

7 The Steady State or Null-Point Technique

One kind of relativistic system (Chapter 4, section 2) is a steady-state or null-point system. It is conveniently exemplified in the area of individual human psychology though similar techniques are used in connection with other kinds of observation in chemistry, physiology and sociology.

According to the classical rules of experimentation, a human subject is isolated as far as possible from extraneous influences and presented with an environment in which he performs a task; generally, a sequential task in which he responds to a series of stimuli. The subject's response is observed as a dependent variable. The task parameters are held constant during an experiment but between experiments some of them are changed, as independent variables, in order to examine the effect of different treatments.

These rules are tailored to fit a mode of thinking in which the subject and the environment are conceived of as finite-state mechanisms. These finite automata may be probabilistic devices; the subject certainly is. But, as an article of faith, the probabilistic automaton may be embedded in a much larger deterministic automaton. Of course, there is no suggestion that the structure of the subject automaton is known. The experimenter might be 'completely ignorant' in the sense that he looks at a 'black box' with input (including parametric input) and output. In *that* sense the experimenter has *no model* (Skinner's (1969) contention, for example). On the other hand, it can be argued that a very definite model exists in so far as the experimenter believes all that *may be known* about the subject is expressible by filling out the 'black box' as a (probabilistic) finite-state machine (or, the equivalent in continuous mathematics, as a generalised transfer function). The argument depends, of course, upon the identification of *input states* and *output states*. It is generally supposed that a stimulus is an input state (meaning 'the only relevant change in input state is the change indexed by the stimulus'). By the same token (and with the same qualification) a response is an output state.

The difficulties over this picture of things can be looked at from more or less fundamental points of view. For example, we might question the idea that the *goal directedness* that impels the subject to fixate attention on the

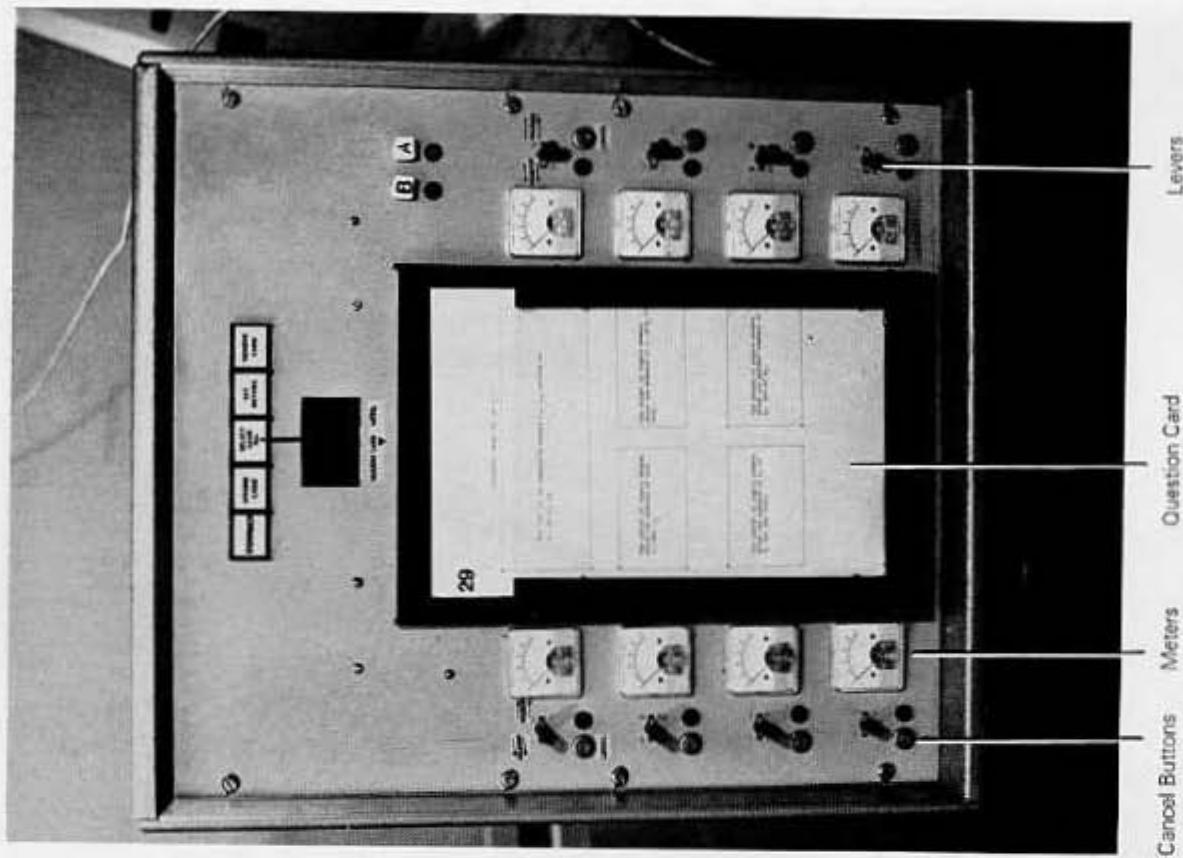


Plate I BOSS (Belief and Opinion Sampling System)

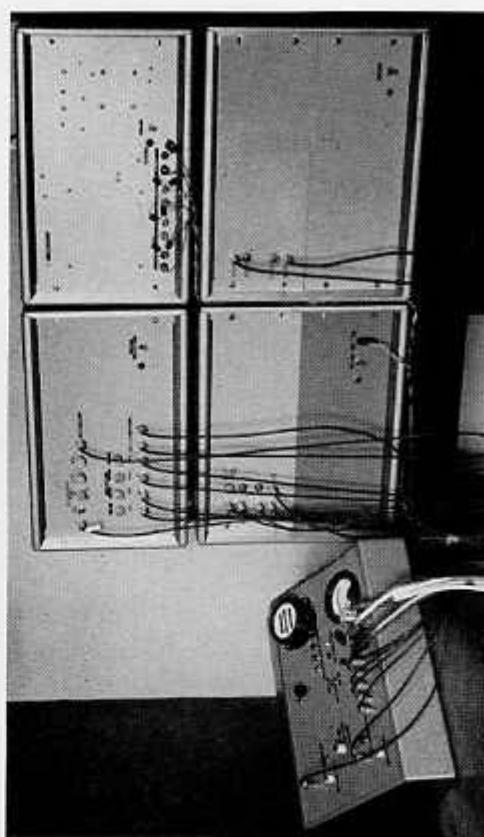
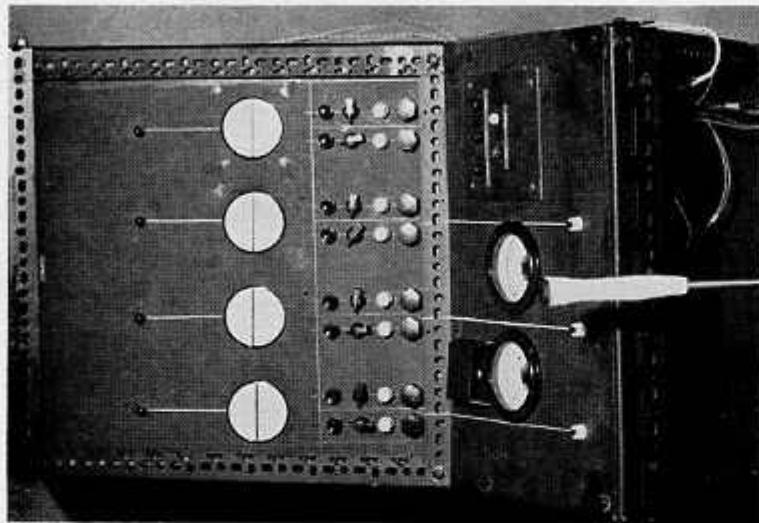


Plate 2a Purchaser decision simulator; respondent's console.

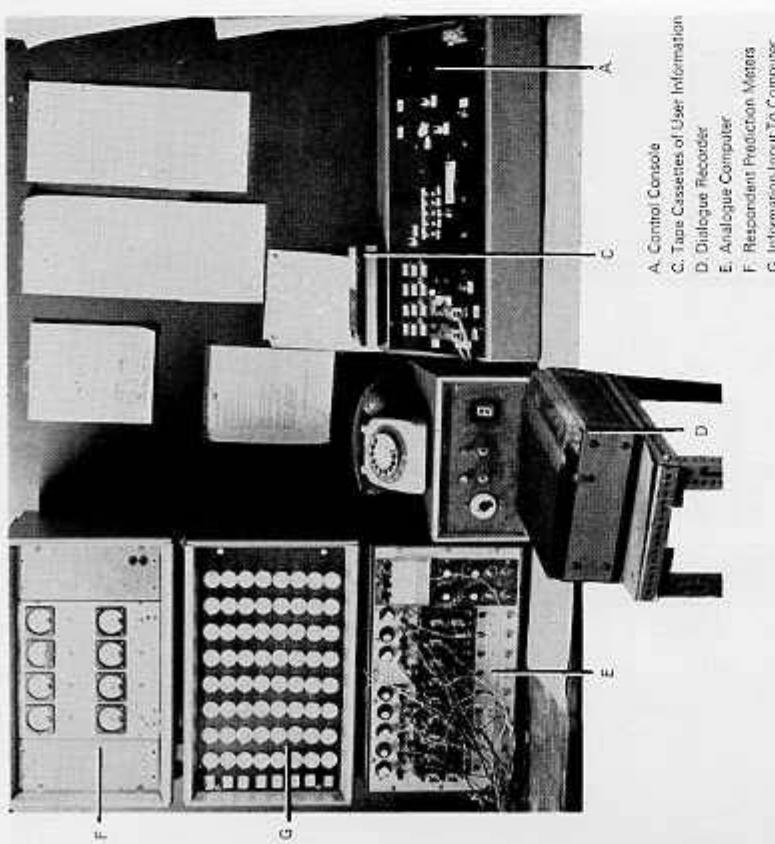


Plate 2b Purchaser decision simulator; experimenter's console.

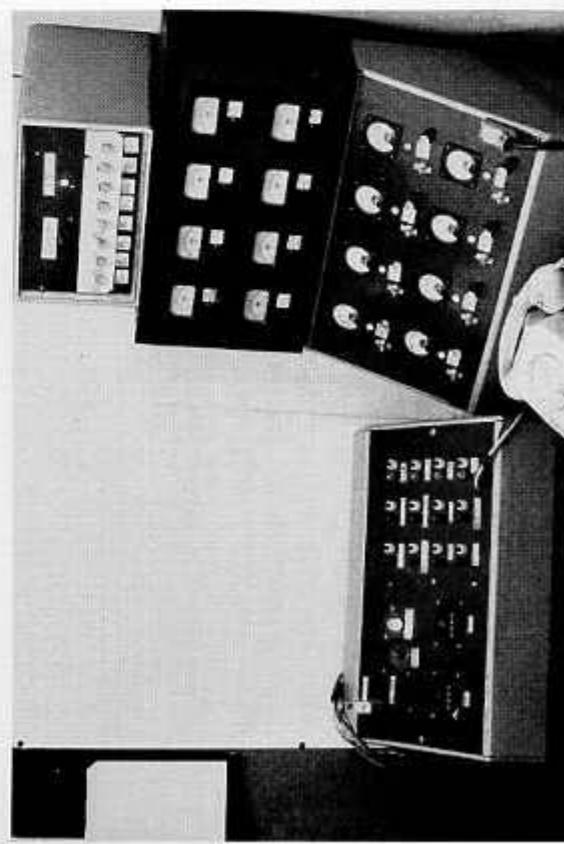


Plate 3a Compensatory tracking; subject's console.

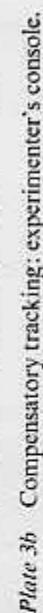


Plate 3b Compensatory tracking; experimenter's console.

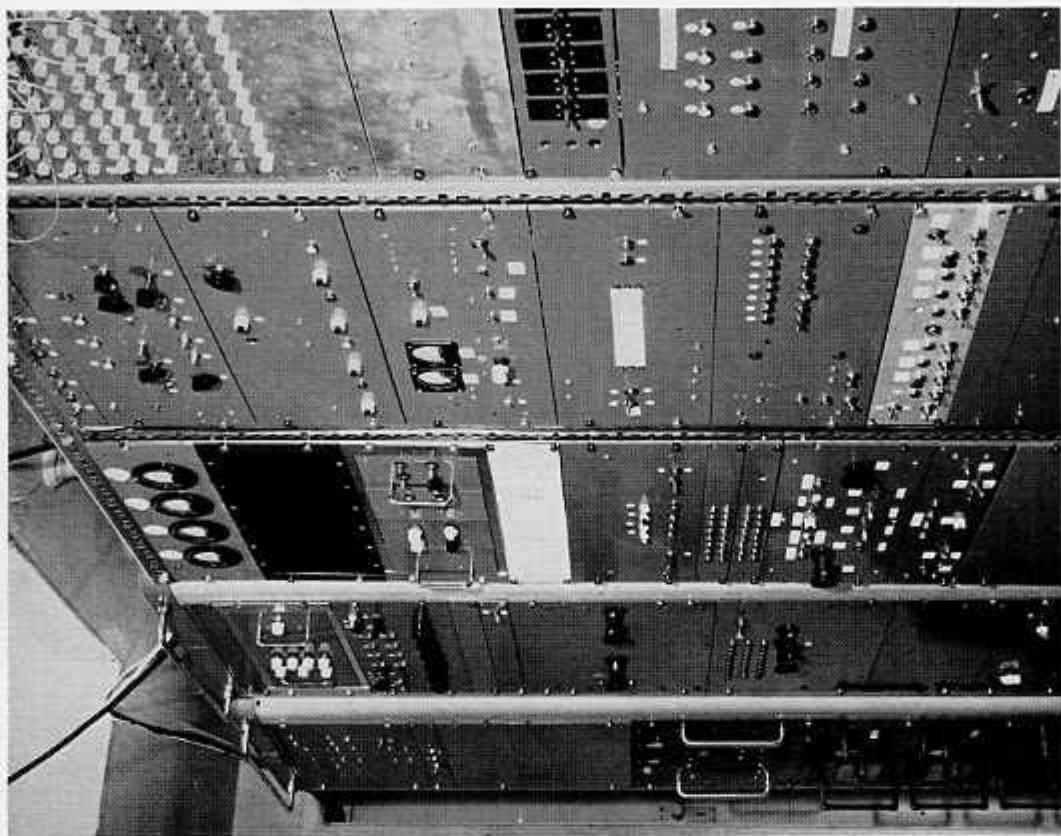


Plate 5 Special purpose computer for adaptively controlled experiments.

Plate 6 Target interception task: subject's console.

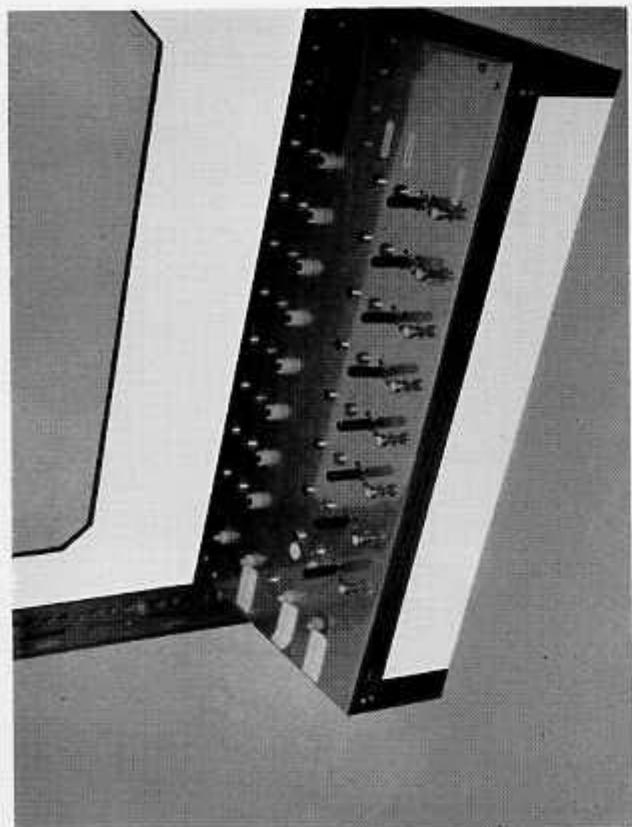
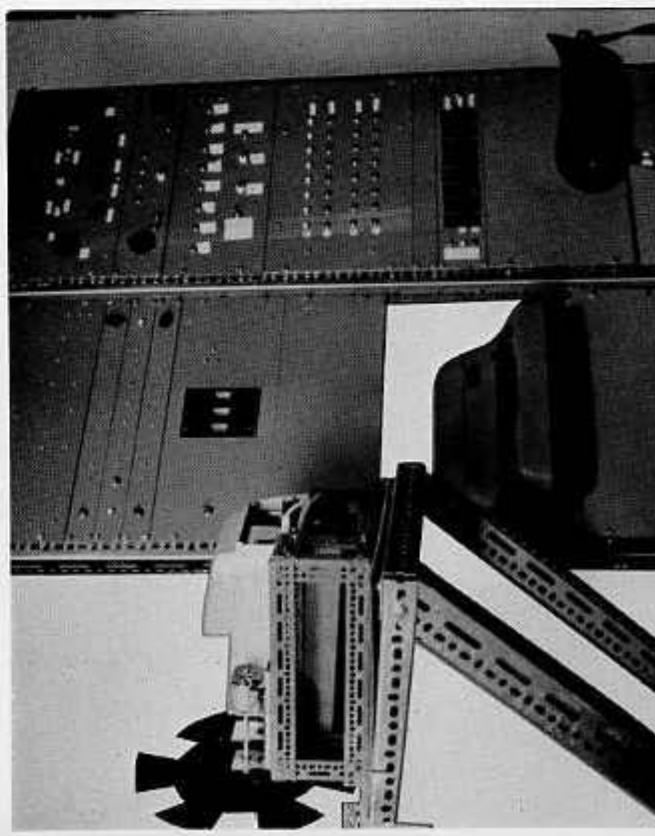


Plate 4a Attribute selection task: subject's console.

Plate 4b Attribute selection task: experimenter's console.



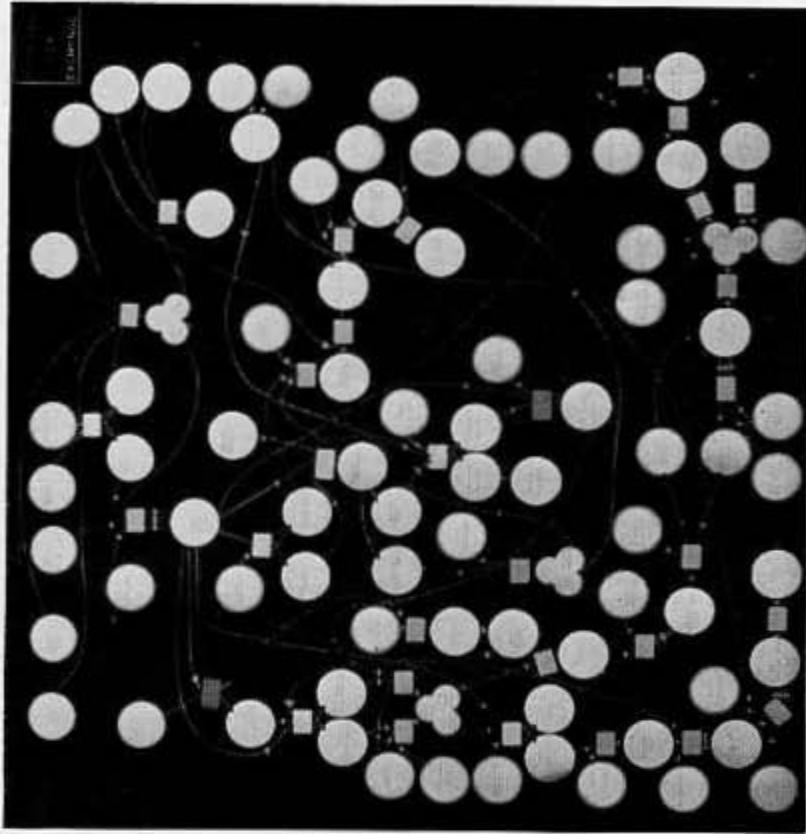


Plate 7 Adaptively controlled perceptual discrimination task with control equipment.



Plate 8 Subroutine flow chart training instrument for a clerical task.

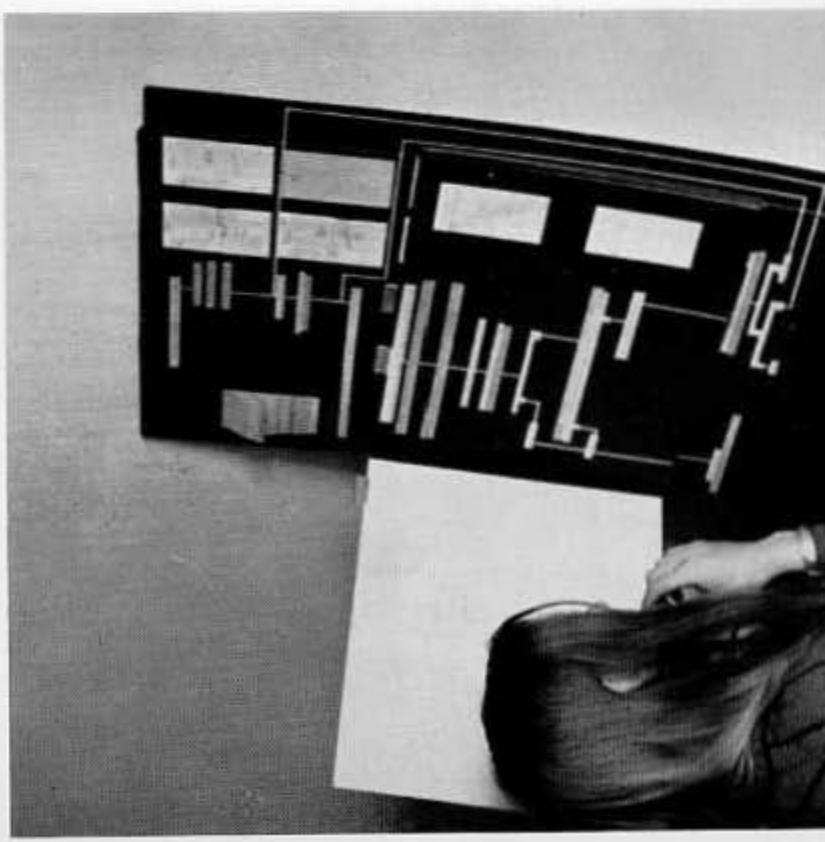


Plate 9 Representation of the process of carrying out a clerical task: that of processing a customer's order for the installation of some equipment. Circular nodes represent conditions; boxes represent transitions; 'cloverleaf' nodes represent conditions requiring information from outside the system in order to determine the next transition.

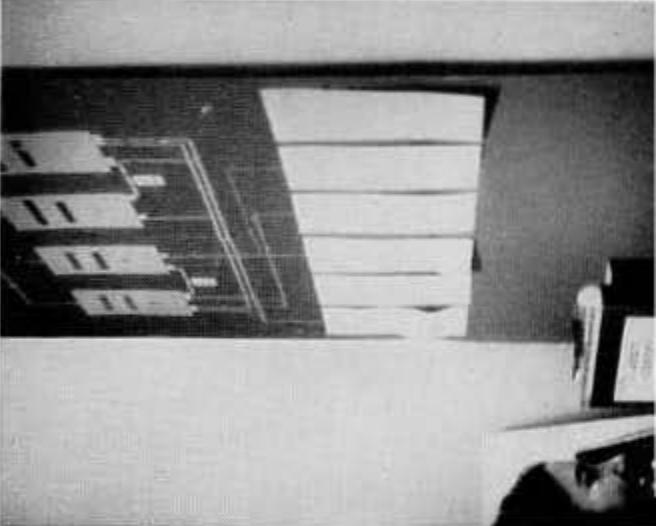


Plate 10 Subject free learning a zoological taxonomy.

task is simply a terse way of expressing a behavioural regularity (the notion that it is so forms the major tenet of behaviourism and underpins the experimental method). Or we might question the idea that *goal directed adaptation* is due to some kind of reinforcement (psychological term) or directed parametric modification (cybernetic term); once again, some such notion underpins the method. Perhaps the least sophisticated comment to make is that human beings, as well as rats and monkeys, are prone to all manner of internal and unobservable changes that alter their behaviour. The reality of coupling these organisms to an environment in such a way that stimuli can be *legitimately* regarded as input states, is dependent upon providing a certain kind of environment and task which places the organism in an *operating region* where the task is neither overloading nor underloading.

The underloading caveat is due to one of the inbuilt propensities for change noted a moment ago; it seems that an organism, and a human being in particular, is built to learn and, as a result, must be given something to learn about if he is to remain *genuinely* coupled to the environment (i.e. if the organism's input states are to be identifiable and in correspondence with the environment's output states and vice versa). If the underload limit is contravened then the human being attends to something else (which is unobservable) even though he may still devote *part* of his attention to the routine task. Conversely, a task that is too difficult is rejected as unintelligible. True, the human being may brood on the matter as a puzzle, but *that* is not observable in the stipulated framework. If either the underload or the overload conditions are contravened, the desired input state/output state/relationship is disrupted. In either case the crucial closure condition which substantiates the coupling of the organism/environment subsystems as *the* system under observation is lost. Moreover, this must be so if man is made to learn, and built, also, with limitations upon capacity, storage and the like.

The *operating region* is not fixed. What does or does not 'load' the human being depends upon his experience and the extent of his adaptation to the environment. Hence, the classical idea of obtaining constant experimental conditions by setting a constant task in a carefully replicated environment from which irrelevant disturbances are excluded, is fated to become an inappropriate paradigm for experiments (or practical test and performance situations) that are at all prolonged. Either the task will overload the subject or it will come to underload the subject; in either case, the whole basis for identification is undermined. The objection, it should be stressed, is inapplicable to many *one shot* experiments (for example, in perceptual psychology) and it is pedantic for short experiments or where problem solving of a particular type is *meant* to be private (for example, experiments on the difficulty human beings experience over complementation and in extracting data from negating statements of 'what is not the case'). In its

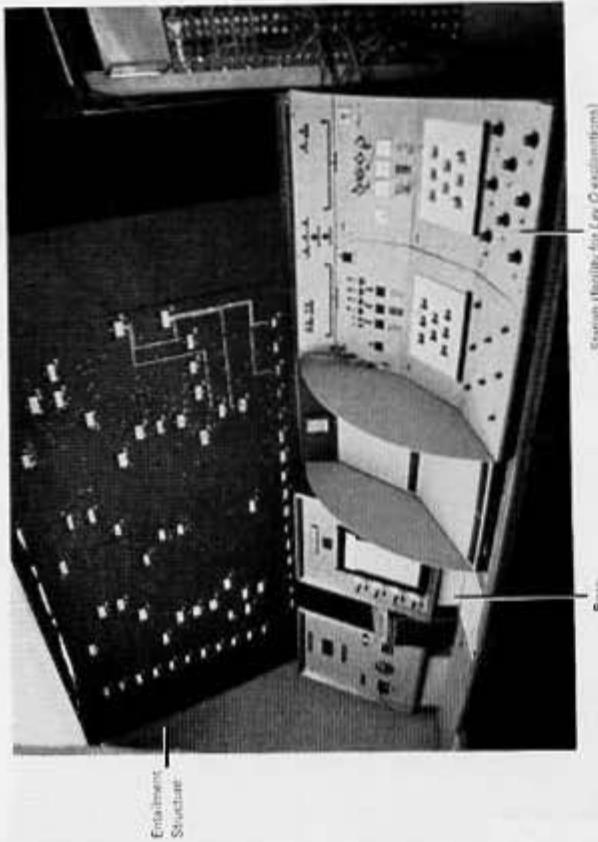


Plate 11 CASTE (Course Assembly System and Tutorial Environment): student system.

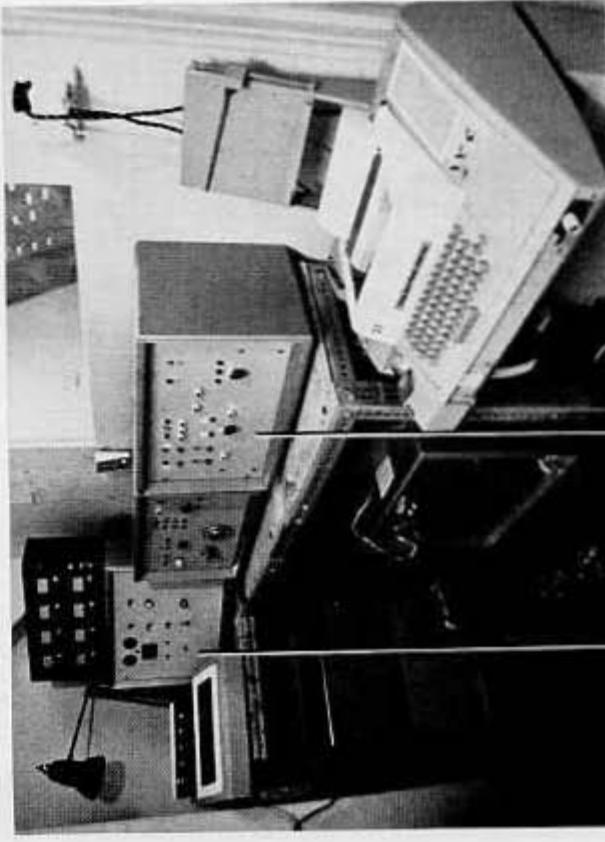


Plate 12 CASTE: operator's station with window overlooking student station.

present form the objection applies only to lengthy experiments. With human subjects though, if a more fundamental point of view is adopted, the principle is universal.

A more appropriate paradigm for securing constant conditions would be an arrangement whereby the subject (in the examples, this is a human being) can be kept working within his individually determined and individualistically altered *operating region*.

The experimenter could, in principle, take the necessary precautions himself, i.e. he might adjust the rate and form of stimulation so that a response criterion, indicative of continuing task performance, is satisfied; for example, by making life more difficult if an index of correct performance increases in value and making life less difficult if the correct performance index decreases in value. To instrument this plan the observer needs a task and an environment which furnishes more than enough variety to overtax the subject at any point in the experiment.

Two problems result from allowing the experimenter to act in this way. First there is a practical administration problem. All the measuring and adjusting is likely to overtax the participant observer so that (even if he circumscribes a compartment of his mind to act as a *regulator* and leaves the rest of it to *observe*) he is likely to find himself so stressed that no energies are left for his main occupation, which is observing. Secondly, there is an experimental design problem due to the fact that the observer is no longer *external* but rather *participant* at least some of the time. It is later shown that the regulatory activities needed to secure the desired consistency 'subject working in operating region' entail other than causal transactions between the subject and the 'regulator', and thus render any 'regulator' a participant most of the time.

The proposed method is as follows. Instead of the classical organism/environment system in Fig. 57, construct the system of Fig. 58 in which the controller is a mechanical or computer programmed device for executing a steady state or null point operation. If ρ is a performance index and if a level of task difficulty is represented by η then the controller in question determines the value of ρ , compares it with a null point value ξ (indicative of satisfactory and correct task performance) and operates upon η , as follows, for experimental trials (or time instants) indexed n .

$$\eta(0) = 0$$

$$\eta(n) = \sum \Delta \eta(n)$$

$$\Delta \eta(n) = \begin{cases} +1 & \text{if, } \rho > \xi \text{ unless } \eta(n) = \eta_{\max} \text{ when } \Delta \eta(n) = 0 \\ 0 & \text{if, } \rho = \xi \\ -1 & \text{if, } \xi > 0 \text{ unless } \eta(n) = 0 \text{ when } \Delta \eta(n) = 0 \end{cases}$$

As a convenience, it is usual to assume that ρ and η are normalised so that

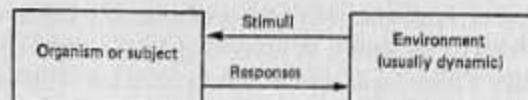


Figure 57 The classical organism/environment system.

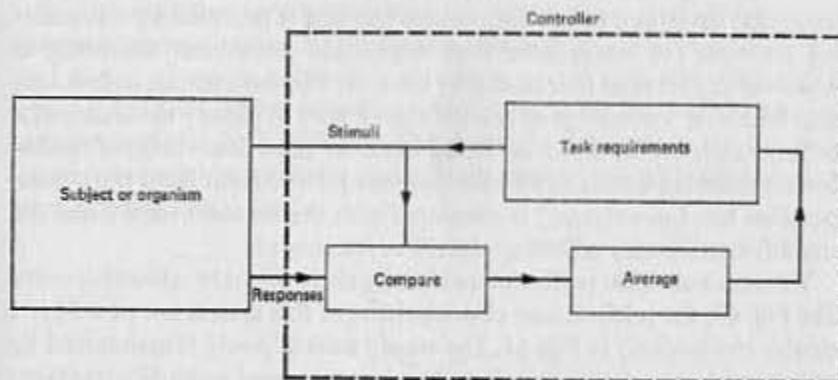


Figure 58 Alternative system to that of Figure 57.

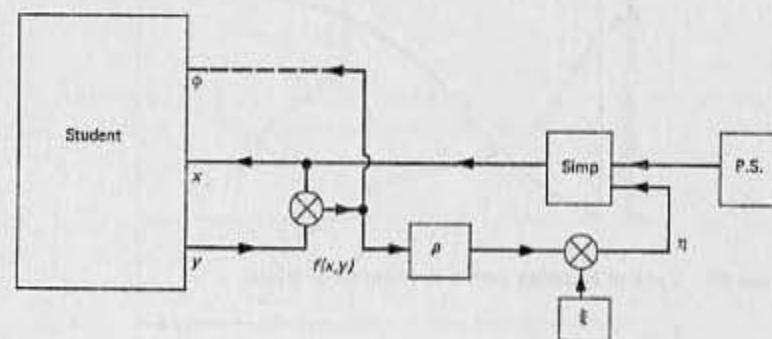


Figure 59 x , stimulus display; y , response; $f(x,y)$, 'correct' comparator output; P.S., problem source; Simp, problem simplification or difficulty variation; ϕ , knowledge of results, if given.

$1 > \rho > 0$ and $1 > \eta > 0$. Since task difficulty is modified, in effect, by simplifying task situations that are inherently too difficult, it is useful to introduce a degree of simplification index $\mu = 1 - \eta$.

The system (Fig. 59) operates in the context of a source of task requirements (PS or problem source since task requirements will later be identified with problems). The subject A receives problems that have been simplified to a variable degree, μ , by the controller, B , which may thus be regarded as

cooperating to a variable extent in task performance. Differently phrased, B adjusts the level of difficulty, η , according to the stipulated algorithm. For this purpose, B must be designed to determine the value of ρ , typically, as a correct response rate or a sequential correct response function (so many correct responses in a row) or a Waldian- or Bellman-type estimate based on correct-response sequences or, for a quasi continuous task, an index of unwanted deviation. Hence the ρ computing box is preceded by a stimulus (x) response (y) comparator that determines correctness according to whatever x, y relation is ordained by the task. This comparator output may also furnish a 'knowledge of results' signal. (Or, to parody the stance of a behaviourist, a 'causative reinforcing stimulus' since 'knowledge of results' has no meaning within strict behaviourism.) The output from the ρ computation box (an averager) is compared with the constant term ξ and the simplifying/difficulty adjusting algorithm is executed.

Whereas a classical performance characteristic would be a learning curve like Fig. 60, the performance characteristics of this system are stated (as a similar idealisation) in Fig. 61. The steady state of $\rho \approx \xi$ is maintained by an appropriate variation in η . Certain advantages, and some disadvantages as well, stem from this fact.

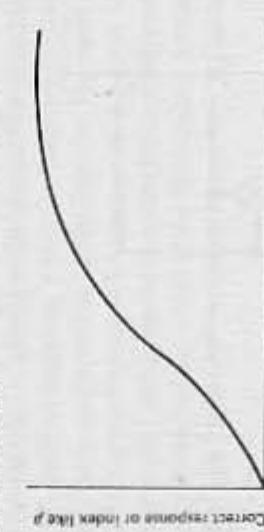


Figure 60 Typical learning curve. n , number of trials.

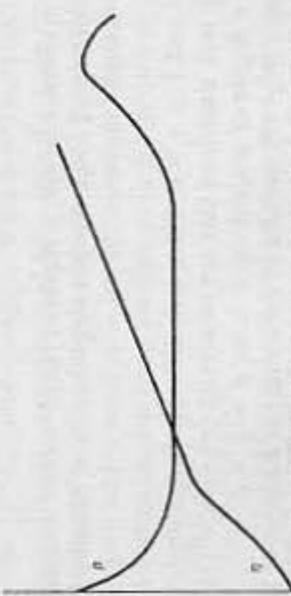


Figure 61 Performance characteristics of a steady-state system. n , number of trials.

1. The advantages are (a) that the subject works in his operating region where he is able to attend to the task which the experimenter deems relevant, and (b) that learning effects in the subsystem A are compensated by the action of B so that the joint system under observation is not a learning system or an adaptive system but a control system pure and simple, the interaction within which is directly observable.
2. The disadvantages are (a) that the joint system is relativistic (in a rather degenerate sense; the observable interactions are of A relative to B and B is a surrogate participant observer), and (b) the advantages only apply if the joint system is stable, which it might not be. The question of stability is obviously crucial. It will be shown, by citing examples, that such arrangements are empirically stable, for a certain to-be-specified class of environments and tasks.

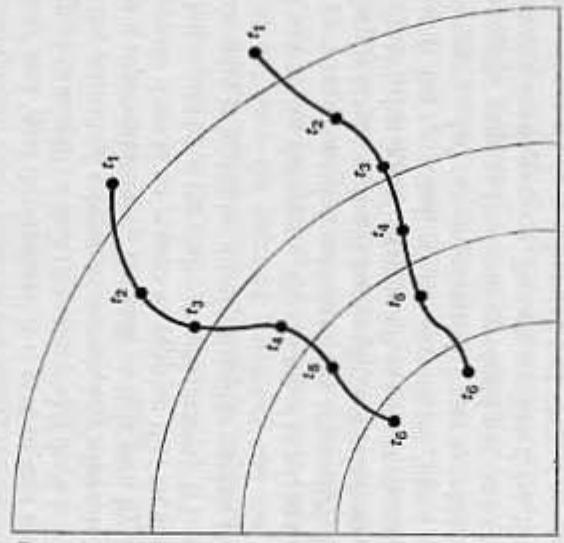


Figure 62 Trajectories for systems that are Lyapunov stable.

The most appropriate formal canon is Lyapunov stability. In the weak sense this condition has the following interpretation. Consider a space with $2m$ dimensions representing the ρ_i and the μ_i , ($i = 1, \dots, m$); $\eta_i = 1 - \mu_i$ with the origin of these coordinates at zero. Since the derivatives of the ρ_i vanish ($\rho_i = \xi$) concentrate on the m dimensional subspace of the μ_i (Fig. 62) on which Lyapunov functions are inscribed as closed surfaces surrounding the origin and nested beneath one another. Represent the system state trajectory (points $\vec{\mu}(n)$ on the surface) as a cyclogram. The

system is Lyapunov-stable if the trajectory eventually approaches the origin and if it never crosses from the inside to the outside of a boundary.

In saying that systems are empirically stable, I mean that their irregularities satisfy this condition, on average, and may be constructed to do so, in all particulars, by the choice of appropriate functions. The proposed method can be elaborated a great deal. The algorithm is extremely primitive, for example, and is easily improved though (as a matter of fact) *most* improvements have little practical value. The controller described above is unidimensional. This restriction is inessential. The majority of steady-state controllers are actually multidimensional devices so that $\rho = \langle \rho_1, \dots, \rho_m \rangle$ and $\eta = \langle \eta_1, \dots, \eta_m \rangle$. The performance criterion is very rigid. That constraint is lifted in at least one of the examples to be discussed.

The crucial and outstanding issues are to do with some half concealed underlying assumptions. Apart from the previously voiced disquiet over the status of 'goals' and 'goal directedness' it is evident that a controller can only be designed to simplify a task (for example) if there is a model for the task and for the controlled subject. Moreover, any adequate model for the task must be a representation of the task as it is seen by the controlled subject (otherwise there is no reason to suppose that an allegedly simplifying operation would have the desired consequences). It is possible to cite numerous examples in which intuitively plausible models incorporated in controllers do *not* have the required properties and the controllers in question fail to stabilise the system. On the other hand, in the examples to be cited the necessary conditions are met quite adequately. Perhaps this is a matter of happy intuition. But it can be phrased otherwise, as follows. If the subject is held in his operating region then he is prepared and he is able to act as though the (physical) stimuli constituted 'input states' and the (physical) responses acted as 'output states'; hence to subserve the external observer's scientific dogma. He could or would not do so unless he were maintained in an operating region. To maintain him there, the experimenter starts with the observer's dogma in mind and designs a controller to hold the subject in an operating region *if* that supposition is valid. In fact, it becomes valid due to the control operation only, i.e. by a boot-strapping process. For the instances to be cited, the boot-strapping works.

1 Error Score Constancy in a Manual Task

The subject is presented with a compensatory display, usually on a cathode ray oscilloscope. In the one-dimensional case shown in Fig. 63 he perceives the locus of a point on a line and is required to adjust a manual control in order to keep the point within a line segment around a fixed position. His manual control determines the acceleration of an idealised vehicle (the

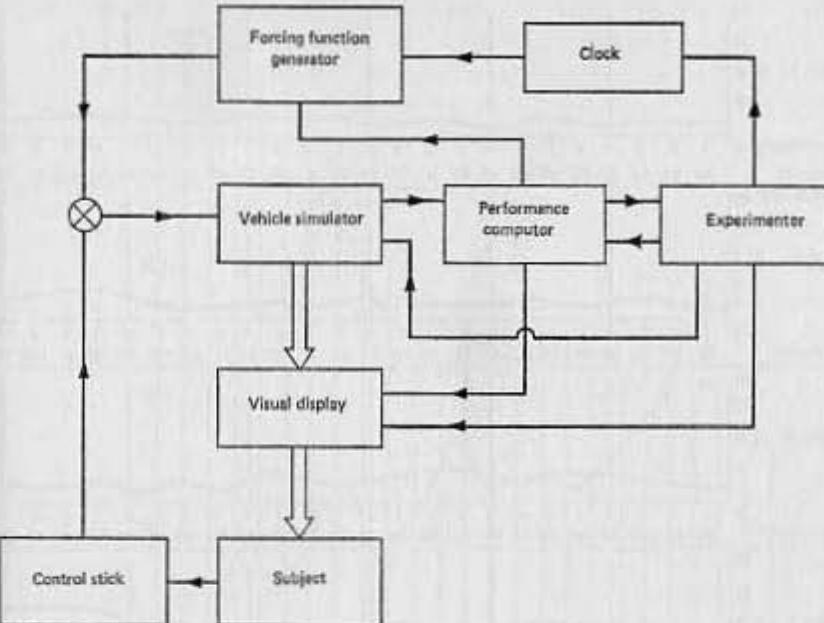


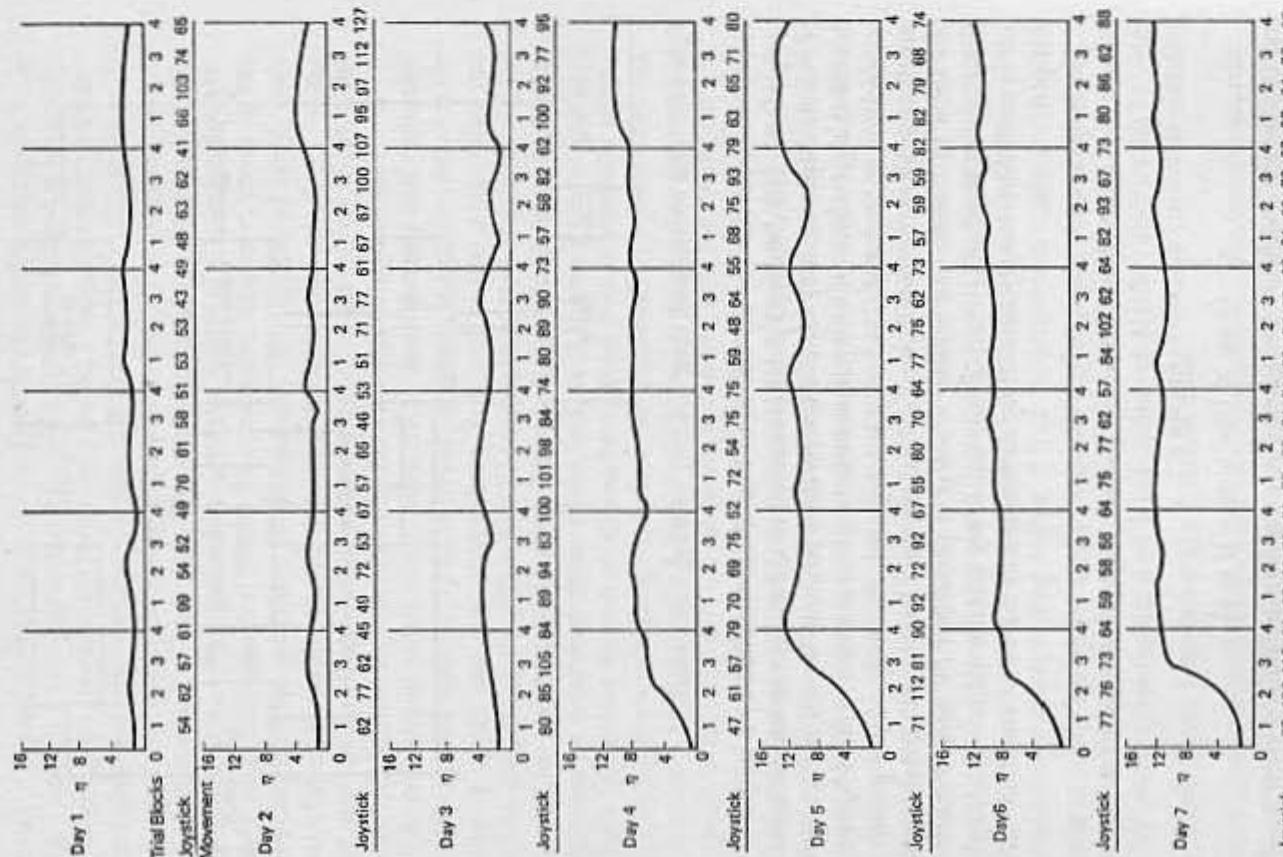
Figure 63 One dimensional, compensatory tracking, task system.

vehicle characteristics of the system in Fig. 63 are specified by a couple of integrators). The displayed point (described to the subject as 'an indicator of vehicle displacement') varies from its fixed position, even if the subject does nothing, because of an input perturbation that is added to the subject's acceleration control signal before it is integrated. The form of this perturbation is unlearnable and hence its value is unpredictable. But the subject can learn to 'handle the vehicle' when it is perturbed by unpredictable disturbances. Further, it is empirically safe to say that the 'vehicle-handling' job is made increasingly difficult by an increase in the mean amplitude of this perturbation.

The subjects error score, γ , (the converse of ρ) is computed either as the average root-mean-square deviation from the fixed point or as the average value of the modulus of this deviation. Let ξ be a critical error score. To maintain constancy of error score, we set

$$\gamma = \text{Const} \times \text{Time Integral } (\xi - \gamma)$$

and choose the (positive) constant so that the man-machine system does not become unstable due to over-compensation. In other words, the task difficulty is continually modulated in order to maintain the chosen form of constancy, that $\xi - \gamma \approx 0$ (or, phrased in a slightly different way, the



(a)

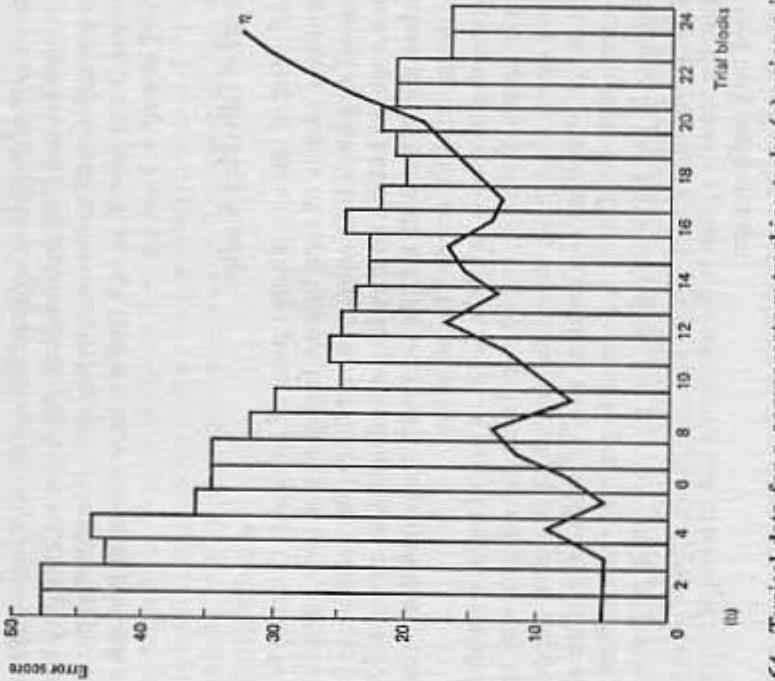


Figure 64 Typical data for a compensatory-tracking task: (a) using a joystick control (Pask *et al.*, 1967); (b) using push-button controls (Pask *et al.*, 1969), with trial blocks administered four per day.

environment is altered in order to maintain a constant relative or subjective difficulty, that is, a 'difficulty seen by a subject' and as indicated by his performance).

Typical response characteristics are shown in Fig. 64 and reveal stabilisation overlaid by a long-term adaptation effect. Systems of this type and their multidimensional analogues have been investigated by Hudson (1962) and Kelly and Prosin (1968) in the United States and, in Britain, by Gaines (1968; 1969b) and myself (Pask *et al.*, 1967; Pask *et al.*, 1969b). One early application, to pursuit radar, is reported in Pask (1957). Only, in this case, the control operation consisted in manoeuvring a vehicle in two spatial dimensions and the manoeuvres were required to satisfy additional conditions regarding a safe attack method; a primitive version of this equipment was demonstrated and gave rise to a good deal of amusement at the British Psychological Society conference in Manchester in 1956.

Plate 3 shows a multidimensional equipment typical of later work – only

two of the four double-beam display tubes were normally used but each tube can accommodate one bidimensional compensatory task (display spot controlled vertically as well as horizontally) or one bidimensional pursuit task (like the radar equipment, except that the tube diameter is much smaller than a PPI screen).

2 Attention Directing Studies

The chief merit of this complicated piece of apparatus (Plate 3) is that it accommodates studies of how human beings direct their attention. The conventional approach is to obtain a direct index; for example, to find out how often a subject notices an intruding stimulus and, as an embellishment, to correlate the index with the physical magnitude (intensity, size) of a variable determining the form of the stimulus.

An alternative method is to set the subject a pair of tasks; a standard task to which he normally attends, and an interfering task on a par with appreciating (and doing something about) the intruding stimulus. Suppose the standard task is performed at a constant level of proficiency in the absence of any interference, i.e. if there is no need to perform the interfering task. By definition, if the skills interfere, the standard task performance is impaired by performing the interfering task, and so it is possible, in principle, to measure the attention given to the interfering task in terms of impairment of standard task performance.

The trouble is that the constant levels of proficiency are rarely achieved; an even more striking difficulty is that interference overloads the subject so that he is displaced from his *operating region*; a level of loading at which it is possible to perform the standard task. Conversely it may happen, for lengthy performances, that the medium sets in and the subject no longer attends to the standard task (once again, he is working outside his operating region but on the opposite or underload side of it).

A scrutiny of Fig. 64 will show that steady-state conditions are maintained by virtue of the controller's compensating action which keeps the subject balanced at a null-point that (for an appropriate choice of ξ) is within his operating region. Moreover, the controller is able to compensate for any disturbance within the compass of the task determined environment, whether it is due to fatigue/adaptation or due to external perturbations that shift the subject's focus of attention. Moreover, within quite wide limits, the controller ensures that the subject works within his operating region for the main task (tracking or the like). Under these circumstances it is possible to have confidence in a differential measure of the attention which is given to an interfering task. Specifically this main task was used, in various experiments, as the

standard task. Other display tubes in Plate 3 were used to display intruding stimuli (meeting with no response) or, in a different study, an input signal for a simple interfering task that the subject was periodically required to perform. Instead of measuring the influence of different intruding stimuli/interfering tasks by a decrease in the level of standard task performance (γ, ρ) the index of impairment was the compensation ($\Delta\eta$) needed to maintain a constant level of performance ($\gamma, \rho \approx \xi$). The curve in Fig. 65 is typical of most subjects.

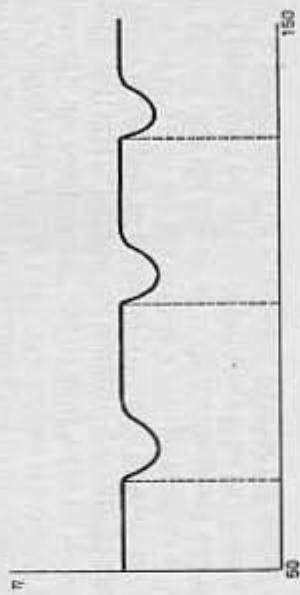


Figure 65 Form of curve for typical subject in compensatory-tracking task showing the effect of intruding/interfering stimuli. Dotted lines indicate onset of interfering stimulus.

The steady-state or null-point method of measurement leads to far more constant results than the conventional techniques; a point made independently and in different contexts by Gaines, Kelly, Seigman, and Sime. In connection with attention directing, it was possible to compare two kinds of measure: (a) a conventional differential measure; increase in r.m.s error (decrease in performance) for an *uncompensated* tracking task, and (b) a steady-state differential measure; decrease in η (to maintain constant r.m.s error on the standard task). The interfering stimuli were the same (moving visual) presentation in each case. Unit scores were taken as the area beneath the deviation in main task performance curve; the unit scores were averaged, for each subject separately, over repetitions of the same kind of interfering stimulus to obtain two mean scores calculated from (a) and (b) for each subject, indicating the attention he had given to each kind of interfering stimulus. For a given subject there is no difference between the ranking of 'attention given' by criterion (a) or (b), i.e. a particular individual consistently treated the interfering stimuli as more or less salient regardless of the method employed. But the mean score variance is very much greater for the (a) scores than it is for the (b) scores and this difference is exhibited consistently by all subjects.

3 Systems for Maintaining a Constant Level of Vigilance

In Fig. 66 the subject looks out for an important or relevant event in a background of irrelevant disturbing events. In my laboratory these events have been specific subpatterns (X_i) in a visually presented series of patterns. A correct response mode is predefined for each subpattern and the patterns may be ranked by importance, as presenting, for example, various levels of hazard. This refinement is not considered in the sequel where stimuli (x) are taken to have equally urgent implications but to differ either in form or in the type of action needed or in both. The subject is required to make the correct response if the event occurs, and he must make no response if there is no relevant event; that is, he must not hallucinate events. The arrangement described below is designed to maintain a relation between the subject and his environment in which the subject has a constant degree of vigilance. It is well known that the degree of vigilance decreases and the subject's selective and perceptual behaviour is impaired if he has been a long while at a job and, in particular, if he is fatigued. Further the degree of vigilance

varies differentially according to the class of relevant event (a subject may be competent to deal with some sorts of events, but incapable of dealing with others).

In Fig. 66 there are m classes of relevant events (X_i ; $i = 1, 2, \dots, m$) each to be dealt with according to a rule \mathcal{R} . In other words, the desired response y , given event x , is $y = \mathcal{R}(x)$ where y is selected from a set Y and where x is selected from a subset X_i of X (the set of all relevant events including the null event).

A basic design assumption underlying this type of system is that off-loading is possible, i.e. there is an auxiliary data processor in parallel with the subject (it may be another human being) which can take over the subject's job. Thus, unknown to the subject, it is occasionally possible to delegate his job to the other data processor and to do without the subject's services. The off-loading intervals must be spaced fairly uniformly in time. They are used for the processing of 'test signals' which, so far as the subject is concerned, are indistinguishable from relevant events. His behaviour with reference to them and to reality is assumed to be and actually is identical.

Constant vigilance control is most readily described under an artificial supposition that the control mechanism is able to inject a test signal at any instant, although it aims to minimise the amount of time spent in this fashion. This assumption is relaxed or even renounced in practical applications of the system; for example, the control of inspection and sampling vigilance.

The control mechanism first computes the conditional frequency estimate,

$$z(n) = \frac{\text{Number of test trials between } n \text{ and } (n - k) \text{ for which } x \in X_i \text{ and } y = \mathcal{R}(x)}{\text{Number of test trials between } n \text{ and } (n - k) \text{ for which } x \in X_i}$$

so that as t increases, $z_t(n)$ approaches the conditional correct response probability $p(n) = P(Y = \mathcal{R}(x)|x)$, $x \in X_i$ for $t = t_n$.

The vigilance at the n th trial (or at the instant t_n) is estimated by

$$V(n) = V(t_n) = 1/m \sum_i z_i(n)$$

on the assumption that the probability, $\pi(n)$ that $x \in X_i$ at the n th test trial approaches $1/m$. This assumption is rendered plausible due to the activity of the control mechanism as it is described below.

The first part of the control operation determines the probability $R(t)$ that some test trial will be made in the interval between t and $(t + dt)$ where dt is short. (If a test trial is made in this interval and if it is the n th test trial then t_n is between t and $t + dt$.) The present controller is designed so that

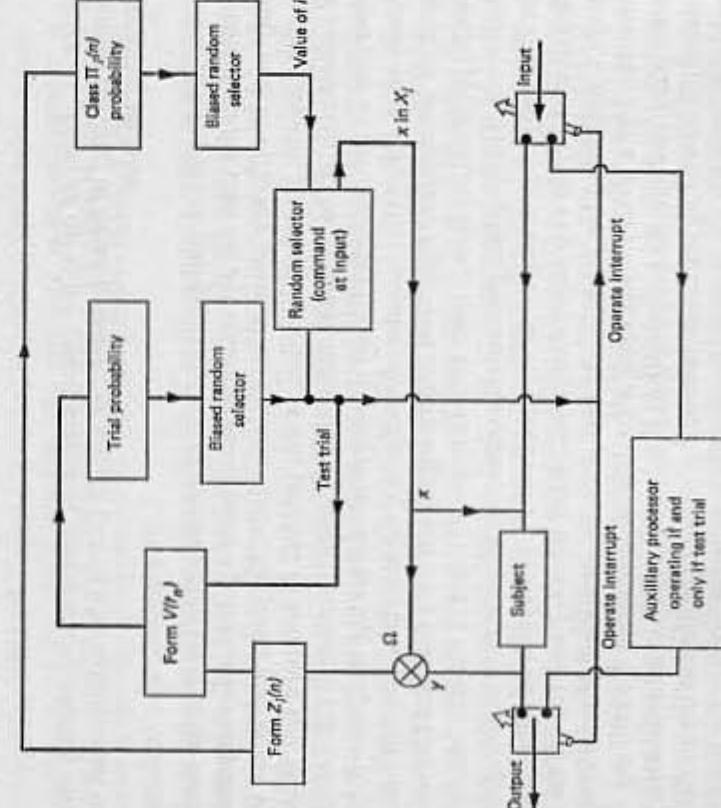


Figure 66 Vigilance system.

$$1. \quad R(t_n) \simeq 1 - V(t_n - 1)$$

The next part of the control operation selects the class of events from which one event will be 'randomly' sampled at $t = t_n$, if some test trial is made on this occasion. The selection probability for the class X_i is denoted π_i ; and the controller of Fig. 66 is designed so that

$$2. \quad \pi_i(n) \simeq 1 - z_i(n)$$

Finally each of the terms $R(t_n)$ and $\pi_i(n)$ is interpreted as a bias applied to a probabilistic selection device (a variable 'window width' random address generator). The effect of equation (1) is to increase the test trial rate as the subject's vigilance decreases, and the effect of equation (2) is to *rehearse*, most frequently, the most neglected events. The latter process maintains the average values of the $\pi_i(n)$ close to $1/m$ as required in the definition of vigilance.

As a theoretical aside, it is worth commenting that the recursive definition embedded in the design is typical of relativistic systems. Here, vigilance is defined on the assumption that the $\pi_i(n)$ are nearly equal. There is a *rule* (namely a psychological principle of familiarisation due to repetition) which, if valid, justifies the last control operation as an expedient for securing equality of the $\pi_i(n)$. Since the *rule* is psychological, the entire operation is a *heuristic* recursion.

4 Maintaining a Constant Proficiency in a Coordinate Transformation Skill

The subject is presented with the display and response board shown in Fig. 67. Each trial, Δt apart, he is presented with a figure in the alphabetic display and required to select the row and column buttons that designate the row and column coordinates of the figure in the 3×4 rectangular display before δt seconds after the stimulus. With δt set at 2.5 seconds, Δt at 5 seconds, and the between-trial interval at 1.5 seconds this is a difficult job, and the novice is unable to do it unless he is provided with cueing information.

The skill is characterised by a couple of error factors (in the sense of Harlow, 1959), namely, a row and column error factor, since the subject's response may be row-correct and column-correct independently. Hence, the cueing information is delivered with reference to the row-response selection and the column-response selection separately. The information concerned is provided by the row lamps and the column lamps of Fig. 67 which are illuminated at a variable interval (within the allowed interval of δt seconds) after the appearance of a stimulus figure in the alphabetic display.

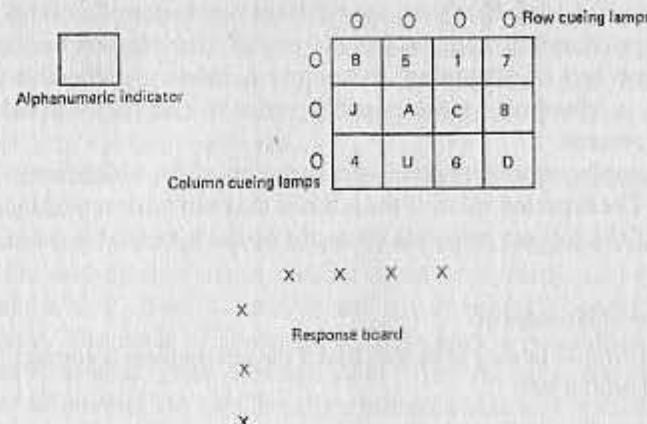


Figure 67 Coordinate transformation task showing display and response facilities.

Let i = 'column' or 'row', and let ρ_i be a proficiency index

$$\rho_i = \text{Average value } (\phi_i)$$

where

$$\phi_i = \begin{cases} +1, & \text{if the } i\text{th type response selection is correct and presented} \\ & \text{before the } i\text{th type cueing information;} \\ 0, & \text{if the } i\text{th type selection is correct but too late;} \\ -1, & \text{if the } i\text{th type response selection is mistaken or absent} \end{cases}$$

The value of +1 in this rule may, with advantage, be replaced by the term $\Delta t - \text{latency}$. Define η_i as the proportion of Δt that elapses before the i th type of cueing information is delivered to the subject. The maximum value of η_i is Δt and if $\eta_i = \eta_{\max}$ then no cueing information of the i th type is delivered.

The control mechanism for variable delay cueing satisfies

$$\eta_i = \text{Constant} \times \text{Time Integral } (\rho_i - \xi) dt$$

where ξ is the required level of proficiency and the constant is a positive rate term. The system is started in the initial condition $\eta_i(0) = 0$. The value of η_i is defined unless or until $\eta_i = \eta_{\max}$.

4.1 Clamping Techniques The assumption underlying simplification by a variable delay cueing procedure, with Δt constant, is that the subject operates decisively (to select responses) at the same rate throughout the experiment. Given this, the cueing information can be delivered at an optimal (constancy-maintaining) instant during the decision process taking

place at a given trial. However, the constant rate assumption is (at most) a crude approximation and the serious use of variable-delay cueing rests upon some sort of 'clamping' technique; a technique whereby the trial duration is 'clamped' to the possibly variable interval occupied by the decision process.

The clamping technique involves a variation of Δt , which is now written as $\Delta(n)$. The expected value of the interval that will be occupied by decision making if the subject responds correctly at the n_0 th trial is estimated for $n_0 > n$ as

$$t^*(n_0) = \begin{cases} \text{mean value of } \\ \{\Delta t(n) - \text{latency at the } n\text{th trial if the } n\text{th response is correct}\} \\ \Delta t(n) \text{ if not} \end{cases}$$

To complete the feedback loop

$$\Delta(n_0) = t^*(n_0) + \gamma$$

where γ is positive (thus increasing the pace as proficiency improves).

Since η_i is defined proportionally, the cueing delay becomes a proportion of $\Delta(n)$ rather than a proportion of Δt and the entire process is variably paced.

4.2 Results Some response characteristics are shown in Fig. 68. There is overall adaptation as the subject comes to grips with the skill but, before that, the interaction between the row and column components of the skill is manifest as a pronounced, but duly compensated, interference effect.

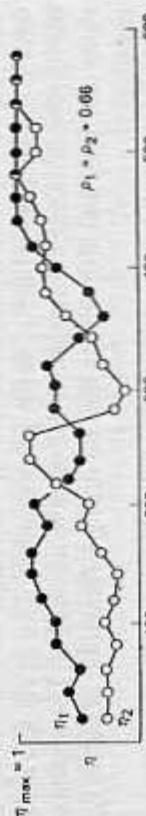


Figure 68 Coordinate transformation task. Response characteristics for row difficulty level (η_1) and column difficulty level (η_2). There is appreciable interference before adaptation.

The clamping operation keeps the subject working within his operating limits with respect to both components of the skill. Under these circumstances it is possible to perform an 'internal' study of attention directing. Either component may be regarded as the *standard* task of the previous example and the other component as the interfering task. Only here there is no initial polarity or task dominance and both tasks (being, in fact, part of the same skill) are performed continually. It is still possible to obtain a

profile of the microstructure of adaptation and to determine how this profile varies from subject to subject, or under the influence of perceptually discriminable changes (making the row lamps brighter than the column lamps), or value assignments telling the subject that a row response is more important than a column response.

All of these measurements have practical consequence. In summary they are as follows.

(a) It has been found that, although all subjects experience some interference, the method they use to combat its adverse effects, as well as the rate at which they do so, is variable and a very useful individual ability discriminant. The mode of combatting interference is generalised for one subject over several quite different skills. For example, the following categories of method are significantly distinguishable and reliable as test indices (predictions between different test tasks, for each individual at the 0.1 per cent level of significance and score differences between groups of individuals using different methods also at the 0.1 per cent level; thirty-five subjects).

Method 1: 'One at once'; an oscillatory behaviour with long cycles (like the data in Fig. 68).

Method 2: 'Assemble subskills'; (the η_1 curve and the η_2 curve increase, more slowly together. At points of difficulty there are very short but high amplitude cyclic fluctuations).

Method 0: 'Bash on regardless' (like the 'assembler method' except that the subject in question apparently has no assembly rule. Adaptation is very prolonged; even for this skill in the order of 1500 trials rather than the 500 to 750 trials of Method 1 or Method 2).

(b) The effect of value assignment instructions is in the expected sense (a subject attends more to the positively valued component) but the notice taken of these instructions differs from subject to subject though (between instructions) it is consistent for any one individual. It is tempting to advance this measure as an index of susceptibility to authority, and I believe it is legitimate to employ it as an index of the extent to which individuals are prepared to accept arbitrary value edicts.

No serious work has been done in my own laboratory upon the influence of perceptual variables such as intensity of illumination. Kelley, on the other hand, has done a great deal and reports on the current state of the art in his compendious manual (Kelley and Kelley, 1970).

4.3 Different Skills It is noteworthy that very similar results are obtained for tasks in which the component subskills are rapidly alternated over blocks of trials rather than performed in the course of the same trial. Figure 69 is data culled from a code-learning study; the two subskills

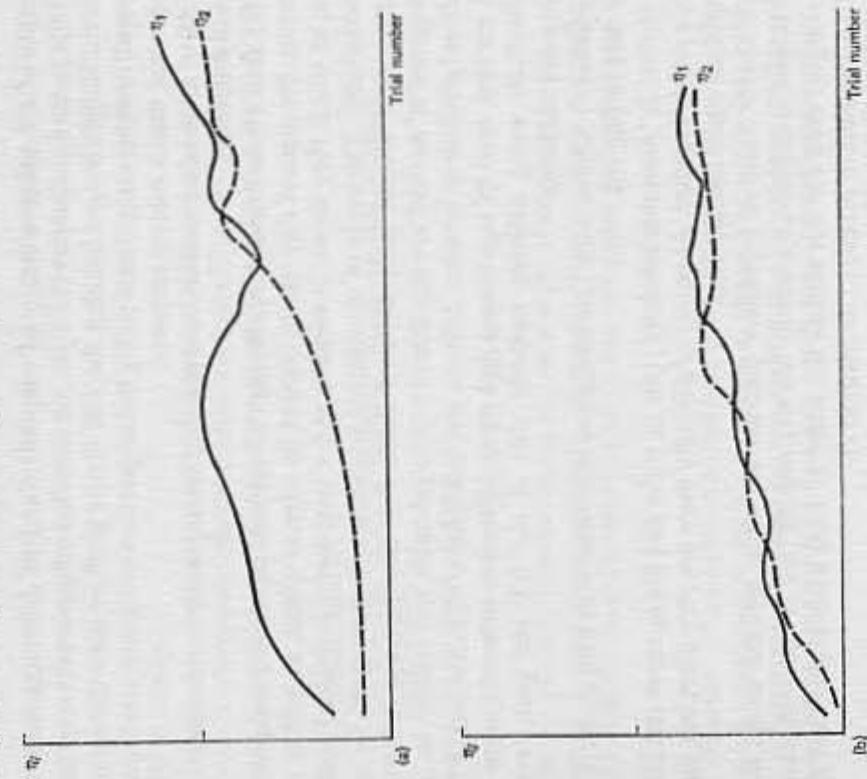


Figure 69 Typical data from code-learning experiments (Pask and Lewis, 1967):
 (a) method 1; (b) method 2.

corresponding to two distinct coding rules, applied alternately. This is one of the test tasks used in order to obtain the results in (a) and there is an obvious difference, even to the eye, between the *Method 1* subject and the *Method 2* subject. Here, the effect of clamping is achieved by a different technique; 'probabilistic alternation' which forces the subject to rehearse the subskill in which he is least proficient with greater frequency (however the rehearsal probability is not allowed to reach 1 or 0; and in the 'minimal condition' both skills are rehearsed with equal frequency).

4.4 Complex Forms The task employed can be quite complex and, if so, much more information can be garnered whilst it is performed. For example, we have carried out experiments (mental tests perhaps) using a

six property classification task for visually displayed ambiguous patterns. The subject is required to assign the value +, -, or null to each of six properties designating features such as 'closure' or 'curvature' of a displayed figure, and to do so for each one of a series of patterns which are presented on a cathode ray tube against a 'noise' background of constant intensity. To aid the subject, he is provided with cueing information in the sense that one feature (the 'upper curvature', for example) is brightened relative to the constant intensity 'noise' background as the pattern is inscribed on the screen. The display and response equipment is shown in Fig. 70. It receives an input from and provides an output to an adaptive device of the kind already described.

Clamping conditions are used, as in the coordinate transformation skill, so that $\Delta(n)$ estimates the subjects expected decision time. Within this

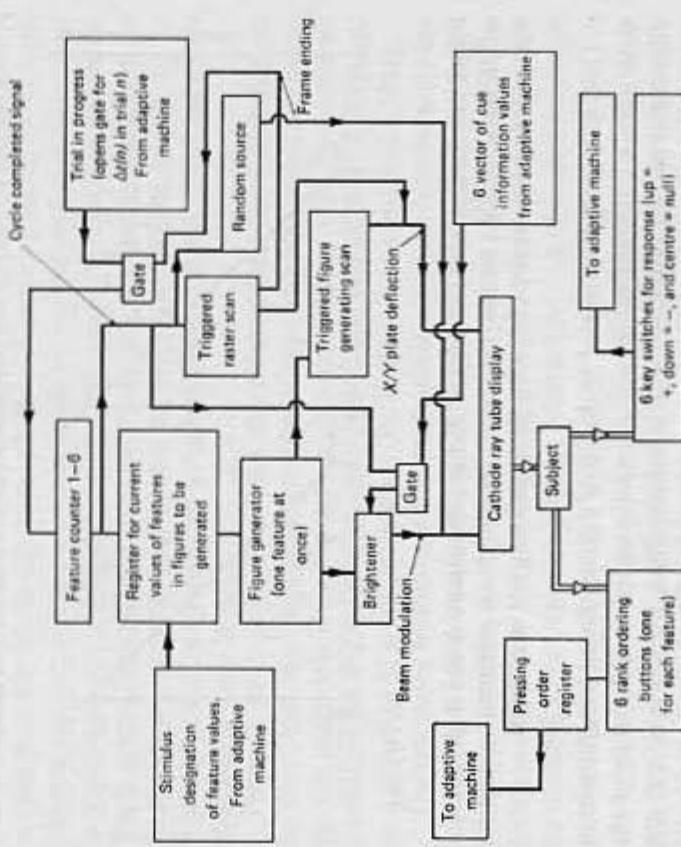


Figure 70 Ambiguous figures task. Figures with six relevant features appear against a random background and are classified by assigning actual values to each feature, via six key switches. For example, the feature *Upper curvature* has possible values: Right (+), Left (-) and none (null, a straight line). If cueing information is delivered in respect of a given feature, the 'writing' beam is 'brightened' whilst that feature is (repetitively) inscribed on all occasions in $t(n)$ subsequent to delivery; hence feature is made prominent against random background.

interval, cueing information is delivered sequentially and delayed with respect to the decision time coordinate. Giving cue information at trial n (with respect to a particular feature/property) means that the feature's inscription is brightened/intensified at the trial in question after and only after a 'cueing threshold' is exceeded. The cueing process is illustrated by Fig. 71 in which the main parameter is the area under the waveform which sequentially switches off the brightener in the cueing display (a response is correct in classifying the i th property if the classification is correct and the response denoting it occurs before the delivery of cueing information about the i th property). Specifically, it is assumed in Fig. 71, (1) that cueing information determining the values of properties a, b, c, d, e and f is removed in alphabetic order, with the f value removed first, and (2) μ is modulated to maintain $\rho = \xi$, where ρ is an average taken over the subject's performance ρ , with respect to each of the properties in the classification task.

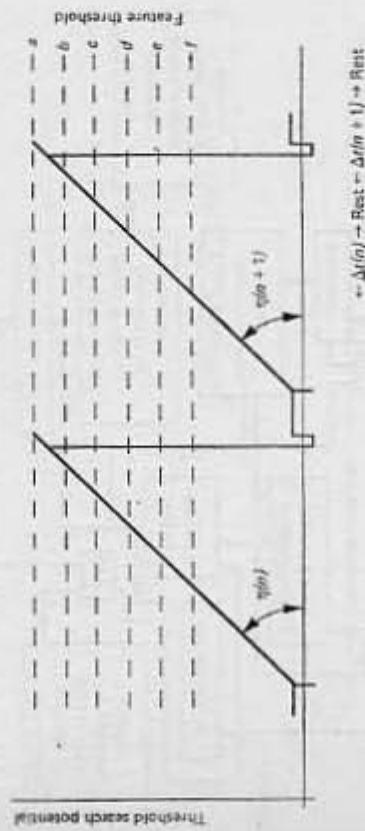


Figure 71 Arrangement for delay of cueing information. A cue is given if and only if threshold for a feature is exceeded. Thresholds are permuted by subject's rank ordering and, on any trial, some thresholds may not be exceeded.

Used as a multidimensional analogue of the coordinate transformation task, the six threshold values are adaptively varied as a function of the differential correct response to each feature. Thus, for $i = a, b, c, d, e, f$, values of proficiency variables ρ_i are computed as functions of specifically correct responses. Specific cueing thresholds $\mu_i = 1 - \eta_i$ (difficulty is increased by lowering the cueing threshold) are regulated individually to values such that $\rho_i = \rho + \text{constant} (1 - \rho)$. In this case, the system yields data of the kind already described.

For a skill of this type, however, the η_i values may be given an interestingly different interpretation; as indices of psychological similarity or dissimilarity. For example, if η_i and η_j ($i, j = 1, \dots, 6$) have the

same value then it can be argued that feature i and feature j are psychologically equivalent. Conversely, if $\eta_i \neq \eta_j$, it can be argued that feature i at 'strength' (or more accurately 'with information') μ_i is psychologically equivalent to feature j at strength μ_j .

4.5 Other Types of Output Data In pursuit of this idea (and as an alternative method) suppose the arrangements for specific η_i (or μ_i) adjustments are deleted, but the subject is allowed to rank order the features or properties a, b, c, d, e , and f so that he receives cueing information sequentially and in his preferred order (rather than an arbitrary, alphabetic order). For this purpose, he is provided with six response buttons which apply electrical charges to capacitors associated with each of the six property specific variables that determine the feature thresholds. The higher the charge on the capacitor, the higher the value of μ_i or, the lower η_i . The subject pays for this facility in terms of a commodity ('money in his bank balance'); in fact, charge on a main capacitor but the 'monetary metaphor' is useful and very intelligible so that he cannot be utterly improvident. These μ_i charges leak away from the capacitors at a fixed rate and in order to maintain a property-ranking the subject must continue to pay for its reinstatement. The entire process is constrained so that the bank balance is kept nearly equal to η .

From the machine's point of view, the informational status is unchanged by the operation of imposing a rank ordering pattern although it is changed by the automatic adjustment of $\mu(\eta)$ to maintain $\rho = \xi$. This example introduces the technique of monetary regulation that has been widely used in the control of group interaction (Pask, 1964; Lewis and Pask, 1964).

The remaining examples of the steady-state techniques have been chosen to show that a performance index need not be *absolute* in the sense that it relies upon an external criterion of rectitude; either a rigid or a probabilistic criterion. All indices like ρ undoubtedly do so. As a technical trick they can be modified to rest upon disjunctive criteria (any one of several responses are correct, not just one response that is correct) but this expedient does not change the picture in an *essential* way. The self-consistency indices used in the sequel are, however, quite distinct from an absolute performance index though their values are used by the controller in much the same fashion as the values of ρ .

The examples also demonstrate that a steady-state system, though in one sense *convergent*, can be used to foster essentially divergent behaviour (invention of classifying attributes; innovation of fresh rules).

4.6 Classification Tasks To begin with, notice a common perceptual oddity of almost any human being. If he is free to classify objects the human



Figure 72 Examples of chequerboard patterns used in the attribute-selection task.

The stimuli were arranged in blocks or subsequences of about fifty items and were all of the same type: chequerboard patterns (Fig. 72). Knowing the type of stimulus and the criteria to be satisfied, the subject first selects the attributes that are to be used by naming response buttons that determine the value +, - or null he assigns to an attribute. The criteria are simply that any chosen classification is self-consistent and informative.

After dealing with a block of stimuli by classifying them, the subject is allowed to rename the classifying response buttons (to redefine his attributes). The experimental data consists in the list of the attribute names that are recorded after each block or sub-sequence of stimuli. Within a particular block of trials, the subject assigns names as labels to a maximum of eight buttons; he may use less than eight if he wishes, providing that his categorisation is informative and also self-consistent. Before the experiment was conducted, we described carefully what is and is not informative and self-consistent. If the subject had any difficulties, these were discussed.

5 Maintaining a Constant Level of Ambiguity of Visual Patterns

Some years ago (Pask, 1964a,b) Lewis and I attempted to induce problem posing or attribute choice in a task of this type by regulating the level of stimulus ambiguity, using the equipment shown in Plate 4.

5.1 An assignment of attribute names is informative if the subject can use the named attributes to divide the stimuli into coherent subsets. An attribute like 'being a chequerboard pattern' is not informative because the

being is prone to use familiar attributes as the basis for his classification. These are idiosyncratic, but tend to be banal and easily named rather than immediately perceived or mentally manipulated, e.g. geometrical attributes like rectangle, square, etc., or exact physical magnitudes. The experimenter can avoid this tendency, by telling the subject in advance that he *must* use particular classifying attributes. In visual classification for example, the subject may be told to classify each member of a sequence of tachistoscopically-exposed pattern stimuli according to attributes such as 'size', 'number of distinct parts', 'circularity', and so on; these names are assigned, in a mechanised experiment, to response buttons that are pressed after each tachistoscopic exposure. This expedient fixes the attribute names and thus prevents an experimenter finding out how the subject chooses new attributes. But, to reiterate the original statement, if a subject is allowed to name the buttons as he desires (to select his own attributes) he nearly always reverts to familiar feature names and, as a result of this, his behaviour exhibits problem-solving according to familiar methods rather than problem-posing or the construction of novel attributes, tests or methods.

These comments about the freely choosing subjects are fairly accurate if the stimuli are unambiguously presented and if the subject is able to use his collection of descriptive attributes in a consistent and successful fashion. If the presentation is ambiguous, the subject may be forced to relinquish his original choice of attributes (because he cannot use them consistently) and to select new ones. Although re-orientation takes place when ambiguity is introduced (conveniently by reducing the tachistoscope exposure interval) there is a delicate balance between the tendency to search for new attributes and a tendency to give up altogether, in the belief that it is impossible to make sense of the environment.

In a free classification task the subject is required to select classifying attributes of his own choice in order to satisfy criteria that depend upon the attributes chosen; for example, to produce an informative and self-consistent classification of the patterns or objects by employing the attributes he opted to specify.

Parenthetically, such tasks have an inherent interest because of their relation to Kelly's (1955) repertory grid technique for eliciting constructs (of which a freely chosen attribute is the limiting case). The constructs, here treated as attributes, are generated by the subject (Bannister and Mair, 1968) not given by the experimenter.

subject knows that all of the patterns belong to the class it defines. Nor is an attribute informative if it can never assume a positive value. Ideally, an attribute should dichotomise the universe of patterns into one class for which it assumes a positive value, and another disjoint and equinumerous class for which it assumes a negative value. There should be a minimal null valued residue. This ideal can rarely be achieved and we did not insist upon a very close approximation.

Self-consistency is explained in commonsensical terms and the subject does not know how it is measured. He is provided with a self-consistency score (the digital indicators in Plate 4) chiefly to convince him that the quantity is not chimerical. In fact, the subject is regarded as self-consistent in his use of an attribute if, when presented with the same *stimulus* upon several occasions, he assigns the same value to each non-null attribute. Any unused attribute, or any attribute present on one occasion and absent on another, is regarded as null-valued. Hence, stimulus repetition (and comparison of the initial and subsequent values of each of the attributes) is the device employed to measure the subject's degree of self-consistency. The machine calculates the self-consistency value by reference to the last few attribute selection values which it stores in respect of repeated stimuli. At any rate on most occasions, the subject is unaware that an identical stimulus is repeated. From his point of view, self-consistency is a matter of using his own selection of attributes in a consistent fashion.

5.2 The calculations needed to obtain a consistency index are straightforward except for the notation involved. To simplify the description, let $y = y_1 \dots y_n$ represent the response elicited with respect to the pattern displayed at the n th trial $x(n)$ which is, in fact, a particular pattern x . Thus $x(n) = x_i$. Let $y^* = y_1^* \dots y_n^*$ be the response to the last *presentation* of x_i (at some previous trial $n - k$; so that $x(n - k) = x_i$). Each entry in y or $y^*(y_j; j = 1, \dots, 8)$ has a value + or - or null, since null is assigned if an attribute button has not been labelled. Let $r_j = 1$ or 0 and let $r_j = 1$ if and only if $y_j = y^*$. A consistency score for the n th trial is obtained by comparing stored response vectors y, y^* , as

$$\frac{y_1 \dots y_n}{y_1^* \dots y_n^*}$$

and $r(n) = \sum_j r_j / 8$. The value $r(n)$ of self-consistency is the sum of these terms,

$$r(n) = \frac{1}{8} \left(\sum_j r_j (n) + r^{\square}(n) \right)$$

Block in
which n
is situated

where $r^{\square}(n)$ is the value of $r(n)$ for the last block. It will be recalled that the subject is permitted or encouraged to relabel his attribute buttons at the end of each block, and this computation is based upon the idea that he does so. A more provident estimate is obtained by carrying out a decremental summation over many blocks (the experimenter enters the names of attribute buttons that have been relabelled at the beginning of each block, so that the machine can count these attributes as null valued for the first of its calculations).

5.3 The 'informativeness' measure was used mainly to ensure non-triviality. For this purpose, it is sufficient to ensure that any labelled attribute is assigned at least one '+' value and (for a different pattern stimulus) at least one '-' value in each block. A check upon this condition is easily introduced by a circuit which detects, for any labelled attribute, the occurrence of a *first* non-null value and the occurrence of a *first different* non-null value. Any attribute button which has both a *first* non-null and a *first different* non-null value in a given block is deemed 'permissible' and the crude 'informativeness' index, displayed to the subject, is merely the number of attributes that are deemed permissible. In practice the subject was required to maintain this index at a value in excess of two.

Various refinements are possible and some of them have been used. For example, it is easy to retain all values of y and y^* over a block and, in calculating $r^{\square}(n)$, to discount the contribution of all other than permissible attributes. Again, it is easy to compute a genuine informativeness index with respect to attributes regarded as *independent*; namely, if N_j^+ is the number of stimuli in a block for which the j th attribute has $y_j = +$, if N_j^- is the number of stimuli in a block for which the j th attribute has $y_j = -$, and if there are N stimuli in a block, the selective information function for each block of trials is

$$Attr\ Info = \sum_j \frac{N_j^+}{N} \log \frac{N_j^+}{N} + \frac{N_j^-}{N} \log \frac{N_j^-}{N}$$

It is more difficult to estimate the informativeness with respect to the subject's statements about patterns (at least, there is more latitude in choosing the appropriate form of expression).

5.4 Several types of steady-state controller were employed (using $r(n)$ in place of $\rho(n)$, and subject to the overriding constraint that the attribute values selected are minimally informative). In one scheme (Pask and Lewis, 1964; Pask, 1964a) the pattern blocks, including repetitions, were predetermined and the slides for each pattern were arranged in a projector cassette.

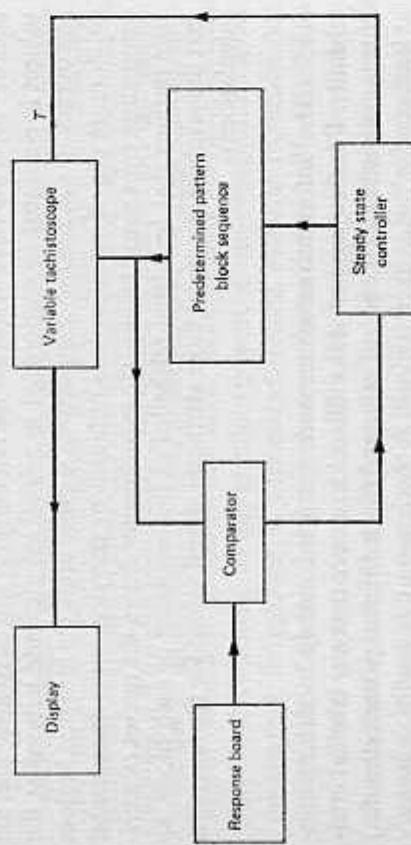


Figure 73 Steady-state control system for maintaining constant level of stimulus ambiguity.

The variable determining the ambiguity (analogous to the difficulty variable η) was the tachistoscopic exposure interval T of each stimulus slide using a minimum value of 200 ms (reached if $\eta = \eta_{\max} = 1$) and a maximum of 2.1 s. T decreased with η so that the exposure at each trial maintained $r(n) \approx \xi$ as in Fig. 73. Under these circumstances, the subject is bound to deal with an environment that becomes increasingly ambiguous as he learns to achieve self-consistency and typically he does so by choosing fresh attribute names. Most subjects, for example, start off by listing the readily verbalised but pedestrian names discussed in the introduction. Later, they progress to names that are idiosyncratic and often fanciful; moreover, they usually name more of the attributes (there is an obvious penalty in terms of consistency score for naming attributes that are not assigned values; hence, for gratuitously maintaining no longer used initial attributes in the current list). This arrangement can be operated within wide limits but there is a maximum of ξ for which it is impossible to catch onto regularities amongst the stimulus patterns and a minimum value of ξ for which the adjustment is swamped by the equivocacy of pattern classification even if the patterns are viewed for as long as desired.

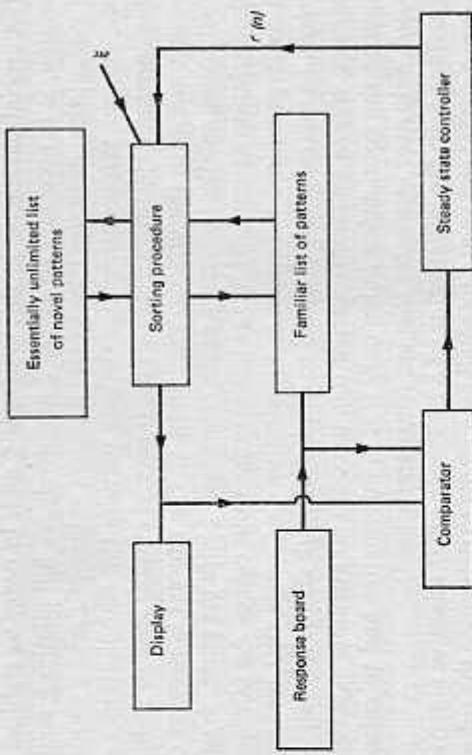


Figure 74 Steady-state controller that maintains constant level of stimulus ambiguity on the basis of differential data about a subject's self-consistency, $r(n)$.

5.5 Another scheme (Fig. 74) is more complicated (Lewis and Pask, 1964; Pask, 1971a). The slide blocks are constructed on the basis of differential data about an individual subject's consistency which is collected, as a set of scores that are averaged with respect to each pattern (over and above the computation of $r(n)$). As a general constraint upon the operation, any slide block that is constructed contains two or more repetitions of each slide so that a consistency value may be defined. Given that, consider the initial block (say with fifteen different patterns and forty-eight slides some containing two or more repetitions). At the end of the block either $r(n) > \xi$ or $\xi > r(n)$. If $r(n) > \xi$ then a sorting algorithm is used to remove a number, ζ , of slides from the block and thus to generate a fresh block. Those slides removed are the slides of stimulus patterns having the highest consistency scores; these are placed on a *familiar* list as candidates for use in a converse process as below. If $\xi > r(n)$ then a number, ζ , of slides from the familiar list (if it is defined) are used to replace the slides of patterns with the lowest consistency score. The sorting operation is repeated after each block of trials. In other words, if the subject detects a regularity in a certain part of his environment, as indicated by a high differential consistency score, the detected regularity, or the data underlying it, is removed from the environment when the overall consistency $r(n)$ is high. Conversely, if $r(n)$ is low, the environment is buttressed by patterns that previously obtained high differential consistency scores and may be expected to make the environment more tractable to the subject.

The crucial parameter in this system is ζ and, in order to stabilise the operation, it is necessary to adjust ζ adaptively. A hill climbing method will suffice; the controller tries widely different values of ζ and ξ choosing as the current value pair whichever one permits the stabilisation $r = \zeta$ by the ongoing adjustment ('rate ζ ') or, if several choices lead to stability, choosing one with the highest 'rate' term ζ . The process is iterated continually using smaller increments or decrements of ζ with the sense, + or -, of the variation determined by the initial or previous essays. An alternative and more provident technique is to select an $\langle \xi, \zeta \rangle$ pair that maximises *Attrinfo*. In this case, it is essential to display *Attrinfo* and to introduce the penalising checks mentioned earlier (i.e. to count an attribute as a contribution to the consistency index only if it is informative and to calculate *Attrinfo* only over attributes that are consistently used).

Provided with the initial value of ζ by one or other adaptive algorithm, the experimenter can vary ξ by small increments to determine the subject's resilience, before stability breaks down. The limiting values of $\langle \xi, \zeta \rangle$ or $\langle \xi, \zeta, Attrinfo \rangle$ provide an index of the subject's tolerances for ambiguity and, it is believed, of his tendency to innovate under circumstances that attach a premium to innovation. The characterisation appears to be stable, so far as an individual is concerned, over a couple of tasks. It was also found that the experience of working for a couple of hours, the length of a typical session, in a system of this kind gives rise to marked emotional effects. Some subjects, usually those who like to innovate, found the experience pleasurable or even euphoric. In contrast, other subjects found it distasteful and, if not allowed to opt out, disorienting (the effect persisted for long enough after the experimental session to prove embarrassing on administrative grounds).

Although these experiments were piloted and initial results were obtained over a decade ago it has been impracticable to pursue the matter in the interim. Work is starting again because there is, at this stage, a sufficiently comprehensive theory of interaction (the conversational theory, discussed later) to account for the affective changes that takes place and the clusters of stable limit values.

5.6 It is of interest to characterise the psychological constancy maintained by this type of SST system. There is a grain of truth in the view that a subject is held on a continuum of which the polar extremities are sensory deprivation (crass underload) and overload pure and simple. The trouble is that these terms are insufficiently exact: overload may mean 'processing rate' for repetitive inputs or 'high complexity', while the meaning of sensory deprivation depends upon the conditions as well as the length of deprivation, i.e. 'sensory deprivation' acts upon several aspects of man;

physiological, psychological, and so on. Moreover, each action has several facets. It is useful to place the constancy on a more specific continuum.

At one extreme, there is a guessing experiment where the subject faces a source which he knows to be random, i.e. he can make no rational sense of its structure. Typically, however, subjects adopt a 'superstitious' behaviour and make sense of the source by imagining a rule or regularity which is locally tenable but, since the source is random, generally unsound. At the other extreme, there are experiments such as those carried out by Warren (1961), Evans et al. (1967), Lilly et al. (1968) and Lilly (1972) in which a subject listens to a monotonous input (say from a tape loop, repeating the word 'kettle'). Under these circumstances, which approximate one kind of deprivation (the minuscule environment could not be *more* regular) the subject hallucinates other words; for example he hears 'pitile' or 'title' instead of 'kettle'. Moreover, these constructions are continued to form chains of fresh words, which suddenly enter perception. In terms of this continuum the SST can be used to balance the subject's operating region between the extremes of 'no rule to learn' and 'no variation at the input'.

6 Commentary on the Components of an SST System and their Interaction

At first sight these case histories bear witness to the success of a ludicrously simplistic method. If viewed as a causal or strictly behaviouristic system the subject has a 'goal' which is simply an observable behavioural regularity and a performance index, ρ , is a measure of the extent to which the actual behaviour fits the regularity the experimenter previously detected as a possible 'goal' and currently has in mind. A difficult index (η) could have a comparable status, i.e. it may have been observed that certain stimuli (or causative input events) give rise to less deviation from the ideal than others and this ordering could be modelled from the *experimenter's* point of view, by something like a nesting of stimulus subsets. If repetition, perhaps with reinforcement, is held to transform this nesting (so that, as a result of adaptation, more difficult stimuli come to act as less difficult stimuli used to act) then the SST is merely a device for maintaining a certain *appearance* of performance by selecting a stimulus sequence, to suit the prevailing level of adaptation, from an appropriate subset.

7 Comment on Goals

These precepts ('goal directed') is an observed regularity, 'performance' is an external observer's construct) have already been questioned and this is an appropriate point at which to assert their insufficiency and counterfactuality. The SST works *in spite* of these suppositions; not because of

them. Moreover, the steady-state system is, in fact brought into existence by quite different means.

(a) The subject is *goal-directed* in the thoroughly full-blooded sense of entertaining an intention or purpose.¹ That is, the subject's brain acts as a general purpose computer which executes goal-directed procedures, alias problem-solving procedures, alias intentions. Generally these procedures are non-deterministic programmes in the sense of Manna (1970) or fuzzy algorithms in the sense of Zadeh (1973).

(b) Under these circumstances, the environment is symbolically interpreted and selected by the subject, so, similarly, is the procedure he executes. The subject attends to a field of attention. The experimenter may determine what he attends to by specifying (as a programming operation) the procedure that he executes. Also, the experimenter may pose problems as conditions in which the goal is not satisfied. But he does not construct an input to which the subject reacts in any other sense.

(c) Under these circumstances the mooted equivalence between difficulty and simplification makes sense. A simplified problem is a partially solved problem and is less difficult than the undiluted problem, in so far as the solution-method used to obtain the partial solution is compatible with the methods realised by the current goal directed procedure.

In general, there is no guarantee that a simplification (partial solution) which seems plausible to the experimenter will serve as a simplification to the subject. In order to satisfy this condition it is necessary to have a model that replicates the subject's problem solving procedure. The cases considered are specially tractable either because the subject's method is dictated or because it is possible to choose an ordering of difficulties (either one ordering indexed by the scalar η or several ordering indexed by η_i) which systematically simplify in respect of *any* method.

(d) The various performance indices, p , are interpreted as estimates of the extent to which an internal-to-the-subject-hypothesis is confirmed or an internal-to-the-subject criterion achieved. All of them are measures of the subject's degree of belief in the rectitude of a goal directed, i.e. intentional or purposeful, action. In addition, of course, they index the extent to which the behaviour under this goal satisfies the external observer's criteria. But, first and foremost, they are estimates of subjective quantities and it is a subjective quantity (the subject's degree of certainty) that is stabilised by the SST. This point is obvious in connection with reflective measures, like r , as these make little sense outside such a framework.

(e) The goal directedness requirements ((a), (b), (c), and (d)) are satisfied by the existence of an experimental contract in which the subject agrees to

1. R. L. Gregory independently makes the same point, in saying that the brain under these circumstances, acts as a *deductive* mechanism.

participate. This is a normative or game-like scheme (not a behavioural paradigm) which is set up by a commanding or programming operation 'giving the experimental instructions'. It is useful to stress that the contract is *agreed* to by the subject. 'Giving instructions' is not a matter of presenting a specially forceful stimulus. The contract has a semantic and a pragmatic aspect. The subject, for his part, agrees to interpret symbols in a particular way; namely, as members of a field of attention. The experimenter, for his part, agrees to provide an SST controller that does so as well. The subject agrees to adopt a *role*, that is, to aim for a certain class of goals and the experimenter, for his part, agrees to furnish a controller that will cooperate (in the agreed domain) by simplifying problems posed under the goal so that the subject may be able to keep his part of the contract. Notably, the controller is designed under the assumption that the contract is kept. It does cooperate and is *able* to cooperate if, and only if, the subject has the intention of keeping the contract in the first place.

(f) Though this aspect is trivial for the case so far cited, the contract also has a syntactic component. The subject and the controller speak the same interaction language (formally an *object language* distinct from the observer's descriptive metalanguage) which has definite grammatical rules, which, for example, determine what is the form of a correct response or the construction of a logical subproblem (subgoal) of the original. The goal-directed procedures are programmed in this language, the rules of the game are stated in it.

(g) Although the theory does not demand a physiological interpretation, it is interesting that the subjective conditions (a), (b), (c), and (d) have reliable neurophysiological concomitants in man. These are described by Grey Walter (1969) who is responsible for much of the original work (this paper is essential reading).

By way of summary, consider a subject as a reactive device in receipt of stimuli that are pure stimuli, and no more than that.

If an auditory stimulus is presented to the subject it is usually possible to record an electrical response from the specialised auditory region of the subject's cortex. The same is true of a visual stimulus, or any other, as each stimulus evokes electrical responses from regions proper to its modality. In addition, for each effective stimulus that does evoke an electrical response, there is a transient electrical response discernible against a background of the usual electrical rhythms, in the frontal regions of the cortex. If traces of activity are laid out in time following the stimulus and if these traces are inscribed in computer storage and their values are averaged over a period of time in which several stimuli are delivered (so as to suppress the unwanted background by cancellation), then the transients are accentuated. But no other changes of response pattern are observed whilst the stimulus remains

literally a stimulus: an event with no symbolic consequence. The picture is altered as soon as the subject is led to anticipate some consequence; either by repetition that indicates a relation between events or a requirement to make a motor response that acts upon the form of future stimuli or (immediately) by verbal mandate, i.e. saying 'this stimulus poses a problem to be solved'. Under these circumstances, the electrical response (sampled after averaging) in the frontal cortex displays a contingent negative variation (CNV) or 'expectancy wave' which increases in magnitude during the problem-solving phase and rapidly decreases as soon as the problem is solved and shortly before a motor response is elicited (supposing that an overt response is called for; in general, the solution may simply be kept in mind). The CNV is due to an interaction between loci dispersed about the frontal cortex and certain arousal systems in the reticular core and those parts of the limbic system responsible for orientation. The CNV can be extinguished by the repetition of stimuli uncorrelated with subsequent events and reinstated either slowly by correlated stimulation or immediately by command. Whilst the CNV is evoked, the brain is acting as a problem solver or symbol processor; it executes a programme with a simple or complex, 'If . . . then . . . else' structure. In these conditions it is possible to speak of information to the system (by the same token of uncertainty, equivocation or redundancy) and the system is goal-directed rather than reactive.

(h) The linguistic transactions of (f) call for a brief comment because 'language' is more than a convenient figure of speech (even though the languages involved in the system so far considered are admittedly quite trivial). Perhaps the most cogent empirical demonstration that a genuine (though trivial) language exists and is not merely 'blessed upon the system, comes from a phenomenon called 'participant interaction' that consists of misusage of the grammar *because* the subject is anxious to say more than he is legally allowed to say.

Under certain conditions, notably when an SST system is operated well below the overload point (low ξ) subjects typically engage in a game-like interaction with the controller. For example, they make responses which they know to be mistaken and which are thus illegal (according to the contract, the subject must make a correct response if he is able to do so, as he is at the moment in question). On examining sequences of illegal responses, or possibly, on interrogating the subject after the experiment, it is clear that the illegal responses are made in an attempt to learn the algorithm in the controller and/or to influence the controller. For example, deliberate mistakes in a unidimensional SST system will reduce η (in a rather complex fashion if ρ is calculated from sequences of correct responses or by using an estimation rule). In a multidimensional system each η_i may be reduced and,

if the subcontrollers responsible for setting $\rho_i = \xi$ by adjusting η_i , interact, as they do in most of the systems to be discussed in the sequel, then elaborate patterns of play and discovery are both possible and observable. If subjects use a system of transactions in an ungrammatical fashion (as they do) either judging from retrospective reports or by inference from records of their behaviour then, trivially or not, there must be a grammar. The construction is buttressed by noting that the phenomenon of illegal usage is inhibited, and the play converted into legal usage, if the subject is provided with transactions that logically allow him to question and issue (conditional) commands to the controller, by means of which he is able to bring about more expeditiously the controller-behaviours he manages to bring about in a very cumbersome manner by illegal usage or misusage of the original transactions. As a point for later reference the augmenting transactions, that allow the subject to interact with the SST controller, operate or comment upon the original problem-solving interaction and are conveniently assigned to a higher level of interaction. In other words, this primitive language is stratified into a lower level concerned with problem-solving and an upper level concerned with comments to the controller about problem-solving.

8 Outline

The next chapter is concerned with an extension of the SST along these lines; namely, adaptive control. One example has already been noted (pattern ambiguity, the adaptive regulation of ξ and Q). The adaptation of the controller, even in this case, compensated for a kind of learning on the part of the subject which could not be satisfactorily represented as a form of adaptation. However, this learning was not, in itself, relevant to the constancy being maintained and it was not emphasised. In the adaptive control systems to be discussed, the constancy is preserved in order to influence the learning process itself (for example, to maximise the individual subject's rate of learning) and the controllers in question are quite explicitly built as teaching and training devices. Since most of them are devices for instructing skills, it is still possible to minimise the attention-directing or exploratory modes of a human being, i.e. to minimise the issue of learning strategy but not to suppress it altogether. Moreover, most of the systems rely upon proficiencies ρ , ρ_i as estimates of an underlying degree of belief.