to the line, get in the way of the operators, and collect data. Then, after arguing about these data, they are likely to tweak one or more of the process inputs. If the fraction nonconforming drops, they take credit for fixing the problem as they race off to the next emergency. If the process continues to be in trouble, then they do another process trial to try to find another input variable to tweak. Occasionally a problem will prove resistant to this sequence of process trials and a more systematic approach will be needed.

BRAKE LININGS

The blueprint specification for the maximum gap between the brake lining and a radius gauge was 9 mils. Only about 50 percent of the product met this specification, so the customer had agreed to a deviated specification of 12 mils. (About 90 percent of the brake linings met this deviated spec.) Since gaps smaller than 12 mils resulted in satisfactory brake shoes, everyone was happy.




Figure 1: Gaps Between Brake Lining and Radius Gauge

However, things changed when the automobile company began a campaign to do "whatever it takes" to improve quality. (Apparently they had no clue that world-class quality had already been defined as "on target with minimum variance.") Under the banner of "whatever it takes" the customer let the supplier know that the deviated specification of 12 mils was going away.

After all, how could they claim to be doing a quality job when half the brake linings failed the blueprint specification? Since the customer also made brake linings, the supplier was afraid that they were about to lose this profitable business. So they formed a project team and told them to improve the conformance with the blueprint specification.

This project team spent the next year running 130 process trials (at a cost of nearly a million dollars). At the end of that year their cause-and-effect diagram contained about 600 causes for excessive gap. And yet only 55 percent of the brake linings met the blueprint specification of 9 mils maximum gap. They had spent a lot of time and money, but no magic bullet had been found that would fix the problem.




Figure 2: Gap Distribution Following 130 Process Trials

Since this was not good enough, the whole project team was sent to my SPC class. Following the class they sat down with their cause-and-effect diagram and identified about 60 input and outcome variables that they wanted to track using average and range charts. The majority of these charts were kept off-line by the project team, who posted them on the walls of a large meeting room. Only a handful of charts were kept on-line by the operators. Since the brake linings were produced in batches, the subgroups on all of these charts were identified by batch number. In this way, when changes were found in the outcomes, they could look back through the charts and read the history of that batch of product.

Over the next nine months they identified fifteen things to change in their process. Twelve of these fifteen things had already been studied in the process trials from the previous year. So why hadn't these inputs been changed earlier? As one engineer put it, "The charts gave us the courage of our convictions. We could see that we really needed to fix these things."

Three of the fifteen changes were complete surprises. These were things that they had not thought to study in their 130 process trials. One of these was discovered when five outcome charts jumped outside their limits simultaneously on Monday morning, January 4, and stayed outside the limits until 4:00 on Thursday, January 7, when all five charts jumped back inside their limits together. Since this jumping back inside the limits coincided with the arrival of the first fresh batches of the year at that step in production, they realized that the product sitting around over the Christmas break was the problem. While the amount of time sitting in the open air might not seem to matter when the next process step is to bake the parts in an oven, it clearly made a difference. Until this time they had no system for controlling the in-process inventory prior to the ovens. Parts were stored on racks with wheels. These racks were simply pushed around within a storage area and racks were pulled out for baking in a haphazard fashion. When they instituted a first-in, first-out inventory control on these racks, with a corresponding minimum and maximum amount of air-dry time, they got more consistent product gaps.

By this time the cause-and-effect diagram had grown to over 1300 causes, but the number of input and outcome charts had dwindled down to about 20, half of which were kept on-line by the operators, and half of which were kept off-line by the quality department.

At this point the supplier went to the customer. With these fifteen process changes they had gotten to the point that 99 percent of their brake linings were meeting the blueprint specification of 9 mils. While this was twice as good as they were doing before, they still were not getting 100 percent conformity. So the customer offered to help with the last one percent by changing the blueprint specification to 13 mils. After all, they knew that gaps smaller than 13 mils did not seriously affect the functionality of the brake shoes!




Figure 3: Gap Distribution After EFAST and 15 Process Changes

By systematically tracking the inputs and outcomes with a collection of process behavior charts this project team substantially improved their production process. Twelve of the fifteen process changes they made had been studied in their process trials. This is typical. About 80 percent of the important cause-and-effect relationships will be known to the process engineers, and about 20 percent of the important cause-and-effect relationships will be unknown. By using input and outcome charts you can discover the impacts of both the known and the unknown causes.

(In my various books I have referred to the use of Input and Outcome charts as “extended monitoring.” I have used the more descriptive name here to differentiate this use from the process monitor charts discussed in July.)

MORE EFFICIENT EXPERIMENTATION

My fellow statisticians would all recognize the limitations of the process trial approach to problem solving. Most process trials turn out to be one-factor-at-a-time (OFAT) experiments where one input is manipulated and the difference in the outcomes is observed. When you pick the right input variable to manipulate, and when you pick the right levels to use, an OFAT experiment will allow you to learn about an important cause-and-effect relationship. But picking the right input variable to manipulate when your cause-and-effect diagram contains hundreds of causes is a bit of a shot in the dark. In addition, as happened in the example above, in the presence of several important cause-and-effect relationships, no single OFAT study may look promising enough to result in people making the needed process changes.

My colleagues usually suggest a multiple-factor-at-a-time (MFAT) strategy. This approach has been thoroughly proven. It consists of using carefully selected experimental runs involving two or more levels for each of several input variables. By observing the outcomes for these experimental runs we can analyze the data to discover both main effects and interaction effects

among the inputs studied. While very efficient experimental designs exist that will handle large numbers of input variables, the complexity of running these experiments in a production environment, combined with the inability to get permission to “change everything at once” will often limit the use of MFAT studies use to research units, pilot plants, and bench units.

A CONTAMINATION PROBLEM

A complex chemical process had a problem with too much contamination in the final product in spite of it being processed through five distillation columns. Since the project team could not experiment with the production unit, they had to settle for using a bench unit. Since this process used hydrogen fluoride as a catalyst, the bench unit which was built inside the high-pressure laboratory for safety. Given the corrosive nature of hydrogen fluoride it took about two to three days after each run to repair the bench unit and set up for the next experimental run. After eight months of OFAT experimentation the team had no recommendations for production that they were confident would work. So the whole team was sent to my class on Industrial Experimentation.

After the class they ran two Plackett-Burman designs (basic and reflected) involving seven input variables. (Running these two designs allowed them to separate main effects from interaction effects.) Out of these sixteen experimental runs they identified two process changes, each of which would reduce the contamination in the final product by an order of magnitude. (One of these changes had been anticipated in the OFAT studies, and the other was an interaction effect they had not thought of examining before.) Thus, in two months time with MFAT they accomplished more than they had in eight months with OFAT. They knew what to recommend to production, and they had confidence in their recommendations.

THE CHOICE OF VARIABLES

As with OFAT, the MFAT approach still depends upon your ability to choose the right input variables. With any experimental study there are only three things you can do with a process input.

1. You can study it in the experiment.
2. You can hold it constant during the course of the experiment.
3. Or you can ignore it.

Before you can put a variable in category 1 it will have to be part of your cause-and-effect diagram. (You cannot study a cause-and-effect relationship that you cannot name.) The variables in category 2 will tend to be those control factors for the process that are not in category 1. And category 3 is going to include all of the remaining known causes plus all of the unknown causes. So, when your cause-and-effect diagram contains hundreds (or even dozens) of causes, category 3 is going to be much more heavily populated than categories 1 and 2. How can you prevent the variation in these category 3 variables from skewing or even obscuring the effects of the experiment?

One hundred years ago the problem of the category 3 variables was a formidable barrier to agricultural and biomedical research. The answer worked out in that environment was to build replication into the study by using many different experimental units for each treatment combination in category 1. By randomly assigning experimental units to treatment combinations

the idea was to average out the extraneous effects of the variables in category 3 within each treatment combination in category 1. This is why agricultural and biomedical experiments tend to be broad and shallow—many experimental units but only one iteration.

However, industrial experimentation is fundamentally different from agricultural and biomedical experimentation—there is only one unit, one line, or one process that we need to get out of trouble. Here the experiments are narrow and deep—one experimental unit and many iterations. In addition we often will have only one or two experimental runs per treatment combination from category 1. This small number of runs per treatment combination does not provide the elbow room needed for the extraneous variables from category 3 to average out, so that the randomization strategy of agricultural experimentation will not work here. (Randomizing the run order in an industrial experiment will not add anything but complexity.)

So, if we pick the right variables for category 1 and category 2, and if we do not get clobbered by unknown or overlooked variables from category 3, MFAT experimentation will work in an industrial setting. Fortunately, there is an easier way.

EFAST STUDIES

In contrast to OFAT and MFAT studies, the use of input and outcome charts allows you to study every-factor-at-the same-time (EFAST). Rather than picking a few factors to manipulate, while other factors are held constant or ignored, an EFAST study lets all of the input variables vary in their routine manner while the process is running. When the process outcomes change and you look for those inputs that also changed, you will have an opportunity to identify those input variables that are of greatest import for your process. If this sounds vaguely Aristotelian, it is. This idea is over 2300 years old. It has a proven track record. And input and outcome charts make EFAST studies feasible for complex problems in a production environment.

Since no manipulation of the inputs is required no effort is needed to keep track of a set of experimental runs. No randomization of run order is needed. No complex analysis is required to disentangle the orthogonal structure of the resulting data. Signals are indicated by points that cross over the line. By focusing on the changes that occur in production you will effectively prioritize the uncontrolled input variables in order of their impact upon the process. (Knowledge that Input A affects Outcome Y is not helpful in troubleshooting when Input A is already tightly controlled.) So, rather than being distracted, and possibly misled, by “what we know” about our process, an EFAST study lets the process tell us *what we still need to learn* about our process. Remember, in the words of Daniel Borstein, “the greatest barrier to discovery is not ignorance, but rather the illusion of knowledge.” An EFAST study helps us discover what causes our process to change, rather than depending upon our understanding to set up an experiment.

Input and outcome charts let us study every-factor-at-the-same-time. We do not have to figure out in advance which input variables are likely to be important. We can wait for dominant cause-and-effect relationships to make their presence known by volunteering. The signals on these charts are easy to detect and easy to explain to others. (Try explaining an ANOVA table to a supervisor.) And when a change is made in the process the charts provide the feedback needed to determine if the change was effective. EFAST studies built on the use of input and outcome charts are a proven path to discovery that can be used in production.