

Optimization and Experimentation

The rest of the story...

Donald J. Wheeler

Experimental designs that result in orthogonal data structures allow us to get the most out of both our analysis and our research budget. As a result, designed experiments have been used for everything from basic research to process optimization. These multiple roles make it crucial to understand the implicit and explicit assumptions behind the use of an experimental approach.

Consider the problem of process optimization. It seems quite simple to carry out a sequence of experiments in such a way that we can move our process toward a set of operating conditions that will optimize the critical outcomes. George Box and Norman Draper came up with the idea of evolutionary operation. Stan Deming taught classes on the use of simplex algorithms for optimization. Over the years many others have followed the same path with various experimental approaches for process optimization, and today you can find software, books, and classes on this topic.

But wait a minute: Was not the purpose of R&D to determine the optimum operating conditions for your process? Weren't they supposed to tell you that if you "do this" you will "get that?" If that is the case, why should you be concerned with further process optimization? Could it be that while the process was "optimized" back at the start of production, it is no longer running optimally today? If this is so, then what happened? In order to answer this question we need to go back and look at the elements of an experimental approach.

When we want to optimize a process relative to even a single process outcome (a response variable) we will have a fairly large number of input conditions (factors) to consider. Fortunately, not all of these factors are equally important. Some factors will have a pronounced effect upon the response variable, while others will have only small effects. As a result these cause-and-effect relationships will usually form a Pareto distribution like the one in Figure 1.

Thus, the first step in the successful optimization of any response variable is to correctly identify those factors that have the dominant effects. It is these factors with the dominant effects that we will want to use as our control factors for production. Of course, as we begin this job of identifying the critical few factors, our knowledge about these cause-and-effect relationships will generally look more like Figure 2 rather than Figure 1. Given this blank slate, our first problem will be the problem of determining which input variables to study in our experiments.




Figure 1: A Few Factors Will Have Large Effects Upon the Response Variable




Figure 2: Our Starting Point: Which Factors Should We Study?

If we wanted to study all twenty factors in Figure 2, and if each factor had only two levels, we would have a total of 1,048,576 factor combinations to consider! Clearly the straightforward approach will not be a viable option. So immediately we look for ways to reduce the size of the problem so we can reduce the size of the experiment. One of the ways this is done is by prioritizing the twenty factors and then studying only those factors we deem most likely to have a large effect upon the response. (Yes, hopefully you are good at guessing which factors are important.)

Say we study the ten factors we think are most likely to have dominant effects using some fractional factorial design and discover the effects shown in Figure 3. At this point we can be reasonably certain that Factors 5, 1, and 7 are part of the critical few. We can also be fairly sure that Factors 2, 3, 4, 6, 8, 9, and 10 belong in the trivial many. But we will still be in the dark regarding the factors not studied.

So, while we do have real knowledge by the time we get to Figure 3, and while we can act on this knowledge, our knowledge remains partial and incomplete. We will know something about the factors we have studied, but we will not know anything about the factors we have not studied. And therein lies the problem with an experimental approach. Regardless of how efficient our experimental design may be, and regardless of how many factors we can study,

there will always be more factors to consider than we can possibly include in our experiment.

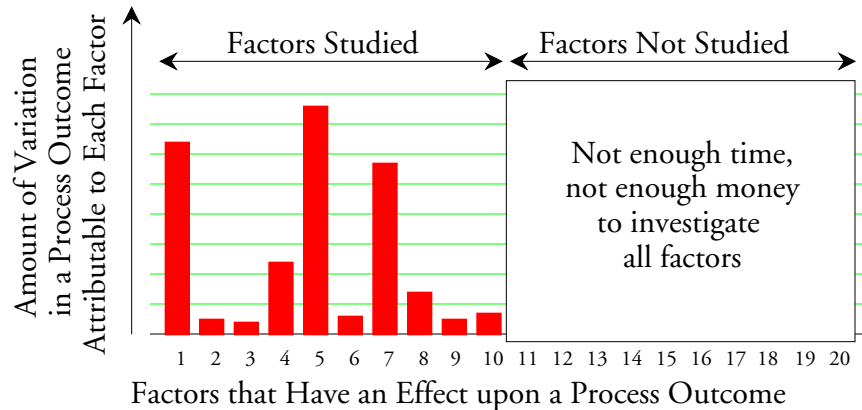


Figure 3: What We Discover When We Study the Ten Most Likely Factors

In any experimental study, there are only three things you can do with an input variable: you can *study it*; you can *hold it constant*; or you can *ignore it*. In agricultural studies, where many different experimental units are used simultaneously, blocks of experimental units are often used to hold certain inputs constant, and then within each block the treatment combinations are randomly applied to the experimental units. The purpose of this randomization is to (hopefully) average out the effects of those factors which are being ignored. In this way we separate the variation due to the factors studied from the variation due to those factors that are not part of the experiment, and thereby manage to get a useful understanding of how the experimental factors affect the response variable. All experiments have this limited scope. They study some factors and seek to block out, or average out, the effects of other factors.

Control Factors 1, 5, and 7 are held constant in production.
All of the product variation comes from the remaining Factors




Figure 4: The Rest of the Story: What You Don't Know Can Hurt You

But what does all this mean in terms of process optimization? Based on the results shown in Figure 3 we might carry out further experiments to determine the optimum levels for Factors 5, 1, and 7 and consider our process to be optimized. But what happens to our process next week

when Factor 14 changes levels? What happens when Factor 16 goes on walkabout next month?

As the levels of Factors 14 and 16 change the process outcomes will change, and soon it will be time to “optimize the process” once again. As one engineer expressed it, “I spend my time working on projects with an average half-life of two weeks, implementing solutions that have an average half-life of two weeks.” Or as a manager once told me, “We currently have a task force on vibration. I have been here 30 years and each year we have had to form a task force on vibration.” As long as your experimental studies continue to exclude Factors 14 and 16 you are likely to also experience this sense of “déjà vu all over again.” Thus, we have the central conundrum of using experimental studies for process optimization. Leave even one dominant factor out of the experiment and your results will be limited, partial, and incomplete. ***You cannot optimize any process until all of the dominant cause-and-effect relationships are known.***

Moreover, all experimental approaches treat your process as if it is static and unchanging. Of course the reality is that processes are never static. Entropy alone guarantees that wear and tear will constantly affect your processes. As a result, the set of critical factors will continually evolve, making yesterday’s experiments less and less applicable to today’s processes. Experimental studies are not enough for process optimization simply because entropy keeps changing the game. ***Entropy prevents us from ever knowing all of the dominant cause-and-effect relationships.***

So what are we to do? How can we escape this endless cycle of continuing reoptimizations of our processes?

When all you have is a hammer, every problem looks like a nail. While you can drive a screw into a wooden plank with a hammer, don’t expect the screw to hold very well. You may have to keep coming back to hammer those loose screws over and over again. But if you change your tool and use a screwdriver instead of a hammer, you get a much more satisfactory and long-lasting result.

As a statistician I was taught how to specialize in experimental studies. My hammer consisted of designed experiments. By training and outlook I was blind to the role of observational studies. Yet it is only by using observational studies that we can break out of the endless cycle of improvement projects that have a half-life of two weeks.

Now it must be said that statisticians who are good experimenters know that it is important to be alert for evidence of important factors that are not part of the experiment. However, in most of the textbooks and classes, statisticians offer no systematic approach for doing this. It is usually left at the level of an art. On the few occasions when this topic is addressed it is done using various techniques and sundry approaches that change from problem to problem. While these miscellaneous techniques make sense to those with the deep understanding that comes from a lifetime of study of the discipline of statistics, these techniques appear mysterious and confusing to our students. They see statistics as a grab bag of techniques having no unifying philosophy. So how can our students practice this important aspect of data analysis? How can they learn to identify the dominant factors that are not included in the experiments? Shewhart supplied the answer with his systematic approach to observational studies.

Observational studies differ from experimental studies in several fundamental ways. The first of these differences is in the type of data used. Experimental studies use data collected under special conditions, so there will always be a fixed and finite amount of data available in an

experimental study. Observational studies use data that are a by-product of operations, so there will usually be more data available over time.

The second difference is in the conditions under which the data are collected. Experimental studies generally compare two or more conditions. These conditions require special set-ups and special effort to change the conditions while collecting the data. Observational data are generally collected under one condition and this condition represents the “run of the mill.”

The third difference has to do with what we expect to find in our data. In an experimental study we are trying to create signals. We have paid good money to create special conditions so we can see if there is a detectable difference between those conditions. In an observational study we are not expecting to find any signals. We expect all of the data to behave in the same way over time.

The fourth difference is how we analyze the data. In an experimental study we analyze the finite data set on a single pass through all the data. Here we must use a one-time analysis technique. In an observational study we perform an act of analysis every time a new value is obtained, so we have to use a sequential analysis technique.

The fifth difference is in our approach to taking risks with our analysis. In the one-time analysis of a finite set of experimental data we can afford to use a traditional or exploratory analysis in order to avoid missing some of those expensive signals we have tried to create. In the sequential analysis of a continuing data stream where we do not expect to find any signals we will want to use a conservative analysis in order to avoid having too many false alarms.

These five differences between experimental studies and observational studies are summarized in Figure 5. An appreciation of these differences is fundamental to the effective use of observational studies.

Experimental Studies	Observational Studies
Finite amount of data	Continuing data stream
Two or more conditions present	One condition present
Expect to find signals	Expect to find no signals
One time analysis	Sequential analysis
Traditional alpha-level	Conservative alpha-level

Figure 5: Experimental Studies versus Observational Studies

Process behavior charts are the premier tool for conducting observational studies. They allow you to study your process without restriction. You do not have to ignore any factors. You do not have to hold any factors constant. You do not have to identify which factors you think are most likely to have a dominant effect upon your response variable. Simply plot your response variable over time (preferably in real time) and pay attention to those points where your process changes. When a change occurs, seek to identify what input variable changed. (When you think you have found that factor, if need be you can always verify it by trying to turn the effect on and off.) Once you have identified a new dominant factor you will want to include it in the set of control factors for your process. Thus, rather than trying to experiment to discover the dominant cause-and-effect relationships, with an observational approach you simply wait until the dominant factors volunteer and make their presence known. And this is how the simple process

behavior chart can be used as the locomotive of process improvement and optimization.




Figure 6: A Process Operating Up to Full Potential

While operating in this way my clients have reported cutting the process variation down to one-third, one-fourth, and one-fifth of what it was initially. They have discovered new technical knowledge about their processes, and they have also discovered some dumb things they were doing. They have taken over markets by offering the best quality at the lowest price while increasing their profit margin at the same time. It is simply a fact of life that when you are working with an existing process, using observational studies to optimize your process is easier, more complete, and more successful than an experimental approach. Moreover, since process behavior charts evolve as the process evolves, they allow you to keep improving the process while operating it optimally.

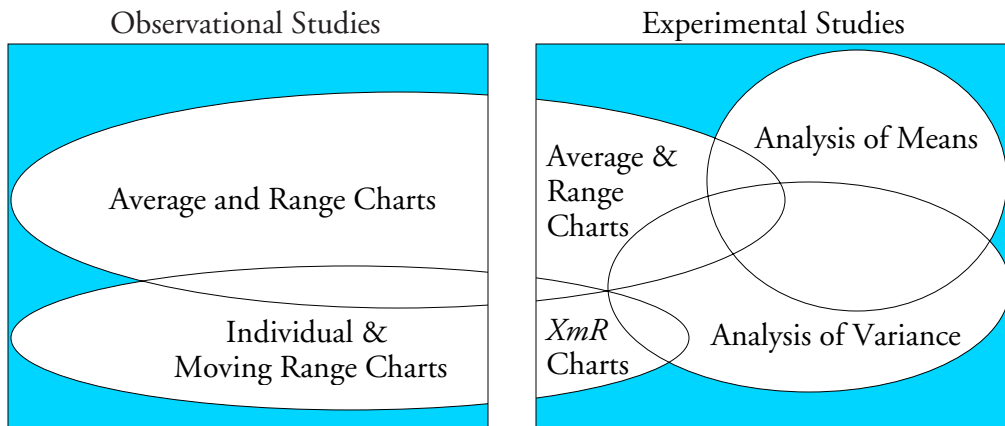


Figure 7: Analysis Techniques for Observational and Experimental Studies

While many different analysis techniques have been developed for experimental studies, the fact that they are all intended for the one-time analysis of a finite data set makes them inappropriate for the sequential analysis of a continuing stream of data. On the other hand, the sequential nature of the process behavior charts allows them to be used with both observational

and experimental studies.

The observational approach is actually older than the experimental approach. It has been part of the empirical tradition from the very beginning. Over 2300 years ago Aristotle taught us that in order to discover the causes that affect a system or process we need to pay attention to those points where the process changes. And this is exactly what a process behavior chart enables us to do.

SUMMARY

Experimental studies are like studying lions at the zoo. There are many things you can learn at the zoo that are true and useful. However, there are things you will never learn about lions at the zoo because the zoo is not their natural habitat. Observational studies are like studying lions in the wild. Both types of studies are valid. Both are needed.

Experiments are very good at establishing the effects of a specific set of factors upon a response variable. But when it comes to process optimization we need to identify the effects of all of the dominant factors. This makes the limited scope of experimental studies a handicap. While we may start off with an experimental approach in order to identify the dominant factors among those that we can think of, and while we can use these dominant inputs as the control factors for our process, experimental studies can only take us so far. They only provide partial solutions because they restrict our attention to only a few factors. In consequence, our experimental results will always be limited by our ability to choose the right factors to study. Moreover, our experimental results will always depend upon all the factors not included in the experiment remaining unchanged. Thus, while we may start off with an experimental approach, we will also need an observational study to complement and complete our experimental approach as we seek to optimize our process.

An optimal process is one that is operated on-target with minimum variance. The only way that a process can operate with minimum variance is for it to be operated predictably, and the only operational definition of predictable operation is a process behavior chart.

Process behavior charts allow you to learn from your process data. They warn you when your process changes. When kept in real time they help you to identify the assignable causes of those process changes. And when you incorporate these assignable causes into the set of control factors you will simultaneously increase your ability to operate at a specific value while you reduce the variation in the process stream.

So, between the limited scope of experimental studies and the necessity of operating predictably to achieve minimum variance, you simply cannot operate your process in an optimal manner over an extended period of time without the use of process behavior charts. They allow you to study your process in its natural habitat where you can discover things you would never find via experimentation.

