
Chapter 1

Introduction

Many of the fundamental ideas and principles of experimental design were developed by Sir R. A. Fisher at the Rothamsted Experimental Station (Fisher, 1926). This agricultural background is reflected in some terminology of experimental design that is still being used today. Agricultural experiments are conducted, e.g., to compare different varieties of a certain crop or different fertilizers. In general, those that are under comparison in an experiment are called *treatments*. Manufacturing processes in industrial experiments and drugs in pharmaceutical studies are examples of treatments. In an agricultural experiment, the varieties or fertilizers are assigned to plots, and the yields are compared after harvesting. Each plot is called an *experimental unit* (or unit). In general, an experimental unit can be defined as the smallest division of the experimental material such that different units may receive different treatments (Cox, 1958, p. 2). At the design stage, a treatment is chosen for each experimental unit.

One fundamental difficulty in such *comparative experiments* is inherent variability of the experimental units. No two plots have exactly the same soil quality, and there are other variations beyond the experimenter's control such as weather conditions. Consequently, effects of the treatments may be biased by uncontrolled variations. A solution is to assign the treatments randomly to the units. In addition to guarding against potential systematic biases, *randomization* also provides a basis for appropriate statistical analysis.

The simplest kind of randomized experiment is one in which treatments are assigned to units completely at random. In a *completely randomized experiment*, the precision of a treatment comparison depends on the overall variability of the experimental units. When the experimental units are highly variable, the treatment comparisons do not have good precision. In this case, the method of *blocking* is an effective way to reduce experimental error. The idea is to divide the experimental units into more homogeneous groups called blocks. When the treatments are compared on the units within each block, the precision is improved since it depends on the smaller within-block variability.

Suppose the experimental units are grouped into b blocks of size k . Even though efforts are made for the units in the same block to be as alike as possible, they are still not the same. Given an initial assignment of the treatments to the bk unit labels based on statistical, practical and/or other considerations, randomization is carried

out by randomly permuting the unit labels within each block (done independently from block to block), and also randomly permuting the block labels. The additional step of randomly permuting block labels is to assure that an observation intended on a given treatment is equally likely to occur at any of the experimental units.

Under a completely randomized experiment, the experimental units are considered to be unstructured. The structure of the experimental units under a block design is an example of *nesting*. Suppose there are b blocks each consisting of k units; then each experimental unit can be labeled by a pair (i, j) , $i = 1, \dots, b$, $j = 1, \dots, k$. This involves two factors with b and k levels, respectively. Here if $i \neq i'$, unit (i, j) bears no relation to unit (i', j) ; indeed, within-block randomization renders positions of the units in each block immaterial. We say that the k -level factor is nested in the b -level factor, and denote this structure by b/k or block/unit if the two factors involved are named “block” and “unit,” respectively.

Another commonly encountered structure of the experimental units involves two blocking factors. For example, in agricultural experiments the plots may be laid out in rows and columns, and we try to eliminate from the treatment comparisons the spatial variations due to row-to-row and column-to-column differences. In experiments that are carried out on several different days and in several different time slots on each day, the observed responses might be affected by day-to-day and time-to-time variations. In this case each experimental run can be represented by a cell of a rectangular grid with those corresponding to experimental runs taking place on the same day (respectively, in the same time slot) falling in the same row (respectively, the same column). In general, suppose rc experimental units can be arranged in r rows and c columns such that any two units in the same row have a definite relation, and so do those in the same column. Then we have an example of *crossing*. This structure of experimental units is denoted by $r \times c$ or row \times column if the two factors involved are named “row” and “column,” respectively. In such a *row-column experiment*, given an initial assignment of the treatments to the rc unit labels, randomization is carried out by randomly permuting the row labels and, independently, randomly permuting the column labels. This assures that the structure of the experimental units is preserved: two treatments originally assigned to the same row (respectively, column) remain in the same row (respectively, column) after randomization.

For example, suppose there are four different manufacturing processes compared in four time slots on each of four days. With the days represented by rows and times represented by columns, a possible design is

1	2	3	4
2	1	4	3
3	4	1	2
4	3	2	1

where the four numbers 1, 2, 3, and 4 are labels of the treatments assigned to the units represented by the 16 row-column combinations. We see that each of the four numbers appears once in each row and once in each column. Under such a design, called a *Latin square*, all the treatments can be compared on each of the four days

as well as in each of the four time slots. If the random permutation is such that the first, second, third, and fourth rows of the Latin square displayed above are mapped to the first, fourth, second, and third rows, respectively, and the first, second, third, and fourth columns are mapped to the fourth, third, first, and second columns, respectively, then it results in the following Latin square to be used in actual experimentation.

3	4	2	1
1	2	4	3
2	1	3	4
4	3	1	2

The structures of experimental units are called *block structures*. Block and row-column designs are based on the two simplest block structures involving nesting and crossing, respectively. Nelder (1965a) defined *simple block structures* to be those that can be obtained by iterations of nesting (/) and crossing (\times) operators. For example, $n_1/(n_2 \times n_3)$ represents the block structure under a *nested row-column design*, where $n_1 n_2 n_3$ experimental units are grouped into n_1 blocks of size $n_2 n_3$, and within each block the $n_2 n_3$ units are arranged in n_2 rows and n_3 columns. Randomization of such an experiment can be done by randomly permuting the block labels and carrying out the appropriate randomization for the block structure $n_2 \times n_3$ within each block, that is, randomly permuting the row labels and column labels separately. In an experiment with the block structure $n_1/(n_2/n_3)$, $n_1 n_2 n_3$ experimental units are grouped into n_1 blocks, and within each block the $n_2 n_3$ units are further grouped into n_2 smaller blocks (often called *whole-plots*) of n_3 units (often called *subplots*). To randomize such a *blocked split-plot experiment*, we randomly permute the block labels and carry out the appropriate randomization for the block structure n_2/n_3 within each block, that is, randomly permute the whole-plot labels within each block, and randomly permute the subplot labels within each whole-plot. Note that $n_1/(n_2/n_3)$ is the same as $(n_1/n_2)/n_3$.

In general, randomization of an experiment with a simple block structure is carried out according to the appropriate randomization for nesting or crossing at each stage of the block structure formula.

Like experimental units, the treatments may also have a structure. One can compare treatments by examining the pairwise differences of treatment effects. When the treatments do not have a structure (for example, when they are different varieties of a crop), one may be equally interested in all the pairwise comparisons of treatment effects. However, if they do have a certain structure, then some comparisons may be more important than others. For example, suppose one of the treatments is a control. Then one may be more interested in the comparisons between the control and new treatments.

In this book, treatments are to have a *factorial structure*: each treatment is a combination of multiple factors (variables) called *treatment factors*. Suppose there are n treatment factors and the i th factor has s_i values or settings to be studied. Each of these values or settings is called a *level*. The treatments, also called *treatment com-*

binations in this context, consist of all $s_1 \cdots s_n$ possible combinations of the factor levels. The experiment is called an $s_1 \times \cdots \times s_n$ factorial experiment, and is called an s^n experiment when $s_1 = \cdots = s_n = s$. For example, a fertilizer may be a combination of the levels of three factors N (nitrogen), P (phosphate), and K (potash), and a chemical process might involve temperature, pressure, concentration of a catalyst, etc. Fisher (1926) introduced factorial design to agricultural experiments, and Yates (1935, 1937) made significant contributions to its early development.

When the treatments have a factorial structure, typically we are interested in the effects of individual factors, as well as how the factors interact with one another. Special functions of the treatment effects, called *main effects* and *interactions*, can be defined to represent such effects of interest. We say that the treatment factors do not interact if, when the levels of a factor are changed while those of the other factors are kept constant, the changes in the treatment effects only depend on the levels of the varying factor. In this case, we can separate the effects of individual factors, and the effect of each treatment combination can be obtained by summing up these individual effects. Under such *additivity* of the treatment factors, for example, to determine the combination of N, P, and K with the highest average yield, one can simply find the best level of each of the three factors separately. Otherwise, the factors need to be considered simultaneously. Roughly speaking, the main effect of a treatment factor is its effects averaged over the levels of the other factors, and the interaction effects measure departures from additivity. Precise definitions of these effects, collectively called *factorial effects*, will be given in Chapter 6.

A factorial experiment with each treatment combination observed once is called a *complete* factorial experiment. We also refer to it as a single-replicate complete factorial experiment. The analysis of completely randomized experiments in which each treatment combination is observed the same number of times, to be presented in Chapter 6, is straightforward. It becomes more involved if the experimental units have a more complicated block structure and/or if not all the treatment combinations can be observed.

When a factorial experiment is blocked, with each block consisting of one replicate of all the treatment combinations, the analysis is still very simple. As will be discussed in Chapter 6, in this case all the treatment main effects and interactions can be estimated in the same way as if there were no blocking, except that the variances of these estimators depend on the within-block variability instead of the overall variability of the experimental units. Since the total number of treatment combinations increases rapidly as the number of factors becomes large, a design that accommodates all the treatment combinations in each block requires large blocks whose homogeneity is difficult to control. In order to achieve smaller within-block variability, we cannot accommodate all the treatment combinations in the same block and must use *incomplete blocks*. It may also be impractical to carry out experiments in large blocks. Then, since not all the treatment combinations appear in the same block, the estimates of some treatment factorial effects cannot be based on within-block comparisons alone. This may result in less precision for such estimates. For example, suppose an experiment on two two-level factors A_1 and A_2 is to be run on two different days with the two combinations (0, 0) and (1, 1) of the levels of A_1 and

A_2 observed on one day, and the other two combinations (0, 1) and (1, 0) observed on the other day, where 0 and 1 are the two factor levels. Then estimates of the main effect (comparison of the two levels) of factor A_1 and the main effect of A_2 are based on within-block comparisons, but as will be seen in Chapter 7, the interaction of the two factors would have to be estimated by comparing the observations on the first day with those on the second day, resulting in less precision. We say that this two-factor interaction is *confounded* with blocks.

When a factorial experiment must be run in incomplete blocks, we choose a design in such a way that only those factorial effects that are less important or are known to be negligible are confounded with blocks. Typically the main effects are deemed more important, and one would avoid confounding them with blocks. However, due to practical constraints, sometimes one must confound certain main effects with blocks. For instance, it may be difficult to change the levels of some factors. In the aforementioned example, if a factor must be kept at the same level on each day, then the main effect of that factor can only be estimated by a more variable between-day comparison.

Often the number of treatment combinations is so large that it is practically possible to observe only a small subset of the treatment combinations. This is called a *fractional factorial design*. Then, since not all the treatment combinations are observed, some factorial effects are mixed up and cannot be distinguished. We say that they are *aliased*. For example, when only 16 treatment combinations are to be observed in an experiment involving six two-level factors, there are 63 factorial effects (6 main effects, 15 two-factor interactions, 20 three-factor interactions, 15 four-factor interactions, 6 five-factor interactions, and 1 six-factor interaction), but only 15 degrees of freedom are available for estimating them. This is possible if many of the factorial effects are negligible. One design issue is which 16 of the 64 treatment combinations are to be selected.

An important property of a fractional factorial design, called *resolution*, pertains to the extent to which the lower-order effects are mixed up with higher-order effects. For example, under a design of resolution III, no main effect is aliased with other main effects, but some main effects are aliased with two-factor interactions; under a design of resolution IV, no main effect is aliased with other main effects or two-factor interactions, but some two-factor interactions are aliased with other two-factor interactions; under a design of resolution V, no main effects and two-factor interactions are aliased with one another. When the experimenter has little knowledge about the relative importance of the factorial effects, it is common to assume that the lower-order effects are more important than higher-order effects (the main effects are more important than interactions, and two-factor interactions are more important than three-factor interactions, etc.), and effects of the same order are equally important. Under such a *hierarchical assumption*, it is desirable to have a design with high resolution. A popular criterion of selecting fractional factorial designs and a refinement of maximum resolution, called *minimum aberration*, is based on the idea of minimizing the aliasing among the more important lower-order effects.

When the experimental units have a certain block structure, in addition to picking a fraction of the treatment combinations, we also have to decide how to assign

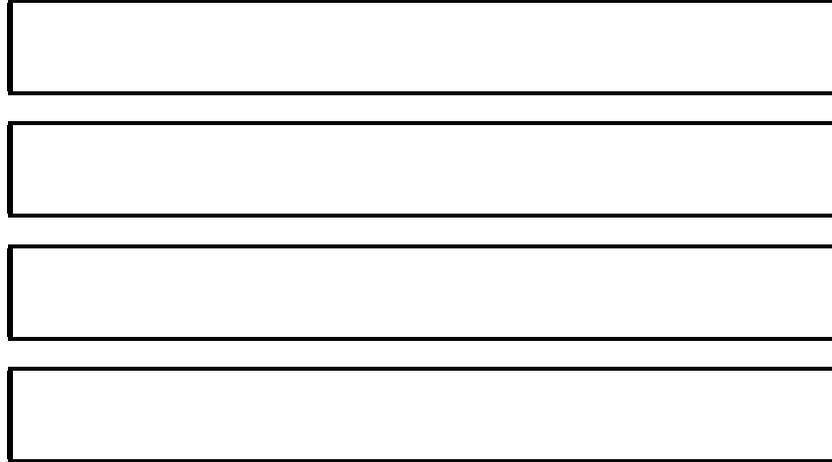
the selected treatment combinations to the units. In highly fractionated factorial experiments with complicated block structures, we have complex aliasing of treatment factorial effects as well as multiple levels of precision for their estimates. The bulk of this book is about the study of such designs, including their analysis, selection, and construction. The term “multi-stratum” in the subtitle of the book refers to multiple sources of errors that arise from complicated block structures, while “single-stratum” is synonymous with “complete randomization” where there is one single error term.

Treatment and block structures are two important components of a randomized experiment. Nelder (1965a,b) emphasized their distinction and developed a theory for the analysis of randomized experiments with simple block structures. Simple block structures cover most, albeit not all the block structures that are commonly encountered in practice. Speed and Bailey (1982) and Tjur (1984) further developed the theory to cover the more general *orthogonal block structures*. This theory, an account of which can be found in Bailey (2008), provides the basis for the approach adopted in this book.

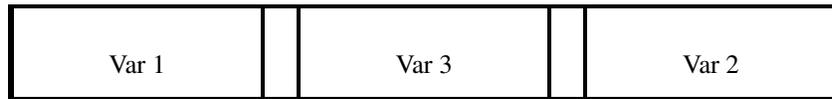
We turn to five examples of factorial experiments to motivate some of the topics to be discussed in the book. The first three examples involve simple block structures. The block structures in Examples 1.4 and 1.5 are not simple block structures, but the theory developed by Speed and Bailey (1982) and Tjur (1984) is applicable. We will return to these examples from time to time in later chapters to illustrate applications of the theory as it is developed.

Our first example is a replicated complete factorial experiment with a relatively complicated block structure.

Example 1.1. Loughin (2005) studied the design of an experiment on weed control. Herbicides can kill the weeds that reduce soybean yields, but they can also kill soybeans. On the other hand, soybean varieties can be bred or engineered to be resistant to certain herbicides. An experiment is to be carried out to study what factors influence weed control and yield of genetically altered soybean varieties. Four factors studied in the experiment are *soybean variety/herbicide combinations* in which the herbicide is safe for the soybean variety, *dates* and *rates* of herbicide application, and *weed species*. There are three variety/herbicide combinations, two dates (early and late), three rates (1/4, 1/2, and 1), and seven weed species, giving a total of 126 treatments with a $3 \times 2 \times 3 \times 7$ factorial structure. Soybeans and weeds are planted together and a herbicide safe for the soybean variety is sprayed at the designated time and rate. Then weed properties (numbers, density, mass) and soybean yields are measured. However, there are some practical constraints on how the experiment can be run. Due to herbicide drift, different varieties cannot be planted too close together and buffer zones between varieties are needed, but the field size is not large enough to allow for 126 plots of adequate size with large buffers between each pair of adjacent plots. Therefore, for efficient use of space, one needs to plant all of a given soybean variety contiguously so that fewer buffers are needed. Additional drift concerns lead to a design described as follows. First the field is divided into four blocks to accommodate four replications:



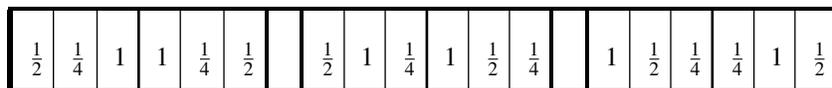
Each block is split into three plots with two buffer zones, and the variety/herbicide combinations are randomly assigned to the plots within blocks:



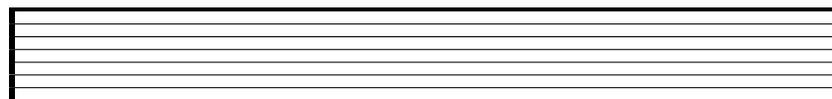
Each plot is then split into two subplots, with application times randomly assigned to subplots within plots:



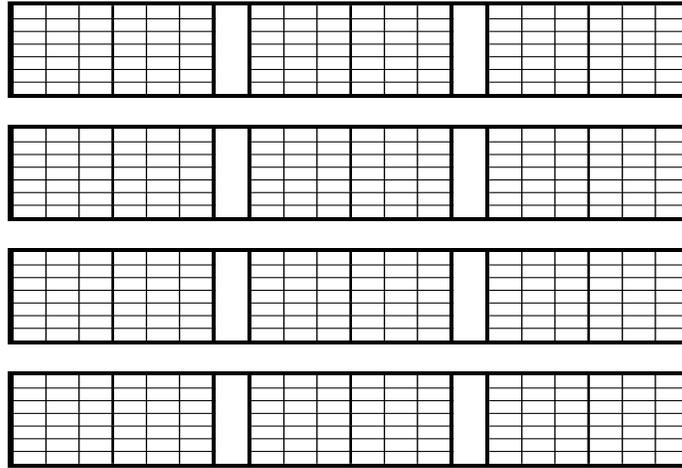
Furthermore each subplot is split into three sub-subplots, with application rates randomly assigned to sub-subplots within subplots:



Each block is divided into seven horizontal strips, with the weed species randomly assigned to the strips within blocks:



We end up with 504 combinations of sub-subplots and strips:



Each of the 126 treatment combinations appears once in each of the four blocks. To summarize, we have four replicates of a complete $3 \times 2 \times 3 \times 7$ factorial experiment with the block structure $4/[(3/2/3) \times 7]$. Both subplots in the same plot are assigned the same variety/herbicide combination, all the sub-subplots in the same subplot are assigned the same herbicide application time, and all the sub-subplot and strip intersections in the same strip are assigned the same weed species. Various aspects of the analysis of this design will be discussed in Sections 12.1, 12.9, 12.10, and 13.10.

In Example 1.1, there are 18 sub-subplots in each block. If different soybean varieties were to be assigned to neighboring sub-subplots, then 17 buffer zones would be needed in each block. With only two buffer zones per block under the proposed design, comparisons of soybean varieties are based on between-plot comparisons, which are expected to be more variable than those between subplots and sub-subplots. The precision of the estimates of such effects is sacrificed in order to satisfy the practical constraints.

Example 1.2. McLeod and Brewster (2004) discussed an experiment for identifying key factors that would affect the quality of a chrome-plating process. Suppose six two-level treatment factors are to be considered in the experiment: A , chrome concentration; B , chrome to sulfate ratio; C , bath temperature; S , etching current density; T , plating current density; and U , part geometry. The response variables include, e.g., the numbers of pits and cracks. The chrome plating is done in a bath (tank), which contains several rectifiers, but only two will be used. On any given day the levels of A , B , and C cannot be changed since they represent characteristics of the bath. On the other hand, the levels of factors S , T , and U can be changed at the rectifier level. The experiment is to be run on 16 days, with four days in each of four weeks. Therefore there are a total of 32 runs with the block structure $(4 \text{ weeks})/(4 \text{ days})/(2 \text{ runs})$, and

one has to choose 32 out of the $2^6 = 64$ treatment combinations. Weeks, days, and runs can be considered as blocks, whole-plots, and subplots, respectively. The three factors A , B , and C must have constant levels on the two experimental runs on the same day, and are called whole-plot treatment factors. The other three factors S , T , and U are not subject to this constraint and are called subplot treatment factors. We will return to this example in Sections 12.9, 13.2, 13.4, 13.5, 13.7, 14.4, and 14.13.

Example 1.3. Miller (1997) described a laundry experiment for investigating methods of reducing the wrinkling of clothes. Suppose the experiment is to be run in two blocks, with four washers and four dryers to be used. After four cloth samples have been washed in each washer, the 16 samples are divided into four groups with each group containing one sample from each washer. Each of these groups is then assigned to one dryer. The extent of wrinkling on each sample is evaluated at the end of the experiment. This results in 32 experimental runs that can be thought to have the $2/(4 \times 4)$ block structure shown in Figure 1.1, where each cell represents a cloth sample, rows represent sets of samples that are washed together, and columns represent sets of samples that are dried together. There are ten two-level treatment factors,

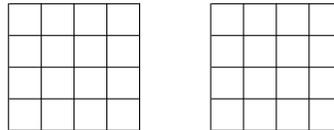


Figure 1.1 A $2/(4 \times 4)$ block structure

six of which (A, B, C, D, E, F) are configurations of washers and four (S, T, U, V) are configurations of dryers. One has to choose 32 out of the $2^{10} = 1024$ treatment combinations. Furthermore, since the experimental runs on the cloth samples in the same row are conducted in the same washing cycle, each of A, B, C, D, E, F must have a constant level in each row. Likewise, each of S, T, U, V must have a constant level in each column. Thus in each block, four combinations of the levels of A, B, C, D, E, F are chosen, one for each row, and four combinations of the levels of S, T, U, V are chosen, one for each column. The four combinations of washer settings are then coupled with the four combinations of dryer settings to form 16 treatment combinations of the ten treatment factors in the same block. An experiment run in this way requires only four washer loads and four dryer loads in each block. If one were to do complete randomization in each block, then four washer loads and four dryer loads could produce only four observations. The trade-off is that the main effect of each treatment factor is confounded with either rows or columns. Construction and analysis of designs for such *blocked strip-plot experiments* will be discussed in Sections 12.2, 12.9, 12.10, 13.3, 13.4, 13.5, 13.6, 13.7, 14.5, 14.6, and 14.14.

Federer and King (2006) gave a comprehensive treatment of split-plot and strip-

plot designs and their many variations. In this book, we present a unifying theory that can be systematically applied to a very general class of multi-stratum experiments.

Example 1.3 is an experiment with two processing stages: washing and drying. Many industrial experiments involve a sequence of processing stages, with the levels of various treatment factors assigned and processed at different stages. At each stage the experimental units are partitioned into disjoint classes. Those in the same class, which will be processed together, are assigned the same level of each of the treatment factors that are to be processed at that stage. We call the treatment factors processed at the i th stage the i th-stage treatment factors. In Example 1.3, levels of the six washer factors are set at the first stage and those of the four dryer factors are set at the second stage. So the washer configurations are first-stage treatment factors and the dryer configurations are second-stage treatment factors. Such an experiment with two processing stages can be thought to have experimental units with a row-column structure.

In Examples 1.1–1.3, the experimental units can be represented by *all* the level combinations of some *unit factors*. In the next two examples, we present experiments in which the experimental units are a *fraction* of unit-factor level combinations.

Example 1.4. Mee and Bates (1998) discussed designs of experiments with multiple processing stages in the fabrication of integrated circuits. Suppose that at the first stage 16 batches of material are divided into four groups of equal size, with the same level of each first-stage treatment factor assigned to all the batches in the same group. At the second stage they are rearranged into another four groups of equal size, again with the same level of each second-stage treatment factor assigned to all the batches in the same group. As in Example 1.3, the groupings at the two stages can be represented by rows and columns. Then each of the first-stage groups and each of the second-stage groups have exactly one batch in common. This is a desirable property whose advantage will be explained in Section 12.13. Now suppose there is a third stage. Then we need a third grouping of the batches. One possibility is to group according to the numbers in the Latin square shown earlier:

1	2	3	4
2	1	4	3
3	4	1	2
4	3	2	1

One can assign the same level of each third-stage treatment factor to all the units (batches) corresponding to the same number in the Latin square. One advantage is that each of the third-stage groups has exactly one unit in common with any group at the first two stages. If a fourth stage is needed, then one may group according to the following Latin square:

1	2	3	4
4	3	2	1
2	1	4	3
3	4	1	2

This and the previous Latin square have the property that when one is superimposed on the other, each of the 16 pairs of numbers (i, j) , $1 \leq i, j \leq 4$, appears in exactly one cell. We say that these two Latin squares are *orthogonal* to each other. If the fourth-stage grouping is done according to the numbers in the second Latin square, then each of the fourth-stage groups also has exactly one unit in common with each group at any of the first three stages. This kind of block structure cannot be obtained by iterations of nesting and crossing operators. To be a simple block structure, with four groups at each of three or four stages, one would need $4^3 = 64$ or $4^4 = 256$ units, respectively. Thus the 16 units can be regarded as a quarter or one-sixteenth fraction of the combinations of three or four 4-level factors, respectively. The following is a complete 2^4 design which can be used for experiments in which the levels of the four treatment factors are set at four stages, one factor per stage: the first factor has a constant level in each row, the second factor has a constant level in each column, the third factor has a constant level in each cell occupied by the same number in the first Latin square, and the fourth factor has a constant level in each cell occupied by the same number in the second Latin square.

0000	0011	0101	0110
0010	0001	0111	0100
1001	1010	1100	1111
1011	1000	1110	1101

We will return to this example in Sections 12.5, 12.10, 12.13, and 13.11.

Example 1.5. Bingham, Sitter, Kelly, Moore, and Olivas (2008) discussed experiments with multiple processing stages where more groups are needed at each stage, which makes it impossible for all the groups at different stages to share common units. For example, in an experiment with two processing stages, suppose 32 experimental units are to be partitioned into 8 groups of size 4 at each of the two stages. One possibility is to partition the 32 units as in Figure 1.1. The eight rows of size 4, four of which from each of the two blocks, together constitute the eight first-stage groups, and the eight columns in the two blocks together constitute the eight second-stage groups. As shown in the following figure, the 32 starred experimental units are a *fraction* of the 64 units in a completely crossed 8×8 square.

*	*	*	*				
*	*	*	*				
*	*	*	*				
*	*	*	*				
				*	*	*	*
				*	*	*	*
				*	*	*	*
				*	*	*	*

As in Example 1.4, the 32 units do not have a simple block structure. An important difference, however, is that in the current setting not all the first-stage groups can meet with every second-stage group, causing some complications in the design and analysis (to be discussed in Section 13.12). The figure shows that the 32 units are divided into two groups of size 16, which we call *pseudo* blocks since they are not part of the originally intended block structure. We will revisit this example in Sections 12.13 and 13.12, and show that it can be treated as an experiment with the block structure $2/(4 \times 4)$. A similar problem was studied in Vivacqua and Bisgaard (2009), to be discussed in Section 14.7.

An overview

Some introductory material is presented in Chapters 2–5. Chapter 2 is a review of some results on linear models, with emphasis on one-way and two-way layout models and a geometric characterization of the condition of proportional frequencies between two factors. Under the assumption of treatment-unit additivity, randomization models are developed in Chapter 3 for some designs with simple block structures, including block designs and row-column designs. In Chapter 4, the condition of proportional frequencies is extended to a notion of orthogonal factors that plays an important role in the block structures studied in this book. Some mathematical results on factors that are needed throughout the book are also gathered there. A condition that entails simple analysis of a randomized design (Theorem 5.1) is established in Chapter 5. This result is used to present a unified treatment of the analyses of three classes of orthogonal designs (completely randomized designs, complete block designs, and Latin square designs) under the randomization models derived in Chapter 3. It is also a key result for developing, in later chapters, a general theory of orthogonal designs for experiments with more complicated block structures.

The treatment factorial structure is introduced in Chapter 6. It is shown how certain special functions of the treatment effects can be defined to represent main effects and interactions of the treatment factors. Unless all the treatment factors have two levels, the choices of such functions are not unique. Several methods of constructing them based on orthogonal polynomials, finite Euclidean geometry, and Abelian groups are presented. The discussion of complete factorial designs is continued in Chapter 7, for experiments that are conducted in incomplete blocks or row-column

layouts, including split-plot and strip-plot designs. In this case, there is more than one error term and some factorial effects are confounded with blocks, rows, or columns. A construction method based on design keys is presented in addition to a commonly used method, and is shown to enjoy several advantages.

Fractional factorial designs under complete randomization are treated in Chapters 8–11. In Chapter 8, the important combinatorial structure of orthogonal arrays is introduced. Some basic properties of orthogonal arrays as fractional factorial designs, including upper bounds on the number of factors that can be accommodated by orthogonal arrays of a given run size, are derived. We also present several methods of constructing orthogonal arrays, in particular the foldover method and the construction via difference matrices. The chapter is concluded with a brief discussion of applications of orthogonal arrays to computer experiments and three variants of orthogonal arrays recently introduced for this purpose. The emphasis of this book is mainly on the so-called regular fractional factorial designs, which are easy to construct and analyze, and have nice structures and a rich theory. In Chapter 9 we provide a treatment of their basics, including design construction, aliasing and estimability of factorial effects, resolution, a search algorithm for finding designs under which some required effects can be estimated, and the connection with the linear codes of coding theory. The criterion of minimum aberration and some related criteria for selecting regular fractional factorial designs are discussed in Chapter 10. The statistical meaning of minimum aberration is clarified via its implications on the aliasing pattern of factorial effects. It is shown that this criterion produces designs with good properties under model uncertainty and good lower-dimensional projections. The connection to coding theory provides two powerful tools for constructing minimum aberration designs: the MacWilliams identities can be used to establish a complementary design theory that is useful for determining minimum aberration designs when there are many factors; the Pless power moment identities lead to the criterion of minimum moment aberration, which is equivalent to minimum aberration. Besides the theoretical interest, this equivalence is useful for analytical characterization and algorithmic construction of minimum aberration designs. A Bayesian approach to the design and analysis of factorial experiments, also applicable to nonregular designs, is presented at the end of Chapter 10. Regular designs are also closely related to finite projective geometries. The connection is made in two optional sections in Chapter 9, and is used to characterize and construct minimum aberration designs in Chapter 10. The geometric connection culminates in an elegant theory of the construction and structures of resolution IV designs in Chapter 11. While foldover is a well-known method of constructing resolution IV designs, many resolution IV designs cannot be constructed by this method. We translate the geometric results into design language, and among other topics, present the methods of doubling and partial foldover for constructing them.

In Chapters 12–14, we turn to factorial designs with more complicated block structures called multi-stratum designs. Some basic results on Nelder's simple block structures and the more general orthogonal block structures are derived in Chapter 12. A general theory for the design and analysis of orthogonal multi-stratum complete factorial designs is developed in Chapter 13. This theory is applied to several

settings, including blocked split-plot designs, blocked strip-plot designs, and design of experiments with multiple processing stages. Chapter 14 is devoted to the construction of multi-stratum fractional factorial designs and criteria for their selection under model uncertainty in the spirit of minimum aberration. The five motivating examples presented above are revisited.

We survey a few nonregular design topics in Chapter 15. Under nonregular designs, the factorial effects are aliased in a complicated way, but their run sizes are more flexible than regular designs. At the initial stage of experimentation, often only a small number of the potential factors are important. Due to their run-size economy, nonregular designs are suitable for conducting factor screening experiments under the *factor sparsity* principle. In this context, it is useful to study the property of the design when it is projected onto small subsets of factors. We also discuss the relevant topics of search designs and supersaturated designs. The objective of a search design is to identify and discriminate nonnegligible effects under the assumption that the number of nonnegligible effects is small. Supersaturated designs have more unknown parameters than the degrees of freedom available for estimating them and are useful for screening active factors. In addition to these and other miscellaneous topics, we show how some of the results presented in earlier chapters can be extended to nonregular designs. For example, coding theory again proves useful for providing a way to extend minimum aberration to nonregular designs.

Throughout this book, the starred sections can be skipped, at least on the first reading. Relevant exercises are also marked with stars.