

OPTimal Strategies for Seeking Information: Models for Statistics, Choice Reaction Times, and Human Information Processing¹

WARD EDWARDS

University of Michigan, Ann Arbor, Michigan

and

System Development Corporation, Santa Monica, California

Models for optional stopping in statistics are also normative models for tasks in which subjects may purchase risk-reducing information before making a decision. A Bayesian model for optional stopping for the two-hypothesis continuous case is developed; it takes explicit account of cost of information, values of the possible outcomes of the final decision, and prior probabilities of the hypotheses.

A nonparametric model for choice reaction times is derived. It makes strong predictions about times and errors; only one quantity in it is not directly observable.

A second example uses the model to design and predict results of a binomial information-purchase experiment.

If information needed for a decision is costly, it may not be optimal to decide in advance how much information to collect. Instead, it may be preferable to decide after each observation whether or not to seek the next one. This strategy permits exploitation of good luck; if a small set of observations happens to be sufficiently conclusive, there is no need to buy more. The stopping rule should be chosen to maximize the expected value of the process, considering both cost of information and costs and payoffs that may be incurred as a result of the terminal decision.

Choice of a stopping rule in sequences of information purchases is a well-known problem of classical statistics. More importantly for psychologists, it is a model for a number of interesting tasks. In a choice reaction task, for example, a subject must decide on the basis of observations which of two (or more) stimuli has been presented to him and make the appropriate response as fast as he can. Instructions to be both

¹ This research was supported in part by Contract AF 19(604)-7393 monitored by the Decision Sciences Laboratory, Directorate of Computers, Electronics Systems Division, Air Force Systems Command, in part by the Advanced Research Projects Agency under Contract SD-97 with the System Development Corporation, and in part by Grant AF AFOSR-192-63, monitored by the Air Force Office of Scientific Research of the Air Force Office of Aerospace Research.

fast and right imply a cost for observing and a payoff for correct decision; in some experiments, those costs and payoffs are explicit. For a second example, direct experiments on information purchase as a preliminary to decision are a part of several current programs of research on human processing of probabilistic information. And the seeking of costly information is a natural function of every real-world system for making important decisions. Although the formal development presented here is statistical, I expect the more important applications will not be, and in fact neither of the applications to be used as examples in this paper are from statistics.

The names used within statistics for this class of problems are optional stopping and sequential decision-making. Wald formulated the problem explicitly in 1947, and Blackwell and Girshick (1954) and many others have contributed to it from the viewpoint of classical statistics. Blackwell and Girshick also consider it from a Bayesian point of view, as do Schleifer (1961) and others. The formal development presented in this paper is elementary, both mathematically and statistically. Nevertheless, I have not been able to find any explicit presentation of it in the statistical books or journals. The problem is formulated, in many far more general forms, in Blackwell and Girshick (1954); one exercise (10.5.5, p. 278) invites the reader to find for himself the solution for the explicitly soluble case of interest in this paper. After a draft of this paper had been circulated for comment, my attention was called to an unpublished Ph. D. thesis by Schleifer (1961). Schleifer presents essentially the same derivation for the case of two normal distributions with equal variance only, and does not discuss the more general problem. Much of his analysis relates his problem to Bayesian analysis of fixed-sample-size experiments. The derivation presented in the first part of this paper is, then, at most a minor extension of information available from Blackwell and Girshick and Schleifer. Nevertheless, it is worth presenting here, both as a basis for the psychological applications to be discussed later in the paper and because the other presentations are fairly inaccessible, because of mathematical difficulty or lack of circulation.

This paper adopts the Bayesian view that probabilities are orderly opinions about uncertain events and that statistics is a set of procedures for revising those opinions in the light of relevant information. (For extensive discussions of this approach to statistics, see Edwards, Lindman, and Savage (1963) and the books and papers they cite.) This paper first presents and solves the formal problem. It next considers two nonstatistical applications of the solution. One uses the model as a descriptive model for choice reaction times. The second uses the model as a tool for the design of information-processing experiments and as a model for the resulting data.

A major difficulty in sequential analysis and study of stopping rules is that observations are discrete. In a Bayesian formulation this means that the posterior odds cutoffs that define the optimal stopping rule cannot in general be precisely reached. Much of the mathematical difficulty of the problem results from this discreteness. This

paper ignores it and instead treats observations as though they were continuous, so that stopping points can be reached precisely. Consequently, its computational formulas should be thought of as approximations, usually excellent ones.

THE FORMAL PROBLEM AND ITS SOLUTION

This paper is about the problem of simple dichotomy. Observations are generated by a well-defined data-generating process that is known to be in one or the other of two states. Hypothesis H_1 asserts that the process is in state θ_1 ; hypothesis H_0 asserts that the process is in state θ_0 .

A Bayesian decision maker will base his decisions on costs, payoffs, and his current probabilities for the two hypotheses. For dichotomies, it is convenient to consider these probabilities in the form of odds. Bayes' theorem says

$$p(H_1 | D) = k p(D | H_1) p(H_1), \quad (1)$$

and

$$p(H_0 | D) = k p(D | H_0) p(H_0), \quad (2)$$

where $p(H_1 | D)$ is the posterior probability of H_1 on the basis of the data, $p(H_1)$ is its prior probability, $p(D | H_1)$ is the probability that the data would have been observed if H_1 were true, k is a constant of proportionality that the likelihood principle of Bayesian statistics (see Edwards, Lindman, and Savage, 1963) tells us is irrelevant to interpretation of data, and the quantities in Eq. 2 have analogous interpretations. Division of Eq. 1 by Eq. 2 yields

$$\Omega(H_1/H_0 | D) = L(D | H_1/H_0) \Omega(H_1/H_0) \quad (3)$$

or, in simpler notation to be used throughout the paper,

$$\Omega_1 = L \Omega_0. \quad (4)$$

The symbol $\Omega(H_1/H_0 | D) \equiv \Omega_1$ stands for the posterior odds in favor of H_1 over H_0 on the basis of datum D ; $\Omega(H_1/H_0) \equiv \Omega_0$ is the prior odds. The word odds (here used as a singular noun) refers to exactly the concept used at the race track. The symbol $L(D | H_1/H_0) \equiv L$ is the familiar likelihood ratio of classical and Bayesian statistics. So Eq. 4 is simply a convenient form of Bayes' theorem for problems of simple dichotomy.

In Bayesian statistics, or at least in the version adopted here, a probability, or an odds, characterizes the opinion of some particular person. That person, called *you* in this paper, is assumed to be remarkably consistent and rational, to have unlimited

costless computational abilities, and to have enough money to be able to pay any losses he might incur.

Continuous sequential sampling. Your prior odds may so favor H_1 or H_0 that, in view of the economic constants of the problem, you will act as though one or the other is true without sampling. Or you may buy an observation or batch of observations; in what follows, batches of observations will be treated as single observations. After observing, you face the same three possible acts as before. Since observing is expensive, some value of the posterior odds will be just small enough so that you will decide immediately for the act appropriate to H_0 , and some value will be just large enough so that you will act appropriately to H_1 . Call these cutoff values Ω_* and Ω^* respectively. If your current odds lies between them, you will continue observing.

The choice of cutoff odds will reflect both cost and payoff considerations, but will be independent of prior probabilities and will not be influenced by n , the measure of how much observing you have already done on this trial and so of how much you have already spent. Money spent observing is down the drain except insofar as the observations purchased have moved the odds; you always consider your immediate situation as though you had not collected any information previously. The point is counter-intuitive, but can be made less so by pointing out that the same thing is true of poker. The amount of money you have already contributed to the pot is entirely irrelevant to whether you should stay or fold; the only relevant considerations are the anticipated value of the pot, your estimate of your chances of winning it, and how much more it may cost you if you lose. None of these arguments apply, in poker or in this paper, if your resources are so limited that you may run out of money or time.

Finding optimal cutoffs. Equation 4 implies the following two equations:

$$\Omega^* = A\Omega_0 \quad (5)$$

and

$$\Omega_* = B\Omega_0 \quad (6)$$

where A is the likelihood ratio just large enough to transform your prior odds into the upper posterior odds cutoff, and B is the likelihood ratio just small enough to transform your prior odds into the lower posterior odds cutoff. The conditional probability that you will accept H_0 when H_1 is correct will be called β . Similarly, the conditional probability that you will accept H_1 when H_0 is correct will be called α . α and β are the operating characteristics of the test, in classical statistical language. It is convenient to compile in a set of equations the relations among the quantities so far discussed.

$$A = \frac{1 - \beta}{\alpha} = \frac{\Omega^*}{\Omega_0} \quad (7)$$

$$B = \frac{\beta}{1 - \alpha} = \frac{\Omega_*}{\Omega_0} \quad (8)$$

$$\alpha = \frac{1 - B}{A - B} = \frac{\Omega_0 - \Omega_*}{\Omega^* - \Omega_*} \tag{9}$$

$$\beta = \frac{B(A - 1)}{A - B} = \frac{\Omega_*(\Omega^* - \Omega_0)}{(\Omega^* - \Omega_*)\Omega_0} \tag{10}$$

Expected sample size. Wald (1947, Eq. 3:57, p. 53) gives a general expression for the expected sample size as a function of A and B . It is:

$$E_\theta(n) = \frac{L(\theta) \ln B + [1 - L(\theta)] \ln A}{E_\theta(z)} \tag{11}$$

$E_\theta(n)$ is the expected sample size as a function of θ , the true mean (or other parameter) appropriate to the state of the world which in fact obtains. $L(\theta)$, the operating characteristic function of the procedure, is the probability that the sequential process will terminate with the acceptance of H_0 when θ is the true parameter. Thus it is $1 - \alpha$ when H_0 is correct and β when H_1 is correct. The symbol z stands for the logarithm of the likelihood ratio for one observation. $E_\theta(z)$ is, then, the expected value of z when θ is the true parameter. Only through $E_\theta(z)$ do the distributional characteristics of the particular problem enter into the selection of Ω_* and Ω^* . But Ω_* and Ω^* can vary without affecting $E_\theta(z)$ in any way; that is, $E_\theta(z)$ is completely determined by the nature of the data-generating process and is not a function of Ω_* and Ω^* . It is therefore appropriate to treat $E_\theta(z)$ as a constant in determining optimal values of Ω_* and Ω^* . Later, the paper examines specific cases and so gives specific definitions of $E_\theta(z)$.

Substitution from Eqs. 7 and 8 into 11 produces:

$$E_{\theta_1}(n) = \frac{1}{E_{\theta_1}(z)} \left[\beta \ln \frac{\beta}{1 - \alpha} + (1 - \beta) \ln \frac{1 - \beta}{\alpha} \right], \tag{12}$$

$$E_{\theta_0}(n) = \frac{1}{E_{\theta_0}(z)} \left[(1 - \alpha) \ln \frac{\beta}{1 - \alpha} + \alpha \ln \frac{1 - \beta}{\alpha} \right]. \tag{13}$$

Selection of α and β . Assume that the payoff matrix for correct decisions and errors is the following:

		State of the world	
		H_0 true	H_1 true
Course of action	H_0	a	b
is appropriate to	H_1	c	d

where a , b , c , and d are amounts of money (assuming your utility function for money to be linear over the relevant range). Now the strategic aspects of the situation are unchanged (assuming your supply of money is large relative to a , b , c , and d) if you

receive a (possibly negative) present of $-a$ dollars whenever H_0 is true and $-d$ dollars whenever H_1 is true, regardless of your decisions; these presents are simply windfall profits or losses, utterly irrelevant to your selection of strategy. Thus the following payoff matrix is strategically equivalent to the previous one:

		State of the world	
		H_0 true	H_1 true
Course of action	H_0	0	$b - d$ (= V_{01})
is appropriate to	H_1	$c - a$ (= V_{10})	0.

For simplicity, the following development uses this second form of the payoff matrix. For convenience, define $V_{10} = c - a$ and $V_{01} = b - d$. The probability of obtaining V_{10} is the prior probability of H_0 , $p(H_0)$, times α . Similarly, the probability of obtaining V_{01} is $p(H_1)$, times β . It is assumed that V_{10} and V_{01} are negative, that is, they are losses, in the formal decision theory sense of that word.

It is convenient to assume that the cost of n observations is a linear function of n , say $t + un$. Like V_{10} and V_{01} , u is assumed negative. Since t is inevitable, it plays no role in strategy selection; the following development simply omits it. If it were carried along, it would drop out of the equations as a result of differentiation. Non-linear costs of observing complicate things, usually intolerably.

Now it is possible to write an equation, putting together the parts of the expression for expected cost. It is

$$E(V) = p(H_0)\alpha V_{10} + p(H_1)\beta V_{01} + up(H_0)E_{\theta_0}(n) + up(H_1)E_{\theta_1}(n). \tag{14}$$

(Note that $p(H_0)$ and $p(H_1)$ are the prior, not the posterior, probabilities of the hypotheses.) To find the optimal values of α and β , it is appropriate to differentiate Eq. 14 with respect to α and β , set the resulting equations equal to zero, and solve them as a system of simultaneous equations in two unknowns. The operations are straightforward and conventional; they need not be exhibited here.

It is convenient to define

$$F = \frac{V_{10}E_{\theta_1}(z)}{u},$$

$$G = \frac{V_{01}E_{\theta_1}(z)}{u},$$

and

$$H = -\frac{E_{\theta_1}(z)}{E_{\theta_0}(z)}.$$

Note that, since V_{10} , V_{01} , and u are all negative, F and G are required to be positive. Since both V_{10} and u are in dollars (or equivalent measures of cost) and z is the logarithm of a ratio, F and G are dimensionless constants. Using these aggregated constants, the final forms of the equations for the posterior odds cutoffs become

$$F - H \ln \frac{\Omega^*}{\Omega_*} + \Omega_* - \Omega^* = 0, \quad (15)$$

$$G - H \left(\frac{1}{\Omega_*} - \frac{1}{\Omega^*} \right) - \ln \frac{\Omega^*}{\Omega_*} = 0. \quad (16)$$

Note that prior odds and prior probability do not appear in Eqs. 15 and 16. They dropped out in the course of the simplifications following differentiation. As asserted earlier, your prior odds are irrelevant to your choice of posterior odds cutoffs.

Equations 15 and 16 are transcendental; explicit solutions are not to be anticipated. Nevertheless, they are reasonably simple as they stand. Trial and error methods will often produce acceptable solutions.²

Simplifications resulting from symmetry. The constant H is, of course, designed to be 1, and will be so in almost all applications. It will be so, for example, in the case of two normal distributions differing only in mean, the case of two binomial distributions whose parameters are located symmetrically around 0.5, and some others. It will often be true also that $V_{10} = V_{01}$, so that $F = G$. In such cases, symmetric with respect to both discriminability and cost-payoff considerations, $\Omega_* = 1/\Omega^*$ and Eqs. 15 and 16 reduce to a single equation

$$F - 2 \ln \Omega^* + \frac{1}{\Omega^*} - \Omega^* = 0. \quad (17)$$

Equation 17, though still transcendental, is so simple that trial-and-error iteration methods will quickly lead to solutions for any values of interest. If Ω^* might plausibly be 100 or so, its reciprocal is negligible, and $2 \ln \Omega^* \simeq 0.869 \log_{10} \Omega^*$, though unfortunately not negligible, is small enough relative to it so that it is appropriate to think of Ω^* as equal to a little less than F . If more refinement is needed, a few trial values will quickly pin it down as precisely as the circumstances require. Table 1 contains values of Ω^* for some values of F between 0.5 and 1000.

Since Eq. 17 does not involve Ω_0 , prior odds need not be symmetric for it to apply. But Eqs. 12 and 13 for the expected numbers of observations, the other important

² For extensive work with asymmetric cases, a table of solutions to Eqs. 15 and 16 is indispensable. A table of such solutions, using as values of F and G all possible pairs of the values of F used in Table 1 of this paper, has been deposited as Document No. 8342 with the American Documentation Institute Auxiliary Publications Project, Photoduplication Service, Library of Congress, Washington 25, D. C. A copy may be secured by citing the Document number and by remitting \$6.25 for photoprints, or \$2.50 for 35 mm microfilm. Advance payment is required. Make checks or money orders payable to: Chief, Photoduplication Service, Library of Congress.

TABLE 1
 Ω^* FOR SELECTED VALUES OF F IN THE SYMMETRIC CASE

F	Ω^*	F	Ω^*
0.5	1.13296	20.0	14.69328
1.0	1.28238	25.0	19.14784
1.5	1.44878	30.0	23.71035
2.0	1.63244	40.0	33.03513
2.5	1.83327	50.0	42.52341
3.0	2.05096	60.0	52.11238
3.5	2.28495	70.0	61.76937
4.0	2.53453	80.0	71.47529
4.5	2.79886	90.0	81.21804
5.0	3.07705	100.0	90.98950
5.5	3.36816	120.0	110.59725
6.0	3.67129	140.0	130.26848
6.5	3.98555	160.0	149.98559
7.0	4.31009	180.0	169.73739
7.5	4.64412	200.0	189.51633
8.0	4.98690	300.0	288.67288
8.5	5.33774	400.0	388.08015
9.0	5.69602	500.0	487.62297
9.5	6.06118	600.0	587.25080
9.96620	6.40737	800.0	786.66566
10.0	6.43268	1000.0	986.21327
15.0	10.41044		

working equations of the model, do depend on prior odds, and important simplifications result if all three symmetry conditions apply. In presenting them, it is convenient to define

$$\frac{p^*}{1-p^*} = \frac{p^*}{q^*} = \Omega^* \quad \text{and} \quad \frac{p_*}{1-p_*} = \frac{p_*}{q_*} = \Omega_*.$$

If $E_{\theta_0}(z) = -E_{\theta_1}(z)$ so that $H = 1$, $V_{10} = V_{01}$ so that $F = G$, and $\Omega_0 = 1$, it follows that $\Omega_* = 1/\Omega^*$, $B = 1/A$, $p^* = q_*$, $q^* = p_*$, $\alpha = \beta$, and $E_{\theta_1}(n) = E_{\theta_0}(n)$. For this utterly symmetric case, it is convenient to define

$$\frac{V_{10}}{u} = \frac{V_{01}}{u} = V, \quad E_{\theta_1}(z) = -E_{\theta_0}(z) = Z, \quad \text{and} \quad E_{\theta_1}(n) = E_{\theta_0}(n) = N.$$

In that notation, Eq. 17 may be written

$$VZ - 2 \ln \frac{p^*}{q^*} + \frac{1}{p^*} - \frac{1}{q^*} = 0, \quad (18)$$

and Eqs. 12 and 13 simplify to

$$NZ = (p^* - q^*) \ln \frac{p^*}{q^*}. \quad (19)$$

The expected sample size function. Consider several situations all having the same cost-payoff structure and the same prior probabilities for the two hypotheses, but differing from one another in $E_\theta(z)$. As the impact of each observation on your opinion increases, other things remaining the same, should you buy more or less information? Intuition is inconsistent. On the one hand, it should take fewer observations to reach a specified level of certainty, so you might expect to purchase fewer observations. On the other hand, observations are an increasingly good buy, so to speak, so you might choose to buy more of them.

Intuition can be clarified by considering extremes. One extreme arises when observations are worthless, that is, $E_\theta(z) = 0$ for both values of θ . Clearly if observations are worthless, you should buy none. The other extreme arises when each observation is utterly definitive, that is, $E_\theta(z)$ approaches plus and minus infinity, with the sign depending on which hypothesis is true. In that case, you should buy exactly one observation (or, in the continuous case being considered here, the number of observations you would buy should approach 0 as $E_\theta(z)$ approaches infinity). Only for intermediate values of $E_\theta(z)$ should you buy larger amounts of information. That suggests that the functions relating $E_\theta(n)$ to $E_\theta(z)$ should each pass through the origin, rise to a maximum, and thereafter fall and approach zero again asymptotically. In general, there will be many such functions. But in the completely symmetric case, it is possible to write the function in a way which makes it unique. Divide Eq. 19 by Eq. 18:

$$\frac{N}{V} = \frac{(p^* - q^*) \ln(p^*/q^*)}{2 \ln(p^*/q^*) - 1/p^* + 1/q^*}. \quad (20)$$

TABLE 2
NZ, VZ, AND N/V FOR SELECTED VALUES OF p^* .

p^*	NZ	VZ	N/V	p^*	NZ	VZ	N/V
0.55	0.020067	0.80538	0.0249162	0.86495	1.355444	9.96258	0.13605346
0.60	0.081093	1.64426	0.0493187	0.865	1.355942	9.96625	0.13605341
0.65	0.185712	2.55676	0.0726356	0.866	1.365945	10.04004	0.1360497
0.70	0.338879	3.59916	0.0941551	0.87	1.406709	10.34480	0.1359823
0.75	0.549306	4.86389	0.1129355	0.88	1.514247	11.18183	0.1354203
0.80	0.831777	6.52259	0.1275225	0.90	1.757780	13.28334	0.1323297
0.85	1.214221	8.95940	0.1355248	0.95	2.649995	24.83625	0.1066987
0.86	1.307009	9.61065	0.1359959	0.97	3.267533	39.25460	0.0832395
0.864	1.346012	9.89337	0.1360519	0.99	4.503218	108.18014	0.0416270
0.86490	1.354946	9.95892	0.13605349				

Since V is a constant in Eq. 20, the maximum value of N will occur where N/V is a maximum. Table 2 presents calculations of the quantities involved for interesting values of p^* , and Fig. 1 presents a plot of N/V as a function of VZ .

Inspection of the table and figure makes it clear that the maximum value of N/V occurs when the cutoff posterior probability is 0.8649, which it will be when VZ is




FIG. 1. N/V as a function of VZ . The indicated points are listed in Table 1. Note that VZ is on a logarithmic scale. This graph is a source of quick approximations; exact work requires more precision.

$9.95892 \approx 10$. This is smaller than the typical VZ of statistical practice, but well within the range of VZ 's likely to be encountered in psychological experiments on information seeking, choice reaction time, and the like. An application of this function to psychological models that consider man as a degrader of data is presented later.

APPLICATIONS

A nonparametric and a normal application: two models for choice reaction times. Recently Stone (1960) has used the classical version of the statistical treatment of optional stopping as a model for choice reaction times. In a choice reaction time experiment, the subject is shown one or more stimuli, and must as soon as possible make the response appropriate to the stimulus presented. He is instructed both to minimize errors and to respond as fast as possible. It is a reasonable idealization of such an experiment to think of the time between the stimulus and the response as being sampling time, in which the subject is gathering information to choose among the several hypotheses about what response is appropriate this time. (Or perhaps it might be better to think of some part of that time as simple reaction time and the rest as sampling time; both models lead to the same equations, though their use may differ depending on which model is being considered.) Stone's treatment is, and this one will be, confined to the case of two simple hypotheses, since the mathematical difficulties of extending the idea to more than two remain an unsolved problem of mathematical statistics.

A Bayesian version of the same kind of model provides explicit and formal treatment of the costs, payoffs, and prior probabilities used in such experiments, in contrast to Stone's informal treatment of these topics. Choice reaction time experiments using explicit costs and payoffs are now becoming more frequent (e.g. Fitts, 1964; Fitts, Peterson, and Wolpe, 1963).

The subject looks at two lights. Between trials, both are off. At the beginning of each trial, one is turned on; the subject must press the button appropriate to it. The probability that each will turn on is fixed and known to the subject. The experimenter has specified and taught the subject about a cost-payoff structure that rewards him for correct responses, punishes him for incorrect ones, and makes it important to be quick rather than slow by charging him for time at a cost that is linearly related to the time used. The subject is assumed to maximize expected financial return.

In this model, H is assumed to be 1. $E_{\theta}(z)$, the ease of discrimination between the two lights, is an intervening variable. It is convenient to call the cost of a second of observing k , and to call the duration of an observation in seconds δ ; by definition, $u = k\delta$. Of course k is specified by the experimenter, but δ and u are unknowable intervening variables. Now, define a new parameter

$$\lambda = \frac{E_{\theta_1}(z)}{\delta} = E_{\theta_1}(z) \left(\frac{k}{u} \right).$$

In the new notation

$$F = \frac{V_{10}}{k} \lambda$$

$$G = \frac{V_{01}}{k} \lambda.$$

F and G are functions of known experimental arrangements and of a single intervening variable λ .

The number of observations, n , multiplied by the duration of an observation in seconds, δ , should give the reaction time in seconds, R , for any single trial, assuming that each observation takes a fixed length of time. So, taking expectations,

$$E_{\theta}(n) = E_{\theta} \left(\frac{R}{\delta} \right) = \frac{1}{\delta} E_{\theta}(R) = \frac{\lambda}{E_{\theta_1}(z)} E_{\theta}(R)$$

Eqs. 12 and 13 may be rewritten

$$E_{\theta_1}(R) = \frac{1}{\lambda} \left[\beta \ln \frac{\beta}{1-\alpha} + (1-\beta) \ln \frac{1-\beta}{\alpha} \right] \quad (21)$$

$$E_{\theta_0}(R) = \frac{1}{\lambda} \left[-(1-\alpha) \ln \frac{\beta}{1-\alpha} - \alpha \ln \frac{1-\beta}{\alpha} \right]. \quad (22)$$

The fact that λ , the single intervening variable of the model, is concerned with both duration and evidential value of an observation means that pure discriminability cannot be disentangled from rate of taking observations in the model; there is no hope, for example, of estimating δ separately and so determining how long it takes the organism to make a single observation. In fact, the notion of a single observation is ill-defined in such continuous-observation models; it is really definable only as the smallest unit of observation large enough so that successive observations may be treated as independent—a rather mystical notion.

The model makes very strong predictions. Once λ is estimated, either from reaction times or from error data, all other dependent variables are precisely predicted. (It is possible, though somewhat complex, to predict standard deviations of reaction times from the same model.) It also makes a number of quite specific predictions that do not depend at all on λ . Any manipulations of costs and payoffs that leave F and G unchanged will also leave reaction times and error rates unchanged. Multiplying all costs and payoffs by a constant is one such manipulation. Furthermore, V_{10} and V_{01} are differences between pairs of payoffs, rather than themselves being payoffs; any transformations on the payoffs themselves which leave these differences unchanged (such as addition or subtraction of a constant amount from all four payoffs) will not affect predicted behavior either.

When $H = 1$, dividing Eq. 15 by Eq. 16 produces

$$\frac{F}{G} = \frac{V_{10}}{V_{01}} = \frac{\Omega^* - \Omega_* + \ln(\Omega^*/\Omega_*)}{(1/\Omega_*) - (1/\Omega^*) + \ln(\Omega^*/\Omega_*)}. \quad (23)$$

This prediction about a relationship between costs of errors and error rates is, of course, not dependent on any estimated parameter.

A similar ratio prediction can be derived from Eqs. 21 and 22; it is

$$\frac{E_{\theta_1}(R)}{E_{\theta_0}(R)} = - \frac{\beta \ln[\beta/(1-\alpha)] + (1-\beta) \ln[(1-\beta)/\alpha]}{(1-\alpha) \ln[\beta/(1-\alpha)] + \alpha \ln[(1-\beta)/\alpha]} \quad (24)$$

Perhaps the most startling prediction of the model arises from the same arguments that lead to Table 2 and Fig. 1. In the completely symmetric case ($\Omega_0 = 1$, $H = 1$, $V_{10} = V_{01}$), the model predicts a maximum expected reaction time that is a function of the cost-payoff conditions but not of λ . Specifically, the maximum expected reaction time in seconds is $0.136V/k$, where V is the cost of an error and k is the cost of a second of reaction time. If $V = k$, a plausible condition, the maximum expected reaction time is 0.136 sec. The prediction is wrong; choice reaction times are not so short.

A slightly more sophisticated model is needed. The observed times can be thought of as sums of sampling times plus simple reaction time. The model as presented is a model for sampling times only; simple reaction time should be added to its predictions before they are compared with observed times. An implication is that experiments designed to test the model should include appropriate conditions for estimating simple reaction time, which may include movement time as well. The simple reaction time should not be orders of magnitude different from 0.1 sec.

Since the most important feature of this model is its intimate link between predictions about errors and predictions about reaction times, it is important in testing the model to make sure that both can be examined. For that reason, the very low error rates that have characterized many previous choice reaction time experiments are undesirable, since they make estimation of Ω^* and Ω_* difficult and inaccurate. Error rates can, of course, be raised by making the cost of time relatively high and the costs of errors relatively low, or by making the two stimuli hard to discriminate, or by doing both. A probability of error of at least 0.1 ($p^* \leq 0.9$) is desirable, though lower probabilities of error can be tolerated if the amount of data available is large.

No model that makes so many specific and easily checkable predictions has any possibility of being consistent with substantial amounts of data; only vague models or models with plenty of fittable parameters survive such confrontations. The functions of models like this one are to provide interesting dependent variables, to stimulate theoretical intuition about what variables are important and how they might work, to provide first approximations to more complex processes—in short to be wrong in interesting ways.

Stone's model treats the two hypotheses being considered as two normal distributions having the same variances but different means on some underlying subjective continuum, in the spirit of signal detectability theory (e.g. Swets, Tanner, and Birdsall,

1961). The same procedure could have been followed here. In that case, Wald (1947) has given the expression for $E_\theta(z)$ (Eq. 3:60, p. 54):

$$E_\theta(z) = \frac{1}{2\sigma^2} 2(\theta_1 - \theta_0)\theta + \theta_0^2 - \theta_1^2.$$

This expression can be much simplified by redefining the origin to be midway between the two means, θ_0 and θ_1 , and by using σ as the unit of measurement.

The assumption of equal variance is equivalent to the assumption that $H = 1$, and the assumption of normality leads to a redefinition of the single parameter of the model, λ , as follows:

$$\lambda = 2\theta_1^2 k/u.$$

The parameter λ still must be fitted from data, and all the working equations of this version of the model remain as before. Nothing has been gained by the extra parametric assumption of normality. In this model, the subject is certain to end up at one of two points on the posterior odds axis. The probabilities of error are the ratios of ordinates at those points; no integration under a distribution function is required to find them, so no assumptions about the form of that function are needed.

A binomial application: design of an experiment. Phillips, Hays, and Edwards (see Edwards and Phillips (1964)), in a large and complex experiment, have shown that subjects are unable to extract as much certainty from information as Bayes' theorem would permit. Since their experiment was large and complex, Edwards and his collaborators (in preparation) asked essentially the same question in a much simpler situation. Subjects saw two bookbags, one containing 60% red and 40% blue poker chips, the other containing 40% red and 60% blue. One was chosen at random. Subjects sampled from it, with replacement, and estimated after each observation the posterior probability that the preponderantly red bookbag had been chosen. The results indicated the same conclusion as those of the Phillips-Hays-Edwards experiment; subjects do not extract as much certainty from information as Bayes' theorem would permit. A Bayesian posterior probability of 0.98 might be produced by information that would lead a subject to estimate 0.80 or so.

But probability estimation is an unfamiliar task. Before firmly concluding that people are unable to extract as much certainty from information as they should, it is appropriate to make sure that the same inability appears in decision making as in probability estimation. In a task where information capable of reducing the risk of a terminal decision can be purchased, the amount of information purchased should be related to the amount of certainty a subject can extract from each item. The bookbag-and-poker-chip situation is well-suited to such an experiment. A subject can be required to decide whether the bookbag is preponderantly red or preponderantly blue, with a reward for being right and a cost for being wrong, and can be permitted

to buy as many samples as he wishes before making his decision. But what costs and payoffs should the experiment use?

First, consider the binomial case of optional stopping. The two bookbags have red chips in them in proportions θ_0 and θ_1 ; abandon for the moment the restriction that $\theta_0 + \theta_1 = 1$. Since the probability of getting a single red chip from the θ_0 bookbag is θ_0 , the probability of getting r red and $n - r$ blue chips in a particular order in n draws is $\theta_0^r(1 - \theta_0)^{n-r}$, and similarly for the θ_1 bookbag. So the likelihood ratio appropriate to the datum r reds in n draws in that order (or regardless of order) is

$$L(r, n) = \frac{\theta_1^r(1 - \theta_1)^{n-r}}{\theta_0^r(1 - \theta_0)^{n-r}}. \quad (25)$$

If, as in the experiment, $\theta_0 + \theta_1 = 1$, Eq. 25 reduces to

$$L(r, n) = \left(\frac{\theta_1}{\theta_0}\right)^{2r-n} = \left(\frac{\theta_1}{1 - \theta_1}\right)^{2r-n}.$$

Note that $2r - n = r - (n - r)$ is the difference between the number of red and the number of blue chips observed; when $\theta_0 + \theta_1 = 1$, only that difference is relevant to posterior opinion, and so 1004 red chips and 1000 blue ones have the same impact on orderly opinion as 4 red chips and no blue ones.

Equation 25 gives the likelihood ratio for r reds in n draws, while $E_\theta(z)$ is the expected value of the logarithm of the likelihood ratio for one draw. If that one draw were a red and the true parameter of the bookbag were either θ_1 or θ_0 , the logarithm of the likelihood ratio would be $\ln \theta_1/\theta_0$. If the draw were a blue, the logarithm of the likelihood ratio would be $\ln [(1 - \theta_1)/(1 - \theta_0)]$. If the true parameter of the bookbag is θ , then the probability that a red will be drawn is θ and the probability that a blue will be drawn is $1 - \theta$. So the expected value of the logarithm of the likelihood ratio for the general two-hypothesis binomial case is:

$$E_\theta(z) = \theta \ln \frac{\theta_1}{\theta_0} + (1 - \theta) \ln \frac{1 - \theta_1}{1 - \theta_0}. \quad (26)$$

A sufficient condition to make $E_{\theta_0}(z) = -E_{\theta_1}(z)$ is that the hypotheses be symmetrically located around 0.5, that is, that $\theta_0 + \theta_1 = 1$. (If the degenerate case of $\theta_0 = \theta_1$ is excluded, that condition is also necessary.) If that condition applies, then Eq. 26 reduces to

$$E_{\theta_1}(z) = (2\theta_1 - 1) \ln \frac{\theta_1}{1 - \theta_1}, \quad (27)$$

$$E_{\theta_0}(z) = -(2\theta_1 - 1) \ln \frac{\theta_1}{1 - \theta_1}. \quad (28)$$

Note the similarity between Eq. 27 and Eq. 19.

As a guide to choice of payoff values we wanted to get a feeling about the relations among possible levels of discriminability, payoff, and errors. Table 3 exhibits some illustrative calculations from Eq. 18 for the case $V_{10} = V_{01}$. Under the further

TABLE 3
 V_{10}/u FOR 12 POSSIBLE EXPERIMENTS

Composition of bookbag	Desired value of Ω^*			
	4	17/3	9	19
60-40	80.43	98.15	163.80	306.27
70-30	19.24	23.48	39.19	73.28
80-20	7.84	9.57	15.97	29.86

restriction that $\Omega_0 = 1$, Table 4 presents the expected number of observations required to reach the posterior odds cutoffs under these various conditions, calculated from Eq. 19. Inspection of these tables and consideration of the fact that a few fixed levels of an independent variable like V_{10}/u are preferable to many, while there is no comparable advantage in fixing levels of so hypothetical a quantity as the optimal Ω^* , led us to use only 60-40 and 70-30 bookbags, to fix $u = \$0.01$, and to use $V_{10} = V_{01}$ values of \$0.50, \$2.00, and \$3.50. (At the time the experiment was designed, we had not yet realized the nonmonotonicity of the relationship between discriminability and expected sample size, and the consequent importance of using values of VZ less than 10 as well as values greater than 10. Only two of the chosen payoffs for the 60-40 bookbag, and none for the 70-30 bookbag, yield VZ values less than 10.)

If men extract information from Bayes' theorem inefficiently, what might they be expected to do in this experiment? One very simple model for this kind of inefficiency is suggested by the results of posterior probability estimation experiments. If such

TABLE 4
EXPECTED NUMBER OF OBSERVATIONS TO REACH Ω^* WITH OPTIONAL STOPPING

Composition of bookbag	Desired value of Ω^*			
	4	17/3	9	19
60-40	10.26	14.97	21.68	32.68
70-30	2.45	3.58	5.19	7.82
80-20	1.00	1.46	2.11	3.19

estimates are reduced to inferred log likelihood ratios (the likelihood ratios necessary to produce those estimates if the estimator were perfectly Bayesian) and then plotted against true log likelihood ratios, the results are often a straight line through the origin with less than the appropriate slope. Such a finding suggests, and would be predicted by, a model which supposes that a subject treats each observation as though it had a fixed percentage less evidential value than it objectively has. In the symbols of this paper, the subject replaces $E_{\theta}(z)$ with $E_{\theta}(z') = w E_{\theta}(z)$; w might be called the inefficiency parameter, and is assumed to lie between 0 and 1. Estimation experiments suggests values of w ranging, for different subjects, from 0.05 to 0.50; 0.20 might be a convenient representative value. Since the experiment being contemplated is fully symmetric, it is convenient to define $Z' = wZ$ in order to replace Z with Z' in Eqs. 18 and 19.

Will a subject characterized, say, by a w of 0.20 buy too much, too little, or just the right amount of information? Fig. 1 says that any of these three possibilities may occur. By means of that figure, it is possible to estimate that such a subject will buy too much information if VZ is greater than 24, just the right amount if VZ equals 24, and too little information if VZ is less than 24. (Actually, the critical value of VZ that goes with a w of 0.20 is slightly less than 24.) Of course critical values of VZ larger than 24 will go with values of w smaller than 0.2, and vice versa. Computational alternatives to estimating critical values from Fig. 1 exist, but are unreasonably tedious, in view of the poor precision with which w can be determined.

It follows that a subject with a w of 0.2 will buy too much information whenever $V > 131$ and too little otherwise for 60–40 situations; the corresponding critical payoff for 70–30 situations is 71. Thus the experimental design actually used should lead to both findings, depending on V —at least for subjects adequately defined by the likelihood ratio degradation model with a w of 0.2. Of course that model cannot be taken seriously; it is far too simple. Still, preliminary results suggest that it is a good first approximation for some subjects.

ACKNOWLEDGMENTS

I am grateful to P. M. Fitts for introducing me to the problem, to L. J. Savage for guidance concerning its Bayesian formulation, to Earl S. Kramer and Tsuneko Fujii for some of the computations, and to W. H. Moore, Jr., for the computer program which produced Table 1 and the auxiliary table deposited with ADI.

REFERENCES

- BLACKWELL, D., AND GIRSHICK, A. *Theory of games and statistical decisions*. New York: Wiley, 1954.
- EDWARDS, W., LINDMAN, H., AND SAVAGE, L. J. Bayesian statistical inference for psychological research. *Psychol. Rev.*, 1963, **70**, 193-242.

- EDWARDS, W., AND PHILLIPS, L. D. Man as transducer for probabilities in Bayesian command and control systems. In M. W. Shelly II and G. L. Bryan (Eds.), *Human judgments and optimality*. New York: Wiley, 1964.
- FITTS, P. M. Cognitive factors in information processing. *Proc. XVIIth Int. Cong. Psychol.*, 1963. Amsterdam: North-Holland, 1964.
- SCHLEIFER, A., JR. Studies in the economics of optimal sample size. Unpublished Ph.D. thesis. Boston, Mass.: Graduate School of Business Administration, Harvard University, 1961.
- STONE, M. Models for choice-reaction time. *Psychometrika*, 1960, **25**, 251-260.
- SWETS, J. A., TANNER, W. P., JR., AND BIRDSALL, T. G. Decision processes in perception. *Psychol. Rev.*, 1961, **68**, 301-340.
- WALD, A. *Sequential analysis*. New York: Wiley, 1947.

RECEIVED: May 11, 1964