Center for Quality and Productivity Improvement
University of Wisconsin
610 Walnut Street
Madison, Wisconsin 53705
(608) 263-2520
(608) 263-1425 FAX
quality@engr.wisc.edu

Report No. 178

Split-Plots For Robust Product And Process Experimentation

George Box and Stephen Jones

April 2000

The Center for Quality and Productivity Improvement cares about your reactions to our reports. Please direct comments (general or specific) to: Reports Editor, Center for Quality and Productivity Improvement, 610 Walnut Street, Madison, WI 53705; (608) 263-2520. All comments will be forwarded to the author(s).

Split-Plots For Robust Product And Process Experimentation

George Box

Center for Quality and Productivity Improvement
University of Wisconsin-Madison
Madison, Wisconsin 53706

Stephen Jones

The Boeing Company
PO Box 3707 MC7L-22
Seattle, Washington 98124

ABSTRACT

Environmentally robust products and processes are designed to be insensitive to variation over the relevant ranges of the environmental conditions in which they need to operate. Split plots frequently provide efficient experimental arrangements whereby environmentally robust process and product designs may be achieved. The various questions that arise in planning and analyzing such experiments are discussed and illustrated with examples.

Split-Plots For Robust Product And Process Experimentation

George Box

Center for Quality and Productivity
Improvement
University of Wisconsin—Madison
Madison, Wisconsin 53705

Stephen Jones

The Boeing Company
PO Box 3707 MC7L-22
Seattle, Washington 98124

Environmentally robust products and processes are designed to be insensitive to variation over the relevant ranges of the environmental conditions in which they need to operate. Split plots frequently provide efficient experimental arrangements whereby environmentally robust process and product designs may be achieved. The various questions that arise in planning and analyzing such experiments are discussed and illustrated with examples.

Keywords: split plots, robust design, environmental sensitivity, choice of experimental arrangements.

Introduction.

Environmental robustness* and optimality are important concerns for improving current products and processes and for designing new ones. Since these two characteristics can rarely be achieved simultaneously, we must aim for a design that is *robust* and *good*. That is to say a design that, although not necessarily best for any particular artificially set of the environmental conditions, is reasonably effective over those conditions in which the product or process will need to operate. In this column we explore the use of split plot and related experimental arrangements to achieve such designs. For a fuller discussion see e.g. Box and Jones, 1992a and 1992b, and the references contained therein.

* A different problem concerns minimizing sensitivity of the performance of an assembled product to variation transmitted by its components. Ways to achieve this kind of robustness were discussed, for example, in a previous column by Box and Fung, 1994 and in the references it contains.

A Reminder About Split Plots.

In an earlier column (Box, 1996) it was explained that for many experiments conducted in industry a completely random experimental arrangement is impractical. This is usually because certain factors (variables) are much more difficult to change than others. In such circumstances split plot experiments can be employed that are often very efficient for exploring the main issues and are much more easy to run.

An example was an experiment to improve the corrosion resistance of steel bars by first applying a surface coating and then baking the bars in a large furnace. In the experiment four different coatings were to be compared at three different temperatures. A completely random arrangement in which the furnace temperature had to be reset repeatedly would have been impractical but a split plot experiment was easily run in the following way: a temperature condition was set and the four different kinds of coated bars were randomly positioned in the furnace and baked for the prescribed length of time. The furnace was then reset at a different temperature to bake a second set of coated bars and so on.

Split plot arrangements were first developed in agriculture, (see Fisher, 1935; Yates, 1937). Using their original terminology repeated furnace heats would correspond to “main plots” and coatings applied to the bars within a single heat to “subplots”. In such an experiment it is essential to understand that there will be two distinct sources of error that affect the results. In terms of this example, the main plot error is associated with the variation between furnace heats at the same temperature setting, and the subplot error with the variation between coatings within a given heat. It is typical of such experiments

that the main plot variation (between heats) can be much larger than the subplot variation (between coatings within heats). Consequently, by using a split plot arrangement it would usually be possible to compare the effectiveness of the different coatings with greater accuracy than with a fully randomized arrangement. This is because, in this arrangement, not only the main effects for coatings, but also *all the interactions of coatings with temperature*, would be estimated with the same smaller within-heats subplot error. Only temperature effects would be estimated with the larger main plot error. It is this property that makes such arrangements of particular value for robustness studies.

Designing products robust to environmental factors using split plots.

A concept described by Michaels (1964), that went largely unnoticed at that time, was the use of split plot experiments to design products robust to environmental variation. This idea of using statistically planned experiments to achieve such robustness was later brought to the attention of many engineers and statisticians by Genichi Taguchi (see Taguchi and Wu, 1980) with many cogent industrial illustrations. However, the importance of exploiting split plot structure to minimize experimental effort and increase statistical efficiency and also the profound effect that such structure has on the correct analysis of the data, were not discussed.

In Michael's example of a split plot design to achieve environmental robustness the product was a detergent to be designed (formulated) so that its washing ability was not much affected by the hardness of the water, the temperature of the wash, etc. The product could equally well be an ignition system to be designed so that its performance was insensitive to the humidity and temperature of the intake air. Also for several

decades the U.S. food industry have used experiments such as is shown in Table 1 to obtain recipes that were robust to departures from recommended cooking instructions. In the discussion that follows whether the objective is to obtain a robust *formulation* for a detergent, a robust *design* for an ignition system, or a robust *recipe* for a cake mix we will in every case describe this activity by the single word *design*. Also, to avoid confusion the structure of the data collection strategy will be described by the words *experimental arrangement*.

In the somewhat simplified example for which the data are shown in Table 1, the manufacturer was seeking a design for a boxed cake mix that would be good and also robust to deviations from the recommended baking temperature (T) and the time of cooking (t). Thus, in this example, T and t were the environmental factors arranged in a 2^2 factorial. The design factors that the manufacturer had under his control—the amounts of Flour (F), Shortening (S) and Egg powder (E) included in the mixture—were arranged in a 2^3 factorial. All 32 combinations of the 8 designs in the 4 environments were run and the results shown are average “hedonic” scores, on a scale of 1 – 7, obtained from a taste panel and measuring “how good” the cakes tasted.

Table 1. Data for the cake mix example

Recipe	Design Variables			Environmental Variables						Std.		Quadrati	
										c			
	<i>F</i>	<i>S</i>	<i>E</i>	<i>T</i> : —	+	—	+	Ave	Range	Added Cost	Std. Dev.	Quadratic Loss	
(1)	—	—	—		1.1	1.4	1.0	2.9	1.6	1.9	0	0.9	119.0
(2)	+	—	—		1.8	5.1	2.8	6.1	4.0	4.3	2	2.0	49.1
(3)	—	+	—		1.7	1.6	1.9	2.1	1.8	0.5	3	0.2	107.3
(4)	+	+	—		3.9	3.7	4.0	4.4	4.0	0.7	5	0.3	36.3
(5)	—	—	+		1.9	3.8	2.6	4.7	3.2	2.8	4	1.2	60.9
(6)	+	—	+		4.4	6.4	6.2	6.6	5.6	2.2	6	1.0	10.5
(7)	—	+	+		1.6	2.1	2.3	1.9	2.0	0.7	7	0.3	101.3
(8)	+	+	+		4.9	5.5	5.2	5.7	5.3	0.8	9	0.4	11.6

F, flour; S, shortening; E, egg; T , temperature; t , time.

As is later discussed the efficiency of the experiment and the analysis of the subsequent data depend critically on how the experiment is actually carried out. We will suppose that each of the eight separate designs (cake mixes) was individually prepared and then subdivided into four batches which were tested at the four environmental conditions of baking time and temperature. In this arrangement therefore the designs were the main plots and the environmental combinations were the subplots.

A first appraisal of the data is gained from Table 1 by looking at the *Average*, *Range* and *Additional Cost* for each of the eight product designs*. The average supplies a measure of the overall goodness of the design and the range is a measure of its robustness over the combinations of the environmental factors tested. The numbers in the "additional cost" column are in fractions of a penny, treating design (1) as a base.

We see, for example, that design number (3) appears to be very robust but uniformly inedible; design number (6) is good on the average but is not particularly robust; design (8) is essentially as good as (6) and is robust (but is the most expensive).

On (more substantial) data and analysis of this kind, management could make informed choices between alternatives that took account of many subjective factors (the perceived quality of the competition, the relative profitability of this particular product etc.).

A Closer Analysis.

* In addition to the average ($\bar{y}$) and the range, the estimated standard deviation ($\hat{\sigma}$) is also shown and may be used instead of the range if preferred. The results are essentially the same. Also shown is the "Quadratic Loss" $L = \sum (7 - y)^2$. This is a somewhat arbitrary combination of $\bar{y}$ and s ; thus this example $L = 4\{(7 - \bar{y})^2 + \frac{3}{4}\hat{\sigma}^2\}$. We think we are better off looking at $\bar{y}$ and $\hat{\sigma}$ separately.

The main effects and interactions for the design and the environmental factors can be obtained by analyzing the data of Table 1 as a 2^5 factorial in the factors F , S , E , T , and t . These calculated effects may then be conveniently arranged in the 8×4 table as shown below in which the symbol I is used to indicate a mean value, see e.g. Box, Hunter and Hunter (1978).

Table 2. Calculated effects for the cake mix data

		Environments			
		I	T	t	Tt
Designs	I	3.45	1.11	0.53	0.07
	F	2.57	0.31	0.01	-0.08
	S	-0.33	-0.92	-0.22	-0.08
	E	1.21	-0.03	-0.08	-0.24
	FS	0.19	-0.17	0.01	0.22
	FE	0.31	-0.26	-0.08	0.08
	SE	-0.43	0.14	0.02	0.01
	FSE	0.28	0.37	0.07	-0.02

We have seen that an important characteristic of split plot arrangements is that the subplot effects and *all their interactions* with the main plot effects are estimated with the same smaller subplot error. Only the main plot effects have the main plot error. Consequently, if the data from a split plot arrangement are analyzed graphically *two* separate Normal plots are needed. In the cake experiment the design factors are associated with the main plots and the environmental factors with subplots. The seven design effects F , S , E , FS , FE , SE , and FSE which have the main plot error, have therefore been graphed in Figure 1(a). The remaining 24 effects associated with the environmental factors and their interactions which have the subplot error are graphed separately in Figure 1(b).

Figure 1. Normal plots: (a) for main plot effects; (b) for subplot effects.

From the gradients of the “error lines” in these two normal plots we see that, as expected, in this experiment the subplot error is considerably smaller than the main plot error. With this arrangement in which main plots are associated with designs, Figure 1(a) tells us which factors control how *good* on the average are the various designs. Thus we see that higher values for Flour (*F*) and Egg Powder (*E*) give cakes that on the whole are better tasting. Figure 1(b) shows which factors affect robustness. If a robust design can be achieved it will be by using the interactions between the design and environmental factors as “levers” to reduce such sensitivity. Suppose the factors are scaled so that minus and plus signs correspond to -1 and $+1$ in the coded variables and that there is an environmental effect *E* that we wish to modify and a design factor *D* that has an

interaction DE with E . Then by setting* the design factor D to a value d , we can change E to a modified value E_m in accordance with the equation

$$E_m = E + (DE \times d). \quad (1)$$

In the cake example, we would like the effect of T , t , and Tt to be as small as possible. But from Figure 1(b) we see that T is not small indicating that the taste of the cake is highly sensitive to baking temperature—an effect clearly seen in the original data. Fortunately there is also a large ST interaction. This can be employed to modify sensitivity to temperature to a value T_m by altering the level of shortening to a value s in accordance with the equation

$$T_m = T + (ST \times s) \quad (2)$$

$$T_m = 1.11 + (-0.92 \times s),$$

where s denotes the amount of shortening on the coded scale.

Ideally then, and if there were no experimental error, sensitivity to temperature, as measured by T_m , could be reduced to zero by setting the design variable S to the value

$$s = \frac{1.11}{0.92} = 1.2.$$

The value $s = 1.2$ is close to $s = 1.0$ and by comparing results in the rows of Table 1 we see that indeed the saw-tooth pattern associated with temperature sensitivity in runs 1, 2, 5 and 6 (when S is at its minus level) essentially disappears in runs 3, 4, 7 and 8 (when S is at its plus level). Finally then, it is very logical that by setting F and E at their

* For the argument, necessary assumptions, and generalizations associated with this concept, see Box and Jones, 1992a; Box and Jones, 1992b; and Box, 1999.

plus levels to get high response, and S at its plus level to minimize sensitivity to temperature, we obtain a cake mix (8) that is both *robust*, as well as *good*.

In this example the environmentally sensitive factor T was not affected by the design variables E and F that changed the average response. When this happy separation does not occur some compromise solution will be necessary.

An estimate of the subplot standard deviation can be obtained from the twelve design $\times$ environment interactions between three factors or more shown below the dotted line in Table 2. The estimate obtained from the average value of the squares of these elements is 0.16.

What? How? & Why?

As always, when considering the results of an experiment we should have in mind the questions: *What* happened? *How* did it happen? *Why* did it happen?

What the various tested cake designs do when exposed to different environmental conditions is clearly seen in Table 1.

How, for example, sensitivity to baking temperature may be reduced is shown by Equation (2).

It would remain to check with the subject matter specialist (the food technologist in this case) *why* would there be such an effect. Good questions are: Would you have expected this? If yes, please explain why. If no, then what do you think might be going on? If your speculation is right and might lead to new ideas for developing the product can we check it out and perhaps move on from there?

Choosing the Best Experimental Arrangement.

Using the cake mix experiment for illustration we have so far supposed that the experiment was conducted with the ($n = 8$) designs as main plots and the ($m = 4$) environments as subplots. We will call this arrangement (a). The types and numbers of operations with this and other possible arrangements are summarized in Table 3.

Table 3. Four possible arrangements for the cake mix example.

		Example		General	
		Mixes	Bakes	Design Models	Environments
Split Plots {	Fully Randomized	32	32	$n \times m$	$n \times m$
	Bakes are subplots (a)	8	32	n	$n \times m$
	Designs (recipes) are subplots (b)	32	4	$n \times m$	m
	Strip Block (c)	8	4	n	m

In a fully randomized design 32 cake mixes would have been prepared independently and each one baked at the prescribed conditions of temperature and time, resulting in 32 separate bakes.

In the split plot arrangement of type (a) already described, 8 separate mixes would be prepared for each of the various combinations of Flour (F), Shortening (S), and Egg (E). Each mix would then be divided into four and the resulting 32 mixes baked independently. In this arrangement the oven is operated 32 times.

In a split plot design arrangement of type (b) 32 cake mixes would be made up independently. These would then be separated into 4 sets each containing all 8 combinations of F , S , and E . The first set of 8 would then all be baked together, with the bake temperature (T) and bake time (t) chosen randomly from the four combinations of

levels. The next set of 8 would be run at a second randomly selected time and temperature, and so on. So in this arrangement the oven is operated only four times.

In the Strip Block arrangement (c), a batch large enough to make four cakes would be prepared for each of the 8 combinations of F , S , and E and a complete set of these eight combinations would be baked at each of the four conditions of Temperature and Time. Again in this arrangement the oven is operated only four times.

It is not necessary, of course, that the design array or the environmental array should themselves be factorials. You could, for example, just test *any* set of product designs in *any* set of environments in a cross product array.

Considerations in Choosing the Experimental Configuration.

Statistical considerations: It is important to be able to study the Design $\times$ Environment interactions individually and it should be noted that *whichever* split plot configuration is used these will all be estimated with the subplot error which is often relatively small.

Split plot arrangement (a) used for the cake mix experiment has the additional advantage that the main effects of all the environmental variables are also estimated with the subplot error. Thus, for example, in relationships such as Equation 2 the coefficients T and ST are both estimated with the smaller subplot error. This advantage is of even greater importance when there is more than one important environmental effect to be minimized and/or more than one design variable that can be adjusted. In such cases Equation 1 is replaced by a set of linear equations (see e.g. Box and Jones, 1992a; Box, 1999). If arrangement (a) is used then all the coefficients in these equations will be estimated with the same subplot error.

We emphasize that it is the manner in which the experiment is actually carried out that decides the appropriate analysis. However, the *numerical* quantities that are needed for *any* analysis of the cake mix data are those in Table 2.

If the experiment had been conducted with split plot arrangement (b) the 7 design effects and the 21 design $\times$ environment interactions would have the subplot error and should appear on the same normal plot. The three environmental affects would have the main plot error.

If the strip block of arrangement (c) had been used there would have been three separate sources of error to consider and if normal plots were used there would be three different plots needed for the analysis: one for the design effects, one for the environmental effects, and one for the design $\times$ environment interactions. The last of these would usually have the smallest error term.

In the paper by Box and Jones (1992b) there is a discussion of detailed models, analyses of variance and the statistical efficiency for each of the various arrangements.

Experimental effort and cost: In practice, as well as statistical considerations, the cost and difficulty of experimentation must be taken into account in deciding whether subplots are associated with the environmental conditions or with the product design factors.

For example, if we were developing a copying machine, arrangement (a) would almost certainly be best from a practical point of view because the production of a number of different types of copier (product designs) would be expensive; but running these copiers at different environmental conditions (ambient temperature, humidity, type of paper, etc.) would be easy and inexpensive. Furthermore if factors, such as

temperature and humidity were the only environmental variables these might be changed en masse in a specially fitted room (a biotron) in which all the copiers could be tested together at each environmental condition in a strip block design. This arrangement is usually so easy to run that a number of repeats in which the environmental conditions are replicated can be made inexpensively thus giving increased precision for the estimated effects and the design $\times$ environment interactions and allowing better estimation of the sources of error.

On the other hand, if the experiment, referred to earlier, to test the corrosion resistance of coated steel bars was part of a robustness study, then it would be the environmental factors (the furnace temperatures) that were difficult and expensive to change. These would therefore be chosen as main plot factors and the product designs (the various coatings applied to the steel bars) would be the subplots; and arrangement (b) would be the preferred arrangement.

The same principles apply to robust *process* design. For example, one of the authors was recently involved in a robust process experiment using arrangement (a). In this study experiment parts were treated in a heated chamber that had gas pumped into it. The problem was to minimize processing differences arising from different orientations and different locations of the parts within the chamber. The design variables included hold temperature, dwell time, gas flow rate, and removal temperature. The environmental variables were the orientation of the part and the location of the part within the chamber. The robust process design finally arrived at allowed the chamber to be run at its maximum capacity.

Extreme Conditions: One technique for reducing the number of environments to be tested that has sometimes been used is to run experiments only at what are judged to be some of the severest environmental conditions. For example, for the cake mix experiment it is reasonable to assume that the effect of baking time and baking temperature would to some extent be mutually compensating. Thus, a mistake in using too high a baking temperature might be roughly compensated by using a shorter baking time. Thus by running the environmental factors only at the combinations $(T, t) = (-, -)$, and $(+, +)$, representing severest conditions, the robustness of the various product designs might be assessed with fewer runs. The same kind of reasoning can be used in other technologies, but the disadvantage of this technique is that it is necessary for the food technologist, engineer, or chemist to be able to guess in advance what the extreme environmental combinations will be—a task that becomes more and more difficult as the number of environmental factors increases. Thus, this approach has serious limitations.

Finding Active Factors: In some instances one of the most difficult problems is to find out *which* of a large number of environmental effects are important and which of a large number of design factors might eliminate them by the use of suitable interactions.

For this purpose, fractional factorials or other orthogonal arrays may be used to carry both the design variables and the environmental variables. For example, using the same overall number of runs as in Table 1, seven design factors might have been accommodated by using a saturated 2^{7-4} fractional factorial for the design array, and three environmental variables by using a saturated 2^{3-1} fractional factorial for the environmental array. Notice that on the assumption that interactions of three or more factors can be ignored, all 21 of the design $\times$ environment interactions could still be

estimated without bias. The main effects for the design variables will then of course be aliased with their two factor interactions, and similarly for the environmental variables. However, in two replications, the signs for the environmental variables could be switched. Then making use of the foldover principle for the complete replicated design, and on the same assumptions concerning higher order interactions, all main effects for the environmental variables would be free of aliases (see for example Box, Hunter and Hunter, 1978, and the references therein). Bear in mind however, that environmental sensitivity might occur because of an interaction between environmental variables; such possible interactions may therefore need to be isolated. Also, while in principle the design variables could be similarly switched in the second replication, this might not be easy to do since it would involve using a further set of different designs that were independently built. The use of fractional factorials and orthogonal arrays would be particularly attractive when the more easily performed strip block arrangements could be used. If these designs are replicated, “partial confounding” could be employed to better isolate interactions. (See e.g. Fisher, 1935; also Bisgaard, 1992 and Kulahci, 1999.)

Finding the Critical Factors: Perhaps most important is to *find* the environmental factors that need to be desensitized and to discover design factors that will interact with them to achieve desensitization. When fractional factorials and Plackett Burman designs are employed to carry the environmental and design factors, they can be used as screens to search both for environmentally sensitive effects and for appropriate design variables that can neutralize them. The projective properties of these designs can then be particularly useful to isolate main effects and their interactions for small subgroups of the important

environmental and design factors. (see e.g. Lin and Draper, 1992; Box and Bisgaard, 1993; Box and Tyssedal, 1996).

When there is one or more environmental factors that have not been neutralized with design factors so far tested, it will be important to make clear to the subject matter specialist the nature of the problems that remain. This is vital to produce new ideas.

A Further Example.

Murphy never sleeps and planned experimentation provides him with an almost irresistible temptation. Fortunately, however, sometimes even an imperfect experiment can yield valuable results. This was true in a study conducted by John Shuerman and Todd Algrim, engineers with DRG Medical Packaging Company in Madison, WI.

In this investigation, the company sought to design a package that would give a good seal over the wide range of sealing process conditions used by its customers. This packaging manufacturer identified five product design factors that might affect the quality of the seal. These were the paper supplier (*PS*), the solids content (*SC*), the cylinder type (*CY*), the oven temperature (*OT*) and the line speed (*LS*), each to be tested at the two levels shown in Table 4. The customer had identified two environmental factors—the sealing temperature (*ST*), which was set at three levels, and the sealing dwell time (*SD*), at two levels.

Table 4. Design factors and their

Factors	Levels	
	—	+
Paper Supplier (<i>PS</i>)	UK	USA
Contents of Solids (<i>SC</i>)	28%	30%
Cylinder (<i>CY</i>)	Fine	Coarse
Oven Temperature (<i>OT</i>)	180°C	215°C
Line Speed (<i>LS</i>)	Low	High

The experiment was conducted in the following manner. An eight-run fractional factorial experiment was run using the five packaging factors, and large quantities of the packaging material were manufactured at each of the eight conditions. The eight batches of material were then shipped to the customer's plant and each batch was subdivided into six sub-batches. The sub-batches were then used to run a 2×3 experiment in the two sealing process factors. This was therefore a split plot arrangement of type (a) in which the designs were main plots and the environments subplots.

Table 5 shows the observed number of defective seals from a sample of 15. This experiment was imperfectly carried out because the two levels of one of the factors were inadvertently switched so that the eight-run experiment as actually carried out was unbalanced. Also there were a number of missing observations in the first run in Table 5 and a complete loss of information for the second run. Nevertheless, it is easy to see from visual examination of the data that product design number (5) behaved almost perfectly under all environmental conditions.

The complete analysis indicated how the manufacturer could produce packaging material that would give a sterile seal under a range of sealing process conditions. This and other experiments led to discoveries that turned out to be worth several million dollars.

Table 5. Number of defects in packaging experiment.

Product Design	Design Factors					<i>ST</i> <i>SD</i>	Environmental Factors					
	<i>PS</i>	<i>SC</i>	<i>CY</i>	<i>OT</i>	<i>LS</i>		—	0	+	—	0	+
(1)	+	—	—	+	+		8	4	*	0	*	7
(2)	+	—	—	+	—		*	*	*	*	*	*
(3)	—	—	+	—	+		3	2	1	1	0	0
(4)	—	—	+	—	—		7	1	1	2	0	4
(5)*	+	+	+	—	+		0	0	0	0	0	1
(6)	—	+	+	+	+		4	0	0	1	1	1
(7)	—	+	—	+	—		9	1	0	4	5	1
(8)	+	+	—	—	—		2	1	3	1	0	0

*Missing observations.

Conclusions.

It is evident that split plot arrangements provide a very efficient means of designing robust products and processes. The configuration of such designs must be made to fit experimental circumstances.

Acknowledgements.

This research was supported by the National Science Foundation under Grant DMI-9812839.

References.

- Bisgaard, S. (1992), "The Design and Analysis of $2^{k-p} \times 2^{q-r}$ Inner and Outer Array Experiments", Technical Report No. 90, *Center for Quality and Productivity Improvement*, University of Wisconsin-Madison.
- Box, G.E.P. (1996), "Split Plot Experiments", *Quality Engineering*, **8**, 515 – 520.
- Box, G.E.P. (1999), "Statistics as a Catalyst to Learning by Scientific Method, Part II, A Discussion", *Journal of Quality Technology*, **31**(1), 16 – 72.

- Box, G.E.P. and Bisgaard, S. (1993), "What Can You Find Out From 12 Experimental Runs?", *Quality Engineering*, **5**, 663 – 668.
- Box, G.E.P., Hunter, W.G. and Hunter, J.S. (1978) (BH^2), *Statistics for Experimenters*, John Wiley, New York.
- Box, G.E.P. and Jones, S. (1992a), "Designing Products that are Robust to the Environment", *Journal of Total Quality Management*, **3**(3), 265 – 282.
- Box, G.E.P. and Jones, S. (1992b), "Split-Plot Designs for Robust Product Experimentation", *Journal of Applied Statistics*, **19**(1), 3 – 26.
- Box, G.E.P. and Tyssedal, J. (1996), "Projective Properties of Certain Orthogonal Arrays", *Biometrika*, **83**(4), 950 – 955. Also CQPI Report #116, May 1994.
- Fisher, R.A. (1935), *The Design of Experiments*, Oliver and Boyd, Edinburgh and London.
- Kulahci, M. (1999), *Applications of Plackett and Burman Designs to Split Plot Designs*, Unpublished Thesis, University of Wisconsin—Madison.
- Lin, K.J. and Draper, N.R. (1992), "Projection Properties of Plackett and Burman Designs", *Technometrics*, **34**(4), 423 – 428.
- Michaels, S.E. (1964), "The Usefulness of Experimental Design" (with discussion), *Journal of Applied Statistics*, **13**(3), 221 – 235.
- Taguchi, G. and Wu, Y. (1980), *Introduction to Off-Line Quality Control*, Central Japan Quality Control Association, Nogoya, Japan.
- Yates, F. (1937), *The Design and Analysis of Factorial Experiments*, Imperial Bureau of Soil Science, Harpenden, U.K.