REPEATED - MLE PROCEDURES FOR STOCHASTIC APPROXIMATION IN QUANTAL RESPONSE PROBLEMS*

Thomas Sellke

Purdue University

Several nonadaptive repeated-MLE procedures for stochastic approximation in quantal response problems are compared. Examination of how the shapes of the efficient score functions of individual observations affect the behavior of these rules shows that the nonadaptive logit-MLE rule is much less susceptible to misleading initial observations than are the nonadaptive probit-MLE rule or a nonadaptive Robbins-Monro rule. A consistency theorem for repeated-MLE procedures in quantal response problems is stated.

1. Introduction and summary.

Let $F: \mathbb{R} \to [0,1]$ be a quantal response curve. Thus, if we choose a stimulus level x_n, we observe a Bernoulli random variable y_n, indicating response or nonresponse to the stimulus, where

$$\Pr\{y_n = 1\} = 1 - \Pr\{y_n = 0\} = F(x_n). \tag{1.1}$$

Examples include bioassay, where $y_n = 1$ indicates the death of a test animal which has been given a dose of level x_n of a toxic substance, and reliability testing of components or materials, where $y_n = 1$ indicates the failure of a test item which has been subjected to a stress or shock of magnitude x_n. The problem

*Research supported by the National Science Foundation under grant #DMS 8401996.

AMS 1980 subject classifications. Primary 62L20; Secondary 62L05.

Key words: stochastic approximation, quantal response curve, logit, maximum likelihood.

of interest here is the estimation of the 100p-th percentile L_p of F, where $p \in (0,1)$ is given. The goal will be to find a scheme for sequentially choosing the stimulus levels $x_1, x_2, \ldots$ so that x_n converges to L_p as quickly as possible. In applications like those mentioned above, the quantal response function F will typically be continuous and increasing with $\lim_{x \to -\infty} F(x) = 0$ and $\lim_{x \to +\infty} F(x) = 1$, or at least with $\lim_{x \to -\infty} F(x) < p$ and $\lim_{x \to +\infty} F(x) > p$. Thus, a unique root L_p of

$$F(L_p) = p \tag{1.2}$$

which also satisfies $F(x) < p$ for $x < L_p$ and $F(x) > p$ for $x > L_p$ will typically exist and be unique. For reasons of tradition and generality, this paper will not assume that F is continuous or increasing, but rather only that there exists a reasonably well defined and unique 100p-th percentile. To wit, we assume the existence of a stimulus level L_p satisfying

$$F(L_p) = p \tag{1.3a}$$

$$\text{For every } 0 < \delta < 1, \sup_{\delta < h < \delta^{-1}} F(L_p - h) < p. \tag{1.3b}$$

$$\text{For every } 0 < \delta < 1, \inf_{\delta < h < \delta^{-1}} F(L_p + h) > p. \tag{1.3c}$$

The stochastic approximation problem described above is obviously one to which Robbins-Monro procedures can be applied. In the present setting, the standard, nonadaptive Robbins-Monro rules for recursively choosing the x_n's take the form

$$x_{n+1} = x_n - (y_n - p)/(nb), \tag{1.4}$$

where b is a positive constant. The Robbins-Monro rule (1.4) will be denoted by RM(b).

The RM(b) rule can be obtained by doing repeated maximum likelihood estimation under a simple linear model. (See Lemma 1 of Lai and Robbins (1979).) To be specific, suppose that our model for the distribution of y_n, given stimulus x_n, is

$$y_n = b(x_n - L_p) + p + \varepsilon_n, \tag{1.5}$$

where $b > 0$ is known and $\varepsilon_1, \varepsilon_2, \ldots$ are i.i.d. $N(0,\sigma^2)$ random variables. Then the x_{n+1} produced by the RM(b) rule is precisely the maximum likelihood estimate for L_p under the model (1.5).

Suppose that $\beta = \frac{d}{dx}F(x)\big|_{x=0}$ exists and is positive. It is well known that the asymptotically best choice of b in (1.4) is $b = \beta$. (See, for example, Sacks (1959), and Lai and Robbins (1979).) Although Lai and Robbins (1979) consider only stochastic approximation with i.i.d. "errors", it seems clear that their results and methods also apply in our setting. Thus, one can attain the asymptotically optimal convergence rate of x_n to L_p by replacing b in (1.4) with certain strongly consistent estimators $\hat{\beta}_n$ of β. Such adaptive RM procedures may to some extent be regarded as repeated MLE procedures based on the model

$$y_i = \beta(x_i - L_p) + p + \varepsilon_i, \tag{1.6}$$

where now both the location parameter L_p and the scale parameter β are unknown, and the errors ε_i remain i.i.d. $N(0,\sigma^2)$. The relationship between adaptive RM procedures on the one hand and repeated MLE (or "iterated least squares") procedures on the other is discussed in Lai and Robbins (1982). See also Wu (1985, 1986).

Wu (1985, 1986) has suggested that the x_n's be chosen to equal the maximum likelihood estimate of L_p under a parametric model which is more appropriate for quantal response that the models (1.5) and (1.6), which at first glance might seem almost grotesquely inappropriate. He has focused most of his attention on the location-scale logit model for F given by

$$(1.7) \qquad F(x) = G(x|\alpha,\lambda) = \{1 + e^{-\lambda(x-\alpha)}\}^{-1}, \ \lambda > 0, \ \alpha \in \mathbb{R}.$$

Wu (1985) has done Monte Carlo simulations to compare the performance of RM(b) and adaptive RM procedures with this "adaptive" (since the scale parameter λ is being estimated) logit-MLE for moderate sample sizes (n = 10 to 35). Five different functions (including a logit curve and a probit curve) were used as the true quantal response functions in these simulations, and Wu claims that a modification of his adaptive logit-MLE method with truncated step sizes generally outperforms RM(b) and adaptive RM procedures. He also claims that his method is asymptotically equivalent to adaptive RM if it is consistent. (Although the model (1.6) may seem "grotesquely inappropriate" for quantal responses, the asymptotic behavior of adaptive RM rules does not seem to be subject to general improvement even here. See Sellke (1986), Section 1, for an example.) However, Wu has not given any rigorous proof of consistency.

Adaptive designs are probably of greater interest and importance than "non-adaptive" (= fixed scale parameter) designs, since only adaptive designs have some hope of matching the asymptotically optimal behavior of adaptive RM when the true value of β is unknown. However, nonadaptive procedures are much easier to study than adaptive procedures, and insights gained from the examination of nonadaptive procedures may to some extent carry over to the much more complicated adaptive designs. Also, the simulations of Wu (1985) show that a well-chosen nonadaptive procedure will often outperform adaptive procedures for moderate sample sizes.

Section 2 will study several nonadaptive repeated MLE procedures by looking at their (normalized) efficient score functions. Comparison of the shapes of the efficient score functions of individual observations shows that the nonadaptive logit-MLE rule performs far better than the nonadaptive probit-MLE rule or the nonadaptive RM rule when the initial observations are taken far from L_p.

Section 3 states a consistency theorem for repeated MLE procedures.

2. Geometry of nonadaptive repeated MLE designs.

Let G be a c.d.f. with an everywhere positive density g, and for which $G(0) = p$. Consider the location model for F given by

$$F(x) = G(x - \alpha), \quad \alpha \in \mathbb{R}. \tag{2.1}$$

Since $G(0) = p$, we have $L_p = \alpha$ under the model. The log likelihood function for α based on data $\{(x_i,y_i)\}_{i=1}^n$ is

$$L^{(n)}(\alpha) = \sum_{i=1}^{n} [\log G(x_i-\alpha)]\, I\{y_i=1\} + [\log\{1-G(x_i-\alpha)\}]\, I\{y_i = 0\}. \tag{2.2}$$

The efficient score function is

$$\dot{L}^{(n)}(\alpha) = \sum_{i=1}^{n} \frac{-g(x_i-\alpha)}{G(x_i-\alpha)} I\{y_i=1\} + \frac{g(x_i-\alpha)}{1-G(x_i-\alpha)} I\{y_i=0\}. \tag{2.3}$$

If we let $\ell_i(\alpha)$ be the efficient score function for the i-th observation, then we can rewrite (2.3) as

$$\dot{L}^{(n)}(\alpha) = \sum_{i=1}^{n} \ell_i(\alpha). \tag{2.4}$$

Of course, the MLE $\hat{\alpha}_n$, if it exists, will be a root of

$$\dot{L}^{(n)}(\alpha) = 0. \tag{2.5}$$

It is desirable that $\dot{L}^{(n)}$ be monotone, so that (2.5) has at most one root. A necessary and sufficient condition for $\ell_i(\alpha)$ to be monotone when $y_i = 1$ (respectively, $y_i = 0$) is that log G (respectively, log (1 - G)) be concave, (c.f. Silvapulle (1981).) If these conditions hold, then it is easily verified that the MLE $\hat{\alpha}_n$ will exist and be unique precisely when

$$\sum_{1}^{n} y_i > 0 \text{ and } \sum_{1}^{n} (1-y_i) > 0. \tag{2.6}$$

One can start using repeated maximum likelihood estimation to choose x_{n+1} only after (2.6) holds. One can either use some other choice rule until some time after (2.6) is satisfied, or one may start with some "fake" initial data for which (2.6) holds and which one may feel reflects prior opinion, so that all "real" observations can be determined according to the MLE rule.

Let us now define a normalized score function. Write

$$r(t) = \frac{p(1-p)}{g(0)} \frac{g(-t)}{G(-t)}, \tag{2.7}$$

$$q(t) = -\frac{p(1-p)}{g(0)} \frac{g(-t)}{1-G(-t)}, \text{ and} \tag{2.8}$$

$$S^{(n)}(\alpha) = \sum_{i=1}^{n} r(\alpha-x_i)I\{y_i=1\} + q(\alpha-x_i)I\{y_i=0\}. \tag{2.9}$$

Then

$$S^{(n)}(\alpha) = -\frac{p(1-p)}{g(0)} \dot{L}^{(n)}(\alpha). \tag{2.10}$$

The MLE $\hat{\alpha}_n$ is of course also a root of

$$S^{(n)}(\alpha) = 0. \tag{2.11}$$

The normalization factor $p(1-p)/g(0)$ was chosen so that $r(0) = 1-p$ and $q(0) = -p$.

Let us examine what the r,q and $S^{(n)}$ functions look like when $p = 1/2$ and G is a logit or a probit c.d.f. The formulas are as follows.

$$\text{Logit: } G_L(t) = (1+e^{-t})^{-1},\ r_L(t) = (1-e^{-t})^{-1},\ q_L(t) = (1+e^{-t})^{-1} - 1.$$
$$\text{Probit: } G_p(t) = \Phi(t),\ r_p(t) = \frac{(2\pi)^{1/2}}{4} \frac{\phi(t)}{1-\Phi(t)},\ q_p = -\frac{(2\pi)^{1/2}}{4} \frac{\phi(t)}{\Phi(t)}.$$

The asymptotic distribution of $n^{1/2}(x_n-L_p)$ under the logit and probit MLE rules is determined by the derivative $r'(0) = q'(0)$. (See Sellke (1986),

Section 5.) If the above r and q functions are rescaled (by a rescaling of the t-axis) so that $r'(0) = q'(0) = 1$, then their graphs are as shown in Figures 1 and 2. Note that r and q are monotone in both cases. (The concavity of log G and log(1-G) is easy to verify analytically.) Figure 5 shows how the score function and the MLE are changed by a single observation under the logit model. The geometric aspects of how the score function and the MLE are updated under the probit model are similar.

If one computes the efficient score function corresponding to the model (1.5), one finds that a repeated MLE rule under this model is equivalent to choosing x_{n+1} to satisfy $S_R^{(n)}(x_{n+1}) = 0$, where

$$r_R(t) = 1 - p + bt, \tag{2.12}$$

$$q_R(t) = -p + bt, \tag{2.13}$$

and where $S_R^{(n)}$ is defined by (2.9). The graphs of r_R and q_R with $b = 1$ are shown in Figure 3. Since $S_R^{(n)}$ is a line of slope (nb) for which $S_R^{(n)}(x_n) = y_n - p$, it is easy to verify that the updating rule (1.4) holds.

The graphs in Figure 1, 2 and 3 can be used to examine how individual observations affect the value of the MLE. The effects of "outliers" are of particular interest. Suppose that $\{(x_i,y_i)\}_{i=1}^n$ have been observed. Let us consider the effect that a particular observation, say (x_1,y_1), has on the value of $\hat{\alpha}_n$ in these three models. In the logit and probit models, $r > 0$ and $q < 0$. Thus, for logit and probit, (x_1,y_1) will "push left" on $\hat{\alpha}_n$ if $y_1 = 1$, in the sense that the MLE would be larger if (x_1,y_1) were deleted from the data. Likewise, (x_1,y_1) will "push right" if $y_1 = 0$, regardless of the value of x_1. The size of the "push" will obviously depend on the value of x_1. If $y_1 = 1$ and $x_1 >> \hat{\alpha}_n$, then (x_1,y_1) will generally have only a tiny effect on $\hat{\alpha}_n$, since $r(\hat{\alpha}_n - x_1)$ will be very small. If $y_1 = 0$ and $x_1 << \hat{\alpha}_n$, the situation is similar. If $y_1 = 1$ and $x_1 << \hat{\alpha}_n$, then $r(\hat{\alpha}_n - x_1) \approx 1$ under the logit model, and $r(\hat{\alpha}_n - x_1) \approx (\pi/2)(\hat{\alpha}_n - x_1)$ under the probit model. Thus, the logit model

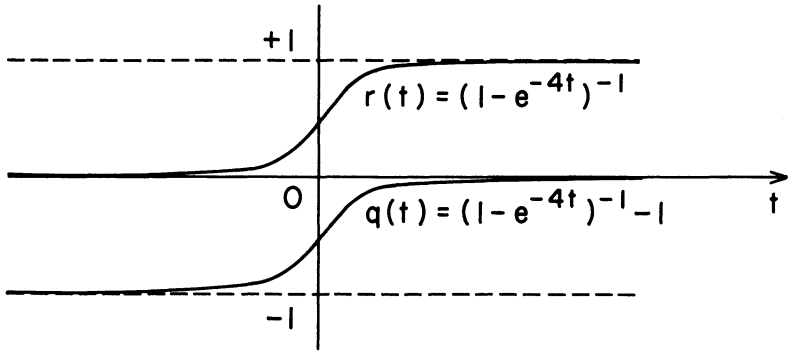

Figure 1. Logit

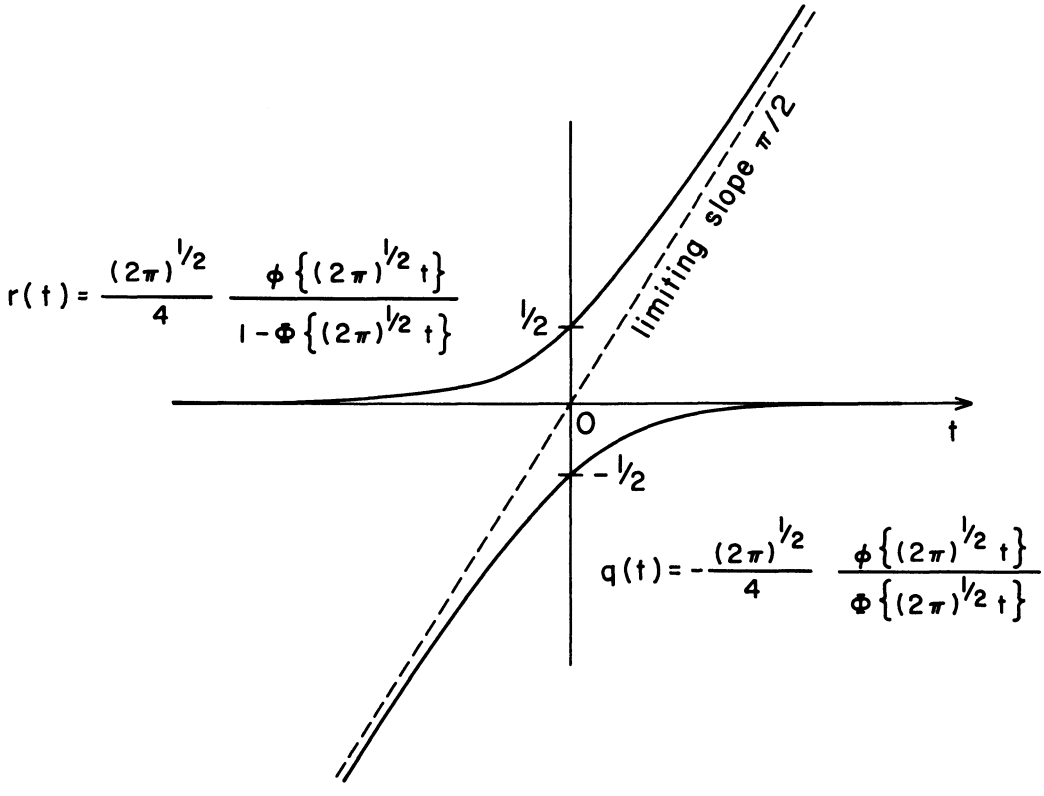

Figure 2. Probit

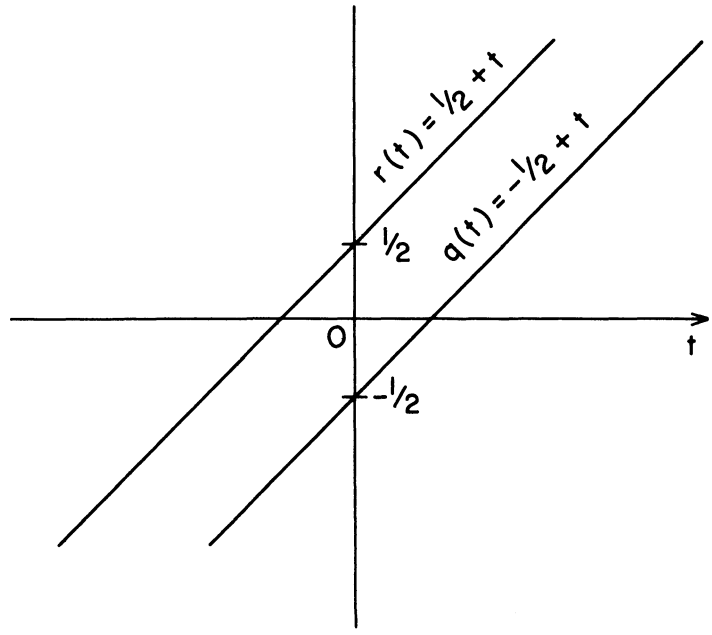

Figure 3. RM (1)

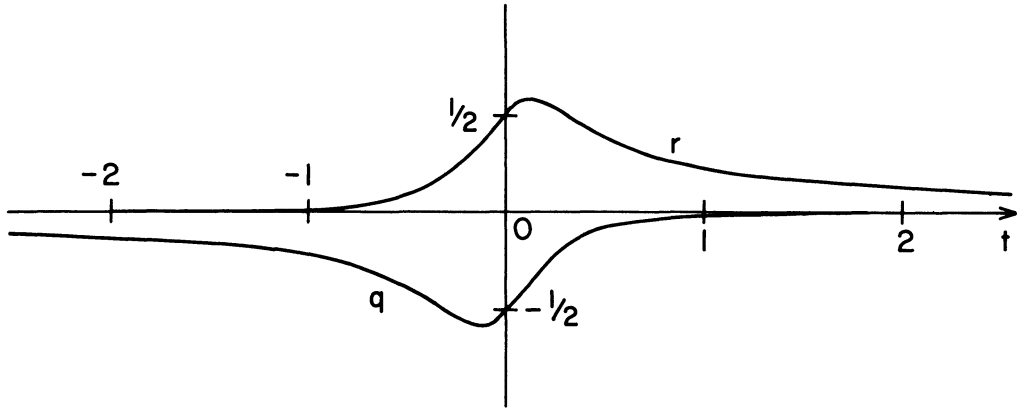

$$r(t) = \frac{1}{\left\{1+(\pi t)^2\right\}\left\{2-(4/\pi)\arctan(\pi t)\right\}}$$

$$q(t) = -r(-t)$$

Figure 4. Cauchy

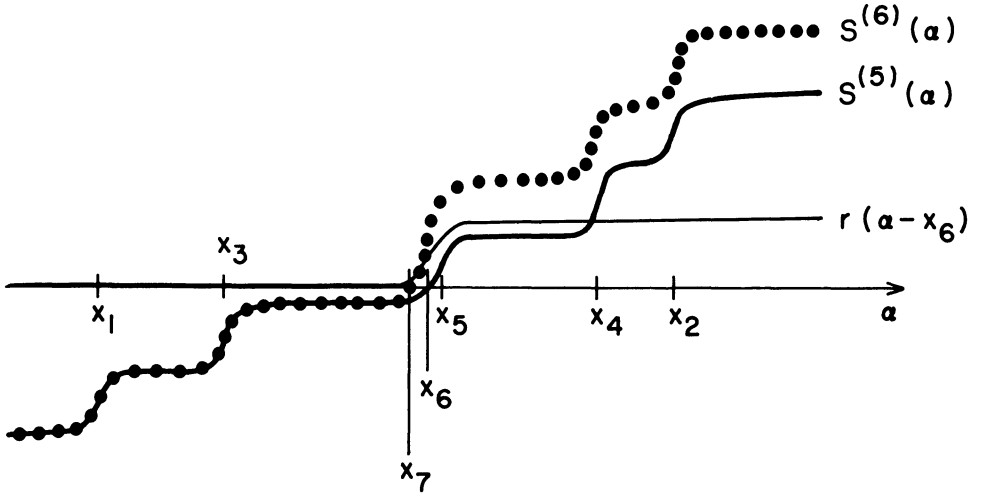

x_6 satisfies $S^{(5)}(x_6) = 0$.

$S^{(6)}(\alpha) = S^{(5)}(\alpha) + r(\alpha - x_6)$, since $y_6 = 1$.

x_7 satisfies $S^{(6)}(x_7) = 0$.

Figure 5. The Logit-MLE rule

is "robust" against outliers in the sense that the presence or absence of a single observation, no matter how extreme, will generally have only a small effect on the value of $\hat{\alpha}_n$ (at least if Σy_i and $\Sigma(1 - y_i)$ are both not small.) The probit model is not robust in this sense: if x_1 is allowed to increase to $+\infty$, with $y_1 = 0$ and $\{(x_i,y_i)\}_{i=2}^n$ held fixed, then $\hat{\alpha}_n$ will also increase to $+\infty$.

The RM(b) rule can produce rather perverse effects in the quantal response setting. If $x_1 > \hat{\alpha}_n + (2b)^{-1}$, then the (x_1,y_1) observation "pushes right" regardless of the value of y_1. If $x_1 < \hat{\alpha}_n - (2b)^{-1}$, then (x_1,y_1) "pushes left." The RM(b) rule is also not robust in the sense used above.

The robustness comparisons of the previous paragraphs are obviously very similar to the standard approach to robustness for M-estimators involving considertion of the influence function. See for example, Huber (1981).

It may be of interest to consider the degree to which a single outlier observation retards the movement of x_n to L_p. Suppose, therefore, that we start with $(x_1,y_1) = (0,1)$ and $(x_2,y_2) = (0,0)$, and that the score function rules of Figure 1, 2 and 3 are used thereafter. Suppose further that $y_i = 0$ for all subsequent observations. Under each of the three rules, $x_3 = 0$ and $x_{n+1} > x_n$ for $n \geq 3$. If $K > 0$, how long will it take x_n to reach K? Under the logit-MLE rule of Figure 1, the size of the steps $(x_{n+1} - x_n)$ rapidly converges down to .560. The first three step sizes are .693, .600 and .571. Thus, it takes about $(.560)^{-1}K = (1.78)K$ steps for x_n to exceed K. Under the probit-MLE rule of Figure 2, it takes

$$\frac{\pi K^2}{16 \log 2}\{1 \pm o(1)\}$$

steps to exceed K, where $\pi K^2/(16 \log 2)$ is probably an underestimate. For the RM(1) rule, for $n \geq 2$,

$$x_{n+1} = (1/2) \sum_{i=3}^{n} (i)^{-1} = (1/2)\{-3/2 + \log(n) + C - (2n)^{-1} + O(n^{-2})\}.$$

where C = .5772 is Euler's constant. (See Apostol (1969), p.618, Exercise 2).

Thus, it takes about

$$\exp\{2K + 3/2 - C + (1/2)\exp(-2K - 3/2 + C)\} - 1$$

additional observations after the first 2 to exceed K. For K = 10, the above estimates for the numbers of additional observations needed for x_n to reach K are 18, 28 and 1,220,814,797 for logit-MLE, probit-MLE,and RM(1), respectively. (The 18 is correct, the 1,220,814,797 is close to being correct, and the 28 is probably much too small.)

What happens if one uses a Cauchy c.d.f. in (2.6)? The resulting r and q functions are shown in Figure 4. Repeated maximum likelihood estimation under this model would have the perhaps desirable property that, the more extreme an outlier, the smaller its influence on $\hat{\alpha}_n$. However, the score function would not necessarily have a unique root, and Proposition 1 of Section 3 does not apply to prove consistency. The author conjectures that this repeated maximum likelihood rule would still be consistent under conditions (1.3a,b,c).

Another way of comparing the robustness properties of the score function rules of Figures 1 through 4 is to consider how fast x_n grows under these rules if the y_n's are all i.i.d. Bernoulli (p_o), with $0 < p_o < p = 1/2$. It can be shown that x_n grows roughly like a multiple of log n in all cases, but the coefficient of log n would decrease in the order of Cauchy, logit, probit, RM(1).

If $p \neq 1/2$, then the graphs of r and q are obtained from the $p = 1/2$ graphs by a horizontal location shift and a vertical rescaling. (Achieving $r'(0) = q'(0) = 1$ would require another horizontal rescaling.) The conclusions reached above concerning robustness and the effects of outliers are generally valid when $p \neq 1/2$.

To summarize, the somewhat qualitative comparisons above lead the author to favor the nonadaptive logit-MLE rule over the competitors. The "bounded influence" property of the logit model would cause x_n to move more quickly to the neighborhood of L_p if one has started far away from L_p than would

the probit or RM(b) rules. The RM(b) rules are clearly the easiest from a computational point of view, but the logit-MLE rule would also be easy to implement. The probit-MLE rule would be slightly more cumbersome, since it involves the normal c.d.f. Even if the Cauchy-MLE rule turns out to be consistent, implementation of this rule would have to contend with the possibility of multiple roots of the score function.

3. Consistency.

It is well known that the RM(b) rules are consistent under conditions (1.3a,b,c). (See, for example, Robbins and Siegmund (1971) for a very elegant proof.) Are other repeated MLE rules also consistent under such weak conditions? One strategy of proof which is suggested by the geometry of score functions is to attempt to show that the relationship of x_{n+1} to x_n for repeated MLE rules is very similar to formula (1.4). If, for large n, $S^{(n)}$ is "approximately linear" near x_n with slope $a_n =: \frac{d}{d\alpha} S^{(n)}_{(\alpha)}\Big|_{\alpha=x_n}$, then

$$x_{n+1} \approx x_n - (a_n)^{-1}(y_n - p). \tag{3.1}$$

Except for the approximation (3.1) has the form of a general Robbins-Monro rule. Furthermore, one would expect a_n to be on the order of n, so that $\Sigma (a_n)^{-1} = \infty$ and $\Sigma(a_n)^{-2} < \infty$. One can in fact apply the non-negative almost supermartingale convergence theorem of Robbins and Siegmund (1971), mimicking their proof of consistency of RM(b) rules, to the approximation (1.3) to establish consistency of the nonadaptive logit-MLE rule. (The proof is not particularly elegant because of various technical distractions.) The argument also works for the nonadaptive probit-MLE once the boundedness of the x_n sequence has been demonstrated.

A different and more general proof of consistency is given in Sellke (1986). Wu (1985a,b) has shown that x_n cannot converge to an incorrect value. In terms of the score function $S^{(n)}$, his argument is as follows. If $x_n \to x^* > L_p$, then by (1.3c) and the continuity of r and q at zero, it follows

that $S^{(n)}$ diverges to $+\infty$ uniformly in a small interval $(x^* - \delta, x^* + \delta)$ about x^*. But this and $S^{(n)}(x_{n+1}) = 0$ imply that $x_n \notin (x^* - \delta, x^* + \delta)$ for sufficiently large n, contradicting $x_n \to x^*$. Sellke (1986) completes the proof of consistency by showing that x_n must converge. The argument is based on an upcrossing inequaltiy. Let $L_p < a < b$ be given, and suppose that $x_m < a$. The upcrossing inequality gives a bound on the probability that $\sup_{n \geq m} x_n > b$ in terms of the difference

$$S^{(m)}(b) - S^{(m)}(a).$$

It follows that x_n cannot cross the interval [a,b] infinitely many times.

The following proposition is a special case of Theorem 1 of Sellke (1986).

PROPOSITION 1. Suppose that F satisfies (1.3a,b,c). Suppose that G is a c.d.f. for which log G and log (1-G) are concave, and for which the second derivatives at zero

$$\frac{d^2}{dx^2} \log G(x)\Big|_{x=0}, \text{ and } \frac{d^2}{dx^2} \log \{1 - G(x)\}\Big|_{x=0}$$

exist and are strictly negative. Suppose also that $G(0) = p$. If x_n is chosen, for all sufficiently large n, to equal the MLE of L_p under the location model (2.1), then x_n converges to L_p almost surely.

Acknowledgement. The author thanks Jeff Wu for a number of helpful comments.

REFERENCES

Apostol, Tom M. (1969). Calculus, Volume II, 2nd ed. Waltham, Mass., Xerox.

Huber, P.J. (1981). Robust Statistics. New York, Wiley.

Lai, T.L. and Robbins, H. (1979). Adaptive design and stochastic approximation. Ann. Statist. 7 1196-1221.

Lai, T.L. and Robbins, H. (1982). Iterated least squares in multiperiod control. Adv. Appl. Math. 3 50-73.

Robbins, H. and Siegmund, D. (1971). A convergence theorem for nonnegative almost supermartingales and some applications. In Optimizing Methods in Statistics. (J.S. Rustagi, ed.) 233-257. Academic Press, New York.

Sellke, T. (1986). Consistency of maximum likelihood recursion in stochastic approximation. Purdue University Technical Report.

Wu, C.F.J. (1985). Efficient sequential designs with binary data. J. Amer. Statist. Assoc. **80**, 974-984.

Wu, C.F.J. (1986). Maximum likelihood recursion and stochastic approximation in sequential designs. This volume.