

Accuracy, Information, and Response Time in a Saccadic Decision Task

B.A.J. Reddi, K. N. Asrress, and R.H.S. Carpenter

The Physiological Laboratory, University of Cambridge, CB2 3EG, United Kingdom

Submitted 16 August 2002; accepted in final form 25 May 2003

Reddi, B.A.J., K. N. Asrress, and R.H.S. Carpenter Accuracy, information, and response time in a saccadic decision task. *J Neurophysiol* 90: 3538–3546, 2003; 10.1152/jn.00689.2002. Reaction times generally follow a simple law economically described by the LATER model, in which a decision signal rises linearly in response to information about a target to a threshold at which a response is initiated, at a rate that varies from trial to trial with a Gaussian distribution. Functionally, LATER may be regarded as an ideal decision mechanism incorporating prior probability, information, and criterion level or urgency; this can be tested quantitatively by seeing whether LATER accurately predicts the effects on latency distributions of manipulating these variables: in this case, information and urgency. We presented subjects with random-dot kinematograms while fixating a central LED. The information content of the display was varied by altering the proportion of the dots moving coherently together either left or right rather than randomly. As soon as subjects detected the direction of coherent movement, they made a saccade in the same direction to one of a pair of LEDs on each side of the fixation target. Subjects responded either carefully, taking time to ensure an accurate judgement, or more hastily and with less regard for accuracy. The distributions of latencies under the different combinations of conditions were found to conform to LATER's predictions. Providing more information or increasing urgency both reduce latency; but they alter the observed distributions in different ways, equivalent to increasing the mean rate of rise on the one hand or reducing the criterion level on the other. Making only simple assumptions about the underlying mechanisms, the observed changes can be accounted for quantitatively.

INTRODUCTION

Reaction times are far longer than would be expected by considering synaptic delays and nerve conduction, reflecting the time taken for higher, cortical levels of the brain to make decisions about whether, and how, to respond to a stimulus (Carpenter 1999a; Schall 1995, 1997, 1999; Thompson et al. 1996). They also vary randomly from trial to trial, and though this variation is surprisingly large, in general, under high-contrast conditions, it follows a simple rule: the reciprocal of the reaction time or latency obeys a Gaussian distribution (Carpenter 1981, 1988). This can be most easily demonstrated by plotting cumulative histograms of such data using a dual transformation of the axes (Fig. 1). In a plot of this kind, the latency is represented on a reciprocal scale, reversed so that latency still increases to the right; rather as with the more familiar logarithmic scale, the interval between successive equal increments steadily decreases, but unlike a log scale the scale terminates at a point representing infinite time. Mean-

while, the cumulative probability is plotted on a probit scale, which has the useful property that a cumulative Gaussian is transformed to a straight line, whose intercept with the 50% axis is the median and whose slope is directly related to the SD. It is convenient to call a distribution whose reciprocal is Gaussian *recinormal* and to refer to a plot of this kind—in which a recinormal distribution becomes a straight line—as a *reciprobit plot*.

The LATER model (Linear Approach to Threshold with Ergodic Rate) provides a simple explanation for the recinormality of reaction times. A decision signal S rises linearly from an initial level S_0 in response to incoming information about a stimulus, until it reaches a criterion or threshold level S_T , at which point a response is initiated (Fig. 1). The distribution is explained by supposing that the rate of rise r varies randomly between trials in a Gaussian fashion with mean μ ; since the latency T_R will be proportional to $(S_T - S_0)/r$, this will make the distribution of T_R recinormal.

A response that lends itself particularly well to reaction time studies is the saccade, the eye movement made to fixate a suddenly presented target. It is fast and stereotyped, and with computer-operated equipment data can be gathered and analyzed very rapidly (Carpenter 1994a). Furthermore, we happen to have a relatively good understanding of the neural basis of the saccadic decision process: in particular, studies in behaving monkeys have revealed activity in neurons of the frontal eye fields that corresponds closely with the rise-to-threshold of LATER's decision signal (Hanes and Schall 1996), with a steady increase in firing rate in response to an appropriate stimulus, the rate achieved at the moment of initiation being relatively constant, but the rate of rise varying randomly from trial to trial.

An attractive interpretation of LATER in terms of decision theory is to identify S with the perceived probability of the existence of a situation demanding a response (Carpenter and Williams 1995; the underlying concepts have been particularly clearly presented by Gold and Shadlen 2001). S_0 then corresponds to the logarithm of the prior probability, μ with the rate of arrival of information, and S_T with a criterion at which a response is initiated—equivalent to a statistical significance level in conventional statistics. LATER can thus be regarded as collecting information until there is enough belief in the hypothesis to justify a response. In real life there would be many competing hypotheses, each presumably with its own LATER unit, which can be imagined to be racing one another to reach their individual threshold first and thus initiate their particular

Address for reprints and other correspondence: R.H.S. Carpenter, The Physiological Laboratory, University of Cambridge, CB2 3EG, UK (E-mail: rhsc1@cam.ac.uk).

The costs of publication of this article were defrayed in part by the payment of page charges. The article must therefore be hereby marked "advertisement" in accordance with 18 U.S.C. Section 1734 solely to indicate this fact.

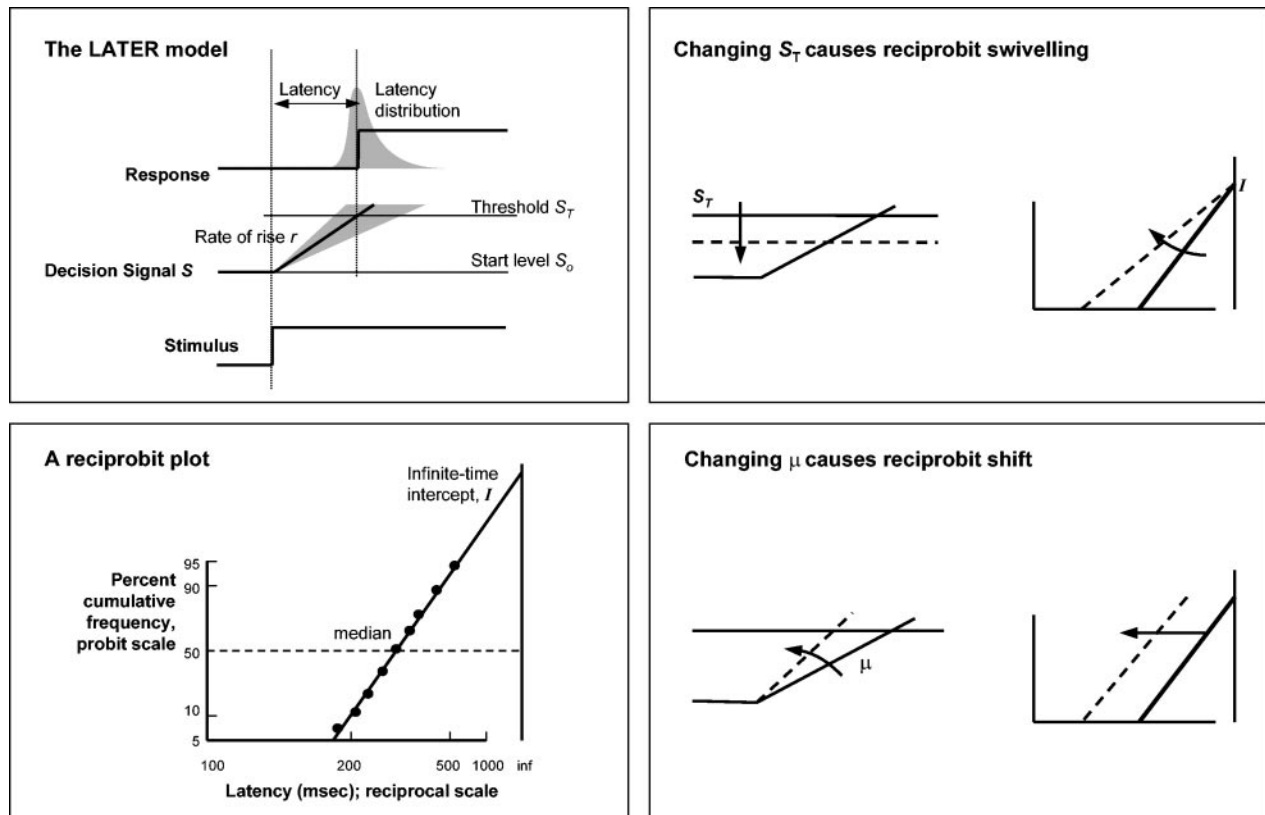


FIG. 1. *Left*: the LATER model. On presentation of a stimulus, a decision signal S rises linearly from an initial level S_0 at a rate r . On reaching the threshold S_T , a response is initiated. On different trials, r varies randomly about a mean μ in a Gaussian manner, such that the distribution of latency is skewed (shaded area) and the reciprocal of latency is distributed as a Gaussian. Hence, plotting the cumulative distribution on a probit scale as a function of the reciprocal of latency (a *reciprob plot*) yields a straight line (*bottom*). This line will intersect the $P = 50\%$ ordinate at the median latency and the $t = \infty$ axis at a point I . Changes in parameters of the LATER model have distinctive effects (*right*). Altering S_0 or S_T causes the line to swivel about the fixed point I on the $t = \infty$ axis, whereas alteration in μ shifts the line horizontally without change of slope.

response. A simple case is that of a pair of units voting for and against a hypothesis, where the log likelihood relationship arises particularly naturally (Gold and Shadlen 2001). Experiments using a precedence task (in which subjects are presented with a pair of targets, with a small temporal offset, and free to choose between them) appear to confirm the general applicability of LATER to such competitive situations (Leach and Carpenter 2001). In addition, recent neurophysiological studies in monkeys have provided supporting though essentially qualitative evidence with regard both to prior probability and to the supply of information: in particular, activity in saccade-related neurons in superior colliculus appears to reflect the likelihood of stimuli before they occur, and indeed their activity varies approximately logarithmically with the prior likelihood (Basso and Wurtz 1997; Dorris et al. 1999), while neurons both in frontal eye fields and in parietal regions show activity that accumulates as a function of the sensory information supplied by a random-dot kinetogram stimulus (Gold and Shadlen 2000; Kim and Shadlen 1999; Shadlen and Newsome 1996) similar to that used in this experiment. But it is a paradox of neurophysiological recording that, while it is often regarded as providing more exact information than behavioral experiments, the functional role of the neurons one is recording from can never be certain. In particular, it is hard to be sure whether they determine decisions or merely reflect them.

Behavioral experiments, such as we describe here, on the

other hand, directly address the model itself, and furthermore are capable of providing a quantity (and therefore in statistical terms, a quality) of data that is relatively difficult to achieve in animal experiments that allow one to distinguish between models that otherwise behave indistinguishably when tested with smaller populations. One example is the diffusion model (Ratcliff 1978; Ratcliff et al. 1999), which like LATER and indeed some other, much earlier, suggestions (for instance Laming 1968; Stone 1960), also envisages a signal embodying the accumulation of information, rising to a threshold. A recent improvement is to add "leakage," so that information decays as well as accumulates (Usher and McClelland 2001). With relatively modest data sets the parameters of such models—which may be quite numerous (six in the case of the leaky accumulator)—can often be adjusted to make them effectively indistinguishable in terms of goodness of fit, but with easily detectable stimuli and large data sets, this is not the case; that is not to say that they may not be a correct description of a different stage of the detection/decision process, which comes into prominence when signal-to-noise ratios are low (see the brief discussion in Ratcliff et al. 2001).

To test LATER thoroughly we need to assess quantitatively the effects of the three fundamental decision variables: prior probability, information supply, and criterion. In the case of both prior probability (Carpenter and Williams 1995) and decision criterion (Reddi and Carpenter 2000), this has already

been done: the model predicts with some accuracy not only how these factors affect the mean reaction time, but—more demandingly—how they affect the shape of the statistical distribution. In this present paper, we attempt to complete the threefold quantitative behavioral testing of LATER by analyzing the effects on reaction time distributions of altering the rate at which information is supplied; we also use a new paradigm to confirm the behavioral effects of altering the urgency of a decision.

METHODS

Five subjects participated with informed consent in these experiments, which had local ethical committee approval. They sat in a darkened room, their heads supported by chin- and headrests, at a distance of 0.5 m from a computer monitor screen subtending 27° horizontally and 22.6° vertically. The display was a modified version of that used by Shadlen and Newsome (1996): the central area was covered by a black rectangular mask 18° wide and 5.4° high, with a central yellow fixation LED flanked on each side at a horizontal distance of 7.2° by two target LEDs (all of luminance 8.1 cd m⁻²). A pattern of 1000 moving white dots could be displayed on the screen, each dot subtending 2 min arc and having a luminance of 1.2 cd m⁻². A predetermined proportion a (8, 16, 32, or 64%) of the dots moved coherently together at 1° s⁻¹ horizontally either to the right or to the left (the direction being chosen randomly for each trial); the others moved horizontally with the same instantaneous velocity of 1° s⁻¹, but for each individual dot the direction varied randomly and independently, with 50% probability, creating an approximation to a horizontal random walk.

Each trial began with the dots turned off. After a brief warning tone, instructing the subject to fixate the central fixation spot, there was a foreperiod randomly and uniformly distributed between 0.5 and 1.5 s, after which the moving dots were turned on; the subject's task was then to make a saccade either to the left or right fixed target, depending on the direction of motion of the coherently moving dots. Forty milliseconds after detection of this saccade, the dots were turned off, and 50 ms later the next trial began. Subjects undertook at least four runs of 100 such trials for every combination of stimulus conditions, with stimuli interleaved, and taking rests as required. They were told to aim broadly for a particular degree of accuracy, and after each run were told the percentage of errors they had in fact made. All subjects found they could modify their speed and accuracy such that error rates within a range of ±4% were obtained in most runs, and these formed the data sets used for subsequent analysis. Each subject provided a set of data for each coherence with at least two different standards of accuracy.

Eye movements were monitored with an infrared oculometer (Carpenter 1988) and recorded and analyzed with the SPIC computer-based instrumentation system (Carpenter 1994a), which also controlled the presentation of stimuli. Saccades were detected in real-time and latencies were assigned to 10-ms bins. After experimental runs, all records were checked manually to eliminate any errors caused by blinks, head movements, or other artifacts.

RESULTS

Reaction times were grouped according to which of the two parameters (information or urgency) was varied. Cumulative histograms of latencies were plotted as reciprobital plots, and examples of sets of plots for different subjects are shown in Fig. 2 (varying information) and Fig. 3 (varying urgency). As described in the Introduction, if the reciprocal of latency follows a Gaussian distribution (as predicted by LATER), the plot should yield a straight line (see Fig. 1). Anticipatory responses

then appear as a secondary distribution deviating systematically from the main distribution (see for example subject K in Fig. 3). Although these deviations can appear substantial in reciprobital plots, because the probit scale greatly stretches out the extreme tails of the distribution, since they constitute only a small percentage of the total, they have only a negligible effect on the fitting procedure (see for example Ratcliff 1993). The same is true of the deviations that can sometimes be seen at the long-latency end, which appear to be the result either of lapses of attention in individual trials, or perhaps, with very long latencies, of the subject's impatience (e.g., subject R in Fig. 2). In fact, all distributions conformed to the LATER model (Kolmogorov–Smirnov one-sample test, $P > 0.1$) (Kolmogorov 1941).

Within each group, the distributions were analyzed to see how well they conformed to the alternative hypotheses that differences in the distributions were due either to an alteration in the mean rate of rise in the LATER model (μ) or in the criterion or threshold, S_T . With a change in μ , the effect on the reciprobital plot will be to move the line horizontally without change in slope (“shift”), whereas in the latter case we expect “swivel”: both the median and slope will alter, in such a way that the intercept I of the line with the infinite-time axis remains fixed (Fig. 1). We tested for shift by finding the best (maximum likelihood) straight-line fit for each distribution, within the constraint of a common value of the slope. In the latter case, testing for swivelling of the distributions, a similar attempt was made, but with the constraint of a common point of intersection I on the infinite-time axis. These are the maximum-likelihood best-fit lines shown in Figs. 2 and 3.

Constant urgency with varying information

Figure 2 shows reciprobital plots of latency distributions when subjects are asked to respond with constant accuracy to varying levels of display coherence. Increased coherence corresponds to increased provision of information, so LATER predicts that the mean rate of rise (μ) of the decision signal should increase, reflected in a parallel shift of the distribution with reduced median but fixed slope. This shift is apparent in the figure; analysis using the Kolmogorov–Smirnov statistic confirmed ($P > 0.1$) that all sets of distributions were compatible with parallel shift. For each subject, an attempt was made to fit the distributions under the alternative constraint of swivel rather than shift, and the total relative support (log likelihood) across all conditions was calculated: these values are shown in Fig. 4. For three of five subjects, shift was very significantly more strongly supported than swivel, while for the other two, the degree of support did not differ significantly. This suggests that altering the amount of available information alters μ , the rate of rise, rather than the threshold, S_T .

Varying urgency with constant information

Figure 3 shows reciprobital plots of latency distributions under conditions in which the coherence (and thus the amount of information) was constant, but the speed and consequently the accuracy of responding were varied. If, as LATER predicts, such alterations in urgency result in a change in S_T , the criterion for making a response, then the reciprobital line should swivel about I , the point where it intercepts the $t = \infty$ axis (Reddi and Carpenter




FIG. 2. Effect of varying information on the distribution of latency. Reciprobit plots are shown for four subjects (B, D, J, and R), the sets of lines corresponding to different degrees of coherence of the kinematograms, (\square 64%, \blacksquare 32%, \circ 16%, \bullet 8%: F, high-urgency-fast, inaccurate-conditions; M, moderate urgency; S, low-urgency-slow and accurate-conditions). The lines represent best linear fits (maximum likelihood) under the constraint of parallelity. Changes in coherence, and hence in the supply of information, would be expected to cause parallel shifts of the distributions rather than swivelling, corresponding to changes in the mean rate of rise, μ , in LATER.

2000). For each subject, lines were fitted using maximum likelihood under the constraint of a fixed I , and in Fig. 3 it can be seen that this prediction appears to be fulfilled. Analysis using the Kolmogorov-Smirnov test showed that, in every case, the differences between the distributions were compatible with swivelling about a fixed I ($P > 0.1$). Again, for each subject, maximum-likelihood best fits were made under the alternative constraint of shift rather than swivel: relative support values are shown in Fig. 4. For all of the five subjects, swivel was significantly more strongly supported than shift. Therefore, and as previously reported in a different kind of task (Reddi and Carpenter 2000), the distributional effect of changing urgency does seem to be as predicted by LATER, the result of a change in S_T , the response criterion level.

DISCUSSION

The LATER model was originally (Carpenter 1981) formulated as a parsimonious way of describing the observed distributions of reaction times in a variety of tasks, including saccades to visual targets. The notion was of a neural signal rising linearly in each trial from an initial level to a threshold that initiates movement, the rate varying in a Gaussian fashion from trial to trial. In a subsequent paper (Carpenter and Williams 1995), this purely descriptive representation was given a func-

tional interpretation by proposing that the instantaneous value of this signal coded for probability or likelihood, specifically on a logarithmic scale; this would then make some sense as a decision mechanism (though the gratuitous variability—at first sight—does not). When thus functionally interpreted in terms of decision-making, LATER makes a number of relatively precise quantitative predictions, and over the last few years our aim has been systematically to test these predictions to see whether the model could—in the Popperian sense—be falsified. More specifically, the points to be tested are as follows.

- 1) Does the starting level actually represent prior probability of the hypothesis?
- 2) Does the final level correspond to the urgency of the task?
- 3) Does the mean rate of rise correspond to the quantity of information supplied?
- 4) Is there evidence for competitive racing between different instantiations of the LATER model?
- 5) Does the rate of rise vary in a Gaussian fashion from trial to trial?

In this sense it is a vulnerable model, and any one of the tests could have turned out negative and invalidated the model; but so far, at least, they have not. Considering these points in the order in which they have been examined:

The Gaussian variability of the rate of rise (point 5) was the original observation that prompted the formulation of the




FIG. 3. Effect of urgency instructions, and therefore accuracy, on the distribution of reaction time. Reciprob plots are shown for different subjects (B, K, J, and R), viewing random-dot kinematograms with different values of coherence (64, 32, 16, and 8%). The lines represent best linear fits (maximum likelihood) under the constraint of a common infinite-time intercept I . For subjects B and K, there are two different degrees of urgency: fast (F, with an increased numbers of errors, as shown) or with greater accuracy and therefore slower (S); these different conditions appear to make the distributions swivel about a fixed point on the infinite-time axis, corresponding to a change in the threshold S_T . For subjects J and R there is in addition an intermediate degree of urgency (M).

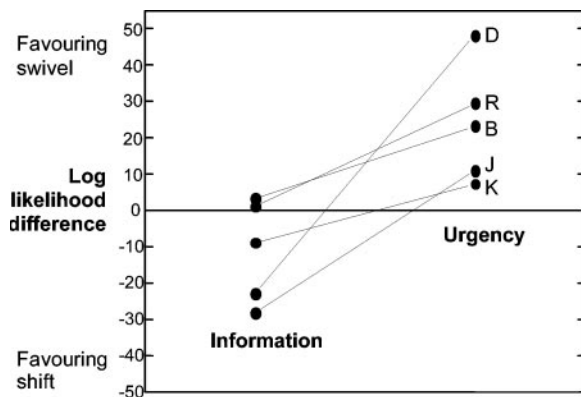


FIG. 4. Likelihood analysis of swivel and shift. Each set of data, corresponding either to different degrees of coherence (information) or different degrees of urgency, were fitted with distributions either having a common intercept (swivel) or having a common slope (shift). The difference in the log likelihood (support) for (swivel versus shift) is shown for each subject and for each data set.

model and is supported by extensive analysis of previously published reaction time distributions that demonstrate that a Gaussian variation of the reciprocal of reaction time provides a better fit than other models that have been proposed, at least under conditions in which the prior stage of signal detection does not contribute significantly to variability (in other words, with high stimulus signal-to-noise ratios). A few examples of analysis of existing data of this kind have already been published (Carpenter 1981, 1988, 1999b); a full set of historical data has been analyzed and submitted for publication, together with comparisons with alternative models. (Functionally, this random variability is a surprising feature, at first sight difficult to reconcile with the idea of an ideal decision maker. It may perhaps be regarded as a “deliberate” feature that has evolved to make behavior less predictable (Carpenter 1999b): by randomizing rate of rise in competing decision units we are in effect randomizing choice, and it is a commonplace both of game theory and of behavioral observation that unpredictability of response is an essential feature of predator-prey relationships.)

The functional significance of the starting level S_0 (point 1) was dealt with in the paper referred to above (Carpenter and Williams 1995), while urgency (point 2) was addressed by Reddi and Carpenter (2000), where it was shown that instructions to subjects, causing them to be either more careful not to make mistakes or to respond as rapidly as possible and to pay less attention to errors, resulted in the swivelling of the reciprocal plots predicted if there is a change of the final threshold criterion. The implied competitive racing by LATER units acting in parallel (point 4) has been addressed partly by a study of countermanding (Hanes and Carpenter 1999), with a data set large enough to establish that LATER could predict not merely mean reaction times and percentages of errors, but the shapes of the distributions as well, and also more directly in the precedence task mentioned earlier (Leach and Carpenter 2001), again with quite satisfactory results.

The remaining point of vulnerability is therefore (point 3), which is what this paper is primarily intended to address. It is perhaps necessary to emphasize that our aim was to do no more than to discover whether this one particular prediction of LATER was correct or incorrect. We do not assert that other possible models of reaction time could not be made to explain these data as well: all we claim here is that—in an experiment that *might* have shown LATER to be completely wrong—the behavior actually observed is consistent with what is predicted. As we have seen, the observations do in fact appear to support the predictions that the changes in the distributions of reaction time in this experiment are qualitatively as would be expected in the LATER model to be the result of alterations in the level of urgency on the one hand and the supply of information on the other. However, in an earlier study of the effect of prior probability (Carpenter and Williams 1995), we were able to show not only that the shapes of the latency distributions altered in the way predicted by LATER, but also that the *magnitudes* of the changes in median latency were what would be expected from theoretical considerations. Can we demonstrate the same for the present data?

Urgency/threshold

In the earlier paper (Carpenter and Williams 1995), the changes in prior probability that gave rise to the observed alterations in reaction time distributions were easy to quantify, since they depended on the proportions of trials on which one or another stimulus appeared, which could be absolutely determined by the experimenter. In the case of urgency, no such quantitative specification is of course possible; we cannot get inside the subject's head to reset the criterion to a different level, so it must be estimated, posthoc, by what must necessarily be a more uncertain process of inference. One possible approach is to make use of the well-known trade-off between accuracy and decision time. When told to perform more quickly, and to sacrifice accuracy for speed, our subjects did indeed make more errors, and the prevalence of those errors could in principle provide some information about how the LATER criterion level was altered. In the functional interpretation of LATER, the decision signal S represents a measure of probability or likelihood: more specifically, it appears to embody the logarithm of the likelihood of a particular hypothesis, for instance, about the existence of a stimulus. Decisions then amount to making comparisons between these likelihoods, for

instance, by estimating log likelihood ratios, or log odds. But when, as normally in the natural world, there are very many possible stimuli, rather than making exhaustive comparisons of every possible pair of stimuli, if all that is required is a selection of the best supported hypothesis, it makes sense—as the LATER model suggests—to run races between signals representing the likelihoods of the various hypotheses, the winner then (leaving aside the gratuitous random element) being the hypothesis with the greatest likelihood. But where, as here, there are just two mutually exclusive hypotheses, say R and L, any sensory information tending to support L must necessarily tend to discount R and vice versa. It would therefore make functional sense to combine the two sources of information and to compute the log likelihood ratio for R versus L. Neurally, this is quite plausible: one of the advantages of using a log scale for likelihood is that the log likelihood of a conjunction such as (R and \sim L) is the sum of the individual log likelihoods, so that the calculation of the likelihoods for composite hypotheses amounts simply to processes of linear summation or inhibition, as has been illustrated particularly clearly by Gold and Shadlen (2001).

In LATER, the median time T_R for S to rise from the initial level S_0 to the threshold S_T is given by

$$T_R = (S_T - S_0)/\mu \quad (1)$$

What does this threshold level represent? Given that S is a measure of probability, the threshold must also represent some kind of probability. If we interpret S as being a log likelihood for R versus L, then S_T represents a criterion log likelihood ratio, analogous to a significance level in statistical evaluation. As in statistical significance tests, it reflects the proportion of occasions that the criterion would be expected to be exceeded when the hypothesis was in fact false. More precisely, S_T should be equal to log of the odds $Q = p/(100 - p)$ of the percentages of correct versus incorrect responses, so that the median time T_R to reach threshold should be given by

$$T_R = (\log Q - S_0)/\mu \quad (2)$$

S_0 represents the logarithm of the prior likelihood ratio, L_0 . But since in this experiment the targets were presented on each side with equal probability, $L_0 = 1$; therefore, $S_0 = \log L_0 = 0$. Therefore the time of rise simplifies to

$$T_R = \log Q/\mu \quad (3)$$

Equation 3 implies a trade-off between accuracy and speed, a phenomenon that has of course often been observed in reaction time or memory retrieval tasks. But it makes the more specific prediction that the increase in latency with increased accuracy should be proportional to the log odds for correct versus incorrect responses. Not many studies have generated appropriate data with which this quantitative prediction might be tested. However, in one such study of latencies using a more complex visual task (Thorpe et al. 1996), the authors noted in passing a similar trade-off between accuracy and latency and tested for a linear relation between the two, yielding a value of $r = 0.62$. However, if their latency data are plotted as a function of $\log Q$, it yields a better fit, with $r = 0.80$.

Information/rate of rise

Can we similarly attempt to make a quantitative prediction of the additional delays expected in this task when the infor-

mation is provided at a lower rate? This implies estimating the rate at which information is supplied by the coherent movement of a proportion a of a set of otherwise incoherently moving dots. Notoriously, quantitative assessment of measures of information in particular cases is often ambiguous and sometimes controversial, and previous attempts to estimate the information content of random-dot kinematograms have come up with different answers: some reasons why real performance may not correspond to the ideal limit have been intelligently discussed by Barlow and Tripathy (1997). The following approach is offered as being relatively simple and direct, with minimal assumptions, specifically addressing the particular display that we used in these experiments and using a local likelihood approach because this matches what we believe to be the neural basis of the underlying decision process (local likelihoods have been similarly used by Weiss and Adelson 1998). It does not claim to embody what is known of the actual neurophysiology of motion-detecting neurons in the visual system but rather to represent an ideal system, which one might reasonably hope such neurons would have evolved to emulate.

Let L be the initial ratio of the likelihoods of the hypotheses H_R and H_L that the pattern as a whole is moving to the right or to the left. Then, by Bayes' theorem, observation E of one dot moving to the right will provide evidence favoring H_R , increasing the ratio to L' , where

$$L' = L \frac{p(E|H_R)}{p(E|H_L)} = LC \quad (4)$$

Here the factor $C = p(E|H_R)/p(E|H_L)$ represents the odds of making the observation under the two hypotheses; thus $\log L$ will be increased by $\log C = \log[p(E|H_R)/p(E|H_L)]$, a quantity—"support"—that is also a measure of the information that the observation has yielded (Edwards 1972).

If out of a total of N dots, a proportion a are moving coherently to the right (hypothesis H_R), then, under the conditions of this experiment, half of the remaining $(N - aN)$ will be moving to the right at any moment and half to the left. Consequently the probability that any dot chosen at random is at some moment moving to the right is given by

$$p(E|H_R) = \frac{aN + 0.5(N - aN)}{N} = 0.5(1 + a) \quad (5)$$

Similarly, if the proportion a are moving to the left (hypothesis H_L)

$$p(E|H_L) = \frac{0.5(N - aN)}{N} = 0.5(1 - a) \quad (6)$$

The ratio $C = (p(E|H_R)/p(E|H_L))$ representing the contribution of any one dot will then be simply

$$C = \frac{1 + a}{1 - a'} \quad (7)$$

(a ratio that can in a sense be regarded as a measure of velocity contrast). Thus, for the N dots constituting the pattern as a whole, and considering the logs of the likelihoods

$$\log L' = \log L + N \log C \quad (8)$$

This implies that the rate of increase of $\log L$, which is also the rate of arrival of information, should be proportional to $\log C$. Therefore if we identify rate of arrival of information with the

mean rate μ with which S rises to threshold, as the LATER model implies, μ should be proportional to $\log[(1 + a)/(1 - a)]$. Such a relationship is consonant with recent neurophysiological studies relating the activity of neurons in frontal cortex to the quantity of information provided, for instance, by withholding part of the information needed for a conjunction visual search task (Bichot and Schall 1999) or by varying the coherence of a random-dot kinematogram (Gold and Shadlen 2000; Kim and Shadlen 1999; Roitman and Shadlen 2002). In the latter case, using changes in saccade metrics rather than latency, the authors' conclusion, in accordance with the LATER model, is that information is linearly accumulated. Although Gold and Shadlen were able to describe this accumulation with a precise mathematical model, it requires rather more free parameters than LATER, and the underlying equation is of a generalized form, including power functions, highly suitable for empirical fitting to data but not intuitively related to the fundamental determinants—expectancy, urgency and information—of such tasks nor of course the result of trying to start from first principles.

Together, the two predictions—concerning urgency and information—yield a formula for the reaction time T in this task for any combination of them

$$T = T_0 + k \frac{\log Q}{\log C} \quad (9)$$

Here $Q = p/(100 - p)$, $C = (1 + a)/(1 - a)$, k is a constant of proportionality, and T_0 is a constant delay term, representing those aspects of the whole process, such as transduction and conduction and the initial detection of motion, that are invariant in this experiment, and which, like k , depend on the individual subject. T_0 in fact conceals the first of the two stages (LATER being the second) that constitute the delay between stimulus and response. This first stage, probably representing mechanisms of local feature detection rather than decision about the stimulus as a whole, introduces a degree of random variability of its own that can be pronounced under conditions of low signal-to-noise, follows a different distributional law, and may introduce errors of a different type (Ratcliff 1978, 1987; Ratcliff et al. 2001; Reddi 2001). But under high-contrast conditions such as these, the extra variability associated with this first stage is small in comparison with what is done by the second; this corresponds to what can be observed in the stimulus- and saccade-related activity of single neurons in the frontal eye field of macaques (Thompson et al. 1996).

ANOVA confirmed that for each subject *Eq. 9* does indeed predict latencies. Table 1 shows the best fit (minimized residual sum of squares) values of k and T_0 for each subject,

TABLE 1. Individual parameter values

Subject	k , ms	T_0 , ms	F	n	P
B	16	316	6.98	8	0.04
R	35	284	72.9	12	<0.0001
K	42	337	8.37	8	0.03
D	27	315	114	8	<0.0001
J	16	286	39.7	12	0.0001

Values are the best-fit (minimized residual sum of squares) of the parameters k and T_0 in *Eq. 1* for each of the five subjects, with corresponding F values, number of data points, and significance

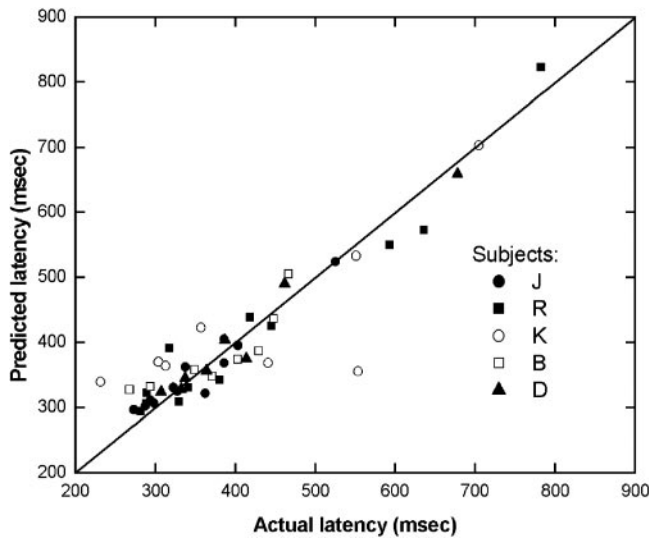


FIG. 5. For each of the 5 subjects, the predicted median latencies, using Eq. 9 and the individual parameter values shown in Table 1, are plotted as a function of the observed latencies for all combinations of conditions: perfect prediction would give the straight line shown.

together with the corresponding F values and significance levels. Putting these parameters into Eq. 9, we can compare the predicted values of median reaction for each subject under all combinations of conditions (Fig. 5): for this complete data set, linear regression analysis yields $r = 0.925$, $P < 0.0001$; a runs test shows no significant deviation from linearity.

Error behavior

Ideally, the LATER model should be able to account for errors and their observed latency distributions as well as correct responses. Unfortunately this is not a straightforward matter, as there are several ways in which errors might occur, and a corresponding range of modifications to the model that could be made to account for them. Errors might, and in many natural circumstances presumably do, occur in the detection stage that precedes the decision mechanism embodied in LATER. They may also occur as a result of competition between regular LATER units acting in parallel with “rogue”

units whose activity is loosely coupled, or not coupled at all, to the detection stage (Fig. 6). An arrangement of this kind has been successful in modeling the early component of the distribution of intersaccadic intervals seen in optokinetic nystagmus (Carpenter 1994b), and it seemed worth seeing whether it might work here as well. The mean rate of rise μ of the “regular” unit, corresponding to the correct direction, was allowed to vary as a function of the coherence of the stimulus, as observed by Shadlen and Newsome (2001) in parietal cortex, and more recently by Roitman and Shadlen (2002) using similar stimuli. (The latter study also demonstrated units with decreasing activity in response to the “wrong” direction, but we did not attempt to incorporate this feature as we wished to keep the number of free parameters as small as possible, and—with activity declining rather than increasing—their influence was likely to be slight.) We used a pair of competing “rogue” units, starting closer to the threshold but with a slower rate of rise than the regular units, indicating a weaker degree of informative support from the stimulus (perhaps corresponding to the fact that there is a common context to each trial that is independent of the direction and coherence of motion). One of these units corresponded to the correct response and one to the incorrect response, and their parameters (threshold S'_T and mean rate of rise μ') were identical. μ' was constrained to be constant for each subject, independent of urgency and information. Within that constraint on μ' , values of S'_T and μ' were then found that provided the best fit for each subject both for the distribution of error response latencies and for the percentage of errors occurring. Of the 44 individual fits that were thus attempted, all but 5 generated distributions that were not significantly different from the observed data (Kolmogorov–Smirnov, $P = 0.05$). Most subjects, under most conditions, generated error responses whose distribution was not markedly different from the correct responses, and generally slightly faster; this is in contrast to some other comparable studies (Ratcliff and Rouder 2000; Roitman and Shadlen 2002) but is not a necessary feature of our model, which can generate errors that are slightly slower or slightly faster, depending on the settings of the parameters. The resultant predictions of error rates, and some examples of predicted and observed error latency distributions, are shown in Fig. 6; applying linear

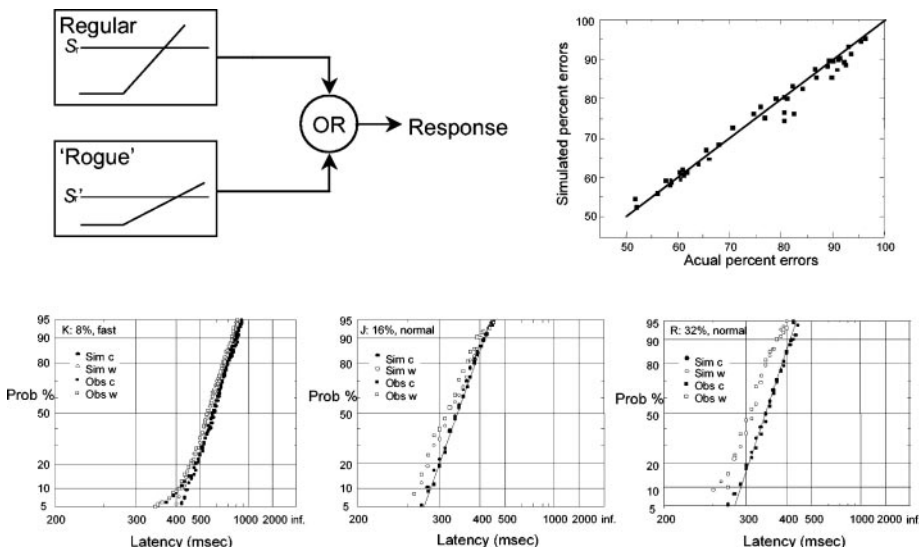


FIG. 6. Analysis of errors. *Above left*: a rogue unit in parallel with a regular unit, having a lower threshold and shallower rate of rise; on some trials a rogue unit may reach threshold before the regular unit, which may then cause an incorrect response. *Top right*: performance of a model of this kind in predicting error rates: simulated percentage errors are plotted as a function of actual percentage errors, for all subjects and conditions combined. *Bottom*: 3 examples of correct and error latency distributions and their simulations with the model, for 3 different subjects and the conditions stated. Filled symbols, correct responses; open symbols, wrong responses; squares, actual observations; circles, simulated.

regression analysis to this complete set of error rates gives $r = 0.991$, $P < 0.0001$, and a runs test shows no significant deviation from linearity.

While this attempt at simulation shows that a configuration of LATER units can provide an explanation for error rates and distributions, we need to emphasize that the particular model we used represents simply an attempt to show that such simulation is possible, rather than the outcome of a systematic search for optimum behavior or functional plausibility—for instance, as noted above, by incorporating competing regular units as well as rogue units, behaving as described by Roitman and Shadlen (2002). There are many variations of this general type that could be explored, and it is highly probable that a more systematic study would yield one that performs better, and this we hope to do. Nor is it generally agreed how best to model the preliminary stage of motion detection (see for example Barlow and Tripathy 1997). But for the moment we are encouraged that—under a variety of conditions—the quantitative details of several aspects of human decision making can be predicted using a simple model with relatively few free parameters.

We thank our subjects for patience in undertaking these lengthy experiments.

DISCLOSURES

This work was supported by a grant to R.H.S. Carpenter from the Wellcome Trust.

REFERENCES

- Barlow H and Tripathy SP.** Correspondence noise and signal pooling in the detection of coherent visual motion. *J Neurosci* 17: 7954–7966, 1997.
- Basso MA and Wurtz RH.** Modulation of neural activity by target uncertainty. *Nature* 389: 66–69, 1997.
- Bichot NP and Schall JD.** Saccade target selection in macaque during feature and conjunction visual search. *Vis Neurosci* 16: 81–89, 1999.
- Carpenter RHS.** Oculomotor Procrastination. In: *Eye Movements: Cognition and Visual Perception*, edited by D. F. Fisher, R. A. Monty, and J. W. Senders. Hillsdale, NJ: Lawrence Erlbaum, 1981, p. 237–246.
- Carpenter RHS.** *Movements of the Eyes*. London: Pion, 1988.
- Carpenter RHS.** SPIC: a PC-based system for rapid measurement of saccadic responses. *J Physiol* 480: 4P, 1994a.
- Carpenter RHS.** Express optokinetic nystagmus. In: *Contemporary Ocular Motor and Vestibular Research*, edited by A. F. Fuchs, T. Brandt, U. Büttner, and D. Zee. Stuttgart, Germany: Georg Thieme, 1994b, p. 185–187.
- Carpenter RHS.** Visual selection; neurons that make up their minds. *Curr Biol* 9: 595–598, 1999a.
- Carpenter RHS.** A neural mechanism that randomises behaviour. *J Consoc Stud* 6: 13–22, 1999b.
- Carpenter RHS and Williams MLL.** Neural computation of log likelihood in the control of saccadic eye movements. *Nature* 377: 59–62, 1995.
- Dorris MC, Taylor TL, Munoz DP, and Klein RM.** Saccadic reaction times are influenced similarly by previous saccadic metrics and exogenous cuing in monkey. *J Neurophysiol* 81: 2429–2436, 1999.
- Edwards AWF.** *Likelihood*. Cambridge, UK: Cambridge, 1972.
- Gold JL and Shadlen MN.** Representation of a neural decision in developing oculomotor commands. *Nature* 404: 390–394, 2000.
- Gold JI and Shadlen MN.** Neural computations that underlie decisions about sensory stimuli. *Trends Cognit Sci* 5: 10–16, 2001.
- Hanes DP and Carpenter RHS.** Countermanding saccades in humans. *Vis Res* 39: 2777–2791, 1999.
- Hanes DP and Schall JD.** Neural control of voluntary movement initiation. *Science* 274: 427–430, 1996.
- Kim J-N and Shadlen MN.** Neural correlates of a decision in the dorsolateral prefrontal cortex of the macaque. *Nat Neurosci* 2: 176–185, 1999.
- Kolmogorov A.** Confidence limits for an unknown distribution function. *Ann Math Stat* 23: 525–540, 1941.
- Laming D.** *Information Theory of Choice Reaction Times*. New York: Academic, 1968.
- Leach JCD and Carpenter RHS.** Saccadic choice with asynchronous targets: evidence for independent randomisation. *Vis Res* 41: 3437–3445, 2001.
- Ratcliff R.** A theory of memory retrieval. *Psychol Rev* 85: 59–108, 1978.
- Ratcliff R.** More on the speed and accuracy of positive and negative responses. *Psychol Rev* 94: 277–280, 1987.
- Ratcliff R.** Methods for dealing with reaction time outliers. *Psychol Bull* 114: 510–532, 1993.
- Ratcliff R and Rouder JN.** A diffusion model account of masking in two-choice letter identification. *J Exp Psychol Hum Percept Perform* 26: 127–140, 2000.
- Ratcliff R, Carpenter RHS, and Reddi BAJ.** Putting noise into neurophysiological models of simple decision making. *Nat Neurosci* 4: 336–337, 2001.
- Ratcliff R, van Zandt T, and McKoon G.** Connectionist and diffusion models of reaction time. *Psychol Rev* 106: 261–300, 1999.
- Reddi BAJ.** Decision making: the two stages of neuronal judgement. *Curr Biol* 11: 603–606, 2001.
- Reddi BAJ and Carpenter RHS.** The influence of urgency on decision time. *Nat Neurosci* 3: 827–831, 2000.
- Roitman JD and Shadlen MN.** Response of neurons in the lateral intraparietal area during a combined visual discrimination reaction time task. *J Neurosci* 22: 9475–9489, 2002.
- Schall JD.** Neural basis of saccade target selection. *Rev Neurosci* 6: 63–85, 1995.
- Schall JD.** Visuomotor areas of the frontal lobe. *Cereb Cortex* 12: 527–538, 1997.
- Schall JD.** Weighing the evidence: how the brain makes a decision. *Nat Neurosci* 2: 108–109, 1999.
- Shadlen MN and Newsome WT.** Motion perception: seeing and deciding. *Proc Natl Acad Sci USA* 93: 628–633, 1996.
- Shadlen MN and Newsome WT.** Neural basis of a perceptual decision on the parietal cortex (area LIP) of the Rhesus monkey. *J Neurophysiol* 86: 1916–1936, 2001.
- Stone M.** Models for choice reaction time. *Psychometrika* 25: 251–260, 1960.
- Thompson KG, Hanes DP, Bichot NP, and Schall JD.** Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. *J Neurophysiol* 76: 4040–4055, 1996.
- Thorpe S, Fize D, and Marlot C.** Speed of processing in the human visual system. *Nature* 381: 520–552, 1996.
- Usher M and McClelland JL.** The time course of perceptual choice: the leaky accumulator model. *Psychol Rev* 108: 550–592, 2001.
- Weiss Y and Adelson EH.** Slow and smooth: a Bayesian theory for the combination of local motion signals in human vision. *MIT AI Lab Memo AI Memo* 1624, 1998.