

AD-770 556

COSTS AND PAYOFFS IN PERCEPTUAL
RESEARCH

Detlof V. Winterfeldt, et al

Michigan University

Prepared for:

Office of Naval Research
Advanced Research Projects Agency

15 October 1973

DISTRIBUTED BY:

NTIS

National Technical Information Service
U. S. DEPARTMENT OF COMMERCE
5285 Port Royal Road, Springfield Va. 22151

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER 011313-1-T	2. GOVT ACCESSION NO.	3. RECIPIENT'S CATALOG NUMBER
4. TITLE (and Subtitle) COSTS AND PAYOFFS IN PERCEPTUAL RESEARCH		5. TYPE OF REPORT & PERIOD COVERED Technical
7. AUTHOR(s) Detlof v. Winterfeldt and Ward Edwards		6. PERFORMING ORG. REPORT NUMBER None
9. PERFORMING ORGANIZATION NAME AND ADDRESS Engineering Psychology Laboratory Institute of Science & Technology University of Michigan Ann Arbor, Michigan 48105		8. CONTRACT OR GRANT NUMBER (s) N00014-67-A-(1)181-0049
11. CONTROLLING OFFICE NAME AND ADDRESS Advanced Research Projects Agency 1400 Wilson Boulevard Arlington, Virginia 22209		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS NR 197-021 ARPA Order No. 2105
14. MONITORING AGENCY NAME AND ADDRESS (if different from Controlling Office) Engineering Psychology Programs Office of Naval Research Department of the Navy Arlington, Virginia 22217		12. REPORT DATE 15 October 1973
16. DISTRIBUTION STATEMENT (of this Report) Approved for public release; distribution unlimited		13. NUMBER OF PAGES 5660
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		15. SECURITY CLASS. (of this report) Unclassified
18. SUPPLEMENTARY NOTES		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
19. KEY WORDS (Continue on reverse side if necessary and identify by block number) Decision Costs Payoffs Flat maxima Expected value Perceptual experiment		Reproduced by NATIONAL TECHNICAL INFORMATION SERVICE U S Department of Commerce Springfield VA 22151
20. ABSTRACT (Continue on reverse side if necessary and identify by block number) A persistent problem in any kind of psychological research that reaches conclusions about inaccessible processes or experiences inside a subject's head is to validate those conclusions--that is, to exhibit persuasive reasons to believe that emitted behavior in some sense faithfully reports inaccessible processes. In the mid-1950s, perceptual researchers widely adopted an approach that might be called validation by cupidity. If the experimenter is willing to define a correct response, he can reward the subject for correct responses and not for wrong ones;		

20.

suitable reward schemes combine with an assumption of rational behavior on the subject's part to permit direct inference of internal processes. However, decision-theoretical maxima are flat, in the sense that seriously inappropriate behavior produces relatively little reduction in the subject's expected payoff. This means that costs and payoffs are rather feeble means of instructing subjects what to do, or of ensuring that he does it.

This argument is made specific in examples drawn from three kinds of perceptual experiments. In some tasks, such as probability estimation, subjects directly estimate subjective quantities, and receive rewards for accuracy of estimate. An analysis of proper scoring rules for probability estimation shows that their maxima are inevitably quite flat. An analysis of a yes-no decision task shows that the incorrect answer produces flat maxima; while the payoff function can be sharpened by increasing the magnitudes of all payoffs, a suitable relative payoff function is intractable. In such yes-no tasks, criterion variability produces even more flatness, so much so that it would be surprising if such variation did not occur in most real experiments. Criterion variability sufficient to produce a 30% reduction in estimates of d' produce only 5% to 8% reductions in expected winnings.

Implications of these results for experimental design, for interpreting experimental results, and for more general decision-theoretical thinking are discussed.

111

COSTS AND PAYOFFS IN PERCEPTUAL RESEARCH

Technical Report

15 October 1973

Detlof v. Winterfeldt and Ward Edwards

Engineering Psychology Laboratory

The University of Michigan

Ann Arbor, Michigan

This research was supported by the Advanced Research Projects Agency and was monitored by the Engineering Psychology Programs, Office of Naval Research, under Contract N00014-67-A-0181-0049, Work Unit Number NR 197-021.

Approved for Public Release;
Distribution Unlimited

Costs and Payoffs in Perceptual Research¹

Detlof von Winterfeldt and Ward Edwards

Engineering Psychology Laboratory

University of Michigan

Ann Arbor, Michigan

Styles in perceptual research change, reflecting the changing styles of psychological research and theorizing. Consider the problem, faced by every perceptual researcher, of validating his subject's responses--that is, of finding persuasive reasons to believe that the behavior emitted by the subjects in some sense faithfully reports the inaccessible processes or experiences that have just gone on inside that subject's head. (This question readily translates into the language of empty-organism psychology, and has the same meaning there; we leave the spelling out of that translation to others more masochistic than we.)

In the days of Wundt and Titchener, the persuasive reason was likely to take the form "The subject was carefully trained in my laboratory to report accurately; moreover I tried it myself and had the same experience." The attack mounted by the Gestalt psychologists and the behaviorists on this sort of truth-by-gentlemen's-agreement

was devastating, and proved fatal. Obviously, in two years you can train a willing, intelligent graduate student to report virtually anything that seems appropriate to you, in whatever language is pleasing to your ear.

The Gestalt psychologists emphasized very simple, intuitively compelling demonstrations that anyone could see. This was fine, as far as it went. But not all perceptual questions can be resolved by means of such demonstrations--and even if a demonstration (for example, of brightness constancy) is intuitively compelling to all, attempts to measure the magnitude of the effect produce the usual individual differences, and so raise the same old question about how to establish an orderly relation between responses and the experiences or processes underlying them.

The Stevens-Garner controversy of the late '40s and early '50s presented a more modern version of the structuralists' dilemma. Responses that demand a very high degree of training to use, such as direct magnitude estimation, did not agree with responses that produce much greater interpersonal agreement, with much less training, such as those based on discriminability. Was the discrepancy induced by the training, or were the two classes of procedures tapping different kinds of psychological mechanisms? In the latter case, which was the "true" representation of, say, loudness? Or, more operationally, which was more useful for applications like design of hearing aids, measurement of industrial noise, etc.?

In the mid-'50s, a new approach to this problem of validation entered the psychological literature. It did not apply to all perceptual issues, but it presented an apparently conclusive solution to the validation problem for all issues to which it did apply. The idea was very simple: arrange the circumstances of the experiment so that it is in S's self-interest to generate a response that depends in an orderly way on the internal experience or process being studied, make sure that he understands the nature of this self-interest, and then assume that he has in fact generated a response appropriate to it. The earliest version of this idea was used in signal detectability experiments (see for example, Tanner and Swets, 1954); by now it has been generalized not only to many other kinds of perceptual experiments but also to a wide variety of non-perceptual ones.

This approach produces what might be called validation by cupidity. It can be used whenever the experimenter is willing to define a function on the product set of all stimulus generation procedures used in the experiment and of all possible responses such that, given the procedures antecedent to an occasion on which a response occurs, all possible responses that might occur then can be ordered in desirability. (Translation, good enough for most purposes: the experimenter knows what the correct response is.) If so, the experimenter simply rewards the subject more for more desirable than for less desirable responses. He assumes that the subject wants to be rewarded (a safe enough assumption), and that he will apply rather sophisticated intellectual tools in order to obtain

as much reward as possible (some problems lie here). Since he finds it boring to observe subjects simply doing what the stimulus tells them to do, he designs experiments so that a subject's access to information about the merit of the response he contemplates making is fallible, and depends as sensitively as he can arrange it on the perceptual or other mechanism he wants to study. Of course he knows that the subject's information about the merits of his contemplated response also depends on other mechanisms, and that the response selection process depends on still others, but he assumes he knows fairly well what those other mechanisms are. By assuming that the subject will extract as much information as he can from the situation and the perceptual experience bearing on what response is best, and then will use that information optimally, or perhaps only systematically, to guide his responding, the experimenter can infer how effectively the perceptual mechanism is providing that information.

From this point of view, the reward structure of an experiment serves four reasonably distinct purposes:

1. Motivation: It encourages the subject to stay awake, pay attention, not goof off, and take the experiment seriously.
2. Instruction: It tells the subject what the relevant features of the experiment are, and how he is supposed to make use of them.
3. Response control: It implicitly specifies for the subject how best to translate his internal experiences or processes into observable responses.

4. Models for data analysis: It permits the use of economic models of response selection to infer the internal processes of interest to the experimenter.

All four of these purposes enter to some extent into what we are calling validation by cupidity, but obviously the last two are the crux of the idea. Although much of this chapter explores difficulties and problems connected with the idea, we might as well declare right now that validation by cupidity is the best form of validation available to psychologists who study internal processes, and indeed has no serious competitors. One consequence, of course, is that we see no way whatever (other than gentlemen's agreements) to validate responses such as magnitude estimates, in which the experimenter has no basis for attaching a value to a given stimulus-response combination.

Validation by cupidity turns perceptual experiments into gambling experiments--or, in less invidious language, into experiments on human decision making under uncertainty. Naturally, the assumption that subjects extract from the perceptual experience as much information as possible bearing on the merits of the responses available to him is a very strong one; the extraction of information from complex signals is a sophisticated and demanding process, and the appropriate mathematics can be extremely complex. Still, the approach has seemed to work. Research on absolute and differential thresholds is now dominated by it, and new applications, or more sophisticated and fancy interpretations of old ones, appear every day.

Few users of these ideas really believe that subjects are, for example, processing input information optimally according to Bayes's theorem. Yet they have no hesitation to use Bayes's theorem in analyzing the results of human information processing experiments. Why? Because it's a good bet that any response-selection procedure that is neither whimsical nor random will produce data indistinguishable from those produced by an optimal response-selection procedure based on somewhat less information. This means that analyses of data based on optimal models will almost always work, in the sense of leading to sensible and reproducible results. Typically, those results cohere fairly well with other results.

In less fancy words: decision analysis is applied to the results of gambling-type perceptual experiments because it works and it would be difficult to devise any orderly response-selecting mechanism for which such analyses would not work.

In 1961, one of us made explicit the by-then-obvious point that costs and payoffs, in perceptual or other experiments, are instructions (Edwards, 1961). In virtually all perceptual experiments, subjects must choose responses in a way that trades off one dimension against another-- false positives against missed signals, cost of errors against cost of more information, and the like. Costs and payoffs explicitly instruct subjects about these tradeoff functions, and no other form of instruction does (or, more carefully, any form of instruction that specifies relevant tradeoff functions is equivalent to a set of costs and payoffs).

This paper picks up where that one left off. Many authors, in many contexts, have pointed out that the instructions specified by costs and payoffs aren't very precise, in the sense that the decision-theoretical maxima are flat. The fact of flat maxima in decision analysis has been rediscovered, with dismay, perhaps a dozen times. This paper rediscovers that fact one more time, but this time in a relatively general treatment.

We should give a very abstract and general verbal statement of our point before going into the technical details. Our point is this: Whenever a continuum (or a dense and closely spaced set of discrete points) enters into the response selection process in an important way, whether because a continuum of responses is available or because response selection depends on cutoff points on an underlying continuum, fairly substantial changes in the location of the point on the continuum that controls responding (the response or the cutoff) will produce extremely small changes in the economic prospects of the responder.

To put it another way, we are suggesting an addition to the message of Edwards's 1961 article. Costs and payoffs, though often the only feasible form of instructions, are almost always rather feeble; although they tell the subject what he should do, they don't punish him much for not doing it.

II. Why costs and payoffs are feeble instructions

A. The mathematical formulation of the problem

The function of this section is to give technical definitions to some terms already used and some others necessary to later sections of this paper.

Consider a subject who does not care about the stimuli or about the payoff structure imposed by the experimenter. He is unmotivated, does not follow instructions, and is uninfluenced, or influenced only capriciously, by variations in costs, payoffs, or prior probabilities. Compare him with a clever and avaricious subject in the same experiment. If the economic prospects of the grasshopper (to use a technical term proposed by Aesop) are not considerably worse than those of the ant, costs and payoffs have not served their purpose. So a generally necessary condition for the effectiveness of costs and payoffs in a given experiment is that economic prospects must change substantially for different responses. We must translate this condition into formal language, defining what we mean by "economic prospects", "substantial change", and "different responses"; to do the last, we must describe responses numerically.

For a given experiment, a single number is associated with each response or set of responses: its expected value (EV). We identify EV with economic prospects, and will hereafter use these two phrases interchangeably.

We should define EV. To do so, we must first define what a stimulus is. Consider Table I, and think of it in the context of a stimulus identification experiment in which a stimulus is presented to the subject, and he must say which of several possibilities it is.

Insert Table I about here

The subject receives a payoff x_{ij} if stimulus S_i has been presented and his response was d_j . Here and throughout this paper, we use phrases like "stimulus S_i was presented" as shorthand for a more complicated idea. Technically, S_i is not a stimulus at all, but rather a set of experimental procedures designed to produce one. So the accurate but tedious phrasing would be "Operations S_i were performed on or by the apparatus, producing a stimulus that was presented to the subject." Similarly, an experiment asking the subject which stimulus was presented is really asking which operations on or by the stimulus-generating apparatus were performed. Of course in many perceptual and other experiments the stimulus produced by operations S_i may vary from instance to instance within the experiment, and may be only imperfectly known to the experimenter.

The presentation of a particular stimulus S_i will give the subject information, perhaps fallible, about which stimulus actually was presented. After the occurrence of S_i , the subject will have a personal probability distribution over the possible stimuli; that distribution associates with each S_i a probability ξ_i . The expected value to the subject of making response d_j is then defined as

$$EV(d_j) = \sum_i \xi_i x_{ij} \quad (1)$$

For stimuli that vary continuously, integration replaces summation in this definition.

We have no reluctance about identifying the economic prospects of response d_j with $EV(d_j)$, thus implying that subjects should select a response-generating procedure that maximizes EV. An antique fallacy questions this recommendation, arguing that EV maximization is a wise strategy only for repeated events. As we see it, the very word "strategy" implies that some principle of response selection will be repeatedly applied; the arguments that make EV maximization optimal don't care whether the conditions are or are not constant from one application of the strategy to the next. Besides, most perceptual experiments repeat the same condition often enough so that the most passionate relative-frequentist would agree that EV maximization is wise.

Though we unhesitatingly identify the subject's economic prospects with EV, the subject may not. Whether men in fact maximize EV or not is a profound, difficult question; to examine it here would take us too far afield. A great deal of experimentation has produced no evidence against that hypothesis; data analyses based on it abound in perceptual experiments, mostly because they work. Moreover, any other consistent decision rule that permits all possible economic outcomes of d_j to enter into assessment of its worth with appropriate sign and with some monotonic variation of the economic consequences of d_j with the magnitude of each payoff will, for virtually any theoretical or practical purpose, be indistinguishable from EV maximization. (Note that we are using the notion of value to

include subjective values as well as dollar payoffs; our arithmetic will be done on dollars, but the nature of our results will imply that it makes little difference whether or not the effective payoff to the subject is or is not linear with its dollar value, so long as it is strictly monotonic).

Next we must give an account of perceptual experiments that permits us to define responses and the strategies that produce them. Our account, naturally, will be decision-theoretical in spirit. Here and throughout, we shall speak of both stimulus-generating operations and responses as being chosen from mutually exclusive, discrete, finite sets. This mild idealization slightly simplifies the mathematics and greatly simplifies the language. Anyone mathematically demanding enough to be bothered by it will also be mathematically skilled enough to see the easy generalizations of the arguments to continuous cases. Following the current conventions of perceptual research, we shall assume that a discrete stimulus-generating operation may lead to any member of a continuously distributed set of alternative stimuli; as we said above, "stimulus S_j was presented" refers to the discrete operation S_j , not to the actual stimulus.

We suppose that before S_j is presented, the subject has a prior probability vector over the set of possible stimuli. Upon observing (the stimulus resulting from operation) S_j , he transforms that vector into a posterior probability vector. The correct rule of transformation, of course, is Bayes's theorem. While we shall use that where relevant, we do not need to assume that the subject uses it; the weaker assumption that the subject

represents his observation internally as a likelihood ratio, or some quantity monotonically related to likelihood ratio, and that he takes prior probabilities into account in response selection is sufficient. Now the subject makes the response from among those available to him that has the best economic prospects. He may or may not receive feedback. In a single-observation experiment, that terminates the trial. In a multiple-observation experiment he may choose or may be required to make more observations. All experiments end each trial with selection of a terminal response, perhaps followed by feedback; some require responses (other than decisions to look at more information) interspersed with the sequence of observations.

A decision rule is a cutoff vector or set of cutoff vectors defined in the set of vectors of posterior probabilities. It specifies which response will be selected given any such vector. In a two-alternative stimulus identification task with a symmetric payoff matrix, for example, the cutoff vector would probably be $(\frac{1}{2}, \frac{1}{2})$. The subject selects response d_j whenever his personal probability that S_j was presented exceeds $\frac{1}{2}$. For an optimal subject, or one whose data are being analyzed as though he were optimal, choice of decision rule depends only on the payoff structure of the experiment, and so can be known before S_j is presented.

In most experiments, decisions depend not only on payoffs and on observations, but also on prior probabilities. Information about the prior probabilities can be combined with information about the payoffs, permitting the decision rule to be redefined (via Bayes's theorem) from the set of vectors of posterior probabilities to a set of vectors of likelihood ratios,

and sometimes further from the likelihood ratios to a set of physical characteristics of the actual stimulus, and occasionally further from the actual stimulus to the stimulus-generating operations S_1 . Any such redefinition of a decision rule out of the set of vectors of posterior probabilities into some other set more directly related to the stimulus presentation conditions of the experiment we shall call a decision function. A decision function partitions the set of possible observations by means of criterion or cutoff points. In signal detection theory, for example, the likelihood ratio criterion β partitions the set of actual observations (insofar as that set can appropriately be mapped into a set of likelihood ratios). Within each class of observations, the same response is appropriate. Obviously any translation from a decision rule to a decision function depends on a model of the sequence of processes beginning with the stimulus-generating operations and ending with the internal process representing the stimulus, and is no more trustworthy than that model. The arguments of this paper complicate that point by showing that the inference from responses to parameters of such models is quite weak; whether that is good or bad depends on the purpose for which the model is being used.

In some experiments, such as those concerned with probability estimation, the response categories available to the subject are continuous, or more often numerous, ordered, and densely spaced. In others, such as stimulus identification experiments, the set of available response categories will typically be sparse. However, the concepts of decision rules and decision functions permit us to think of the response-selecting process

as either continuous or as having numerous ordered and densely spaced alternatives in such experiments also. So, in order to examine the specificity with which costs and payoffs control responding, we can simply plot the expected value of the response, decision rule, or decision function, as appropriate, against a numerical representation of that response, decision rule, or decision function. Our main interest, of course, is in the shape of this function around its maximum, since we assume that subjects, motivated by cupidity, try to select their responses in a fairly optimal way.

The thesis of this paper can now be more exactly stated: the expected value of responses, decision rules, or decision functions changes only slightly with large deviations from optimal values. Consequently the economic prospects of the grasshopper may be only slightly worse than those of the ant; economic prospects often do not change substantially for different response selection rules.

B. Expected value as a function of decisions

Consider a recognition task in which the subject must estimate on a 0-to-100 scale the probability that the current stimulus is old; that is, has been presented before. The experiment has two stimulus-generating conditions, old and new, a dense, orderly set of available responses, and a payoff defined for each stimulus-response combination. Let ξ_j be the response the subject chooses to make on a given trial. Then $x_{0,j}$ is the

payoff for it if the stimulus is old, and $x_{n,j}$ is the payoff for it if the stimulus is new. Let ξ_t be the subject's actual personal probability on this trial that the stimulus is old; of course, nothing guarantees that $\xi_j = \xi_t$. The subject can evaluate any response ξ_j by its expected value:

$$EV(\xi_j) = x_{o,j} \xi_t + x_{n,j} (1 - \xi_t). \quad (2)$$

For each value of ξ_t , a value of ξ_j will exist for which $EV(\xi_j)$ will be maximized. This maximal expected value EV^* will be a convex function of ξ_t . Often, the experimenter will try to encourage the subject to report his true opinions by using a function to specify the x 's such that EV is maximized whenever $\xi_j = \xi_t$. Such functions are called proper scoring rules; for leads into the extensive literature about them, see Aczel and Pfanzagl (1966), Murphy and Winkler (1970), or Savage (1971).

A typical function plotting maximum expected value given optimal choice of ξ_j , EV^* , against ξ_t is shown in Figure 1. It represents the quadratic scoring rule, one of the two most frequently used proper scoring

Insert Figure 1 about here

rules, where $x_{o,j} = 1 - (1 - \xi_j)^2$; $x_{n,j} = 1 - \xi_j^2$. The two lines show the EVs of responses ξ_1 and ξ_2 as a function of ξ_t . Of course ξ_1 would be the optimal action if the subject's actual probability were ξ_1 , and similarly for ξ_2 ; this follows from the definition of proper scoring rules. Now, assume that the subject's actual probability is ξ_1 , but that he nevertheless chooses ξ_2 as his response.

His economic prospects are diminished by the difference between the EVs of ξ_1 and ξ_2 given actual probability ξ_1 : in Figure 1 that difference is labelled Δ . It will typically be very small in relation to the total EV if the difference between actual and optimal response is not too large, especially if the optimal EV function is itself rather flat in the region of ξ_t . The nearer ξ_t is to the minimum of the function, that is, the more uncertain a subject is about what response is best, the less he will suffer as a consequence of suboptimal decisions.

The fact that proper scoring rules have this undesirable property of flatness is well known (see, for example, Murphy and Winkler, 1970). However, the ubiquity of proper scoring rules is less well known. First, notice that the label on the response is irrelevant; proper scoring rules are not confined to situations in which the response is an explicit probability estimate. A useful distinction can be made between conditions that must be satisfied to generate a proper scoring rule and conditions that must be satisfied just to recognize one. It is easy to recognize a proper scoring rule: any list of acts that includes none that are dominated, stochastically dominated, or duplicated is based on a payoff matrix that is an extract from a proper scoring rule. That sentence sounds fancier than it is. All it means is that if each act has the property that some set of vectors of probabilities makes it optimal, then choice of that act in effect signals estimation of a vector within that set. Any response is a probability estimate. Moreover, the linearity of the equation for EV guarantees that if

an act is optimal given more than one vector of probabilities, then all vectors for which it is optimal will be adjacent to one another. That word "adjacent" is here used in the rather special sense that all these adjacent vectors will fall within a closed convex region of a regular hypertetrahedron of dimensionality one less than the number of probabilities in the vector.

(The usual sets of conditions that define proper scoring rules are more complicated, because they typically are designed to ensure that different vectors produce different scores. But such issues are irrelevant here; we are only interested in the fact that different acts identify different sets of vectors.)

This argument for the ubiquity of scoring rules means little for payoff matrices in which the number of acts approximates the number of states. But for payoff matrices in which the number of acts is very much larger than the number of states, and yet no acts are dominated or stochastically dominated, the argument for flat maxima given above become increasingly applicable. If the act space is continuous or acts are densely distributed over a continuum, the argument for flat maxima applies with full force.

C. Expected value as a function of decision rules

So far we have discussed reductions in EV that result from a single nonoptimal act. Now we will examine the effect of consistently applied nonoptimal decision rules. Consider a task in which the subject must discriminate two objects according to brightness. He can make either of

the following two responses:

d_1 = Stimulus 1 is brighter than stimulus 2

d_2 = Stimulus 2 is brighter than stimulus 1

Correspondingly, on each trial, there are two possible states of nature; either S_1 or S_2 is physically brighter. The experimenter defines a payoff matrix as in Table II, where a and d can be thought of as payoffs, b and c as costs. We assume that the subject must perform this discrimination

Insert Table II about here

task repeatedly for different pairs of stimuli. After presentation of a pair of stimuli, the subject will have some probability distribution over the two states of nature:

$$\Pr(S_1) = \xi, \Pr(S_2) = 1 - \xi$$

The optimal strategy for such a task can be derived easily. The subject should choose d_1 , whenever

$EV(d_1) > EV(d_2)$, d_2 otherwise. I. e., he should choose d_1 if

$$\xi a + (1-\xi) c > \xi b + (1-\xi) d \quad (3)$$

or

$$\xi > (d-c) / [a + d - b - c] = p^* \quad (4)$$

This result is represented in Figure 2.

The EVs of the two decisions are linear functions of ξ and their intersection defines the cutoff point p^* . Therefore optimal strategy in this problem is totally determined by p^* . Now assume that the subject in fact adopts a strategy $p \neq p^*$, i.e., he chooses decision 1 whenever

Insert Figure 2 about here

$\xi > p$ and decision 2 whenever $\xi \leq p$. How will the EV of this decision rule compare with the EV of the optimal strategy? It depends on the subject's prior opinion about the posterior distribution of ξ over trials. Assume for a moment that the subject considers all values of ξ to be equally likely. The EV can be expressed in terms of the cutoff point p .

$$\begin{aligned} EV(p) &= \int_{\xi=p}^1 EV(d_1) d\xi + \int_{\xi=0}^p EV(d_2) d\xi = \\ &= 1/2 p^2 [a + d - b - c] + p[d - c] + 1/2 [a + c] \end{aligned} \quad (5)$$

which is a quadratic function of p , whose parameters are determined by the costs and payoffs involved. Its first derivative is

$$EV'(p) = -p [a + d - b - c] + [d - c] \quad (6)$$

From this it follows that the maximum EV is obtained when we set

$$p = p^* = (d - c) / [a + d - b - c] \quad (7)$$

which we saw in Eq. (4).

The first derivative of the function $EV(p)$ tells us how steep $EV(p)$ is around the optimal value of p^* . Equation (6) shows that the steepness of $EV(p)$ depends on the values of all payoffs. The larger the costs and payoffs, the steeper this function will be--which is not a surprise. Two examples with values typically used in psychological experiments will give an idea of how flat this function will be in most situations. The results for values of $a = d = +1\phi$, $b = c = 0\phi$ (1) and of $a = 1.5\phi$, $b = c = -.5\phi$, $d = +.5\phi$ (2) can be seen in Figure 3.

Insert Figure 3 about here

A more intuitive way of looking at the expected value as a function of a response strategy p results from an inspection of the areas in Figure 2. The area under the two heavy lines defines the average win under an optimal strategy p^* . The shaded area represents the average loss due to a suboptimal strategy p . Graphs like this can help the experimenter to gain insight into the effectiveness of his payoff structure.

Mathematically the description of the expected value of a strategy p as a function of p and the outcome structure is very convenient. It is easy to see that by multiplying all outcomes by a factor greater than 1, the experimenter can steepen this function as much as he wishes. However, increasing the steepness of the $EV(p)$ function may not change the relative loss of any non-optimal strategy, expressed as a percentage of the optimal EV. Consider the payoff matrix displayed in table II and the resulting

EV function over strategies p . If we multiply all outcomes by a constant $g > 1$, we will get

$$EV_g(p) = g EV(p) \quad (8)$$

which is steeper than the original EV function. Now consider the relative quantity

$$REL(p) = 100 [EV(p^*) - EV(p)] / EV(p^*) \quad (9)$$

where REL stands for relative expected loss and is measured in percent (see Edwards, 1956). If REL is the value with which subjects are concerned, multiplying all outcomes by a constant does not affect the motivational effect of the costs and payoffs, since

$$REL(p) = REL_g(p) \quad (10)$$

The experimenter will often be concerned with manipulating the REL function rather than the expected value function. This is mathematically simple but unpleasant computationally. Let

$$U = [a + d - b - c] \quad (11a)$$

$$V = [d - c] \quad (11b)$$

$$W = [a + c] \quad (11c)$$

From this it follows that

$$p^* = V/U \quad (12)$$

$$EV(p^*) = 1/2 [w + v^2 / U], \quad (13)$$

$$REL(p) = 100 [1 + U^2 p^2 / (UW + v^2) - 2UVp / (UW + v^2) - UW / (UW + v^2)]. \quad (14)$$

In the foregoing two examples we find that the two REL functions are similar in flatness to the EV functions (Figure 4). For example, for

- - - - -
Insert Figure 4 about here
- - - - -

the symmetric (first) payoff matrix, the subject would not lose more than 8% of the expected value of the optimal strategy $p^* = \frac{1}{2}$ for any value of p between $\frac{1}{4}$ and $\frac{3}{4}$. Similar analyses are mathematically fairly easy to develop for the more general case of n decisions and m states but they lack the simple graphical interpretation displayed here--unless you happen to be good at visualizing convex regions within regular hypertetrahedrons.

REL functions like those shown in Figure 4 are extremely useful ways of examining the properties of payoff schemes, but they need careful interpretation. This part of the paper will consider only symmetric 2×2 payoff matrices and prior odds of 1:1; similar but somewhat more complicated arguments apply to more complicated cases. In symmetric cases, the subject can guarantee that he will be right half the time simply by flipping a coin. Any non-perverse strategy must be at least as good as that. So the maximum feasible REL is defined as $1 - (1/2p_c)$, where p_c is the probability of being correct if the subject uses the optimal strategy. If p_c is 0.6, the maximum feasible REL is $1/6$; if p_c is 0.9, it is $4/9$.

From the subject's point of view, the REL expresses how much difference in economic prospects exists between an optimal ant and a feckless grasshopper, for various degrees and kinds of fecklessness.

If a grasshopper can earn $5/6$ as much as an ant does without even noticing the stimulus, he doesn't have much incentive to notice it, much less to think about its meaning.

From a somewhat different point of view, the distance between the optimal REL of 0 and the maximum possible REL defines the range of degrees of success the experimenter can have in inducing his subject to perform in an ant-like rather than a grasshopper-like way, for fixed stimuli, responses, and payoffs. The nearer the subject gets to 0 REL, the more successfully the training and economic pressures are combining to produce ant-like behavior. From this point of view, the available REL range, numerically equal to the maximum feasible REL defined above, might be taken as 100% and the percentage deviation of REL from 0 on this transformed scale might be taken as an index of fecklessness.

Of course REL is calculable only if p_c , the probability of correct response if the subject uses the optimal strategy, is known. Often, it won't be.

In cases using asymmetric payoff matrices, unless they are balanced by prior odds asymmetric in the other direction, the subject can gain much more than half of the available payoff by a stimulus-ignoring strategy: simply pick the act that has the larger sum of payoffs over states. If the asymmetric payoffs are exactly balanced by priors asymmetric the other way, the situation is reduced to the symmetric case discussed above. And if the priors are even more asymmetric than the payoffs, once again the

subject can gain much more than half the available payoff by a stimulus-ignoring strategy. So the general tendency of such asymmetric arrangements is to reduce the advantage of the ant over the grasshopper-- provided that the grasshopper at least notices and exploits the payoffs and prior probabilities.

Nonuniform distributions over ξ will change the foregoing analysis, but as long as values of ξ close to p^* or far away from it are most probable, nonuniformity will only increase the flatness of the expected value or the relative expected loss function. In other words, subjects will suffer less from using nonoptimal strategies if their task is very difficult or very easy than if it is of moderate difficulty.

For tasks of moderate or severe difficulty, the subject's prior expectations about the relative frequency of the various stimulus conditions bear importantly on posterior values of ξ . During the early trials of an experiment, the subject's distributions over ξ may be fairly uniform. This means that he will face EV and REL functions like those in Figure 3 and 4-- relatively steep. As a result of his experience over a sequence of trials, his prior and consequently his posterior distribution over ξ on each trial will sharpen considerably around the values implied by the experimenter's stimulus presentation frequencies and the difficulty of the task. The experimenter typically does not want the subject to find the task easy, so these values are likely to lie in the region between 1/4 and 3/4. An experimental procedure that leads to values of ξ in this range will flatten out the EV and REL functions considerably over those of Figure 3 and 4. The

Effect will be especially marked if only a few distinct stimulus relative frequencies and task difficulties are used.

D. Expected value as a function of decision functions

In decision problems with observations, flat maxima appear when we plot EV against decision functions, or, equivalently, against decision criteria. Consider a simple auditory detection task (a yes-no task) in which a pure tone may or may not be embedded in white noise and assume a general signal detection model (see for example Green and Swets, 1966). The subject must perform a decision task based on a single observation. According to the general signal detection model, the subject's observation is a random variable generated either from a signal distribution $f(y|S)$ or from a noise distribution $f(y|N)$. The subject is assumed to apply some likelihood ratio criterion in order to decide whether he should choose act 1 (observation was generated by the noise distribution) or act 2 (observation was generated by the signal distribution). Assuming that k is monotone with β , a likelihood ratio criterion β generates the following decision function δ :

$$\delta : Y \rightarrow D = \{d_1, d_2\} \text{ with}$$
$$\delta(y) = d_1 \quad \text{whenever } y < k$$
$$\delta(y) = d_2 \quad \text{whenever } y \geq k$$

where k is the solution to

$$f(k|S) / f(k|N) = \beta \tag{15}$$

The general formula for the EV of such a decision function for a continuous state and decision space is

$$EV(\delta) = \int_S \int_Y V(s, \delta(y)) f(y|S) \psi(s) ds dy \quad (16)$$

where $V(s, \delta(y))$ is the outcome associated with a particular state value s and the decision defined by $\delta(y)$. $\psi(s)$ is the prior distribution over the states. Solving for the particular 2x2 situations of this signal detection task, we obtain for a payoff matrix like the one in Table II and prior $\psi(N) = \xi$

$$EV(\delta) = \Pr(Y < k|N) \xi(a - b) + b\xi + \Pr(Y < k|S) (1 - \xi) (c - d) + d(1 - \xi) \quad (17)$$

Since this expectation is solely a function of k , we can write

$$EV(\delta) = EV(k) = EV(\delta)$$

From first derivative of this function it can easily be shown that $EV(k)$ is maximized for $k = k^*$, where k^* is the solution to

$$f(k^*|S)/f(k^*|N) = [(1 - \xi)/\xi] [d - c]/(a - b) = \beta^* \quad (18)$$

This is, of course, a familiar result in signal detection theory. β^* is the optimal likelihood ratio criterion specified by the payoffs and prior probabilities of the experimental task. How EV changes as a function of changes in k or β depends on the conditional distribution of Y . Specific results can be derived only if we restrict ourselves to particular distributions. However, the expected value of extreme policies can easily be

derived:

$$EV(k = +\infty) = a\xi + (1 - \xi)c \quad (19)$$

$$EV(k = -\infty) = b\xi + (1 - \xi)d \quad (20)$$

Thus, as a first check, without assuming any specific distribution, the experimenter can analyze the maximum differences in expected value generated by different response strategies (excluding, of course, a diabolical subject who would use an optimal decision criterion but reverse his decisions.)

Next, consider

$$EV(k^*) - EV(k) = \left[\int_k^{k^*} f(Y|N) dY - \beta^* \int_k^{k^*} f(Y|S) dY \right] \xi(a - b), \quad (21)$$

the loss in EV caused by a nonoptimal strategy k (without loss of generality we assumed that $k < k^*$). Again without assuming any distributions, some major implications can be drawn, which are illustrated in Figure 5. First, note that the difference in EV will be linear in the prior probability and the payoffs.

Insert Figure 5 about here

For extreme values of ξ or for large costs and payoffs the loss defined by Eq. (21) will increase. Second, assume that both conditional distributions have a small amount of overlap. Then the expression in the brackets of Eq. (21) will typically be very small, since the integrals are represented by the areas shaded in Figure 5 around the optimal cutoff point k^* . In other words, a large d' would therefore mean that suboptimal values of β would result in little less EV than would be produced by β^* . Similar conclusions can be drawn for a large amount of overlap, i.e., a small d' . If β^* is near 1,

the expression in the brackets of Eq. (21) will become small, since both integrals will have nearly equal values. If the separation is not extreme and if k^* is shifted in the direction of the mean of either distribution by changing prior probabilities or payoffs, one of the two integrals in Eq. (21) will usually become very small, while the other increases. The relative change in expected value, however, will be smaller than $\xi(a - b)$. This intuitive analysis permits the experimenter to visualize the effects of cost and payoffs. Values of ξ near 1/2 and quite large or quite small separations of the conditional distributions tend to produce flat expected value and relative expected loss functions. In Figure 6 the expected value is plotted

Insert Figure 6 about here

as a function of k and d' . Here we made the assumption that both conditional distributions are normal. (see also Chinnis, 1971). The results are plotted for the symmetric payoff matrix. Figure 7 shows the relative expected loss version of the same information. It highlights the fact that intermediate values of d' produce most steepness. You can see that for all levels of d'

Insert Figure 7 about here

a difference of ± 1 standard deviation between the optimal and the actual decision criterion would result in a loss of only about 10% or less. The equation relating β , k , and d' is

$$\ln \beta = k d'. \quad (22)$$

In other words changes in k are linearly related to changes in $\ln \beta$, and the larger the value of d' the larger is the change in $\ln \beta$ associated with a given change in k .

Other decision functions, the so-called stopping rules, can represent response sets whenever stimuli from the same stimulus-generating procedure are presented in sequence and the subject can decide after seeing each stimulus (or perhaps less frequently) if he wants to make a response or if the stimulus presentation should be continued, perhaps at a cost. Rapaport and Burkheimer (1970, 1971) showed that EV-functions are also flat over stopping rules. DeGroot (1970) argues that costs and payoffs for final responses have much less influence on stopping rules than the cost of observations, if the latter cost is relatively small. The flatness of the EV-function over stopping rules is important to studies of multiple observations in signal detection theory, to studies of choice reaction time, and to studies of information-purchasing decisions.

E. Expected value as a function of criterion variability.

So far we assumed that the subject's responses were generated by a systematic procedure, such as the use of fixed cutoff points. How will the subject's economic prospects change, if he varies his decision rule or decision function randomly?

As an example, we assume that a subject varies his p criterion in a simple decision task randomly according to some uniform distribution between values p_0 and p^0 , both equally far from the optimal criterion p^* . We know from the previous discussion about expected value functions over decision rules that

$$EV(p) = -1/2p^2 U + p V + 1/2 W \quad (\text{see equation 5})$$

Since p has a uniform distribution, we can infer the following expectation:

$$E (EV(p)) = \int_{p_0}^{p^0} [EV(p) / (p^0 - p_0)] dp \quad (23)$$

Figure 8 shows how this expectation varies as a function of p^0 for the two payoff matrices used in Figure 3. Note that for the symmetric payoff matrix (1), a subject will not lose more than 12% of the maximal EV, if he should choose a decision rule at random from the interval $[0, 1]$ at each trial. Of course this expectation will always be larger than the EV of p_0 and p^0 and smaller than the EV of p^* . Thus the flatness of the expected value function as a function of p determines the flatness of the expectation considered here. Furthermore, the assumption of a uniform distribution makes this expectation artificially small. In most situations, we might expect the subject to vary his response criterion according to some bell shaped function around a criterion p , which might even be the optimal criterion p^* . This in turn will increase the expectation substantially-- that is, it will make the EV of the random strategy even closer to the EV of the optimal strategy p^* .

It is well known that criterion variability, within any signal detection model, will result in lowered estimates of d' (or equivalent quantities), but the magnitude of the effect is not so well known. Unfortunately, the function that relates criterion variability to decrease in d' is not at all flat; it is virtually linear over interesting parts of its range.

Table III is based on our standard symmetric payoff matrix, and assumes two normal distributions differing only in mean. It shows the

Insert Table III about here

effect of various amounts of criterion variability at various levels of true d' , along with the economic consequences, in EV and REL, of this kind of suboptimal behavior. It matters very little what the form of the criterion variation function is, so long as its mean is at k^* and its standard deviation is the one listed in the table. The computations in Table III assumed normally distributed variations in k ; very similar computations would apply for other reasonable assumptions. Criterion variability sufficient to produce a 30% reduction in d' will cost the subject only .05, .08, and .08 of a cent, for a one-cent difference between the value of right and of being wrong; these numbers all correspond to RELs of less than 10%. If it costs so little to let his criterion vary by one standard deviation or more, why should a subject bother to hold it constant? Nor will he lose much sleep over the plight of the poor experimenter, who thinks d' is 1.4 when it is actually 2.0.

III. Consequences for research, in perception and elsewhere

We have spent many pages discussing flat maxima by means of examples but we never answered the obvious question: How flat is flat? We might choose some criterion of flatness--and indeed our examples suggest candidates. But without behavioral evidence of the effects of flatness and of steepness, any such choice would be arbitrary. We might instead define flatness by means of its behavioral results, and let the experimental literature tell us how flat is flat. But the literature is silent; the obvious experiments haven't been done. So we choose to leave the definition of flatness to your intuition--and, we hope, to the data that our sad story leads you to collect.

We are both decision theorists, and much of the preceding discussion shows the feebleness of decision theory. A natural conclusion for us is to order two steins of hemlock. While we sip, we can amuse ourselves by observing the asymmetry between those who wish to use decision theory and those who worry about the fact that they and others aren't using it. Those who wish to use decision theory, whether as a basis for experimental design or as a practical tool in real contexts, should be seriously disconcerted by our whole line of reasoning, since our fundamental conclusion is that decision-theoretical niceties, such as eliciting exactly appropriate values of probability or payoff, are unimportant. An escape from this conclusion may be that EV, rather than REL, is what really counts, and a 1% decrease of EV in a $\$10^9$ decision is still a $\$10^7$ loss. That thought helps consultants more than it does experimenters.

Those who worry about the obviously non-optimal behavior of decision-makers, and wonder how the world manages to hold together in spite of human inefficiency, should look at the previous pages with quite different eyes. Our argument is readily interpretable as a theory about the answer to that question--and a persuasive one. The world holds together and functions as well as it does because major stupidities, inefficiencies, selfishnesses, and the like on the part of decision-makers produce only minor losses in EV. The physical facts of life are typically far more important than the subtleties of human decision-making in controlling the EV of an action, for two intertwined reasons. First, the outcome of any significant decision depends on Nature (or chance, or the opponent, or whatever) as well as on the act chosen. A good decision can lead to a bad outcome, and often does. Second, the flatness of EV functions means that the best decision may be hard to discover.

Moreover, since wide variations in decision lead to similar EVs, men are encouraged to attribute success to their own insight in varying their decisions, rather than to chance. (On this state of mind rests the state of Nevada.) We attribute success, in ourselves or in others, to good management rather than to good luck--and then wonder what is wrong when our luck changes.

The entire preceding discussion should be immensely reassuring to those who worry about the fate of man and the vagaries of politics. But it does not speak to the problems of the researcher on perception who, having been offered and having accepted decision theory as a major research tool, now finds out that it is a feeble one.

Our comfort for him, if not cold, is no better than luke-warm. We do not believe he has any alternative to decision-theoretical experimental designs. Validation by cupidity, though less precise than you and we might like, is the only intellectually acceptable form of validation we know of. And validation by gentlemen's agreement, its only serious competitor, is no validation at all. (Except, perhaps, for in-groups of agreeing gentlemen.) And validation by convergent operations is simply a more complicated and elaborate form of validation by gentlemen's agreement, unless those converging operations give some good reason, other than cooperativeness and good-fellowship, for subjects to search their souls and report their true opinions.

But we do believe that perceptual experiments built on decision-theoretical ideas can be improved. Indeed, we have in this paper presented the tools we regard as most useful for making such improvements. The experimenter should know how effective his payoff arrangements can be, and he should pay at least as much attention to improving them as to improving his stimulus-generating apparatus.

A. What to do: increasing the effectiveness of costs and payoffs

Experiments that use costs and payoffs are gambling experiments; the key to improving them is manipulation of their gambling characteristics. The experimenter will typically know exactly what his costs and payoffs are; he is likely to know somewhat less about the prior probabilities affecting subjects' choices, and considerably less about posterior probabilities. But

he at least knows a lot about the variables, such as stimulus parameters, that affect posterior probabilities.

The graph of the EV function over measures representing responses or response sets characterizes the effectiveness of a payoff function, other experimental parameters held constant. The experimenter wants to make that graph as steep around its maximum as he can manage.

The most obvious and most successful way of doing this is simply to increase the absolute values of both costs and payoffs. This leaves REL unaffected, but strongly affects EV--and no one pays off in RELs. The main problem in doing so is that large costs and payoffs introduce money-management problems into the experiment. However, such problems can at least be palliated; often, they can be solved. For an example of how far one can go with college student subjects, see Swensson and Edwards (1971) and Swensson (1968). The latter experiment used costs and payoffs up to \$10 for a single response.

In money management of gambling experiments, one principle outweighs all others: Cheat! Adjusting EV to come out right is no problem, no matter how large the stakes. But increasing stakes increases the variance of total earnings, and that is the problem. An experimenter who is a careful cheater will arrange his experimental conditions so that most subjects win less than they should--and will add further stimuli at the end of the experiment designed to bring them up to the desired pay rate. In perceptual experiments, this is especially easy, since manipulations that make the subject right nearly 100% of the time are ridiculously easy to perform and hard to detect. Inflexible

rules about the minimum rate of pay for each subject will help get such research strategies past Human Subject Committees.

If the experimenter is concerned about steepening the REL as well as the EV function, he faces a more difficult problem. Equation 14 helps-- a little.

Variation in the physical values of stimulus will produce variation in posterior probabilities; to some extent this helps control the distribution of over trials, and that in turn enters into the sharpness or flatness of the REL function. Translated out of the mathematics, this says that the experimenter can steepen the REL function considerably by mixing up his stimuli. Uniformly hard-to-discriminate stimuli (in a discrimination experiment) will lead to a peaked function and so to a flatter REL function. Whether this prescription for peaking up the REL function conflicts with the experimenter's wish to get as much information as possible from each trial of his experiment will depend on the experiment. Often, the conflict will exist and a compromise will be called for. As has already been pointed out, skewed payoff matrices are a poor idea unless they are balanced by prior distributions skewed the other way. Extreme priors make the EV function more skewed, but probably cost too much in reducing the informational content of the experiment.

In signal detection experiments of the yes-no type, how does d' relate to EV? Figure 6 shows that different values of d' will result in

differences in the maximum attainable EV, and therefore in the flatness of the EV function over values of k or β . For symmetric situations, and assuming normal distributions, a d' between 2 and 3 is best.

Variability in cutoffs flattens the EV function. Any procedure that steepens the EV function will of course steepen the function resulting from a specified cutoff variability. In addition, cutoff variability may be influenceable by instructions, and will certainly be a function of EV variance; the more EV a subject stands to lose by varying the cutoff, the less likely he is to do so.

Not much can be done for probability estimates or for sequential-sampling experiments, other than raising the stakes.

Our crucial point here is that the experimenter should look at the appropriate function relating EV or REL to both optimal and other behavior, and use that function as a tool of experimental design. He can also use it as a tool of data interpretation if he is willing to make some fairly plausible and straightforward assumptions about what the subject has done in the face of the experimental conditions.

An important way in which an experimenter can improve his experiments is to measure the extent to which the payoffs he intends to use in fact lead to optimal or near-optimal behavior. For instance, in a signal detection experiment he can use known stimuli, such as numbers sampled from one of two normal distributions, in preliminary calibration sessions, along with the payoff structure intended for the main experiment. This procedure

would at least give the experimenter some idea of how much criterion variability and suboptimal responding to expect. It might also suggest modifications of the payoff structure.

In experiments in which trial-to-trial fluctuations of stimuli can be physically measured (e. g. by means of a microphone in or near the ear), they should be. Such measurements can suggest inferences about the nature of the effective stimulus on this trial, and thus inferences about what the subject is actually doing.

B. What to do: interpreting experiments

Typical payoffs in perceptual experiments have been a few cents, or even fractions of cents or points. This does make management easier--but we hope that by now you are chewing your fingernails about the flatness of the maxima that result.

If the subject's economic prospects are essentially unaffected by what he does, what will he do? He might behave randomly, especially if he doesn't know he is doing so. Often, other facts about the experiment imply non-monetary payoffs, and the subject is likely to pay most attention to these. He likes to be right, so he may not respond well to asymmetric payoffs. He may minimize effort, perhaps by as extreme a procedure as ignoring the stimulus or perhaps by hunting for some simple strategy. He may look for hidden meanings in the experiment, and so develop self-instructions and idiosyncratic response strategies. (The famous probability matching phenomenon in repetitive-binary-choice experiments, to the extent that it occurs at all, seems to be a result

of this kind of self-instruction to vary responses.) Or he may just be bored.

We wish we could cite a string of experiments all showing the effects of flat maxima. We can't. The difficulty is that most perceptual (and other) experiments are not designed to disentangle such decision variables from the perceptual variables being studied. The preceding section contained some comments about how to avoid this confounding, but the existing literature contains few instances of successful avoidance.

The effect of flat maxima may be to produce random errors, systematic errors, or both. Several signal detection, recognition, and discrimination tasks show that subjects adopt response strategies that lie between optimal and error minimizing strategies. In signal detection experiments, Green (1960), and Swets, Tanner, and Birdsall (1961) found that subjects tended to deviate from optimal cutoffs in the direction of making both responses more nearly equally often. Green concluded that "The way in which the expected value changes for various criterion levels is the crux of the problem;" that is, that he had flat-maximum trouble. Others have proposed that subjects regard such tasks as being sensorily oriented, or that subjects do not in fact maximize EV. These hypotheses cannot be studied without larger payoffs and steeper EV functions. Uhlela (1966) also reported nonoptimal decision criteria in an experiment requiring recognition of tilted lines. Again, the subjects tended toward an error minimizing strategy.

Similar findings have been reported for other tasks by Uhlela and Schmaltz (1966) and by Lieblich and Lieblich (1969a, 1969b).

Virtually all of these experiments combine skewed payoff matrices with quite difficult discriminations. This produces a probability distribution over ξ that is peaked at .5. Skewed payoff matrices produce optimal decision rules such that p^* is well away from .5. As arguments above show, such procedures lead to very flat maxima in the region of p^* . This point implies an explanation of Uhlela and Schmaltz's (1966) findings that presentation rates (priors) had a larger effect on decision criteria than did costs and payoffs. Changing costs and payoffs only controls the location of p^* . But changing the prior both changes the location of the criterion (if it is defined in likelihood ratio or related ways) and changes the flatness of EV over other criteria in the region of p^* .

Criterion variability is very important to the interpretation of signal detection experiments. The formal arguments summarized in Table III suggested that abundant amounts of it should be present in typical experiments. Hammerton (1970), using a simulated signal detection task in which subjects made inferences about parameters of normal distributions, showed that they did not adopt stable criteria. Galanter and Holman (1967) also reported that subjects used decision strategies inconsistently when faced with different payoff matrices. Where they have been studied, the implications of Table III seem to be experimentally confirmed. We conclude that most values of d' reported in the literature are substantially depressed--sufficiently so that reported values of d' in most experiments should be regarded only as lenient lower bounds on true values.

Can this distressing picture be alleviated by increasing costs and payoffs? Some data report invariance of d' for different payoff structures and levels of motivation (for a summary see Green and Swets, 1966; for specific studies, see Swets and Sewall, 1967, and Lukaszewski and Elliott, 1962). On the other hand Watson and Clopton (1969) and Calfee (1970) found noticeable effects of costs and payoffs on the detection rate and on d' respectively.

No experiments have studied the effect of changes in costs and payoffs on the effects of proper scoring rules for probability estimation, so far. The scoring rule maxima are so flat that we are very pessimistic--yet the studies need to be done. In their absence, there seems to be little point in using proper scoring rules to improve probability estimates--except in the sense that the scoring rules are really instructions about the meaning of probability estimates, and so may have merits not dependent on their property of rewarding most the man who makes the most truthful estimate.

References

- Aczel, J. and Pfanzagl, J. Remarks on the measurement of subjective probability and information. Metrika, 1966, 11, 91-105.
- Calfee, R. Effects of payoff on detection in a symmetric auditory detection task. Perceptual and Motor Skills, 1970, 31, 895-901.
- Chinnis, J. The problem of insensitivity in the use of costs and payoffs in psychological research. Unpublished manuscript, Engineering Psychology Laboratory, University of Michigan, 1971.
- DeGroot, M. H. Optimal statistical decisions. New York: McGraw-Hill, 1970.
- Edwards, W. Reward probability, amount, and information as determiners of sequential two-alternative decisions. Journal of Experimental Psychology, 1956, 52, 177-188.
- Edwards, W. Costs and payoffs are instructions. Psychological Review, 1961, 68, 275-284.
- Galanter, E. and Holman, G. L. Some invariances of the isosensitivity function and their implications for the utility function of money. Journal of Experimental Psychology, 1967, 73, 333-339.
- Green, D. M. Psychoacoustics and detection theory. Journal of the Acoustical Society of America, 1960, 32. 1189-1203.
- Green, D. M. and Swets, J. A. Signal detection theory and psychophysics. New York: Wiley, 1966.
- Hammerton, M. An investigation into changes in decision criteria and other details of a decision making task. Psychonomic Science, 1970, 21, 203-204.
- Lieblich, A. and Lieblich, I. Arithmetical estimation under conditions of different payoff matrices. Psychonomic Science, 1969 (a), 14, 87-88.

- Lieblich, I. and Lieblich, A. Effects of costs and payoff matrices on arithmetic estimation tasks: an attempt to produce rationality. Perceptual and Motor Skills, 1969 (b), 29, 467-473.
- Lukaszewski, J. S. and Elliott, P. N. Auditory threshold as a function of forced choice techniques, feedback and motivation. Journal of the Acoustical Society of America, 1962, 34, 223-228.
- Murphy, A. H. and Winkler, R. L. Scoring rules in probability assessment evaluation. Acta Psychologica, 1970, 34, 273-286.
- Rapaport, A. and Burkheimer, G. J. Sequential decision making: descriptive models, sensitivity analysis, and numerical results. L. L. Thurstone Psychometric Laboratory Report No. 83, University of North Carolina, 1970.
- Rapaport, A. and Burkheimer, G. J. Models for deferred decision making. Journal of Mathematical Psychology, 1971, 8, 508-538.
- Savage, L. J. Elicitation of personal probabilities and expectations. Journal of the American Statistical Association, 1971, 66, 783-801.
- Swensson, R. G. The elusive trade-off: speed vs. accuracy in choice reaction tasks with continuous cost for time. Technical Report No. 13, Human Performance Center, University of Michigan, 1968.
- Swensson, R. G. and Edwards, W. Response strategies in a two-choice reaction task with a continuous cost for time. Journal of Experimental Psychology, 1971, 88, 67-81.
- Swets, J. A. and Sewall, S. T. Invariance of signal detectability over stages of practice and levels of motivation. Journal of Experimental Psychology, 1963, 66, 120-126.
- Swets, J. A., Tanner, W. P. Jr., and Birdsall, T. G. Decision processes in perception. Psychological Review, 1961, 68, 301-340.

- Uhlela, Z. J. Optimality of perceptual decision criteria. Journal of Experimental Psychology, 1966, 71, 564-569.
- Uhlela, Z. J. and Schmaltz, E. Probability and payoff effects on conceptual decision criteria. Unpublished report presented at the Bayesian Conference, University of Michigan, Ann Arbor, 1966.
- Watson, C. S. and Clopton, B. M. Motivational changes of auditory sensitivity in a simple detection task. Perception and Psychophysics, 1969, 5, 281-287.

Footnotes

¹This research was supported by the Advanced Research Projects Agency of the Department of Defense and was monitored by ONR under Contract No. N00014-67-A-0181-0045. We are grateful to Robyn M. Dawes, Paul Slovic, and Robert L. Winkler and all of the members of the Engineering Psychology Laboratory for their insightful comments, which all too often we obstinately ignored.

Figure Captions

- Fig. 1. Maximal EV as a function of the actual probability for a quadratic proper scoring rule.
- Fig. 2. EV as a function of probability for two actions.
- Fig. 3. EV as a function of the cutoff probability for two different payoff matrices.
- Fig. 4. REL as a function of the cutoff probability for two different payoff matrices.
- Fig. 5. Hypothetical conditional probability distributions of the random observation variable Y .
- Fig. 6. EV as a function of the decision criterion for a symmetric payoff matrix and for different values of d' .
- Fig. 7. REL as a function of the decision criterion for a symmetric payoff matrix and for different values of d' .
- Fig. 8. EV as a function of the upper bound of a probability distribution over the decision rule.

TABLE I.
PAYOFF MATRIX FOR A 3 × 3 STIMULUS
IDENTIFICATION EXPERIMENT

		STIMULI		
		1	2	3
RESPONSES	1	x_{11}	x_{21}	x_{31}
	2	x_{12}	x_{22}	x_{32}
	3	x_{13}	x_{23}	x_{33}

TABLE II.
TYPICAL PAYOFF MATRIX IN A TWO STATE,
TWO ACT DECISION PROBLEM

STIMULI

	S_1	S_2
d_1	a	c
d_2	b	d


RESPONSES

TABLE III.

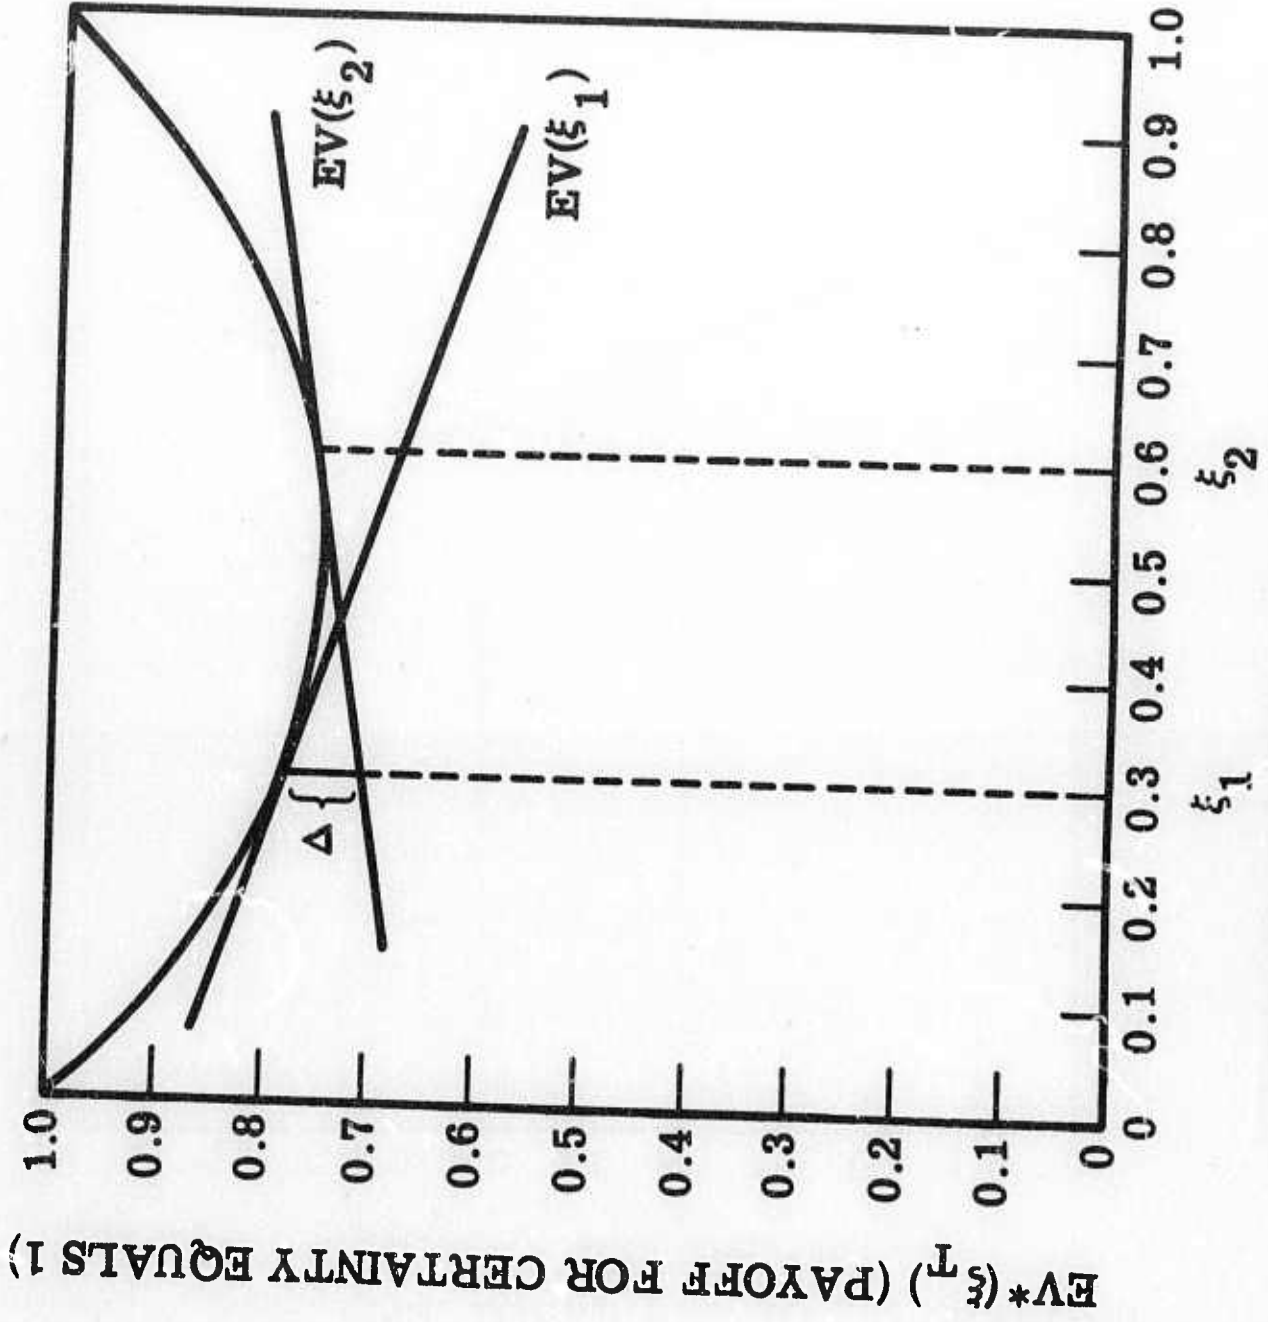
RELATIONSHIP BETWEEN VARIABILITY IN THE
 DECISION CRITERION, REDUCTION IN d' , EXPECTED
 VALUE, AND RELATIVE EXPECTED LOSS

SD(k)	Reduction in d' (%)	True $d' = 1$		True $d' = 2$		True $d' = 3$	
		EV (cents)	REL (%)	EV (cents)	REL (%)	EV (cents)	REL (%)
0.00	0	0.691	0.0	0.841	0.0	0.933	0.0
0.20	2	0.689	0.5	0.836	0.5	0.929	0.5
0.40	7	0.677	2.0	0.822	2.0	0.918	1.5
0.60	14	0.666	3.5	0.805	3.0	0.899	3.5
0.80	22	0.652	5.5	0.782	7.0	0.879	6.5
1.00	30	0.640	7.5	0.761	9.5	0.855	8.5
1.50	45	0.610	12.0	0.712	15.0	0.797	14.5
2.00	55	0.587	15.0	0.670	19.5	0.749	19.5
2.50	63	0.574	17.0	0.644	23.5	0.712	23.5
3.00	68	0.564	18.5	0.626	25.5	0.682	27.5


49

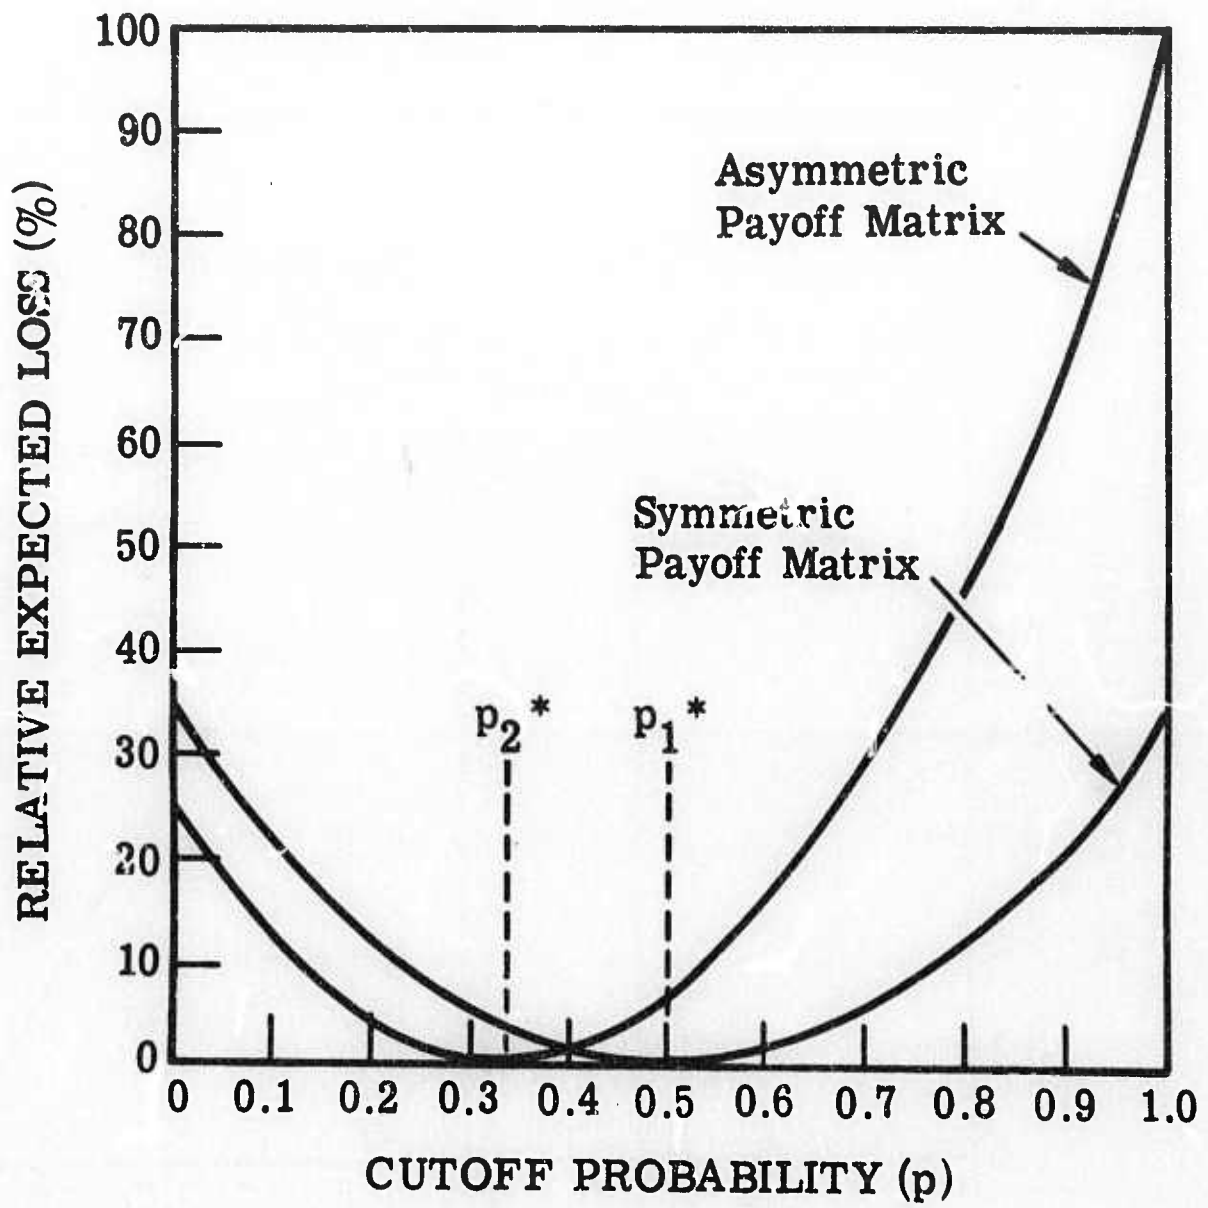



PROBABILITY OF THE HYPOTHESIS THAT MAKES d_1 THE APPROXIMATE RESPONSE (ξ)



PROBABILITY THAT STIMULUS IS OLD (ξ)








PROBABILITY OF Y

$f(Y|S)$

$f(Y|N)$


k k^*

Y



DECISION CRITERION (k)

54



DECISION CRITERION (k)

53

