10

Reactions to reward change

IN establishing that the locale system is integrally involved in reactions to environmental change we purposely omitted any detailed treatment of reactions to changes in biologically meaningful items, though this is central to much of the behaviour we have considered. It is clear from much of the previous discussion that shifts in reward contingencies can often lead to aberrant behaviour in hippocampal animals. In the present chapter we shall provide the basis for understanding these effects.

A wide variety of changes fall under the rubric of reward shifts: changes in the amount of reward, its quality, its locations, the objects with which it is associated, and so on. Though we have already discussed many studies incorporating such changes we have not focused on reward change *per se.* Certain studies are directly concerned with this problem; the most obvious of these involve the complete removal of reward. This leads to what is called the *extinction* of the behaviour previously based on that reward.

According to the present theory, both the locale and taxon systems contain mechanisms which are sensitive to changes in reward conditions and which could lead to the cessation of inappropriate behaviours. As a way of approaching the question of the specific effects of hippocampal lesions upon reactions to reward changes we shall (1) examine the concept of extinction and the general problem of persistence, (2) consider these in terms of the locale and taxon systems, (3) discuss the lesion data concerned with extinction, and (4) turn to the general case of any change in reward contingencies.

10.1. Extinction and persistence

After reviewing the extant theories of extinction, Mackintosh (1974) pointed out that

> 'traditional S-R theory is … left with no account of the learning process underlying extinction' (p. 418).

Partly, this results from the tendency to consider extinction a unitary process and to attempt to explain all extinction phenomena within a single

theoretical framework. Such an approach is likely to be no more successful than its counterpart in learning theory. More specifically, this failure follows from the emphasis of S-R theory upon the inhibition of old learning without any detailed consideration of the new learning which can occur during extinction.

The inhibition conceived by traditional theory flows from any of several sources: (1) *reactive inhibition* (Hull 1943) follows the occurrence of any response, exhibiting both a transient and a more permanent phase; (2) *generalization decrement* occurs because the extinction situation is distinguishable from training, and the trained behaviour is not conditioned to this new situation; (3) *conditioned inhibition* results from the association of non-reward with the trained behaviour; (4) in a less direct way *interference* results from the exploration elicited by the new situation and from competition with any new behaviours. Only the last of these hints at the new learning that can occur during extinction, and it is a rather indirect mechanism. Though all of these forms of inhibition no doubt occur they cannot, singly or in combination, account for all the facts of extinction.

Effects upon extinction are usually gauged in terms of *resistance to extinction,* or *persistence,* which refer to the extent to which behaviour remains unchanged in the face of the removal of reward. As we have seen, abnormal persistence is one of the dominant features of the syndrome produced by hippocampal dysfunction, and an understanding of the processes underlying extinction would shed considerable light on hippocampal function. For this reason we shall consider the problem of persistence in some detail.

In describing the properties of the various behavioural hypotheses we raised the possibility that different hypotheses might be differentially persistent; that is, they would be used with varying degrees of rigidity. Place hypotheses were assumed to be minimally persistent because they are based upon a type of information structure which can be rapidly and easily amended when the environment changes. The unexpected absence of reward should elicit exploration, the map changed, and the previous place hypothesis rendered inoperable. This is a type of change which falls outside the inhibitory mechanisms postulated by S-R theory. Maps and place hypotheses are not inhibited; they are changed as a function of new learning. Taxon hypotheses were assumed to be more persistent and the possibility was noted that guidance and orientation hypotheses might differ in this regard. The persistence of a taxon hypothesis should be directly related to its strength, which in turn would be a function of the stereotypy with which it is deployed. Orientation hypotheses, by definition, are most likely to be used in a stereotyped fashion and could demonstrate the most marked persistence.

These postulated differences in the persistence associated with the various types of hypotheses suggest that shifts in persistence brought about by certain experimental treatments, including for instance brain lesions, might reflect changes in the use of particular hypotheses. Thus, any treatment inducing a shift from the use of a place hypothesis to a taxon hypothesis would be expected to increase persistence.

10.1.1. PERSISTENCE SHIFTS

A variety of experimental treatments affect the persistence with which behaviours are employed, some producing increases, some decreases. The stereotyped behaviour described by Maier (1949) in response to an insoluble problem is an example of an increase in persistence elicited by particular training conditions. Conversely, increased training can sometimes elicit a paradoxical decrease in persistence either in reversal, the *over-training reversal effect,* or in extinction, the *over-training extinction effect.* The bulk of experimentation aimed at an understanding of extinction and persistence phenomena has been concerned with another treatment which affects persistence–the use of random intermittent reward during training. In most situations such partial reward training serves to increase resistance to extinction, and a discussion of this phenomenon would serve as a means of exploring the possibility that persistence shifts can accompany changes in hypothesis behaviour.

10.1.1(a). The partial reinforcement effect. There are several major theories of the partial reinforcement effect, which is one of the most widely studied phenomena in experimental psychology:

(1) The *discrimination hypothesis* suggests that partial reward training is more similar to extinction than is continuous reward training; thus the onset of extinction cannot be easily discriminated. This notion is derived from the generalization decrement hypothesis of extinction.

(2) The *memory after-effects hypothesis* (Capaldi 1966), another application of the generalization decrement notion, suggests that behaviour during training becomes conditioned to the after effects produced by reward or non-reward. During continuous-reward training animals learn to respond only in the presence of reward after effects; thus extinction is relatively rapid. During partial-reward training animals learn to respond in the presence of both reward and non-reward after effects; thus extinction is prolonged.

(3) The *frustration hypothesis* (e.g. Amsel 1958) suggests that animals are frustrated by their failure to find food (or water) on non-reward trials, and that behaviour gradually becomes conditioned to the cues associated with frustration. Then, the frustration elicited during extinction is compatible with continued responding.

None of these theories, according to Mackintosh (1974), accounts satisfactorily for the partial reinforcement effect, though all the processes

postulated could play some role in its development. In addition to these factors we would suggest another which could contribute to the unusual persistence seen after partial-reward training: a shift in the nature of the hypotheses underlying behaviour.

(1) In response to the frustration associated with unpredictable non-reward, taxon hypotheses could become increasingly stereotyped in the fashion described by Maier (1949). This factor could be partially responsible for the *partial reinforcement acquisition effect* which involves the gradual development of supra-normal performance during partial-reward training (but see (3) below). This is manifested as faster running in an alley or as increased response rates in the operant chamber.

(2) In response to the uncertainty associated with random reward place hypotheses might become inoperable. This follows from a set of connected considerations: *(a)* any variability in the environment elicits, via the locale system, exploratory behaviour; (b) random intermittent reward constitutes a source of continuous variability; (c) the continuous elicitation of exploration interferes with goal-directed behaviour; thus (d) partial-reward training biases strongly against the use of place hypotheses. If animals discard a place hypothesis in the face of variability this should happen within a few trials of non-reward and should lead to the use of a more persistent taxon hypothesis. This factor could account for the partial reinforcement effect observed after only a few training trials, a phenomenon which cannot be explained in terms of the gradual conditioning of behaviour to the cues provided by frustration.

(3) As a corollary to the shift from place to taxon hypotheses, the variability associated with intermittent non-reward could devalue the output from the locale mismatch system, even when taxon hypotheses are being used. When extinction begins the animal would not react immediately with exploration, and this would have the effect of increasing persistence.

The factor leads to the general postulate that any form of environmental variability during training could increase persistence, a postulate similar to one forwarded recently by Amsel (1972).* The importance of this idea is that it is the variability inherent in partial reward, rather than non-

* According to this view
'the maintenance of any ongoing behaviour in the face of a stimulus that evokes an orienting-disruptive response increases resistance to extinction ... The mechanism responsible for this increased persistence is a more general form of the counterconditioning mechanism in the frustration theory of the partial reinforcement effect' (Amsel et al. 1973, p. 176)
This latter explanation becomes unnecessary when one assumes that in order to maintain behaviour in the face of stimuli which would otherwise elicit exploration, the animal must ignore its locale mismatch system.

reward *per se,* which primarily produces increased persistence. It is important to note in this regard that the variability must be truly unpredictable in order to elicit increased persistence; neither regularly patterned partial reward nor a constant delay of reward, both of which could be considered frustrating, produce increased persistence.

Thus, the increased persistence seen during extinction after partial-reward training can be ascribed to several factors: (1) a shift in the hypotheses underlying behaviour occasioned both by frustration and variability, leading towards the use of more persistent hypotheses; (2) the neglect of the exploration system which normally acts to facilitate extinction; (3) the conditioning of behaviour to the context provided by non-reward. Certain of these would appear in hippocampal animals, while others would not, in that they depend upon the presence of the locale system. We shall return to the partial reinforcement effect, and its manifestation in hippocampal animals, later.

10.1.1(b). Extinction variables. Resistance to extinction can also be affected by conditions during extinction itself. Two such factors are of particular interest in the present context; the interval between repetitions of the trained behaviour, and the availability of alternative behaviours. In many situations it appears that the use of shorter inter-trial intervals facilitates extinction. This would seem to be a function of the rapid build-up of reactive inhibition which serves directly to inhibit the previously trained behaviour. Under these conditions a rest interval allows the inhibition to dissipate and the behaviour *spontaneously recovers.* According to our model both the effect of massed extinction trials and spontaneous recovery should be evident in hippocampal animals, in as much as they depend upon factors working within the taxon system. In fact, the normal animal might show less sensitivity to the effects of inter-trial interval during extinction when place hypotheses are being used for extinction would not depend upon the accumulation of reactive inhibition.

The availability of alternative behaviours can also affect extinction rates. Reversal can be viewed as an extinction situation (for one behaviour) which incorporates a highly likely alternative behaviour. The extinction of certain avoidance tasks, superimposed upon approach tasks, offers another example of a situation where suitable alternatives could help the animal extinguish old behaviours. It is obvious that the hippocampal animal, lacking exploratory mechanisms, might be crucially dependent upon the presence of such alternatives.

10.1.1(c). Extinction reconsidered. The foregoing discussion indicates that variations in persistence can result from several factors: (1) shifts between

hypotheses of varying persistence;* (2) the neglect of the misplace system; (3) the action of variables influencing the build-up of inhibition. The latter factor depends upon processes occurring within the taxon systems, while the former two factors depend upon interactions between systems. Place hypotheses are most easily extinguished; merely demonstrating the absence of reward should elicit exploration, the development of up-dated maps, and the abandonment of the now inappropriate place hypothesis. The phenomenon of latent extinction (eg. Seward and Levy 1949) represents this form of extinction in a relatively pure state. Taxon hypotheses are more persistent, though in the intact animal the existence of the locale system provides for exploration, which plays an important role in the extinction of these hypotheses partly by allowing for the development of competing behaviours.

Thus, the animal without a hippocampus suffers from a dual deficit in extinction situations. First, the use of minimally persistent place hypotheses during training is impossible. Second, the facilitating effect of exploration should be absent. This suggests that strong deficits in extinction could occur in lesioned animals. Further, it suggests an unusual dependence upon those factors affecting persistence within the taxon systems. On the other hand, a number of phenomena associated with extinction could be intact in hippocampal animals. That portion of the partial reinforcement effect unrelated to the shift from place to taxon hypotheses, or the blocking of exploration, should appear; spontaneous recovery should be normal; the action of reactive inhibition should be observed.

10.2. Extinction after hippocampal damage

According to our model the effect of hippocampal lesions should be to make the animal dependent upon the taxon systems both for the behaviour learned during training and the means by which this behaviour is changed once extinction has begun or when some other change in the reward contingencies has been instituted. We have already described some evidence for this; changes in the location of reward, or its associations with other objects, can be reacted to in a normal or abnormal fashion, depending upon the extent of locale involvement in the normal animal. The introduction of shock in the alley approach experiment shows both normal and abnormal behaviour developing in the hippocampal animal at the same time. The intact rat, behaving on the basis of a place hypothesis, reacts to the introduction of shock by altering its behaviour

* Aside from the partial reinforcement effect the over-training extinction effect might also be partially explicable in terms of shifts between hypotheses. In this case, however, the shift would be from a more persistent to a less persistent hypothesis. We cannot examine his possibility in detail here, but it is consistent with the fact that the over-training extinction effect is most prevalent in situations (eg. runways) allowing for a shift from taxon to locale hypotheses, presumably when considerable training has generated extensive reactive inhibition (cf. Kendrick 1958).

at a remove in the start area. The lesioned rat, using a taxon hypothesis, cannot inhibit its running behaviour in the start area but can avoid contact with the object associated with shock. Similarly, when reward is moved from one place to another, the intact rat quickly switches its behaviour to accommodate the change; the hippocampal rat cannot manage the same. Such examples predict that deficits will often, but not always, be seen in hippocampal animals during extinction.

The results of studies concerned with extinction in hippocampal animals are presented in Table A26. It is clear from the table that profound deficits are seen in a variety of situations. However, two important exceptions to this pattern emerge. First, in a few studies which have manipulated inter-trial intervals (eg. Jarrard, Isaacson, and Wickelgren 1964; Jarrard and Isaacson 1965) large deficits were observed only when spaced trials were used. Second, nearly normal extinction rates, with certain qualifications, are often seen in avoidance tasks, operant tasks, and classical conditioning situations. These facts indicate that extinction deficits are not an inevitable accompaniment of hippocampal damage. Before turning to a closer analysis of these exceptions we should note a few important observations concerning the nature of the hippocampal animal's behavioural response to the withdrawal of reward.

Cohen (1970) observed that while control rats sniff, scratch, and attempt to get out of the empty goal box, hippocampal rats remain placid, seemingly oblivious to the absence of reward. This failure to react to non-reward could be interpreted in several ways. First, it could imply that these animals generally lack a frustration response to the withdrawal of reward. We think this is unlikely, and shall discuss the question of frustration later. Here, it suffices to say that hippocampal animals display what appear to be either increased or at least normal frustrative reactions in a number of other situations. Second, the absence of a reaction to non-reward in the runway could indicate the absence of the specific expectation of reward in the goal box; that is, no mismatch occurs. This seems quite plausible to us, and it is consistent with the fact that hippocampal animals do not show the characteristic elevation of plasma corticosterone upon the withdrawal of reward (Coover, Goldman, and Levine 1971b), a function apparently associated with the occurrence of environmental mismatch (see pp. 360-1). The absence of such a mismatch, however, does not mean that extinction becomes impossible. In fact, in both the Cohen (1970) and the Coover *et al.* (1971b) studies the hippocampal animals showed extinction, albeit somewhat more slowly than did controls. The absence of the mismatch function does imply, however, that extinction might proceed in different ways in intact and hippocampal animals.

Kimble (1969) trained rats on a brightness discrimination task in a Y-maze for water reward. He ran the animals to a satiation criterion;

that is, he continued to train the rats until responding for water stopped spontaneously. Both control and hippocampal rats stopped drinking at about the same time. A few trials later the control rats stopped running; the hippocampal rats ran throughout 100 trials, though they did not drink. The same pattern of persistent responding was seen in the hippocampal rats when a standard extinction procedure was used. These data suggest that the hippocampal rats were not running for water *per se.* Their behaviour bore all the marks of a fixation which, initially dependent upon reward, had become autonomous. In this situation, as Kimble has shown, running could persist almost indefinitely. Another aspect of the data from the Kimble study confirms the suggestion that the hippocampal animals were responding on the basis of a persistent taxon hypothesis; they responded almost entirely to the cue that had been associated with reward. Similar fixation of choice during extinction in hippocampal animals has been reported by Niki (1965) and by Douglas and Pribram (1966). Thus, in situations where profound extinction deficits are reported, there are clear indications that hippocampal animals perform on the basis of persistent taxon hypotheses. This provides a basis for understanding the important effect upon extinction of those factors to which the taxon system is most sensitive: the rate of repetition of non-rewarded responses, and the availability of alternative behaviours.

10.2.1. VARYING INTER-TRIAL INTERVAL

For the normal animal in the runway inter-trial interval seems relatively unimportant during extinction; this is reasonable if we assume that extinction in these situations is based predominantly on changes in the locale system. For the hippocampal animal, however, inter-trial interval appears critical; massed extinction trials can lead to relatively normal extinction rates in hippocampals. We assume that these data reflect the way in which persistent taxon hypotheses can be altered; massed repetitions without reward lead to a rapid build-up of inhibition. Thus, these results parallel those presented for habituation (p. 249) and reversal (p. 382) in suggesting that when the same behaviour is repeated (and unrewarded) at short intervals it will drop out in hippocampal animals. It is important to note that the behaviour in question must not constitute a long chain of sequential acts. We noted in our discussion of habituation and hyperactivity that hippocampal animals decrease their activity at a seemingly normal rate only in small boxes where the same responses are repeated at a rapid rate, and often fail to decrease their activity in larger boxes where repetition rates for the same response are slow. A parallel to this effect in extinction is provided by the study of Jarrard and Lewis (1967); they showed that massed extinction trials in a complex maze failed to lead to normal extinction rates in hippocampal rats, presumably because the rate at which particular responses in the maze were repeated was too slow.*

Much the same analysis can be applied to extinction in operant situations. These can be viewed, as we have pointed out, as fairly pure examples of taxon function. Thus, we would expect that extinction in both normal and hippocampal animals might be relatively lengthy, but that no major differences should appear between the groups. Most operant situations leave response rate at the discretion of the animal; thus, inter-response times should be quite short during the early stages of extinction. Extinction in operant tasks typically proceeds at about the same rate in normal and hippocampal animals, with small deficits sometimes appearing in the lesioned subjects. The large deficits reported in the runway, where the normal animal can avail itself of place learning and its associated rapid extinction, find no parallel in the operant situation.

10.2.2. THE AVAILABILITY OF ALTERNATIVE BEHAVIOURS

As we have seen, the 'extinction' of responding to the previously rewarded stimulus in a non-spatial discrimination reversal task is relatively rapid in hippocampal animals (p. 283); they quickly adopt an alternative, though usually maladaptive, behaviour which is attractive in that it is at least sometimes rewarded (Silveira and Kimble 1968). This effect holds even when trials are spaced. In contrast, the total removal of reward in the same situation (Kimble and Kimble 1970) results in a profound extinction deficit which is characterized by abnormally stereotyped responses to the previously rewarded stimulus.

This comparison brings out the crucially important role that an alternative behaviour can play in extinction for hippocampal animals. The distinguishing feature of those aversively motivated tasks which have produced normal extinction rates in hippocampal animals is that they provide highly likely alternative behaviours for the animals. Thus, in the case of the conditioned emotional response, the aversively controlled behaviour (suppression) is superimposed upon an appetitive habit. When the aversive contingencies are removed the animal can re-engage in this appetitive behaviour; hippocampals do so at least as readily as do normals. In this situation the alteration in aversive contingencies is clearly signalled to the animal; shock is no longer presented after the conditioned stimulus. In most avoidance situations this is not the case; as long as the animal

* A curious aspect of the data in the Jarrard and Isaacson (1965) study concerned the effect of switching the inter-trial interval in the course of extinction. Hippocampal rats extinguished initially with massed trials showed normal extinction rates; when they were switched to spaced trials the running behaviour recovered. Normal rats did not show this recovery effect. These data strongly suggests that the two groups had extinguished in different ways. We suspect that the recovery of responding in the hippocampal rats reflected the dissipation with time of some of the short-term inhibition built up during the massed extinction trials. The absence of this effect in normal rats indicates that extinction in these animals was based on something other than, or in addition to, this type of inhibitory effect.

continues to avoid it fails to detect the removal of shock at the onset of extinction. This, of course, partly accounts for the well-known difficulty associated with extinction of avoidance behaviours. Lovely *et al.* (1971) have shown how this can affect extinction rates in hippocampal rats trained on a two-way avoidance task. They found an extinction deficit in their lesioned rats when one considered the number of trials on which a response was made after shock had been made contingent upon response; this is consistent with what we have already seen concerning the learning of this task. However, if extinction was measured in terms of the number of responses made *after* an animal had once failed to run, there was no difference between the groups. Most of the hippocampal rats never responded again, after once performing this alternative behaviour.

10.2.3. EXTINCTION IN NORMALS AND HIPPOCAMPALS COMPARED

The foregoing has documented the case that hippocampal animals are often deficient at extinction, but that they can show `normal' extinction rates when given either massed unrewarded trials or a strong alternative response. This sensitivity to factors which often fail to influence the rate of extinction in normal animals indicates that extinction proceeds in different ways in the intact and hippocampal animal; this implication is strongly supported by a study reported by Warburton (1972). Rats were given repeated acquisition-extinction cycles on an operant task. The hippocampal rats had a small, but persistent, deficit in extinguishing this behaviour. Warburton then tested the effects of the intraperitoneal injection of an anti-cholinergic drug upon the extinguished behaviour. This drug elicited a recovery in the normal animals but had no effect upon the hippocampal animals. These data suggest that normals and hippocampals had extinguished in quite different ways, using brain systems dependent upon different neurotransmitters. They also support the suggestion, made earlier, that normal animals can extinguish through the action of the locale system when taxon mechanisms alone would not have supported extinction. Removal of the locale system releases the taxon systems and the behaviour 'recovers'. This analysis can be applied to an experiment reported by Niki (1962). He showed that hippocampal lesions produced after extinction can lead to the partial recovery of the previously extinguished behaviour, in the absence of any post-operative training. Here again, extinction within the locale system seems to have occurred before sufficient inhibition had been generated in the taxon system, and this disparity was brought out by the subsequent removal of the locale system.

The deficits reported in extinction in hippocampal animals, then appear to be a function of the different ways in which the locale and taxon systems respond to the withdrawal of reward. At the same time a few

aspects of reasonably normal extinction performance can be observed in hippocampal animals, presumably reflecting functions mediated solely by the taxon systems. Thus, spontaneous recovery, assumed to reflect the dissipation of short-term inhibitory effects, remains intact in hippocampal animals (Schmaltz and Theios 1972). Similarly, the effect of work decrement on extinction, a factor directly related to the inhibition generated by the performance of a response, remains the same in normal and hippocampal animals (Van Hartesveldt 1973).

10.2.3(a). The partial reinforcement effect after hippocampal damage.

Several studies have been reported which investigated the role of the hippocampus in the partial reinforcement effect. We have suggested that this effect rests partly on shifts from less persistent place hypotheses to more persistent taxon hypotheses, partly on the shutting off of the exploratory system, and partly on the effects of frustration. Animals with hippocampal lesions should show a partial reinforcement effect, but should not show that part of it related to the persistence induced by hypothesis shifts involving the locale system, or that part due to the shutting off of the mismatch system. The absence of the locale system makes original learning more prone to persistence, even without partial reward. Thus, the partial reinforcement effect might, in hippocampals, be superimposed upon the deficit in extinction shown by continuously rewarded hippocampal animals.

The available data generally support this position. Bloom and McFarlain (1971) tested rats in a runway under continuous or partial (50 per cent) reward conditions. Although the hippocampal rats had a deficit in extinction they displayed a partial reinforcement effect. That is, continuously rewarded hippocampals extinguished more rapidly than did partially rewarded hippocampals. On the other hand, Franchina and Brown (1970) showed that hippocampal rats trained with either patterned or random partial reward in a runway extinguished at the same rate, which was roughly equivalent to that seen in intact rats given random partial-reward training. The authors concluded that the partial reinforcement effect was absent in the hippocampal rats; this conclusion would seem to us to be unwarranted. The control groups in this study learned to pattern their running speeds in the patterned partial-reward condition; they ran fast after non-reward trials and slow after reward trials. They did not show the partial reinforcement effect in this situation, as we have already noted. However, the hippocampal rats tested under this patterned partial-reward condition failed to pattern their responses in accordance with the alternating schedule. They ran rapidly on all trials, as did the hippocampal rats given random partial rewards. The fact that both these hippocampal groups showed extinction rates similar to the random partial-reward control animals raises a sticky interpretive problem. Did the

hippocampals show a partial reinforcement effect in both cases, or did they simply have an extinction deficit and no partial reinforcement effect at all? Two points suggest that the former is the more likely explanation. First, the Bloom and MacFarlain study, and others to be noted, indicate that the partial reinforcement effect can be demonstrated in hippocampal rats. Second, Franchina and Brown used massed extinction trials (20 s inter-trial interval), a procedure which should have produced normal extinction rates in their hippocampal animals (see above). Thus, the prolonged extinction in both hippocampal groups would seem to reflect the presence of a partial reinforcement effect, rather, than its absence.

Williams (1971) trained rats in a runway under continuous or partial reward using spaced trials. Hippocampal rats showed the expected extinction deficit in this situation, whether they had been continuously or partially rewarded. However, in terms of the number of trials required to reach two different criteria of extinction (the first or the third extinction trials with a latency of 60 s), the partially rewarded hippocampals took reliably longer than did the continuously rewarded hippocampals. Thus, in the presence of an extinction deficit it is still possible to detect a partial reinforcement effect in hippocampal animals. More recently, Brunner *et al.* (1974) have reported a partial reinforcement effect in rats with X-ray-induced damage in the hippocampus.

The most convincing data have been provided by Amsel *et al.* (1973). These investigators showed that the increased persistence selectively elicited by stimulus variability during training (a tone presented at the end of an FR-21 chain) was absent in hippocampal rats. This is consistent with the notion that this increase in persistence is related to a shift in hypotheses, or to a shutting off of the exploratory system, and differs somewhat from the persistence elicited by the frustrative aspects of random non-reward. The latter, in view of the studies discussed above, appears intact in hippocampal animals.*

10.3. Reaction to reward change

Extinction, as we noted at the beginning of this chapter, represents a drastic case of a class of situations involving a change in reward contingencies. The abnormal persistence seen in hippocampal animals in extinction is often seen in other situations involving such changes, as we have noted at several points in this section of the book. In the remainder of this chapter we shall attempt to account for the general character of this

* A study by Rabe (1968) showed that abnormal persistence in hippocampal rats *depended* upon frustrative non-reward in an operant situation. Rats were trained an FR-20 schedule. Following this, a second lever was introduced, responses to which were rewarded either on an FR-5, FR-10, or CRF schedule. Both normal and hippocampal rats showed considerable, though not different, persistence to the less rewarding FR-20 lever. This is markedly different from what is seen in an analogous runway task (Gaffan 1972) where normals switch but hippocampals do not.

persistent behaviour as it applies to any situation involving shifts in reward contingencies.

Less drastic situations involving reward shifts include the shifts in operant schedules discussed earlier (p. 322) and changes in the quantity or quality of reward given for a particular response, amongst others. Through all these situations the common thread of frustration appears. This concept, discussed briefly above, is worth elaborating upon because it has been suggested that the hippocampus is centrally involved in frustrative reactions (eg. Gray 1970, Glazer 1974a,b).

Frustration is assumed to be mildly aversive, to energize behaviour, and to provide cues to which behaviours can be conditioned. The classical paradigm for demonstrating frustration makes use of its energizing property. Rats are trained to run in an alley with a goal box at the end and another goal box midway down the alley. Removal of reward from the initial goal box triggers frustration, which speeds up running to the final goal box (Amsel and Roussel 1952). Other tests for frustration include shifts in operant schedules and the partial reinforcement effect. The general problem of reactions to reward shifts, then, is often intertwined with the notion of frustration, except in those cases where a shift involves an increase in reward.

10.3.1. THE EFFECTS OF HIPPOCAMPAL LESIONS UPON FRUSTRATION

Notwithstanding the claim that the hippocampus is important to frustration mechanisms, the available evidence indicates that frustrative effects are largely intact in hippocampal animals. In the most direct test of this notion, Swanson and Isaacson (1969) showed that hippocampal rats were indistinguishable from control rats in the double-alley described above. In operant situations, as we have seen, frustrative shifts often elicit overresponding in hippocampal animals (e.g. Swanson and Isaacson 1967, Rabe and Haddad 1968). In fact, Swanson and Isaacson suggested that the

'hippocampal rats may have been uniquely sensitive to the frustrative conditions of reinforcement withdrawal in the sense that their frustration response was more intense' (p. 34).

In contrast to this demonstration of 'increased' frustration, a study of reward shift in a runway failed to find significant performance changes in hippocampal rats following either an increase or a decrease in reward quantity (Franchina and Brown 1971). Further, Niki (1962) has shown that hippocampal rats fail to decrease running speeds in an alley when a delay is interposed between response and reward,* while two studies

* Mikulka and Freeman (1975) have shown that this condition retards the learning of a Y-maze position habit in hippocampal rats.

have demonstrated deficits in hippocampal rats in a situation requiring slow running (Posey 1972, Rickert and Bennett 1972). Similarly, Murphy and Brown (1970) failed to find a reliable contrast effect* in hippocampal rats switched from high to low-concentration sucrose solutions (however, see Gaffan (1973) for a study reporting normal contrast in fornical rats). Most of these studies demonstrate decreased frustration reactions in hippocampal animals. Thus, one can observe either increased or decreased frustrative reactions in response to various treatments. It seems unlikely that a direct effect upon frustration is one of the effects of hippocampal lesions.

10.3.2. THE BASIS OF ABNORMAL SHIFT BEHAVIOUR

The thread tying the various shift studies together is neither abnormal frustrative reactions, though these may sometimes occur, nor an inability to respond to altered reward contingencies, though this too is sometimes seen. Rather, it appears to be an inability to change certain forms of behaviour. As a number of studies have made clear, hippocampal animals are sensitive to different levels of reward when such differences are present from the onset of training (e.g. Murphy and Brown 1970, Franchina and Brown 1971, Kramarcy, Mikulka, and Freeman 1973) in that their asymptotic performance rates reflect these different reward levels. Thus, it cannot be the case that alterations in the quantity or quality of reward go unnoticed. Whether or not the hippocampal animal can alter its behaviour in the face of these changes, however, seems to depend upon the nature of the behaviour itself. Thus, appropriate shifts in behaviour have been reported in some operant situations, or in situations where consummatory responses alone are measured, or in some extinction tests. On the other hand, decreasing reward, or delaying it, in a runway often fails to elicit a marked decrease in running speed. This difference between spatially extended runways and operant situations has been noted before (see p. 342). Intact rats can alter their behaviour in the start box in such situations, while lesioned animals seem to have to wait until they catch sight of the objects of reward themselves, though this does not always lead to normal performance (particularly in extinction).

This difference between normal and hippocampal reactions to reward shifts shows up clearly in Posey's (1972) study of runway behaviour. His rats were first trained to receive reward regardless of runway speed. Following this, reward was only given on those trials on which a run took at least 5 s; this required slowing down. The hippocampal animals were poor at this task, but what is particularly interesting is the distribution of changes in running speeds within the runway. Lesioned rats were

* This refers to an unusually large downward shift in performance after a downward shift in reward such that shifted animals perform worse than controls always tested at the lower reward level (cf. Crespi 1942).

faster than controls in the start area, but slower in the goal area. This suggests that the control rats master the task primarily by hesitating in the start box, a feat beyond the powers of the lesioned rats.*

It seems clear that changes in reward are often noticed, if responded to inappropriately, in hippocampal rats. A study reported by Gaffan (1972) provides evidence that hippocampal disruption does not eliminate sensitivity to reward shifts, even in runways, though it might leave animals unable to alter their behaviour rapidly. Rats were trained to turn left on a T-maze to obtain four small food pellets. Following the achievement of stable performance on this task, eight pellets were made available in the right arm. Normal rats rapidly learned to turn right in order to get the eight pellets first; rats with fornix lesions never did this. However, they demonstrated clearly that they had learned about the eight pellets by decreasing their latency to the right arm (via the left arm and its four pellets) dramatically.**

In a recent study (Van Hartesveldt 1973) the effects of variations in reward size upon the acquisition, performance, and extinction of both continuous and partial reinforcement operant schedules were explored. When reward size was decreased on the CRF schedule both normal and hippocampal rats increased their response rates. Similarly, both groups responded similarly to the partial reward schedule (FR), including the number of sessions it took the rats to extinguish when the ratios was set at 105 responses per reward.*** Van Hartesveldt suggested that abnormal increases in response rate were not seen in her hippocampal rats upon reward level decreases because the decreases occurred gradually. Increases were seen, but they did not differ from the increases observed in normal animals. As we noted earlier (p. 322) hippocampal rats adjust normally to shifts in operant schedules only when the shifts are not abrupt. Abrupt shifts are likely to elicit exploration in intact animals, which would militate against the marked response rate increases observed in hippocampal animals.

The data from all these studies suggest that the basis for persistence in hippocampal animals resides in their use of stereotyped taxon hypotheses, without the possibility of disrupting these through exploration. Normal animals either use place hypotheses, which are minimally

* The Hullian concepts of backward chaining and the backwards breaking down of an alley response during extinction might apply quite nicely to hippocampal rats, if not to normal ones.

** This study shows that the effects of hippocampal damage can be equally marked in reactions to increased, as well as decreased, reward. Thus, it constitutes *prima facie* evidence against a faulty frustration reaction interpretation.

*** Perhaps the most important finding in this study was the fact that extinction rates were similar for the two groups. Extinction in this situation, as the author pointed out, was a function of the inhibition

'generated by increasing the amount of work required to obtain a reward' (p. 356).

This type of inhibition corresponds to that postulated for the taxon systems and it is clear from these data that this form of inhibition remains intact after hippocampal lesions.

persistent and easily extinguished, or they use taxon hypotheses which can be disrupted by exploration. Only when the experimental situation is arranged such that taxon hypotheses must be used and exploration has been eliminated, as in partial reinforcement training for instance, will the normal animal be reduced to the persistent behaviour seen in hippocampal animals.

10.3.3. CONCLUSIONS

The above discussion indicates that reactions to reward changes are often abnormal in hippocampal animals because they lack the locale mechanisms available to intact animals; their performance reflects solely the properties of the taxon systems. The obvious corollary to this position, as we just noted, is that intact animals should perform like lesioned animals when they use taxon hypotheses and when exploration has not been allowed.

The central point to be made concerning an understanding of extinction, both in normal and lesioned animals, involves what takes place prior to the removal of reward: how has the animal learned the task? In some situations a variety of hypotheses can be used, while in others performance requires the use of a specific hypothesis. A proper evaluation of extinction effects requires the prior analysis of learning. The means by which the former occurs will depend upon the hypothesis used during learning. The withdrawal of reward will be rapidly registered in the locale system, and this will lead to the rapid extinction of inappropriate place hypotheses. The cessation of behaviour, however, will still depend to some extent upon the strength of the taxon hypothesis mediating it, for most tasks incorporate habits requiring response hypotheses. One-trial extinction is relatively rare. The locale system contributes to the extinction of these taxon habits by triggering exploration, such that the full persistence associated with taxon hypotheses is rarely seen. This delicate interplay between locale and taxon systems accounts for the difficulty in providing a neat analysis of extinction, in normal as well as in hippocampal animals. Yet it offers the possibility of understanding the persistence often seen in hippocampals, not as a function of some inhibitory capacity of the hippocampus which has been disrupted but rather as a function of the properties of what remains after this disruption. As such, it might point the way to a more comprehensive theory of extinction, one which sees this phenomenon as occurring in any of several ways, involving both new learning and the inhibition of old learning.

© John O'Keefe & Lynn Nadel

You may copy or reproduce any part of this file for teaching purposes or personal use. The original text and figures should not be altered in any way. Permission should be obtained in writing from one of the authors if all or part of any of the figures or text is to be used in a publication and the source should be acknowledged:

John O'Keefe & Lynn Nadel (1978) The Hippocampus as a Cognitive Map , Oxford University Press.

You may redistribute the file electronically providing you do not modify it in any way.