Biocybernetics and Biomedical Engineering
2012, Volume 32, Number 2, pp. 49–63

Analysing the Detail of Saccadic Reaction Time Distributions

R. H. S. CARPENTER*

Department of Physiology, Development and Neuroscience, University of Cambridge, United Kingdom

Measuring saccadic reaction time distributions is an increasingly popular technique, making it possible to obtain a large amount of data non-invasively in a short period of time. Such distributions can often be encapsulated with just two parameters, representing the mean and variance of the rate of rise in the LATER model. For many purposes, both scientific and clinical, this is enough. But for normal as well as pathological subjects, particularly when using more complex tasks, one may often see features that cannot be explained by a simple LATER model. These include early and express saccades, 'late' saccades, and (in tasks such as go/no-go and antisaccades) more complex modifications. These features can be explained relatively easily by introducing extra LATER units, and enable them to be quantitatively parameterised; this potentially offers much more precise ways of quantifying the effects of such clinical conditions.

K e y w o r d s: reaction time, saccades, latency, LATER, antisaccades1. Introduction

1. Introduction

For those of us who have the privilege of teaching eager and intelligent medical students, a frequent cause of sadness is that a passion for neuroscience when they are undergraduates is by no means always followed by an equal passion for clinical neurology. Why is this? Partly because they soon perceive that it is a specialization where it is relatively difficult to offer much help to the patient, but also one where the pleasures of making deductions from intelligent, active observation are increasingly being sidelined by the beguiling simplicity of ordering a brain scan. Behind this trend is of course the depressing implication that ultimately it is structure that

* Correspondence to: R.H.S. Carpenter, Department of Physiology, Development and Neuroscience, The Physiological Laboratory, University of Cambridge, CB2 3EG, United Kingdom, e-mail: rhsc1@cam.ac.uk
Received 14 November 2011; accepted 06 February 2012

counts, not function, and that observation of behaviour is no more than a roundabout and inefficient way of pin-pointing a location in the brain. We are back, in fact, to the captivating naivety of phrenology: encouraged to think of 'functions' that are 'localised' in different regions of the brain, we are discouraged from asking too closely what either function or localization might actually mean. As Sherrington wrote [1], over 70 years ago, "Facts rebut the over-simplified conceptions such as to ascribe to separate small pieces of the roof-brain, wedged together like a jigsaw puzzle, separate items of highly integrated behaviour. A special place for comprehension of names, a special place for arithmetical calculation, a special place for musical appreciation, and so on. … Rather, we may think, the contributions which the roof-brain, in collaboration with the rest of the brain and spinal cord, makes toward integrated behaviour will, when they are ultimately analysed, resolve into components for which at present we have no names. To state the organization of the mind in terms of roof-brain activities is a desideratum not in sight". To put it another way, where *is not how*.

Why have other specializations, such as cardiology or endocrinology not fallen victim to a correspondingly simplistic and unscientific approach? An important factor is that to a very large extent neurology is a number-free zone. One cannot help being reminded of a technique popular with doctors in the middle ages: uroscopy, diagnosis by staring at the patient's urine. It lingered on until the mid-nineteenth century; then, more-or-less overnight, it was killed by the development of quantitative chemical analysis. But if uroscopy has become extinct, *neuroscopy* is alive and well. This, the most complex discipline of all, where we most desperately need detailed data, is the one where we have almost no quantitative measures of performance at all. As Lord Kelvin appreciated, without numbers there can be no genuine science: "When you can measure what you are speaking about, and express it in numbers, you know something about it; but when you cannot measure it, when you cannot express it in numbers, your knowledge is of a meagre and unsatisfactory kind" [2]. One could well argue that it is the lack of something that will give quantitative measures of brain disorder that has held back neurology for so long. A telling parallel is the development of the clinical thermometer, invented nearly 400 years ago by Santorio, turned into a practical clinical device 150 years ago by Sir Clifford Allbutt. A century and a half later, could saccadometry – the analysis of saccadic latencies – provide the magic diagnostic wand we need, as it were a *neurometer* that could match the thermometer?

Now of course when we measure a patient's temperature, their systolic pressure, or their plasma sodium, what we get is a single number. If we make repeated measurements, we can expect the values we obtain to change only slowly. Saccadic latencies are quite different: they vary dramatically from one trial to the next, even when separated by only a few seconds. However much we simplify the task, by stripping it back to its fundamentals, this remains true. The simplest of all saccadic tasks is the step task, in which a central visual target, after a random delay, jumps

unpredictably a fixed distance to the right or left. Clinicians sometimes call this a 'reflexive' movement (a regrettable terminology: Pavlov used to fine his students for using the word 'voluntary'; clearly we should do the same for 'reflex', or at least insist that those who use it should be prepared to define what they think it means). A clear indication that this response is not in any meaningful sense a 'reflex' is both its very long latency (typically some 200 ms), and even more significantly, its huge random variability, with 10% of responses usually lying outside a range of half the median latency. The saccadometer tends therefore to overwhelm us with information – an *embarrass de richesses* – measuring, in a typical clinical examination, perhaps 200 latencies in 8 minutes. This variability turns out to be a virtue, because its precise characteristics appear to reflect an essential aspect of how the brain is functioning [3]. So – almost uniquely in clinical science – this is a field where it is important to have ways of quantifying not just the *mean* values of what we are measuring, but the *distributional parameters*: and this in turn means that we have to think a little bit about the statistical analysis of the data the saccadometer generates, and what we hope to achieve by doing it.

Now there are two reasons for doing statistics at all. The first is purely pragmatic: to reduce a set of data to a smaller number of parameters by which it can be summarised. The second is that by modelling the data and its variation we can learn something about the underlying mechanism. Even in the first case there is still always a model, even though most people forget that it is there: even just reporting a mean and standard error, by implication we think that there is a metaphysical 'true' value which has been corrupted by evil noise. The purpose of modelling is of course to exorcise the evil noise by trying to explain it away.

2. The Later Model

A well-known feature of all latency histograms is that they are skewed, with a stretched-out tail of longer latencies, whose shape unfortunately fails to correspond to any of the standard mathematical stochastic functions. It was this failure that triggered the conception of LATER, and suggested that the conventional view of latency was fundamentally misleading. It is natural to think of latency in terms of a clock ticking away, since that is how we measure it. But if for a moment we concentrate not on the measurement process, but on what is likely to be going on in the system itself, what comes immediately to mind is that there must be some process, initiated by the stimulus, that continues at a certain rate until it reaches completion, in a way analogous to a chemical reaction. And as with chemical reactions, the most obvious explanation for why reaction times vary is simply that this underlying *rate* varies. So we need to think all the time in terms of reciprocal time, of $1/t$ and not t: we then find – quite excitingly – that the reciprocals of latencies follow a Gaussian or normal distribution (the latencies can then be said to follow a *recinormal* distribution),

a sign that we have reached a fundamental stochastic phenomenon. A convenient way of demonstrating this graphically is to plot cumulative latency histograms on a probit scale, with a reciprocal time-axis – a *reciprobit plot* (Fig. 1) [4]. If the reaction times really are recinormal, we should then get a straight line. We can represent the underlying mechanism in terms of a very simple functional model, LATER (Linear Approach to Threshold with Ergodic Rate)[5]. Here, we postulate a decision signal S, that in response to a stimulus rises linearly at a rate r from an initial value S_0, until it reaches a threshold or criterion level S_T, at which point a response is initiated. If we suppose r to vary as a Gaussian, with mean μ and variance σ^2, then the well-known skewness is immediately explained (Fig. 1).

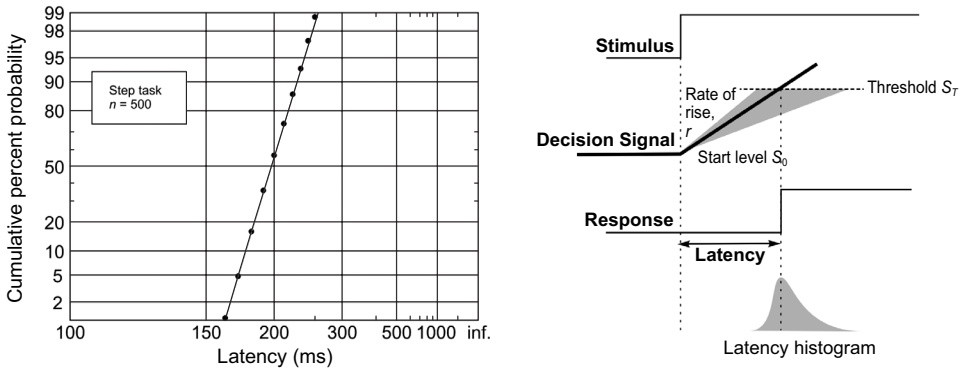

Fig. 1. The LATER model. Left, when latency distributions are plotted cumulatively using a probit scale, as a function of reciprocal latency (a reciprobit plot), a straight line is typically generated, implying that the reciprocal of reaction time follows a Gaussian or normal distribution. Right, this can be explained most easily by the LATER model, in which a decision signal S rises linearly in response to a stimulus at a rate r from an initial value S_0 until it reaches a pre-determined threshold value S_T, at which point a response is initiated; on different trials, r varies randomly as a normal variate (mean μ, variance σ^2). μ corresponds to the median of the reciprobit plot, and σ to its slope

An immediate benefit is that we can now encapsulate the behaviour of a subject with just two parameters, μ and σ. From a clinical point of view this can in itself be useful. An example is from one of the very first applications of saccadometry, the effect of sedative doses of volatile anaesthetic [6, 7]. This study was prompted by the need to know whether a patient who had just undergone day-surgery, and might therefore have a functionally significant concentration of anaesthetic remaining in their blood, could be safely discharged. In a double-blind experiment, subjects breathed known concentrations of anaesthetic (at concentrations well below the minimum anaesthetic concentration, or MAC), whilst their saccadic latencies in a step task were being measured. The results showed a clear and linear effect of sevoflurane concentration on the LATER parameter μ, at doses as small as 0.1 MAC, supporting its potential clinical utility (Fig. 2). And this would clearly be true even if we have not the slightest idea what μ actually *is*: LATER is useful even if it doesn't mean anything.

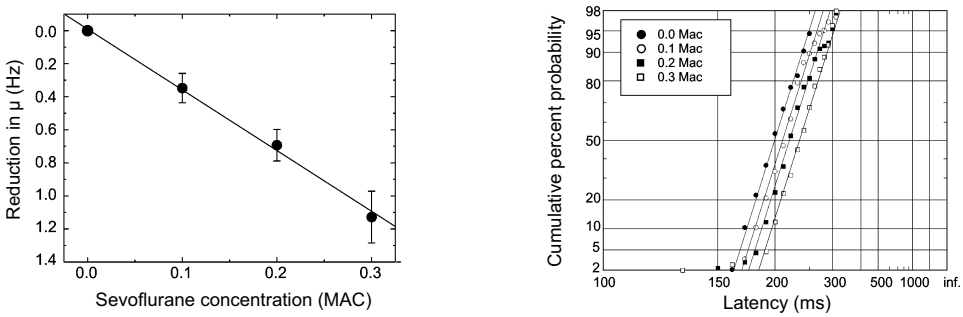

Fig. 2. Sedation and latency. Left, the LATER parameter μ falls linearly with increasing anaesthetic concentration: average of six human subjects, performing a saccadic step task while breathing low concentrations of a volatile anaesthetic, sevoflurane (MAC = minimum anaesthetic concentration). Right, reciprobit plots of saccadic latency in one subject at different anaesthetic levels, showing parallel shift of the curves, implying that the rate of information processing is impaired by the anaesthetic (Data from [7])

If LATER were merely arbitrary and empirical, it would have little scientific interest: but the underlying model is of course intended to have functional meaning. The argument is simply that if one accepts that reaction time is dominated by decision time (rather than being due mostly to such low-level phenomena as synaptic delay, conduction velocity, etc), then one might hope that the characteristics of RTs would in some way reflect what would be expected from an ideal decision-making mechanism. Decisions are necessarily probabilistic, and its elemental processes must therefore (leaving aside considerations of utility) be Bayesian if they are to be rational [8]. In essence this means combining prior probabilities or expectations with information received from sensory systems, in the form of likelihoods, to form updated, posterior probabilities: if the latter are sufficiently large, they may justify a response. Since priors and likelihoods combine by multiplication, it makes sense to imagine that these computations are performed using a logarithmic transformation of probabilities into degrees of neuronal activation, and recent experiments in monkeys trained to make dynamic probabilistic judgments demonstrate that this is indeed the case [9]. We can therefore postulate an ideal neuronal decision-maker, whose activity S represents the log probability of some hypothesis, perhaps the existence of a particular goal in the outside world. Initially, its initial activity level S_0 represents log prior probability; as information arrives from a stimulus, S rises steadily until it reaches a criterion level S_T – analogous to a significance level in statistics – at which the probability is so high that a response can be justified [5]. It should be apparent that this is exactly what is described by the LATER model (Fig. 1): perhaps surprisingly, our theoretically ideal decision mechanism happens to be indistinguishable from the simplest empirical description of how reaction times vary. Of course, there is one aspect of this functional interpretation that is not ideal, and that is the Gaussian variation itself. By neuronal recording from neurons in frontal cortex[10], this variability can

be demonstrated to be injected gratuitously within the brain rather than reflecting noise in the outside world: there are good reasons why this might be an evolutionary desirable feature, even if it is not Bayesian [11].

This may seem an ambitious claim, but it is firmly supported by human experiments in which the underlying Bayesian factors (expectation, rate of information supply, urgency) are manipulated, and the effects on the shapes of reaction time distributions are compared with what is actually observed [5, 12, 13]. An example of how this knowledge may be employed in a clinical situation comes from the same data concerning sedative levels of volatile anaesthetics that was presented earlier: if one looks at what the anaesthetic is doing to the distributions themselves (Fig. 2), it is clear that they are being displaced sideways in a parallel fashion; this is exactly what would be predicted from LATER if it were the case that the effect of sevoflurane is to reduce the rate at which information is being supplied, rather than, for instance, to reduce the general level of excitability.

A sceptic might well feel all that is too good to be true. How well does LATER in fact represent real, observed distributions of reaction time? Often we do indeed get very nice distributions that follow what LATER predicts quiet precisely, as for example in Figure 2. But sometimes, even though the task is still extremely simple, like the step task described earlier, we see significant and persistent deviations from this ideal behaviour. And sometimes we get complex distributions because the task itself is complex. In what follows we will look at each of these situations in turn: those where the peculiarities are idiosyncratic and due to the *subject*, and those where they are due to the *task*.

3. Idiosyncrasy

Figure 3 shows a representative set of reciprobits from 48 students, closely matched in age and educational attainment, performing a completely standardised saccadic step task. It is clear first of all that despite the homogeneity of both the population and the task there is considerable inter-subject variability in terms of μ and σ. In the study from which this data comes, we also measured μ and σ in a manual task, the subjects pressing buttons in response to the same step movements of the targets: the values were equally variable from subject to subject, and interestingly showed no correlation with the saccadic parameters. It is not simply, therefore, that each person has a basic 'speed' that applies to both kinds of response, though recent data shows that neurodegenerative disease as well as interventions such as deep brain stimulation can alter both manual and saccadic responses in a comparable manner [14].

But it is clear in Figure 3 that in addition to exhibiting inter-subject differences in the LATER parameters, many subjects' distributions are not in fact of a simple linear form, but show significant deviations from it, mostly in the short-latency region, but sometimes also in the long-latency tail. At present we tend just to ignore these

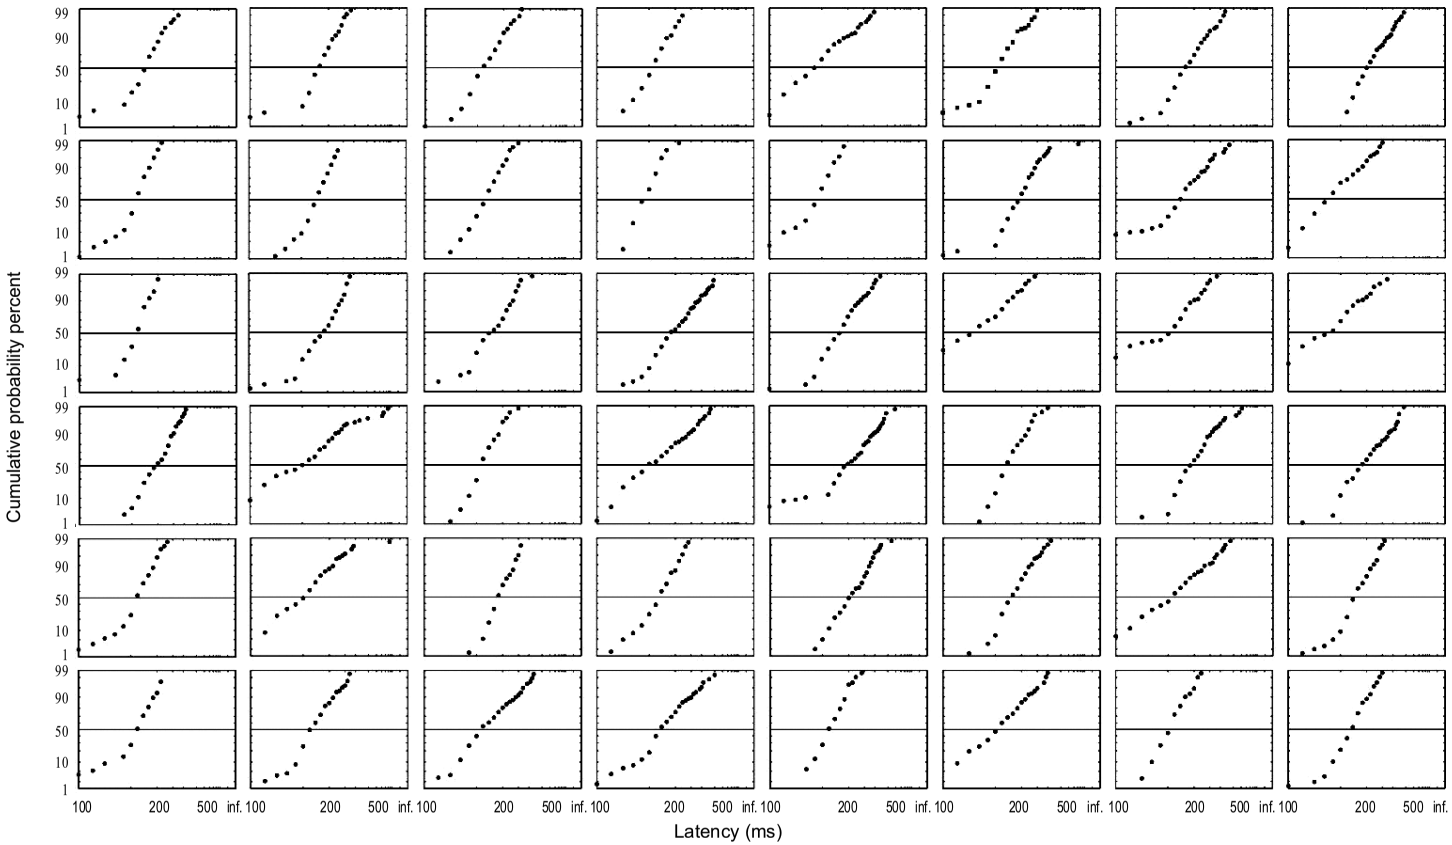

Fig. 3. Idiosyncrasy of saccadic latency distributions. Reciprobit plots are shown for 48 students of similar age and educational attainment, performing an identical saccadic step task (200 trials each)As well as differences in μ and σ, in many cases early, express and late components can be seen

deviations, treating them as part of the 'evil noise' that we can do nothing about: but – particularly when they persist with repeated measurement – these details of the distribution are probably telling us something important. There is an analogy here with colour, and the difference between colorimetry and spectroscopy. We perceive colour because we have three sets of receptors in our retina, that perform a crude analysis of the spectrum of the light reflected from an object, broadly by measuring the relative amounts of energy at long, short, and intermediate wavelengths. Because of the breadth and overlap of three systems, colours that look identical to us may have very different spectra: a yellow formed by mixing broad-band red with broad-band green may be indistinguishable from the yellow from a sodium lamp, yet their spectra are utterly different: one is smooth and wide, the other a single monochromatic spike. Thus spectral analysis provides detailed information that is completely invisible to our eyes; in the same way, limiting ourselves to simple low-order statistical measures, such as μ and σ, blinds us to the 'spectral' complexity exhibited by many individuals' reaction time distributions. These extra components can be found at both ends of the distributional spectrum, generating unexpected large numbers of either fast or slow saccades. We shall examine each of these in turn.

The fast ones have given rise to a great deal of confusion over the years. It was Burkhart Fischer who first drew attention to a sub-population of fast saccades that appeared to be distinct from the main population, and these are generally known as *express saccades*: they are more prominent in tasks such as the gap task in which there is advance warning of the stimulus, and in other situations of enhanced expectation [15, 16]. The defining feature of express saccades is that they create an overall distribution that is bimodal, resulting in a reciprobit plot that shows two relatively steep regions separated by one that is flatter, or even horizontal (Fig. 4). One might think that this behaviour could be modelled simply by having two LATER units in parallel, a fast one that generates the express saccades, and a slower one that accounts for the main part of the distributions. However, such an arrangement will not in fact

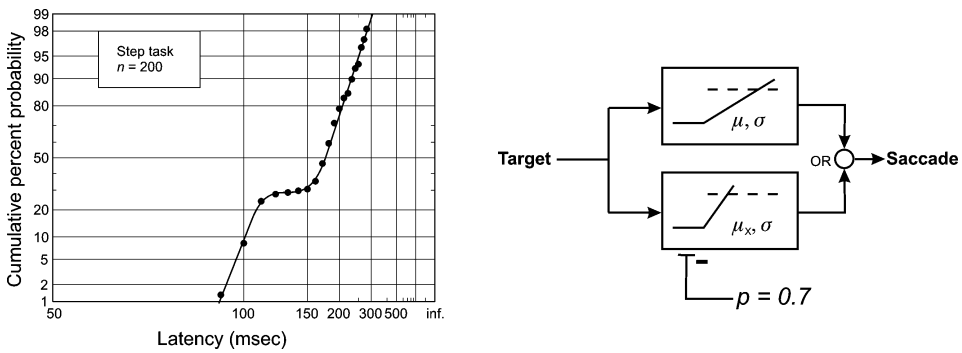

Fig. 4. Express saccades. Left, reciprobit plot of a subject generating a sub-population of express saccades in a step task, thus showing bimodality. The line was generated by a model of the form shown on the right, in which a faster LATER unit is tonically inhibited, but randomly activated on some trials (in this case, with a frequency of 0.3)

generate bimodality, because the faster unit is almost always going to win against the slower one. What is needed is for the faster units to be tonically inhibited, the inhibition being lifted only on occasional trials, at random. Such a model can simulate the observed distributions very accurately (Fig. 4). It introduces two new parameters: the mean rate of rise of the express unit, and the probability of its being activated; in general the σ of the express unit can be taken to be identical with that of the main unit. What is not satisfactory, though frequently to be seen in published papers, is to adopt some fixed criterion, for example that any saccade shorter than 150 ms is 'express'. It is obvious from the stochastic nature of the distributions that saccades generated by the main unit may easily happen to have latencies as short as this, and in addition a crude criterion of this kind makes no allowance for the very great variability of median latency that is seen across subjects (Fig. 3). The only satisfactory way to report the incidence of express saccades is to estimate the parameters of the underlying model.

The other kind of short-latency response is the *early saccade*, actually considerably more common than express saccades in the step task. The distinction with express saccades is that early saccades do not in themselves give rise to bimodal distributions: in reciprobit plots their existence is immediately obvious as they form a sub-population lying on a straight line of shallower slope, that usually extrapolates back to the origin (infinite time and 50% probability) (Fig. 5). Like express saccades, they become more prominent under conditions of high expectation, and also when a subject is under pressure to respond quickly (urgency), or when there is an element of distraction from another demanding task [12, 17]. They are also a prominent and invariable feature of the distribution of latencies (i.e. the interval between the end of one saccade and the beginning of the next) in spontaneous viewing, most notably in reading, and in nystagmus [18–21], very likely because in such cases there is a particularly high expectation of the appearance of the new target [21]. Early saccades turn

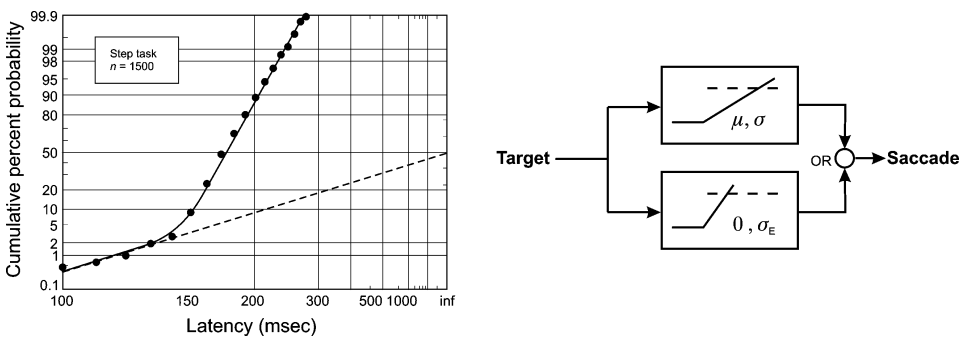

Fig. 5. Early saccades. Left, reciprobit plot of a subject generating a sub-population of early saccades in a step task, that tend to lie on a straight line extrapolating to the origin (infinite latency, 50% probability). The line was generated by a model of the form shown on the right, in which a 'rogue' LATER unit with $\mu = 0$ and an increased σ competes with the main unit that generates the majority of the responses (Data from [18])

out to be particularly easy to model: what is needed is again a LATER unit in parallel with the main unit, but this time it has a large value of σ, but a μ of zero – generating the observed extrapolation to the origin. It is, in other words, a kind of maverick, rogue unit that because of its wildness sometimes manages to overtake the main unit (Fig. 5). A plausible explanation, not yet proved, is that early saccades (and possibly express saccades as well) are due to a fast brainstem or collicular route, normally kept in check by descending tonic inhibition from the cortex that is weakened under conditions of high expectation, urgency or cognitive distraction [12, 17].

At the other end of the saccadic distribution, a phenomenon equivalent to express saccades is when there exists a sub-population of saccades of unusually long latencies, their median often twice as long as that of the main distribution, that may give rise to bimodality (Fig. 6). They are often noticed by the subject, who realizes that he has failed to make a response and performs a 'double take'; they are frequently a sign of sleepiness and inattention. These *late saccades* are also easy to model, in a similar way to express saccades: a second LATER unit, in parallel with the main unit, has a much smaller value of μ and would therefore not normally give rise to a response. But in a random proportion of trials the main unit suffers some kind of descending inhibition and fails to operate, and it is therefore this second 'late' unit that generates the response. This kind of inattention is seen in manual responses as well as saccades, and may in itself be an indication of pathology.

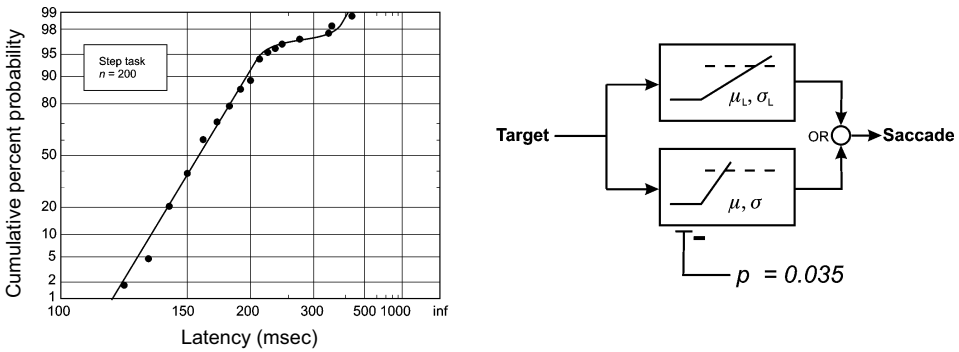

Fig. 6. Late saccades. Left, reciprobit plot of a subject generating a sub-population of late saccades, generating a degree of bimodality. The line was generated by a model of the form shown on the right, in which a slower LATER unit is in parallel with the main unit, the latter randomly failing to be activated on a certain proportion of the trials (in this case with a frequency of 0.035)

4. Complex Tasks

The other way in which complex latency distributions may be observed is when the task itself is complex. These situations are not merely of abstract scientific interest, as in clinical settings such tasks could be even more informative than the simple

step task. A good example is the antisaccade task, in which a subject is instructed to make a saccade in the opposite direction to the presented target. In some conditions, for instance Parkinson's Disease or Huntington's Disease, there is reason to think that performance in the antisaccade task may be a more sensitive indicator of pathological impairment than the ordinary step task [22, 23]. However, at present this potential usefulness is to an extent cancelled out by the fact that performance in the task is typically evaluated relatively crudely, often, in effect, qualitatively rather than quantitatively. In the absence both of a standard experimental paradigm for such tests, and of a quantitative model by which underlying parameters can be estimated, it may not be as informative as the step task, which is typically performed all over the world in an absolutely standardized way, and for which the parameters can be estimated precisely enough to enable quantitative measures of disease progression and therapeutic success to be estimated. However, this situation is likely to change once quantitative models of the type to be described below begin to be deployed more widely, and we acquire greater experience of how the underlying parameters might relate to the disease condition itself.

It is helpful to start with one of the simpler complexities, the *countermanding task*. A typical implementation consists of the basic step task, except that in a certain number of trials, chosen at random, after a delay D the central fixation spot reappears in addition to the target, and acts as a 'stop signal' that instructs the subject to withhold the impending saccade (Fig. 7). As might be expected, if D is long this signal arrives too late, and the saccade is nevertheless made; with smaller values of D there is an increasing likelihood that the countermanding succeeds. The behaviour is stochastic, in that for a given value of D there is a particular probability of success, and this relationship provides quantitative information from which one can postulate quite a simple model: that there is a race between the LATER-like mechanism

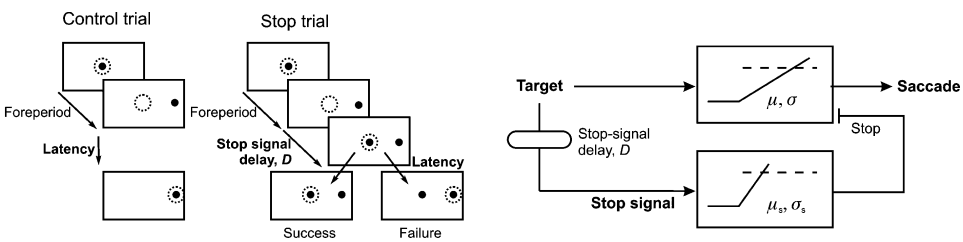

Fig. 7. Countermanding. Left, in an otherwise conventional step task the central fixation spot occasionally reappears after a delay D, a 'stop signal' that instructs the subject to withhold the impending saccade that would otherwise have been made: with decreasing D, the ability to countermand successfully steadily diminishes. (The dashed circle represents the position of the gaze). Right, a model that can successfully simulate both the error rates in this task, and the distribution of latencies when countermanding has been unsuccessful. A main LATER unit initiates saccades in the usual way; but it is supplemented by a STOP unit, also LATER-like, whose rise is triggered by the appearance of the stop signal. If the STOP unit reaches its threshold before the main unit, the impending saccade is cancelled. For simplicity, fixed conduction delays have been omitted, and lateral inhibition

that is initiated by the original stimulus and would normally culminate in a saccade (the GO process), and a STOP process, also LATER-like, that is initiated by the stop signal and cancels the saccade if it reaches completion before the GO process does [24–26] (Fig. 7). Such an arrangement generates good simulations both of the distributions of latencies for different values of *D*, and the frequencies with which countermanding is successful. The fact that we can probe the system by altering *D* is attractive from the scientific point of view, but is a drawback for possible use in clinical practice as it considerably lengthens the duration of the test. Furthermore, because the majority of trials are controls (and even on test trials, those in which countermanding successfully occurs generate no reaction time data), this is in many ways an inefficient experimental paradigm.

A classic variation of this is the *Wheeless task* [27]. As in countermanding, the basic paradigm is of a conventional step task, but on a certain number of randomly-chosen trials, after a delay *D* from the presentation of the target, it jumps a second time to the opposite side. When *D* is long, the subject makes an initial saccade to the first position of the target, and a subsequent one to the second position (a Type A response). But – rather as with countermanding – as *D* is reduced, the behaviour becomes stochastic: the subject is increasingly likely to saccade straight to the second target without an intermediate saccade to the first position (a Type B response). There is a more-or-less linear relationship between *D* and the probability of making a Type A response (Fig. 8). When the latency distributions for the first saccade in Type A and Type B responses are examined, they show a significant feature: the distribution of Type A responses tend to start by coinciding with the controls, but then cease; there is then a gap in which no responses are made at all, followed by the Type B responses (Fig. 8). This suggests strongly that – rather as in countermanding

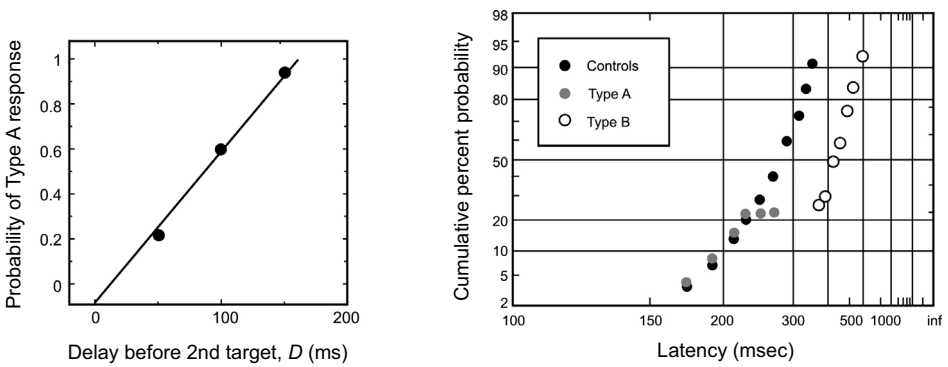

Fig. 8. The Wheeless task. In a certain number of trials during a conventional step task, after a delay *D* the target jumps to the opposite side. The first subsequent saccade may then be either to the original target (Type A response) or to the new one (Type B): with increasing *D*, the probability of Type A increases roughly linearly (left) (Data from [30]). Analysis of the distributions (right: in this case *D* = 100 ms) shows that the extra target jump induces bimodality, suggesting that it causes the underlying LATER unit to be stopped (Data from [27])

– there is a STOP unit that is activated by the appearance of the second target, that cancels the impending first response, and also initiates the linear rise of a second LATER unit that eventually generates the saccade to the second target.

Another popular task, related to Wheeless, is the *Go/No-go* task. In one version, a conventional step task is performed, except that on some trials the target is red, whereas on others it is blue: the subject is instructed to saccade to red targets but ignore blue ones. In practice, mistakes are made, as the appearance of the wrong colour target is nevertheless sufficiently compelling to evoke a saccade, despite the prohibition. We can then measure latency distributions for correct responses (i.e. to the red target) and error responses (to the blue), making this a relatively efficient way of gathering data, especially as there is no *D* that needs to be varied. A model of similar type to the Wheeless one, with a STOP process that cancels any impending response and initiates the rise of a LATER units for the correct one, does a good job of predicting both the observed latency distributions, and the error rates [28]. In this model the STOP unit is assumed to be activated by colour information, arriving after a delay relative to the simple message that a target is present; such a delay can be seen in recordings from neurons in monkey frontal eye fields in other tasks requiring a colour discrimination to be made [10].

Finally, the *antisaccade task*, described earlier. Like Go/No-go, this has the advantage that there is no *D* to prolong the duration of experimental trials, and from this point of view has the additional benefit that all trials produce some kind of response (correct antisaccade or incorrect prosaccade) whose latency can be measured and analysed. As in all these more complex tasks, the choice of response, as well as its timing, is stochastic, and can be modelled using LATER-like units combined with a STOP unit, responding again with a delay that could be taken to represent the added processing time needed to implement the underlying, unnatural transformation of

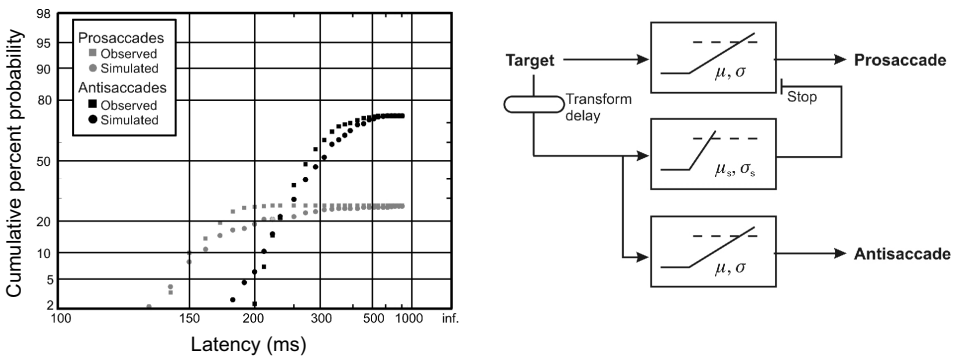

Fig. 9. Antisaccades. In an antisaccade task, subjects are instructed to make their saccades in a direction opposite to that of the target. Left, observed correct (antisaccade) and incorrect (prosaccade) responses made by a subject, together with results of a simulation using the model similar to the one on the right, in which the target activates a Stop unit to inhibit the more natural prosaccadic response, as well as a unit corresponding to the correct, antisaccadic, response (Data from [29])

stimulus position into response position. Preliminary work [29] suggests that this model can do a good job of modelling both latency distributions and error rates (Fig. 9), with the implication that it should be possible to determine which parameters are affected in pathological conditions that impair antisaccade accuracy or speed.

5. Conclusion

Being able to model quantitatively these more complex responses could potentially greatly improve our ability to apply saccadometry to the diagnosis of disease and the monitoring of its progression, to providing precise quantitative evaluations of the efficacy of different kinds of treatment, and even more particularly, to the identification of sub-types of such over-generalised categories as Parkinsonism, and thus the ability to match therapies much more effectively to conditions than is currently the case. Many of the problems in these areas are due to personal variation or idiosyncrasy, and saccadometry has now reached the point where rather than ignoring this factor we can actually willingly embrace it.

References

1. Sherrington C.S.: Man on his Nature. Cambridge: Cambridge University Press 1940.
2. Thomson W.L.K.: Electrical units of measurement. Popular Lectures and Addresses 1883, vol. I, 73.
3. Carpenter R.H.S.: The saccadic system: a neurological microcosm. Advances in Clinical Neuroscience and Rehabilitation 2004, 4, 6–8.
4. Carpenter R.H.S.: Oculomotor Procrastination. In: Fisher D.F., Monty R.A., Senders J.W., (Eds): Eye Movements: Cognition and Visual Perception. Hillsdale: Lawrence Erlbaum 1981, 237–246.
5. Carpenter R.H.S., Williams M.L.L.: Neural computation of log likelihood in the control of saccadic eye movements. Nature 1995, 377, 59–62.
6. Carpenter R.H.S., Descamps M.J.L., Morley C.H., Leary T.S., Jones J.G.: The effect of low dose sevoflurane on saccadic eye movement latency. Anaesthesia 2002, 57, 855–859.
7. Nouraei S.A.R., de Pennington N., Jones J.G., Carpenter R.H.S.: Dose-related effect of sevoflurane sedation on the higher control of eye movements and decision-making. British Journal of Anaesthesia 2003, 91, 175–183.
8. Good I.J.: Rational decisions. Journal of the Royal Statistical Society Series B. 1952, 14, 107–114.
9. Yang T., Shadlen M.N.: Probabilistic reasoning by neurons. Nature 2007, 447, 1075–1080.
10. Thompson K.G., Hanes D.P., Bichot N.P., Schall J.D.: Perceptual and motor processing stages identified in the activity of macaque frontal eye field neurons during visual search. Journal of Neurophysiology 1996, 76, 4040–4055.
11. Carpenter R.H.S.: A neural mechanism that randomises behaviour. Journal of Consciousness Studies 1999, 6, 13–22.
12. Reddi B., Carpenter R.H.S.: The influence of urgency on decision time. Nature Neuroscience 2000, 3, 827–831.
13. Reddi B.A.J., Asrress K.N., Carpenter R.H.S.: Accuracy, information and response time in a saccadic decision task. Journal of Neurophysiology 2003, 90, 3538–3546.

14. Antoniades C.A., Carpenter R.H.S., Temel Y.: Deep brain stimulation of the subthalamic nucleus in Parkinson's disease: similar improvements in saccadic and manual responses. Neuroreport (in press).
15. Fischer B., Boch R.: Saccadic eye movements after extremely short reaction times in the Monkey. Brain Research 1983, 260, 21–26.
16. Fischer B., Boch R.: Express saccades of the monkey: a new type of visually guided rapid eye movement after extremely short reaction times. In: Gale A.G., Johnston F. (Eds) Theoretical and Applied Aspects of Eye Movement Research. Amsterdam: Elsevier 1984, 403–408.
17. Halliday J., Carpenter R.H.S.: The effect of cognitive distraction on saccadic latency. Perception 2010, 39, 41–50.
18. Carpenter R.H.S.: Express optokinetic nystagmus. In: Fuchs A.F., Brandt T., Büttner U., Zee D. (Eds) Contemporary ocular motor and vestibular research. Stuttgart, Georg Thieme 1994, 185–187.
19. McDonald S.A., Carpenter R.H.S., Shillcock R.C.: An anatomically-constrained, stochastic model of eye movement control in reading. Psychological Review 2005, 112, 814–840.
20. Carpenter R.H.S., McDonald S.A.: LATER predicts saccade latency distributions in reading. Experimental Brain Research 2006, 177, 176–183.
21. Roos J.C.P., Calandrini D.M., Carpenter R.H.S.: A single mechanism for the timing of spontaneous and evoked saccades. Experimental Brain Research 2008, 187, 283–293.
22. Everling S., Fischer B.: The antisaccade: a review of basic research and clinical studies. Neuropsychologia 1998, 36, 885–899.
23. Hutton S.B., Ettinger U.: The antisaccade task as a research tool in psychopathology: a critical review. Psychophysiology 2006, 43, 302–313.
24. Hanes D.P., Carpenter R.H.S.: Countermanding saccades in humans. Vision Research 1999, 39, 2777–2791.
25. Hanes D.P., Schall J.D.: Countermanding saccades in macaque. Visual Neuroscience 1995, 12, 929–937.
26. Logan G.D., Cowan W.B., Davis K.A.: On the ability to inhibit responses in simple and choice reaction time tasks: a model and a method. Journal of Experimental Psychology: Human Perception and Performance 1984, 10, 276–291.
27. Wheeless L.L., Boynton R.M., Cohen G.H.: Eye-movement responses to step and pulse-step stimuli. Journal of the Optical Society of America 1966, 56, 956–960.
28. Noorani I., Gao M.J., Pearson B.C., Carpenter R.H.S.: Predicting the timing of wrong decisions. Experimental Brain Research 2011, 209, 587–598.
29. Noorani I., Carpenter R.H.S.: The Timing of Antisaccades. Journal of Physiology Proceedings 2011, 23, PC295.
30. Becker W., Jürgens R.: An analysis of the saccadic system by means of double-step stimuli. Vision Research 1979, 19, 967–983.