SCHOOL OF OPERATIONS RESEARCH
AND INDUSTRIAL ENGINEERING
COLLEGE OF ENGINEERING
CORNELL UNIVERSITY
ITHACA, NEW YORK 14853

TECHNICAL REPORT NO. 614

December 1983

A SURVEY OF LITERATURE ON ESTIMATION METHODS FOR QUANTAL RESPONSE CURVES WITH A VIEW TOWARD APPLYING THEM TO THE PROBLEM OF SELECTING THE CURVE WITH THE SMALLEST q-QUANTILE (ED100q) (Preliminary Report)

by

Ajit C. Tamhane*

Research supported by

U.S. Army Research Office - Durham Contract DAAG-29-81-K-0168

at Cornell University

Approved for Public Release; Distribution Unlimited

*Department of Industrial Engineering and Management Science, Technological Institute, Northwestern University, Evanston, Illinois.

TABLE OF CONTENTS

Abstract

A survey of the literature is presented on various statistical inference problems and procedures associated with the following experimental setting: Any level of dose of some treatment can be administered to each experimental unit, and the observed response on each unit is binary ("success" or "failure") with the unknown success probability being a nondecreasing function of the dose-level. This function is referred to as a quantal response curve. Both parametric and nonparametric models for a quantal response curve are considered. In each case both single-stage and sequential estimation procedures are reviewed in detail.

This survey was motivated by the following statistical selection problem: Suppose that there are available several treatments, and that each one has associated with it a completely unknown quantal response curve. For each curve define its qth quantile (ED100q) as the dose-level which corresponds to a "success" probability of q where $q \in (0,1)$ is specified. For any prespecified q we wish to select the treatment associated with the smallest ED100q. Some difficulties inherent in solving this selection problem, and possible approaches to devising statistical procedures for achieving the stated goal are discussed. We have not found satisfactory solutions to these difficulties, and the problem is essentially open.

Keywords: Dose-Response; Quantal Bioassay; Sensitivity Experiments; Logistic Model; Probit Model; Maximum Likelihood Estimation; Weighted Least-Squares Estimation; Minimum (Modified) Chi-Square Estimation; Isotonic Estimation; Nonparametric Estimation; Robust Estimation; Design of Experiments; Stochastic Approximation; Up-and-Down Method; Sequential Designs; Multiple Comparisons; Ranking and Selection; Indifference-Zone Approach.

1. Introduction

In many types of experimental research the treatment factor is quantitative in nature and an experimental unit to which a certain "dose" of the treatment is administered either responds (a "success") or does not respond (a "failure") to the given dose, i.e., the response is quantal or binary in nature. In biological applications such experiments are known as quantal response assays (Finney 1978, Chs. 17-20) or sensitivity experiments.

In a sensitivity experiment the probability of response is some unknown function of the dose-level x. We shall denote this probability by $p(x)$ and refer to it as the quantal response curve or the dose response curve. In many applications it is reasonable to assume that $p(x)$ is non-decreasing (at least over the range of x of practical interest) with $p(-\infty) = 0$ (note that usually the logarithm of the dose-level is used as the x variable and thus negative x values are possible) and $p(+\infty) = 1$. One can conceive that each experimental unit in a population has some fixed unknown tolerance and the administration of the treatment to that unit results in a "response" if the dose-level exceeds its tolerance. In this context $p(x)$ can be thought of as the cumulative distribution function (cdf) of the tolerance values for the population and therefore it is sometimes referred to as the tolerance distribution.

In some applications it may be desired to estimate or make other inferences concerning the entire curve $p(x)$ while in others the interest may center on some selected point(s) on the curve. One such point on the curve is the dose-level μ that corresponds to 0.5 probability of response; depending on the practical context, μ is referred to as the median effective dose (ED50) or the median lethal dose (LD50) or the median tolerance. More

generally, the experimenter may be interested in the dose-level that stimulates a response with probability q (the qth quantile of the tolerance distribution) for some specified $q \in (0,1)$; we shall refer to this dose-level as ED100q and denote it by $\mu^{(q)}$ where

$$\mu^{(q)} = \inf \{x : p(x) \geq q\}; \tag{1.1}$$

note that $\mu = \mu^{(0.5)}$. Sometimes the mean tolerance θ is used instead of the median tolerance μ where

$$\theta = \int_{-\infty}^{\infty} x \, dp(x); \tag{1.2}$$

if $p(x)$ is skew-symmetric and if θ exists then $\theta = \mu$.

We are interested in a selection problem that arises when we have several different quantitative treatments with associated unknown quantal response curves and it is desired to select the treatment which has the "best" (in some appropriate sense) curve. For example, several drugs may be available for curing a certain physiological disorder, the outcome on a patient administered a certain drug being simply recorded as a "success" (cured) or a "failure" (not cured). Although the therapeutic effectiveness of each drug (as measured by its "success" probability) may be an increasing function of its dose-level (at least over some restricted range) high doses may be undesirable because of accompanying toxic side-effects. Therefore a medical researcher may wish to select that drug which induces a success rate of 50% (say) at the lowest dose-level, i.e., the drug that is associated with the smallest ED50.

More formally, let $\Pi_1, \Pi_2, \ldots, \Pi_k$ denote $k \geq 2$ different populations (treatments, drugs, etc.). Let $Y_i(x)$ denote an observation on Π_i at dose-level x, which takes on the value 1 (if "success" or "response") or 0 (if "failure" or "no response") with probabilities

$$\left.\begin{aligned} & P\{Y_i(x) = 1\} = p_i(x) \\ \text{and} \qquad & \\ & P\{Y_i(x) = 0\} = 1 - p_i(x) \end{aligned}\right\} \qquad (1 \le i \le k).$$

Each $p_i(x)$ $(1 \le i \le k)$ is assumed to be nondecreasing and continuous in x with $p_i(-\infty) = 0$, $p_i(+\infty) = 1$, but is otherwise assumed to be completely unknown. For specified $q \in (0,1)$ let $\mu_i^{(q)}$ denote the ED100q (cf. (1.1)) associated with $\Pi_i (1 \le i \le k)$ and let $\mu_{[1]}^{(q)} \le \mu_{[2]}^{(q)} \le \ldots \le \mu_{[k]}^{(q)}$ denote the ordered values of the $\mu_i^{(q)}$. The correct pairing between the $\mu_{[i]}^{(q)}$ and $\Pi_j (1 \le i,j \le k)$ is assumed to be completely unknown. The goal of the experimenter is to select the population (assumed to be unique) associated with, say $\mu_{[1]}^{(q)}$, which is here referred to as the "best" population. If the decision procedure selects the "best" population then we say that a correct selection (CS) is made.

Adopting the indifference-zone approach of Bechhofer (1954) to this selection problem, we can state a requirement on the probability of a correct selection (P(CS)) as follows: Let $\{\delta^*, P^*\}$ be preassigned constants where $\delta^* > 0$ and $k^{-1} < P^* < 1$. Consideration is restricted to those procedures which guarantee the probability requirement

$$P(CS) \ge P^* \qquad \text{whenever } \mu_{[2]}^{(q)} - \mu_{[1]}^{(q)} \ge \delta^*. \qquad (1.3)$$

A bit of reflection shows that there are many deep and subtle difficulties associated with this selection problem some of which we now discuss.

(i) First (and the most obvious) point to note is that if the quantal response curves cross each other, i.e., if there is treatment-dose interaction, then the population that is the "best" for one specified value of q

need not be the "best" for some other specified value of q. However, if it can be assumed that the $p_i(x)$ do not intersect (as would happen, e.g., if the $p_i(x)$ are members of a location parameter family and thus can be obtained by parallel shifts of the same basic curve) then we have a uniformly (for all $q \in (0,1)$) unique best population, say, $\Pi_{[1]}$. In this case we have for all $x \in \mathbb{R}$ that $p_{[1]}(x) \geq p_i(x)$ for all $i \neq [1]$.

(ii) A second point to note is that the parameters of interest $\mu_i^{(q)}$ must be estimated "backwards" by first estimating $p_i(x)$ at selected values of x and then "interpolating" at a given value of q by a suitable method. If the choice of x-values at which observations are made happens to be unfortunate (e.g., if all of the x-values lie on one side of the unknown $\mu_i^{(q)}$) then it would be impossible to obtain a reliable estimate of $\mu_i^{(q)}$ unless some additional assumption regarding the functional form of $p_i(x)$ is made, e.g., one might assume the logistic or the probit model; see §2. A more serious difficulty which arises because of this estimation feature of the problem is the following.

(iii) If the quantal response curve $p_i(x)$ is almost "flat" in the region of $\mu_i^{(q)}$ then the estimate of $\mu_i^{(q)}$ would be highly unreliable. Thus, even if the $\mu_i^{(q)}$-values for two populations differ greatly, say, $\mu_{[2]}^{(q)} - \mu_{[1]}^{(q)} >> \delta^*$, they would be virtually indistinguishable on the probability of response scale if the corresponding $p_i(x)$ curves are almost flat in the interval $(\mu_{[1]}^{(q)}, \mu_{[2]}^{(q)})$ as shown in Figure 1.

Figure 1

Almost Flat Quantal Response Curves

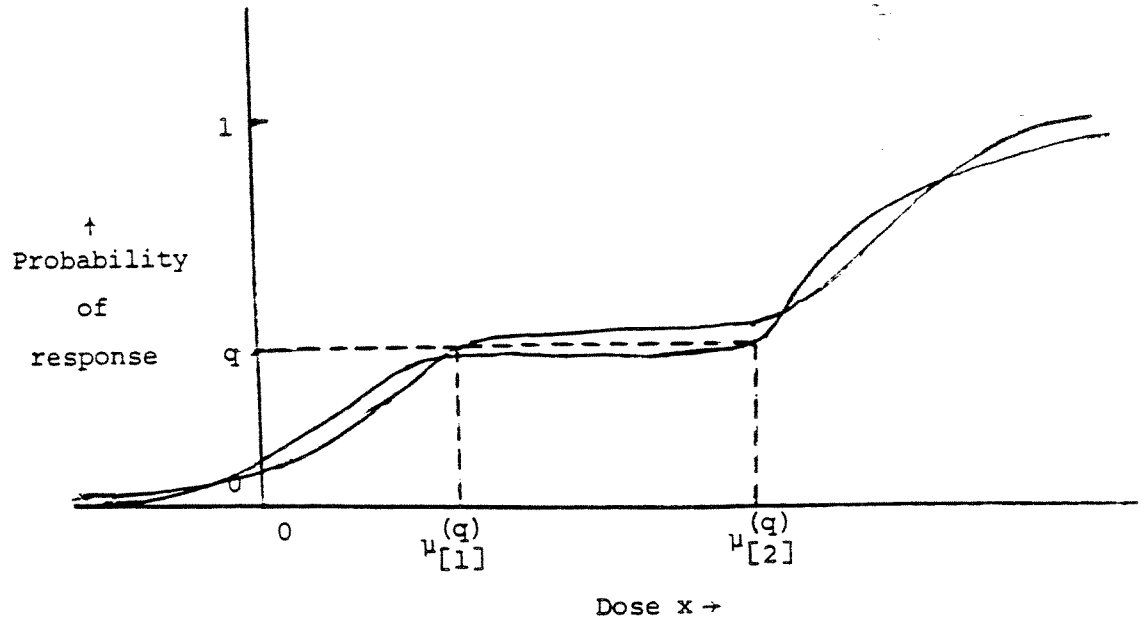

To avoid this difficulty, an indifference-zone formulation analogous to the one proposed by Sobel (1967) may be required. In this alternative formulation a threshold separation is postulated not between $\mu_{[k]}^{(q)}$ and $\mu_{[k-1]}^{(q)}$, but between the corresponding quantal response curves along the vertical scale in the neighborhood of q.

If some specific functional form is assumed for the quantal response curves then the results obtained from sampling outside the "trouble region" of dose-level values can presumably be used to obtain information relevant to selecting the population associated with $\mu_{[1]}^{(q)}$. By using an appropriate adaptive sequential sampling procedure it may be possible to avoid the "trouble region" during the major (later) portion of sampling from each population. However, if no specific functional form can be assumed for the curves then one must minimally assume some positive lower bound on the slopes of the curves in the region of interest, namely $(\mu_{[1]}^{(q)}, \mu_{[k]}^{(q)})$.

(iv) It is interesting to note that if the quantal response curves do not cross each other then it is not necessary to estimate the $\mu_i^{(q)}$ for identifying the "best" population, and therefore the difficulties discussed under (ii) and (iii) do not arise. The selection problem in this case can be "solved," at least in principle, by choosing an arbitrary dose-level x_0, and then using an appropriate Bernoulli selection procedure for identifying the population associated with $p_{[1]}(x_0) = \max_{1 \le i \le k} p_i(x_0)$. This situation is depicted in Figure 2.

Figure 2

Nonintersecting Quantal Response Curves

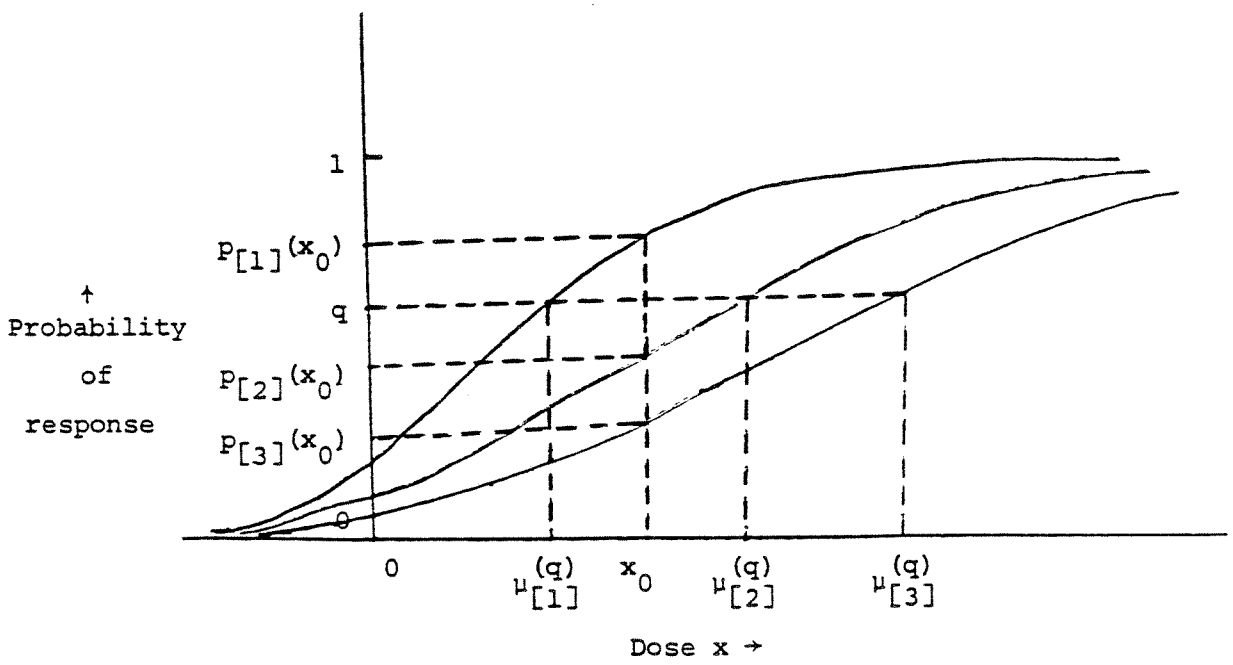

Since the existence of a uniformly "best" population is guaranteed (and, in fact, the same ordering between the $\mu_i^{(q)}$ is maintained for all values of q; see Figure 2) under the assumption of nonintersecting curves it suffices to take observations on all populations at the same dose-level x_0. The populations can then be thought of as Bernoulli with success probabilities $p_i(x_0)$, $i = 1,2,\ldots,k$. Thus the selection problem reduces to the classical

one of selecting the Bernoulli population associated with the largest success probability.

A large number of procedures have been proposed in the literature for the classical Bernoulli selection problem, any one of which could be applied in the present context; we refer the reader to the article by Bechhofer and Kulkarni (1982) for a review of this literature and a complete bibliography. In particular, a new closed adaptive sequential procedure proposed in this article appears particularly attractive because of its many desirable properties.

The main obstacle in implementing the approach just outlined is the choice of x_0. With an unfortunate choice of x_0 it is possible to have $p_{[1]}(x_0)$ differ very little from the next largest $p_i(x_0)$ making it impossible to guarantee (1.3). Ideally one would like to choose x_0 so that $p_{[1]}(x_0) - p_i(x_0) \geq \Delta^*$ for all $i \neq 1$ whenever $\mu_{[2]}^{(q)} - \mu_{[1]}^{(q)} \geq \delta^*$ where $\Delta^* > 0$ is known. Then one could use a Bernoulli selection procedure which guarantees the $\{\Delta^*, P^*\}$-requirement (see (5) of Sobel and Huyett (1957)). However, in general, it is not possible to choose x_0 in such a manner without any knowledge of the $p_i(\cdot)$'s.

(iv) Finally a key feature common to many practical problems involving inferences concerning one or more quantal response curves is that the observations are costly. As a result, experimenters must work with relatively small samples. However, with small samples it is not feasible to obtain good nonparametric estimates of the $p_i(x)$. Thus one may be forced to assume some parametric model for the quantal response curves. Even for the problem of estimation of a single parametric quantal response curve, most of the theoretical results are available only for the large sample case; thus they are often irrelevant for intended applications.

We have pointed out above that the particular selection problem of interest to us (as stated in its most general form) is fraught with many difficulties. To date we have not found a satisfactory solution to it. However, as a first step towards studying this problem in some depth, we attempted a systematic review of the pertinent literature on estimation problems for quantal response curves. That therefore is the main objective of the present article. We do discuss, however, in the last section of the paper, some tentative approaches to the selection problem which serve to highlight some of the specific difficulties involved.

The outline of the paper is as follows. We begin the review of the literature with the problem of estimation of a single quantal response curve or some specified point, e.g., the ED50, for a single such curve. In §2 we introduce the two commonly used parametric models -- the logistic and the probit. §3 considers single-stage procedures for such models. The first part of §3 deals with the case of parametric models while the second part deals with the nonparametric case where techniques of isotonic estimation and classical nonparametric estimators of the mean tolerance, e.g., the Spearman-Kärber estimator, are reviewed. §4 considers sequential procedures which are useful when the result of administration of a dose becomes known without a prolonged time delay. The main emphasis in §4 is on stochastic approximation procedures of the Robbins-Monro type; the up-and-down method is also briefly reviewed. §5 reviews some available literature on multiple comparisons among several quantal response curves. Finally in §6 we treat the selection problem in some detail. We give an explicit asymptotic solution for a single-stage selection procedure under certain restrictive assumptions on the quantal response curves.

2. Parametric Models for a Quantal Response Curve

Throughout this and the next two sections we consider the case of a single quantal response curve which we denote by $p(x)$. Various mathematical models have been proposed in the literature for $p(x)$ to portray the usual S-shape associated with quantal response curves. Many of these models are of the form

$$p(x) = F(\alpha + \beta x) \tag{2.1}$$

where $\alpha \in \mathbb{R}$, $\beta > 0$ are unknown parameters and $F(\cdot)$ is a cdf having a specified form. If ξ_q denotes the qth quantile of $F(\cdot)$ then the ED100q is given by

$$\mu(q) = \frac{\xi_q - \alpha}{\beta} \qquad (0 < q < 1). \tag{2.2}$$

Usually $F(\cdot)$ is chosen so that it possesses a density $f(\cdot)$ which is symmetric around zero. In that case (and more generally if $F(0) = 1/2$) $\xi_{0.5} = 0$ and the ED50 is given by

$$\mu = \mu(0.5) = -\frac{\alpha}{\beta}. \tag{2.3}$$

Often the model (2.1) is written in the form

$$F^{-1}\{p(x)\} = \alpha + \beta x \tag{2.4}$$

showing explicitely that the inverse F transform of the quantal response curve is a straight line. The logistic and probit models correspond to specific choices of $F(\cdot)$.

2.1 Logistic Model

For the logistic model, $F(\cdot)$ is the logistic cdf

$$F(x) = \frac{1}{1 + \exp(-x)} \tag{2.5}$$

and thus

$$F^{-1}\{p(x)\} = \ln\left\{\frac{p(x)}{1-p(x)}\right\} = \alpha + \beta x . \tag{2.6}$$

The quantity $\ln\{p(x)/(1-p(x))\}$ is known as the logit transform of $p(x)$. For the logistic model the ED100q is given by

$$\mu(q) = \frac{\ln\left(\frac{q}{1-q}\right) - \alpha}{\beta}. \tag{2.7}$$

A general reference on the logistic model is a book by Ashton (1972).

2.2 Probit Model

For the probit model, $F(\cdot)$ is the standard normal cdf

$$\Phi(x) = \int_{-\infty}^{x} \phi(u)du = \int_{-\infty}^{x} \frac{1}{\sqrt{2\pi}} e^{-\frac{u^2}{2}} du \tag{2.8}$$

and thus

$$F^{-1}\{p(x)\} = \Phi^{-1}\{p(x)\} = \alpha + \beta x \tag{2.9}$$

where $\Phi^{-1}\{p(x)\}$ is known as the probit transform of $p(x)$. (In traditional practice, the term probit model is used to refer to the transformation $\Phi^{-1}\{p(x)\} + 5 = \alpha + \beta x$. To distinguish between this model and (2.9), the latter is sometimes referred to as the normit model.)

If z_q denotes the qth quantile of the standard normal distribution then the ED100q for the model (2.9) is given by

$$\mu(q) = \mu + \frac{1}{\beta} z_q = \frac{z_q - \alpha}{\beta} . \tag{2.10}$$

A general reference on the probit model is a book by Finney (1971).

Cox (1970, §2.7) and Finney (1978, §17.14) have shown that the quantal response curves resulting from the above two models match very closely in their central portion ($0.2 \le p(x) \le 0.8$) as do several other models such as the angular (arcsin transform) and the linear. Thus the choice of a parametric model is usually dictated by other considerations such as convenience of computation and statistical analysis, ease of interpretability, etc. Based on these considerations the logistic model is preferred in a large majority of situations. We review inference procedures for both the logistic and the probit model here, the latter being considered mainly for historical reasons. Prentice (1976) has proposed a general four parameter model of which the logistic and probit models are special cases.

3. Single Quantal Response Curve: Estimation Using Single-Stage Procedures

We consider a single-stage procedure which uses $m \ge 2$ fixed dose-levels and which takes a predetermined number n_j of observations at dose-level $x_j (1 \le j \le m)$ where $x_1 < x_2 < \ldots < x_m$. Let r_j denote the number of successes obtained at dose-level x_j $(1 \le j \le m)$. Based on these data we wish to estimate the dose-response curve $p(x)$. In §3.1 and §3.2 we discuss two popular methods of estimation for the parametric model (2.1). After discussing the methods of estimation for general $F(\cdot)$ (of specified form), we specialize the results to the logistic and normal cdf's.

For convenience, we denote $p(x_j)$ by p_j, and let

$$\hat{p}_j = r_j/n_j \qquad (1 \le j \le m). \tag{3.1}$$

3.1 Parametric Estimation: Maximum Likelihood Method

The likelihood function is

$$L = \prod_{j=1}^{m} \binom{n_j}{r_j} p_j^{r_j} (1-p_j)^{n_j - r_j}. \tag{3.2}$$

Setting partial derivatives of $\ell n\ L$ with respect to α and β equal to zero, the following equations are obtained for the maximum likelihood estimators (MLE's) $\hat{\alpha}$ and $\hat{\beta}$ of α and β, respectively:

$$\frac{\partial\ \ell n\ L}{\partial\alpha} = \sum_{j=1}^{m} \left\{ \frac{r_j - n_j p_j}{p_j(1-p_j)} \right\} f(\alpha + \beta x_j) = 0$$

$$\frac{\partial\ \ell n\ L}{\partial\beta} = \sum_{j=1}^{m} \left\{ \frac{r_j - n_j p_j}{p_j(1-p_j)} \right\} x_j f(\alpha + \beta x_j) = 0. \tag{3.3}$$

Here $f(\cdot)$ is the probability density function (pdf) corresponding to $F(\cdot)$.

In general, some iterative algorithm is needed for solving (3.3), the Newton-Raphson algorithm being the one most commonly employed. After computing $\hat{\alpha}$ and $\hat{\beta}$, the MLE of $\mu^{(q)}$ is computed by using

$$\hat{\mu}^{(q)} = \frac{\xi_q - \hat{\alpha}}{\hat{\beta}} \tag{3.4}$$

where ξ_q is the qth quantile of $F(\cdot)$.

To implement the Newton-Raphson algorithm a current estimate of the Hessian matrix

$$\begin{bmatrix} \dfrac{\partial^2 \ell n\ L}{\partial\alpha^2} & \dfrac{\partial^2 \ell n\ L}{\partial\alpha\partial\beta} \\ \dfrac{\partial^2 \ell n\ L}{\partial\alpha\partial\beta} & \dfrac{\partial^2 \ell n\ L}{\partial\beta^2} \end{bmatrix} \tag{3.5}$$

is needed at each step; expressions for the second partials can be easily obtained from (3.3). The asymptotic variance-covariance matrix of $(\hat{\alpha},\hat{\beta})$ is given by -1 times the inverse of the expected value of the matrix (3.5). Using the delta method, the asymptotic variance of $\hat{\mu}^{(q)}$ is found to be

$$\mathrm{Var}(\hat{\mu}(q)) \simeq \left(\frac{\alpha_q}{\beta}\right)^2 \left\{\frac{\mathrm{Var}(\hat{\alpha})}{\alpha_q^2} + \frac{\mathrm{Var}(\hat{\beta})}{\beta^2} - \frac{2\,\mathrm{Cov}(\hat{\alpha},\hat{\beta})}{\alpha_q \beta}\right\} \tag{3.6}$$

where $\alpha_q = \alpha - \xi_q$ $(= \alpha$ for $q = 1/2)$.

Expressions for the asymptotic $\mathrm{Var}(\hat{\alpha})$, $\mathrm{Var}(\hat{\beta})$, and $\mathrm{Cov}(\hat{\alpha},\hat{\beta})$ for the logistic model are given in (3.9); for the probit model a slightly different parametrization is used and the corresponding expressions are given in (3.19). The asymptotic variance-covariance matrix of $(\hat{\alpha},\hat{\beta})$ can be consistently estimated in large samples $(n_j \to \infty\ \forall j)$ using the matrix obtained by inverting the estimate of (3.5) at the final iteration of the Newton-Raphson algorithm.

We now specialize the above method to the logistic and probit models.

3.1.1 Logistic Model

Using the fact that

$$f(x) = F(x)\{1 - F(x)\} \tag{3.7}$$

for the logistic model and hence that $f(\alpha + \beta x_j) = p_j(1 - p_j)$ $(1 \le j \le m)$, we see that the equations (3.3) reduce to

$$\begin{aligned} \sum_{j=1}^{m} n_j p_j &= \sum_{j=1}^{m} n_j \hat{p}_j \\ \sum_{j=1}^{m} n_j x_j p_j &= \sum_{j=1}^{m} n_j x_j \hat{p}_j \end{aligned} \tag{3.8}$$

where the $\hat{p}_j$ are defined in (3.1). The necessary and sufficient conditions for the existence of a unique solution to (3.8) were given by Silvapulle (1981).

The expressions for the asymptotic variances of $\hat{\alpha}$ and $\hat{\beta}$, and their

covariance are given by

$$\text{Var } (\hat{\alpha}) = \frac{1}{\sum_{j=1}^{m} w_j} + \frac{\bar{x}^2}{\sum_{j=1}^{m} w_j (x_j - \bar{x})^2},$$

$$\text{Var } (\hat{\beta}) = \frac{1}{\sum_{j=1}^{m} w_j (x_j - \bar{x})^2}, \tag{3.9}$$

and

$$\text{Cov } (\hat{\alpha}, \hat{\beta}) = \frac{-\bar{x}}{\sum_{j=1}^{m} w_j (x_j - \bar{x})^2}.$$

Here

$$w_j = n_j p_j (1 - p_j) \qquad (1 \le j \le m) \tag{3.10}$$

are the <u>unknown</u> weights, and

$$\bar{x} = \frac{\sum_{j=1}^{m} w_j x_j}{\sum_{j=1}^{m} w_j} \tag{3.11}$$

is the <u>unknown</u> weighted mean. Using (3.6) it can be shown that the asymptotic variance of $\hat{\mu}^{(q)}$ is given by

$$\text{Var } (\hat{\mu}^{(q)}) \simeq \frac{1}{\beta^2} \left\{ \frac{1}{\sum_{j=1}^{m} w_j} + \frac{(\mu^{(q)} - \bar{x})^2}{\sum_{j=1}^{m} w_j (x_j - \bar{x})^2} \right\}. \tag{3.12}$$

This variance can be consistently estimated in large samples ($n_j \to \infty \ \forall j$) by replacing $\alpha, \beta, \mu^{(q)}$ by their MLE's and estimating the w_j and $\bar{x}$ by

$$\hat{w}_j = n_j \hat{p}_j (1 - \hat{p}_j) \qquad (1 \le j \le m), \tag{3.13}$$

and

$$\hat{\bar{x}} = \frac{\sum_{j=1}^{m} \hat{w}_j x_j}{\sum_{j=1}^{m} \hat{w}_j}, \tag{3.14}$$

respectively.

3.1.2 Probit Model

It is more common to use the parametrization

$$\Phi^{-1}\{p(x)\} = \frac{x - \mu}{\sigma} \tag{3.15}$$

for the probit model. Here $\mu = -\alpha/\beta$ is the ED50 and $\sigma = 1/\beta$; see, e.g., Golub and Grubbs (1956). The following equations for the MLE's $\hat{\mu}$ and $\hat{\sigma}$ are easily obtained:

$$\frac{\partial \ell n\, L}{\partial \mu} = \frac{1}{\sigma} \sum_{j=1}^{m} \left\{ \frac{(n_j - r_j)\phi_j}{1-\Phi_j} - \frac{r_j \phi_j}{\Phi_j} \right\} = 0$$

$$\frac{\partial \ell n\, L}{\partial \sigma} = \frac{1}{\sigma} \sum_{j=1}^{m} \left\{ \frac{(n_j - r_j) t_j \phi_j}{1-\Phi_j} - \frac{r_j t_j \phi_j}{\Phi_j} \right\} = 0 \tag{3.16}$$

where

$$t_j = \frac{x_j - \mu}{\sigma}, \quad \phi_j = \phi(t_j), \Phi_j = \Phi(t_j) = p_j \quad (1 \le j \le m). \tag{3.17}$$

The MLE of the ED100q for arbitrary $q \in (0,1)$ is given by

$$\hat{\mu}^{(q)} = \hat{\mu} + z_q \hat{\sigma}. \tag{3.18}$$

The asymptotic variances of $\hat{\mu}$ and $\hat{\sigma}$, and their covariance are given by

$$\mathrm{Var}(\hat{\mu}) = \frac{\sigma^2 B}{AB-C^2}, \quad \mathrm{Var}(\hat{\sigma}) = \frac{\sigma^2 A}{AB-C^2}, \quad \mathrm{Cov}(\hat{\mu},\hat{\sigma}) = \frac{-\sigma^2 C}{AB-C^2} \tag{3.19}$$

where

$$A = \sum_{j=1}^{m} \frac{n_j \phi_j^2}{\Phi_j(1-\Phi_j)},$$

$$B = \sum_{j=1}^{m} \frac{n_j t_j^2 \phi_j^2}{\Phi_j(1-\Phi_j)}, \tag{3.20}$$

$$C = \sum_{j=1}^{m} \frac{n_j t_j \phi_j^2}{\Phi_j(1-\Phi_j)}.$$

Using (3.18) and (3.19), the asymptotic $\mathrm{Var}(\hat{\mu}^{(q)})$ can be easily derived. Consistent estimates of the variances can be obtained in large samples by replacing the unknown parameters by their MLE's.

3.2 Parametric Estimation: Weighted Least-Squares (Minimum Modified Chi-Square) Method

The basic method is as follows: Compute the empirical transforms

$$u_j = F^{-1}(\hat{p}_j) \qquad (1 \le j \le m). \tag{3.21}$$

Ignore for the moment the possibility that $\hat{p}_j$ can be 0 or 1 which can make $u_j = -\infty$ or $+\infty$, respectively. (Applying a continuity correction can help overcome this difficulty; see (3.27).) Comparing (3.21) with (2.4) we see that α and β can be estimated by fitting a linear regression of the u_j on the x_j. We should not use ordinary least-squares, however, since the

variances of the U_j (where U_j is the random variable corresponding to u_j) are not constant. In fact, using the delta method we see that

$$\mathrm{Var}(U_j) \cong \left\{ \frac{1}{f(\bar{F}^{-1}(p_j))} \right\}^2 \frac{p_j(1-p_j)}{n_j} \tag{3.22}$$

$$= \frac{1}{w_j} \qquad \text{(say)} \qquad (1 \le j \le m).$$

Thus we should use the weighted least-squares method to obtain the estimates of α and β with the w_j as the weights. But the w_j are unknown and must be replaced by their estimates (referred to as empirical weights)

$$\hat{w}_j = \{f(u_j)\}^2 \frac{n_j}{\hat{p}_j(1-\hat{p}_j)} \qquad (1 \le j \le m). \tag{3.23}$$

Thus the weighted least-squares estimates of α and β are obtained by minimizing

$$\sum_{j=1}^{m} \frac{n_j\{f(u_j)\}^2}{\hat{p}_j(1-\hat{p}_j)} \cdot (u_j-\alpha-\beta x_j)^2 . \tag{3.24}$$

Note that (3.24) corresponds to the minimization criterion proposed by Neyman (1949) to obtain what he called the minimum modified chi-square estimates of α and β. These are thus identical to the weighted least-squares estimates. These estimates are explicitely given by

$$\hat{\alpha} = \hat{\bar{u}} - \hat{\beta}\hat{\bar{x}} \text{ and } \hat{\beta} = \frac{\sum_{j=1}^{m} \hat{w}_j(u_j-\hat{\bar{u}})x_j}{\sum_{j=1}^{m} \hat{w}_j(x_j-\hat{\bar{x}})^2} \tag{3.25}$$

where

$$\hat{\bar{x}} = \frac{\sum_{j=1}^{m} \hat{w}_j x_j}{\sum_{j=1}^{m} \hat{w}_j} \quad \text{and} \quad \hat{\bar{u}} = \frac{\sum_{j=1}^{m} \hat{w}_j u_j}{\sum_{j=1}^{m} \hat{w}_j} . \tag{3.26}$$

It should be noted that the minimum modified chi-square estimates are usually recommended only for large samples.

We now specialize the above results to the logistic and probit models.

3.2.1 Logistic Model

Here corresponding to (3.21) we compute the empirical logistic transforms

$$u_j = \ell n \left\{ \frac{\hat{p}_j + 1/2n_j}{1-\hat{p}_j + 1/2n_j} \right\} \qquad (1 \le j \le m) \tag{3.27}$$

where $1/2n_j$ is the continuity correction. From (3.7) we see that $f(u_j)$ $= F(u_j)\{1-F(u_j)\} = \hat{p}_j(1-\hat{p}_j)$ which when used in (3.23) yields

$$\hat{w}_j = n_j \hat{p}_j (1-\hat{p}_j) \qquad (1 \le j \le m). \tag{3.28}$$

Once again, the $\hat{w}_j$ may be adjusted by adding $1/2n_j$ to the $\hat{p}_j$ and $1-\hat{p}_j$ terms in (3.28). The estimates $\hat{\alpha}$ and $\hat{\beta}$ can now be computed using (3.25) and (3.26).

Berkson developed the weighted least-squares method for the logistic model; he called it the minimum logit chi-square estimation method. An account of his work with reference to earlier work can be found in Berkson (1955).

Considerable controversy has existed concerning the relative merits and demerits of the maximum likelihood and the minimum chi-square methods of estimation; the reader is referred to Berkson (1980) and the accompanying discussion for relevant references and points at issue. As a practical

matter, the fact that the minimum modified chi-square estimates can be explicitely computed using (3.25) and do not require recourse to an iterative algorithm as do the MLE's might be a consideration in some situations. But in most other situations (particularly those involving small samples) the MLE's would be preferred to the former.

Recently Cobb and Church (1983) have proposed some method-of-moments estimators for the location parameter of a location-scale family of quantal response curves of the type (2.1). When applied to the logistic family, these estimators performed far better in small samples than the MLE and the minimum modified chi-square estimators in a numerical comparison of their exact distributions and mean squared errors. A more detailed study of these new estimators is needed before they can be recommended for a wider use.

3.2.2 Probit Model

Here no special simplifications result as in the case of the logistic model. Thus the procedure is the same as the general procedure described earlier. The standard normal pdf $\phi(\cdot)$ is substituted for $f(\cdot)$ when computing the empirical weights $\hat{w}_j$ by (3.23).

3.3 Nonparametric/Robust Estimation

In this section we no longer assume the parametric model (3.1). We assume only that $p(x)$ is nondecreasing in x with $p(-\infty) = 0$ and $p(+\infty) = 1$. In §3.3.1 we discuss the problem of estimating $p(x)$ and in §3.3.2 we discuss the problem of estimating the mean tolerance θ (which equals the ED50 μ if $p(x)$ is skew-symmetric).

3.3.1 Isotonic Estimation

Ayer et al. (1955) showed that the maximum likelihood estimates of the parameters $p_j = p(x_j)$ (i.e., the p_j-values which maximize the likelihood function (3.2)) under the order restriction

$$p_1 \le p_2 \le \ldots \le p_m \tag{3.29}$$

are given by

$$\tilde{p}_j = \max_{1 \le u \le j} \min_{j \le v \le m} \left[\frac{\sum_{i=u}^{v} r_i}{\sum_{i=u}^{v} n_i} \right] \qquad (1 \le j \le m). \tag{3.30}$$

The $\tilde{p}_j$ are known as the isotonic estimates of the p_j. For a general reference on the topic of estimation under order restrictions, see Barlow et al. (1972).

The estimates (3.30) are equal to the proportions $\hat{p}_j$ if the latter satisfy the order restriction (3.29); if not then (3.30) can be computed by a stepwise pooling process as follows: For any adjacent proportions not satisfying the order constraint, say, for $\hat{p}_j > \hat{p}_{j+1}$, find a common pooled adjusted estimate $\tilde{p}_j = \tilde{p}_{j+1} = (r_j + r_{j+1})/(n_j + n_{j+1})$. Continue this process until all the adjusted estimates satisfy (3.29). These final adjusted estimates are the same as those given by (3.30).

An isotonic estimate of $p(x)$, say $\tilde{p}(x)$, is obtained by connecting the successive points $(x_j, \tilde{p}_j)$ by straight lines. An estimate of the ED50 is then given by $\tilde{p}^{-1}(0.5)$.

3.3.2 Nonparametric/Robust Estimation of the Mean Tolerance

Various nonparametric/robust estimators of the mean tolerance θ have been proposed in the literature. These are often used as estimators of the ED50 with the implicit assumption that $p(x)$ is skew-symmetric. We refer the reader to articles by Miller (1973), Hamilton (1979), and Miller and Halpern (1980) for comparisons among these estimators. Here we provide only a brief summary of the estimators which are commonly used and/or which per-

formed well in the comparisons carried out in the aforementioned articles.

3.3.2.1 Spearman-Kärber Estimator

The Spearman-Kärber estimator is given by

$$\tilde{\theta}_{SK} = \int_{-\infty}^{\infty} x \, d\tilde{p}(x)$$

$$= \sum_{j=1}^{m-1} (\tilde{p}_{j+1} - \tilde{p}_j)(x_{j+1} + x_j)/2 \qquad (3.31)$$

where the $\tilde{p}_j$ are given by (3.30). The classical version of the Spearman-Kärber estimator $\hat{\theta}_{SK}$ uses the $\hat{p}_j$ in place of the $\tilde{p}_j$ in (3.31) but in that case it is possible to have some of the $\hat{p}_{j+1} - \hat{p}_j < 0$. From (3.31) it can be seen that $\hat{\theta}_{SK}$ is not an unbiased estimator of θ but rather of the discretized means of the tolerance distribution $p(x)$; the same is true of $\tilde{\theta}_{SK}$ asymptotically ($n_j \to \infty \ \forall j$). Thus it is appropriate to use $\tilde{\theta}_{SK}$ or $\hat{\theta}_{SK}$ to estimate θ only if the grid of the x_j values is very fine. The exact variance of $\hat{\theta}_{SK}$ (which is also the asymptotic variance of $\tilde{\theta}_{SK}$) is given by

$$\mathrm{Var}(\hat{\theta}_{SK}) = \frac{p_1(1-p_1)}{n_1}\left(\frac{x_2+x_1}{2}\right)^2 + \frac{p_m(1-p_m)}{n_m}\left(\frac{x_m+x_{m-1}}{2}\right)^2$$

$$+ \sum_{j=2}^{m-1} \frac{p_j(1-p_j)}{n_j}\left(\frac{x_{j+1}-x_{j-1}}{2}\right)^2. \qquad (3.32)$$

When the doses are equally spaced with $x_{j+1} - x_j = d > 0$ $(1 \le j \le m-1)$, we find that (3.31) simplifies to

$$\tilde{\theta}_{SK} = \left(x_m + \frac{d}{2}\right)\tilde{p}_m - \left(x_1 - \frac{d}{2}\right)\tilde{p}_1 - d\sum_{i=1}^{m}\tilde{p}_i. \qquad (3.33)$$

The performance of the Spearman-Kärber estimator is superior (based on the usual criteria) to two other nonparametric estimator of θ namely the Reed-Muench and the Behrens-Dragstedt estimators that have been proposed (Finney 1978, Miller 1973). For additional properties of the Spearman-Kärber estimator, see Brown (1961), Church and Cobb (1973), Chmiel (1976) and the references listed therein.

3.3.2.2 Robust Estimators

Many robust estimators of θ, have been proposed and compared by Hamilton (1979) and Miller and Halpern (1980). These are analogues of the commonly used robust estimators of location of a symmetric distribution such as L-estimators (of which the trimmed mean is a special case) and M-estimators. The studies by Hamilton, and Miller and Halpern show that the trimmed Spearman-Kärber estimator with about 10% trimming is the preferred choice among many alternatives that they studied. Here a 100α% trimmed Spearman-Kärber estimator $\tilde{\theta}_{SK}^{(\alpha)}$ is given by the following: For $0 < \alpha < 1/2$ let $\tilde{\mu}^{(\alpha)}$ and $\tilde{\mu}^{(1-\alpha)}$ be such that $\tilde{p}(\tilde{\mu}^{(\alpha)}) = \alpha$ and $\tilde{p}(\tilde{\mu}^{(1-\alpha)}) = 1-\alpha$ where $\tilde{p}(x)$ is the estimate of $p(x)$ obtained by linearly interpolating between the $\tilde{p}_j$. Then

$$\tilde{\theta}_{SK}^{(\alpha)} = \frac{1}{1-2\alpha}\int_{\tilde{\mu}^{(\alpha)}}^{\tilde{\mu}^{(1-\alpha)}} x\, d\tilde{p}(x)$$

$$= \frac{1}{2(1-2\alpha)}\left\{(\tilde{\mu}^{(\alpha)} + x_{\ell_\alpha+1})(\tilde{p}_{\ell_\alpha+1} - \alpha)\right.$$

$$+ \sum_{j=\ell_\alpha+1}^{u_\alpha-2} (x_j + x_{j+1})(\tilde{p}_{j+1} - \tilde{p}_j)$$

$$+ (\tilde{\mu}^{(1-\alpha)} + x_{u_\alpha-1})(1-\alpha-\tilde{p}_{u_\alpha-1}) \Bigg\} \qquad (3.34)$$

where $\ell_\alpha = \max\{j:\tilde{p}_j \le \alpha\}$ and $u_\alpha = \min\{j:\tilde{p}_j \ge 1-\alpha\}$.

The Monte Carlo comparisons of Hamilton (1979) and the theoretical asymptotic efficiency calculations of Miller and Halpern (1980) also showed that the performance of the estimators given in §§3.1 and 3.2 for parametric models is very poor if the actual tolerance distribution has heavy tails.

3.4 Design Aspects

Not much research has been carried out concerning the question of how to design an experiment "optimally," i.e., at what dose-levels should the observations be taken, and what proportion of the total sample size should be allocated to each dose-level, to estimate some selected parameter(s) associated with a quantal response curve. Brown (1966) gives some practical guidelines for designing an experiment for estimating the ED50. Bayesian design criteria have been considered by Freeman (1970), Tsutakawa (1972,1980) and Leonard (1982a,b). Some of the designs considered in these papers are sequential in nature. Here we discuss non-Bayesian approaches proposed by Hoel and Jennrich (1980) and Abdelbasit and Plackett (1983).

Hoel and Jennrich consider the problem of the "optimal" extrapolation design for estimating $p(x^*)$ (and also $p(x^*)-p(0)$; note that here x is not transformed to the logarithmic scale) when the observations are to be taken in the interval $[a,b]$ with $0 < x^* < a < b$. These estimation problems arise

when the results of an experiment to study the effect of a suspected carcinogen on laboratory animals at high dose-levels, are to be extrapolated to low dose-levels. Hoel and Jennrich assume the model

$$p(x) = 1 - \exp(-\sum_{i=0}^{p} \beta_i x^i) \tag{3.35}$$

which is an exponential model. They determine the "optimal" design which minimizes the asymptotic variance of the weighted least-squares estimate of $p(x^*)$; this is done by using the known results on optimal extrapolation designs for polynomial regression (Hoel 1966). The optimal design is supported on $p+1$ points. A method of determining these points (using the Chebyshev polynomials) and the proportion of observations to be allocated to each point is given in the Hoel-Jennrich article.

Abdelbasit and Plackett (1983) considered for the logistic model the problem of maximizing the determinant of the information matrix of $\mu = -\alpha/\beta$ and β using two and three-point designs. They considered only symmetric designs (i.e., $p(x_1) = 1-p(x_2)$ for a two-point design, and $p(x_1) = 1-p(x_3)$, $p(x_2) = 0.5$ for a three-point design with equal number of observations at each x_i for each design). Clearly, given that $p(x_1) = p$ for some specified $p \in (0,1)$, we need some prior estimates of μ and β to determine the design points. Abdelbasit and Plackett determine the optimal values of p for two and three-point designs. They study a multistage procedure wherein the results up to a given stage of the experiment are used to update the estimates of μ and β and thence to determine new design points for the next stage of the experiment.

4. Single Quantal Response Curve: Estimation Using Sequential Procedures

In this section we discuss two sequential procedures (i)

the Robbins-Monro stochastic approximation method, and (ii) the up-and-down method. Both these procedures are nonparametric in that they do not assume a specific model for the quantal response curve. A good elementary reference for both the procedures is Wetherill (1975, Ch.10). An advanced reference for stochastic approximation procedures is a book by Wasan (1969) and a review article by Schmetterer (1961). These references also contain extensive bibliographies.

4.1 Robbins-Monro Stochastic Approximation Method

The Robbins-Monro (1951) stochastic approximation method can be applied to the problem of estimating the ED100q = $\mu^{(q)}$ of a quantal response curve. We describe the method for this problem: Fix a positive constant $c > 0$. Denote the outcome of a single Bernoulli trial at dose-level x_i by $y(x_i) = 1$ ("success") or 0 ("failure"). Perform the initial trial at some dose-level x_1 (chosen randomly or nonrandomly). In general, given the outcome of the ith trial, $y(x_i)$, choose the dose-level x_{i+1} for the $(i+1)$th trial by the recursive relation

$$x_{i+1} = x_i - \frac{c}{i}\{y(x_i)-q\} \quad (1 \le i \le n) . \tag{4.1}$$

Stop the experiment after n observations are made where n is usually fixed in advance, and then use x_{n+1} (i.e., the dose-level for the next trial) as an estimate of $\mu^{(q)}$.

Sacks (1958) showed that asymptotically $(n \to \infty)$, under suitable regularity conditions $\sqrt{n}\,(x_n - \mu^{(q)})$ is normally distributed with zero mean and variance given by

$$\frac{c^2 q(1-q)}{\{2c\gamma^{(q)}-1\}} \tag{4.2}$$

provided that $\gamma^{(q)} > 1/2c$ where

$$\gamma^{(q)} = \frac{dp(x)}{dx}\Bigg|_{x=\mu^{(q)}} \tag{4.3}$$

is the slope of the quantal response curve at $\mu^{(q)}$. The optimal choice of c that minimizes (4.2) is

$$c = \frac{1}{\gamma^{(q)}}, \tag{4.4}$$

and the corresponding minimum variance is

$$\frac{q\,(1-q)}{\{\gamma^{(q)}\}^2}. \tag{4.5}$$

Note that (4.5) is the Cramer-Rao lower bound on the asymptotic variance of the MLE of $\mu^{(q)}$ and thus x_n is asymptotically efficient if (4.4) is used in the recursion (4.1). Note also that for the logistic model (2.6),

$$\gamma^{(q)} = \beta q\,(1-q) \tag{4.6}$$

while for the probit model (3.15),

$$\gamma^{(q)} = \frac{1}{\sigma}\,\phi\left(\frac{\mu^{(q)}-\mu}{\sigma}\right). \tag{4.7}$$

In practice, of course, $\gamma^{(q)}$ is unknown and therefore the optimal choice (4.4) cannot be implemented. It is readily seen that if because of not knowing $\gamma^{(q)}$, one uses a c-value in (4.1) which is r times $(0 < r < \infty)$ the optimal value (4.4) then the resulting asymptotic variance is $r^2/(2r-1)$ times the minimum value (4.5). An examination of this factor shows that moderate errors in prior guesses of $\gamma^{(q)}$ (and hence c) do not have very adverse effects on the asymptotic efficiency of the Robbins-Monro estimator.

However, a poor guess of $\gamma^{(q)}$ can lead to a substantial loss in the asymptotic efficiency.

All of the above discussion is concerned with large sample properties of the Robbins-Monro method. Small sample properties of this method are studied in Wetherill (1963), Cochran and Davis (1965) and Davis (1971). These authors also suggest several modifications of the basic method to improve its performance in small samples, e.g., to reduce the bias (which can be particularly severe when estimating $\mu^{(q)}$ for an extreme value of q such as 0.1 or 0.9) that can result because of a poorly chosen starting value x_0, and the much larger small-sample variance than (4.5) predicted by the asymptotic theory.

4.2 Adaptive Modifications of the Robbins-Monro Method

To obviate the difficulty of not knowing the optimal value (4.4) of c to use in (4.1) and also to correct the drawback of the classical Robbins-Monro method that it does not provide a good estimate of $\gamma^{(q)}$, several adaptive modifications of the Robbins-Monro method have been proposed. We review these modifications in the present section.

Anbar (1977,1978) proposed an adaptive procedure which estimates $\gamma^{(q)}$ at each stage of the experiment. This is done by first using a linear approximation to the quantal response curve in the region of $\mu^{(q)}$ and then applying the usual least-squares formula to obtain

$$\hat{\gamma}_i^{(q)} = \frac{\sum_{j=1}^{i} (x_j - \bar{x}_i)(y_j - \bar{y}_i)}{\sum_{j=1}^{i} (x_j - \bar{x}_i)^2} \qquad (2 \le i \le n). \tag{4.8}$$

Here $y_j = y(x_j)$, $\bar{y}_i = \Sigma_{j=1}^{i} y_j/i$ and $\bar{x}_i = \Sigma_{j=1}^{i} x_j/i$. To avoid the instability associated with initial observations it may be desirable to use, instead of (4.8), a "moving average" estimate based on some fixed number of the

latest observations. Another alternative suggested by Anbar is to specify two positive constants $L < U$ and truncate the estimate (4.8) below and above by L and U, respectively. Then (4.1) is modified to

$$x_{i+1} = x_i - \frac{c_i}{i}\{y(x_i) - q\} \qquad (1 \le i \le n) \tag{4.9}$$

where c_i is the inverse of the truncated estimate of $\gamma^{(q)}$ computed as just described. Anbar (1978) showed that the resulting estimator is strongly consistent and $\sqrt{n}(x_n - \mu^{(q)})$ is asymptotically normal with zero mean and variance $= q(1-q)/\{\gamma^{(q)}\}^2$; i.e., the estimator is asymptotically efficient. (Similar results were obtained earlier by Venter (1967) and Fabian (1973) using a different estimator of slope which requires taking observations at a pair of x-values at each stage.)

Lai and Robbins (1978) proposed an adaptive procedure which is very similar to Anbar's but their analysis was done for the case in which the $Y(x_i)$ have a constant variance.

In all of the above procedures no explicit use is made of the non-decreasing property of p(x); also the final estimate x_{n+1} can be thrown off by a few wild observations near the end. To obviate these difficulties Mukerjee (1981) proposed a modification of the Robbins-Monro method which constructs an isotonic estimate of p(x) (see §3.3.1) at each stage of the experiment based on all of the data collected up to that stage. Let $\tilde{p}_i(x)$ denote such an estimate at the ith stage of the experiment $(i = 1,2,\ldots)$. Then, roughly speaking, the next observation is taken at $\tilde{p}_i^{-1}(q)$, i.e., $x_{i+1} = \tilde{p}_i^{-1}(q)$. Mukerjee restricted the choice of dose-levels $\{x_i\}$ to a fixed grid of equispaced points which is also the case with the up-and-down method discussed in the following section. Under certain regularity conditions, Mukerjee showed the almost sure convergence to $\mu^{(q)}$ of his

sequential estimation procedure.

Wu (1983) made a proposal similar to Mukerjee's proposal but which uses a parametric model to determine the dose-level at which the next observation is to be taken. In particular, Wu proposed that at the ith stage of the experiment, an estimate $\hat{p}_i(x)$ of the unknown quantal response curve $p(x)$ be made assuming the logistic model (2.6) for $p(x)$. Wu points out that such a parametric assumption is required because with small samples it is not feasible to obtain a good smooth nonparametric estimate of $p(x)$. At the ith stage $(i = n_0 + 1, n_0 + 2, \ldots n)$ of the experiment Wu estimates $p(x)$ by

$$\hat{p}_i(x) = \frac{1}{1 + \exp\{-(\hat{\alpha}_i + \hat{\beta}_i x)\}} \tag{4.10}$$

where $(\hat{\alpha}_i, \hat{\beta}_i)$ is the MLE of (α, β) at that stage and n_0 is some fixed initial sample size; the dose-levels for the first n_0 observations are chosen based on whatever prior knowledge the experimenter may have about $p(x)$. The MLE $(\hat{\alpha}_i, \hat{\beta}_i)$ is computed by iteratively solving the equations (3.8) (after making appropriate changes in notation). The dose-level for the $(i+1)$th stage is then obtained by solving the equation $\hat{p}_i(x_{i+1}) = q$ which yields (cf. (2.7))

$$x_{i+1} = \frac{\ell n\left(\frac{q}{1-q}\right) - \hat{\alpha}_i}{\hat{\beta}_i} . \tag{4.11}$$

Wu noted that (4.11) can result in an unduly large change from x_i to x_{i+1} in certain "ill-posed" cases. Therefore he suggested a truncated version of (4.11) which works as follows: Let c_i be the solution to the equation

$$\frac{\ell n\left(\frac{q}{1-q}\right) - \hat{\alpha}_i}{\hat{\beta}_i} = x_i - \frac{c_i}{i}\{y(x_i) - q\}.$$

Then the (i + 1)th dose-level is chosen as

$$x_{i+1} = x_i - \frac{c_i^*}{i} \{y(x_i) - q\} \tag{4.12}$$

where for some prespecified positive constant C,

$$c_i^* = \max\{-C, \min(c_i, C)\}.$$

Note that specifying a large C amounts to virtually no truncation.

Wu carried out a simulation study for comparing the performance of various sequential designs for small samples ($12 \le n \le 35$) for the problem of estimating the ED50 of a quantal response curve with an unknown functional form. He considered the Robbins-Monro design (4.1), the adaptive Robbins-Monro design (4.9) (with a slightly different truncation scheme than the one proposed by Anbar), Wu's design (4.12), and the up-and-down design described in the following section. Each design was used with selected values of its design parameter, viz., c for the Robbins-Monro design, C for (4.12) and so on. The mean square error of the estimator was used as the performance criterion in each case. In these simulations the design (4.12) with a choice of moderate to large C performed much better than its competitors and is thus a serious choice in practical work.

4.3 Up-and-Down Method

The up-and-down method proposed by Dixon and Mood (1948) for estimating $\mu = \mu^{(0.5)}$ and $\gamma = \gamma^{(0.5)}$ for a quantal response curve operates as follows: Fix an equispaced grid of dose-levels, say, $L = \{\ell_i = c + id, i = 0, \pm 1, \pm 2, \ldots\}$ for some prespecified values of c and d. Choose a starting dose-level x_1 from this grid at the best prior guess of μ available and then perform the trials sequentially. In general, at the ith trial if $y(x_i) = 1$ (resp.,0) then take the next observation at the next lower (resp.,higher)

level, i.e., $x_{i+1} = x_i - d$ (resp., $x_{i+1} = x_i + d$). Stop after a predetermined number n of trials and then estimate μ and γ by the method described below.

Dixon and Mood proposed the method of maximum likelihood for estimating μ and γ assuming the probit model for the quantal response curve. This method results in equations (3.16) for the estimates of μ and σ (note from (4.7) that for the probit model $\gamma = (\sigma\sqrt{2\pi})^{-1}$) and in equations (3.19) for their asymptotic variances. Dixon and Mood suggested a simplified version of this method which is based on the fact that in the up-and-down method the number of successes r_i (say) at any dose-level ℓ_i do not differ by more than unity from the number of failures s_{i-1} (say) at the next lower dose-level ℓ_{i-1} and therefore for large n one can set $r_i = s_{i-1}$ for all i with negligible loss in information. Wetherill (1963,1975 Ch.10) proposed the use of the logistic model instead of the probit model. Brownlee, Hodges and Rosenblatt (1953) proposed a very simple estimate of μ namely the arithmetic average of the dose-levels used (excluding x_1 but including x_{n+1}), i.e., $n^{-1} \sum_{i=2}^{n+1} x_i$; this estimate is asymptotically $(n \to \infty)$ equivalent to the MLE of μ proposed by Dixon and Mood.

Derman (1957) generalized the original up-and-down method to the nonparametric setting and to the problem of estimating $\mu^{(q)}$ for arbitrary $q \in (0,1)$. He proposed the following recursive scheme for selecting the dose-levels: Let x_1 be a starting dose-level in grid L and let

$$x_{i+1} = \begin{cases} x_i - d & \text{wp } \frac{1}{2q} \text{ if } y(x_i) = 1 \\ x_i + d & \text{wp } 1 - \frac{1}{2q} \text{ if } y(x_i) = 1 \\ x_i + d & \text{wp } 1 \text{ if } y(x_i) = 0 \end{cases}$$

where it is assumed that $1/2 \leq q < 1$. After n trials, $\mu^{(q)}$ is estimated by the mode of the realized process $x_2, x_3, \ldots, x_{n+1}$. Wetherill (1963, §12) has suggested some alternative up-and-down methods for estimating $\mu^{(q)}$; these have not been studied in detail.

Brownlee et al. (1953) pointed out that the up-and-down method gives a very poor estimate of the slope γ unless n is very large. This was confirmed in the sampling studies of Wetherill (1963) who noted that if the product $d\gamma$ is small then the $\mathrm{Var}(\hat{\gamma})$ can be arbitrarily large. Thus a large d is required to obtain a reasonably precise estimate of γ. However, a large d can lead to a large bias in the estimate of μ. Wetherill (1963) proposed a two-phase sequential procedure which attempts to obviate these difficulties. In the first phase of the up-and-down method a large value of d is employed. This phase is continued until k runs of "successes" or "failures" are observed where k is a number fixed in advance. Based on the results of this phase a preliminary estimate of μ is computed. The second phase is started at this estimated value of μ with a grid d which is half the value of d used in the first phase. The second phase is terminated after a fixed number of trials are performed and then the final estimate of μ is computed using one of the methods described earlier. This modification (referred to by Wetherill as the "k-changes rule") performs well in sampling experiments but its theoretical properties are unknown. Hsi (1969) has proposed a modification of the up-and-down method which takes multiple observations at each stage; his modification also permits estimation of the ED100q for arbitrary $q \in (0,1)$.

5. Several Quantal Response Curves: Multiple Comparisons Under the Logistic Model

In this section we return to the set up of §1 where we had $k \geq 2$ populations with $p_i(x)$ as the quantal response curve associated with the ith population Π_i $(1 \leq i \leq k)$. Very little literature exists on the problem of

comparing two or more quantal response curves simultaneously, one major exception being analysis of dilution assays (Finney 1978).

Reiersøl (1961) considered the goal of making all pairwise comparisons between k populations under the logistic model

$$p_i(x) = \frac{1}{1+\exp\{-(\alpha_i + \beta_i x)\}} \quad (1 \le i \le k). \tag{5.1}$$

In particular, he considered three separate simultaneous pairwise hypotheses testing problems:

$$\left.\begin{aligned} H_{ii'}^{(a)}: &\quad \beta_i = \beta_{i'}, && (5.2a)\\ H_{ii'}^{(b)}: &\quad (\alpha_i + \beta_i x) = (\alpha_{i'} + \beta_{i'} x) \text{ for specified } x \in \mathbb{R} && (5.2b)\\ H_{ii'}^{(c)}: &\quad \mu_i^{(q)} = \mu_{i'}^{(q)} \text{ for specified } q \in (0,1) && (5.2c) \end{aligned}\right\} \quad (1 \le i < i' \le k)$$

where $\mu_i^{(q)}$ for population Π_i is defined by (2.7) $(1 \le i \le k)$. Reiersøl derived Scheffé-type simultaneous tests for (5.2a), (5.2b), and (5.2c) when a single-stage experiment of the type described at the beginning of §3 is performed independently for each population. We now describe his procedure in detail.

Choose dose-levels $x_{i1} < x_{i2} < \ldots < x_{im_i}$ for experimenting with population $\Pi_i (1 \le i \le k)$. Take n_{ij} observations at dose-level x_{ij} for population Π_i; let r_{ij} be the corresponding number of successes and let $\hat{p}_{ij} = r_{ij}/n_{ij}$ be the proportion of successes $(1 \le j \le m_i,\ 1 \le i \le k)$. By analogy with (3.27) and (3.28) define

$$u_{ij} = \ell n \left\{ \frac{\hat{p}_{ij} + 1/2\, n_{ij}}{1-\hat{p}_{ij} + 1/2\, n_{ij}} \right\} \quad (1 \le j \le m_i, 1 \le i \le k), \tag{5.3}$$

and

$$\hat{w}_{ij} = n_{ij}\hat{p}_{ij}(1 - \hat{p}_{ij}) \qquad (1 \le j \le m_i, 1 \le i \le k), \tag{5.4}$$

respectively. Then by analogy with (3.25) compute

$$\hat{\alpha}_i = \hat{\bar{u}}_i - \hat{\beta}_i \hat{\bar{x}}_i \text{ and } \hat{\beta}_i = \frac{\sum_{j=1}^{m_i} \hat{w}_{ij}(u_{ij} - \hat{\bar{u}}_i)x_{ij}}{\sum_{j=1}^{m_i} \hat{w}_{ij}(x_{ij} - \hat{\bar{x}}_i)^2} \quad (1 \le i \le k)$$

where

$$\hat{\bar{x}}_i = \frac{\sum_{j=1}^{m_i} \hat{w}_{ij}x_{ij}}{\sum_{j=1}^{m_i} \hat{w}_{ij}} \text{ and } \hat{\bar{u}}_i = \frac{\sum_{j=1}^{m_i} \hat{w}_{ij}u_{ij}}{\sum_{j=1}^{m_i} \hat{w}_{ij}} \quad (1 \le i \le k). \tag{5.5}$$

Also compute

$$\begin{bmatrix} a_i & c_i \\ c_i & b_i \end{bmatrix} = \begin{bmatrix} \sum_{j=1}^{m_i} \hat{w}_{ij} & \sum_{j=1}^{m_i} \hat{w}_{ij}x_{ij} \\ \sum_{j=1}^{m_i} \hat{w}_{ij}x_{ij} & \sum_{j=1}^{m_i} \hat{w}_{ij}x_{ij}^2 \end{bmatrix}^{-1} \quad (1 \le i \le k).$$

Reirsol proposed the following procedure for testing (5.2a):

$$\text{Reject } H_{ii'}^{(a)}, \iff (\hat{\beta}_i - \hat{\beta}_{i'})^2 > \chi^2_{\alpha,2k-2}(b_i + b_{i'}) \quad (1 \le i < i' \le k)$$

where $\chi^2_{\alpha,2k-2}$ is the upper α-point of the chi-square distribution with $2k-2$ degrees of freedom. For testing (5.2b) he proposed the procedure:

$$\text{Reject } H_{ii'}^{(b)}, \iff \{\hat{\alpha}_i + \hat{\beta}_i x - (\hat{\alpha}_{i'} + \hat{\beta}_{i'}x)\}^2$$

$$> \chi^2_{\alpha,2k-2} \{a_i + b_i x^2 + 2c_i x + a_{i'} + b_{i'} x^2 + 2c_{i'} x\} \quad (1 \le i < i' \le k).$$

For testing (5.2c) he proposed the procedure

$$\text{Reject } H^{(c)}_{ii'} \iff \{(Q - \hat{\alpha}_i)\hat{\beta}_{i'} - (Q - \hat{\alpha}_{i'})\hat{\beta}_i\}^2$$

$$> \chi^2_{\alpha,2k-2} \min_t G_{ii'}(t) \quad (1 \le i < i' \le k)$$

where the minimum is taken over $-\infty < t < +\infty$,

$$G_{ii'}(t) = (a_i + a_{i'})\{t\hat{\beta}_i + (1-t)\hat{\beta}_{i'}\}^2$$

$$+ (b_i + b_{i'})\{t\hat{\alpha}_i + (1-t)\hat{\alpha}_{i'} - Q\}^2$$

$$-2(c_i + c_{i'})\{t\hat{\beta}_i + (1-t)\hat{\beta}_{i'}\}\{t\hat{\alpha}_i + (1-t)\hat{\alpha}_{i'} - Q\},$$

and

$$Q = \ell n\left(\frac{q}{1-q}\right) .$$

All three procedures follow from Scheffé's projection method. Under the hypothesis that all k quantal response curves are identical (i.e., $\alpha_1 = \ldots = \alpha_k$ and $\beta_1 = \ldots = \beta_k$) asymptotically ($n_{ij} \to \infty \ \forall i,j$) the probability that any of the above tests makes a false rejection is no greater than α.

Jensen (1976) deals with the problem of comparing several dose-response curves with a standard; however, he assumes continuous normally distributed responses.

6. Several Quantal Response Curves: Selection Problem Under the Logistic Model

6.1 A Single-Stage Procedure

We begin by considering the simplest setup namely that of §5 wherein

we assumed the logistic model (5.1) for each population. Here we further assume that

$$\beta_1 = \beta_2 = \ldots = \beta_k = \beta \text{ (say)}, \tag{6.1}$$

(thus the k quantal response curves are members of a location parameter family) and that the common value of β is known.

As formulated in §1, the goal of the experimenter is to select the population associated with $\mu_{[1]}^{(q)} = \min_{1 \le i \le k} \mu_i^{(q)}$ where $q \in (0,1)$ is specified. From (2.7) we see that because of the assumption (6.1), the goal of selecting the population associated with the smallest $\mu_i^{(q)}$ for any $q \in (0,1)$ is equivalent to selecting the population associated with the largest α_i.

Let $\alpha_{[1]} \ge \alpha_{[2]} \ge \ldots \ge \alpha_{[k]}$ be the ordered α_i's so that $\alpha_{[i]}$ is associated with $\mu_{[i]}^{(q)} (1 \le i \le k)$. The probability requirement (1.3) can be rewritten as

$$P(CS) \ge P^* \quad \text{whenever} \quad \alpha_{[1]} - \alpha_{[2]} \ge \beta \delta^* \tag{6.2}$$

where now the "correct selection (CS)" refers to selecting the population associated with $\alpha_{[1]}$.

We propose the following single-stage natural selection procedure: For each population Π_i choose equispaced dose-levels $x_{i1}, x_{i2}, \ldots, x_{im}$ where $x_{i,j+1} - x_{ij} = d_i > 0$ $(1 \le i \le k, 1 \le j \le m-1)$. For each population Π_i take n independent observations at each x_{ij} and let $N = mn$ denote the total number of observations per population. Let r_{ij} denote the number of successes for population Π_i at dose-level x_{ij} and let $\hat{p}_{ij} = r_{ij}/n$ $(1 \le i \le k, 1 \le j \le m)$. Compute the MLE $\hat{\alpha}_i$ by solving the equation (which corresponds to the first of the two equations (3.8)):

$$\sum_{j=1}^{m} p_{ij} = \sum_{j=1}^{m} \hat{p}_{ij} \quad (1 \le i \le k)$$

where

$$p_{ij} = \frac{1}{1 + \exp\{-(\alpha_i + \beta x_{ij})\}} \quad (1 \le i \le k, 1 \le j \le m). \qquad (6.3)$$

Alternatively compute the weighted least squares estimate of α_i by

$$\hat{\alpha}_i = \hat{\bar{u}}_i - \hat{\beta}\hat{\bar{x}}_i \quad (1 \le i \le k)$$

where the quantities $\hat{\bar{u}}_i$ and $\hat{\bar{x}}_i$ are given by (5.5) and the u_{ij} and $\hat{w}_{ij}$ required in their computation are given by (5.3) and (5.4), respectively. Select the population yielding $\max_{1 \le i \le k} \hat{\alpha}_i$ and assert that it is associated with $\alpha_{[1]}$ (i.e., with $\mu_{[1]}^{(q)}$).

We take the set of dose-levels $\{x_{ij} \ (1 \le i \le k, 1 \le j \le m)\}$ as given and address the problem of determining the smallest sample size N per population (or equivalently the smallest $n = N/m$) which will guarantee (6.2) when used in the selection procedure just described. We provide a <u>large sample</u> ($m \to \infty$, $n \to \infty$) solution to this problem.

First note that for large n, the MLE (and also the weighted least-squares estimator) $\hat{\alpha}_i$ is approximately normal with

$$E(\hat{\alpha}_i) = \alpha_i, \ \mathrm{Var}(\hat{\alpha}_i) = \frac{1}{\sum_{j=1}^{m} w_{ij}} \quad (1 \le i \le k)$$

and, of course, the $\hat{\alpha}_i$ are independently distributed. The formula for the variance follows by a calculation analogous to that made in (3.9) but here it results in a simple one term expression because β is assumed known and

hence is not estimated. If m is large and d_i is small $(1 \le i \le k)$ then we can write

$$\sum_{j=1}^{m} w_{ij} = n \sum_{j=1}^{m} p_{ij}(1 - p_{ij})$$

$$= \frac{n}{d_i} \sum_{j=1}^{m} p_i(x_{ij})\{1 - p_i(x_{ij})\}d_i$$

$$\tilde{=} \frac{n}{d_i} \int_{-\infty}^{\infty} p_i(x)\{1 - p_i(x)\}dx. \qquad (6.4)$$

In the above p_{ij} is given by (6.3) and $p_i(x)$ is given by (5.1) with $\beta_i = \beta \ \forall i$. For the logistic model we obtain from (3.7) that

$$p_i(x)\{1 - p_i(x)\} = \frac{1}{\beta} \cdot \frac{dp_i(x)}{dx} \quad (1 \le i \le k). \qquad (6.5)$$

Substituting (6.5) in (6.4) and noting that $dp_i(x)/dx$ is a (logistic) density function which integrates to unity, we obtain

$$\sum_{j=1}^{m} w_{ij} \tilde{=} \frac{n}{\beta d_i} \quad (1 \le i \le k).$$

If a common spacing between dose-levels is used for all k populations, i.e., if $d_1 = d_2 = \ldots = d_k = d$ (say) then we have for large m and n that

$$\hat{\alpha}_i \sim N(\alpha_i, \frac{\beta d}{n}) \quad (1 \le i \le k).$$

Thus the estimators of the α_i are (approximately) normally distributed with a common known variance $\beta d/n$ and are independent. Based on the known results for the problem of selecting the largest normal mean when the populations have a common known variance (Bechhofer 1954) we can conclude that the

infimum of the large-sample P(CS) over the part of the parameter space where $\alpha_{[1]} - \alpha_{[2]} \geq \beta\delta^*$ is attained at the slippage configuration $\alpha_{[1]} - \delta^* = \alpha_{[2]} = \ldots = \alpha_{[k]}$ and this infimum is given by

$$\int_{-\infty}^{\infty} \Phi^{k-1}\left(x + \delta^*\sqrt{\frac{\beta n}{d}}\right) d\Phi(x). \tag{6.6}$$

Equating (6.6) to P^* we find that the desired n which guarantees (6.2) is given by

$$n = \left\langle \left(\frac{c_{k,P^*}}{\delta^*}\right)^2 \frac{d}{\beta} \right\rangle \tag{6.7}$$

where $\langle x \rangle$ denotes the smallest integer $\geq x$ and c_{k,P^*} is the solution to the equation

$$\int_{-\infty}^{\infty} \Phi^{k-1}(x + c)\, d\Phi(x) = P^*.$$

Selected values of c_{k,P^*} have been tabulated by Bechhofer (1954) and Gupta (1963). We note from (6.7) that n is decreasing in β (i.e., n increases as the slope of the quantal response curve decreases) and increasing in the common spacing d between dose-levels.

6.2 Discussion

The intention of the above exercise was not to provide a realistic solution to the selection problem but simply to illustrate the difficulties that arise even under very highly restrictive assumptions (e.g., the logistic dose-response curves with a common known slope, equispaced dose-levels, etc.). An exact small-sample solution is hard to obtain, and so it

was necessary to employ asymptotics in both the number of observations n at each dose-level and the number of dose-levels m. On the other hand, in case of nonintersecting quantal response curves (Figure 2) the exact small-sample procedure proposed by Bechhofer and Kulkarni (1982) for the underlying Bernoulli selection problem requires no additional assumptions but, as pointed out in Section 1, it does not necessarily guarantee the probability requirement (1.3).

Clearly, if the shapes of the quantal response curves are assumed to be completely unknown then fully sequential procedures with large samples must be used. Gupta and Huang (1975) have proposed such a Robbins-Monro type sequential sampling procedure followed by a natural selection terminal decision rule. But their results also are asymptotic. They also assume known lower bounds on the slopes of the quantal response curves. In practice only small samples are available from each population, and the work of Wetherill (1963) has shown that the asymptotic theory is not a good guide in that case. Thus the selection problem is still essentially wide open.

Acknowledgment

The author is extremely grateful to Robert Bechhofer who conceived the selection problem which motivated this literature survey, pointed out various difficulties (discussed in §1) associated with its solution, read the earlier drafts of this article and made suggestions for improvement, and through many discussions helped clarify a number of issues.

REFERENCES

1. Abdelbasit, K. M. and Plackett, R. L. (1983), "Experimental design for binary data," J. Amer. Statist. Assoc., 78, 90-98.

2. Anbar, D. (1977), "The application of stochastic approximation methods to the bio-assay problem," J. Statist. Planning and Inference, 1, 191-206.

3. Anbar, D. (1978), "A stochastic Newton-Raphson method," J. Statist. Planning and Inference, 2, 153-163.

4. Ashton, W. D. (1972), "The Logit Transformation," New York: Hafner Publishing Co.

5. Ayer, M., Brunk, H. D., Ewing, G. M., Reid, W. T., and Silverman, E. (1955), "An empirical distribution function for sampling with incomplete information," Ann. Math. Statist., 26, 641-647.

6. Barlow, R. E., Bartholomew, D. J., Bremner, J. M., and Brunk, H. D. (1972), "Statistical Inference Under Order Restrictions," New York: John Wiley and Sons, Inc.

7. Bechhofer, R. E. (1954), "A single-sample multiple decision procedure for ranking means of normal populations with known variances," Ann. Math. Statist., 25, 16-39.

8. Bechhofer, R. E. and Kulkarni, R. V. (1982), "Closed adaptive sequential procedures for selecting the best of $k \geq 2$ Bernoulli populations," Proc. Third Purdue Symp. on Statistical Decision Theory and Related Topics (eds. S. S. Gupta and J. Berger) New York: Academic Press, I, 61-108.

9. Berkson, J. (1955), "Maximum likelihood and minimum χ^2 estimates of the logistic function," J. Amer. Statist. Assoc., 50, 130-162.

10. Berkson, J. (1980), "Minimum chi-square, not maximum likelihood!," with discussion by B. Efron et al., Ann. Statist., 8, 457-487.

11. Brown, B. W., Jr. (1961), "Some properties of the Spearman estimator in bio-assay," Biometrika, 48, 293-302.

12. Brown, B. W., Jr. (1966), "Planning a quantal assay of potency," Biometrics, 22, 322-329.

13. Brownlee, K. A., Hodges, J. L., and Rosenblatt, M. (1953), "The up-and-down method with small samples," J. Amer. Statist. Assoc., 48, 262-277.

14. Church, J. D. and Cobb, E. B. (1973), "On the equivalence of Spearman-Kärber and maximum likelihood estimates of the mean," J. Amer. Statist. Assoc., 68, 201-202.

15. Cobb, E. B. and Church, J. D. (1983), "Small-sample quantal response methods for estimating the location parameter for a location-scale family of dose-response curves," J. Amer. Statist. Assoc., 78, 99-107.

16. Cochran, W. G. and Davis, M. (1965), "The Robbins-Monro method for estimating the median lethal dose," J. Roy. Statist. Soc., Ser. B, 27, 28-44.

17. Cox, D. R. (1970), "The Analysis of Binary Data," London: Methuen and Co.

18. Davis, M. (1971), "Comparison of sequential bio-assays in small samples," J. Royal Statist. Soc., Ser. B, 33, 78-94.

19. Derman, C. (1957), "Nonparametric up-and-down experimentation," Ann. Math. Statist., 28, 795-797.

20. Dixon, W. J. and Mood, A. M. (1948), "A method for obtaining and analyzing sensitivity data," J. Amer. Statist. Assoc., 43, 109-126.

21. Fabian, V. (1973), "Asymptotically efficient stochastic approximation; the RM case," Ann. Statist., 1, 486-495.

22. Finney, D. J. (1971), "Probit Analysis," 3rd Edition, Cambridge: Cambridge University Press.

23. Finney, D. J. (1978), "Statistical Methods in Biological Assay," 3rd Edition, New York: Hafner Publishing Co.

24. Freeman, P. R. (1970), "Optimal Bayesian sequential estimation of the median effective dose," Biometrika, 57, 79-89.

25. Golub, A. H. and Grubbs, F. E. (1956), "Analysis of sensitivity experiments when the levels of stimulus cannot be controlled," J. Amer. Statist. Assoc., 51, 257-265.

26. Gupta, S. S. (1963), "Probability integrals of multivariate normal and multivariate t," Ann. Math. Statist., 34, 794-828.

27. Gupta, S. S. and Huang, W. T. (1975), "Some selection procedures based on the Robbins-Monro stochastic approximation," Tamkang J. Math., 6, 301-314.

28. Hamilton, M. A. (1979), "Robust estimates of the ED 50," J. Amer. Statist. Assoc., 74, 344-354.

29. Hodges, J. L. and Lehmann, E. L. (1955), "Two approximations to the Robbins-Monro process," Proc. Third Berkeley Symp., 1, 95-104.

30. Hoel, P. G. (1966), "A simple solution for optimal Chebyshev regression extrapolation," Ann. Math. Statist., 37, 720-725.

31. Hoel, P. G. and Jennrich, R. I. (1979), "Optimal designs for dose response experiments in cancer research," Biometrika, 66, 307-316.

32. Hsi, B. P. (1969), "The up-and-down method in bio-assay," J. Amer. Statist. Assoc., 64, 147-162.

33. Jensen, D. R. (1976), "The comparison of several response functions with a standard," Biometrics, 32, 51-59.

34. Lai, T. L. and Robbins, H. (1978), "Adaptive design in regression and control," Proc. Nat. Acad. Sci., 75, 586-587.

35. Leonard, T. (1982a), "An inferential approach to the bio-assay design problem," MRC Tech. Summay Rep. #2416, Mathematics Research Center, University of Wisconsin, Madison, WI.

36. Leonard, T. (1982b), "A small sample evaluation of a Bayesian design method for quantal response models," Unpublished manuscript.

37. Miller, R. G. (1973), "Nonparametric estimators of the mean tolerance in bio-assay," Biometrika, 60, 535-542.

38. Miller, R. G. and Halpern, J. W. (1980), "Robust estimators for quantal bio-assay," Biometrika, 67, 103-110.

39. Mukerjee, H. G. (1981), "A stochastic approximation by observations on a discrete lattice using isotonic regression," Ann. Statist., 9, 1020-1025.

40. Neyman, J. (1949), "Contribution to the theory of the χ^2 test," Proc. First Berkeley Symp., 1, 239-273.

41. Prentice, R. L. (1976), "Generalization of the probit and logit methods for dose-response curves," Biometrics, 32, 761-768.

42. Reiersøl, O. (1961), "Linear and non-linear multiple comparisons in logit analysis," Biometrika, 48, 359-365.

43. Robbins, H. and Monro, S. (1951), "A stochastic approximation method," Ann. Math. Statist., 22, 400-407.

44. Sacks, J. (1958), "Asymptotic distribution of stochastic approximation procedures," Ann. Math. Statist., 29, 373-405.

45. Schmetterer, L. (1961), "Stochastic approximation," Proc. Fourth Berkeley Symp., 1, 587-609.

46. Schumacher, M. (1980), "Point estimation in quantal response models," Biom. J., 22, 315-334.

47. Shuster, J. J. and Yang, M. C. K. (1975), "A distribution-free approach to quantal response assays," The Canadian J. Statist., 3, 57-70.

48. Shuster, J. J. and Dietrich, F. H. (1976), "Quantal response assays in inverse regression," Commun. Statist., Ser. A, 293-305.

49. Silvapulle, M. J. (1981), "On the existence of maximum likelihood estimators for the binomial response model," J. Roy. Statist. Soc., Ser. B, 43, 310-313.

50. Sobel, M. (1967), "Nonparametric procedures for selecting the t populations with the t largest α-quantiles," Ann. Math. Statist., 38,1804-1816.

52. Tsutakawa, R. K. (1972), "Design of experiment for bio-assay," J. Amer. Statist. Assoc., 67, 584-590.

53. Tsutakawa, R. K. (1980), "Selection of dose levels for estimating a percentage point of a logistic quantal response curve," Appl. Statist., 29, 25-33.

54. Venter, J. H. (1967), "An extension of the Robbins-Monro procedure," Ann. Math. Statist., 38, 181-190.

55. Wasan, M. T. (1969), "Stochastic Approximation," Cambridge University Press: Cambridge.

56. Wetherill, G. B. (1963), "Sequential estimation of quantal response curves," J. Roy. Statist. Soc., Ser. B, 25, 1-48.

57. Wetherill, G. B. (1975), "Sequential Methods in Statistics," Second Edition, London: Methuen.

58. Wu, C. F. (1983), "Efficient model-based sequential designs for sensitivity experiments," Math. Sciences Research Institute, MSRI 042-83, Berkeley, CA.