

SENSITIVITY ANALYSIS FOR ROBUST PARAMETER DESIGN EXPERIMENTS

Joseph R. Litko

Engineering Management and Systems Department
300 College Park
University of Dayton
Dayton OH 45469-0236, U.S.A.

ABSTRACT

Robust parameter design experiments lead to products and processes that are insensitive to the effects of noise. These experiments reveal the interaction of the noise sources with design or control factors usually allowing creation of products that are relatively immune to noise. Finding truly optimal settings for design factors depends on the noise in the lab being representative of the actual operating environment and assumes potential product users all see the same noise conditions. This paper shows how basic design solutions can be shaped when multiple populations see different noise conditions and when typical assumptions on noise sources are violated.

1 INTRODUCTION

Taguchi is generally credited with making robust design a standard product or process design practice in many industries. See Box (1988) and Phadke (1989) for a critique and explanation of Taguchi's contributions. The general idea is to use an experiment to find design factor settings that make a product insensitive to the effects of noise. The response or objective measure is some physical characteristic that is closely related to the primary function of the product or process. For example, brake torque or stopping distance might be the response that is measured to judge the quality of a brake system design.

Noise can be internal, external, or caused by unit-to-unit variation in product components. Wear and tear over the life of the product (e.g. worn vs. new brake pads) is an example of internal noise. External noises are typically customer usage patterns (heavy application of brakes) and environmental conditions (wet road and brake pads).

A robust parameter design experiment controls noise levels in a laboratory setting to explore and measure their relation to settings of the design factors. In actual operation by customers or users, noise levels must be viewed as random variables. When design factor levels can be found that give consistent performance (i.e., low response variance

over all noise conditions), the parameter design experiment has been successful. When the product gets into the hands of users the noise sources will not be controlled. In principle, the intelligent choice of design parameters has made the product immune to the effect of noise. The obvious benefit is that no extra cost (beyond the experiment) is needed to control or compensate for noise.

Since its introduction and popularization by Taguchi, parameter design has evolved and been incorporated with some modification as a mainstream statistical tool. Nair (1992) is a panel discussion of parameter design that includes several points of view and approaches to parameter design. One principal modification is a move away from Taguchi's use of signal to noise ratio and a move toward a response function model (RFM) approach. The RFM is a direct model of the quality characteristic (response) of interest in terms of the design and noise factors. The primary objective remains minimizing the variance of the response with respect to the noise factors. See Miller (2002) or Miller and Wu (1996) for good discussions of the RFM approach and comparisons to Taguchi's original S/N ratio approach to parameter design.

The RFM is based on a designed experiment – often a two-level fractional factorial of some kind. That is the context for this paper. The issue we focus on is the problem of setting noise levels for the experiment. There are two points to be made here. First, relative levels of various noise factors may be different across the population of intended users. Typical experiments view users as a single population and select noise levels accordingly. Second, good analysis ought to keep in mind uncertainty in estimates of the noise distribution parameters. In effect, we ought to employ a sensitivity analysis. Simulation in conjunction with response surface modeling is one way to do this sensitivity analysis.

2 DESCRIPTION OF THE PROBLEM

Response function methods create a predictive model for the response (quality characteristic to be measured) as a

function of control and noise factors. Noise factors are explored in the same experimental design array (usually called a combined array) as the control factors. If we set up an experiment of the proper resolution, we can get a direct estimate of the main effect of the noise factors and also estimate noise factor interactions with the control factors. See Borkowski (1997) for a discussion of possible experimental designs. The response model to be estimated is shown in equation (1), where the x_i represent design factors and the z_j represent noise factors.

$$y^{(s)} = f(x_1, x_2, \dots, x_{N_x}) + \sum_{j=1}^{N_z} \delta_j z_j + \sum_{j=1}^{N_z} \sum_{i=1}^{N_x} \gamma_{ij} x_i z_j \quad (1)$$

The function $f()$ represents the main effect of control factors on the response. The parameters δ_j and γ_{ij} are to be estimated from the experiment and will determine control factor settings that minimize the effects of the noise variables – the z_j .

A transmitted variance model (i.e., propagation of error) is used to derive control factor settings that minimize response variance based on the estimates of the coefficients in the model. This optimization might include constraints on factor settings and minimization of variance subject to a constraint on acceptable mean response, for instance. This could be expressed as in Problem $P^{(s)}$ below.

Problem $P^{(s)}$

$$\begin{aligned} \text{minimize: } & \sum_{j=1}^{N_z} \left(\delta_j + \sum_{i=1}^{N_x} \gamma_{ij} x_i \right)^2 \sigma_{z_j}^2 + \\ & 2 \sum_{j=1}^{N_z} \sum_{k < j} \rho_{jk} \left(\delta_j + \sum_{i=1}^{N_x} \gamma_{ij} x_i \right) \left(\delta_k + \sum_{i=1}^{N_x} \gamma_{ik} x_i \right) \sigma_{z_j} \sigma_{z_k} \\ \text{s.t. } & \sum_{i=1}^{N_x} a_{ik} x_i \leq c_k \quad k = 1, 2, \dots, K \end{aligned}$$

where ρ_{jk} represents the z_j, z_k correlation

$\sigma_{z_j}^2$ represents the z_j variance.

In $P^{(s)}$ the constraints on the x_i (design factor settings) might represent a minimum acceptable value for the response, limits on the ranges of the x_i , or mandatory relations between the x_i . The objective function represents $\text{var}_z(y)$ – variance of the response with respect to the noise sources. Although the levels of noise factors used in the experiment are not crucial to the estimates of the δ and γ coefficients, the minimization is not meaningful unless these levels are representative of the noise one encounters in the real world, e.g., do all users see the same distribution of noise, and is the effect of individual noise factors on the response linear. Solution to the quadratic problem in $P^{(s)}$ is straightforward from an optimization viewpoint.

To illustrate the procedure consider a case with three noise and three design parameters. Suppose the matrix in Table 1 represents the δ and γ coefficients estimated from a factorial experiment. The values shown directly below the x_i represent the optimal solution for the given matrix. Note two points. First, if there were no off-diagonal elements, then each control factor setting for $[x_1, x_2, x_3]$, could be used to neutralize a particular noise source, i.e. if $\text{Gamma}(Z_2, X_1)$ were zero, the optimal solution by inspection is $x_1 = 1, x_2 = -1, x_3 = -1$. Here ‘neutralize’ means to minimize the variance which is the objective function of problem $P^{(s)}$.

Table 1: Example of Main (Z) and Interaction (Z by X) Coefficients

			X1	X2	X3
		$\gamma(i,j)$.735	-1.00	-1.00
δ_1	3.0	Z1	-3.0	0.0	0.0
δ_2	3.0	Z2	3.0	3.0	0.0
δ_3	6.0	Z3	0.0	0.0	6.0

Second, estimates of Δ_1 and $\text{Gamma}(1,1)$ can be accurate relative to each other regardless of whether the Z1 noise levels used in the experiment are totally representative of the real world. But when no set of design values exists that can totally neutralize noise, the best solution does depend on the relative magnitude of the noise sources. Essentially, when compromises must be made, it is critical to know the real variability users will see from each noise source – how else to minimize their effects.

It is also very important that the response varies linearly with respect to the noise factor as this is an assumption of the response function model approach. Further, if distinct segments of the population see different noise distributions (e.g. hot vs. cold weather or low altitude vs. high altitude customers) not everyone receives the same benefit. Though you can minimize aggregate variability, it may be interesting to know the effect on the various customer populations to see if good compromises exist..

These considerations highlight the problem of uncertainty in the estimate of the range (variance) of each of the noise factors. Such estimates are often done without benefit of rigorous study or dedicated experiment. In spite of uncertainty in variance estimates or multiple population segments, some benefit in variance reduction is to be expected from robust design. But perhaps the improvement can be increased by studying the sensitivities involved. This prompts the construction of a Monte-Carlo experiment as a framework for a variety of purposes all related to making the robust design process itself more robust.

3 TWO CONSTRUCTED EXAMPLES

Implementing correlations in the Monte-Carlo experiment is not difficult, but we focus on a simple scenario. As a simpli-

fying assumption then, let the noise factors be uncorrelated. Let the vector of δ and matrix of γ coefficients be as in Table 2. In practice these coefficients would have been estimated from an experiment as described in the previous section. For specificity, assume it was a fractional factorial experiment in the X (design) and Z (noise) factors.

In example one, we assume that noise factor two has a nonlinear response beyond the range ($\pm \sigma$) used in the original experiment. In the second example we use the same set of coefficients in Table 2 and consider a population composed of two distinct sets of users, each of which sees a different set of variances for the first two noise factors.

Table 2: Coefficients for Examples One and Two

		$\gamma(i,j)$	X1	X2	X3
δ_1	2.0	Z1	-0.5	0.6	0.0
δ_2	2.0	Z2	0.5	0.4	0.0
δ_3	2.0	Z3	0.0	0.0	2.0

In examples one and two we use Arena to simulate the true state of affairs – some nonlinear behavior of the noise factors beyond the range used in the factorial experiment, and two distinct populations of users. Response surface explorations use a two variable central composite design represented in Arena’s Project Analyzer. Optimal settings for x_1 and x_2 are confirmed using OptQuest. The optimal settings for x_3 is obviously $x_3 = -1$ and that variable is omitted from the response surface.

Note that in this example there are no settings for x_1 and x_2 in the experimental region that can completely neutralize the effects of noise. Good settings can reduce noise, however. You can see this by plugging in the values for δ and γ into problem $P^{(s)}$ and experimenting with different choices for the design variables – the x_i .

The Arena model itself is very simple. In example two, entities from each population are created and a response for each entity is estimated. That response includes draws of random variables for each of the noise distributions being used. No restriction is imposed on the form of these distributions. A separate response surface is estimated for each population in terms of mean response and different measures of variability captured across the entities. We can then look at tradeoffs in design factor settings which lead to acceptable solutions for both populations.

In example one, the diminishing effect of noise on response is simulated by a conditional choice of the δ coefficient values – conditioned on the random variable representing that noise source. This serves to capture the effect of the nonlinear relation of noise on response that went unrecognized in the original factorial experiment.

3.1 Example One – Nonlinear Behavior of Response over Noise Source Two (Z2)

In the Arena model the noise variables are drawn from a standard normal distribution. The physical robust design

experiment would use fixed noise values of $\pm \sigma$. The numerical value of σ would be an estimate, perhaps based on some preliminary or previous experiment. This would be a standard robust design approach. In the Monte-Carlo experiment nonlinearity is introduced by changing the coefficient ($\delta_2 = 2$) of the noise variable Z2. Under linear behavior this coefficient would be constant regardless of the value of Z2. Here we impose possible breakpoints at $\pm \sigma$. Mimicking a sort of diminishing response, $\delta_2 = 3$ for $Z2 < -\sigma$ and $\delta_2 = 1$ for $Z2 > +\sigma$. So in any Monte-Carlo trial for the response, a test is made to check the value of the second noise variable and the δ coefficient is adjusted.

Four cases are considered: no breakpoints (linear behavior), breakpoints both above and below, breakpoint above, and breakpoint below. Comparison of the cases is shown here in terms of overlapping contour plots -- minimum overlapping maximum response (Figures 1 and 2) and average response overlapping standard deviation of response (Figures 3 and 4). The idea is that acceptable performance may require a combination of conditions to be met. Jointly acceptable values are in distinctly different regions depending on whether breakpoints exist.

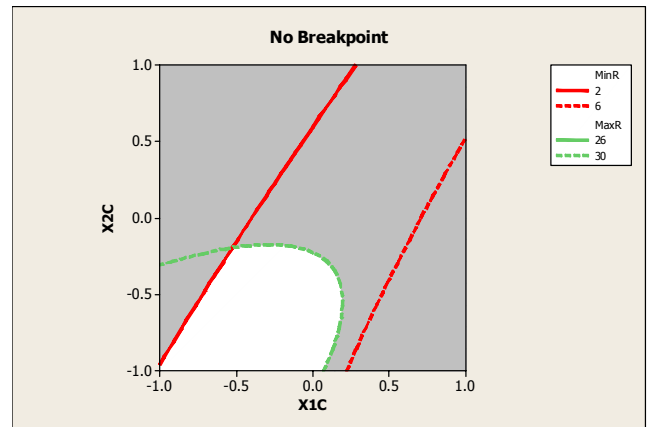


Figure 1: Acceptable Response Ranges with No Breakpoints

The non-shaded region on each plot represents acceptable minimum and maximum response for the system being designed. Acceptable ranges are defined identically for figure 1 vs. figure 2 and for figure 3 vs. figure 4. The non-shaded regions are substantially different, however.

Figures 3 and 4 illustrate a similar point regarding acceptable ranges for the mean and the standard deviation of response. Although the definition of acceptable ranges remains the same in both figures, the region that satisfies these definitions depends critically on the behavior beyond the $\pm \sigma$ points.




Figure 2: Acceptable Response Ranges with Upper Breakpoint Defined.

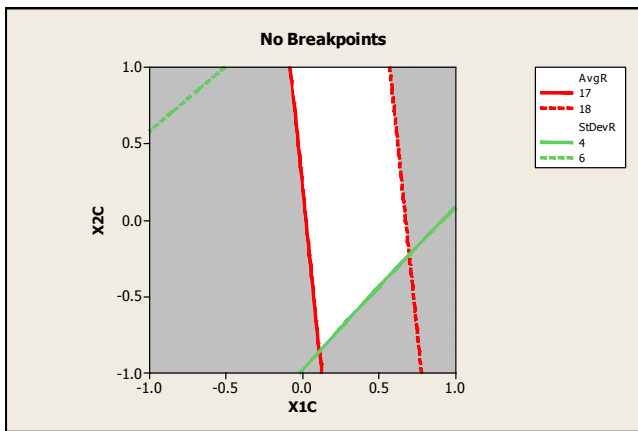


Figure 3: Acceptable Ranges for Mean and Standard Deviation of Response with No Breakpoints

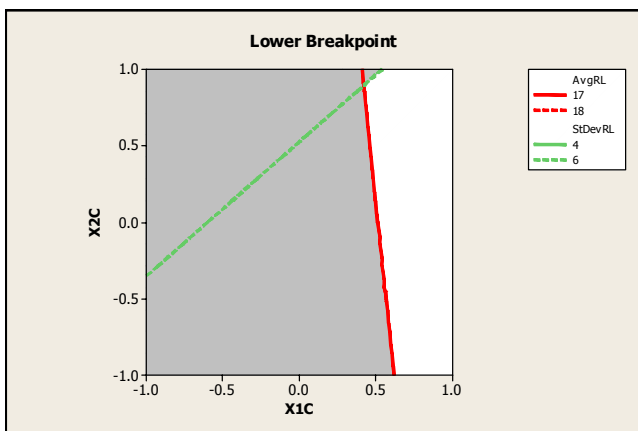


Figure 4: Acceptable Ranges for Mean and Standard Deviation of Response with Lower Breakpoint

These figures illustrate the point that what happens beyond the $\pm \sigma$ points affects how the product would **operate in the customer environment**. A simulation to test the solution against various possibilities augments the benefits and knowledge gained from the original experiment.

3.2 Example Two – Multiple Populations

Here we assume that two separate subpopulations exist. We refer to these as population one and two. Each population is characterized by the noise factor variance it will actually encounter in operation, i.e., outside the lab. We look at the behavior of each population and compare it to aggregate behavior using response surface models.

Both populations share the same set of design factor values and model coefficients. These would have been estimated in the original robust design experiment. These values are shown in Table 3. Somewhere in the design region are the values that minimize aggregate variance, i.e. if we treat this as one homogeneous population. Here we assume that population one and two each represent equal slices of the overall population. Equal numbers of population one and two entities are created in the Monte-Carlo experiment. Both populations see identical variance on noise factor three (Z3), but a different variance on the first two noise factors (Z1 and Z2).

Table 3: Variance Experienced by Population Group

	Z1	Z2	Z3
P1	3.0	1.0	2.0
P2	1.0	3.0	2.0

In this case the point being made is best illustrated by a trio of contour plots that represent the standard deviation of response for the individual populations and for the aggregate population. Contours for average response are unremarkable and reflect a gradient in the X1 direction for both populations – large main effect of the X1 design factor on mean response. These are not shown here.

In a typical robust design application one would try to minimize variance over the feasible design region while maintaining an acceptable mean response. Just how that is accomplished and how effective it will be depends on which set of contours you examine for the standard deviation. Certainly the perspective is different for each population and an alternative to using the contours for the joint or overall population is to consider these contours for the individual populations.

In this example there are no settings for X1C and X2C that will entirely zero out the variance of the response. Generally speaking, population one favors decreasing X2C and increasing X1C to reduce response variability. This choice follows from the relative size of the variances seen by that population (variance of Z1 vs. variance of Z2) and from the

particular δ and γ coefficients used in this example. The same δ and γ coefficients apply to population two, but the experienced variances are not the same – see Table 3. Thus, population two favors increasing X2C and decreasing X1C.




Figure 5: Contours for Subpopulation One Standard Deviation of Response

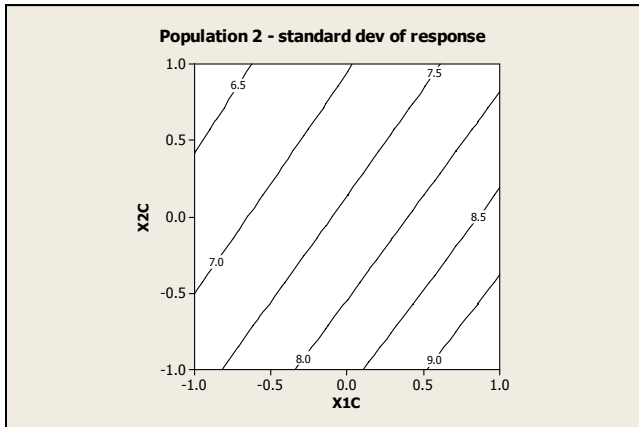


Figure 6: Contours for Subpopulation Two Standard Deviation of Response

Figure 7 shows the response surface that would be estimated if both populations are combined and treated as one undifferentiated population. This solution favors smaller values of X1C and X2C. The compromise hides some useful information on the particulars for the two populations. The search for a good operating point could consider the average response, as well as the separate models of standard deviation for both populations.

4 DISCUSSION AND CONCLUSIONS

Robust design searches for parameter settings that make products and processes immune to noise sources. It is an opportunity to guarantee consistent performance, possibly

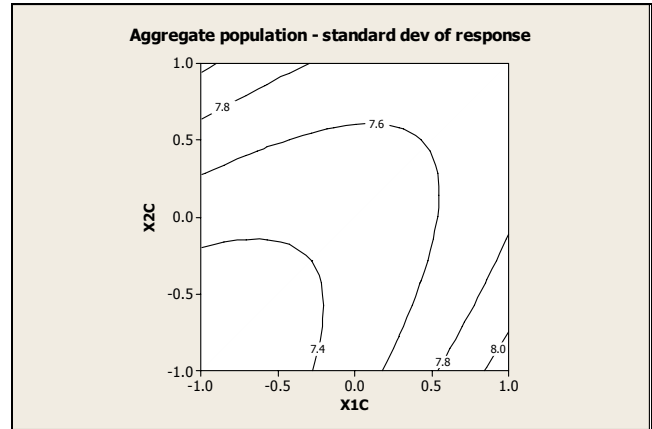


Figure 7: Combined Population Contours for Experienced Variation

at little or no extra cost. Robust design looks for interactions between noise sources and design factors. Advantageous settings of design factors neutralize the effect of noise.

The wild card in this methodology is that noise sources will not be under controlled laboratory conditions in the user environment. If extensive experiments to model noise are not to be done and if noise sources are not thoroughly understood, then sensitivity analysis ought to look at alternate models.

Simulation can remove some of the assumptions needed in an analytic approach to optimizing design over noise. Two examples are shown here. The first looks at what happens beyond the noise values used in the original design experiment, the second looks at what happens when the overall population has some distinct subpopulations.

This is not the end of the possibilities. Other interesting cases that could be examined are interactions among the noise variables and different distributions for the noise sources. Even a simple sensitivity analysis on misspecification of the $\pm \sigma$ points for each noise source is useful. All of these excursions can be built into a single Monte-Carlo simulation model which centers on the parameters estimated in the original robust design physical experiment.

In the cases reviewed here, a look at multiple populations reveals behavior very different than the aggregate. Subpopulations might be expected to exist when noise factors represent environmental conditions, for instance. The more detailed model can reveal whether different versions of the product or operating instructions are necessary, or whether a single version can accommodate all users. The first example looks at how the solutions are affected by more complicated behavior, e.g., nonlinearity, outside the limits of the laboratory experiment.

A simulation model represents an inexpensive way to do this sensitivity analysis. It can answer the question of

how robust a design is with respect to the specification of the factors in the robust design experiment. Considering the expense of a physical experiment the simulation represents a flexible, straightforward way to do sensitivity analysis. Future work will lead to development of a more general test bed for the Monte-Carlo experiment.

REFERENCES

- Borkowski, J., and Lucas, J. (1997). Designs of Mixed Resolution for Process Robustness Studies. *Technometrics* 39: 63-70.
- Box, G., Bisgaard S., and Fung, C. (1988). An Explanation and Critique of Taguchi's Contribution to Quality Engineering. *Quality and Reliability Engineering International*.
- Koksoy, O., and Doganaksoy, N. (2003). Joint Optimization of Mean and Standard Deviation Using Response Surface Methods" *Journal of Quality Technology*, 35: 239-252.
- Miller A., and Wu, C.F.J. (1996). "Parameter Design for Signal-Response Systems: A Different Look at Taguchi's Dynamic Parameter Design". *Statistical Science*, 11: 122-136.
- Miller A. (2002). Analysis of Parameter Design Experiments for Signal-Response Systems. *Journal of Quality Technology*, 34: 139-151.
- Nair, V. N. (1992). Taguchi's Parameter Design: A Panel Discussion. *Technometrics* 34: 127-161.
- Phadke, Madhav S. (1989). *Quality Engineering Using Robust Design*. Prentice Hall, Englewood Cliffs, New Jersey.

AUTHOR BIOGRAPHY

JOSEPH R. LITKO is an associate professor in the Engineering and Management Systems Department at the University of Dayton. His research interests include robust parameter design experiments and the application of robust techniques to optimization and simulation models. He is a member of INFORMS and ASQ. His e-mail address is litko@udayton.edu.