Preface

In July 1992, 45 researchers gathered at Mt. Holyoke College for a week long conference on adaptive designs. Papers were presented by both well-known and new researchers, and covered such broad areas of applications as engineering, clinical trials, pharmacologic studies, quality control, and computer science. Twenty of those papers are presented in this volume. Each paper was refereed for merit by researchers in similar areas.

The conference on adaptive designs was one of a series of summer conferences, co-sponsored by the American Mathematical Society, the Institute of Mathematical Statistics, and the Society for Industrial and Applied Mathematics. A generous grant was received from the National Science Foundation, which covered travel expenses for the participants. The organizing committee consisted of Professor Stephen Durham of The University of South Carolina (co- chair), Professor Nancy Flournoy of The American University (co-chair), Professor Gordon Simons of The University of North Carolina, and Professor Michael Woodroofe of The University of Michigan.

An adaptive design is an experimental design that chooses future points of experimentation based, wholly or in part, on responses observed at previous points of experimentation. While similar in spirit to the Bayesian philosophy of incorporating prior information in an experiment, adaptive designs can be either Bayesian or frequentist, as evidenced by the papers in this volume. While sequential in nature, adaptive designs can fit into the traditional sequential analysis framework, but many do not, as also evidenced by the papers in this volume. Hence adaptive designs are best categorized as an independent subdiscipline of the experimental design discipline.

Because adaptive designs rely on previous responses, stochastic processes are generated which require the full arsenal of results on dependent observations. In particular, the rich theory of martingales is used in many of the papers in this volume. While the theory of i.i.d. sampling has been the primary focus of statistics until recently, we now have the tools for estimation and inference to fully develop designs which impose dependencies. Potentially more efficient, more ethical, and perhaps even more sensible designs are being developed under the adaptive framework. These papers show progress toward this goal.

The first section of this volume deals with adaptive randomization designs in clinical trials. As emerging information on patient response becomes available during the course of a clinical trial, ethical reasons may demand that patients not be assigned in equal numbers to the treatment arms. For example, if one treatment arm is performing particularly well, it might be desirable to skew the treatment allocation probabilities to favor that treatment, so that more patients in the trial will be assigned to it. Adaptive designs in this framework are only attractive if estimation and inference can be performed as efficiently and convincingly as with an equal allocation design. This is so, in many cases, and the papers in this section demonstrate that well.

Rosenberger discusses recent controversies regarding the use of adaptive designs in clinical trials and develops nonparametric inference techniques for the well-known randomized play-the-winner (RPW) rule, for binary responses, and a new design for continuous responses. Smythe and Rosenberger discuss the relationship among the RPW rule, an urn model, and the theory of Markov branching processes. They use these relationships to develop a central limit theorem for the proportion of patients assigned to each treatment. The RPW is just one member of a class of adaptive biased coin rules. In his paper, Bather discusses the potential for selection bias when these rules are implemented. Melfi develops renewal-theoretic tools to analyze such designs. Eisele reviews to adaptive biased coin designs and gives applications to estimation and testing problems.

Bandit problems originated in gambling, but can be applied to randomization in clinical trials. Hardwick compares a bandit rule with the RPW rule and another rule in her paper. Jones, Lewis, and Hartley describe a multicriteria bandit in the Bayesian context.

Most clinical trials methodology has been developed for the comparison of two treatments. However, many clinical trials examine multiple treatment arms. Coad and Palmer independently describe approaches to the problems associated with allocating multiple treatments. Coad develops a fully sequential procedure where the error probabilities are independent of the adaptive allocation rule used and Palmer examines three elimination procedures in the context of multiple comparisons.

Finally, multistage designs have been developed for designing a clinical trial. Cheng and Berry examine the optimality of a multistage design with dichotomous response. In their design, information is updated after each stage using Bayes theorem.

The second section deals with adaptive dose-response designs and quantile estimation. The problem is similar to that of allocation of treatments in clinical trials, except that dose levels are allocated rather than treatments. Adaptive designs can be employed to avoid assigning patients to toxic dose levels. Another application is in engineering, where stress levels may be applied to some device in a quality control experiment. The inherent ordinality in the dose levels can be exploited to develop reasonable allocation rules. This is done via an up-and-down design in the paper by Durham and Flournoy, which characterizes the asymptotic distribution of proportions of patients assigned to each dose level for their design, and the paper by Durham, Flournoy, and Montazer-Haghighi, which derives the exact moments of the proportions of patients assigned to each treatment.

While the principal goal of pharmaceutical dose-response studies is to describe the relationship between dose level and toxicity, Li, Durham, and Flournoy, develop an adaptive procedure which also incorporates the efficacy (in terms of curing the disease) of the dose level. They examine the optimality of such a contingent binary response design.

Sequential estimation for dose-response studies is examined in the paper by Govindarajulu. He compares three methods, including stochastic approximation, for sequentially estimating a quantile corresponding to a target probability of response.

The third section concerns applications of adaptive designs in engineering, quality control, and linear modeling. Page examines the estimation of functions of parameters and how best to allocate in order to achieve optimality. Hardwick and Stout discuss computational strategies for exact analyses of allocation rules in both an industrial and a clinical trials scenario. They use dynamic programming to develop efficient algorithms.

Quality control in industry is one area of potential application for adaptive designs. Rather than ethical cost being the primary focus, as in human experimentation, efficiency and cost are the predominant factors in the industrial setting. Aras develops an optimal age replacement policy for a system when the underlying model for machine failure is unknown. Hsu develops new procedures for a group-testing problem, where the goal is to identify defective units in a quality control setting. Soyer and Vopatek present a Bayesian framework for accelerated life testing. Optimal designs are derived for exponential life models.

In the last paper in this section, Schwabe demonstrates that adaptive designs improve the performance of classical linear models.

We close by expressing our appreciation to Michelle Beattie, for organizing and typing many of the manuscripts, to Dr. Leonard Hearne for tirelessly reading the Latex manual and solving many computer systems problems, and to the referees who volunteered their time; many of the papers in this volume were substantially improved thanks to their thoughtful comments.

N. F.
W. F. R.

Washington, DC
December 1994