Foraging in the Barbary Dove:
Evolution, optimisation, and rules of thumb

by

Clive D. L. Wynne

Doctor of Philosophy

University of Edinburgh

1986

DECLARATION

This thesis is my own composition, and the work presented in it is my own.

ABSTRACT

Barbary Doves *(Streptopelia risoria)* were tested in two simple foraging experiments based on concurrent interval schedules of reinforcement. In both experiments food reinforcement was made available, independently on each option, for the first response to a schedule after some time had elapsed since reinforcement had last been delivered by that option. The two schedules were independent, and the mean interval between reinforcements on each was not always the same. For one experiment the response required was depressing one or other side of a floor panel, and the intervals between reinforcements were drawn from an exponential distribution, so that they were as variable as possible. In the other experiment the response requirement was a key peck; and intervals between reinforcements were either drawn from an exponential distribution; from a semi- variable arithmetic distribution; or from a completely regular set of fixed intervals.

Results from these experiments are compared to the predictions of optimal foraging theory, assuming that the subjects are trying to maximize net energy intake. No agreement with prediction was here found for the floor panel experiment. For the key peck experiment it was found that visits to the better key were longer than to the key with the lower rate of payout, but that visits to the latter key were not instantaneous as had been predicted; nor was the interval between pecks uniform.

Results from both experiments were also analysed in terms of the matching relation, though the utility of this procedure is here disputed.

Data from the key peck experiment was tested to see if the Doves were using either of two 'rules of thumb' (behaviour generating rules) that have been proposed to account for behaviour on these schedules. No support at all was found for the notion that the proportion of reinforcements from the two options in the past was controlling the proportion of responses given to each option in the future (known as 'melioration'). Evidence for the theory of 'momentary maximising', predicting that each response would be placed on the option offering the highest momentary probability of reward, was hard to assess, because of the presence of a changeover delay in the key peck experiment. It seemed that, in general, responses were not placed so as to maximize momentary probability of reinforcement, but that a certain

subset of responses, the first response to a new schedule, was well placed in this respect.

These results are discussed in the context of a theory of the role of learning and behaviour in evolution known as 'Evolutionary Epistemology' which was not, however, directly tested.

ACKNOWLEDGEMENTS

TABLE OF CONTENTS

CHAPTER 1
INTRODUCTION

> Once upon a time, there were two monsters. Instead of devouring young maidens or destroying New York, they spent their time worrying about how animals behave. Godzilla put animals in boxes and gave them electric shocks. His aim was to establish general laws, and to this end the natural history of the species in the box was irrelevant. In marked contrast to Godzilla, the Creature from the Black Lagoon used to dress up in gum boots and an anorak to crawl around in the mud, watching the sexual habits of animals with the aid of his binoculars. Godzilla would attack the Creature's ignorance of statistics and inability to control any of the important variables; the Creature would retort with charges that Godzilla's experiments were unnatural.

Houston 1980:297

The aim of this thesis is easier summarized than explained. It is to explore the implications of evolution for individual behaviour, especially learnt behaviour. In chapter 2 I criticize some of the approaches that exist in the literature to the problem of learning. Optimal foraging theory, whilst a useful component of a study of behaviour, I find inadequate as a theory of what to expect when looking at behaviour, especially flexible behaviour. To hold that one should expect to observe optimal behaviour ('literal optimising' as Herrnstein & Vaughan [1980] call it), is to confuse the aim of a process with the state actually observed. In human affairs we do not confuse someone's aims with their actual behaviour, and it is strange that in animal research such a view has had any support. Of course every extant species must be keeping up a certain standard of performance in an activity critical to survival, like foraging, but there is some distance between saying that, and demanding *optimal* behaviour. When we consider learnt behaviour within this framework the situation becomes even more problematic. In principle it should be possible to identify what appear to be the constraints on behaviour for the organism under study, and formulate the optimal strategy in a similar way as for inflexible behaviour. In practice the need to account for the constraint of gaining knowledge of the environment, makes it too easy to take the observed pattern of behaviour as a constraint, and then whatever the behaviour maximizes is what the subject is doing optimally. I admit that this is a caricature of the way that most researchers proceed, and certainly optimality analysis has a role to play in helping us to find the constraints under which an organism is working, and the functions its behaviour serves. None the less its

utility is limited, and cases where observers record that an organism is foraging 'optimally' tell us very little, and raise the suspicion that the investigator has implicitly adopted an 'inverse optimality' approach, and worked backwards from the observed behaviour to the function it must be optimising. (Explicit use of an inverse optimality approach is different- to demonstrate the function that observed behaviour maximizes [or minimizes] can have its uses, but it would never lead a researcher to conclude that behaviour is optimal, since optimality of behaviour is one of the premises of such an approach).

A critical attitude to optimal foraging theory prompted the development of the floor panel experiment described in chapters 3 and 4. In 1981 Houston and McNamara developed a model specifying optimal behaviour on the popular psychological paradigm known as a concurrent variable interval schedule. In the form used in the floor panel experiment this meant that so long as one or other side of the floor panel was depressed, reward would appear at a certain average rate; while when the other side of the floor panel was depressed reward would occur at an average rate associated with that side. Whilst a side of the panel was not depressed a reward could still come to be set up, ready to be delivered, but could not be obtained until the relevant side was depressed. The optimal policy in such a situation can be derived, and predictions more detailed than had previously been subjected to test from an optimality account, are explored in chapter 4.

Although it is certainly a fair criticism of this experiment that it is not obviously like any natural foraging situation, as I argue in chapter 4, it is not clear that it is so much more unnatural than many of the experimental set-ups that have produced evidence in favour of optimal foraging theory. None the less the results of this experiment were so far from those predicted from optimality considerations that it suggested the Doves were entirely failing to comprehend what the experiment required of them.

In the latter half of chapter 2 I consider psychological theories of flexible behaviour, and conclude that they have still to come to terms with the role of behaviour in evolution, and the so-called 'biological boundaries' to learning that became apparent twenty years ago. The suggestion that comparative psychology and ethology could be married by ethology providing the functional (optimality) theories, whilst psychology provides theories of the mechanisms by which behaviour is controlled, I do not consider likely to be a very fruitful

union. Foremost amongst my misgivings is the simple fact that few psychological theories are, indeed, theories of mechanism. On balance I would estimate that ethology actually has more theories of behavioural mechanism than animal psychology.

This dissatisfaction with extant attempts to cope with learnt behaviour within an evolutionary framework leads me, in chapter 3, to outline a theory of animal behaviour based on the ideas known as 'Evolutionary Epistemology'. This account requires a radical reassessment of what we mean by 'evolution' so that we can see evolution, as it proceeds through 'blind' variation and natural selection, as itself a learning process. When we see evolution in this way the evolutionary function of individual learning becomes clearer, and a new approach to animal behaviour, which I just sketch the perimeter of, can be developed.

In chapter 3 I claim that this new approach applies to behaviour of all kinds in all situations; however, this does not mean to say that it is equally easy to study under all conditions. In order to make as fair a test as possible of optimal foraging theory it seemed important to test the subjects until I could be sure that they were fully settled into the experimental conditions. This required the development and application of a criterion of stability that the subjects' behaviour had to fulfil. The net effect of this constraint was that the floor panel experiment took around five months to run. In principle one could have analysed the variation and selection of behaviour (as the theory developed in chapter 3 would advocate) over this length of time, in practice the task would be impossible.

I resolved the dilemma— whether to continue to investigate optimal foraging, or to explore evolutionary epistemology in more detail, in favour of optimal foraging. I reasoned that it would be interesting to compare the predictions of Houston & McNamara's optimality model with behaviour on an experiment in which pigeons (at least) are generally considered efficient foragers, although the predictions of this optimality model had never been tested. I refer to this as the 'key peck experiment'(set out in chapters 6 and 7). Its structure is the same as that of the floor panel experiment except that instead of standing on one side or other of a floor panel, the birds pecked at two keys. Pecking the keys collected rewards that had set up in the same way that moving onto one side of the floor panel had done. Obviously there was

no analogue to standing on the floor panel, so the birds had to keep pecking at the keys in order to obtain all the available rewards. As I suspected, under these conditions performance showed some sensitivity to the dimensions important for efficient responding. At the same time the behaviour was in no way perfectly optimal.

In chapter 2 I complain that 'rules of thumb' (rules that animals could actually be using to control their behaviour), in ethology are just "pulled out of a hat". In chapter 8 I pull two rules of thumb out of my hat. My justification for doing so was pragmatic. These rules have generated a lot of interest– I wanted to see if they fitted my data. The results were not uninteresting, I feel, but they were not clear cut. If it had been possible to say conclusively what rule of thumb the birds had been using to control their behaviour, it would then have been possible to try and discover how this rule developed, and this might have tied in with the ideas of 'variation and selection' in behaviour, developed in chapter 3. The last point where I thought it might have proved possible to follow through the ideas of variation and selection was in the consideration of the fine temporal patterning of pecking in chapter 9. This was necessary in any case as part of the optimality analysis. Again, however, apart from being sub–optimal, the pattern of behaviour was not clear, probably because of limitations in the recording equipment.

Although the theorising in chapter 3 is blatantly incomplete, and, in the end, was not testable in the experiments described in chapters 4 through 9, I do not feel the need to apologize for it. Criticism of animal experiments from the public at large at the present time is very vocal, and not entirely incoherent. If the study of animal minds cannot be presented as having some semblance of applicability to our own human concerns, then I am not surprised that people get impatient with animal researchers. Evolutionary epistemology, though it may be (in my hands at least) a faltering and incomplete theory, addresses an important question, How can we understand our own cognitive processes as a product of evolution by natural selection? For this reason at least it should be tolerated.

CHAPTER 2
APPROACHES TO ANIMAL LEARNING AND BEHAVIOUR

It is not always easy to explain to people who are not academics why the field of animal behaviour is made up of two distinct disciplines. In the 'good old days' when all learnt behaviour progressed through the entirely even-handed formation of associations, and inflexible 'Fixed Action Patterns' were triggered by 'Innate Releasing Mechanisms', perhaps the partitioning of animal behaviour into two disciplines was understandable. But if this description of the two fields of animal behaviour was ever more than a caricature, it is most certainly inadequate now. Two recent collections of papers testify to the enthusiasm of research workers on both sides of the psychology/ ethology divide to unite in the search for a wholistic science of behaviour. One of these collections is the set of responses in 'The Behavioural and Brain Sciences' to Fantino and Abarca's (1985) paper suggesting how their own (psychological) 'delay reduction hypothesis' could integrate with (ethological) 'optimal foraging theory'; the other example is the proceedings of the Dahlem Conference on the 'Biology of Learning' (Marler & Terrace 1984), at which psychologists and ethologists got together to try and find a common plan for the study of animal learning and behaviour. I find, however, that I must agree with Hailman (1985), that enthusiasm is not itself sufficient, and that the new synthesis 'still remains elusive'.

It is my belief that an integration of all forms of animal behaviour study into a single conceptual schema can only be achieved by a radical reassessment of what we mean by learning, and of the role of the acquisition of knowledge in evolution. The position I advocate here is that known as 'Evolutionary Epistemology', associated in particular with the social psychologist, Donald Campbell. The purpose of this chapter is to argue the need for such a reassessment, and in chapter 3 I outline the theory and its implications.

Firstly let us consider the positions of the psychological and ethological camps, and the extant attempts at reconcilation.

2.1 Psychological Approaches

2.1.1 General Process Theory

'Tis sufficient to observe, that there is no relation which produces a stronger connexion in the fancy, and makes one idea more readily recall another, than the relation of cause and effect betwixt their objects

David Hume

The historical development of comparative psychology out of philosophy will be sketched in chapter 3, section 3.1.3, for now let us simply summarize the conceptual framework of the traditional behaviourist position, much as one might find it in any introductory textbook. This is the position known as 'General Process Theory' (Seligman 1970).

Behaviourism grew out of the British empiricist philosophy of the 18th and 19th centuries. Learning, according to these philosophers, consists of the formation of associations between causes and effects. The quote from Hume at the head of this section is actually the frontispiece to a recent text on animal learning theory (Dickinson 1980). Classical, Pavlovian conditioning showed that animals could associate stimuli with their effects; operant conditioning showed that they were capable of associating actions with their consequences. Classical conditioning was the preferred mode of adjustment of the smooth muscles and glandular systems under the control of the autonomic nervous system; while skeletal muscle control was learnt by operant conditioning. In its strong form behaviourism was a totally environmental theory − it followed Locke's doctrine that the mind at birth was a blank slate *tabula rasa*, on which the chalk of experience then wrote. This view led the behaviourists to uphold that the environment, through differential reinforcement of different actions, had complete control of the organism's behaviour. Of course there were behaviours to which this description did not apply − these were the instincts, and they were not the domain of psychology because they were stereotyped and uninfluenced by the environment. If a behaviour was at all influenced by environmental feedback its conditioning would follow the same associative rules as the conditioning of any other behaviour the organism (or any other organism) was capable of; and it could be associated with any reinforcer.

As for the fitness value of behaviour, comparative psychologists of the old

school adopted a progressive, anagenetic view of evolution (Gottlieb 1985). Once the ability to learn (by association) had appeared in the evolutionary tree, all animals above that point had got it. Associative learning (or at least classical and operant conditioning considered separately) was inherited as a package – if you had it, you had all of it, otherwise you had to make do with instincts, which, being unresponsive to environmental feedback, were clearly second best.

2.1.2 Limitations of General Process Theory

2.1.2.1 Biological Constraints

In the 1960s a number of chinks appeared in the all encompassing armour of general process theory. These 'biological constraints' (Shettleworth 1972) have been extensively reviewed many times before (Staddon 1983; 1986; Bolles 1985; Garcia, Clarke & Hankins 1973; Shettleworth 1972; Seligman 1970; Terrace 1984; McFarland 1985) and so they will only be discussed here very briefly.

The Brelands (Breland & Breland 1961; 1966) set up in business to train animals by operant methods to perform tricks for commercial purposes – typically advertisements. They found that although their subjects would initially perform as expected, after a period of training they might begin to engage persistently in behaviours that had never been reinforced, and that destroyed the performance they were being trained towards. An oft repeated example is the story of the pig that was being trained to pick up tokens and deposit them in a piggy bank. At first all went well but over a period of weeks the pig developed the tendency to drop and root the tokens repeatedly instead of simply putting them in the bank. The Brelands termed this phenomenon 'Instinctive drift'.

The notion that all stimuli were equally associable with all responses was refuted by the work of Garcia & Koelling (1966; see Garcia et al 1973 for review). They demonstrated that rats would learn very much more readily, and over longer delays, that sickness was associated with flavoured water, than that it was associated with water accompanied by flashing lights or sounds. However when the consequence of drinking the water was electric shock, the relationship was reversed, and the flavouring was a poor cue compared to the presence of lights or noise with the water.

Similar conclusions as to the selective importance of certain cues for particular consequences can be drawn from the work of Bolles (1970;1971) in aversive conditioning. Bolles concluded that rats can only learn to terminate an aversive electric shock if the response they are required to make contains elements of flight (moving to another chamber from where shocked), simulated flight (running in wheel), or a partial sort of flight (moving across to the other side of a shuttle box). To train a rat to avoid shock by pressing a bar could only be achieved by three out of twenty four subjects, at even a moderate level after 1000 trials (D'Amato & Schiff 1964 in Bolles 1971). Bolles coined the term 'Species Specific Defence Reactions' to describe the set of responses which could be conditioned to electric shock.

The third problem to arise for general process theory was the procedure known as 'auto-shaping'. Traditionally pigeons had been trained to peck keys in Skinner boxes by a process of reinforcing successive approximations to the required key peck response – this was known as 'hand shaping'. The experimenter would initially operate the food hopper when the bird was anywhere near the key that it was to peck. For the next reinforcement the experimenter would wait until the bird was slightly closer to the key, and so on; reinforcement each time contingent on a slightly closer approximation to the required key peck response. Brown & Jenkins (1968) performed an experiment presenting pigeons with regular food reinforcement in a Skinner box, quite independent of their behaviour. This Skinner box also contained a response key that was illuminated just before food was delivered. It was found that the pigeons quickly developed the habit of pecking the illuminated key, despite the irrelevance (from an external viewpoint) of this act. The situation, clearly, had been changed from an operant conditioning one, to a classical conditioning experiment. The lit key was acting as a conditioned stimulus, producing key pecking as a conditioned response. Although this went against the prevailing beliefs as to the type of responses that could be classically conditioned, at least it could still be accommodated with the general process account. The work of Williams & Williams (1969) however, put an end to that. They adjusted the auto-shaping procedure so that now, instead of being just irrelevant, pecking the key actually stopped the delivery of reinforcement. The scenario had been changed back to operant conditioning. Contrary to the plentiful evidence that key pecking was an easily conditioned response, Williams & Williams found that autoshaping under these conditions still

occurred. The pigeons still developed the tendency to peck the key; the response contingent omission of reinforcement slowed responding, but did not abolish it.

These examples from the psychological laboratory, plus the burgeoning literature on learning from field studies that could not easily be incorporated into the associationist program (e.g. imprinting, Lorenz 1969; Bird song learning, Marler 1970; Bee learning, von Frisch 1967) led some authors to demand additional principles beyond those in standard behaviourism. These principles would take account of the special demands that adaptation to different environments put on different organisms. It was argued that animals had "species' specific reaction patterns" (Bolles 1971); that different behaviours had different 'preparedness' to be associated with different consequences (Seligman 1970); or that there were a 'multiplicity' of learning processes for a multiplicity of species and niches (Shettleworth 1972).

2.1.2.2 Computational Constraints

An additional set of problems that has not received any attention in the animal literature are the epistemological shortcomings of associative learning theory. Associationism has not had any support in philosophy since the turn of the century (Russell 1914 was the last work to use associationism, according to Oldroyd 1980). The problem with associationism is that without any additional rules to restrict search it must lead to an overloading of the subject's memory system. The injunction to 'associate' is impossible to obey – what should one associate with what? (Compare Popper's discussion of the injunction to 'observe', 1972:259). We must either accept that deciding what would be associated is as important a part of a learning process as the procedure of association itself, or give up the attempt at developing an associative account of learning. Seligman's suggestion that a certain preparedness exists for certain actions to be associated with particular consequences, is therefore by no means peripheral phenomenon – a footnote to behaviourism – but rather a *sine qua non* for a functioning associative theory. It was no surprise, therefore to find that Thorndike (1911; 1932) recognized a concept of 'belongingness' to describe the prepotency of certain actions to be associated with certain results. The notions of preparedness and belongingness will be returned to in chapter 3, section 3.2.3.2.

Riedl (1985) points out that if we were to assume that there were only 16 synaptic junctions between the auditory nerve of a dog and the nerve ending activating the salivary gland, then the number of possible permutations of connection is 16!, about $2 * 10^{13}$. This is about 10^5 times longer than the lifespan of a dog in seconds – so the chance of the correct association, from sound to salivation, being hit by chance once in a whole lifetime is insignificant.

2.1.3 More Modern Approaches

Contemporary psychological approaches to animal behaviour can be distinguished by the positions adopted on two main dimensions of disagreement. These are specialization versus generalization (Plotkin & Odling-Smee 1979) and associationism versus selectionism. Since it is a form of selectionism that is to be advocated here, discussion of theories scoring significantly on this dimension will be delayed until after the exposition of the principles of Evolutionary Epistemology in chapter 3. Specialization versus Generalization refers to whether the authors concerned see particular instances of learning in individual behaviour as special adaptations to particular environmental features, or as a generalized ability to learn from experience.

2.1.3.1 Neo-behaviourism

There is no shortage of researchers tracing their conceptual allegiances back to Skinner, Hull or Guthrie, with little attempt to take stock of the discoveries of the intervening years sketched above. This 'head in the sand' attitude does not lead to an enthusiasm for theoretical statements, and so it is difficult to point to references. Any recent issue of the 'Journal of the Experimental Analysis of Behavior' contains a number of papers whose pedigree is impeccable general process theory – Wearden (e.g. 1983); Herrnstein (e.g. 1982) and Davison (e.g. 1985) are some of the most prolific authors of this school at the present time. To these theorists the biological constraints are a very peripheral phenomenon. Domjan & Galef (1983) suggest that the impact of the 'biological constraints' on general process theory is not great, and that these phenomena can be incorporated into general process theory by the addition of the principles of 'belongingness' (Revusky 1977), and by including stimulus similarity as a factor in Pavlovian conditioning (Rescorla & Furrow 1977; Testa & Ternes 1977). Long delay learning, potentiation in classical

conditioning, and constraints on punishment and avoidance learning, do not require the addition of any new postulates according to Domjan (1983). The holders of this view argue that since the laws of association still hold in a great many cases, that is quite good enough for them, they are happy to experiment in contexts where general process theory still holds true without particular anxiety over the need for *ad hoc* principles to encompass the biological constraints literature.

2.1.3.2 Ecological Psychology

At the alternative extreme of the specialization – generalization continuum to the neo-behaviourists are those who adopt a position emphasizing the many sided nature of the interface between organism and environment, and proposing a multitude of learning mechanisms to parallel this (e.g. Shettleworth 1972; Johnston 1981; 1985; Gottlieb 1985; Reed 1985; Miller 1985; Kalat 1985).

The ecological approach can be distinguished from the biological constraints viewpoint (to be discussed next) on a number of points. Firstly theorists of an ecological persuasion resent the implication that biology sets limits on an otherwise unfettered learning process (Johnston 1985). The ecological approach borrows ideas from ethology and Gibsonian ecological psychology to formulate a program of research emphasizing the need for (at least) three phases to animal learning research. Initially they stipulate a phase of naturalistic observation of organisms in their natural habitats (Johnston 1981;1985; Miller 1985). Johnston advocates that this description should follow Gibson's (1977; 1979) plan of specifying the 'affordances' of the ecosystem for the organisms concerned. Here we see also the influence of ethology, where naturalistic observation has always been considered very important. This initial phase would lead into the specification of a 'task description' – What has to be learned by an animal in its daily life in the natural habitat? Only after this second phase is complete are they willing to countenance experimental manipulation of parameters in order to investigate how the particular examples of learning proceed.

The possibility that some general principles might come out of this enterprise is not ruled out of court, but the question is considered an open one, that must await the results of empirical investigation.

The problem with this approach to animal learning is simply that an

empirical program does not constitute a theory. Domjan & Galef (1983) have pointed out, and Kalat (1985) has conceded, that the ecological critiscms of general process theory have not led to a theory of their own. 'Belongingness', for example, does not seem to be specifiable *a priori*, it, "is not... a precise theory that allows one to predict new results with any confidence" (Kalat 1985). Although Johnston (1985:16) argues that, "it is unlikely that associative principles will contribute as broadly to an ecological view of learning mechanisms as they do to the non-ecological view, because so many of their fundamental assumptions are explicitly opposed to the ecological perspective", he is forced to concede that, "on the other hand, it is unlikely that associative mechanisms play no role in learning under natural circumstances...". Shettleworth (1984) is surely right to be suspicious of the ecological program's ability to deliver when she writes "Johnston... seems to suggest that it [the ecological approach] can be productive in the absence of any pre-existing system for classifying and analysing examples of learning". It does not seem clear to me why we should accept the ecological psychologists' faith in their method to deliver a new paradigm for animal learning research, when after two decades of more ecologically minded investigation on animal learning, associationism remains the only conceptual framework available to order our knowledge of animal learning.

2.1.3.3 Biological Constraints Theorists

A small group of theorists (Bolles 1971; 1985; Lea 1984; Revusky 1977; 1984; 1985; Terrace 1984; and perhaps Shettleworth 1984) whilst fully cognizant of the impact of the biological constraints research on learning, retain none the less, the opinion that learning is a general process. This group therefore lies between the ecological school, and neo-behaviourism, and merges at the edges with both of these. Plotkin & Odling-Smee (1979) point out that to jettison the idea of there being something that makes all instances of learning similar is to jettison parsimony in a big way. Bolles (1985), Revusky (1984; 1985), Terrace (1984) and Shettleworth (1984) are agreed that there is not yet sufficient justification to take this step. Shettleworth expresses well this reluctance to throw out the baby with the bathwater when she suggests the two different ways that learning could come to be a general process:

> Associative learning capabilities may at first operate only on certain ecologically relevant combinations of events..[and]...gradually the ability to form associations becomes

more general...An alternative view is that pain-avoidance learning and other examples of association learning with specialized functions are manifestations of a generally accessible associative mechanism, although it has to be granted that the general mechanism is more responsive to some combinations of events than others

Shettleworth 1984:423

Either way associationism remains the general principle into which the biological constraints must be fitted. In contradistinction to the neo-behaviourists however, Shettleworth, Terrace and Bolles (at least) emphasize the importance of the organism's habitat in understanding instances of selective association. They see the niche as important to understanding behaviour, but not as something from which predictions can be made. Bolles (1985) asks, "How can a description of the niche enable us to specify anything about the psychology of an animal?". This is the most marked difference between the biological constraints approach and the ecological viewpoint. Johnston (1985) and Reed (1985) argue, not just that the environment can tell us something about the organism, but that (following Gibson 1977) the only useful description of the environment is in terms of the affordances it offers to the organism.

2.1.3.4 Selectionist Accounts

Discussion of a fourth type of theory – selection theory – which is similar to the theoretical viewpoint to be developed in chapter 3, will be deferred until after the principles of Evolutionary Epistemology have been put forward.

2.2 Ethological approaches

An important part of the role of ethology in the development of the present situation in animal behaviour science has already been discussed in section 2.1.2.1, in the context of the problems of behaviourism. Ethology, however, has an independent existence as well. The ethological approach to behaviour emphasizes painstaking observation of animals in their natural settings. Where comparative psychology has focused on environmental factors, ethology has emphasized innate, heritable behaviours. The most important explanatory principle to ethologists is that "as biologists thinking in terms of evolution

through natural selection, they are aware that an understanding of the diversity in the animal kingdom, requires that the behaviour of each species be seen in relation to the environmental context to which it has been adapted", Hinde (1982).

In his seminal work 'The Study of Instinct' Tinbergen (1951) distinguishes two forms of explanation of importance to ethologists. These are the functional and the mechanistic forms of analysis. Functional analysis refers to the teleonomy of behaviour – the (ultimately evolutionary) ends towards which it is directed. A functional analysis addresses the 'Why?' question in behaviour. A mechanistic analysis is directed at finding the immediate causal factors that are controlling a behaviour pattern. The question being addressed here is that of 'How?' the behaviour is controlled.

The major difference between ecological psychology and ethology seems to be that, as I mentioned above, ethologists emphasize both functional and mechanistic analysis, whereas the Gibsonian school of psychology argues that all behaviour can be explained in functional terms (Reed 1985). In practice this difference of opinion is not great, because with the success of optimality theories, functional analysis has come to dominate behavioural ecology over the last twenty years (see, for example, Krebs & Davies 1984).

2.2.1 Functional Analysis

> The entire modern deification of survival *per se*, survival returning to itself, survival naked and abstract, with the substantive excellence in *what* survives, except the capacity for more survival still, is surely the strangest intellectual stopping-place ever proposed by one man to another
>
> William James (quoted without reference to source in Hofstadter 1955:201)

In a post Darwinian age, functional explanation of behaviour is synonymous with a discussion of the fitness or adaptedness of behaviour. This has meant that ethologists have run straight into the problem of how to define individual fitness non-tautologically. Darwin himself preferred to talk of 'Natural Selection' and only adopted Herbert Spencer's phrase 'Survival of the Fittest' on the prompting of Thomas Huxley, who argued it would overcome the problem some of their contemporaries had in coming to terms with the idea of nature

doing something so apparently purposeful as 'selecting' (Gould 1980; Reid 1985). The above quotation (probably from one of James' many tirades against Spencer) draws attention to the tautology that arose once one started to speak of the 'survival of the fittest' – where 'the fittest' could only be defined as 'those that survive'.

The neo-Darwinian redefinition of natural selection as differential survival solved the tautology, but produces a criterion of fitness that can only be applied at a multi-generational timescale. At a certain perspective, it might be possible to observe that a certain group of animals behaving in a certain way had become extinct, or, conversely, had come to dominate a new habitat at the expense of other groups. But this perspective would require research over eons; at a day to day level differential survival does not offer much for the scientist interested in the behaviour of organisms with lifespans over half an hour. This may well be one of the reasons why comparative psychology has been happy to stand outside evolutionary theory till recent years.

2.2.1.1 Optimality Theory

In 1966 MacArthur and Pianka published a paper suggesting that because foraging was an activity that was critical to survival and reproduction (and hence inclusive fitness), a selection pressure existed such that competition between individuals would lead to a state where the observed behaviour would maximise energy intake, in a manner corresponding to what would obtain if the animals concerned had been built by an omnipotent designer. The incorporation of the mathematical maximization models of economics into ethology appears, at least superficially, to have been a massively successful enterprise. Review after review has praised the effectiveness of 'Optimal Foraging Theory' (as this approach is known) to produce a framework for the functional study of behaviour (Krebs 1973; 1978; Krebs & McCleery 1984; Pyke, Pulliam & Charnov 1977; Kamil & Roitblat 1985).

Do we really expect evolution to produce organisms whose behaviour conforms to what would be chosen by an omnipotent designer? Some evolutionary theorists have argued that this is the case (e.g. Oster & Wilson 1978; Maynard Smith 1978); on the other hand the point is not uncontroversial. D'Arcy Thompson recognized the strangeness of rejecting a supernatural designing deity for a process that had the net effect of producing the same

results as the deity would have done:

> To buttress the theory of natural selection the same instances
> of adaptation are used, which in an earlier but not distant age
> testified to the wisdom of the Creator.

D'Arcy Thompson (1917:672) cited in Reid (1985:34)

Gould and Lewontin (1979; Lewontin 1979; Gould 1982) have recently reopened the old sore of what they term the 'Panglossian Paradigm'. (Dr Pangloss was the character in Voltaire's 'Candide' who believed that everything is made for the best possible purpose). Gould's argument (1980:27) is that the historical nature of living things means that the designer's optimal solution is precluded. He points to the opportunist nature of the evolutionary process, cobbling today's survival together out of the parts left over from yesterday's. Natural selection does not set goals, it sets criteria – 'pass marks' if you like. So long as an individual does better than some minimum it will survive to reproduce, and its progeny will form the next generation. Evolutionary theory does not postulate an overseer who inspects ecosystems to make sure they are up to scratch. The view that individuals need only do just well enough to get by is known as 'satisficing' (Simon 1956), and it is an argument many theorists want to distance themselves from. Dawkins (1981), for example, concedes that the designer's optimal solution is ruled out by the opportunist nature of selection. ("Nature does not have the foresight to put together a sequence of mutations which, for all that they may entail temporary disadvantage, set a lineage on the road to ultimate global superiority" [ibid:46]). On the other hand he argues that 'satisficing' is not strong enough. "Living things are not selected for their capacity simply to stay alive"; he writes (ibid:45), "they are staying alive in competition with other things". Dawkins proposes the term 'meliorizing' (meaning 'bettering') to cover this viewpoint. I feel that this additional term is unnecessary. Simon's formulation predates the recent discussions of optimality in behaviour and seems to confuse the issue we now understand as 'satisficing' with what are now known as 'rules of thumb'. Rules of thumb are proximate behavioural mechanisms that could be aimed towards achieving either an optimal, or a 'satisfactory' goal state. I see no reason therefore why we should not simply recast 'satisfying' to allow that organisms must do satisfactorily in the context of competition from other organisms. We could call this 'satisficing under competition', to show that we are thinking in terms of a base line exceeded, rather than a goal attained (as is implied by optimisation). Dawkin's new term is unfortunate, at least within the context of this thesis,

because it is also the accepted name of a theory of behaviour on concurrent variable interval schedules (Herrnstein & Vaughan 1980), discussed in chapter 8.

The distinction between satisficing and optimising is in some ways equivalent to the physicist's concepts of stable and unstable equilibria. An optimal policy might be expected to become dominant in a population because, being by definition the best possible balance of costs and benefits for each individual, it cannot be beaten by any mutant strategy (Houston 1986; though cf Harley 1983). This is equivalent to the physicist's concept of a stable equilibrium. There is another form of equilibrium, however, the unstable form. This describes systems that are not in their most stable situation, but that, nevertheless, may endure for a very long time. In biology this is equivalent to saying that a population may be vulnerable to invasion by a mutant strategy, but that such a mutant either does not arise, perhaps because variation including the possibility of such an individual does not occur; or that the mutant's improved efficiency on the dimension under study entails an unacceptable loss of efficiency on some other fitness-related parameter. Even if the new mutant thrives, its increased fitness may have such a small effect that it takes an indefinitely long time to invade the population. This time scale could be slower than the rate of exogenous environmental change, and so our chances of observing optimal behaviour in all but a minority of a population will be nill.

One of Ollason's (1986) arguments against optimal foraging is that since Hume has shown that it is not possible to demonstrate divine design in the functioning of living things, it must be equally impossible to demonstrate design by natural selection – even if we conceded that natural selection could be a designing agent.

One further argument against observing optimal behaviour can be developed from Plotkin & Odling-Smee's (1979) discussion of the distinction between adaptedness and adaptability. In a stable environment, they suggest specialization will be selected for, and optimal adaptedness to the contemporary environment becomes a possibility. In a changing environment, however, the generalist ability to respond flexibly to changing circumstances (adaptability) will be at a premium. The alternative strategy of possessing multiple specializations, would involve expensive redundancy. In a sense there must be interference between adaptedness and adaptability such that an

individual cannot be both at once. The net effect of this argument is that if we expect to find optimal foraging anywhere it must be amongst evolutionarily conservative organisms that have specialized in niches that have remained steady and predictable across evolutionary time. On the other hand we can predict that opportunist organisms, consuming a varied diet, in a continuously changing environment would have been selected to have sacrificed optimality in any one situation for behavioural flexibility in a variety of contexts. Organisms that follow man, such as feral pigeons (*Columba livia*) and rats (*Rattus rattus*) would fit into this latter category.

Even if we accept that natural selection could tend to optimise behaviour, there are a number of reasons why we could not expect to observe optimal behaviour (cf Pierce & Ollason 1986; Ollason 1986; McFarland & Houston 1981; Staddon 1983:493).

An animal may be marginalized by competitor pressure, evolutionary lag following environmental change, or some other factor, out of the environment in which its foraging strategy is optimal to a niche in which its foraging is less efficient. McFarland & Houston (1981) refer to this as the case where the animal's objective function (the function it is maximising) is distinct from the cost function (the function it ought to maximise to get maximal fitness in this environment). McFarland & Houston further point out that it will rarely be the case that these two functions are the same.

A point that has not been emphasized in the foregoing account of optimal foraging theory is that, although evolution might be assumed to design optimally behaving fitness maximisers, our investigation of this phenomenon can only commence once we have identified a more proximate currency to stand in for fitness in our calculations. Typically in foraging studies net energy gain is used, but there are other possibilities - e.g. minimisation of the probability of starvation, levels of certain critical nutrients, minimisation of exposure to predation etc, etc. In general we can never be certain that the currency we have chosen to maximise is the same as that being maximised by the subject of our studies. And, more problematic still, except in (very) simple restricted environments, it must be the case that our subjects will be obliged to maximise all these functions at once. Foragers must not just maximise energy intake, but minimise risk of predation and maintain water balance and the levels of critical nutrients etc, etc all at the same time. I'm not sure if such

an analysis is impossible in principle, however, in practise it must be.

An approach that inverts the logic of the argument sketched above is the 'inverse optimality approach' (McFarland & Houston 1981:153ff). For an inverse optimality analysis, optimality of behaviour is assumed, and the investigator's interest lies in deducing the objective function which the behaviour maximises. This approach offers the possibility of uncovering the function the organism is really maximising, as opposed to the function the investigator assumes is being maximized. Such an approach is dangerous for a number of reasons. McFarland & Houston (1981) consider two possibilities. Firstly, there is no guarantee that the function obtained is actually governing the organism's behaviour. Houston (1986) goes further and states "The first law of model building is that for any set of data there is at least one model (Kalman 1968). The relevant question is therefore not 'can a model be found?' but 'what do I want from a model?'" The danger here is that even if the animal in question does not proceed by maximising a function, it will none the less be possible to deduce a function it could be maximising.

The second problem with the inverse optimality approach is that the mapping from behaviour to functions may not be continuous (McFarland & Houston 1981:154). The gist of this is that small changes in behaviour can require massive changes in the function assumed optimized.

If optimal foraging theory has so many conceptual problems the question naturally arises How come it is so successful? One line of defence of optimal foraging theory has been that its utility justifies its use. Ollason (1986) argues that the success of optimal foraging theory is due to it having been set up in a tautological and therefore unassailable position. He suggests that combination of prior optimality analysis with inverse optimality analysis has allowed a situation to develop where an investigator (Ollason's 1986 example is Cowie 1977) can initially perform an optimality analysis, and then measure some relevant parameters of the behaviour. Finding a discrepancy between prediction and observation, the investigator then works backwards from the behaviour to the function it must be maximising, and/or simultaneously adding in constraints to the prior optimality analysis until the two directions of optimality analysis meet up. This creates a circularity in the argument, and produces optimality as the inevitable result.

The other way that optimality tends to be an assured result is that the prior predictions have often been so weak.

> Unfortunately, most of the predictions thus far derived from optimal foraging theory are merely qualitative. For example, a forager with a choice of different quality prey should prefer the better ones and should be more selective when the rate of encountering good prey is higher. A forager searching for patchily distributed prey should stay longer in patches with more prey and should be more reluctant to leave a patch of given quality when overall prey density is lower or travel time is longer. These predictions are so weak that they can hardly fail to be correct.

> Green 1985:338

Although optimal foraging theory may have been abused, is it still possible that it might have a legitimate function? Although its predictions have been weak, is it possible they might be made stronger? I think the answer to both these questions is a qualified 'yes'. The second question is the easier to deal with so I will take it first. True, the marginal value theorem (Charnov 1976; probably the first optimal foraging theory to tested widely), does make fairly intuitive predictions, but in principle optimal foraging theories may be as complex as we care to imagine. It is more the case that researchers have tended to water down the predictions of optimal foraging theories in order to increase the apparent agreement between observation and theory (this is discussed in the context of laboratory test of optimal foraging in the introduction to chapter 4). The optimal foraging analysis presented in chapter 5, for example, is fairly complex and precise in its predictions.

To the first question above (does optimal foraging theory have a legitimate function?) the answer is a more qualified affirmation. I do not believe that optimal foraging theory is a theory of behaviour. I think there are enough reasons why we cannot expect animal behaviour to be optimal with respect to any prior definable function, that we must concede that predictions about behaviour cannot be drawn from optimal foraging theory. There may be some restricted conditions where optimality might be expected (species specializing in very conservative habitats), but I am not convinced that we can specify in detail what these conditions are. On the other hand I do see a function for optimal foraging theory (there is no paradox here − many theories that are not theories of behaviour have a role to play in the study of behaviour; theories of

physics, biochemistry, physiology etc). To ask, 'What should this animal be doing, if it is going to maximise energy intake?' is an informative question. The discrepancy between optimality and actual behaviour is an interesting one, and one that should not be brushed under the carpet, but brought out and examined. In some cases an inverse optimality analysis may help to define the difference between the observed behaviour and that which seemed *a priori* to the investigator to be optimal. This could provide insight into the functions of the behaviour and raise suggestions as to the mechanisms of the behaviour. In fact, optimality theory was used implicitly in this way in early psychological experiments. In maze running, for example, the notion of the most direct path was always defined before ever a rat was run.

From this perspective then the important thing is to set up optimality analyses with as few constraints as practicable, and work out the predictions of the analysis *prior* to observation. What we must *not* do is to incorporate knowledge of our subject's behaviour as constraints into the optimality formulation without first attempting an unconstrained optimality analysis. The informal incorporation of constraints makes the subjects behaviour seem much more optimal than it really is (Houston & McNamara 1985). In analysing the results of subsequent experiment or observation it is important to emphasize the *differences* between theory and observation. Similarity tells us very little; it is the departures from optimality that raise questions as to how the organism is actually controlling its behaviour (Staddon 1983:6). It is this form of optimality analysis that is attempted in chapters 5 and 7.

2.2.2 Mechanistic Analysis

Although, as mentioned above, ethology recognizes two forms of analysis, it is functional analysis, in the form of optimality theory, that has dominated behavioural ecology in the last decade. It is only quite recently that it has been recognized that an optimality analysis leaves unaddressed the mechanistic question as to How the behaviour is controlled (Harley 981; Houston 1980). A number of 'rules of thumb' (as theories of mechanism that could approximate optimal foraging are known), have been proposed to account for behaviour in particular situations (e.g. Ollason 1980; McNamara & Houston 1985; McFarland & Houston 1981; Ydenberg 1984). Rules of thumb can be subdivided into those that do not change strategy with increasing experience (which I shall call 'static rules'), and rules that do so change ('learning rules'). These can both be dealt

with fairly briefly because there is only one theory of rules of thumb. This is the hypothesis that the rules constitute the optimal solution to the problem of finding the best overall return of benefits on costs, after allowing for the cognitive costs of calculating a more perfect solution, as opposed to some simple rule.

2.2.2.1 Optimal Rules of Thumb

Considering static rules of thumb first, a typical pattern, unfortunately, is for the experimenter to incorporate the observe patterning of behaviour into an optimality model as a constraint, and thereby the behaviour becomes optimal by fiat (Ollason 1986). Since there is no way of estimating the cognitive cost to the subject of evaluating alternative strategies of behaviour which it is not capable of, there is no non-tautological way that optimality considerations can be applied to static rules of thumb. All that can be done is to argue that an alternative strategy that the subject is not performing, but appears capable of, would offer better returns than the behaviour observed. This could be interesting, particularly where behaviour in the situation observed follows a more or less optimal rule than behaviour in some other, apparently similar, situation (say, in cases where choice between patches of water was more optimal than between patches of food). It should be emphasized, however, that such an approach could never be prescriptive.

For learning rules the situation is very similar. If the observed learning rule is incorporated as a constraint into the optimality analysis, then the behaviour is almost inevitably optimal. Learning rules could be ranked in terms of how close they come to the globally optimal behaviour. This could be a worthwhile exercise, though since different rules might assume different amounts of prior knowledge, a fair comparison may not be possible. Since we have no knowledge (nor could we have), of the relative cognitive costs of different learning rules to any particular subject (except, perhaps, ourselves), it is necessarily impossible to do a full cost-benefit analysis to calculate the truly optimal learning rule.

Ollason (1986) has argued that optimal foraging theory should be scrapped in favour of simply generating and testing rules of thumb. Most authors, however, seem to prefer a flawed conceptual schema to no schema at all.

2.3 Attempts at Reconciliation

As noted at the beginning of this chapter there is now evidence of a willingness on the part of both psychologists and ethologists to unite together in a search for a wholistic science of animal behaviour. I am aware of only two ways in which this has been attempted. The first lies in the development of ecological psychology.

Many of the concepts and approaches advocated by the ecological psychologists (section 2.1.3.2 above) are not alien to ethology. The early ethologist Jacob von Uexkull defined an organism's 'Umwelt' as that part of its environment to which it is receptive. He wrote;

> The sum of the stimuli affecting an animal forms a world in itself. The stimuli...form certain indications, which enable the animal to guide its movements, much as the signs at sea enable the sailor to steer his ship. I call the sum of the indications the *world-as-sensed* [Umwelt].

> von Uexküll 1926:123

This emphasis on the econiche, combined with Lorenz's advocation of some of the principles of Gestalt psychology (Lorenz 1971/1959), generates a position having many points of contact with the ecological psychology of J.J. Gibson (compare Gibson's notion of 'affordance' with Koffka's 'valences'. Gibson 1977).

Difficulties for this approach include the absence of any suggestion on the part of the ecological psychologists to cover the inadequacies of the theory of optimality in ethology. Ecological psychology can be characterized as ethology without any theory.

The other approach to reconcile ethology and psychology is less an attempt at integration, as at demarcation.

Fantino & Abarca (1985) proposed that a psychological theory (such as their 'delay reduction hypothesis') represented a rule of thumb that could account for observed behaviour at a more proximal level of explanation than the equally valid functional explanation of optimal foraging theory ("A model of natural selection for foraging efficiency makes several predictions consistent with a hypothesis of a more proximate phenomenon (reduction in time to reinforcement)"; Fantino & Abarca, 1985:319). The majority of commentators on their target article agreed with this view (Barnard 1985; Gass 1985; Green 1985;

Houston 1985; Lea 1985; McNair 1985; Mellgren 1985; Shettleworth 1985a).

If psychological and ethological theories did occupy different levels of description then it would be possible that they could be simply sandwiched together. I do not think that the resulting composite would be a very satisfying, conceptually integrated, theory of behaviour, but it certainly might be a possibility. The problem here is that, in fact, the delay reduction hypothesis, like the matching law and most psychological accounts of behaviour, is a functional theory – not a mechanistic one. Fantino & Abarca concede this at one point in their article, and two of the commentators (Branch 1985; Kruse 1985), also pick up on this. Staddon (1983:503) writes, "Fantino's model and any optimality analysis, is parameter free, but says nothing about the means animals use to adapt to these procedures".

The absence of a tradition of explicitly distinguishing functional and mechanistic theories of behaviour in psychology has meant that theorists have been able to switch between the two levels at will. The delay reduction hypothesis is a good example. Stated in words it appears to be a mechanistic hypothesis – a statement as to How behaviour is controlled vis:

> A stimulus correlated with a greater reduction in time to food presentation will be a stronger reinforcer (as measured in a choice test) than one correlated with a longer reduction in time to food.

Fantino & Abarca 1985:316

When formulated mathematically, however, it simply represents a functional generalization that could ·be attained by many different true rules of thumb. The delay reduction hypothesis is not unique in this respect. Herrnstein's 'matching law' (discussed in detail in chapter 6), was initially formulated as a principle sounding like a rule of thumb (Herrnstein 1961). Matching stipulated that the subject should distribute it's responses across a number of options, in proportion to the reinforcements it obtained from those options. From this proximate principle an equation was derived stating that the ratio (or proportion) of reinforcements at the end of a session should equal the ratio (or proportion) of responses made to the available options. Herrnstein (1982; Herrnstein & Vaughan 1980) now concedes that matching as originally formulated is not a mechanistic theory of behaviour. Two examples of mechanistic theories are discussed in more detail in chapter 8. My purpose here has simply been to argue that psychological theories are not in general,

though in particular cases they may be, mechanistic theories; and so the suggestion that they could provide the mechanistic side of a demarcation of theorising between psychology and ethology, is not tenable.

2.4 Conclusion

This chapter has been largely negative. I have argued, on the one hand, that comparative psychology has not succeeded in absorbing the 'constraints on learning' that became apparent in the 1960s; but that, at the same time, ethology has no adequate conceptual framework to cope with behaviour that responds flexibly to the environment. Although optimal foraging has a role to play in helping us to understand behaviour, it does not, in and of itself, constitute a theory of behaviour. This shows itself most acutely in the absence of any viable theory of rules of thumb. At present rules of thumb in ethology are simply pulled out of a hat.

What I feel is required, and what I shall argue for in the next chapter, is a functionally transparent theory of rules of thumb. By 'functionally transparent' I mean that it should be possible to see function through it. Although it will not be a functional theory of behaviour (predictions about behaviour will not be made from solely functional considerations), it will shed light on the evolutionary function of behaviour. This will become clearer (I hope) once the theory is set out.

CHAPTER 3
EVOLUTIONARY EPISTEMOLOGY:
A NEW WAY FORWARD

The idea I want to explore in this chapter is not original. It has come largely from the minds of two men; Donald Campbell- social psychologist, and Sir Karl Popper- philosopher. The notions I will explore here may seem at first to be at an altogether too general level to have anything to do with animal behaviour. In the previous chapter I attempted to show that none of the extant theories of flexible animal behaviour were fully satisfactory, and, in so doing, to prepare the reader for the idea that a radical reassessment of basic principles was in order. The latter half of this chapter will come back down to a more directly applicable level of theorising.

3.1 Outline of Evolutionary Epistemology

Let us consider three basic questions. What do we mean by evolution? What do we mean by learning? How are these two processes related? Consider Campbell's answer:

> <u>Hypothesis:</u> All knowledge gain is by Blind Variation and Selective Retention (BV&SR).
>
> <u>Corollaries:</u>
>
> 1. A BV&SR process is fundamental to all inductive achievements, to all genuine increases in knowledge, to all increases in fit of system to environment.
> 2. In such a process there are three essentials; (a) Mechanisms for introducing variation; (b) Consistent selection processes; and (c) Mechanisms for preserving and/ or propagating the selected variants.
> 3. The many processes which shortcut a more full BV&SR process are in themselves inductive achievements, containing wisdom about the environment achieved originally by BV&SR.
> 4. In addition, such shortcut processes contain in their own operation, a BV&SR process at some level, substituting for overt locomotor exploration or the life- and- death winnowing of organic selection.
>
> from Campbell (1974)

Campbell's schema, then, is for a universal biology of knowledge. One process, Blind Variation and Selective Retention, is not only responsible for

individual learning, but for 'all increases in fit of a system to its environment', in other words for evolution itself.

I will attempt to make this proposal clearer by exploring some of the objections that may be raised to it. I will consider the following: How can evolution be considered a learning process? (section 3.1.1). What entitles us to call 'Blind Variation and Selective Retention' a knowledge gaining process? (section 3.1.2). Can a psychological theory be built from a philosophical theory? (section 3.1.3). Finally I consider; How can such a global schema have any implications for individual learning? (section 3.2).

3.1.1 How can evolution be considered a learning process?

> A remarkable parallel which I think has never been noticed, obtains between the facts of social evolution on the one hand, and of zoölogical evolution as expounded by Mr. Darwin on the other.

> James 1979/1890: 163

William James appears to have been the first of a string of writers who have noticed parallels between individual and/ or social learning, and the primary process of "zoölogical evolution as propounded by Mr. Darwin". Campbell (1974) cites 42 sources of this idea; some of the more recent advocates include Pringle (1951); Russell (1962); Staddon (1975; 1983); Popper (1972/1961; 1972/1965); Lorenz (1975/1941; 1977); Riedl (1984/1981); Vollmer (1975); Wuketits (1986).

It seems fairest to follow James' account, since he has historical precedence.

> I shall try to make plain....that the features of our organic mental structure cannot be explained at all by our conscious intercourse with the outer environment, but must rather be understood as accidental in the first instance ([James' footnote:] 'Accidental' in the Darwinian sense, as belonging to a cycle of causation inaccessible to the present order of research), but then transmitted as fixed features of the race.

> James 1981/1890:1212 (p613 of earlier editions)

James makes these remarks in the context of a discussion over where our a priori knowledge comes from. 'A priori' is the term used by the german philosopher Kant for knowledge of the world that cannot be deduced from

experience of the world. A priori knowledge was considered to cover such things as the axioms of geometry and logical principles which appeared to be true even though no experience could prove them so. James argued that human knowledge of this kind had become fixed 'in the race' by the same Darwinian process of blind variation and selective retention as any other adaptation.

Fifty years later Konrad Lorenz considered the same question and came to a very similar conclusion.

> One familiar with the innate modes of ... reaction of subhuman organisms can readily hypothesize that the *a priori* is due to hereditary differentiations of the central nervous system which have become characteristic of the species, producing hereditary dispositions to think in certain forms. One must realize that this conception of the *a priori* as an organ means the destruction of the concept: something that has evolved in evolutionary adaptation to the laws of the natural external world has evolved *a posteriori* in a certain sense, even if in a way entirely different from that of abstraction and of deduction from previous experience. We believe we can demonstrate the closest functional and probably genetic relationship between these animal *a prioris* and our human *a priori.*

Lorenz 1975/1941

Lorenz, I feel, weakens his argument by his insistence that the learning of the species in evolution is 'entirely different' from the learning of individuals 'by deduction from previous experience'. This demarcation stems from the need he felt to distinguish clearly between the field of ethology and that of comparative psychology (e.g. Lorenz 1971/1935 cited in Hailman 1985).

The reader may perhaps be surprised by the absence of scare quotes round the word 'learning' in the previous paragraph, where I have spoken of 'species learning'. The reason for this is that I accept Hull's (1982) argument that species are not 'natural kinds' (a philosophical abstraction characterized as eternal, immutable, discrete and important); but rather are better considered as spatio-temporally localized individuals – historical entities which exist only in one time and place (Hull 1982: 279ff). If the reader does not agree with this view, and objects to the idea of an abstraction like the genome 'learning', he or she may simply imagine scare quotes round the word 'learning' when it is applied to species.

Plotkin & Odling-Smee (1979) formulate a distinction which is useful in this

context. They talk of the 'order' possessed by an environment, to which the 'organization' of an organism is an adaptation. To the extent that an individual's organization matches elements of the order of the environment, that individual can be said to have knowledge of the environment. To the extent that the process of evolution by natural selection consists of the tracking of environmental order by biological organization, it can be considered a knowledge gaining (learning) process.

Now that we have reconsidered evolution as a knowledge gaining process this provides our first clue that evolutionary epistemology might help us to integrate individual learning into the evolutionary synthesis.

3.1.2 What makes BV&SR a knowledge gaining process?

To those familiar only with learning within an associative paradigm, the idea that 'Blind Variation and Selective Retention' could be called a learning process may seem foreign. In the previous section we have noted that evolution may be considered a knowledge gaining process; and of course it is well known that evolution proceeds by the production of variants, from which the survivors are selected. There are other theoretical reasons why we can consider BV&SR to be a knowledge gaining process.

The philosopher Karl Popper set out in his classic 'Logic of Scientific Discovery' (1980/1934) a novel principle for scientific progress.

> According to to my proposal, what characterizes the empirical method is its manner of exposing to falsification, in every conceivable way, the system to be tested. Its aim is not to save the lives of untenable systems but, on the contrary, to select the one which is by comparison the *fittest* by exposing them all to the fiercest *struggle for survival*....We choose the theory which best holds its own in competition with other theories; the one which by *natural selection* proves itself the *fittest to survive*.

Popper 1980/1934:108 my emphasis

Popper, at this time, was highly critical of evolutionary theory (see Popper 1972:241), and so the fact that he used an evolutionary analogy in his original exposition of his theory of knowledge is all the more striking evidence for the homology between inductive progress and evolution.

The critical components of a Popperian learning system can be summarized in three points. Firstly the learner must generate the widest possible range of

hypotheses concerning the nature of the problem under investigation. Secondly these hypotheses must be subjected to the strongest possible attempts at refutation (not confirmation). And thirdly, from those theories that survive our sternest efforts at refutation a new range of hypotheses must be generated, in the manner of the first stage- but this time informed by the attempted refutations of the second phase. This process, clearly, is iterative and can be repeated ad infinitum; each cycle producing slightly more complete knowledge than the cycle before.

It is conventional to emphasize the 'refutational' aspect of Popper's theory of knowledge, as opposed to the principle of 'confirmation' that had previously been advocated in theories of scientific progress. There are two additional points I wish to draw attention to.

Popper's development of an evolutionary epistemology out of his theory of knowledge gain gives a certain legitimacy to the claim that BV&SR can be considered as a learning process, but there are distinctions between Popperian epistemology and evolutionary epistemology. It is important to draw attention to these in order to be able to argue that the evidence that individual and social learning in scientific communities and elsewhere does not process in a Popperian manner, is not relevant to a discussion of the merits (or otherwise) of evolutionary epistemology. The first of these is simply that Popper's theory of scientific progress is a normative theory of how scientists should behave; not a sociological description of how they actually act. In so far as they are solving problems, Popper claims that his method is the only one they can be using; but they may be doing this only a small proportion of the time, and quite unconsciously. Evolutionary Epistemology, on the other hand, is held to apply to evolving entities at all times- it is not a normative theory, but a prescriptive theory of what is actually going on.

3.1.2.1 Learning From Mistakes

A second point, which I do not think has been previously emphasized sufficiently, is that a Popperian epistemology requires a learner to 'learn from its mistakes'. Although Popper's epistemology is often characterized by the phrase 'conjectures and refutations', in the development of his evolutionary epistemology the importance of refutation- or learning from one's mistakes- has been diluted. I emphasize this point because Popper himself has

advocated an evolutionary epistemology very similar to Campbell's, in three essays; 'Of Clouds and Clocks' (1972/1965'; 'Evolution and the Tree of Knowledge (1972/1961); and 'Evolutionary Epistemology' (1983). Popper argues that organisms are continuously involved in problem solving by trial and error. Error elimination proceeds by killing off unsuccessful forms and so the living organism is always at the spearhead of its phylum (accounting for the 'progressive' appearance of evolution). The pattern Popper envisages parallels the three phases of induction outlined above. A problem is followed by a set of tentative solutions (in the form of a variety of organisms); this is followed by a phase of error elimination (selective death and failure to reproduce of some organisms); and then a new problem presents itself; and so the cycle continues. Popper integrates his evolutionary epistemology with his theory of human knowledge with the thesis:

> Our schema allows for the development of error-eliminating controls (warning organs like the eye; feedback mechanisms); that is, controls which can eliminate errors without killing the organisms; and it makes it possible, ultimately, for our hypotheses to die in our stead.

> Popper 1972:244

But this does not allow the different between zoological and individual learning- that individuals really can learn from their mistakes, whereas, as Lorenz (1977/1973) points out, the genome learns from its successes. In order to learn from one's mistakes it is necessary to keep some kind of record of previous hypotheses that failed the test of refutation. The genome only keeps copies of successful alleles, no record is kept of failures (except so far as they become recessive- and then no record is kept of the fact that they were once refuted). The genome can and does make the same 'errors' repeatedly. This difference is not problematic for evolutionary epistemology; it is simply a point on which it does not offer a prediction, and on which it differs from a truly Popperian epistemological system.

3.1.3 Can a psychological theory be built from a philosophical one?

Although I have tried to show in the above two sections that changes have had to be made to Popper's epistemology before it could be considered a viable theory of biological knowledge (although these are largely unacknowledged); the idea of tapping a philosophical theory as the basis of a

biological theory may still seem strange. This has however, happened once before, when learning theory in psychology was formed out of British empiricist philosophy. A brief consideration of that history (drawn, almost entirely from secondary sources) may firstly serve to legitimize the present exercise, and secondly accentuate the inadequacies of associationism.

John Locke (1632 to 1704) and David Hume (1711 to 1776) are generally considered to be the originators of associationism.

> Locke maintained that knowledge consists of simple ideas which are not analysable into other ideas and complex ideas which are composed of simple ideas linked in temporal trains or in synchronous complexes. In this proposition Locke asserted the central tenet of associationism.
>
> O'Neil 1968:26

Hume's view of the importance of the relation between cause and effect has already been cited in chapter 2, section 2.1.1.

David Hartley's (1705 to 1757) role was also important.

> Hartley, in his "Observations on Man" (1749) had worked out a psychological theory – associationism – which depicted the human mind and human knowledge as being built up by the operation of a few psychological laws upon the materials presented in sensation. Not at first creating any stir, associationism was kept alive by a few, although ardent, supporters until James Mill seized upon it and developed it in his "Analysis of the Phenomena of the Human Mind" (1829), which his son was later (1869) to edit and annotate.
>
> Passmore 1957

James Mill's (1773 to 1836) filling out of this plan begins to sound like Watson's Behaviorism (1924).

> With procrustean efficiency (James) Mill fitted all mental phenomena to this frame. Language is the association of words with ideas of objects; predication is the association of subject and predicate; memory is the evocation, in accordance with associations, of ideas of formerly sensed objects; and so on.
>
> O'Neil 1968: 32

Boakes (1984), Thomson (1968), Flugel (1951) and O'Neil (1968) recount how

this continuous philosophical tradition was absorbed into the fledgeling empirical science of psychology. The Scots philosopher Alexander Bain (1818 to 1903) was a student of John Mill's, and adopted his associationism into a set of strictly psychological principles. Herbert Spencer did something similar in his 'Principles of Psychology' (1855).

William James summarized these ideas in his "Principles of Psychology" (1981/1890). James' position was a subtle one. As we have seen, for the genesis of ideas he adopted what we may call an Evolutionary Epistemological position and was scathing of the notion that ideas could be constructed by the Spencer– Bain law (as the principles of associationism were known).

> Mr Spencer's opinion that our consciousness of classificatory, logical, and mathematical relations between ideas is due to the frequency with which the corresponding 'outer relations' have impressed our minds, is unintelligible.

James 1981/1890:1268

However, he was perfectly happy to consider *habits* to be learnt in this way. His reformulation of the 'Spencer– Bain' law into a physiological language was the key to applying these ideas to other species than mankind.

> When two elementary brain-processes have been active together or in immediate succession, one of them, on re-occurring, tends to propagate its excitement into the other.... The amount of activity at any given point in the brain cortex is the sum of the tendencies of all other points to discharge into it, such tendencies being proportionate (1) to the number of times the excitement of each other point may have accompanied that of the point in question; (2) to the intensity of such excitements; (3) to the absence of any rival point functionally disconnected with the first point, into which the discharges might be diverted.

James 1961/1896: 123-4.

A large part of what made Associationism useful to early psychologists was that it could be considered to work at many different levels. James made a move away from mentalistic structures by referring to "Brain-Processes", even though he talked of them in much the same way as James Mill had talked of ideas. The influence of the purely speculative theory of associationism on Thorndike is very clear in the introduction to his original puzzle box paper: "Animal Intelligence: An Experimental Study of the Associative Processes in Animals"

> The importance to comparative psychology in general of a more scientific account of the associative process is evident....The main purpose of the study of the animal mind is to learn the development of mental life down through the phylum, to trace in particular the origin of human faculty. In relation to this chief purpose of comparative psychology the associative powers acquire a role predominant over that of sense powers or instinct, for in the study of associative powers lies the solution of the problem...the cognitive side of consciousness has changed not only in quantity but also in quality...The association of ideas proper, imagination, memory, abstraction, generalization, judgement, inference, have here their source...For the origin and development of human faculty we must look to these processes of association in lower animals.

> Thorndike 1967/1898

Thorndike summarized his work in the famous "Law of Effect", which can be seen as the foundation stone as behaviourism. Although he claimed that this was a generalization from empirical evidence, Boakes (1984) points out that many of the experiments for such a generalization were not to be done for many years after Thorndike's work; but that his law is a clear repetition of the associative laws of Spencer, Bain and James, whose pedigree goes back to Locke and Hume.

3.2 Evolutionary Epistemology and individual learning

We have now a theory of sorts. But does it mean anything? Can anything be predicted from such a vague and general principle? Let us go over the main points of evolutionary epistemology once more, and then consider these questions.

3.2.1 Recapitulation of Principles of Evolutionary Epistemology

Although evolutionary epistemology is encapsulated in just one principle – that all knowledge gain is by a process of BV&SR– it is easier conceptualized as three principles.

Firstly, that evolution is itself a learning process.

Secondly, that all learning processes operate by initially proposing a set of variants, from which survivors, representing the best fit with environment, are

selected.

Thirdly, therefore, that evolution in its dealings with the environment through BV&SR is learning in exactly the same sense as an individual learns, and that the individual's knowledge and ability to learn have not only evolved in primary zoological evolution, but also operate in the same manner as zoological evolution.

This third point is clearly critical to understanding individual learning within an evolutionary epistemology framework.

3.2.2 Levels of knowledge gain

The question of levels of knowledge gain is important, but I will not go into much detail here because the question has already been thoroughly discussed by Plotkin & Odling-Smee (1979; 1981a; 1981b). They argue that individual learning evolved as an information gaining 'sub-routine', out of the action of BV&SR at the primary level of the genome. This in itself is not novel (e.g Tinbergen 1951). What is original is that, within an E.E framework, emphasizing the organism's need to match its organization to environmental order, they are able to identify, in broad terms, the form of environmental changes that would require the evolution of such subroutines. Such environmental changes are "predictably unpredictable ... changes that are pegged within certain limits, but within these limits can vary randomly" (Plotkin & Odling-Smee 1979:31). Plotkin & Odling-Smee (1981a & b) go on to investigate the different levels at which the learning process operates. They consider three different sources of variants which form four levels of knowledge gain (called 'referents'); the genetic source; the phenotypic source as it proposes variants that are selected through variable epigenesis; the nervous system as it proposes behavioural and cognitive variants selected through individual learning; and the nervous system as it interacts at the cultural level. I am not sure that a listing like this is, or could be, exhaustive. The source of variation and selection for the learning of fine temporal control of movement would seem to be the cerebellum (Eccles 1977; which, by the way, is itself a theory of learning by selection). It may prove to be the case that every individual instance of learning has a unique source of variants (or unique combination of sources), and that a simple schema with a small number of levels may be untenable. The important lesson from Plotkin &

Odling-Smee's plan is the manner in which the higher levels (be they many or just a few) stand on the shoulders, so to speak, of the more basic levels. The pattern of individual behaviour we expect from evolutionary epistemology is not a clear demarcation between instinct on the one hand, and associative learning on the other, but a continuum of degrees of flexibility of variation and selection. At the simplest level, an organism may produce a single, inflexible response to a problem situation. The only source of variation here is the genetic variation between individuals; and the only selection that of natural selection. A second case covers an organism possessing a small restricted set of responses to a situation, from which the most effective is selected by feedback sensors within the organism. This could cover imprinting where following a variety of objects rapidly reduces to following one- and, as we will see below- bird song learning. At the most flexible extreme are organisms whose set of variants is very wide, and whose selection mechanisms are elaborate and keep track of past errors to inform future production of variants. We like to think of ourselves in the way, solving problems by hypothesizing a potentially unlimited set of variants form which we select the correct one by an elaborate empirical testing process. This is also the behaviourists' conception of the organism, able to associate any action with any consequence. Of course no organism actually works this way. BV&SR at higher levels is always constrained by more basic levels. We can't hypothesize 'anything' because, as James said 'features of our organic mental structure...[are] transmitted as fixed features of the race' (1981/1890:1212).

3.2.3 Variation & Selection as an empirical program

An empirical program is not just a set of predictions, it is also a way of viewing the world. The theory that all knowledge gain is by BV&SR may be wrong (to be self-consistent, it must acknowledge itself to be at most a step on the road to truth); it might be (though I don't think it is) tautological; but even if it were both these things I think it would still be a very useful way of viewing behaviour, because it emphasizes some interesting aspects of the way that organisms interact with their environment through behaviour, that extant theories, both psychological and ethological, have overlooked. Let us consider briefly the questions that an analysis of behaviour prompted by the evolutionary epistemology paradigm would have to address.

3.2.3.1 Questions of Variation and Selection

1. What is the nature of the variants here? Are they variants of simple action (form of movement, muscle groups utilised, temporal patterning of movement etc, etc); of object oriented towards (e.g in imprinting, early phase of shaping of operant behaviour etc, etc); hypothesis of cause (e.g in classical conditioning); hypothesis of effect (e.g in operant conditioning)?
2. What is the source of the variants? Are they genetic; morphological; or from the central nervous system? (or a combination?)
3. How are the variants generated? This will be related to the previous question, but in cases where the variants are produced by a nervous system, this question will direct research to the effects of prior experience on the production of new variants.
4. What is the locus of selection? Natural selection; cognitive selection; developmental selection?
5. What is the nature of the selection process? Does the selection mechanism keep any records; are failures recorded; is the context of the failure recorded? How does the feedback influence the next stage of variation?

This last point is the one where the most striking differences are to be expected. As Lorenz (1977:24) points out, primary genetic learning differs from (some) individual learning because the genome only learns from its 'success'. In so far as a learning system keeps a record of its errors it will be able to learn faster, and cope better with new information. A BV&SR model becomes similar to an associative model if a learner keeps a record of each variant paired with the consequence of its use. Associationism, then, is a limiting case of BV&SR, but evolutionary epistemology recognizes a wider range of ways of proceeding, and does not for a moment suggest that associationism is a general process. All a learner might associate are variants and the effect of selection on them, selection mechanisms could employ other forms of 'record keeping' than association, and much may be learnt without any record keeping at all.

Mike Harris (in preparation) has developed a computer simulation that can learn the relative ordering of a set of items presented two at a time. On any trial one item only is rewarded, such that a consistent pattern of which item is better than which others may be deduced. Initially the model possess no hypotheses, and it chooses items simply by guessing. Harris' simulation learns the series by initially guessing which item to take. From the results of a guess it generates a hypothesis about the item taken, of the form 'take x', or 'avoid x', depending on whether taking the item was rewarded or not. Hypotheses become ordered in a stack so that the first successful hypothesis gets tested first on each trial, the second hypothesis to be generated is tested in second place, and so on. The learner searches the stack from the top and applies the

first rule that relates to one of the items presented on the trial. If no rule applies it guesses as it did when it had no rules in the stack. If a rule fails, the newest rule to have been added in the bottom of the stack gets thrown out, and a new rule relating to the object just taken, gets added onto the bottom of the stack. Such a model keeps no record of its failures. If a rule is always incorrect it will sooner or later fall off the bottom of the stack, but this in no way stops it reappearing again if other rules fail. The variants here are of the form of hypotheses as to whether to take or leave objects. Their source is the computer's 'central nervous system'. This is also their sight of storage. The generation rule for variants is that they are deduced from the consequence of taking an time (if **x** is rewarded then generate rule 'take **x**'; if x is not rewarded then generate rule 'avoid **x**'). The selection rule for variants is 'If rule fails jettison last rule from bottom of stack'.

The point I want to draw from this model is that the form of learning it achieves is considered quite complex (transitive inference learning, as it is known, is considered unique to primates). Nevertheless Harris' model learns the ordering of items implied by the repeated pair wise comparisons, by a simple variation and selection rule, without having recourse to associations at any stage.

3.2.3.2 Belongingness and Prepardness

With respect to associationism two points should now be clear. Firstly association is not necessary for learning. Although we may notice organisms learning 'associations' (that is to say, relations between causes and effects), this may often proceed without the need for association on the part of the subject. Secondly association, where it does occur, will be just one of a number of ways that a learner could log the fates of its variants, and will always necessarily be constrained by the set of variants generated. Since the set of variants produced in response to a problem situation will itself be a product of BV&SR in the past, qualities of 'belongingness' and differential 'prepardness' to associate certain causes and consequences are inevitable. We can predict that an organism can only learn the relationship between a cue and its consequence (classical conditioning) if the cue enters the set of cues the subject 'hypothesizes' as the cause of the event. The set of causes a subject will hypothesize for an event will be those cues that have in the past (including evolutionary past), been good predictors of events of the type under

consideration (where the event has a significance in the subjects natural habitat for its inclusive fitness). For a very flexible organism with an advanced nervous system those qualities of cue that make it a good predictor may be solely formal ones (qualities leading to Bayesian inference; Staddon 1986). For an organism adapted to a less varied environment the set of suitable cues may be much more restricted. This prediction is open to empirical test.

In an operant learning context we can predict that an organism will learn to link its behaviour with consequence if that consequence is a member of the set of consequences that such behaviour could have generated in evolutionary history in the natural habitat. Again, if the organism has a very flexible nervous system the set of consequences for any one behaviour may be very large. A comparison of organisms from different habitats would be predicted to find that only those possessing flexibility of response in the natural ecosystem would be able to form wide ranging associations in the operant laboratory. A proviso has to be made here to acknowledge that some organisms may have redundant flexibly of responsiveness. I am not arguing here that a subject would *never* learn to associate production of food (say), with a behaviour not normally related to food acquisition in its natural environment. Again, as for classical conditioning, the temporal relationship of an action and environmental effect may be enough of a cue to ensure the relationship is learned.

3.2.3.3 Ecological Studies

BV&SR provides a focus for interrelating psychological studies with field work. As mentioned above, predictions from evolutionary epistemology depend on the behaviour observed in the natural environment. If a subject needs to recognize a variety of cues for a certain consequence in its daily life, then we expect it to demonstrate that in laboratory conditions. To that extent evolutionary epistemology coincides with the ecological approach to animal behaviour. The BV&SR approach to behaviour does not suggest, however, that atheoretical observation of organisms in their natural habitats will ever lead to a science of behaviour. It would be internally inconsistent if it did.

3.2.3.4 Definition of learning

An intentional oversight in the above account has been an avoidance of defining learning. Revusky (1985) has reviewed the problems that attach to

customary definitions of this process. I propose that BV&SR, as an all encompassing description of learning, be used as a definition. Such a definition would exclude the passive modifications with experience that are difficult to leave out with other definitions. It would bring in evolution as a learning process, and also other knowledge gaining processes, like the immune system, that are typically left out.

3.2.4 Theories of Variation and Selection

Rather than attempt here to develop a complete theory covering all extant instances of learning within a framework of BV&SR, I will briefly discuss the theoretical positions of a number of authors who show a strong affinity with BV&SR.

3.2.4.1 Krechevsky

Over fifty years ago Krechevsky (1932; 1967/1932) noticed that rats in discrimination maze learning tasks seemed to possess 'hypotheses' about the nature of the correct solution to the problem, and work through these until they found the right one. The task they were set involved walking along a maze alleyway which was punctuated by a number of choice points at which the rat could attempt to go left or right. Of the two alternative alleyways confronting the subject, one was a cul-de-sac, and the other, identified by a hurdle, was the through route onto the next choice point, and ultimately the goal box. Krechevsky found that many rats (36 out of 40) made repeated attempts at what appeared to be incorrect, but feasible, solutions to the task. A typical such 'hypothesis' was consistently taking the right fork at each choice point. Krechevsky (1967/1932) concluded that "the tendency to build up a systematic series of responses during the so-called period of chance is characteristic of most of the rats studied in this experiment".

Krechevsky attempted to incorporate his results into Tolman's schema of 'means-end readiness' (Tolman & Krechevsky 1933). Levine (1967/1959) later incorporated Krechevsky's ideas into Harlow's learning-set theory. Unfortunately nobody thought to consider these experiments as paradigmatic in their own right. Krechevsky's observation, of a set of hypotheses getting selected down to one successful strategy, is exactly the kind of experiment evolutionary epistemology would treat as a paradigmatic case of learning.

3.2.4.2 Staddon

Staddon (1975) has noted that when first placed in a Skinner box, and particularly after first delivery of reward therein, a subject's behaviour will be highly variable. This post-reinforcement variability remains even when the subject is highly experienced in the experimental chamber. Activities closer in time to the reward (assuming reward delivery is temporally regular), become much more stereotyped. Staddon argues that a biologically meaningful stimulus creates a range of behaviours. Increasing experience of the environment causes the subject to select from amongst its behaviours those that most efficiently fulfil the response contingency laid down by the experimenter. This selection of behaviours includes not just what to do, but also when to do it. In this way, at times when food is likely, relevant behaviours are concentrated, but when food is unlikely a greater variety of behaviours are engaged in to keep the animal responsive to its environment.

3.2.4.3 Angermeier

A model possessing many points of contact with Staddon's (1975), has been developed by Angermeier (1984). Angermeier considers the simplest operant experiment, where a subject learns to operate a manipulandum to gain a reward. Angermeier and his co-workers studied a number of species and noted the wide range of behaviours initially displayed, and the process of selection of the relevant, rewarded, behaviour from the initially variable set of possibilities. A parametric model, capable of accounting for the rate of decay of irrelevant behaviour, was developed.

3.2.4.4 Marler

A source of optimism for evolutionary epistemology is that, as well as being able to account for associative learning, the theoretical account put forward by Marler (1984) for bird song learning, is also compatible with BV&SR. Marler envisages an interaction between auditory templates and the auditory input to a juvenile song bird. As I understand it, these innate templates guide the selective attention and selective memorization of the auditory input by the bird. This selective learning guides the shape of the bird's own song.

3.2.4.5 Physiology and Anatomy

I am getting even deeper out of my depth here, but I am excited to find that a number of authors in the neurosciences have emphasized the selective nature of learning and memory mechanisms. These include Young (1976; 1979), Changeux, Heidmann & Patte (1984) and Mountcastle & Edelman (1978). These models all share the property that it is selection, from the myriads of possible interconnections between neurons, that constitutes learning. Variation does not seem to be mentioned explicitly, but is implicit in the multitude of possible connections potentially existing in a nervous system. Mountcastle & Edelman's notion of 'microcolumns' of neurons may well fulfil the requirement of 'reverberatory circuits of neurons' that Pringle (1951) postulated would generate variation and selection in behaviour.

3.3 Conclusion

In this chapter and the preceding one I have tried to set the ground for, and very briefly outline, a theory which I believe has the potential to be a complete theory of animal behaviour. This exposition has been blatantly incomplete, and further, the experiments that follow were much to slow for analysis of variation and selection of behaviour to be a reasonable proposition. None the less, the stakes are high enough that I feel the undertaking was worth attempting.

CHAPTER 4
FLOOR PANEL EXPERIMENT:
METHOD AND MOLAR RESULTS

4.1 Introduction

This experiment has three aims. Firstly I test the predictions of a particularly rigorous optimal foraging model (Houston & McNamara 1981). Secondly I hoped that it might prove possible to develop a conceptually simple two patch foraging environment which would lend itself readily to being employed by a variety of species and subjects at different developmental phases. In addition to providing opportunities for future ontogenetic and phylogenetic research it was considered important that the paradigm should possess as much external validity as could be made compatible with computerized control and consequent detailed event recording.

The first reports of the use of continuous choice procedures in operant chambers (concurrent scheduling) appeared in the late 1950s (Ferster & Skinner 1957; Findley 1958). Although it is not easy to claim that this form of experiment is an obvious analogue to any form of foraging normally performed by columbine birds, it would be short-sighted to ignore the evolution of the operant chamber as experimental environment that had occurred since, at least, Skinner (1938). A standard dual key experiment will be discussed in chapters 6 to 9; the present experiment utilised a floor panel that could rest in two positions to collect reinforcements that had been set up by two independent variable interval schedules.

A variety of simple experimental set ups have been developed in ethology in the last decade, mainly to test the predictions of optimal foraging theory – though 'rules of thumb' that could be governing the observed behaviour have also been considered. Interestingly many of these have given subjects choices between sources of food which payed out with different probabilities; this form of choice paradigm had a period of popularity in psychology, but has largely fallen from favour in the last ten years (see e.g. Bitterman 1965; Warren 1974; Mackintosh 1974 for reviews; and Couvillon & Bitterman 1985 for a recent study).

Some ethological studies have, as one would expect, possessed greater face validity as foraging environments for bird subjects than the little boxes used by psychologists. Krebs & McCleery (1984) and Kamil and Roitblat (1985) are recent reviews of this field.

Krebs, Ryan & Charnov (1974) initiated the use, in an aviary, of a number of small 'trees' (made of wood), each containing 'pine cones' (small blocks of wood drilled with holes), filled with either a small (1) or high (3) number of meal worms. Chickadees (*Parus atricapillus*) were used individually in this apparatus to test Gibb's (1962) hunting by expectation hypothesis, and the predictions of the marginal value theorem of optimal foraging (Krebs et al 1974; and see Charnov 1976). Hunting by expectation predicted that because there were many more 1-prey than 3-prey patches, the predators would develop the expectation of finding only 1 prey item in each patch, and therefore would not bother to look for and remove the other two meal worms in the richer 'pine cones'. From the marginal value theorem it was predicted that the predator should leave a patch when its rate of food intake from this patch fell below the average level for the environment as a whole. Since all prey were of similar handling time this led to the prediction that there should be a constant giving up time for each patch. Contrary to Gibb's hypothesis, but broadly in line with the marginal value theorem, the Chickadees' time spent in searching each patch was independent of the number of prey in any particular patch, but was constant for all patch types within an environment, and inversely related to the average capture rate for the environment.

A similar set up was subsequently used by Cowie (1977) on Great Tits (*Parus major*) to test optimal foraging with different 'travel times' (simulated by putting lids on the 'pine cones'). It was predicted that with longer travel times the subjects would spend longer in each patch so that the patch was still depleted to the average rate available within the environment as a whole. Increasing the 'travel times' effectively reduces the average rate of prey intake available in the environment. Cowie found that when the 'travel time' requirements were longer, the subjects spent longer in each patch as predicted, but that he could only get a good quantitative fit to his data by adding in the relative metabolic costs of flying and feeding as post hoc parameters to the model.

Zach & Falls (1976) found results that contradict Krebs et al (1977) and

Cowie (1977). In their study oven birds (*Aves: Paralidae*) foraged over a patch containing many holes, some of which hid freeze-killed flies. No evidence was found for hunting by expectation, but neither was there any sign that giving up time depended on overall prey density, as would be predicted from the marginal value theorem.

Other studies have used simpler two choice environments more akin to typical operant studies. Krebs, Erichson, Webber & Charnov (1977) presented Great Tits with prey of two different values passing by them on a conveyor belt. With this arrangement it was possible to alter the encounter rate with each prey type and compare the subjects' choice with optimality predictions. Krebs et al (1977) found that choice of prey depended on the encounter rate with large prey as predicted by optimal foraging, but that the change from a mixed distribution to an all or none pattern of choice was not a step function of encounter rate as had been predicted.

A paradigm intermediate between using the Krebs et al (1977) conveyor belt and the use of 'trees' in an aviary is the employment of two food hoppers in an aviary (Krebs, Kacelnik & Taylor 1978; Kacelnik 1979; Kacelnik & Krebs 1985; Kacelnik, Krebs & Ens 1986 in press). This has been used as a probability choice experiment to test optimality predictions and to investigate the balance between sampling of an environment and exploitation thereof. This paradigm is discussed in chapter 8; section 8.3.

Larkin (1981) tested Barbary Doves (*Streptopelia risoria*) in a variety of choice experiments, one of which was very similar to the experiment described in chapter 6. Larkin's birds foraged on two depleting schedules in a Skinner box. The longer that was spent responding to a schedule the less it paid out. When a bird returned to a schedule after a period of time on the other one, the newly chosen alternative would reset to its highest rate of reinforcement. Larkin found that when the travel time requirement (imposed by a partition in the chamber) was large, the birds spent less time in each visit to each patch than was optimal; and when the travel time was short, the birds spent a longer than optimal time in each patch. These results were found for both food and water reinforcement.

An entirely different approach to testing accounts of foraging in the laboratory, has been adopted by those researchers who believe that Skinner

box scheduling is substantially comparable to 'real world' foraging. On this assumption elaborate multiple concurrent schedules of reinforcement have been constructed, containing a number of phases, each with its own completion requirement and named after some component of natural foraging. The first experiment of this type was performed by Lea (1979). A pigeon (*Columba livia*) in a Skinner box was confronted with two keys. Initially only one of them was illuminated (white), but after completion of a fixed interval (FI) requirement on this key (named 'search phase') the other key would come on. The probability was p that this key would be red, $1-p$ that it was green. During the next phase, called the choice state, the original white key was still illuminated, but the new, coloured key was also lit. The subject had to choose whether to peck either the white key and return to the search phase again, or respond to the coloured key and proceed to the 'handling' state. The handling phase was another FI schedule, which differed in magnitude depending on the colour of key light that had preceded it. At the end of the handling phase a reward phase was instituted during which the subjects had access to grain.

This type of procedure enables the systematic variation of each particular component of the foraging cycle. By varying the duration of the FI search requirement, with varying proportions of good (short handling requirement) and less good (long handling requirement) prey, Lea (1979) was able to show qualitative agreement with the marginal value theorem. Choice of prey types depended on the density of better prey only. Following Krebs et al (1977), however this change in preference for prey requiring a longer handling time, was not found to be a step function.

This procedure has been modified and extended by Abarca and Fantino (1982). In the handling phase they replaced Lea's FIs with variable interval (VI) schedules, and following Lea (1979), varied the length of the FI requirement in the search state. Their results paralleled Lea's (1979). Fantino & Abarca (1985) review the results of a variety of modifications of Lea's (1979) procedure performed in Fantino's laboratory. These include the effects of changes in search and handling times and different reward quantities (Ito & Fantino 1983 cited in Fantino & Abarca 1985), and comparing performance of subjects maintained on closed and open economies (Fantino & Abarca 1985).

Lea's (1979) concurrent chains procedure has also served as the basis for research by Collier (1983; Collier & Rovee-Collier 1981), who used ratio instead

of interval schedules for the various components. Peden & Rohe (1984) did likewise, having first recorded the baseline levels of responding and eating in their subjects, to test whether they minimized number of pecks per food delivery and maintained their baseline intake of food. Although baselines of eating were maintained during the experiment, this was not done with a minimum expenditure of effort.

The major problem with studies of this type, concerns the meaningfulness of assigning names, taken from the foraging literature, to components of schedules operating on two keys in a Skinner box. Shettleworth (1985a) asks whether "signalled delays to reinforcement are psychologically equivalent to handling times, and thus whether performance in a situation where handling times are simulated by signalled delays differs in any important way from choice among real items requiring real handling (husking, killing, chewing, etc)." This problem seems to have been largely overlooked by the advocates of these schedules. The simplest and most common form of discontinuity built into a typical concurrent schedule experiment is the changeover delay (COD). This is usually referred to as a 'time out', because nothing is programmed to occur during it. There is plenty of evidence, however, that although we may call this phase a 'time out', subjects certainly do not take 'time out' from responding during it. In fact reports that have studied responding during COD in detail, have found higher rates during COD than when the schedules are operative (Silberberg & Fantino 1970; Pliskoff 1971; Scoun 1983). If this is true of a relatively simple schedule manipulation, what grounds are there to accept that these more elaborate structurings of the options confronting a subject are perceived by it in the manner the experimenter intends? In other words, the fact that we as experimenters recognize a formal equivalence between, for example, a delay before reinforcement presentation, and the time a bird spends handling a food item before eating it in the wild, does not imply that a functional equivalence exists between the two states, nor therefore that the responses they produce from a hungry subject will be equivalent.

For present purposes recourse to solely formal equivalences between the world outside the laboratory and the experiment will be avoided, so far as possible, subject to the aims already stated above. No changeover delay, for example, is here programmed.

The experimental chamber to be described here is similar to that used by

Baum & Rachlin (1969). The floor of that chamber was roughly twice as long as it was deep, such that a pivot point across the middle split it into two equally-sized segments that would tip (.5cm) when the subject stood on them. At each end of the chamber were hoppers that were controlled by two independent VI schedules that ran continuously, but would only pay out when the subject depressed the respective side of the chamber (concurrent variable interval; VI:VI). The major difference between Baum and Rachlin's (1969) design and the present experiment is that here a single, centrally placed food hopper is employed. In this respect it is more akin to Wheatley & Engberg (1978). They used a similar treadle manipulandum which could be moved around in the chamber by the experimenter so that two treadles could be presented, or one treadle and one key. Results from both Baum & Rachlin's and Wheatley & Engberg's studies, however, were only assessed in terms of the ratios of time spent on each option compared to the ratios of reinforcements from each option. Both sets of results confirmed to the prediction of Herrnstein's matching law (Herrnstein 1961; 1974; Baum 1974) in that these ratios were approximately equal. Matching is discussed in more detail in chapter 6.

Shettleworth (1985b) utilised a shuttlebox very similar to that of Baum & Rachlin (1969) to test pigeon's choice between a few large food items and many small items of the same total weight (and nutritional value). Initially concurrent VI schedules were used, as will be here; in a subsequent experiment Shettleworth employed variable ratio (VR) schedules. In both cases she found that the birds preferred many small items over a few large ones, even when schedules were programmed such that choosing the items with the longer handling time (literally- many small items) substantially reduced the rate of food intake.

A matching analysis will be presented below, solely to verify the compatibility of these data with previous studies. My real interest lies in the opportunity to test a rigorous optimality model which makes very clear quantitative (rather than just qualitative) predictions (Houston & McNamara 1981). I had intended to also investigate the 'rules of thumb' that the subject might be using, but from the results of the optimality analysis it will become clear that such an analysis would be rather pointless. The experiment described in chapters 6 to 9, though more of a traditional psychological paradigm, and less amenable to optimality analysis, in fact offers better opportunities for the testing of rules of thumb.

4.2 Method

4.2.1 Subjects

Six Barbary Doves (*Streptopelia risoria*) from the departmental colony started the experiment. They were all over 18 months old, were caged singly with water and grit continuously available; lights were kept on 14 hours a day. Prior to the experiment they were maintained on a diet of 50% wheat, 25% millet and 25% cut maize, and the same mixture of grains served as reinforcement and additional feeding where needed after testing to maintain their weights at 90% +/- 5% of free feeding levels. All had been hopper trained 9 months previously. Two (Birds 88 and 96) had some experience of concurrent key peck interval schedules in a Skinner box (as described in chapter 6) but had been removed after they developed a manner of using the food hopper that led to them damaging their wings.

4.2.2 Apparatus
Figure 4.2 (p58)
shows the experimental chamber. It was a locally constructed box made of a steel tubular frame with a chipboard base and ceiling, and sides of wire mesh covered by paper. The floor was 35 cm square, and the roof was 60 cm above the floor. The floor itself was raised 8 cm above the base to accommodate the pivot and microswitches. The 5 cm square hopper opening was situated 9cm above the floor in the centre of the back wall (i.e. above the pivot point). A mass of 140g (the mean 90% free feeding weight of the birds in this experiment), at a distance of 10 cm from the pivot point was sufficient to depress the floor panel the 2 cm required to activate the microswitch. All panel and hopper movements were recorded, to the nearest .01sec, by a Acorn BBC Model B micro-computer, for subsequent transfer to the university's ICL 2988 computer for analysis. Illumination of the experimental chamber was provided partially by the lighting in the laboratory and also by four 24V bulbs within the chamber itself. A white light in the centre of the ceiling was illuminated at all times; a red light came on in the left centre of the ceiling when the left side of the floor panel was depressed; and likewise a red light illuminated when the right side of the floor panel was depressed. During reinforcement the hopper opening was lit by a white bulb. White noise to mask extraneous sounds was broadcast around the experimental chambers during the experiment.

4.2.3 Procedure

Without additional pre-training subjects were entered onto the first condition of a concurrent variable interval schedule. Timers for each side of the panel ran independently of each other, and of the bird's position throughout the experiment, but reinforcement could only be collected from a schedule when the side of the floor panel associated with the schedule was depressed. The schedules were designed such that the minimum interval between reinforcements followed a negative exponential distribution. A set of 1000 intervals was constructed according to the approximation of Flescher & Hoffman (1962). After reinforcement, an interval was selected at random from this set, and the next reinforcement to the same side was then delivered once this interval had timed out, and the respective side of the floor panel was depressed. No COD was programmed, the only delay in changing over was the time required by the subject to swing the floor panel across from one microswitch to the other. Reinforcement was 3.5sec access to mixed grain. Sessions of approximately 45 min each were run six days a week, approximately 90 mins after lights on. (Programming problems made it difficult to ensure the duration was always exactly as intended). Supplementary feeding, if necessary, was given immediately after the experiment. Each subject was run 6 days a week until a stability criterion was fulfilled. A new criterion was used in this experiment to overcome three problems.

Firstly that a criterion should stipulate a condition for consecutive days of testing (Davison 1972, for example, does not) - it is not clear in what sense a criterion can be considered to be testing stability of choice if this condition is not fulfilled. Secondly the response ratios were logged (to base 10) to normalise them. Ratios were used because it is from them that the matching equation is calculated. Thirdly, spectral analysis of early results suggested a weekly cycling, so the criterion was calculated on 6 day groups to control for this weekly effect. The criterion was that the coefficient of variation (mean/ standard deviation) of the logged (base 10) response ratios should exceed 5 (or in cases where the log mean was small, that the s.d. be less than .10) on six consecutive occasions. The value of 5 was chosen after analysing some early data to maintain comparability with other methods of calculating stability (e.g. Davison 1972)

Bird 88 was removed after 20 sessions when it damaged its wing. Birds

105 and 269 were removed from the experiment after failing to fulfil the stability criterion after 48 and 46 sessions respectively. Three conditions were run, concurrent VI 120s: 120s; VI 240s: 30s; VI 60s: 180s (in the rest of the thesis the 's' for seconds will be omitted when identifying schedule values. Conditions and numbers of sessions for the three remaining subjects (Birds 92, 96 and 120) are shown in table 4.1.

Table 4.1: Conditions and numbers of session for floor panel experiment

Subject	Condition- VI:VI	120:120	240:30	60:180
92		26	16	28
96		24	32	33
120		29	16	14

4.3 Results

Table 4.2 shows in seconds the length of time each side of the panel was depressed (T_n, n=1..2); the number of times each side of the panel was depressed (N_n); the time spent without either side of the panel depressed (CO time- X); and the number of reinforcements from each side, for the last five sessions of each condition. In cases where the actual duration of the experiment differed by more than 5s from that intended (2700s), the actual session duration is given in brackets.

Table 4.2: Number and duration of visits to each side of floor panel

Subject	Condition	T_1	N_1	T_2	N_2	X	Rf_1	Rf_2
92	120:120	109	222	554	597	1920	7	22
(duration)		76	246	425	519	2083	12	17
		68	248	418	465	2096	13	16
		33	240	533	455	2017	14	15
		30	182	276	387	2278	11	18
92	240:30	20	217	810	655	1638	8	50
		44	200	804	702	1615	1	58
		12	118	816	550	1651	4	51
		22	168	756	676	1754	5	57
		3	105	557	537	1936	6	50
92	60:180	268	386	612	210	1704	24	5
		54	567	186	224	2308	33	5
		33	395	168	170	2347	32	6
		171	613	277	247	2086	36	5
		135	346	406	201	1999	35	5
96	120:120	27	97	741	334	1816	15	14
		66	97	805	228	1760	7	20
		15	174	716	343	1813	18	21
		13	85	273	159	2306	11	16
		62	113	860	273	1661	11	18
96	240:30	142	141	1199	322	1119	2	58
		92	138	921	262	1499	1	46
		45	112	967	286	1348	10	75
		45	112	950	272	1509	5	44
		7	42	616	128	1933	5	31
96	60:180	708	416	99	164	1777	25	4
(3325)		695	492	108	196	2406	25	4
		701	583	62	189	1769	32	10
(3291)		1009	624	113	209	1997	39	4
		555	278	15	79	2014	21	8
120(2514)	120:120	535	864	638	747	1185	18	21
(2421)		495	850	591	760	1175	19	21
		554	872	548	722	1434	21	20
		590	855	648	584	1350	11	17
		460	779	536	365	1560	20	16
120	240:30	33	256	729	698	1638	10	66
		37	249	662	679	1650	11	78
		76	282	745	730	1575	6	70
		134	522	674	809	1588	4	72
		72	576	652	982	1624	11	77

55

Table 4.2: continued

| Subject	Condition	T_1	N_1	T_2	N_2	X	Rf_1	Rf_2
120(1973)	60:180	365	911	306	695	1126	33	11
(2236)		495	870	440	741	1145	30	9
(1757)		364	831	387	799	926	13	7
(1842)		258	785	415	828	1021	30	7
(1907)		386	868	446	759	983	21	2

Figure 4.1 shows the log ratio of time spent on each side of the chamber, against the log ratio of reinforcements obtained from each side of the chamber. Table 4.3 shows the parameters of the best fitting (matching) lines to these data. (These are slightly non-standard in format, because each of the last five sessions is treated independently, compared to the normal practice of averaging over sessions before plotting).

Table 4.3:
Parameters of Matching Regressions

Subject	Intercept (log b)	Gradient (c)	Tail Prob	r^2
92	-.916	.557	< .1%	.629
96	-.212	1.021	< .2%	.574
120	-.295	.606	< .1%	.655
Group	-.485	.746	< .1%	.484

Matching is discussed in chapter 6; for now I note the following points. With the possible exception of subject 96 these values of the bias (b) and sensitivity (c) parameters (Baum 1974), are within the range to be expected in concurrent operant experiments (see Baum 1979; Wearden & Burgess 1982; Taylor & Davison 1983 for empirical reviews). If there is any departure here from standard linear matching, this is probably due to treating each of the last five sessions as a separate data point.

One consideration that is not apparent from these graphs is that the number of visits to the more preferred side of the floor panel is often double that to the less preferred side, implying a large number of 'dithers', i.e. visits to the centre zone followed by returns to the same side of floor panel. It can be seen that the subjects spend approximately half of each experimental session in the centre of the experimental chamber (not including time required to collect reinforcement). Considering that reinforcements can only be obtained whilst one or other side panel is depressed, this represents a substantial

figure 4.1   MATCHING LINES FOR PANEL DATA   56
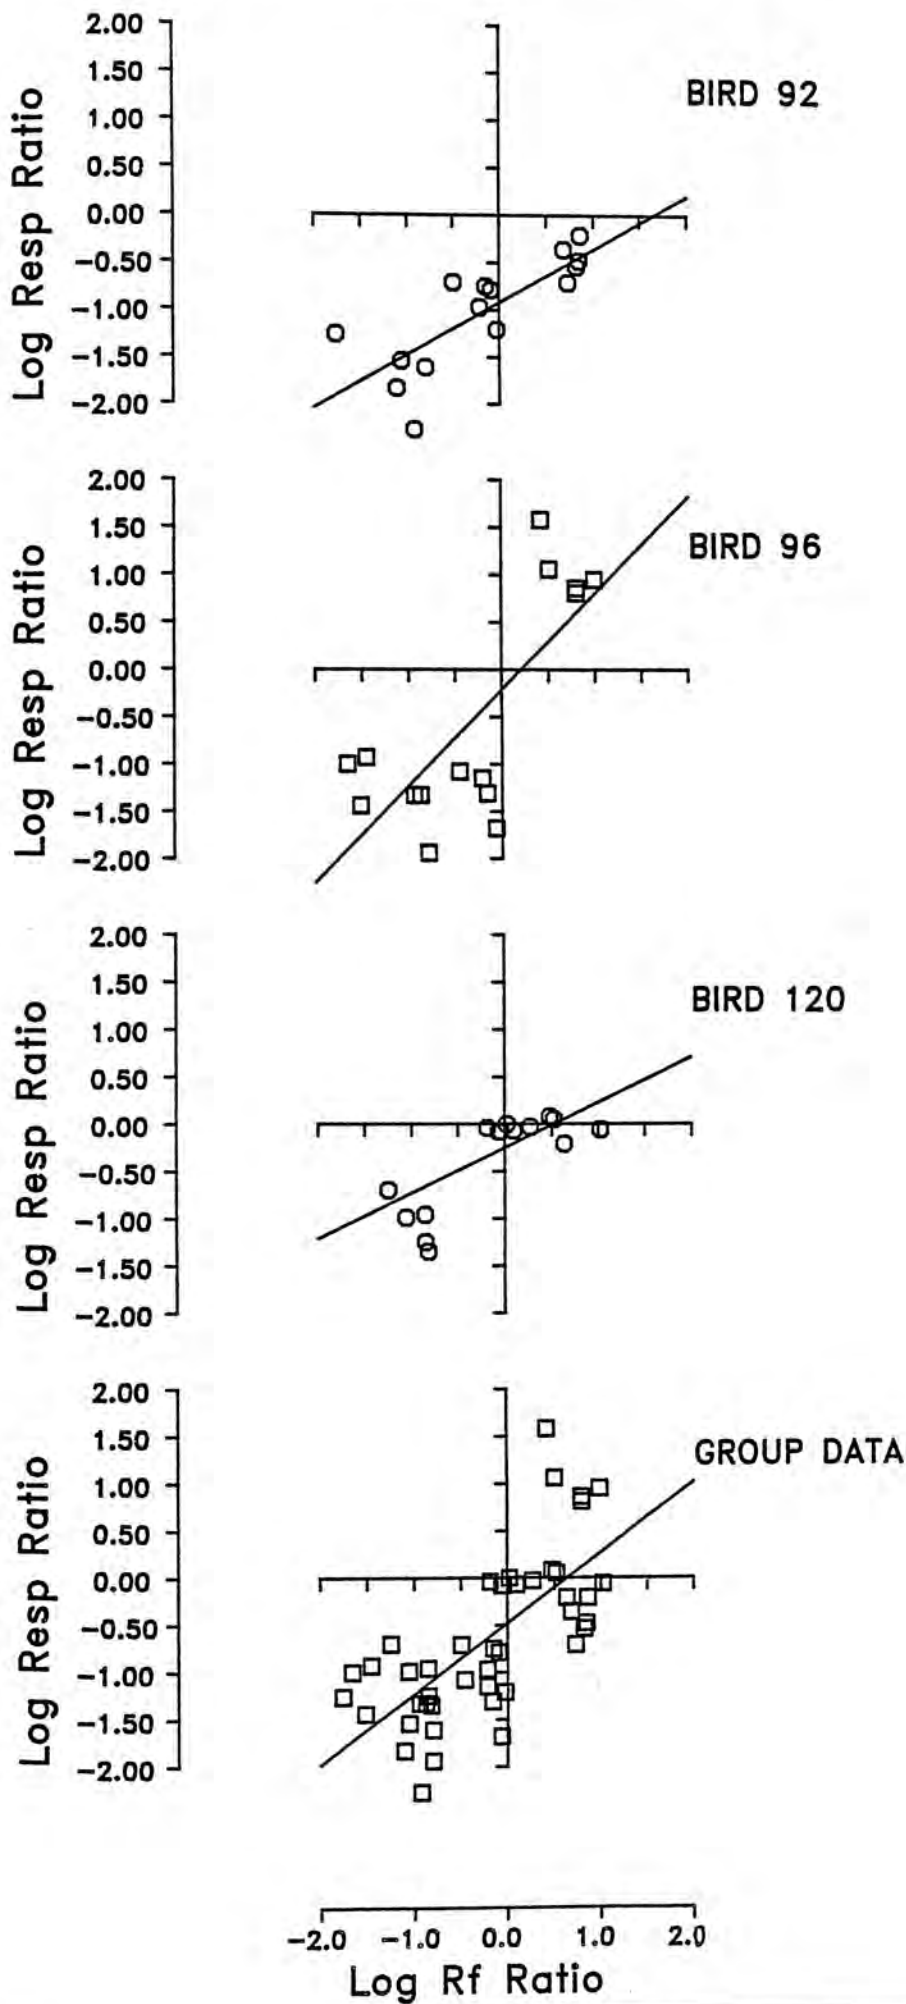
inefficiency. This supports the impression gained from informal observation of subjects whilst in the experiment, that they tended to base themselves in the centre of the chamber - near the hopper, and rather gingerly depress one or other side of the floor panel briefly, before returning to the centre to check the hopper. At this point it seemed, very informally, that they were just as likely to return to the side of the panel they had previously depressed, as to continue on over to the other side of the chamber.

Mean times in each area of the experimental chamber will be investigated in more detail in the next chapter, and discussion of these results will be held in abeyance until those results have been presented.

Figure 4.2

hopper
opening

pivot
point

micro-
switch
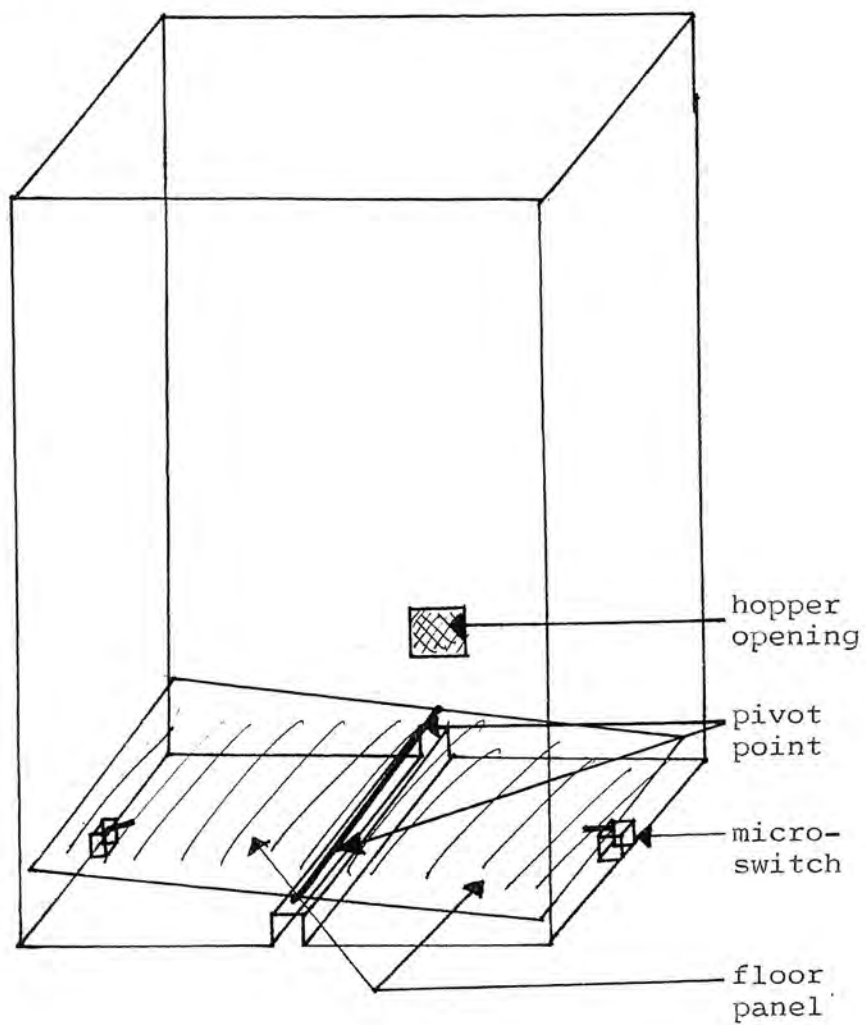
floor
panel

CHAPTER 5
FLOOR PANEL EXPERIMENT: OPTIMALITY ANALYSIS

5.1 Introduction and Method

As outlined in chapter 4, part of the rationale behind the present experiment was that a very thorough optimality analysis has been performed for it that makes predictions that are not only quantitative, but also precise enough to be unlikely from any other account. This optimality analysis (Houston & McNamara 1981) is of the unconstrained form advocated, as a preliminary to further analysis, in chapter 2. There have been previous attempts to deduce the optimal response pattern on concurrent variable interval (VI) schedules, but these have concentrated on key pecking experiments and are therefore discussed in chapter 7.

Houston and McNamara's analysis considers the allocation of time between the two alternatives that will maximise reward rate. It ignores, therefore the relative costs of different behaviours within the experiment and the contribution to inclusive fitness from activities servicing other needs. For deprived animals in brief experiments of this type this assumption is probably reasonable. The account below follows Houston & McNamara (1981) except where stated otherwise. References given are additional sources.

On a schedule constructed from inter-reinforcement intervals forming a negative exponential distribution (as used here), the probability of a reward not having set up after some time, t in which the schedule has been left running unattended is $e^{-\lambda t}$, where λ is the mean rate of reward on that schedule (= 1/ programmed mean interval between rewards). Therefore the probability that a reward has set up after a time t is $1-e^{-\lambda t}$ (Flescher & Hoffman 1962; Cox & Lewis 1966; Staddon, Hinson & Kram 1981). This is true for both schedules at once because the timers for the two schedules are independent and run concurrently. When a reward is due the timer for that schedule stops, and the hopper will be activated immediately, if the bird is on the correct side of the floor panel, or as soon as it next visits that side of the chamber.

When the subject switches from one side of the panel to the other it necessarily experiences a delay of duration τ seconds during which it cannot

collect rewards from either schedule. We are assuming that the subject knows the values of the mean rates of payout on both schedules, and the value of T. This is another reason why in this study the change over delay was simply the time the subject needed to traverse the central area of the chamber, and a rigorous stability criterion was enforced.

Houston & McNamara consider a cycle of behaviour in the experiment starting on schedule 1, moving directly to schedule 2, (thereby experiencing a delay (T)), staying on schedule 2 for a period a_2, then returning to schedule 1, again experiencing the delay T, and finally staying on schedule 1 for time a_1, before the cycle starts again. They do not consider the case of a bird entering a changeover phase and then returning out of it back to the same schedule, but since by definition the changeover delay is a period when reinforcement is impossible, it is obvious that this can only reduce net reward rate.

Firstly Houston & McNamara demonstrate (1981 appendix 1) that the reward rate for variable values of a_1 and a_2 will always be bettered by fixed values of these parameters. This then leaves two possibilities, either never switch (i.e. stay all the time on the better option); or stay for fixed times a_1 and a_2 on the two schedules.

Returning to the cycle described above, it is clear that during the times spent on each of schedules 1 and 2 the reward to be expected is simply the reward rate of the schedule multiplied by the time spent on it: $\lambda_1 a_1$ and $\lambda_2 a_2$ respectively. The probability of no reward having set up at the point of returning to schedule 2 after a period away from it of $a_1 + 2T$ is $e^{-\lambda_2 (a_1 + 2T)}$ so the probability that a reward has set up is $1 - e^{-\lambda_2 (a_1 + 2T)}$. Similarly at the end of the cycle, when the subject returns to the first schedule, the probability of a reward having set up is $1 - e^{-\lambda_1 (a_2 + 2T)}$. Over the total cycle of two stay times a_1 and a_2 and two changeovers ($2T$), the mean reward to be expected R is therefore given by:

Equation 5.1

$$(a_1+a_2+2T)R =$$

$$(1 - e^{-\lambda_1(a_2+2T)}) + (1 - e^{-\lambda_2(a_1+2T)}) + \lambda_1 a_1 + \lambda_2 a_2$$

At this point the mathematics can be simplified by taking the mean rate on the better schedule (say λ_1) to be the unit of time. The rate on this schedule therefore becomes 1, and λ_2 becomes equal to λ_1/λ_2 which we can call simply λ. T must be rescaled as a proportion of the mean interval between reinforcements on the better schedule.

To find the extreme values of R we must differentiate equation 5.1 with respect to a_1 and a_2 and put both equal to zero. After the above substitutions equation 5.1 becomes:

Equation 5.2

$$(a_1 + a_2 + 2T)R = 2 - e^{-\lambda(a_1+2T)} - e^{-(a_2+2T)} + a_1 + a_2\lambda$$

differentiating with respect to a_1 gives

Equation 5.3

$$(a_1 + a_2 + 2T)(\delta R/\delta a_1) + R = 1 + \lambda e^{-\lambda(a_1+2T)}$$

and for a_2

Equation 5.4

$$(a_1 + a_2 + 2T)(\delta R/\delta a_2) + R = \lambda + e^{-(a_2+2T)}$$

setting $\delta R/\delta a_1 = \delta R/\delta a_2 = 0$ in equations 5.3 and 5.4 gives:

Equation 5.5

$$R = 1 + \lambda e^{-\lambda(a_1+2T)} = \lambda + e^{-(a_2+2T)}$$

Houston and McNamara investigate the consequences of equation 5.5 in detail. When the schedule values are equal they show that the optimal stay time on each schedule, a^*, is approximated by,

Equation 5.6

$$a^* \simeq (2/3)T^2$$

for small T (keeping in mind that T is scaled in terms of the better schedule value) (Houston & McNamara 1981 appendix 2).

When the schedule rates are unequal it is fairly intuitive that the subject should spend longer on the better schedule. Houston & McNamara demonstrate this analytically in their appendix 3. They go on to show that for the majority of conditions the optimal policy involves staying for time a_1^* on the better schedule, and just switching ($a_2^* = 0$) to the other schedule instantaneously. Only if λ is greater than about .84 (decreasing with increasing T above $T= .92$), is it optimal to stay for a time greater than zero on schedule 2. At the other extreme, if the COD is very large, never switching to the second schedule will be optimal. They show that for this to be the case T has to be larger than .92 (increasing with increasing λ above $\lambda= .84$). In order therefore, for anything but 'Stay-switch' to be the optimal policy, the schedules either have to be very similar (e.g. VI 60: 61 with 4s COD) for stay-stay to be optimal; or possess an enormous changeover delay (e.g. VI 30:240 with 28s COD) in which case never switching is optimal. It is immediately clear, therefore that no experiments have ever been run (with unequal schedule rates) where anything but stay-switch would be the optimal policy. Unfortunately a_1^* for different values of T and λ cannot be calculated directly from any one equation. However Houston & McNamara (1981 table 6) provide a number of values, and also their figure 3 enables the estimation of values of a_1^* for three values of λ.

In conclusion, there are two types of prediction we can make, depending on the assumptions we are willing to entertain about the subjects' knowledge of the experimental conditions. If we assume only that they are attempting to maximise net reward rate, and that they are aware that this is a typical concurrent VI experiment, then we can predict they should:

1. Minimize the duration of each change over. This is not considered in the above analysis, but follows simply from the fact that the COD is a period during which reinforcement is impossible.
2. Constrain visits to constant durations. (Houston & McNamara 1981 appendix 1)
3. Minimize the duration of visits to the schedule with the lower rate.

If we are also willing to assume that they are aware of the value of T (which does not seem unreasonable given that it is a consequence of their own

behaviour); and that they know the value of λ (which may not be unreasonable when the stability criterion is fulfilled) we can additionally predict the exact duration of visits to the preferred schedule.

The rest of this chapter will test these predictions.

5.2 Results

Table 5.1 contains the means and standard deviations of the times spent by each subject in each phase of the experiment, for each of the last 5 sessions of each condition. The phases consist of: right side of the panel (microswitch 1) depressed (t_1); left side of panel (microswitch 2) depressed (t_2); neither microswitch depressed, where microswitch just released differs from microswitch next depressed (true changeover; **CO**); neither microswitch depressed, where microswitch just released is the same as the microswitch next to be depressed (dither; **di**).

Table 5.1: Mean and standard deviations of time in each component

Subj	Cond	t_1	sd_1	t_2	sd_2	CO	sd_{co}	di	sd_{di}
92	120:120	.49	1.51	.93	1.40	2.12	3.10	2.50	2.06
		.31	1.61	.82	1.32	2.48	2.73	2.95	3.57
		.28	1.37	.90	1.58	2.69	2.63	3.16	2.86
		.14	1.36	1.17	2.68	2.47	4.37	3.36	2.64
		.16	1.32	.71	1.39	3.52	3.65	4.24	4.38
	240:30	.22	.51	1.15	1.84	1.81	1.36	1.78	2.11
		.10	.78	1.48	2.03	1.82	3.32	2.65	8.20
		.13	.81	1.12	1.66	1.66	1.23	2.09	3.64
		.03	.96	1.04	1.73	2.46	2.40	3.09	4.39
		.09	.84	1.20	1.80	1.80	1.46	1.76	1.98
	60:180	.69	1.45	2.92	21.86	1.87	2.02	1.75	3.21
		.10	.98	.83	.95	3.87	9.58	2.43	4.36
		.08	1.21	.99	1.24	5.27	9.37	3.48	7.18
		.24	.98	.74	1.18	3.25	2.55	1.93	3.36
		.69	1.54	1.85	2.90	4.21	2.75	5.72	19.52
96	120:120	.28	1.57	2.22	2.03	3.65	5.88	4.73	5.76
		.68	2.17	3.53	9.46	3.23	7.78	3.18	5.79
		.09	1.40	2.09	2.59	3.57	5.94	3.55	3.12
		.16	1.85	1.72	2.29	9.17	19.93	9.73	14.91
		.54	4.62	3.15	5.28	3.70	3.56	4.74	4.85
	240:30	1.01	.95	3.73	4.90	2.04	2.02	2.68	2.93
		.67	.73	3.52	4.79	1.42	1.31	2.03	2.41
		.40	1.33	3.38	5.03	2.61	2.26	3.22	3.05
		.40	1.12	3.49	4.81	2.01	1.86	3.54	15.52
		.18	1.62	4.81	5.92	4.54	12.49	16.32	96.82
	60:180	1.70	2.76	.60	.82	1.47	1.53	2.23	3.41
		1.41	2.12	.55	.73	1.56	1.51	1.84	2.07
		1.20	1.87	.33	.99	1.62	1.67	1.97	2.81
		1.62	2.49	.54	.71	1.64	2.06	1.73	2.25
		2.00	2.87	.19	1.28	7.84	65.59	3.74	12.19

Table 5.1: continued

Subj	Cond	t_1	sd_1	t_2	sd_2	CO	sd_{co}	di	sd_{di}
120	120:120	.62	.88	.85	1.08	.95	1.18	.42	1.18
		.58	.92	.78	1.04	.89	1.02	.51	1.49
		.64	.99	.76	1.00	1.06	1.34	.66	1.61
		.69	.92	1.11	1.26	1.06	1.26	.75	1.66
		.59	.97	1.47	1.40	1.24	1.45	1.49	2.46
	240:30	.13	.89	1.04	1.53	1.43	1.75	1.95	2.71
		.15	.95	.97	1.57	1.61	2.80	1.89	2.26
		.27	.78	1.02	1.55	1.35	1.69	1.72	2.13
		.26	.48	.83	1.34	1.18	1.28	1.21	1.98
		.13	.58	.66	1.30	1.18	1.43	.91	1.54
	60:180	.40	.91	.44	.63	1.02	1.23	.36	.90
		.57	.92	.59	.68	.92	.90	.46	1.16
		.44	.64	.49	.58	.83	.82	.25	.62
		.33	.87	.50	.56	.85	.81	.37	1.15
		.45	.75	.59	.82	.87	1.17	.33	.80

Comparison of each mean stay time in this table with its corresponding standard deviation shows that, in the vast majority (89%) of cases, the standard deviation is larger than the mean. In no case is the s.d. less than 65% of the relevant mean value, and all cases with smaller s.d.s than means are for changeovers and dithers. The very large values of many s.d.s suggest that these stay times may well follow an exponential distribution; a method that could be used to investigate this question is presented in chapter 9. However, for present purposes sufficient has already been presented to refute the hypothesis that these subjects are operating a policy of fixed stay times on each panel.

5.2.1 Stay times on VI 120:120

Next I consider whether the mean stay times on each panel are close to optimal despite being highly variable. Table 5.2 is constructed from the session means of time spent on each panel (ignoring therefore within session variability), and gives the means and standard deviations of time spent on each panel for each subject under concurrent VIVI 120:120. Table 5.3 shows the results of an analysis of variance for repeated measures on these data. It can be seen that there is a significant difference in stay times between birds ($p < .5\%$) and also between panels ($p < .1\%$), despite the programmed rates being the same. There is also a significant interaction between these factors ($p < .1\%$).

Table 5.2: Mean and standard deviations of mean times on each panel under VIVI 120:120

Cell Means

Bird-Panel	92	96	120	Mean
1	0.276	0.350	0.624	0.417
2	0.906	2.542	0.994	1.481
Mean	0.591	1.446	0.809	.949

Standard deviations

	92	96	120
1	0.140	0.252	0.044
2	0.170	0.763	0.301

Table 5.3: Analysis of Variance on mean stay times under conc VIVI 120:120

Source	sum of squares	d.f.	mean square	F	tail prob
Mean	27.00	1	27.00	140.98	<.1%
Bird	3.95	2	1.98	10.31	<.5%
1 Error	2.30	12	0.19		
Panel	8.49	1	· 8.49	120.00	<.1%
Panel-Bird	4.86	2	2.43	34.32	<.1%
2 Error	0.85	12	0.07		

Inspection of table 5.1 shows a consistent preference on the part of all three subjects for panel 2 over panel 1, though post hoc comparisons (Scheffe's method) only found the difference for subject 96 to be significant ($p < 5\%$).

The combined existence of dithers, substantial within session variability in stay times, and significant differences in mean stay times across panels, implies that there is almost certainly no validity in calculating theoretical mean stay times from equation 5.6. None the less, the differences between subjects might conceivably be elucidated by such a procedure, so in table 5.4 I present the derivation of values of a^*, the optimal stay time on the better panel. Remember that in equation 5.6 a^* and T are measured in terms of the schedule value, and so in table 5.4 the average T for each subject must first be scaled in terms of that value; and the value of a^* must, after calculation, be scaled up into seconds.

Table 5.4: Derivation of a^* $(=2/3T^2)$

Bird	T/s	T/120s	a^*/120s	a^*/s
92	2.656	.525	.000326	.0391
96	4.614	.0385	.000986	.118
120	1.040	.00867	.0000500	.00601

It can be seen by comparing tables 5.4 and 5.2 that the optimal stay times are orders of magnitude different from the obtained values, and that they do not even form the same rank ordering.

There does not seem to be any aspect of responding on concurrent VIVI 120:120, therefore, that is consistent with the predictions of optimality considerations.

5.2.2 Stay times on VI 240:30 and VI 60:180

Table 5.5 gives the mean and standard deviations of the session average stay times for each panel on VIVI 240:30 and VIVI 60:180. It was predicted above that visits to the side of the panel with the lower rate should be constrained to a minimal duration. From table 5.1 an operational definition of a minimum duration could be about .5s.

Table 5.5: Means and standard deviations of time on each panel under conc VIVI 240:30 and 60:180

Cell Means

Condition	-240:30------			------60:180------			
Bird-Panel	92	96	120	92	96	120	Mean
1	0.106	0.532	0.198	0.356	1.586	0.430	0.546
2	1.194	3.782	0.955	1.458	0.434	0.518	1.405
Mean	0.650	2.157	0.576	0.907	1.010	0.474	0.976

Standard Deviations

Condition	-240:30------			------60:180------		
Bird-Panel	92	96	120	92	96	120
1	0.075	0.319	0.080	0.310	0.302	0.094
2	0.173	0.587	0.100	0.932	0.176	0.073

Inspection of table 5.5 suggests that under VIVI 240:30 reduction of minor schedule visits to a minimal duration might be taking place, but only subject 96 spends longer on the major schedule for conc VIVI 60:180.

Table 5.6 presents the results of an analysis of variance for repeated measures on the data summarized in table 5.5.

**Table 5.6: Analysis of Variance on mean stay times
under VIVI 240:30 and 60:180**

Source	sum of squares	d.f.	mean square	F	tail prob
Mean	53.35	1	53.35	320.63	<.1%
Condition	1.58	1	.1.58	9.47	<1%
Bird	11.77	2	5.89	35.38	<.1%
Cond–Bird	5.29	2	2.64	15.90	<.1%
1 ERROR	3.83	23	0.17		
Panel	10.54	1	10.54	95.04	<.1%
Panel-Cond	10.23	1	10.23	92.24	<.1%
Panel-Bird	1.30	2	0.65	5.87	<1%
P*C*B	3.98	2	6.99	63.00	<.1%
2 ERROR	2.55	23	0.11		

It can be seen that the effects of Bird, Experimental Condition, and Panel are all significant (at least $p < 1\%$), as are all possible interactions ($p < 1\%$ or better).

With the same provisos as were made when optimal stay times were calculated for VIVI 120:120, the optimal stay times for the better schedule (a_1^*) are presented in table 5.7. The values for VIVI 60:180 were read off Houston & McNamara's 1981 figure 3 (this figure gives t_1^* which is $= a_1^* + T$), and the approximate values for VIVI 30:240 were obtained from their table 6. Because, as mentioned in the introduction to this chapter these values cannot be calculated directly, Houston & McNamara give a selection of values of a_1^* for different T and λ. However since T in this experiment was set by the subject's own behaviour they are not exactly equal to any value given by Houston & McNamara. Values of a_1^* in table 5.7 are therefore only given to the nearest second.

Table 5.7: Optimal stay times on conc VIVI 240:30 and VIVI 60:180

Bird	Cond	T/s	T	a^*	t^*/s	a^*/s
92	240:30	1.91	.064	1		30
	60:180	4.49			26	21
96	240:30	2.52	.084	1.21		36
	60:180	2.83			15	12
120	240:30	1.15	.038	.47		14
	60:180	.90			4	3

This is quite sufficient to demonstrate that, as for VIVI 120:120 the predicted values are an order of magnitude different from those observed. This time, in contradistinction to the values found when the schedule values were equal, the predicted values are larger than the obtained ones. Under VIVI 240:30 the rank ordering of the predicted and obtained values are the same, but this is not the case on VIVI 60:180.

5.2.3 Reward Rates

Taking 1.5s as a value of T that all subjects should be able to attain (see table 5.1), table 5.8 takes the nearest value to this available in Houston & McNamara (1981) and calculates the maximal reward rate (R^*) obtainable by adopting the optimal panel stay times. For equal schedule values R^* is calculated from their equation A2.7 and 10. For VIVI 240:30 and VIVI 60:180 the two nearest values to 1.5s were obtained from Houston & McNamara's table 6.

Table 5.8: Maximal reward rates for 1.5s COD

Condition	T/s	T	R^*	R^*/s
120:120	1.5	.0125	1.975	.0165
240:30	1	.0333	1.117	.030
240:30	2.5	.0833	1.105	.030
60:180	1	.0166	1.322	.0126
60:180	2	.0333	1.312	.0219

From table 5.8 we can see that R^* changes so slowly with increasing T for VIVI 240:30 that the value of .030 is clearly good for T=1.5s. For VIVI 60:180 the position is slightly less clear, but a value of R^* = .020 would appear to be a reasonable approximation for a 1.5s COD. Table 5.9 presents the actual reward

rate, **R/s** averaged over the last five sessions of each condition (calculated from the number of reinforcements and durations given in table 4.2), compared with the optimal rate of reinforcement (R^*), and the number of rewards lost in each 45min session by not using the optimal stay times.

Table 5.9: Obtained reward rates compared to optimal values, assuming $T = 1.5s$

Bird	Cond	R/s	R^*/s	Lost Rfs/Session
92	120:120	.011	.017	44.6- 29.7= 14.9
	240:30	.021	.030	81 - 56.7 = 24.3
	60:180	.014	.020	54 - 37.8 = 16.2
96	120:120	.011	.017	44.6- 19.7= 24.9
	240:30	.021	.030	81 - 56.7 = 24.3
	60:180	.012	.020	54 - 32.4 = 21.6
120	120:120	.014	.017	44.6- 37.8= 6.8
	240:30	.030	.030	81 - 81 = 0
	60:180	.017	.020	54 - 45.9 = 8.1

This table shows that the losses to the subjects are real and substantial, even after allowing a rather generous CO time. Only subject 120 loses less then 10 reinforcements per session, and this is simply because his actual T was less than the 1.5s allowed here (see table 5.1).

5.3 Discussion

The results presented here have shown a substantial failure to optimise on the part of Barbary Doves foraging in a simple experimental environment. The subjects fail to maximise their returns in five ways.

1. Stay time on each side of the floor panel are highly variable, rather than being constrained to single values.
2. Many trips to the centre of the panel are concluded by returns to the same side of the panel (dithers).
3. The amount of time spent changing over is not restricted to a minimum.
4. Average stay times on each side of the panel are orders of magnitude different from those predicted. For VIVI 120:120 observed times are too long; for VIVI 240:30 and VIVI 60:180 the stay times are too short.
5. Stay times on the less good schedule do not appear to be as vanishingly brief as was predicted; on VIVI 60:180 they are not shorter than visits to the better schedule.

A number of explanations of this failure to optimise are possible.

Firstly it might be suggested that the system is so unresponsive to changes in subjects allocation of time within the experiment that even the substantial deviations presented here do not lead to sufficient losses of reinforcement, or excess effort on the subjects' part, to be worth worrying about. In table 5.9 above it was demonstrated that two out of three subjects were obtaining between 15 and 25 fewer reinforcements per 45 min session than was possible; and that this involved the additional costs of more rapid changing over implied in the existence of dithers and excessively short stay times. 20 reinforcements in 45 mins is a rate of .074Rf/s or VI 135s. This is within the range of the schedules programmed in the experiment and cannot be considered trivial.

A popular argument to dismiss unpleasant results from experiments using operant scheduling is to point out that the experimental environment is unnatural. I am immediately willing to concede that this experiment was 'unnatural'. It is, however, important that we should be able to identify what it is that makes the experiment unnatural, apart from the fact that the subjects did not forage optimally in it. What is it that makes an experiment in which a dove (that was reared in a cage), has to walk around a rather similar cage in order to operate a reward producing mechanism, so much more unnatural than an experiment in which a feral born great tit chooses food items off a conveyor belt? (Krebs et al 1977). That different spatial locations should be associated with different densities of food, even in a small area is trivially obviously a feature of the world outside the laboratory; whereas, outside a few university canteens, conveyor belt food delivery systems are not reported.

On reflection, however, there is a substantial artificiality to this experiment that appeared, from informal observations of subjects (as noted in chapter 4) to lead to the birds structuring the task differently than was assumed in the optimality analysis. There are two components to this. Firstly, that the hopper was situated centrally; and secondly that the doves appeared to be unsettled by the movement of the floor panel. The net effect of these two factors was that the birds stopped in the centre of the apparatus to examine the hopper regularly, and only rather gingerly and briefly depressed one side of the panel before returning to the middle of the chamber again.

If this explanation is correct then we can predict that a functionally equivalent set up, but using a separate hopper at each end of the chamber to pay out rewards collected for standing at that end of the chamber, and a stable

floor, would improve performance.

Shettleworth (1985b) ran an experiment using a chamber possessing a floor panel with hoppers at each end and only tilting .5cm (as per Baum & Rachlin 1969). The purpose of her study (discussed in the introduction to chapter 4) was to investigate the preference of pigeons for different numbers of food items under different delays. She also ran two conditions (first and last conditions of experiment 1) of concurrent equal VI schedules (90:90 and 60:60 respectively) without any other complicating factors. Unfortunately it is not clear whether the schedule values were constant throughout the three sessions of the first of these conditions. That leaves a set of data averaged over the last three of ten sessions on conc VIVI 60:60. Table 5.10 presents the mean time per session on the left (T_L) and right (T_R) sides of the floor panel, the mean numbers of reinforcements (R_R and R_L) and the mean number of changeovers (N_{co}) per session, extracted from Shettleworth 1985b table 1. From these data, and the information that the experiment used a 4s changeover delay the optimal stay times (a^*) and maximal reward rates (R^*) were constructed. These are .178s and .0312Rf/s respectively.

Table 5.10: Responding on conc VIVI 60:60 from Shettleworth 1985b

Bird	Session		Means		N_{co}	duration	Rf/s	Lost Rfs
	T_L	T_R	R_L	R_R				
2	210.5	199.7	20.3	19.6	194	1312.7	.0305	1.8
5	195.7	167.3	18.6	21.3	247	1481.5	.0270	11.3
9	117.2	157.9	21.0	19.0	238	1349.3	.0296	4.2
10	157.4	144.9	19.0	21.0	241	1390.3	.0288	6.4

Assuming there were no 'dithers' (no mention is made of this possibility) the mean stay times per visit were calculated and these are presented in table 5.11.

Table 5.11: Shettleworth 1985b: Mean stay times per visit

Bird	T_L	T_R
2	1.085	1.029
5	.792	.677
9	.492	.663
10	.653	.601

It is clear that although the stay times are again longer than the optimal value, the number of reinforcements that would be lost as a consequence of this over a 45 min session (sessions were in fact run until 40 reinforcements had been received) is much smaller.

Although the use of a different subject species is a confounding variable in comparing Shettleworth's results with the experiment reported here, some independent support is provided for the idea that the central location of the food hopper, and rather large swing of the floor panel, were factors restricting the efficiency of the doves' behaviour. This underlines the point made in chapter 2 that the way the experimenter parses an environment is not guaranteed to be the way the subject conceives of that same environment. It also emphasizes the importance of recognizing that although subject species may not have evolved to live in operant experiments, these experimental environments have evolved to be efficiently utilizable by the subjects. It was mentioned in the introduction to chapter 4 that it had been hoped that the experimental chamber described here, by virtue of its apparent simplicity and ease of operation, might provide a useful basis for research involving phylogenetic and ontogenetic comparisons – that promise has not been fulfilled. The simple modification of providing two food sources and an electronic position-sensing system might well create an environment in which such research could be conducted.

The experiment to be described in chapters 5 to 9 employs a more traditional manner of presenting a bird subject with a concurrent two way choice.

CHAPTER 6
KEY PECK EXPERIMENT:
INTRODUCTION AND EXPERIMENTAL METHOD

6.1 Introduction

The behaviour of animals searching for food rewards in an operant chamber by responding on two manipulanda, each of which reinforces the first response after some time interval has passed since it last payed out, (concurrent interval schedules) has been the subject of at least twenty or more studies over the last thirty years (see Taylor & Davison 1983; Baum 1979; Wearden & Burgess 1982; Myers & Myers 1977 for references). Most work within this paradigm has followed Herrnstein's (1961) demonstration of a simple relation between the ratio of the total numbers of responses to each option over a session (or usually pooled across several sessions), and the ratio of reinforcements obtained due to this responding. This result, known as Herrnstein's 'Matching Law' can be expressed:

Equation 6.1

$$\frac{P_1}{P_2} = \frac{R_1}{R_2}$$

where **P** and **R** are the numbers of responses and reinforcements respectively, and the subscripts identify the two manipulanda.

Baum (1974) generalized this relation thus:

Equation 6.2.1

$$\frac{P_1}{P_2} = b\left[\frac{R_1}{R_2}\right]^c$$

where **b** (known as bias) and **c** (sensitivity) are constants close to unity. In this form it has been found to account for data from a wide range of studies on concurrent interval and ratio schedules with a number of subject species (Baum 1979; Wearden & Burgess 1982). Clearly, when **b** and **c** are both unity, this equation reduces to equation 6.1. Baum (1974) argued that bias represented a consistent preference for one schedule over the other due to un-programmed

assymmetries in the experimental conditions and that only changes in the sensitivity parameter (**c**) were of theoretical interest. Houston & McNamara (1981, appendix 7) have demonstrated that optimal behaviour on concurrent exponential VIs can require a non unity bias parameter without the need to assume any unobserved sources of reinforcement.

It is my contention that the generalized matching law is at least in part artefactual, and that its apparent success has been a hinderance to understanding both the mechanisms and functions of choice in the Skinner box, and to foraging studies in general.

Criticisms of the matching law have been of two types. Firstly it seems that matching is dependent on the techniques that have evolved to produce it (Scoun 1983). The use of a changeover delay (COD; dead time following first response to a new schedule during which no response can be reinforced) was found by Herrnstein (1961) to be necessary for matching to occur; he argued that it served to prevent reinforcement of chains of responses across the two keys. But far from abolishing responding, Silberberg & Fantino (1970); Pliskoff (1971) and Scoun (1983) have found that responding is more rapid during the COD. Scoun (1983) reports that with a 2sec COD on concurrent VI 120:60 around 25% of responses occur during the COD. Brownstein & Pliskoff (1968) ran a preliminary experiment with a variety of CODs and chose for their main study the value that had given the clearest matching.

The procedure for estimating the parameters of the generalized matching equation (Equation 6.2.1) is also questionable. Taking logs of that equation yields a linear relation:

Equation 6.2.2

$$\log\left[\frac{P_1}{P_2}\right] = c.\log\left[\frac{R_1}{R_2}\right] + \log b$$

To find estimates of **b** and **c** the log response ratios are regressed against log obtained reward ratios. This procedure is not strictly legitimate because neither log ratio constitutes a dependant variable. Furthermore data are usually pooled over the last five sessions before being plotted, which reduces the variance in the scores and will tend to produce spuriously high proportions of variance accounted for.

The second set of problems with the generalized matching equation is that it does not determine a pattern of behaviour. That is to say, as a molar generalization, matching could be produced by a wide variety of behaviour generating mechanisms; it is, as noted in chapter 2, a functional theory, not a mechanistic one. Early expositions of matching (e.g. Herrnstein 1961) carried the implication that molar matching in totals of responses and reinforcers was necessary and sufficient evidence of a process whereby obtained reinforcement proportions controlled response choice by the subject. It is now widely understood that this is not the case (e.g. Herrnstein 1982). Not only can matching be produced by mechanisms other than active equalizing of response proportions to reinforcement proportions by the subject, but it can also be produced by optimal responding (consisting of fixed inter switch times; Houston & McNamara 1981 and see below chapter 7); by momentary maximising (Staddon, Hinson & Kram 1981 and see below chapter 8);and even by random responding as would be performed by a perfectly elastic ball thrown into a Skinner box (Shimp 1982c).

These two shortcomings of molar matching theory are, I submit, responsible for the ubiquity of generalized matching with values of c (in equation 6.2; slope of matching line) close to unity on concurrent interval schedules. The only systematic difference in matching parameters that has been observed in 4 review papers is a suggestive difference in values of the c parameter between research performed in New Zealand by Davison and his associates (c≈.9); and work in the United States by Herrnstein, Baum, and others (c≈1.0). Baum (1979) first pointed out this difference and suggested it was due to unrecorded systematic differences in procedure. Taylor and Davison (1983) have re-analysed concurrent VI data from 18 studies to suggest that this difference may be due to the use of VI schedules built on different forms of progression. Davison's group use schedules based on an arithmetic progression of intervals, whilst American investigators use VIs based on an exponential progression creating an approximation to a Poisson distribution of inter-reinforcement intervals (Flescher & Hoffman 1962; Cox & Lewis 1966). Only one prior study has included both arithmetic and exponential VI schedules (Trevett, Davison and Williams 1972); this study was not designed to compare the different forms of schedule, and unfortunately the exponential VI was based on just 14 intervals, probably presented repeatedly in the same order. Subjects were tested on the exponential VI for just two conditions and these were not at the same schedule

values as for the arithmetic schedules. Both types of schedule were pooled in the results presented and the authors make no reference to any differences between the schedule types.

More recent work has emphasized that an analysis of mechanisms of behaviour requires a more fine grained level of data analysis than the totals of responses and reinforcements used in assessing matching. Real (1983; Real & Dreyfus 1985) and Hinson & Staddon (1981,1983) have analysed intact time series of events (switches and responses respectively) to test 'melioration' and 'momentary maximising' respectively (see chapter 8). However it is not just theories of mechanism that require more detailed data to be effectively tested. Recent functional analyses (Staddon, Hinson & Kram 1981; Houston & McNamara 1981) also argue that optimality in behaviour can only be effectively tested by considering the absolute lengths of visits to each alternative and not just the ratios of visit durations.

The present study on barbary doves compares performance on concurrent exponential VI, (VIVIE), concurrent arithmetic VI (VIVIA) and concurrent FI schedules (FIFI) to attempt to answer three questions.

1. Is there any support for Taylor & Davison's (1981) suggestion that different forms of VI schedule produce different forms of matching? (this chapter)
2. How well does the behaviour of ring doves approach optimality on different forms of concurrent interval schedules? (Chapters 7 & 9)
3. Can the behaviour be accounted for by either of two theories of mechanism in the literature ('melioration' and 'momentary maximising') ? (Chapter 8)

A variety of forms of interval schedule were tested here because, as well as being the only variable implicated in altering the c parameter in the matching equation, it seemed intuitively reasonable to expect the temporal patterning of food occurrences to have an impact on the temporal patterning of responding. The different effects of fixed and variable interval schedules of reinforcement on response patterning were first recorded in the work of the Skinnerian school (Ferster & Skinner 1957; Catania 1966). Although arithmetic and exponential V.I. schedules are normally considered equivalent, the temporal patterning of inter-reinforcement intervals created by each is quite different. A well randomised exponential VI schedule gives intervals that follows a Poisson distribution. No segment of such a sequence can provide information about the time to the next event because the probability of an event is always constant (Cox & Lewis 1966). Typical arithmetic VIs on the other hand, are not

only constructed from a progression that would not ensure a constant probability even if well randomised, but also only usually employ a small set of intervals presented repeatedly in the same order (Flescher & Hoffman 1962; Ferster & Skinner 1957).

6.2 Method

6.2.1 Subjects

Six experimentally naive Barbary Doves (*Streptopelia risoria*) housed under the conditions described in chapter 4, aged approximately 9 months at the start of testing, served as subjects. They were maintained at 90% +/- 5% of free feeding weight throughout the experiment by supplementary feeding when required. Birds were assigned randomly to two groups of three: birds 100, 130 and 144 forming group A; birds 65, 106 and 149 formed group B.

6.2.2 Apparatus

All conditions were run in locally built standard two key operant chambers of internal dimensions 35.5 by 35.5 by 30 cm. Two side walls and the ceiling were made of Perspex, and the back was constructed of wire netting to facilitate video analysis; the front wall was sheet metal and contained two keys 1.5 cm in diameter (requiring a minimum force of .05N to operate), 11cm above the floor and 18cm apart, each trans-illuminated by a light emitting diode. The hopper opening (a 5cm square window with hopper of radius 1.25cm at a depth of 2.5cm) was situated in the centre of the back wall, equidistant between the keys and 9.3cm above the floor. Figures 6.1 - 6.3 show the experiment in situ. Reinforcement was 3.5 seconds access to mixed grain during which a hopper light came on. Each experimental chamber was kept in a separate room (approximately 170 by 300 by 200 cm, illuminated by 80-Watts of fluorescent tubing), visually and auditorially isolated, with an Acorn BBC model B micro-computer alongside it for experimental control and recording of every peck and hopper movement to the nearest .01s, subject to a minimum interpeck time of .3s. Data was later transferred to the university's ICL 2988 computer for analysis.

Skinner box in situ

figure 6.1
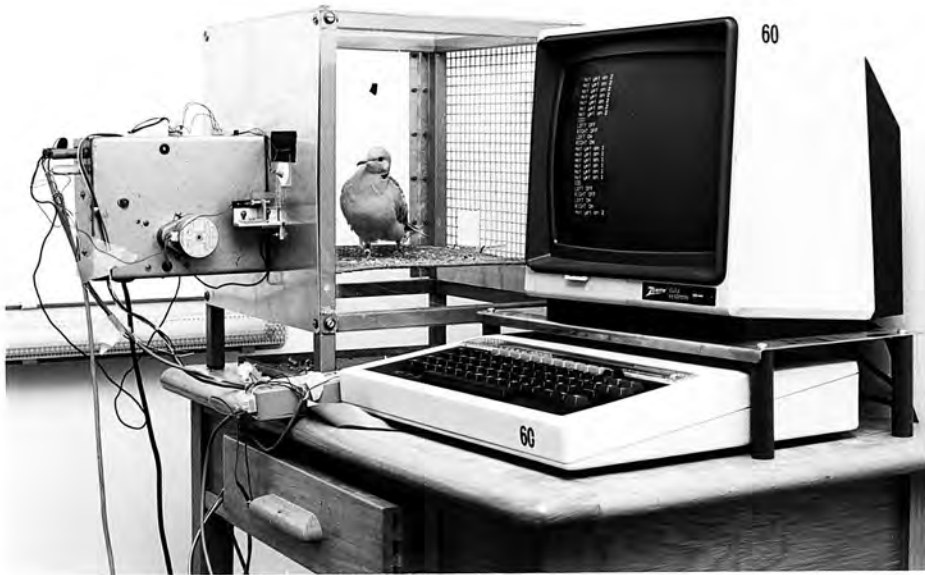
figure 6.2
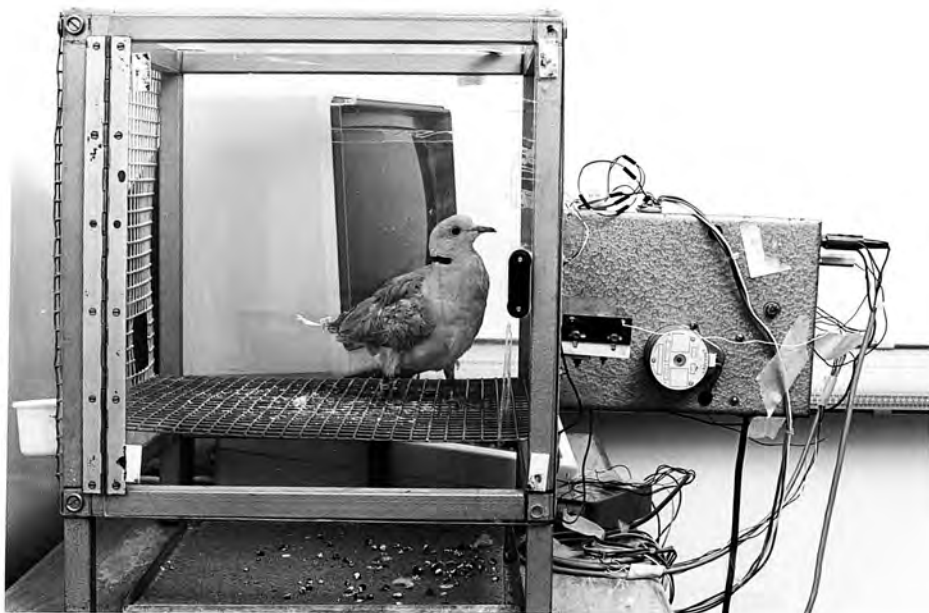
figure 6.3

6.2.3 Schedules

A changeover delay of 2 secs was used for all schedules.

Exponential variable interval schedules (VIVIE) were constructed according to the method of Flescher and Hoffman (1962) and intervals were selected as required at random from a virtual set of 1000.

Arithmetic VI schedules (VIVIA) were constructed from a set of twelve intervals from a progression of the form a + n.r where a was 1/12th the mean inter-reinforcement interval (iri), r was 1/6th the mean iri, for n = 1 to 12. The order of this twelve interval series was randomised before any testing commenced but subsequently always presented in the same sequence. The series was therefore equivalent to a loop of tape. Due to a programming error no pecks were recorded during COD.

6.2.4 Procedure

Experimental conditions and numbers of sessions are shown in Table 6.1.

Table 6.1:
Conditions and Number of Sessions

Group **B:**   Schedule – **VIVIE**

Bird	120:120 Number Sessions	Major Key	240:30 Number Sessions	Major Key	180:60 Number Sessions
65	57	2	25	1	36
106	51	1	46	2	56
149	44	2	54	1	89

Group **B:**   Schedule – **VIVIA(B)**

Bird	120:120 Number Sessions	Major Key	240:30 Number Sessions	Major Key	180:60 Number Sessions
65	31	2	12	1	17
106	34	2	11	1	12
149	19	2	12	1	25

Group **A:** Schedule – **FIFI**

Bird	120:120 Number Sessions	Major Key	240:30 Number Sessions	Major Key	180:60 Number Sessions
100	53	1	29	2	43
130	35	1	19	2	53
144	54	2	34	1	42

Group **A:**   Schedule – **VIVIA(A)**

Bird	120:120 Number Sessions	Major Key	240:30 Number Sessions	Major Key	180:60 Number Sessions
100	36	2	15	1	13
130	17	2	19	1	13
144	18	2	11	1	13

Each subject was tested on a condition six days a week so far as possible until it fulfilled a criterion of stability. The rationale behind this criterion was described in chapter 4. It stipulated that the coefficient of variation (mean/ standard deviation) of the logged (to base 10) response ratios should exceed 5 (or in cases where the logged mean was small, that the s.d. be less than .10) on 6 consecutive occasions.

Pre-training consisted of sessions lasting 15 hours each. Each subject received 5 sessions of irregular unconditional reinforcement, followed by up to 5 sessions of autoshaping and one 15 hour session where reinforcement was obtained for the first peck after crossing over from one key to the other.

The first 8 sessions of VIVIE 120:120 lasted 240mins and involved a 4sec change over delay (COD) (signalled by extinguishing key lights and sounding a

buzzer on the micro computer) and 6sec access to grain as reinforcement. These sessions have been excluded from the analyses. All other conditions used a COD of 2secs. Reported sessions lasted 45 mins except for a few early ones of VIVIE lasting 60 mins.

Group A experienced three conditions of FIFI, followed by three conditions of VIVIA. Group B first received three conditions of VIVIE followed by the same three conditions of VIVIA. Group A were therefore progressing from completely stable interval schedules to moderately variable interval schedules, whilst Group B experienced firstly completely random schedules before moderately variable ones. This partial crossing over was designed to make it possible to look at the effect of prior experience of temporal variability on present performance.

Video tapes of 18 early sessions were analysed for the occurence of the following 9 behaviour patterns. These were chosen with the aim of estimating the extent to which a subject was directing its behaviour towards the experiment.

1. Proximity to – Back Wall
2. " – Left Wall
3. " – Front (hopper) wall
4. " – Right Wall
5. Head oriented to – Left Key (Key 1)
6. " – Hopper
7. " – Right Key (Key 2)
8. Wing Flap
9. Preening

6.3 Molar Results and Discussion

Video analysis was too sporadic to be considered separately but will be mentioned below where it illuminates results from other data.

Time series results will be presented and discussed separately for each method of analysis in the following chapters. This chapter address molar results (totals of events at the end of a session) only.

6.3.1 Sessions to Stability

One unexpected effect I wish to draw attention to is the difference in numbers of sessions required to fulfill the stability criterion under different experimental conditions. An analysis of variance on the durations of each condition, presented in table 6.1, found an overall significant difference between the means (F= 21.6, df 3,8; p<.1%) and post hoc comparisons (Scheffe's method) showed differences at the same level of significance in each pair-wise comparison of different schedule types, but no difference between the two groups of subjects on the VIVIA condition.The trend of the mean number of sessions required for stability is VIVIE > FIFI > VIVIA (means 50.9, 40.2, and 18.2 respectively). (11 sessions are required to fulfill the stability criterion.) I am not aware of any published research that has reported differences in numbers of sessions required to reach stability under different schedule conditions. It is of note however that all the work from Davison's laboratory (where arithmetic VI schedules are used exclusively) makes use of a relatively rigorous criterion of stability (that the median response proportion of the last five sessions should fall within 5% of the median of the previous five sessions on 5, not necessarily consecutive occasions, for the whole group of subjects in an experiment; Davison 1972), whilst work by other authors utilizing exponential schedules is more variable in the methods employed to determine when to change conditions. Some decide arbitrarily the number of sessions they wish to run (Hinson & Staddon 1983; Green, Rachlin & Hanson 1983); others choose by inspection of the choice proportions (Baum 1982; Moore 1984; Pliskoff, Cicerone & Nelson 1978; Vaughan & Miller 1984; Baum 1976; Bradshaw, Szabaldi & Bevan 1976). Where criteria have been employed these have tended to be weaker than that of Davison (1972). The method of McSweeney (1975; Norman & McSweeney 1978), that the overall rate of responding in the last five sessions should fall within the range set by the earlier sessions, fails to relate to response allocation at all. Heyman (1979) required that both the relative response rate and the overall probability of a change over should not show an extreme value for five sessions. Although this relates to the right variables it is not clear how firm a criterion is set by simply insisting that data do not exceed the previously defined range. On balance therefore, it seems possible that a difference in the number of sessions required to reach stability when using different types of schedule, has been obscured by the different practices routinely used in different laboratories to ascertain when response choice

should be considered stable.

That choice in a randomised environment (VIVIE) takes longer to stabilize than in less variable environments would be expected from any account of behaviour that stresses an organisms active searching for regularities in its environment to guide patch selection. However from this perspective it would seem strange that the FIFI condition requires longer for stability to be attained than the VIVIA condition. Two explanations can be put forward for this. Firstly the results presented here could be contaminated by order effects. For both groups of subjects the VIVIA was the last schedule type run and so the birds may have been getting quicker at reaching stability independent of the schedule type presented. Inspection of the data for each condition in table 6.1 does not suggest any trend in sessions required within the schedule types. Alternatively the critical factor in the VIVIA could be that because the twelve component intervals were shorter than the FI value, and their relative sizes and order of presentation did not vary from condition to condition – possibly this may have provided an opportunity for cognitive set learning (Harlow 1959). This possibility will be explored further in chapter 7.

6.3.2 Matching Parameters

The parameters of the best fitting generalized matching equation (Equation 6.2.1) were calculated by the standard method of a linear regression of the log ratios of pecks on log ratio reinforcements pooled over the last five sessions; the analysis was also performed including each session's ratios as separate points. Figure 6.4 shows pooled data and regression lines. Table 6.2 gives the parameters of the regression equations for both pooled (type 'P') and separate ('S') analyses.

figure 6.4   MATCHING LINES FOR POOLED DATA   86

Figure 6.4: Matching lines for pooled data — plots of Log Resp Ratio versus Log Rf Ratio for FIFI, VIVIE, VIVIA (GROUP A), and VIVIA (GROUP B).
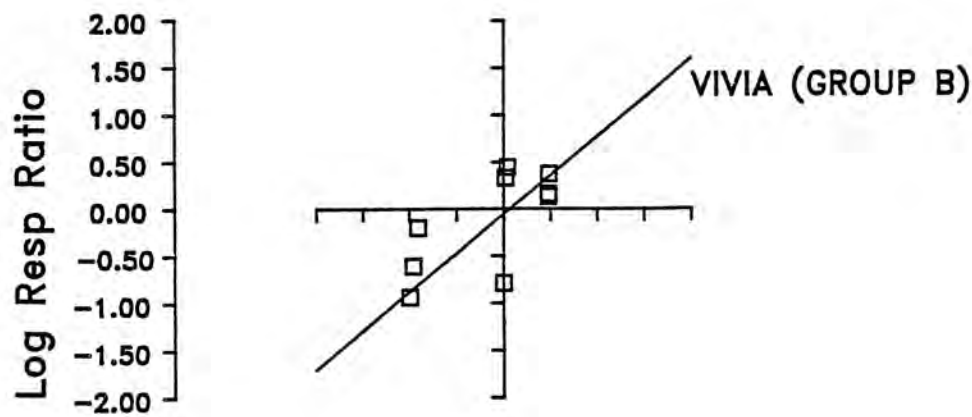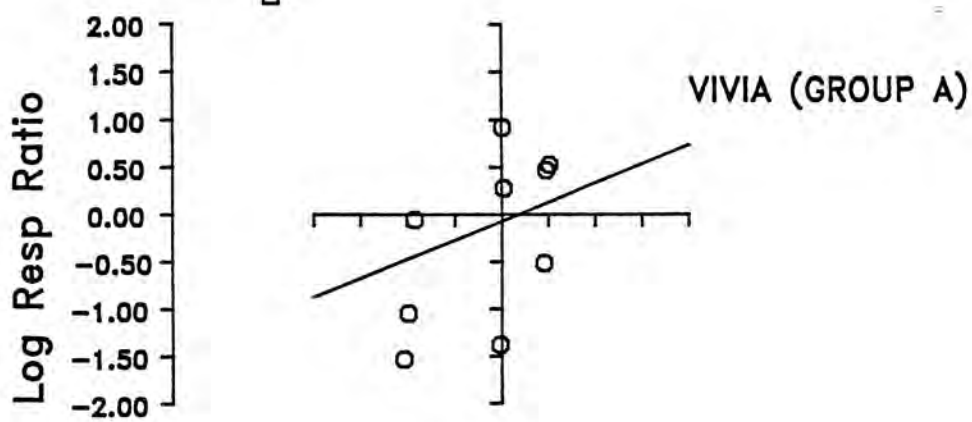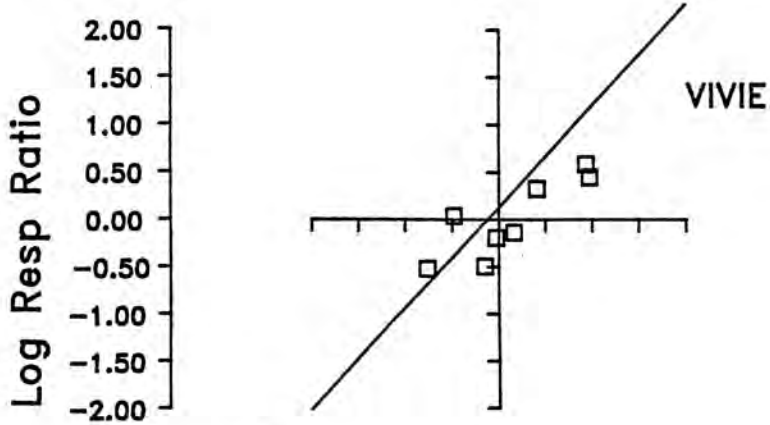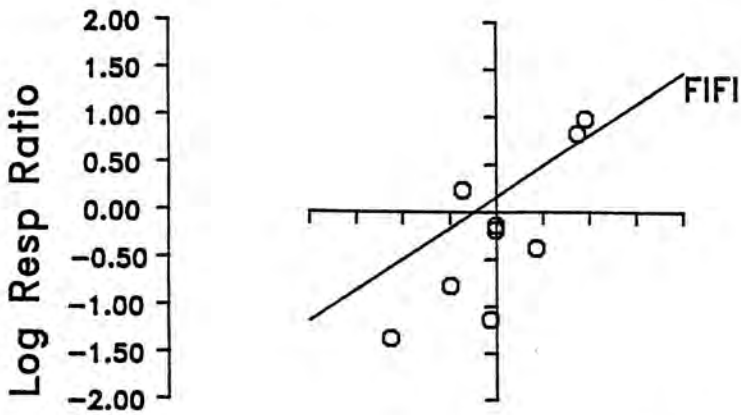
Table 6.2:
Parameters of Matching Regressions

Conditions & Groups	Intercept (log b)	Gradient (c)	Tail Prob	r^2	Type
TOTAL	.0395	.586	< .1%	.508	S
"	.0421	.729	< .1%	.588	P
FIFI	.163	.687	< .1%	.679	S
"	.158	.659	<1%	.650	P
VIVIE	.140	.668	< .1%	.686	S
"	.135	1.071	< .1%	.996	P
VIVIA (A)	-.071	.362	< .1%	.280	S
"	-.066	.402	NS	.311	P
VIVIA (B)	-.059	.797	< .1%	.468	S
"	-.053	.825	<5%	.481	P

Performing regressions on the data from separate sessions increased the number of data points five fold from the usual methods of estimation and made it feasible to perform an analysis of variance to compare the regression lines. The F statistic used was the ratio of the mean square deviations over groups divided by the mean square deviance within groups (BMDP Program 1R; Dixon et al 1983). For pooled data $F_{6,\ 28}$ = 1.392, NS; but for each session separately $F_{6,\ 172}$= 3.494, p<5%. The results of analyses of variance for each pair-wise comparison of groups counting each session separately are shown in table 6.3;

Table 6.3:
ANOVAs on regression equations

Conditions & Groups	F- stat $df_{2\ 86}$	2- Tail Probability
FIFI, VIVIA(A)	5.121	<1%
FIFI, VIVIA(B)	3.971	<5%
FIFI, VIVIE	.038	NS
VIVIA(A),VIVIA(B)	3.660	<5%
VIVIA(B), VIVIE	2.977	<10%

In table 6.2 the gradient of the regression lines (sensitivity parameter **c** in equation 6.2) were usually closer to unity for pooled than for separate data, similarly the r^2s were greater for pooled data. This confirms Scoun's (1983) argument that the methods of assessing matching are predisposed towards producing linear matching.

In terms of the sensitivity parameter **c** the results can be ordered thus:

```
    NS          *      **
VIVIE >   VIVIA(B) > FIFI > VIVIA(A)
1.07      .83        .66    .40
```

(where '*' indicates a difference significant at the 5% level; '**' significance at 1%.)

These differences in the sensitivity parameter run in the counter direction to those found by Baum (1979) and Taylor & Davison (1983) by comparing results from different studies. Slight over-matching ($c > 1$) is here found with the exponential VI schedules, and under-matching ($c < 1$) with the arithmetic VI schedules. The analysis of variance used here is a relatively robust form of comparison compared with the non-parametric test used by Taylor & Davison (1983) which made no allowance for the enormous differences in intra-study variance (Taylor & Davison 1983; table 1). I suggest that the differences found between the studies from different laboratories are more likely to be due to unrecorded differences in the manner of measuring matching (e.g. extent to which data are pooled before calculation), or in the manner of signalling change over delay, than to any difference in schedule types. One way in which matching parameters can be altered by changing the method of analysis has been demonstrated above. Scoun (1983) found that responding during COD is indifferent to the schedules, and that this responding is abolished if the operant chamber is blacked out during COD. The manner of signalling the COD is therefore a factor that differs between laboratories, and would be likely to affect matching parameters.

Over-matching implies that the subject's preference for a schedule is in excess of the proportion of reward gained from it - such behaviour is equivalent to what is termed 'Win Stay' behaviour in spatial choice conditions (Olton 1978). Arithmetic VI and FI schedules contain periods during which the probability of reward is zero (see chapter 7); this may provide an affordance (Gibson 1977) to which the subjects are sensitive, that enables them to interweave responding on the two keys efficiently, and even respond disproportionately on the less good schedule without losing much reinforcement. On an exponential VI condition the subjects, lacking any clear affordance to guide choice may revert to a simpler 'win-stay' type strategy tending to move preference disproportionately to the better schedule. This interpretation is consistent with the finding that preference took longer to

stabilize on the more variable schedules.

CHAPTER 7
KEY PECK EXPERIMENT: OPTIMALITY ANALYSIS

7.1 Introduction and Analysis Methods

Optimal foraging theory has been discussed in chapter 2, and experimental tests of it in chapter 4. In brief it is an attempt to find the best returns in terms of a fitness related parameter (typically net energy gain) available to a predator searching for food. The classical optimal foraging model considers a forager searching for prey in discrete patches which deplete as the food items are removed from them (Charnov 1976). Concurrent interval schedules differ from this pattern in that the 'patches' or schedules do not deplete as 'prey' or reinforcements are removed.

7.1.1 Concurrent Exponential VI

The first attempts to deduce optimal responding under VIVIE (Rachlin 1978; Rachlin & Burkhard 1978; Staddon & Motheral 1978; Staddon 1980) involved assuming a schedule feedback function to relate average reinforcement rate to response rate on each key. These two schedule functions were then combined into a single utility function which was maximized mathematically to find the best possible return of reinforcements on responding. Since the analysis was couched in terms of relations between mean molar rates of responding and reinforcement it lent itself readily to being expressed in terms of the matching relation. For schedules of the same type simple linear matching (Equation 6.1) was supported (Rachlin 1978; Staddon & Motheral 1978).

Critiscms of this form of analysis fall into two groups. Firstly in specifying a priori a relation between response rate and reinforcement rate it is already assuming part of what is under test. How are we to know that the assumed relationship is itself optimal? Houston & McNamara (1985) discuss the limitations of prior constraints in optimal foraging analyses. The schedule feedback functions also incorporate nothing to distinguish different forms of interval schedule. We shall see below that the different forms of concurrent interval schedule demand, in fact, quite different optimal strategies. The second problem with this form of analysis is that in combining the two schedule functions it makes no allowance for the change over delay and implicitly allows

responses to occur on both keys simultaneously (Houston & McNamara 1981).

The optimality analysis of Heyman & Luce (1979) deals explicitly with the two ways in which reinforcements can be obtained – by staying on a key or crossing over – but still does not allow for the COD. Their model also assumes that the rate of responding is high enough for reinforcements to be collected as soon as set up; and takes the observed constant probability of switching (Heyman 1979; Hinson & Staddon 1983) as a constraint on possible strategies.

The approach presented above for concurrent VT schedules (chapter 5 from Houston & McNamara 1981) imposes no prior constraints on the behaviour. However, that analysis cannot be extended to response dependant VI schedules without assuming that the rate of responding is high enough to pick up each reinforcement as soon as it is set up. To assume this, as Houston & McNamara (1981) point out, is to assume that the organisms are behaving optimally – thereby invoking an inverse optimality argument (McFarland & Houston 1981), and assuming what we are trying to test. This nullifies any attempt to extend the length of stay component in the analysis of chapter 5 to response dependant schedules. The prediction that visits to the minor schedule should be instantaneous, however, can still be upheld because the question of response rate cannot apply to an instantaneous visit. The duration of visits to the minor key is therefore our first measure of optimality on VIVIE.

McNamara & Houston (1983) consider the optimal response rate on an exponential VI. They demonstrate firstly that the optimal policy involves a constant response rate. Figure 7.1 (from McNamara & Houston [1983] fig 1) shows this graphically. Time since last response is plotted against reinforcement probability. The solid line shows the relationship between the two ($y=1-e^{-\lambda x}$). It is clear that the convexity of this curve implies that the net gain from a mixed strategy of, say, equal proportions of inter-response times $x+\varepsilon$ and $x-\varepsilon$ can always be improved upon by a pure strategy of inter-response times x. McNamara & Houston (1983) go on to demonstrate that the net reward rate for a constant inter-response time x (R_x) is given by:

figure 7.1
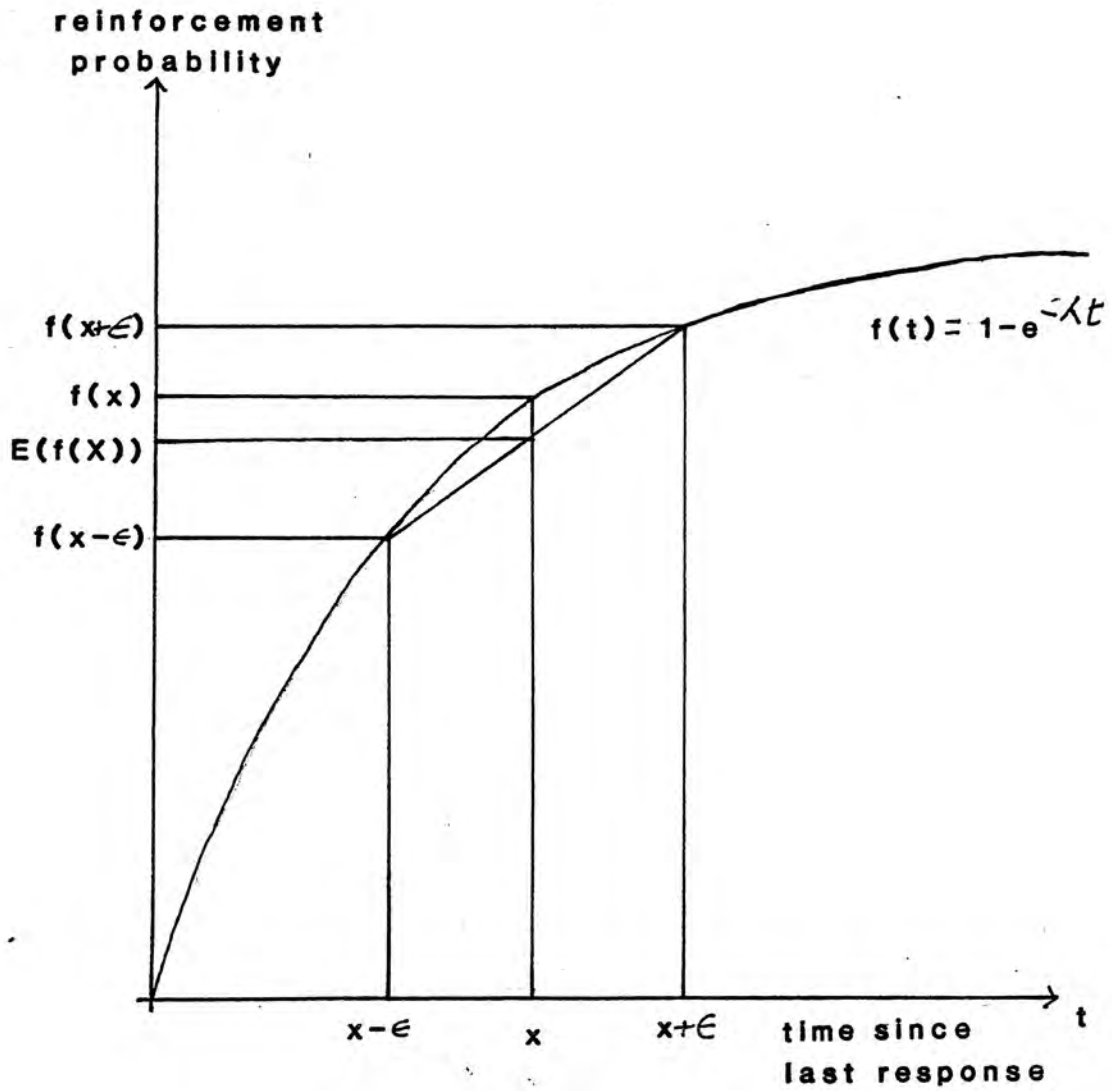
from McNamara & Houston (1983)

$$R_x = M \, \frac{1 - e^{-\lambda x} - c}{x}$$

where **M** is the value of a reinforcement, **c** is the response cost as a proportion of reinforcement value, and λ is the programmed rate of reinforcement. The value of **x** that maximizes R_x, x^*, they show is described by the following equation:

$$\lambda x^* \simeq \sqrt{2c}$$

for very small c; and so

$$\frac{1}{x^*} \simeq \frac{\lambda \sqrt{M}}{\sqrt{(2c_0)}}$$

where c_0 is the absolute response cost. From this it can be seen that the optimal response rate for a constant response cost and reinforcement value is directly proportional to the programmed rate of the VI schedule. Patterns of responding are analysed in detail in chapter 9 but in so far as these results relate to the optimality analysis they will be discussed in this chapter.

7.1.2 Concurrent Fixed Interval

Fixed interval schedules differ in kind from the exponential VI schedules described above. Whereas the probability of reinforcement on an exponential VI is a continuous function in time ($R_x = 1 - e^{-\lambda x}$ where x is the time since last response; see figure 7.1); on an FI schedule the probability of a reinforcement is zero until the programmed FI has elapsed, it then becomes unity until the reinforcement is collected, when it returns to zero again. The probability function is therefore discontinuous. The optimal strategy on such a schedule is trivial: wait until the fixed interval has elapsed, respond once, and then wait again. When two independent FI schedules are running concurrently the optimal policy remains similar, fitting the COD and response on one schedule into the interval during which no responding is required on the other. In practice the perfect timing required by this strategy is never found and responding begins before the inter-reinforcement interval has expired (Ferster & Skinner 1957). Schneider (1969) defined a post reinforcement pause on FI schedules as the intersection of the two best fitting straight lines on a plot of response rate

against time since reinforcement. He found that the value of the FI schedule was the almost exclusive determinant of the duration of post reinforcement pausing. Here I will use a simpler criterion – the interval between a reinforcement and the first response after it to the same schedule – as a measure of the sensitivity of the subjects to the duration of the programmed inter-reinforcement intervals.

7.1.3 Concurrent Arithmetic VI

An arithmetic VI schedule lies somewhere between the two forms of schedule described above. McNamara & Houston (1983) considers the case of an arithmetic VI constructed of sufficient intervals to create a continuous probability function. Unlike an exponential VI, where the probability of reinforcement is a function of time since last response, on an arithmetic VI the probability of reinforcement is an exponentially increasing function of time since last reinforcement. The precise form of the schedule used in this experiment (following Ferster & Skinner 1957:327 and all subsequent research) does not contain enough intervals to justify being considered a continuous schedule. Instead, since it is composed of twelve intervals presented repeatedly in the same order, I consider that it is more a set of FI schedules chained together – hence the reinforcement probability on this schedule is a discontinuous function. For a time interval t_1 the probability is zero; at the end of this interval it is unity; at the first response after this interval it falls to zero again for a time t_2; and so on through the twelve intervals until the cycle repeats. To assess the optimality of responding on VIVIA the post-reinforcement pauses (as defined above for FI schedules) will be calculated for each of the twelve component inter-reinforcement intervals of the schedules and tested for sensitivity to this dimension.

7.2 Results

7.2.1 Concurrent Exponential VI

7.2.1.1 Key stay times

For the first and last ten sessions of each condition the stay times on the two keys were assigned to one of three classes.

1. Shorter than COD (ineffective on either key).
2. Exceeding the COD by one peck (optimal on the minor key).
3. Exceeding the COD by more than one peck (sub-optimal on the minor key; optimal on the major key).

Figure 7.2 shows the proportion of visits to the minor key (shaded) and major key (unshaded) falling into these categories pooled over all three subjects. For the first condition, VIVIE 120:120, there is of course no major or minor key, so the shading here identifies key 2.

It can be seen that at the beginning of this condition approximately half of the visits to both keys were shorter than the COD and therefore had no effect. By the end of VIVIE 120:120 the proportion of ineffectively brief stay times had fallen to around 30% with approximately equal gains to each of the other two classes. This effect is consistent across the subjects.

At the beginning of VIVIE 30:240 the proportion of ineffectively short visits to the major key has already begun to drop – a fall which by the last sessions of that condition has reduced these visits to less than 10% of the total for that key. Most of the visits to the major key (around 80%) are now in class 3. On the minor key the proportion of visits in class 1 actually rises consistently across subjects at the beginning of VIVIE 30:240. For two subjects it continues to rise to the end of this condition (Birds 65 and 106) while subject 149 shows a small drop. The proportion of intervals for the minor key in class 2, the optimal class, shows no consistent movement over this condition; most of the gains to class 1 are made up from class 3.

For VIVIE 60:180 the order of the keys (in terms of major versus minor) was reversed (see table 6.1). None the less the profile for the beginning of this condition differs little from that at the end of VIVIE 30:240, implying a rapid crossing over of preference. The exception is that for subject 65 – minor key –

Proportion of Runs in Different Duration Classes: All Birds
figure 7.2

First Ten Sessions   Last Ten Sessions

120:120

30:240

60:180
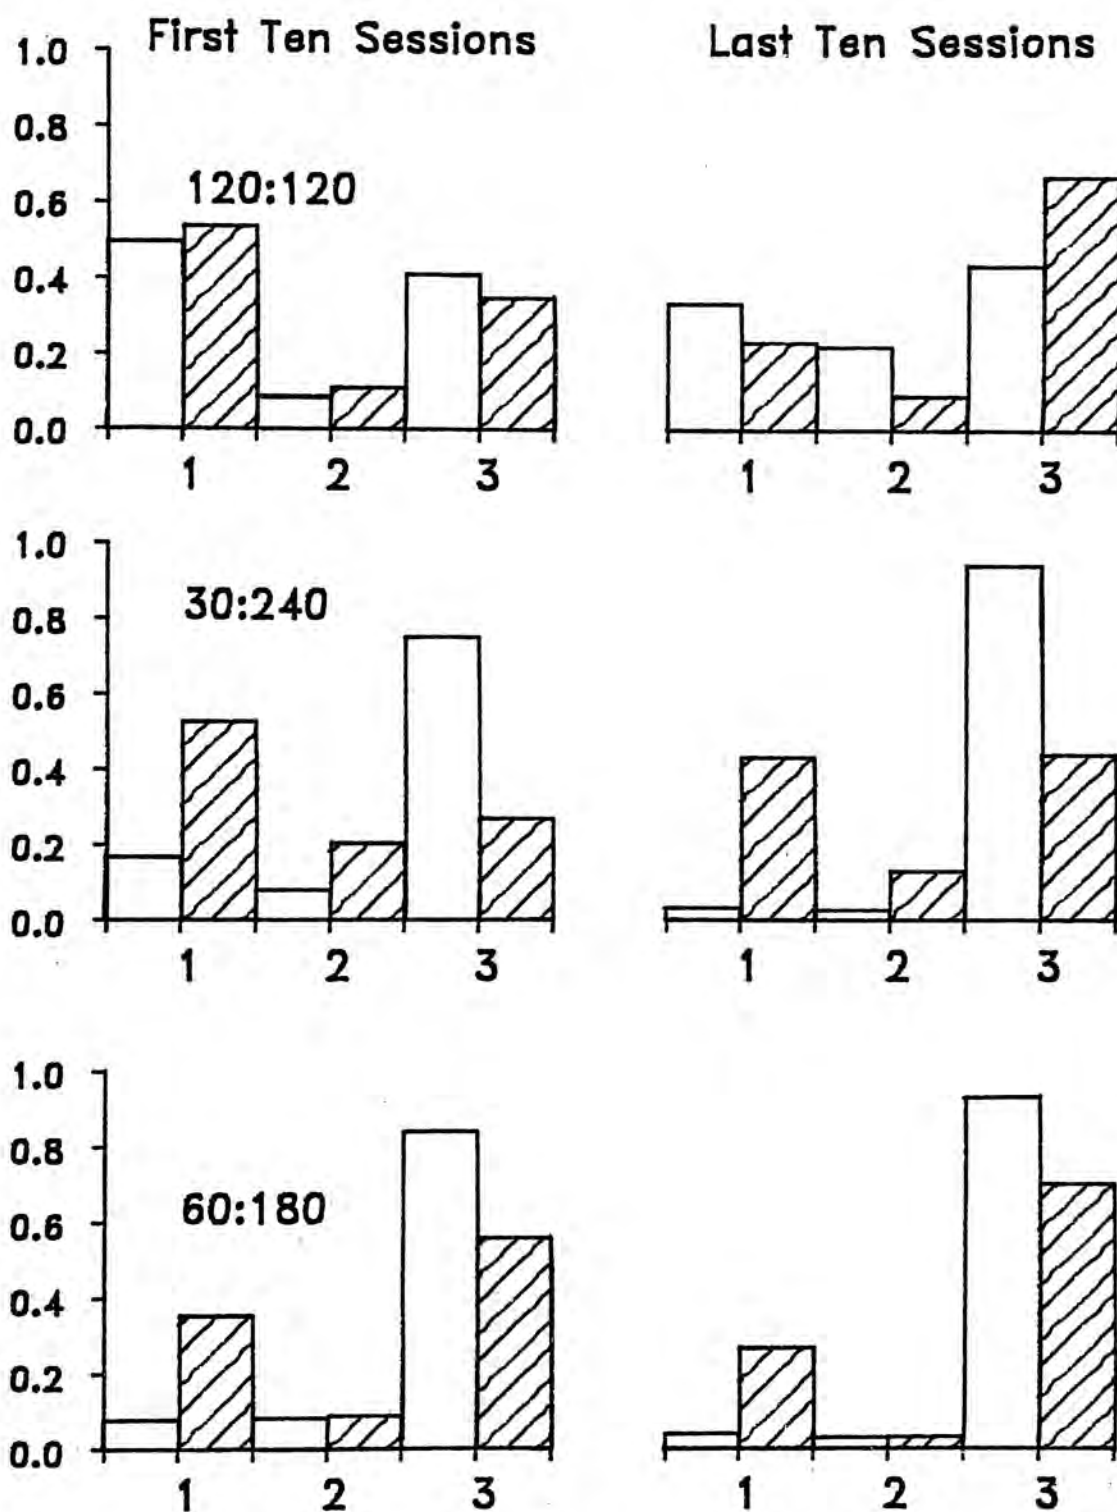
Duration classes. 1: Shorter than C.O.D.
2: C.O.D. plus 1 peck.
3: C.O.D. plus over 1 peck.
Unshaded: Major key.   Shaded: Minor key.

the proportion of visits in class 3 increased substantially. By the end of this condition almost all (>90%) of responses to the major key were falling into class 3 – the optimal class on that key, for all subjects. The minor key shows slightly greater inter-individual differences but it is still the case that class 2 contains an insignificant proportion of visits for all subjects. For two subjects (65 and 149) class 3 holds the largest number of minor key stay times, while for Bird 106 class 1 still holds the longest proportion of visits.

Table 7.1 shows the mean number of pecks per visit (excluding COD) for the last five sessions of each condition for each subject. Due to the unintended exclusion of responses during COD, a visit shorter than the COD is recorded as 1 peck. The results of an ANOVA for repeated measures on these data for VIVIE 240:30 and VIVIE 60:180 are summarized in table 7.2. The effects of subject, condition (240:30 vs 60:180) and key (major vs minor) are all significant at 5%, as are all possible interactions. Post hoc comparison (Scheffe's method) found the mean number of responses per visit were significantly greater (at the same significance level) for the major than the minor key for every subject and condition except Bird 149 on VIVIE 60:180, where the observed difference was not significant.

Table 7.1:
Mean number of pecks per visit (excluding COD).

Subj	120:120		240:30		60:180	
	Key 1	Key 2	Major	Minor	Major	Minor
65	8.67	2.45	29.5	3.55	9.71	5.00
	8.53	2.14	56.5	2.46	9.64	.26
	7.84	2.10	28.2	2.18	10.3	5.71
	8.73	2.33	45.4	1.84	10.7	5.99
	7.28	2.49	32.4	1.49	8.83	5.85
106	2.97	1.74	75.0	1.00	6.93	2.27
	3.01	1.78	162	1.00	7.00	2.32
	4.47	3.50	908	1.00	7.80	2.74
	5.18	3.22	422	1.00	8.08	3.32
	5.52	3.46	117	1.00	7.47	4.04
149	2.37	3.51	9.95	3.21	7.18	7.50
	2.94	3.41	7.58	2.98	7.95	9.25
	2.97	3.92	11.1	2.94	8.01	9.05
	2.45	3.79	10.5	2.65	7.64	9.03
	2.45	3.65	9.96	2.82	8.73	8.10

Table 7.2:
ANOVA for Number of Pecks per Visit (excluding COD).

Source	Sum of Squares	df	F	Tail Prob.
Mean	78310	1	7.81	<1%
Subject	78172	2	3.90	<5%
Condition	50738	1	5.6	<5%
Key	62315	1	6.21	<5%
Sub-Cond	84144	2	4.19	<5%
Sub-Key	84466	2	4.21	<5%
Cond-Key	57307	1	5.71	<5%
Sub-Cond-Key	81288	2	4.05	<5%
Error	481551	48		

7.2.1.2 Rates of Responding

The patterning of responding under the various schedules is analysed in detail in chapter 9. It is sufficient for our present purposes to note that the results of those analyses show:

1. The data are in all cases parsimoniously described by exponential distributions.
2. That in some cases the data are best described by combining more than one distribution.
3. That the means of the distributions are insensitive to the schedule values.

7.2.2 Concurrent FI

Linear regressions (BMDP Program 1R; Dixon et al 1983) were performed to relate the post reinforcement pauses to the scheduled minimum inter-reinforcement intervals, for the last two sessions of each condition for each subject. Pauses briefer than one second were excluded from the analysis on the grounds that they were more likely to form part of a bout of responding during the programmed period of reinforcement (which, due to dead time in the operation of the machinery was about .5s longer than the period during which reinforcement was actually available to the subject). These amount to less than 4% of the total number of responses analysed. Results of this regression are in table 7.3. Although the relationships are slight they are all significant ($p < .1\%$ for t-tests on the gradients, and for regression ANOVAs).

Table 7.3:
Regressions of Post Rf Pauses vs Schedule Values

Bird	n	Gradient	Intercept	r^2
130	343	.0573	977.8	.038
144	224	.2616	1150.4	.059
100	322	-.0374	1615.4	.072

7.2.3 Concurrent Arithmetic VI

Similar linear regressions found the relationship between the post reinforcement pauses on concurrent arithmetic VI and the programmed minimum interval till next reinforcement (i.e. the next interval in the scheduled cycle) was usually statistically significant though of very small magnitude. Table 7.4 gives the regression parameters for each subject pooled over the last two sessions of each condition.

Table 7.4:
Regressions of Post Rf Pauses vs Inter-reinforcement intervals

Bird	n	Gradient	Intercept	R^2	T(For-gradient)	p (2-tail)
100	293	.0494	733.7	.0762	4.90	<.1%
130	309	.0050	655.3	.009	1.67	<10%
144	300	.0849	898.3	.068	4.67	<.1%
65	317	.0687	471.0	.126	6.75	<.1%
106	307	.0332	598.0	.101	5.85	<.1%
149	308	.0022	835.4	.001	.39	NS

7.3 Discussion

7.3.1 Concurrent Exponential VI

The above results suggest firstly that the birds are sensitive to the need to control the length of their visits to a schedule, dependant on the context in which the schedule is found. Secondly the changes in proportions of visits during VIVIE 30:240 suggest that the subjects are first reducing the duration of minor key visits to a minimum, but then towards the end of VIVIE 30:240, and continuing into VIVIE 60:180 the ineffectiveness of visits shorter than the COD causes the lengthening of minor key stay times. Since class 2 (COD plus 1 peck) covers a more limited range of possibilities than the other two stay time classes, it is not surprising that, as the minor key stay times increase on the

last condition, class 2 fails to increase its share of the visits (it can be seen however from table 7.1 that the minor key visits are shorter than major key ones). The absence of a significant difference between the major and minor key run lengths for Bird 149 on VIVIE 60:180 is the only item against this argument.

An alternative explanation might be proposed that the subjects are making all visits to the minor key as short as possible (class 1), but that their perception of relative schedule values is dependant on how different the schedules are. This would explain why the difference in run lengths (table 7.1) is more pronounced for the VIVIE 240:30 condition than for VIVIE 60:180. It could also account for the proportion of minor key visits in class 1 on VIVIE 60:180 being less than on VIVIE 240:30. However such an account would have difficulty explaining why the proportion of visits in this class for the minor key was no greater on the final condition than on VIVIE 120:120; nor can it account for the consistency in major key proportions in class 3 from the second condition to the third.

7.3.2 Concurrent FI and Arithmetic VI

Although the minimum time till next reinforcement accounts for only a small fraction of the variance in post reinforcement pausing on FIFI and VIVIA (<13%), the effect is none the less statistically very significant in most cases. This compares with $\simeq$ 80% variance accounted for in an analysis of variance for FI, using a more robust definition of post reinforcement pausing (Schneider 1969). If it is accepted that had the pausing here been defined by Schneider's (1969) method, the same level of relationship would have been found, then it is noteworthy that there appears to be no trend here in the r^2s between concurrent FI and the two groups on VIVIA. This is very surprising since the arithmetic VI requires the subjects to learn the relative values of twelve intervals per schedule, whilst the FI only involves one interval. The intercepts, however, are smaller for the VIVIA conditions than for FIFI; suggesting that pauses were on average shorter on the former condition than the latter (cf Catania & Reynolds 1968). Although the idea of using linear regression in this context seems new, previous research has shown that post-reinforcement pausing is affected by the type of interval schedule (Catania & Reynolds 1968 [pigeons]; Harzem, Lowe & Priddle-Higson 1978 [rats]) to the extent that pausing is greatest on FI schedules, reduced but present on arithmetic VI schedules, and absent on exponential VI schedules. Table 7.4 is the first

evidence, however, that the length of a pause is dependant on the component interval of the schedule that is currently elapsing. Experiments on memory for time intervals in animals have usually been restricted to three intervals or less (e.g. Church 1980). It is possible that the ability of birds to discriminate relative durations of intervals presented repeatedly in the same order has been seriously underestimated.

These results are in marked contrast to those of Lejeune & Richelle (1981), who compared the performance of barbary doves, pigeons and rats on simple FI schedules. They found that the performance of the doves was much worse than that of pigeons or rats, to the point that they claim that there was no evidence of post-reinforcement pausing at all for some of their dove subjects. There seems to be no obvious explanation of this difference in results: an attempt at replicating their results would offer the only resolution.

7.3.3 General Discussion

Two different, though not necessarily mutually exclusive forms of optimality analysis have been here applied to three different schedule types. In all cases some sensitivity to the relevant dimensions has been demonstrated.

It has been found that under concurrent VIVIE the subjects reduce the length of visits to the minor key, as predicted, but to such a great extent that under early conditions they were making these visits shorter than the COD and thereby totally ineffective. Rates of responding were widely variable.

Assessments of the magnitude of post reinforcement pausing in the literature (e.g. Toal & Leslie 1984; Harzem et al 1978; Catania & Reynolds 1968) have not attempted regressing the pauses onto the inter-reinforcement intervals. This method has shown up a slight but significant sensitivity of most subjects on VIVIA to the twelve component intervals of that schedule. Although the importance of such a slight effect must not be exaggerated, it does support the suggestion in chapter 6 that the regularity in these schedules provides an affordance for the subjects to interweave responding on the schedules (producing under-matching) and stabilize performance more quickly.

The degree of agreement with optimality predictions found here is insufficient to justify characterizing the behaviour as 'optimal'; this should not be surprising. On the other hand the behaviour is more sensitive to the

relevant dimensions than was found in the floor panel experiment in chapter 5. This suggests that neither regular nor irregular replenishing patches are totally alien to the phylogenetic history of the subjects used here. The following two chapters will attempt to analyse the mechanisms by which foraging in this task was accomplished.

CHAPTER 8
KEY PECK EXPERIMENT:
RULES OF THUMB

8.1 Introduction and Data Analysis Methods

Even if the optimality analysis in chapter 7 had accounted for all the variance in the relevant data, this would still have given us no idea how the doves were actually controlling their behaviour. Since, in fact, the comparison against optimality considerations found quite a loose fit, the question that remains unanswered is all the more substantial. How is the behaviour being controlled? The matching analysis in chapter 6 has provided some clues, but the number of degrees of freedom remaining unaccounted for in that form of analysis are too great for it to be satisfying. In this chapter I consider two theories that have been proposed to account for behaviour at a finer level of description. These are momentary maximising and melioration.

8.1.1 Momentary Maximising

Perhaps the most demanding theory of mechanism in the literature (in terms of processing requirements on the subject), is momentary maximising (Shimp 1966; 1969; Staddon, Hinson & Kram 1981; Hinson & Staddon 1981; 1983). The decision rule for momentary maximising is that the subject should respond on the alternative offering the highest momentary probability of reward. In order to do this it has to be aware of the type of schedule it is on, the nature of the relation determining reinforcement probability on the schedule, and then somehow perform an operation equivalent to estimating the probabilities of reinforcement on the options available. Since on FI and arithmetic VI schedules the reward probability functions are discontinuous, the question of momentary maximising reduces to the problem of timing explored in chapter 7. On concurrent exponential variable interval schedules (VIVIE), however, momentary maximising is not so simple. Firstly, it may not be obvious how momentary maximising differs from global optimisation discussed in chapters 4 & 7. In general an optimisation (or global maximization) account is looking for the pattern of responding that provides the greatest return of reinforcement on responding by the end of a session. A momentary maximising

strategy, on the other hand, requires that each response maximize the immediate probability of reward. In this particular context Houston & McNamara (1981, appendix 9) have demonstrated that changing over as soon as the immediate probability of reward is higher on the alternative schedule leads to briefer visits to the keys (more rapid switching) than would be found if subjects changed over so as to maximize pay-off in the long term.

Shimp's (1966) method for testing momentary maximising involved the production of tables of sequential probabilities "that are as difficult to interpret as they are inconvenient to compile" (Hinson & Staddon 1983). A major break-through for the analysis of momentary maximising was the recognition that the probability of reinforcement at any moment in time on an exponential VI schedule is fully determined by the time since last response to that schedule (Houston & McNamara 1981; Staddon, Hinson & Kram 1981). On an exponential VI then, the probability of reward at time **t** since last response to that schedule is given by:

Equation 8.1

$$P(Rf|t) = 1 - e^{-\lambda t}$$

where λ is the mean rate of reinforcement (= 1/ scheduled inter-reinforcement interval). By demanding that a subject respond on the alternative offering the highest probability of reinforcement, momentary maximising is stipulating that to respond on key 1 (say), the following condition be fulfilled:

Relation 8.2

$$P(Rf|t_1) > P(Rf|t_2)$$

$$1 - e^{-\lambda_1 t_1} > 1 - e^{-\lambda_2 t_2}$$

$$\lambda_1 t_1 > \lambda_2 t_2$$

$$\frac{t_1}{t_2} > \frac{\lambda_2}{\lambda_1} \text{ or } \frac{T_1}{T_2}$$

where $T = 1/\lambda$

This seems to imply the need for a substantial timing ability by the subject; the criterion requires the comparison of the ratio of times since last response on both keys with the ratio of the schedule values. However, if subjects responded at a constant rate, momentary maximising would define the number of responses that should be made to each schedule. Since, for typical schedule

values, minor key visits under momentary maximizing should be instantaneous (Houston & McNamara 1981) (as they would under optimisation), then a good approximation to momentary maximising would require only a suitable length of visit to the major key. This simplification would not appear to be available to the subjects since they do not respond at a constant rate (chapter 9), and it seems they are left with no alternative but to calculate the moment by moment probabilities on the basis of the times since the last response on each schedule, and then place their responses accordingly.

Hinson & Staddon (1981; 1983) argue that momentary maximising leads to a simple measure, **m** that can be calculated for each session of an experiment. **m** is the proportion of pecks placed on the key where the highest probability of reinforcement was to be found at that moment in time, weighted by the magnitude of the discrepancies between the probabilities of reward on the key chosen, compared to the other key. For each response in a session the absolute difference in probability between the two keys is calculated and assigned to a cell of a 2 by 2 matrix, according to whether it was on key 1 or key 2; and whether it was a correctly or incorrectly placed response. **m** is then the proportion of the sum of the probability differences that was due to correctly placed pecks.

Such a method has three drawbacks. Firstly, as mentioned above, one can find very good momentary maximising from fixed patterns of responding; Hinson & Staddon's method draws no distinction between different ways of attaining a value of **m**. Secondly, it is not possible in principle to compare probabilities of reward for the current key at this moment in time, with the probability at the same time on the alternative key, because the alternative key can only be reached after some time during which reinforcement is unobtainable. Although Hinson & Staddon's study did not program a changeover delay (COD), it is none the less apparent from inspection of their figures 3 and 4 (Hinson & Staddon 1983) that there was a minimum interval of about .7 to .8s between the last peck on one key and the first peck on the other. This interval cannot be simply incorporated in the calculation of the reward probability on the alternative key, because the question then becomes one of comparing a certain probability of reinforcement here and now, with another probability of reinforcement after a delay. This is a subjective choice and leads us into the realm of temporal discount functions (for a discussion of discount functions in foraging see Real 1980).

The third problem with this method is that it only assesses whether a peck was placed on the correct key, not whether it was right to peck at all. The variability in rate of responding makes some responses 'pointless', in the sense that they occur so soon after the previous responses (relative to the average rate of responding) that the increment in reinforcement probability on the present schedule is totally insignificant.

The method that will be adopted here is to plot the ratios of the probabilities of reinforcement on the two keys at the time of each response, against the time of response. This time series approach, though it does not yield a summary statistic for a benchmark of momentary maximising, does enable one to identify the relative efficacy, in terms of reinforcement probability, of responses at different points in time.

Data will also be presented in the form of Hinson & Staddon's (1981;1983) 'Clock-space'. The clockspace was developed to display the temporal patterning of responses so that the degree to which the data supported momentary maximising could be easily seen. Since the critical factor (ignoring COD) in momentary maximising is the time since last response to each key, the clock space plots for each response, time since last key 1 peck (t_1) on the x-axis, against time since last key 2 response (t_2) on the y-axis. Figure 8.1 shows three sample points plotted in the clockspace. Let us consider for simplicity the case where the schedule values are equal. Point A represents a response that occurs 7 secs after the last key 1 response, and 5 secs after the last key 2 response. Since $t_1 > t_2$ this response will momentary maximise only if it is a key 1 response. Similarly point B, where $t_2 > t_1$, will only be momentary maximising if it is on key 2. Point C, and all other points on the line $t_1 = t_2$, represents the transition between the states where key 1 responding is required for momentary maximising (below the line), and key 2 responding is correct (above the line). Hinson & Staddon call this the 'switching line', and from relation 8.2, it can be seen that its gradient is T_1/T_2.

Figure 8.2 (from Hinson & Staddon 1983 figure 2) shows how three simulated patterns of responding appear in the clock space. The top row shows responding that is indifferent to momentary probabilities of reward (but that matches). Responding which fulfils the momentary maximising criterion – always responding on the key offering the highest momentary probability of reward is shown in the middle row, while the lowest panels show momentary

figure 8.1

figure 8.2

Key 1   **Key 2**

t2

t2
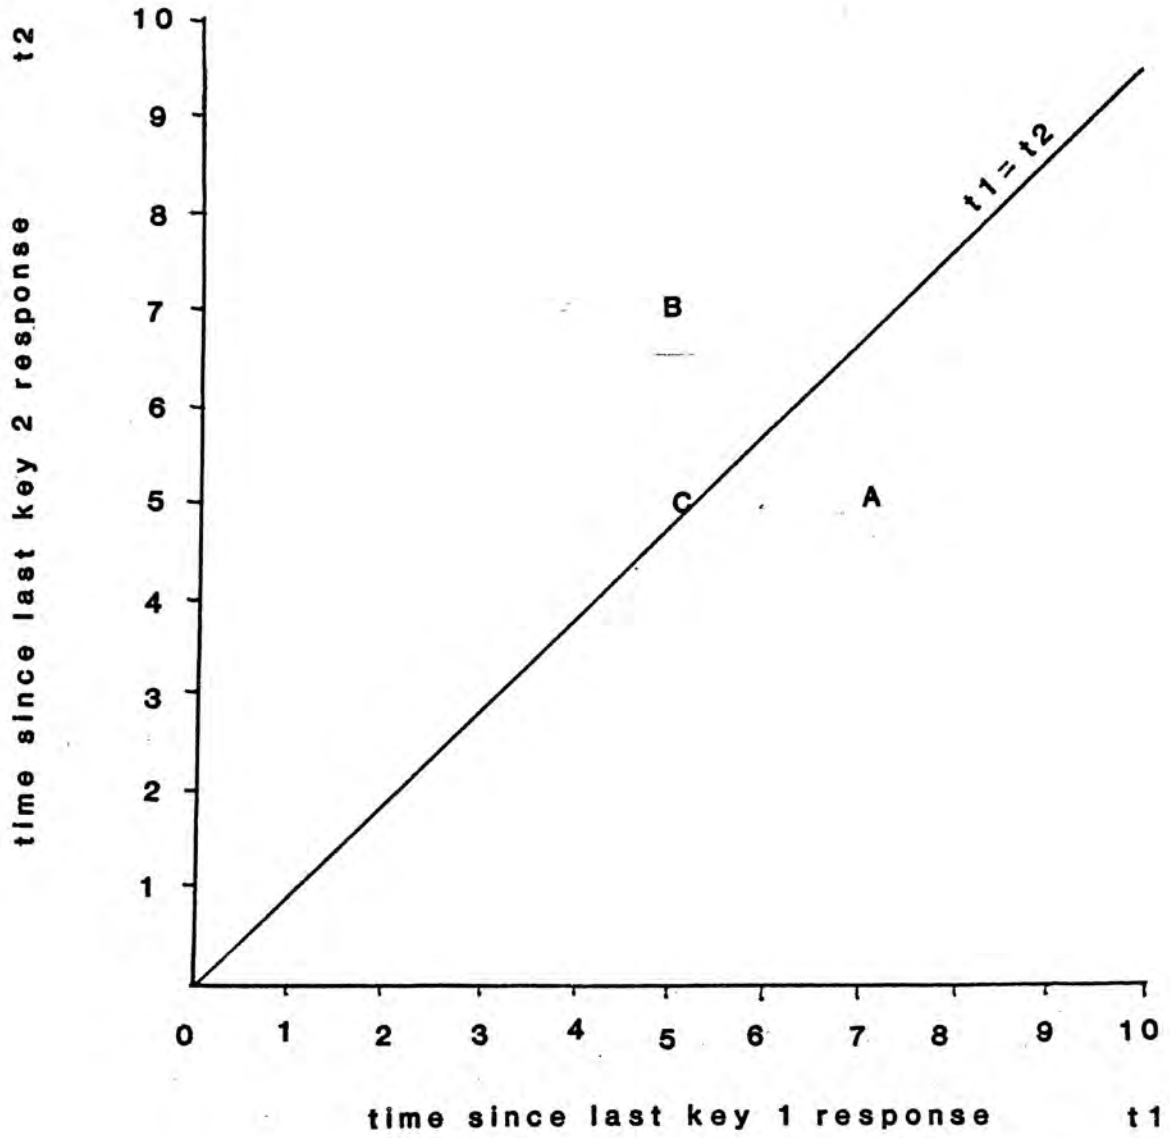
t2
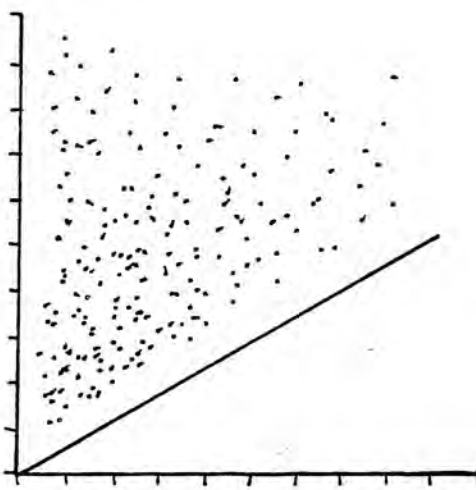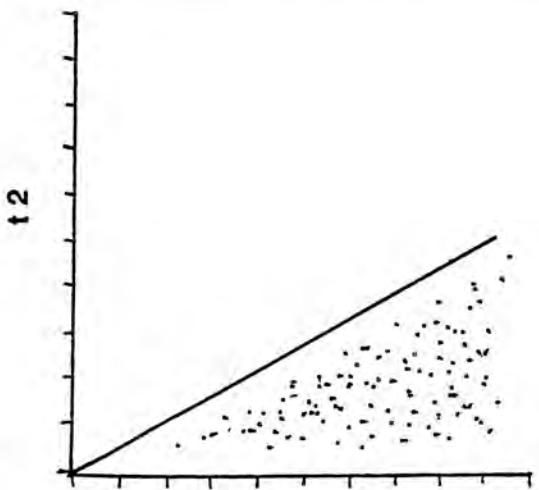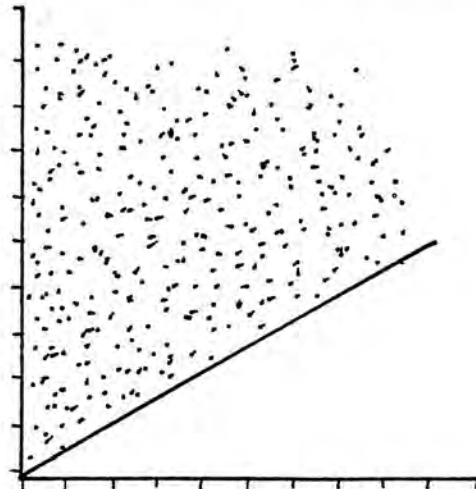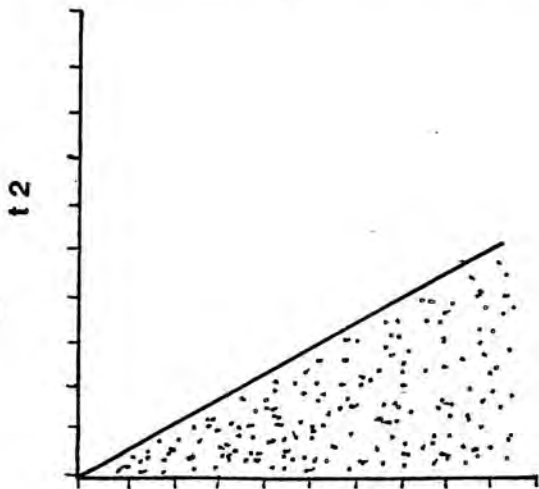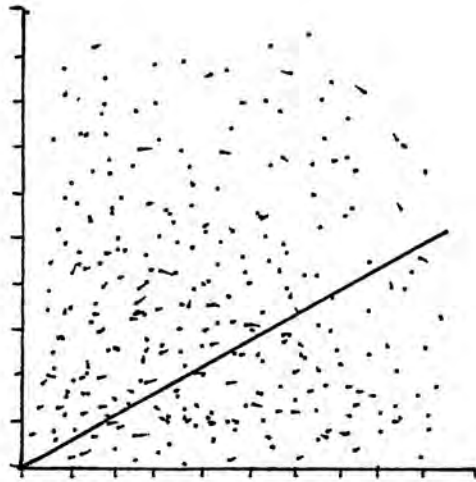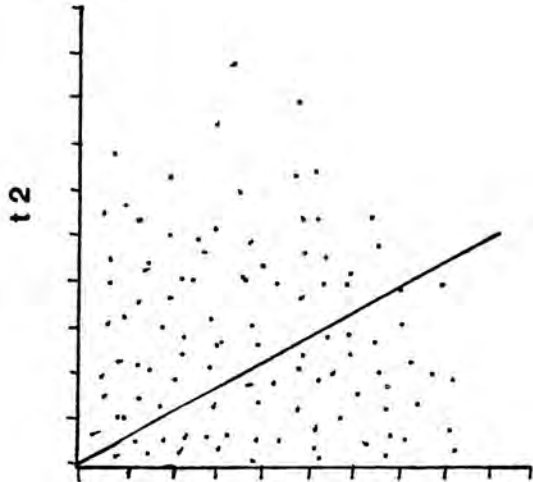
t1   t1

maximising patterns, allowing a minimum time spent in changing over from one key to the other. It should also be noted that responding that always occurs at a constant ratio of reward probabilities can only arise through a constant ratio of times since last response to each key, and so will fall along a line through the origin of the clockspace. Pecks at a constant interpeck interval will form a line parallel to the axis representing time since last peck to the alternative key.

8.1.2 Melioration

Recognizing that matching did not identify a process, but just a molar average result, Herrnstein & Vaughan (1980; Vaughan 1981; Herrnstein 1982) proposed a process of 'melioration' (meaning 'bettering') that would approach matching as a molar result, and maintain the philosophy of matching transposed to a more mechanistic level. A similar theory has also been proposed within the ethological literature (Harley 1981; Harley & Maynard Smith 1983; Lester 1984) as a process to approximate optimal foraging.

In Herrnstein & Vaughan's formulation, a parameter R_D is identified, which is the difference in local rates of reinforcement.

Equation 8.3

$$R_D = \frac{R_1}{t_1} - \frac{R_2}{t_2}$$

where R and t are in this case the local numbers of reinforcements and times spent responding on each of 2 alternative schedules. "Melioration requires a subject to redistribute responses associated with lower to those with higher local rates of reinforcement. Equilibrium is attained when the alternatives reinforce at equal rates, which is formally equivalent to matching" (Herrnstein & Vaughan 1980). That is to say R_D should be minimized. Aside from predicting matching as a molar finding, melioration is so vaguely stated that it is not at all clear what its local level predictions are.

The theories of 'Dynamic Matching' (Lester 1984) and 'Relative Pay-off Sum' (Harley 1981; Harley & Maynard Smith 1983) expand the same idea as melioration, but with greater mathematical rigour. 'Dynamic matching' is Lester's simplification of Harley's 'relative pay-off sum' model, and forms a link between the sketchy formalization of the melioration account, and the elaborate multi-parameter 'relative pay-off sum' theory.

Harley suggests that an animal in an environment consisting of discrete patches, allocates its time to each patch according to the proportion of rewards it obtained in each patch in the interval since last decision:

Equation 8.4

$$\frac{T_{i,n+1}}{\sum_{i=1}^{k} T_{i,n+1}} = \frac{R_{i,n}}{\sum_{i=1}^{k} R_{i,n}}$$

where $T_{i,n+1}$ is the time spent in patch i during the n+1th time bin; $R_{i,n}$ is the number of rewards obtained in patch i during time bin n, and k is the number of patches.

In a two patch environment this implies that the ratio of time spent in the two patches in the (n+1)th time bin is equal to the ratio of rewards obtained in each patch in the previous time bin.

Equation 8.5

$$\frac{T_{1,n+1}}{T_{2,n+1}} = \frac{R_{1,n}}{R_{2,n}}$$

It can be seen that as R_0 in equation 8.3 tends to zero, so that equation and equation 8.5 become similar. There is, however, a difference. Lester's dynamic matching rule contains subscripts that distinguish responding in the future, from the rewards in the past that influenced it. Equation 8.3 contains no such subscripts. From Herrnstein & Vaughan's written explanations of the theory it would appear that this omission is an oversight. The section quoted above from Herrnstein & Vaughan (1980) is certainly expressing the idea that future responding is a consequence of past reinforcements. Henceforth in this account I will assume melioration to be identical to Lester's dynamic matching rule.

Herrnstein & Vaughan's failure to mathematically characterize their model correctly has led to confusion as to the predictions of the theory. Real & Dreyfus (1985), for example, found that local response proportion on concurrent VIVI oscillated, rather than stabilizing to the session average value, and took this to refute Herrnstein & Vaughan's statement that the subject should equalize response and reinforcement proportions. Lester (1984) has shown however that the model of equations 8.4 and 8.5 predicts oscillation of preference as a consequence of present response proportion being based on reward proportion in the past. If the patches deplete, then oscillation in response proportion would continue indefinitely. Although, on the average, the

VI schedules used by Real & Dreyfus would not change in their payout, on smaller timescales the availability of prey would certainly be expected to vary substantially due to the stochastic nature of the schedule.

Harley's relative pay-off sum model can be obtained from Lester's dynamic matching rule by substituting

$$\sum_{j=1}^{n} b^{n-j} R_{i,j} \quad \text{where } 0 < b < 1$$

for $R_{i,n}$ in equation 8.4 (Lester 1984). b is a weighting parameter determining the effect of more distant rewards (at lag j) on behaviour.

The right hand side of equation 8.4 thereby becomes:

$$\frac{\sum_{j=1}^{n} b^{n-j} R_{i,j}}{\sum_{i=1}^{k} \sum_{j=1}^{n} b^{n-j} R_{i,j}}$$

That is to say, response proportion is now determined not just by reward proportion in the previous time bin, but by the weighted sum of rewards over all the previous time bins. These are weighted so that the most recent time bins have the most substantial impact on response proportion.

To make the model stochastic rather than deterministic, Harley considers the left hand side to be the probability of a patch i choice, $P_{i,n+1}$, rather than the proportion of patch i choices: mathematically these are equivalent.

The final elaboration that Harley introduces into the model is to include a quantity e_i representing the subjects prior expectation of gain in patch i. These are constants that sum onto the effect of immediately previous feeding experience. The complete model can now be expressed:

Equation 8.6

$$P_{i,n+1} = \frac{e_i + \sum_{j=1}^{n} b^{n-j} R_{i,j}}{\sum_{i=1}^{k} \left(e_i + \sum_{j=1}^{n} b^{n-j} R_{i,j} \right)}$$

Lester shows that although Harley's model is more elaborate, its parameter independent predictions are in line with those of the simpler melioration/dynamic matching rule.

It might appear that a melioration type rule cannot be tested in a free operant environment without difficult, perhaps critical assumptions about the

interval between decision points. The general principle of whether the distribution of reinforcements is influencing the moment to moment distribution of responding in the future can be tested by the method of cross-correlational analysis. Cross-correlation involves the correlation of each point in a time series with points in another time series. These may be points in the future (positive lags) or past (negative lags), and by exploring the cross-correlation between two time series at a variety of lags we can see whether one series is tending to recapitulate the same pattern as the other. This method only has validity if each series does not have a significant tendency to oscillate or cycle in its values. If the two series cycle, then significant cross-correlation may occur just because the cyclings are accidentally in phase at certain points. Autocorrelation provides a technique to investigate cyclings in a series. In autocorrelational analysis each item in a time series is correlated with its neighbour a certain number of items away (these are best calculated after partialling out the effect of autocorrelations at intermediary lags (partial autocorrelations)). If there are significant autocorrelations this means that a series is tending to repeat itself (Chatfield 1975).

Following Real & Dreyfus (1985), we need firstly to identify the minimum period over which a proportion of responding can be defined. This is, of course, a cycle of a visit to one key, followed by a visit to the other key: i.e. two consecutive change overs. We can then add together the pecks and reinforcements occurring in different numbers of consecutive cycles and calculate response and reinforcement proportions for different levels of aggregation. In the present study aggregates of up to 40 cycles were investigated. Autocorrelational analysis (BMDP Program 2T; Dixon et al 1983) was employed to investigate cyclings within each series; and then cross-correlations (ibid) were used to identify any tendency for changes in one series to be reflected at some point in the other series.

8.2 Results

8.2.1 Momentary Maximising

8.2.1.1 Time Series Analysis of Reward Probability Ratios

For reasons that will become clear, I firstly consider only those pecks that initiated visits to a key. Figure 8.3 shows the ratios of the reward probability on key 1, over the reward probability on key 2 (calculated from equation 8.1) at the time of changing over, for ten minute segments from the first session of the first condition, and the last sessions of all VIVIE conditions for subject 65. Figures 8.4 and 8.5 present the same data for subjects 106 and 149. Logs have been taken of these ratios to make them symmetrical. Circles identify key 1 responses; asterisks key 2. It can be seen that the reward probability ratios are consistently on the correct side of the line representing equality of probabilities, and therefore that the reward probability ratios at change-over are different for the two keys.

Figures 8.6 - 8.8 contain all the responses occurring in 5 minute segments from the same sessions as figures 8.3 to 8.5. This set of graphs is characteristic of all sessions in that most responses fall the wrong side of the line of equality of reinforcement probabilities. One subset of responses however, the responses that were shown in figures 8.3- 8.5, the first responses to a new key (identified in figures 8.6 - 8.8 by super-imposed squares), always occurs on the correct side of the equality line. Noting that this study programmed a COD but that responses within it were not recorded, it can be seen that subjects are responding too rapidly for too long a time once on a key (such that the probability of reward on the other key was higher); but that by the time they do change over they always do so when the local probabilities of reward strongly favour such an action. This is what one would expect given the unquantifiable, but discouraging, impact of a COD on crossing over.
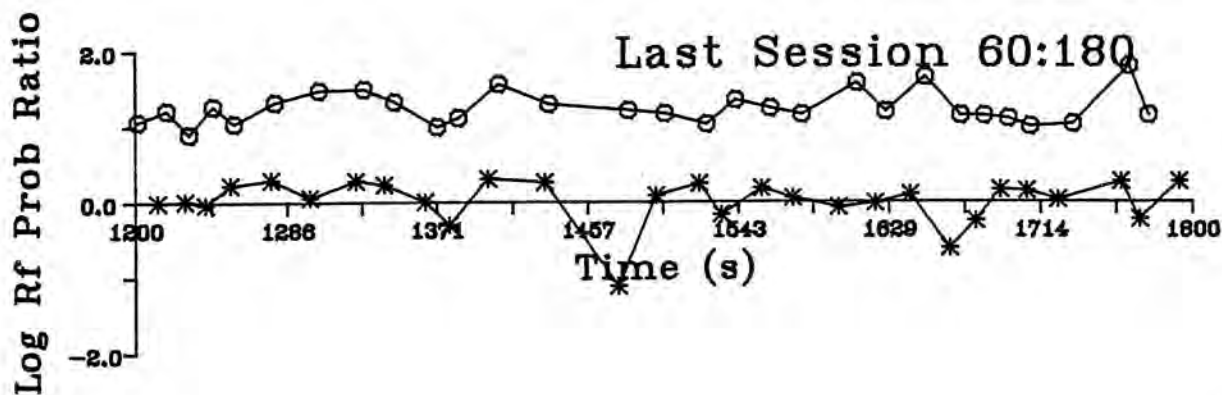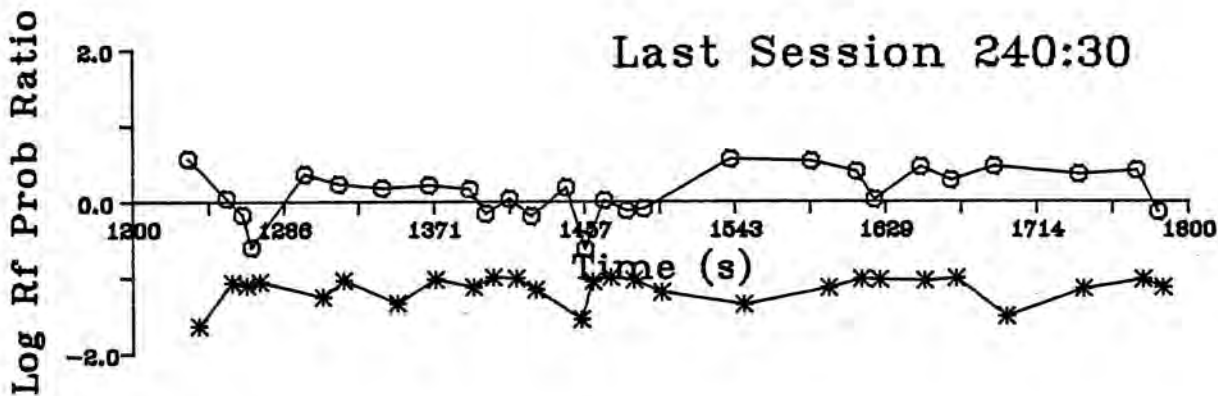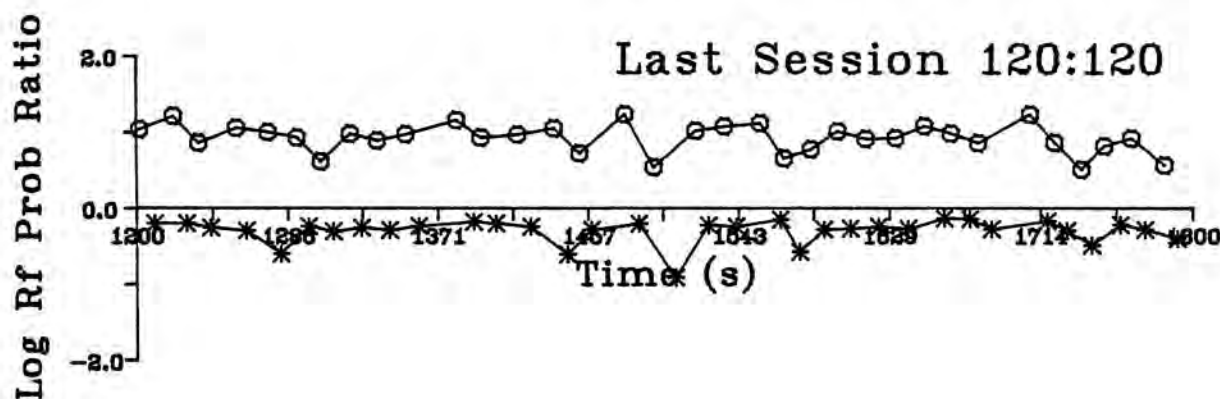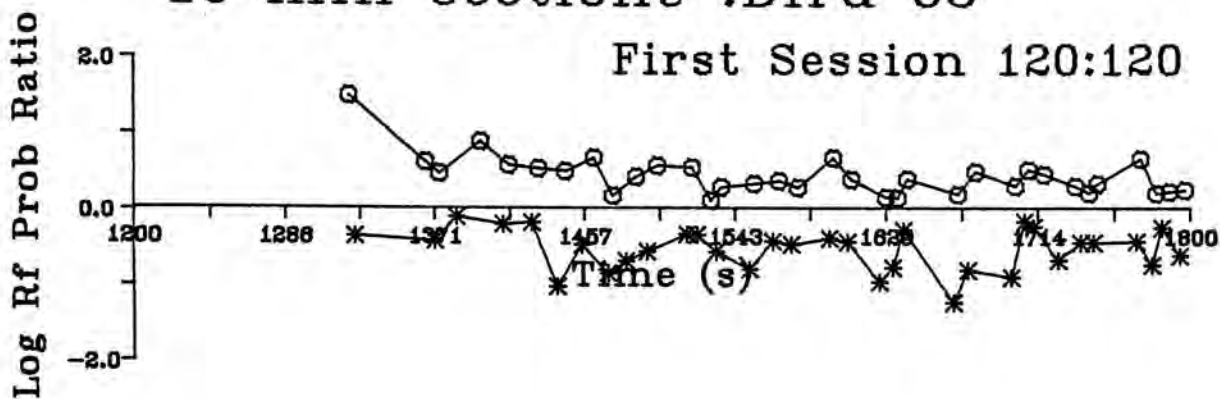
figure 8.3

116

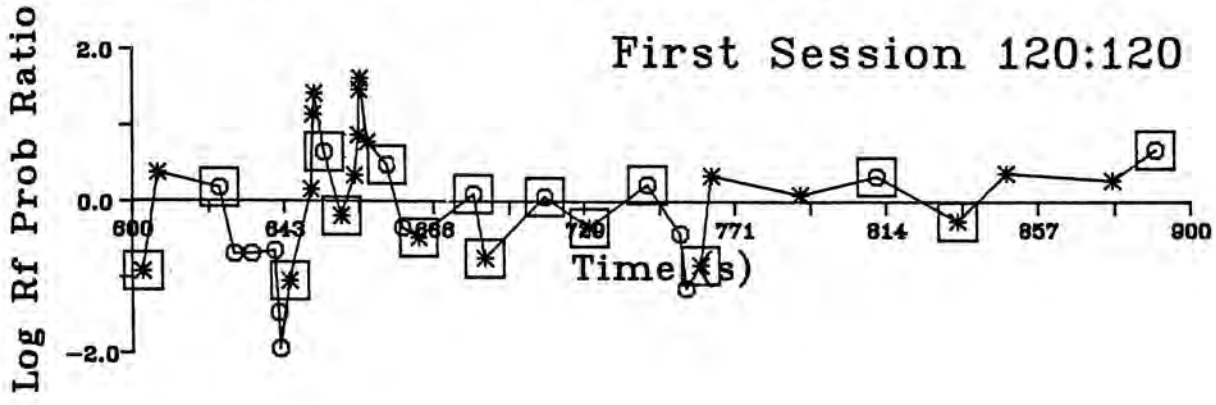# Log Reward Prob. Ratio for
CO responses only
10 min sections :Bird 65

First Session 120:120

Last Session 120:120

Last Session 240:30

Last Session 60:180

figure 8.4   1 17

Log Reward Prob. Ratio for
CO responses only
10 min sections :Bird 106

First Session 120:120

Last Session 120:120

Last Session 240:30

Last Session 60:180
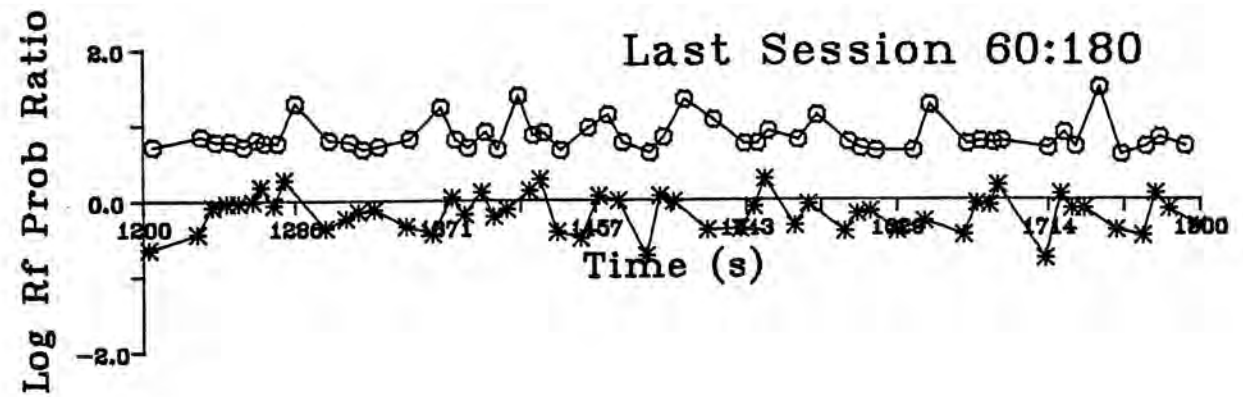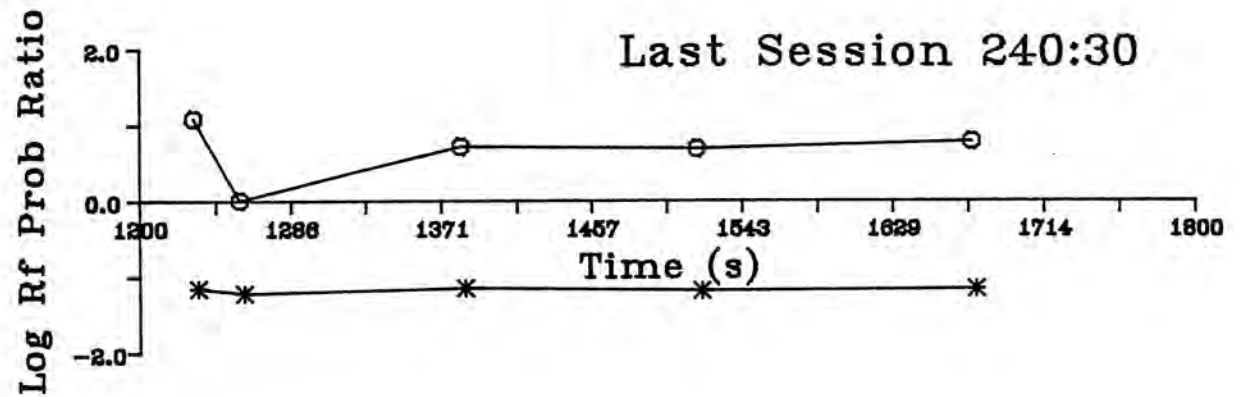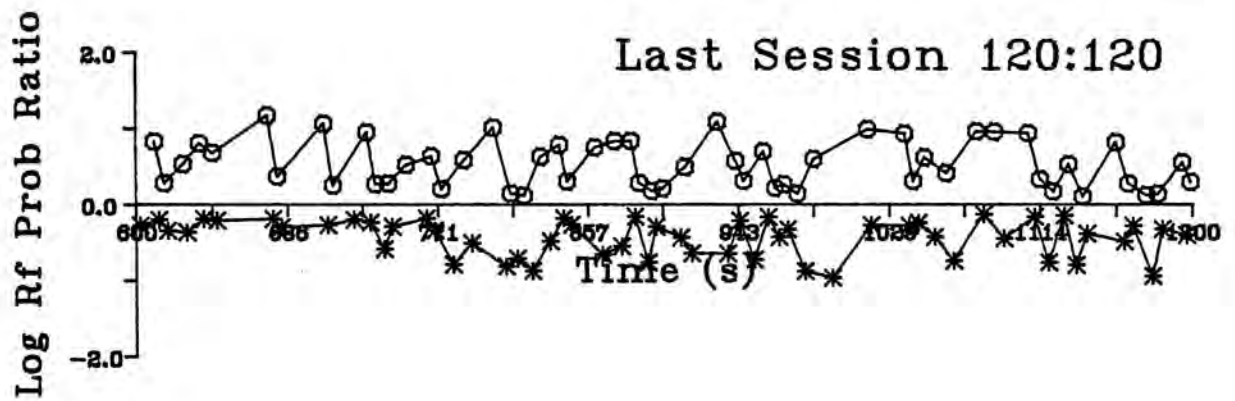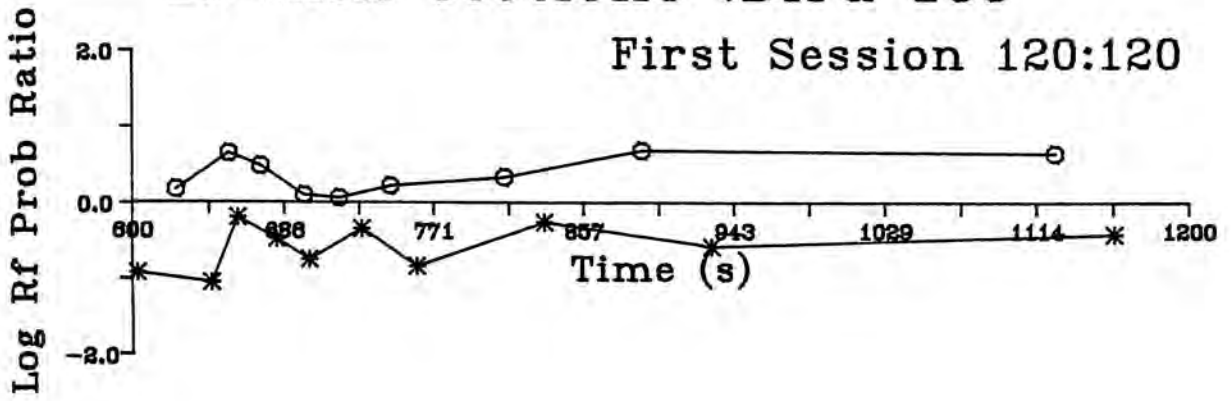
figure 3.5   1 1 3

Log Reward Prob. Ratio for CO responses only
10 min sections :Bird 149

First Session 120:120 insufficient data
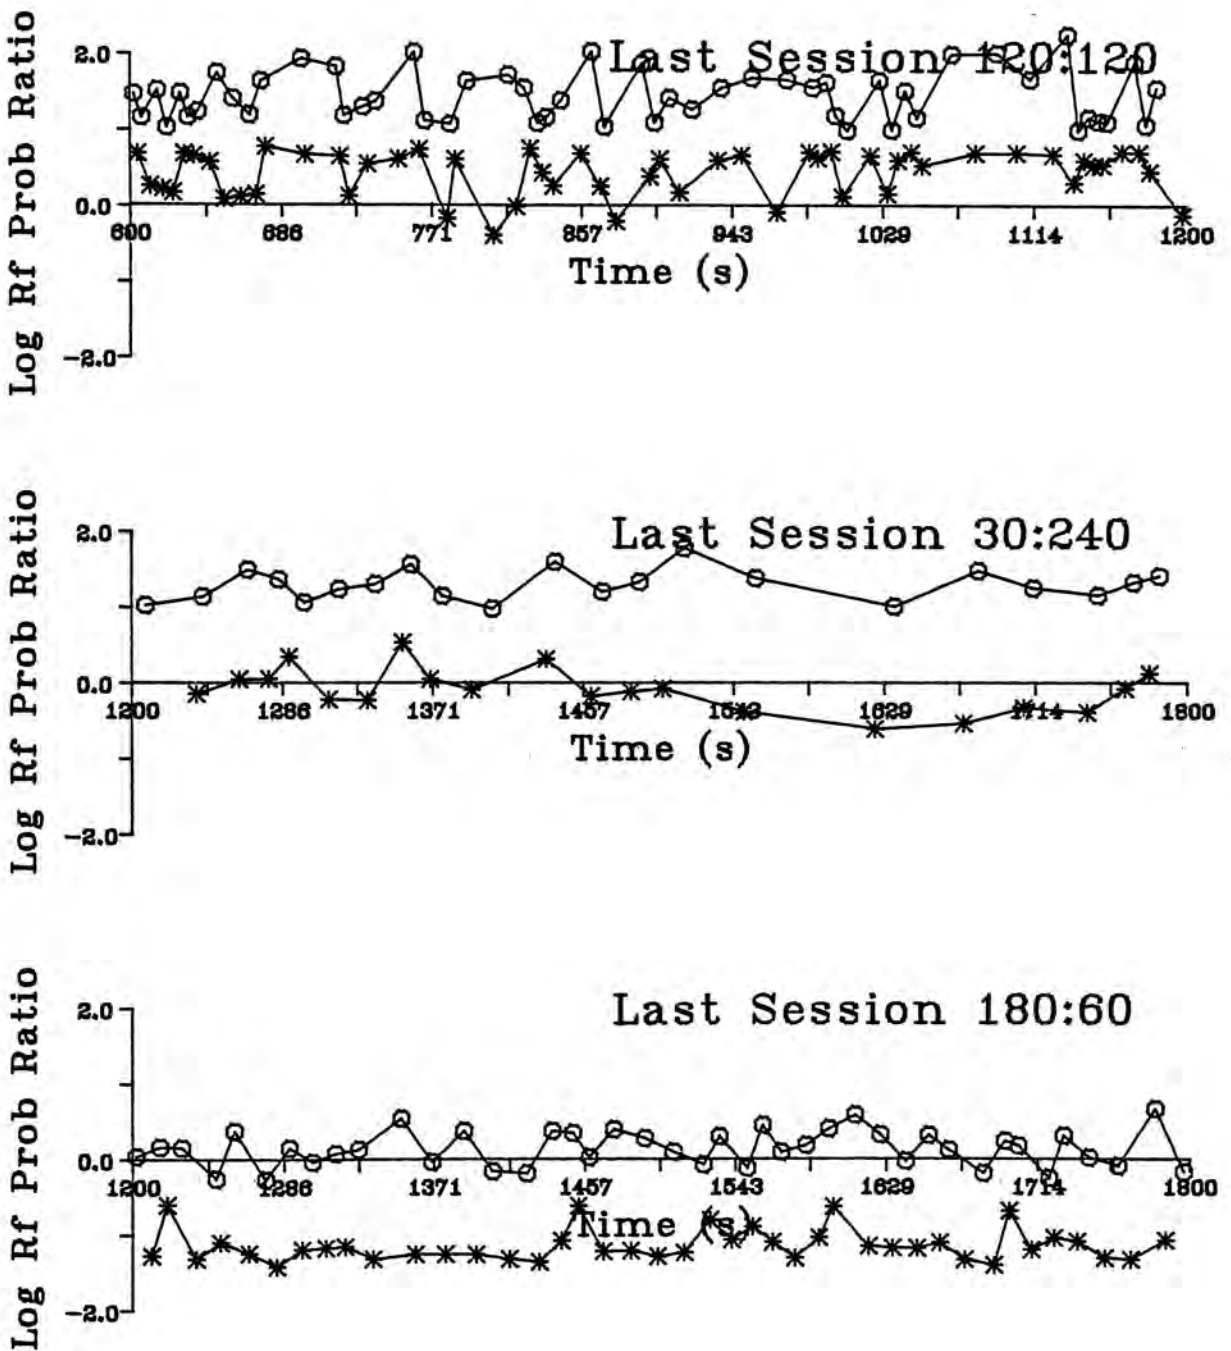
Last Session 120:120
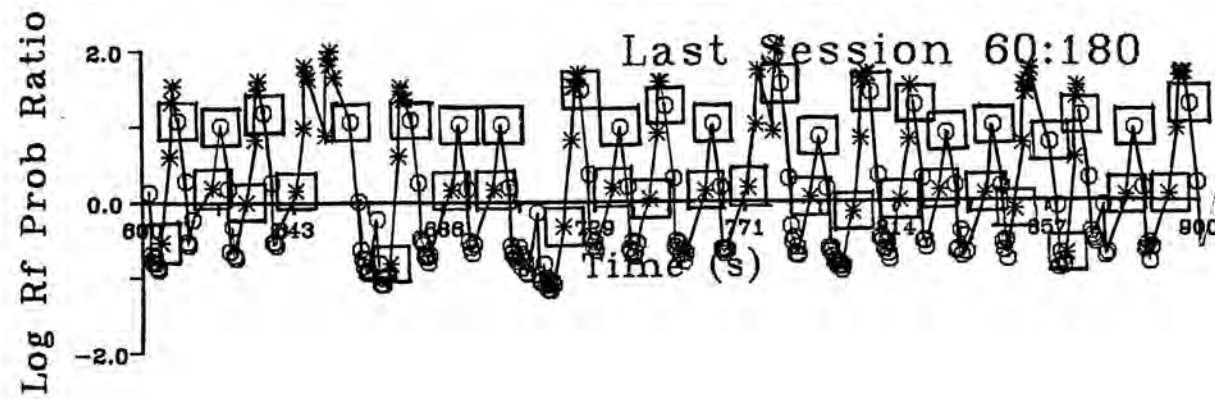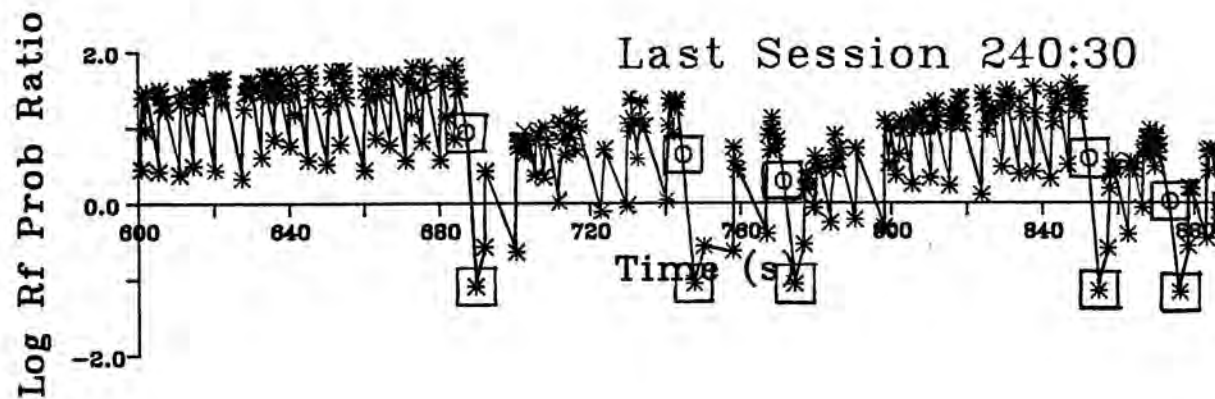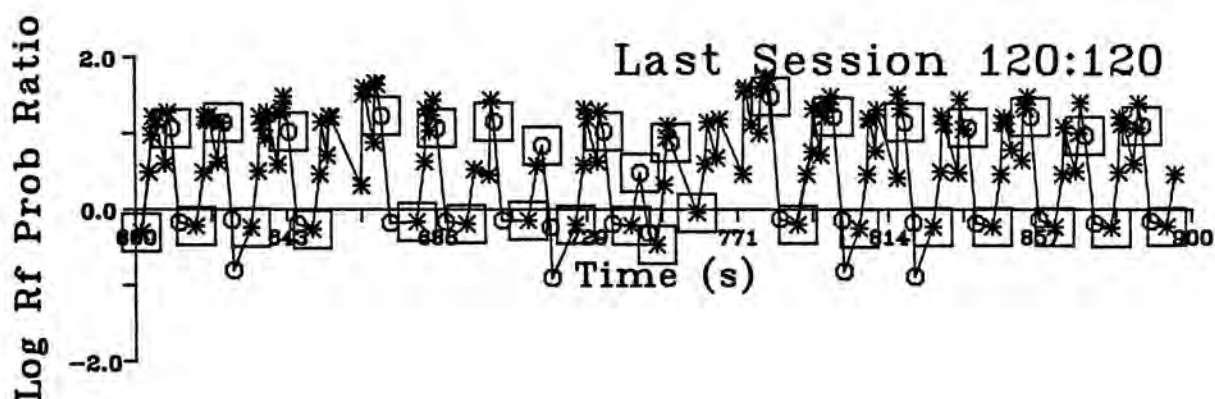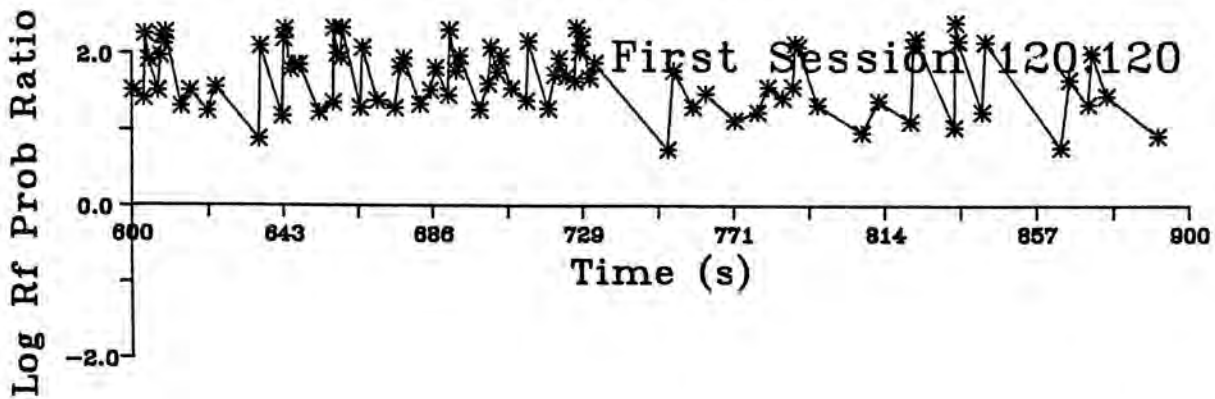
Last Session 30:240

Last Session 180:60

figure 8.6

Log Reward Prob. Ratio for
5 min sections :Bird 65

First Session 120×120

Last Session 120:120

Last Session 240:30

Last Session 60:180

figure 8.7

Log Reward Prob. Ratio for
5 min sections :Bird 106

First Session 120:120

Last Session 120:120

First Session 240:30

First Session 60:180
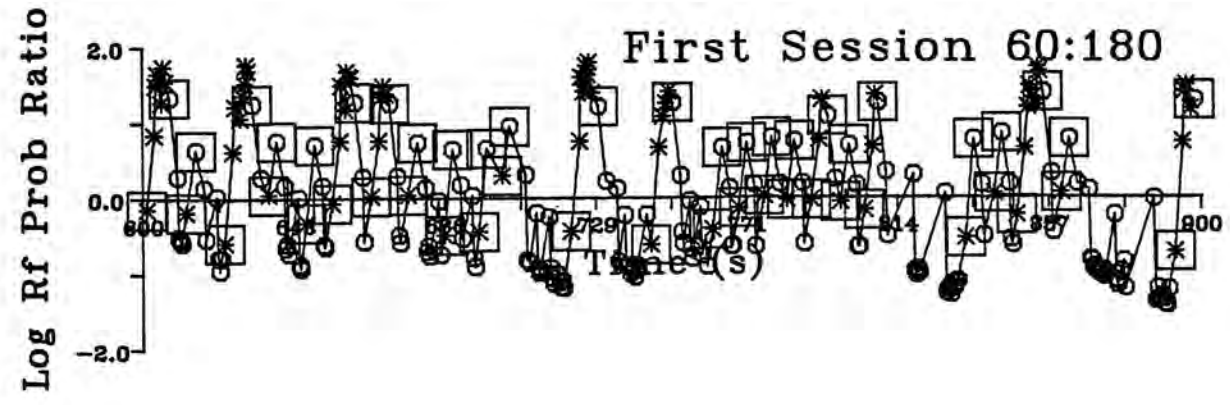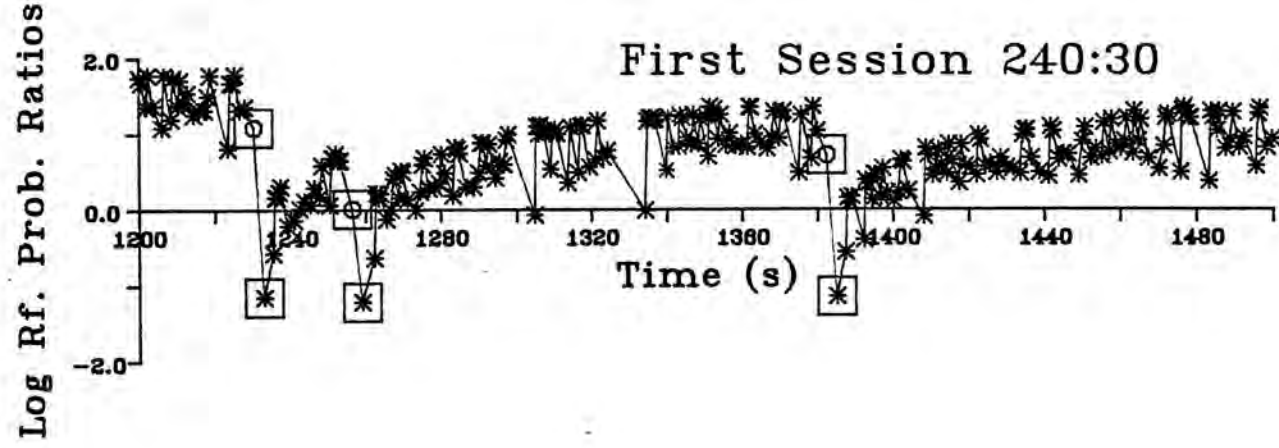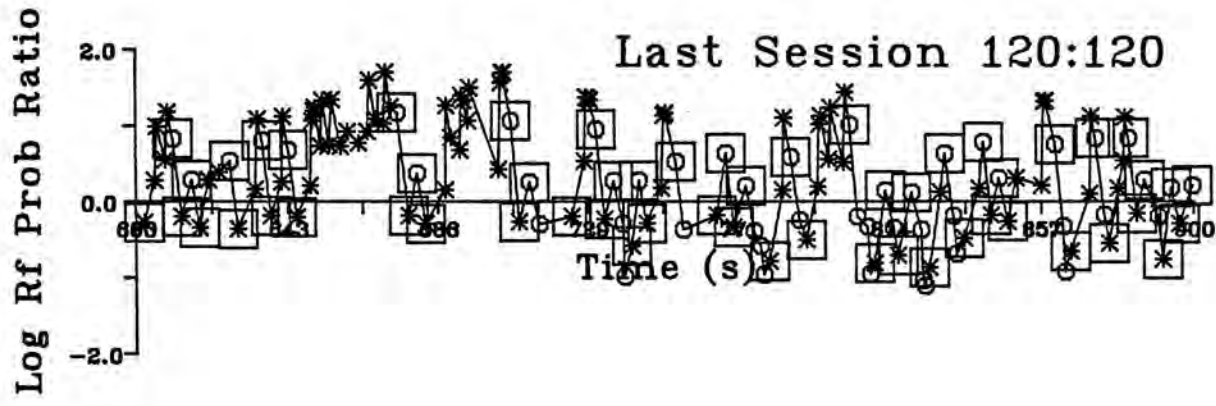
figure 8.8

Log Reward Prob. Ratio for
5 min sections :Bird 149
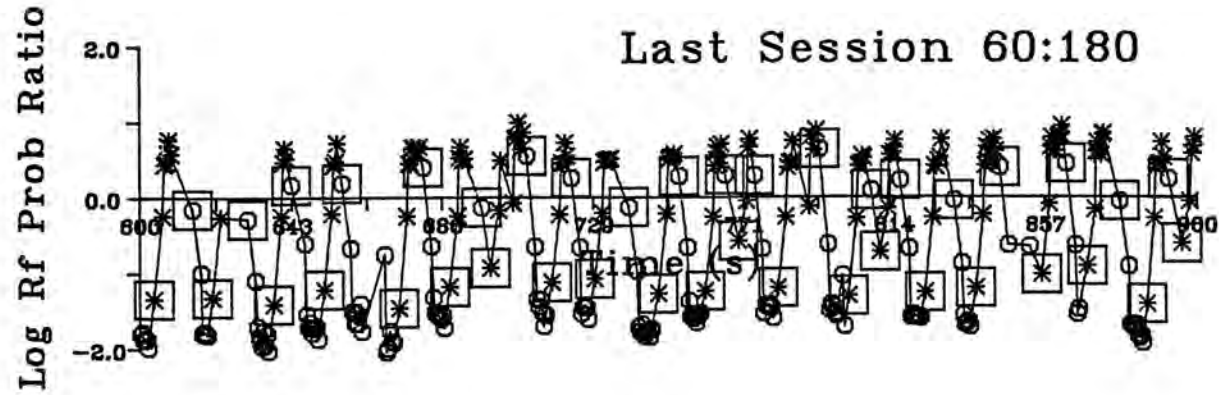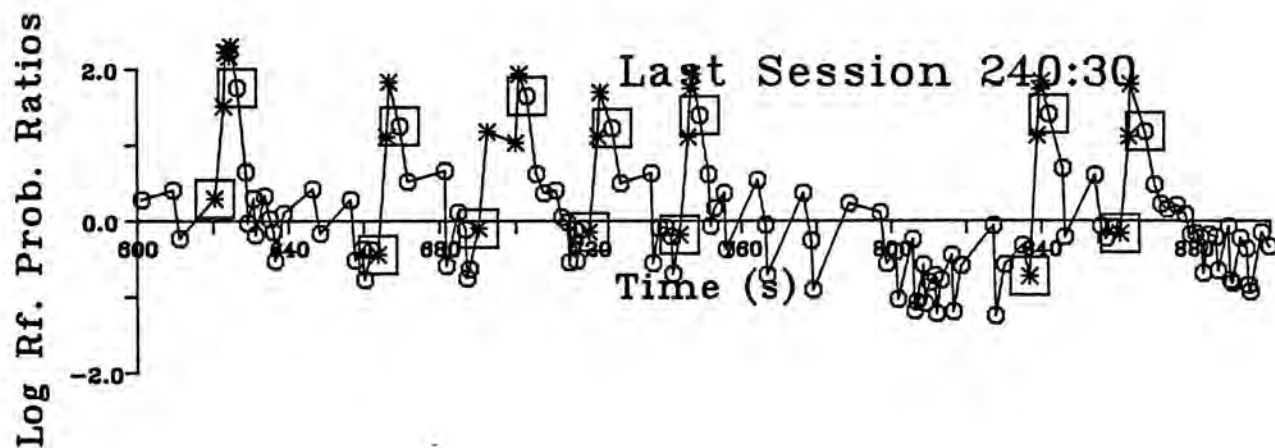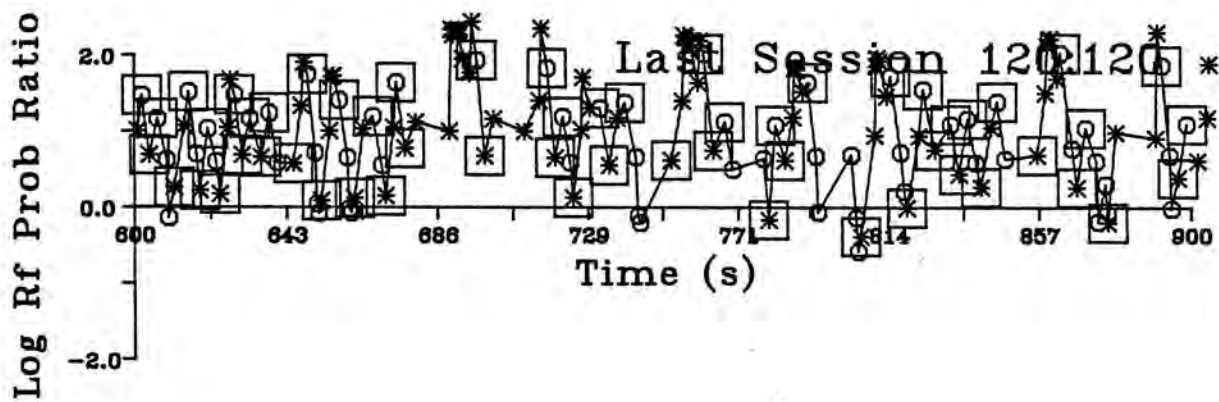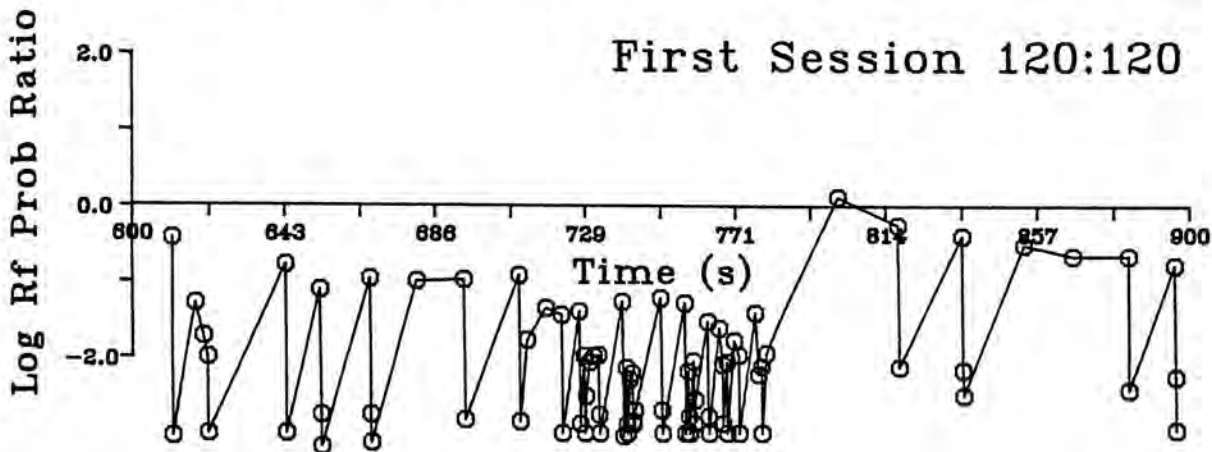
8.2.1.2 Clockspace

Figures 8.9 - 8.11 contain, for the same sessions as figures 8.3 to 8.5, results presented in the clockspace format. Remember that, for momentary maximising, key 1 responses should fall below the switching line, and key 2 responses above. Table 8.1 contains the mean percentage of responses on the correct side of the switching line, for the first five sessions of VIVIE 120:120; and the last five sessions of all VIVIE conditions.

Table 8.1:
Percentage responses on the correct side of the switching line.
Mean of five sessions

Condition	Key	Bird- 65	106	149
120:120	1	36.2	41	23.6
Begin	2	32.8	27.8	54.4
120:120	Major	43	39.6	37.6
End	Minor	11.6	25	27
240:30	Major	12.2	2.8	46
End	Minor	38.8	84.6	16.8
60:180	Major	29.2	24.6	27.4
End	Minor	8.6	30.4	7.6

As could be predicted from figures 8.6 to 8.8 these proportions are small (with the exception of the anomalous case of bird 106, minor key under VIVIE 240:30). Figures 8.12 - 8.14, following figures 8.3 - 8.5, present in clockspace form just the first responses to a new key (here, because there is so little overlap keys 1 and 2 are plotted on the same graphs, as circles and asterisks, respectively). With the exception of the case with the most widely differing schedule values (VIVIE 240:30), these responses are far more often on the correct side of the switching line. Note that whereas from figure 8.3 to 8.5 we can only identify that the first pecks to a key are at a fairly constant reward probability ratio, from figures 8.12 - 8.14 it can be seen that this is achieved even though time since last peck to each key varies substantially.

figure 8.9

123.

Clockspace representation
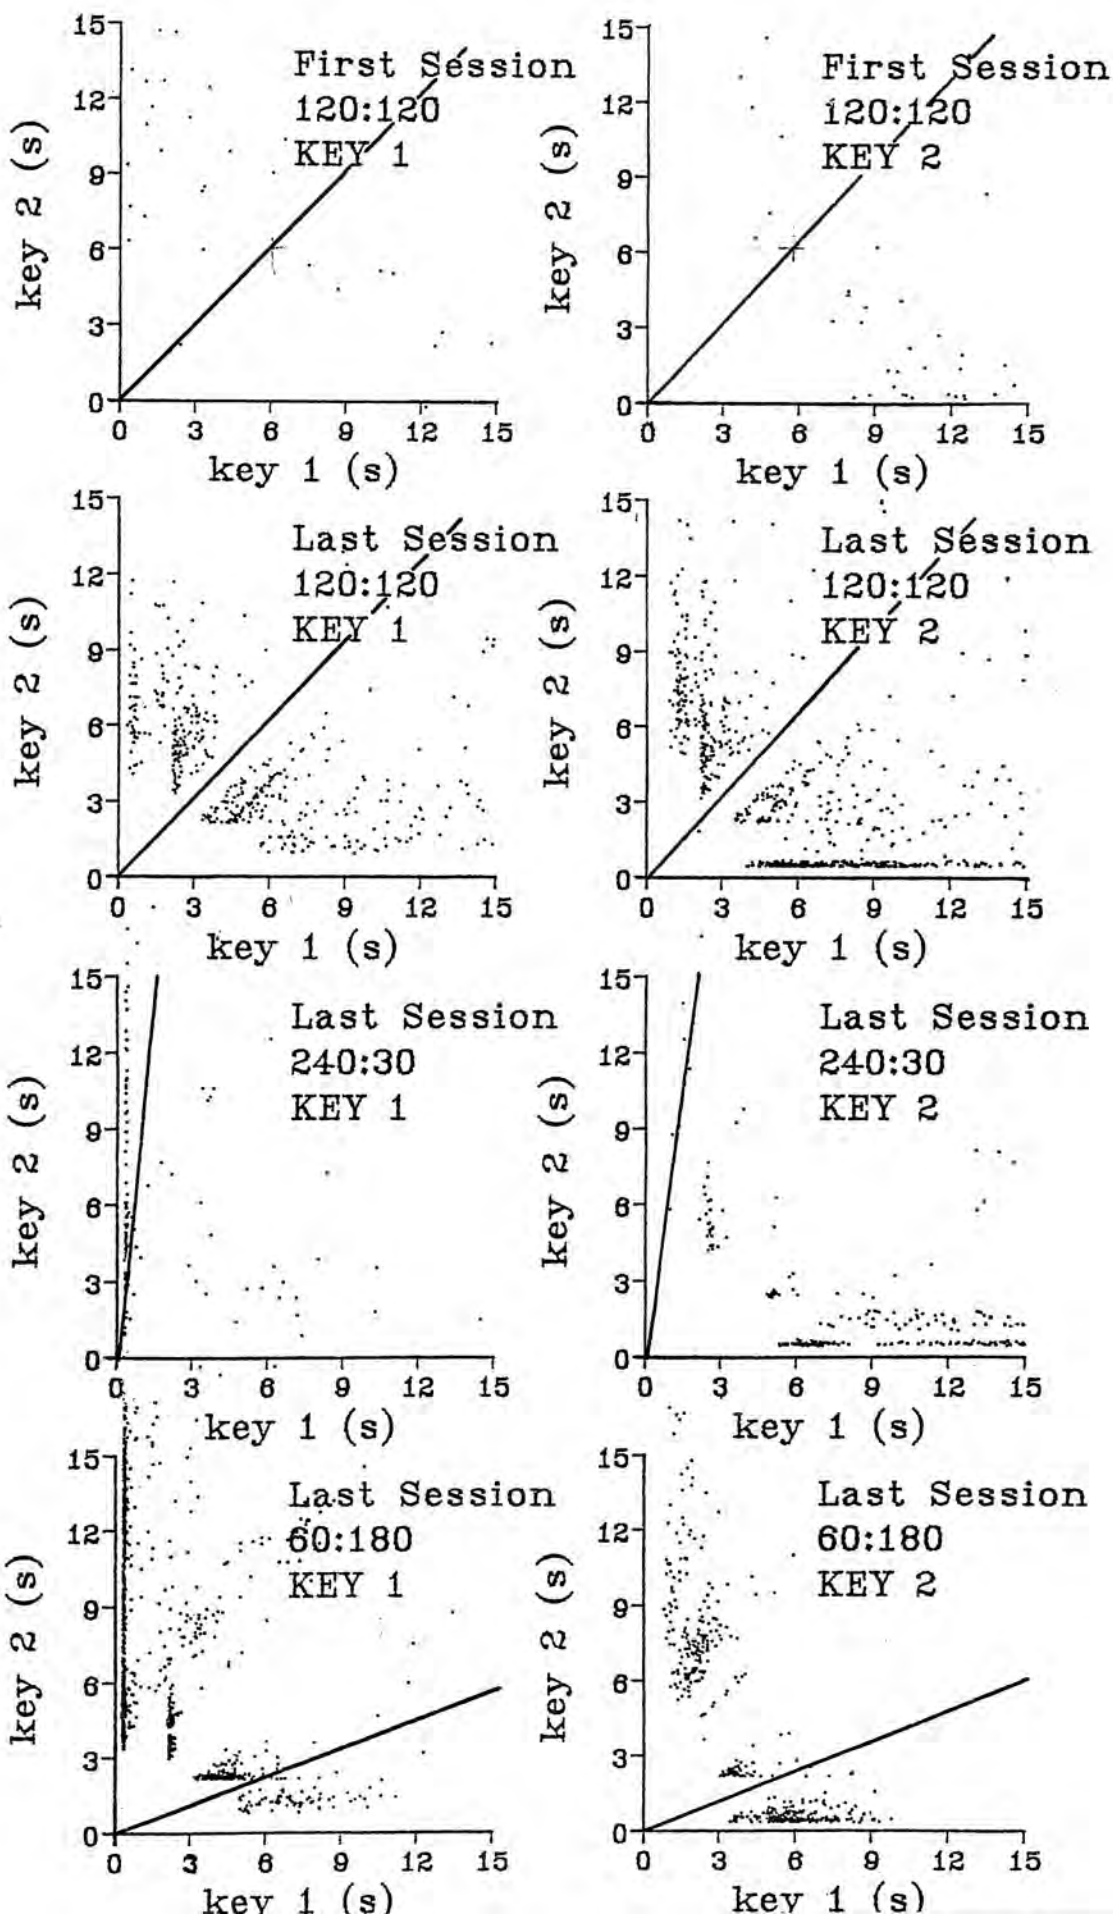# Bird 65

figure 8.10   124
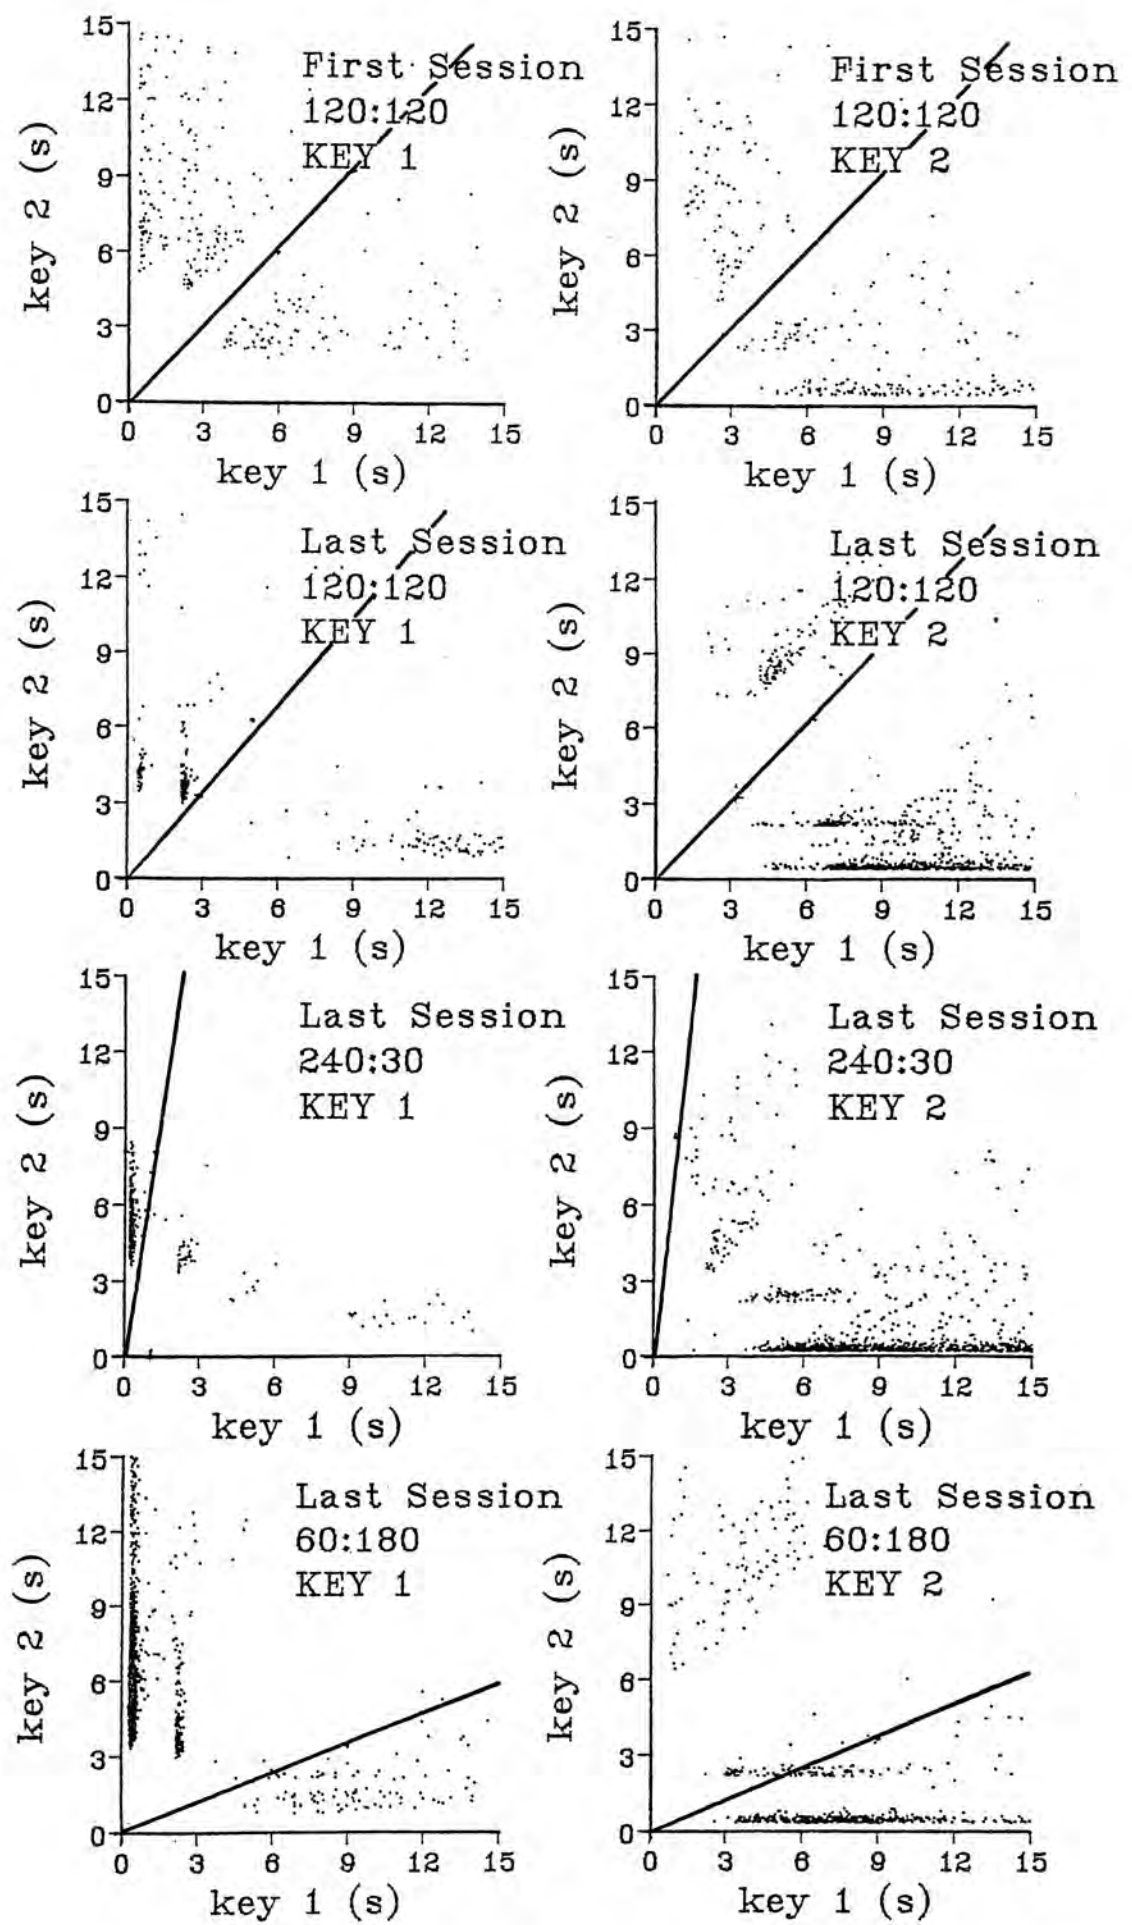
Clockspace representation
Bird 106

figure 8.11
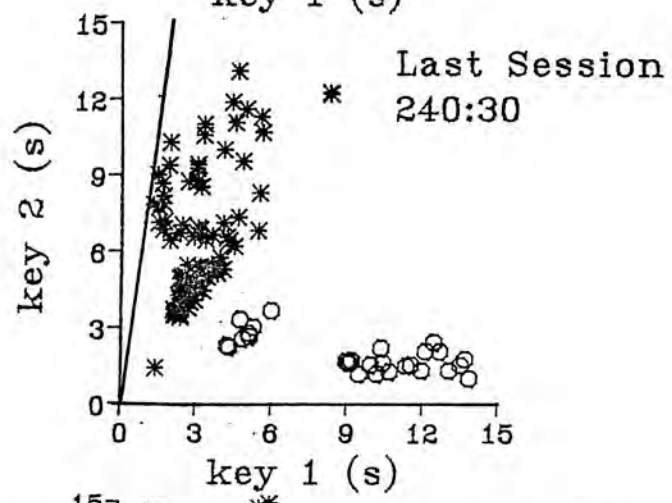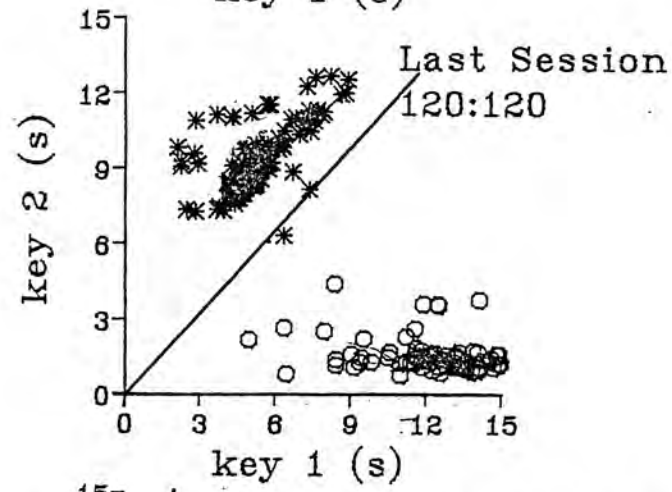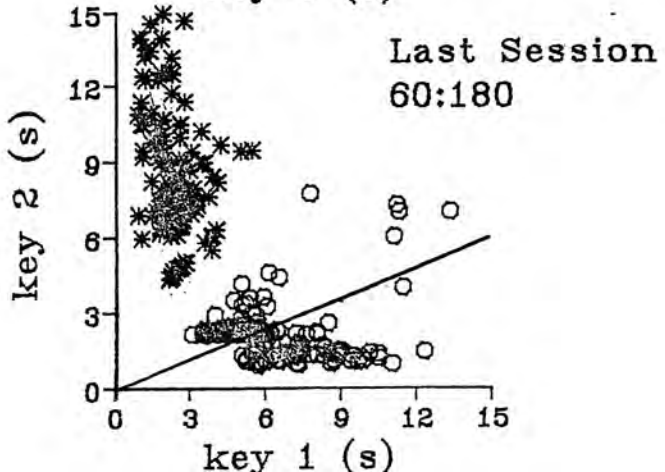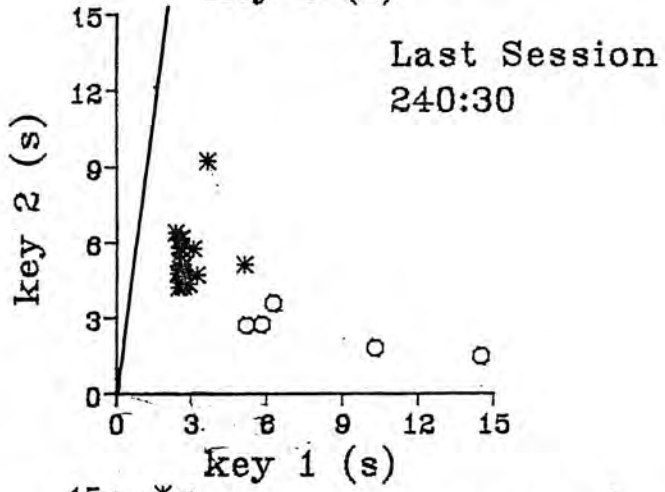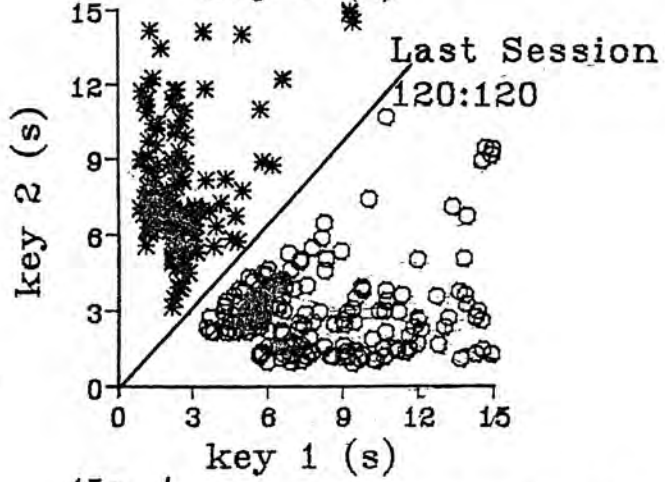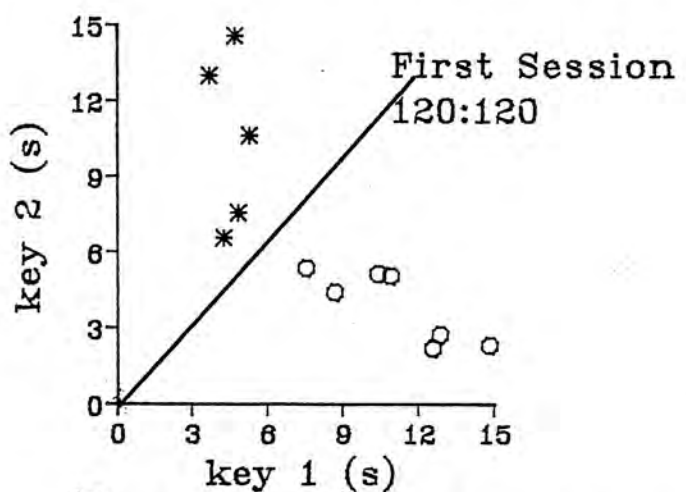# Clockspace representation
Bird 149
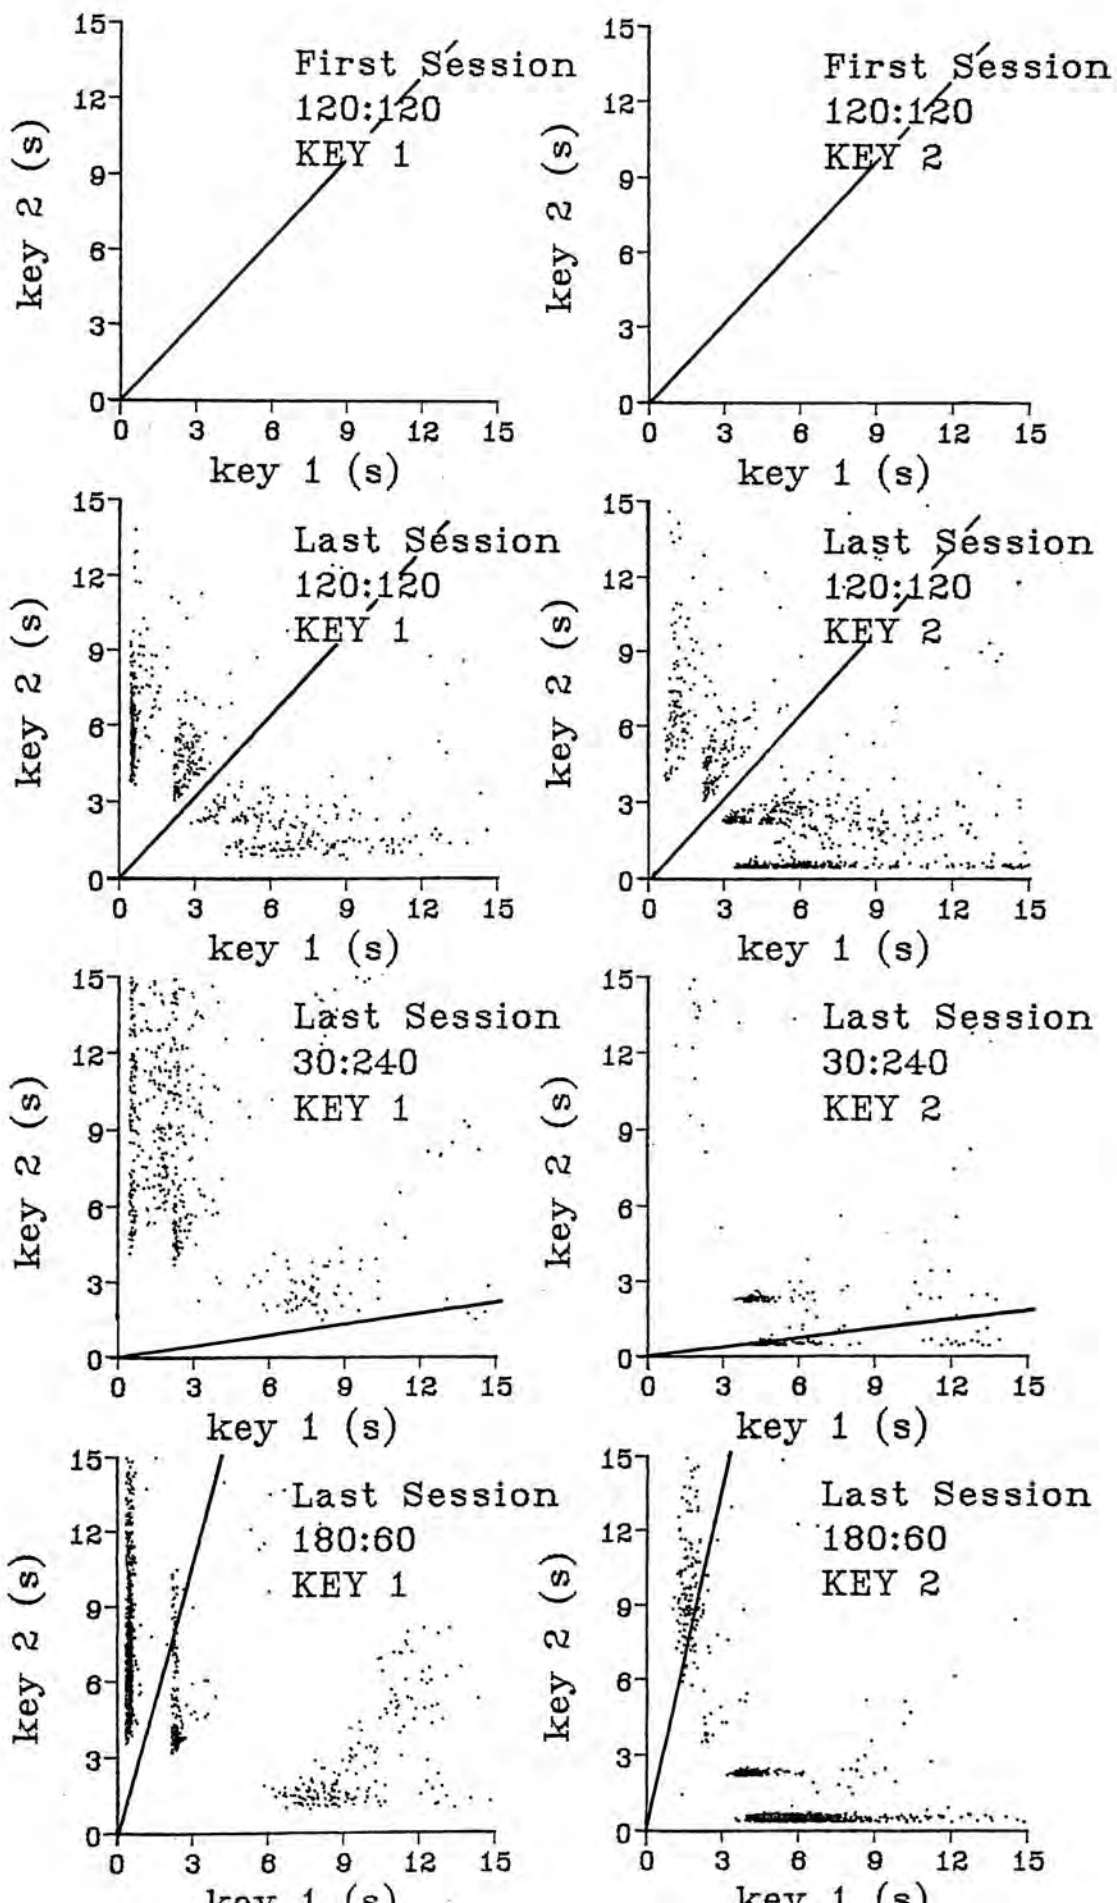
figure 8.12
126

Clockspace representation
Bird 65: CO responses only

First Session
120:120

Last Session
120:120

Last Session
240:30

Last Session
60:180

figure 8.13

127
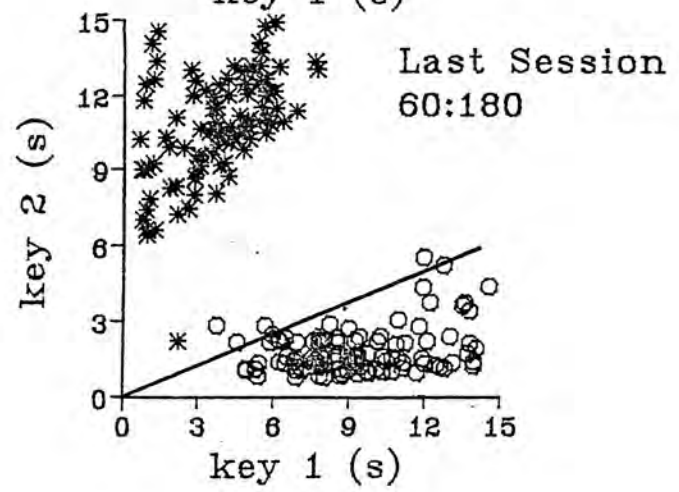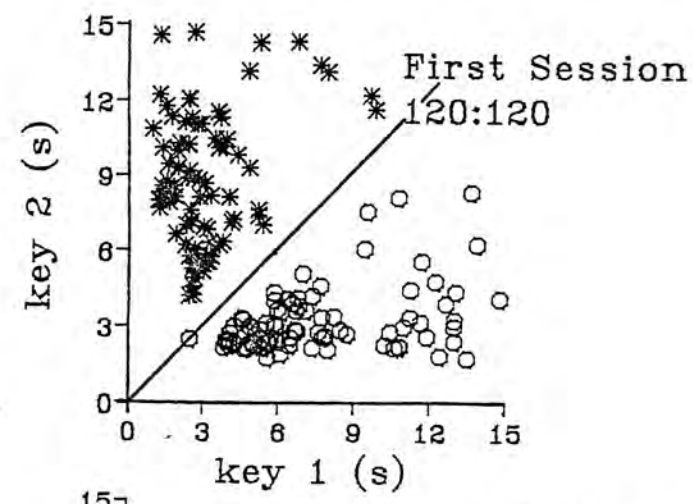
Clockspace representation
Bird 106: CO responses only
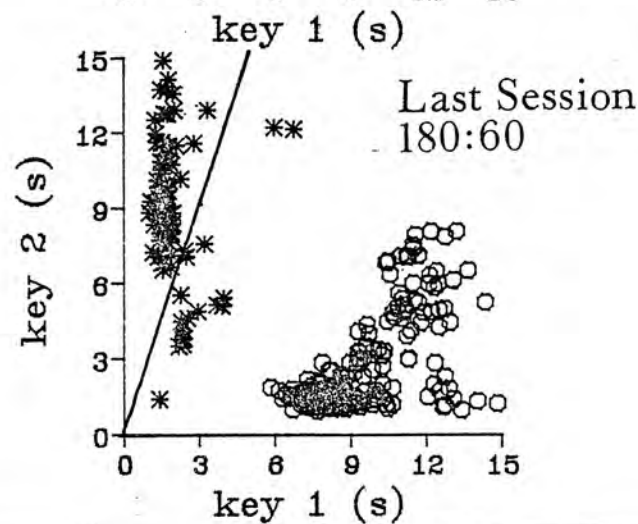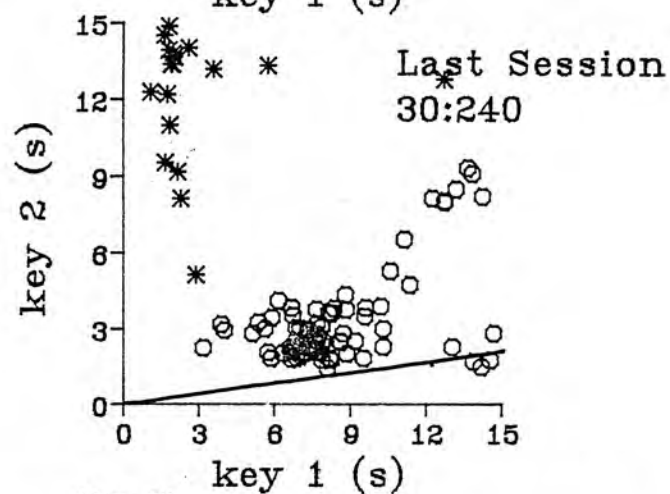
figure 8.14
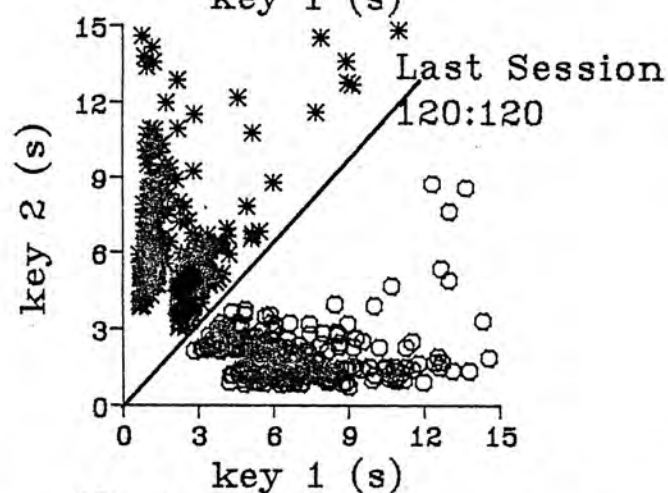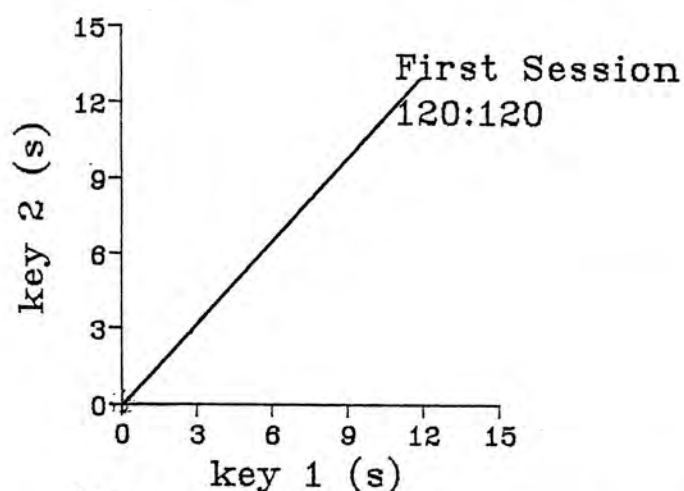
Clockspace representation
Bird 149: CO responses only

8.2.2 Melioration

8.2.2.1 Response and Reinforcer Proportions

As described in the introduction, response and reinforcer proportions were calculated for different numbers of cycles aggregated together (where a cycle was defined as two consecutive changeovers). For aggregates of less than 5 cycles the number of reinforcements occurring on either or both keys was too often zero for the analysis to be worth proceeding with: aggregates of over 30 cycles usually smoothed out to the session average values. Figures 8.15 to 8.18 show the first 100 response and reward proportions (on key 1) for the last sessions of a selection of conditions and subjects at aggregates of 5, 15 and 25 cycles. Figure 8.15 shows data from the last session of bird 130 on FIFI 120:120; 8.16 shows bird 65 VIVIE 240:30, last session; 8.17 shows bird 100 VIVIA 240:30, last session; and 8.18 shows the last session of VIVIA 120:120 for bird 106. Aggregates are calculated for each cycle, so neighbouring data points overlap substantially in the cycles they are calculated from. Response proportions are circles joined by solid lines; reinforcement proportions are triangles joined by dashed lines.

These four figures show clearly that there are fluctuations in both series. These oscillations appear larger in the reinforcement proportion series, especially at aggregates of 5 cycles where the absence of reinforcement on one or other key forces the proportion to its boundary values.

8.2.2.2 Time Series Correlations

Figure 8.19 shows partial autocorrelograms and cross-correlation plots, as described in section 8.2.1.1 for aggregates of 5 cycles from the penultimate session under FIFI 120:120 for bird 130, and the last session under VIVIE 120:120 for bird 106. The value of the correlation on the y-axis is plotted against lag (in cycles) on the x-axis. The 5% significance level for these correlations is indicated by the dashed line. The graphs shown are typical in that there are often a few significant autocorrelations, almost always at lag one, and often scattered unsystematically at larger lags. The existence of the significant autocorrelations means that the cross-correlations must be treated with caution. Correlograms of more than 5 cycles produced graphs smoother than these, usually with a maximum cross-correlation at lag zero. If as melioration predicts, the previous reinforcement proportion is influencing the

figure 8.15

Bird 130. Conc. F.I.:F.I. 120:120
Response – circles. Reinf – triangles.

figure 8.16
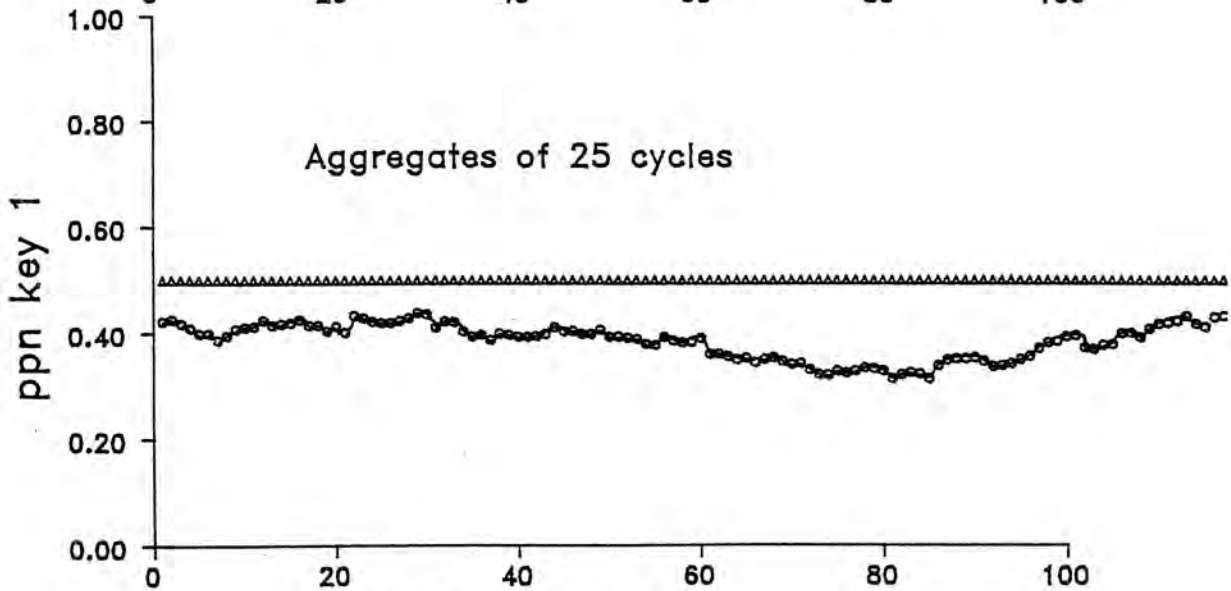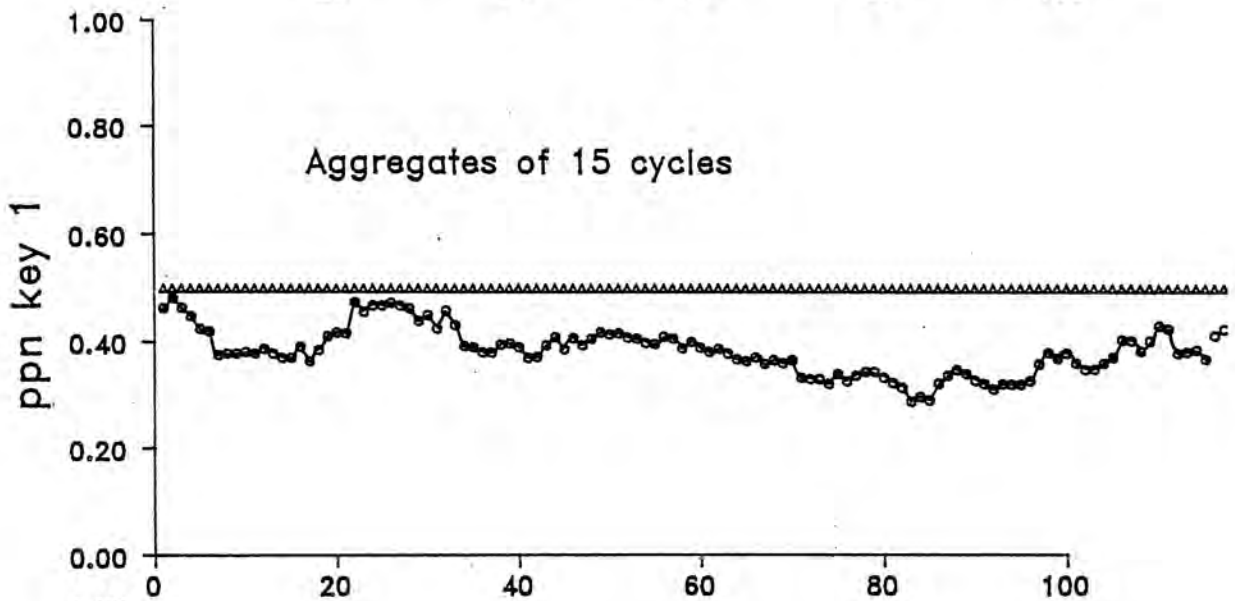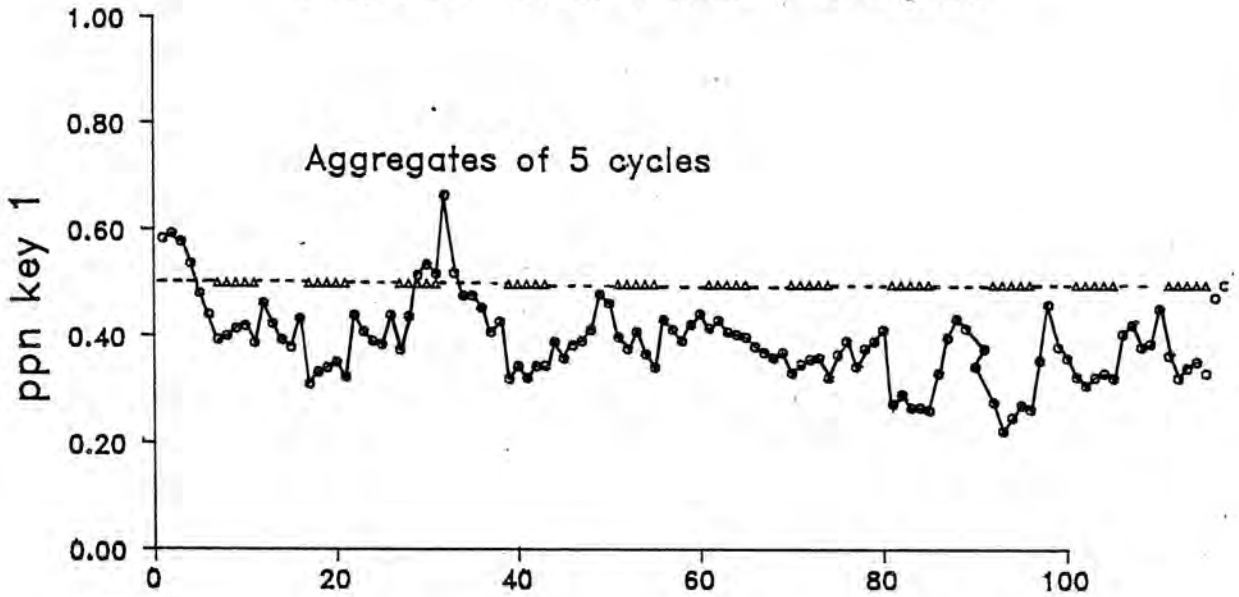
Bird 65. Exponential V.I.:V.I. 240:30
Response — circles. Reinf — triangles.

figure 8.17
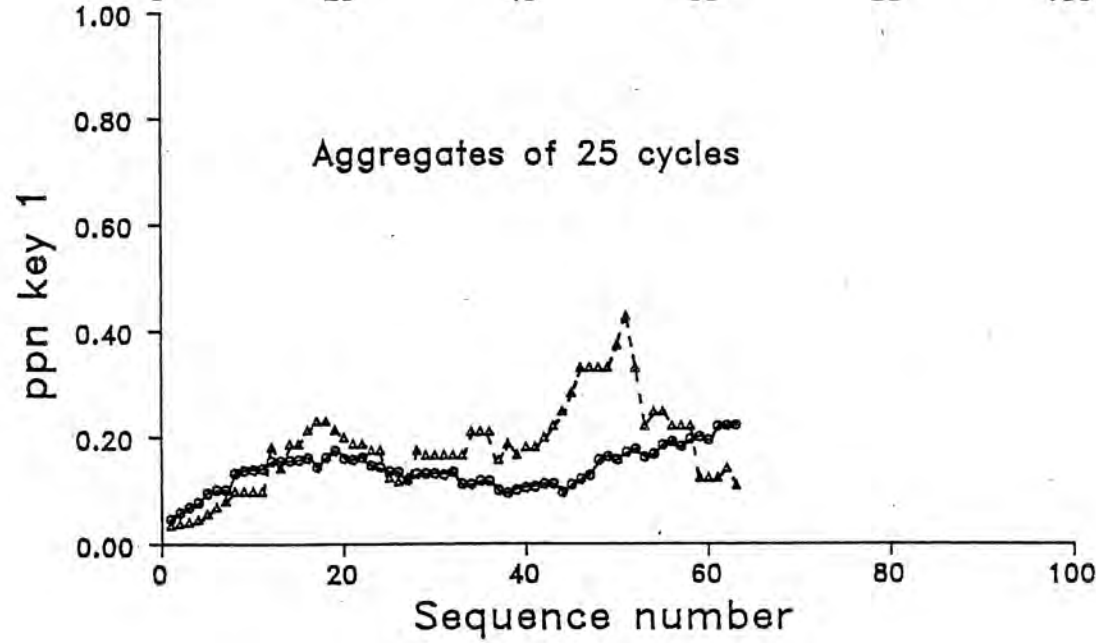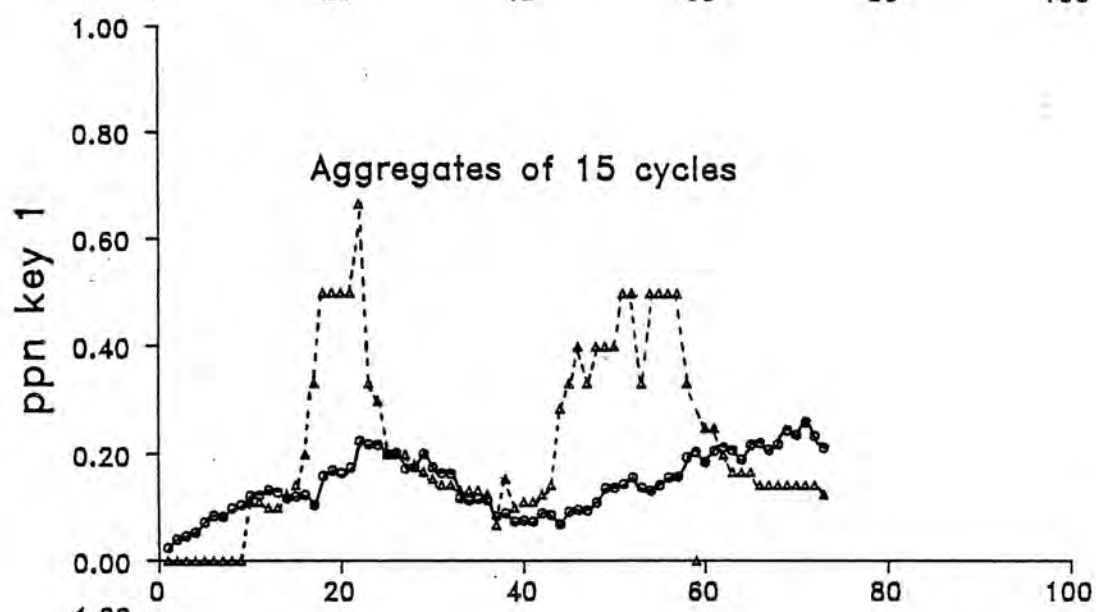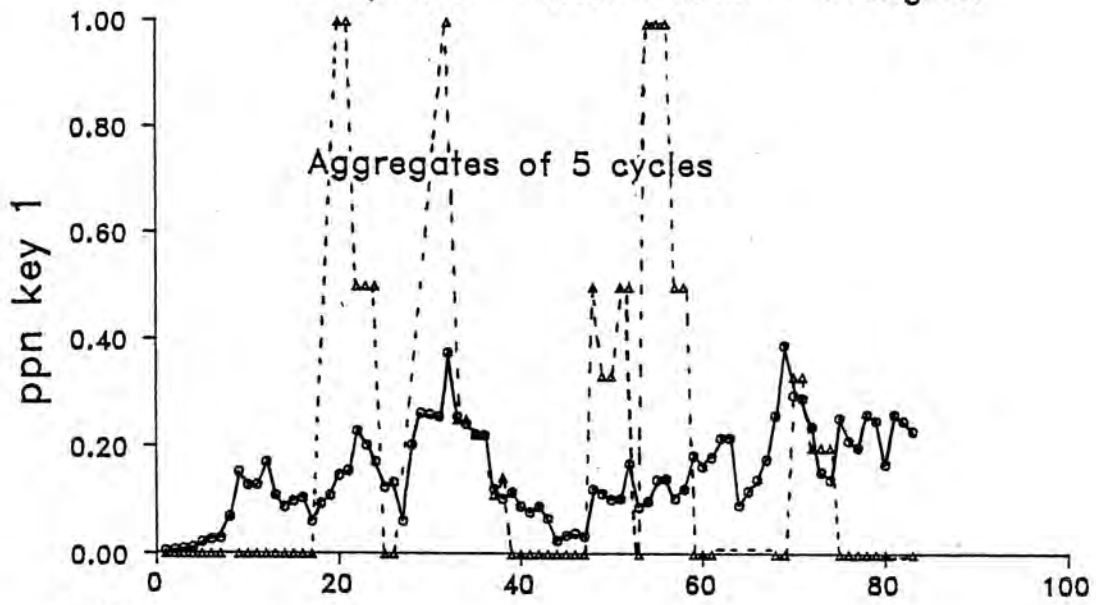
Bird 100. Arithmetic V.I.:V.I. 240:30
Response — circles. Reinf — triangles.

figure 8.18
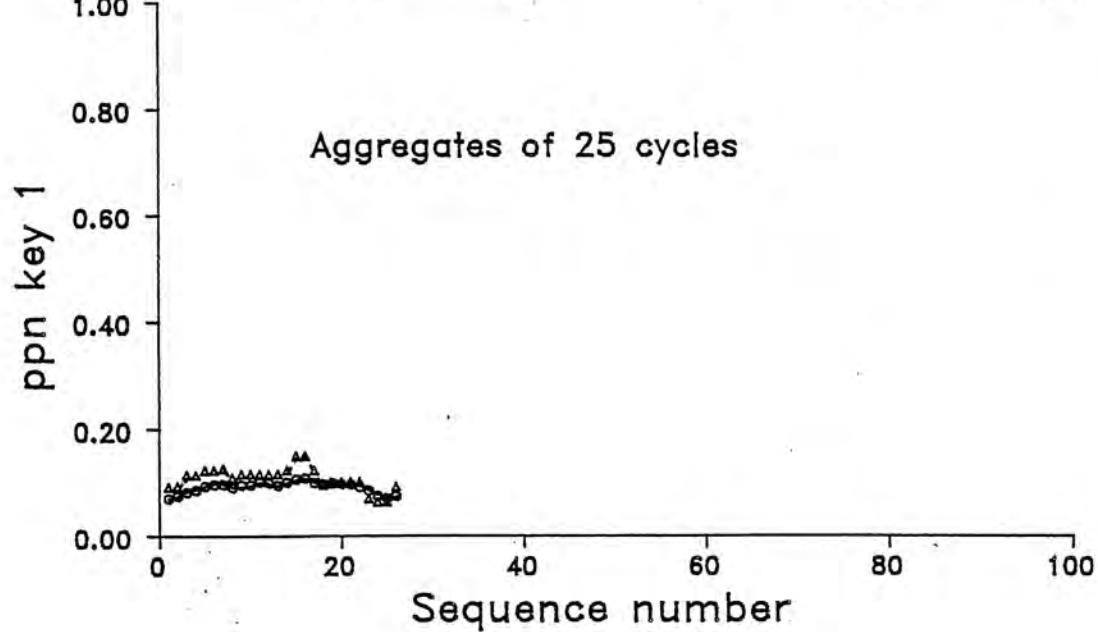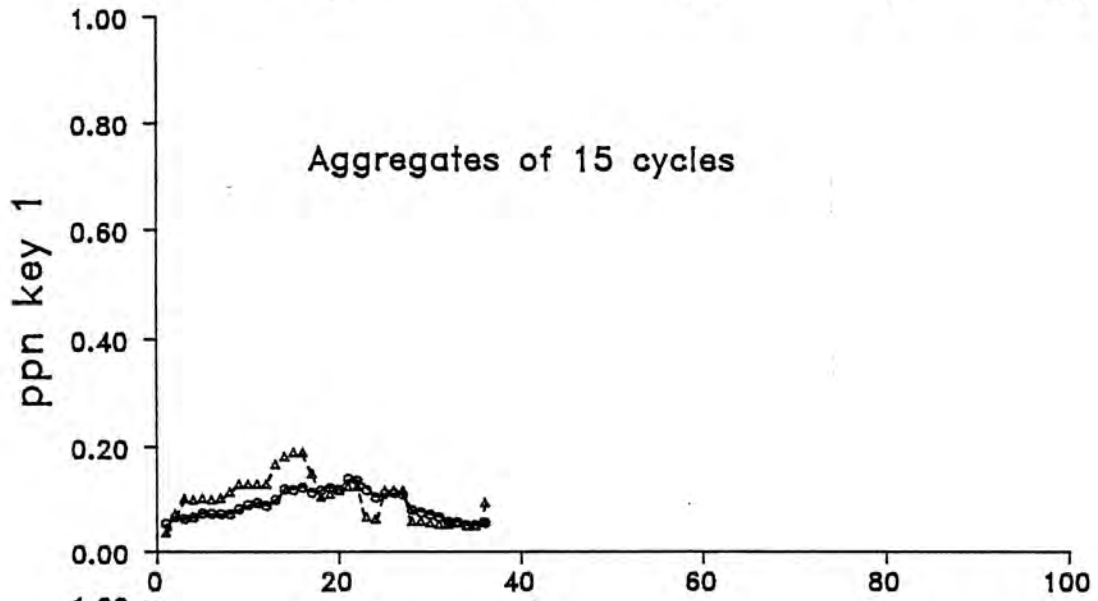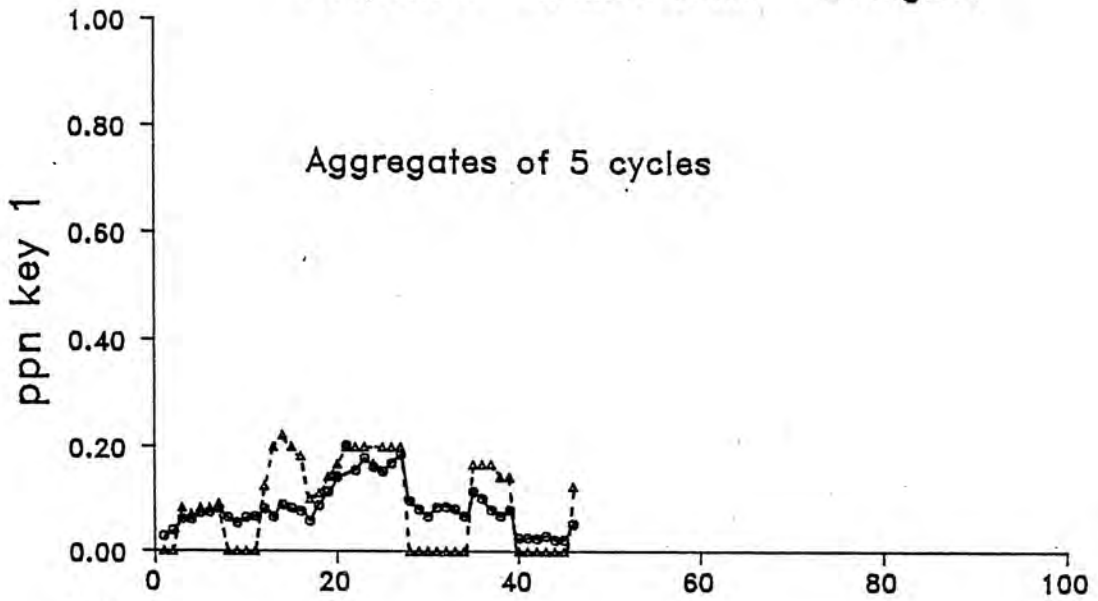
Bird 106. Arithmetic V.I.:V.I. 120:120
Response — circles. Reinf — triangles.

figure 8.19

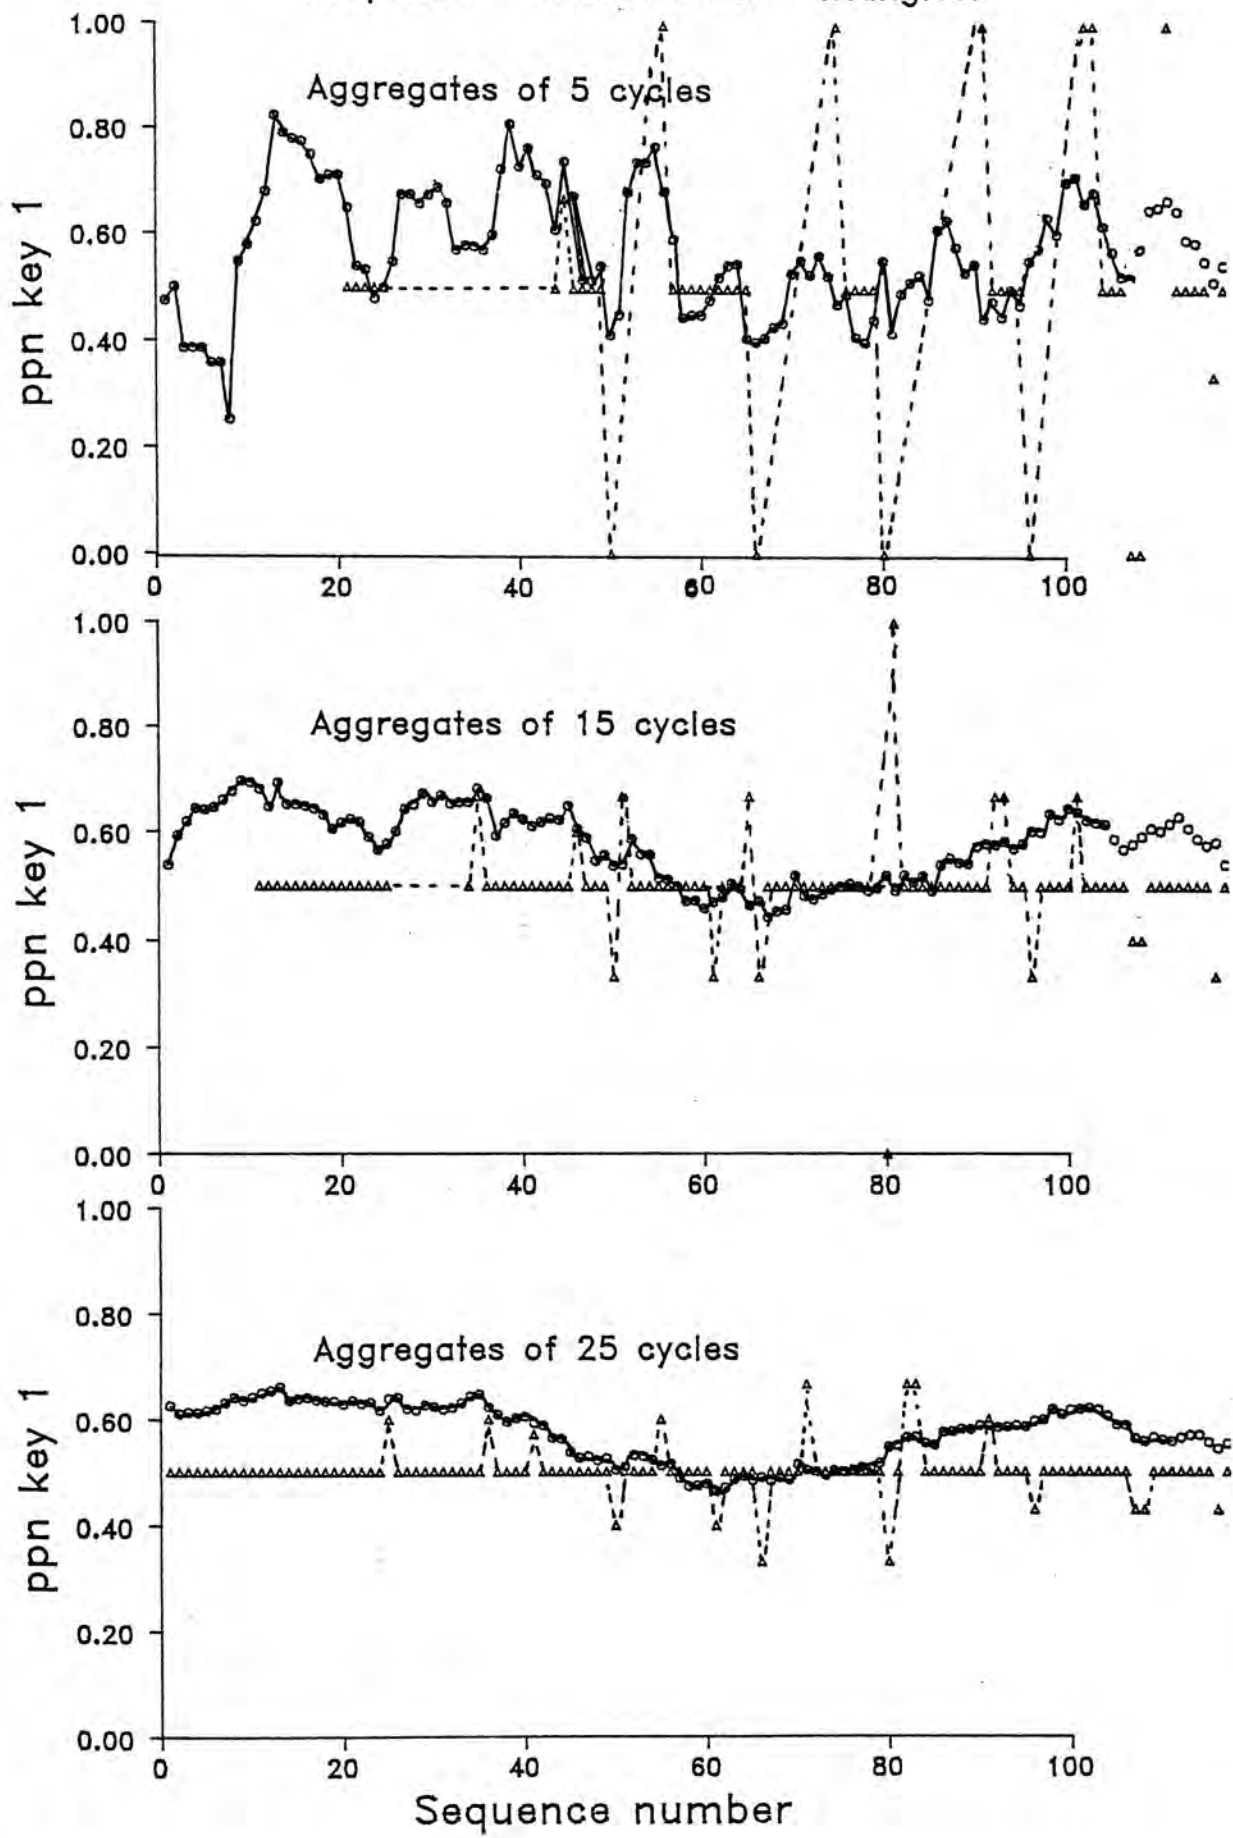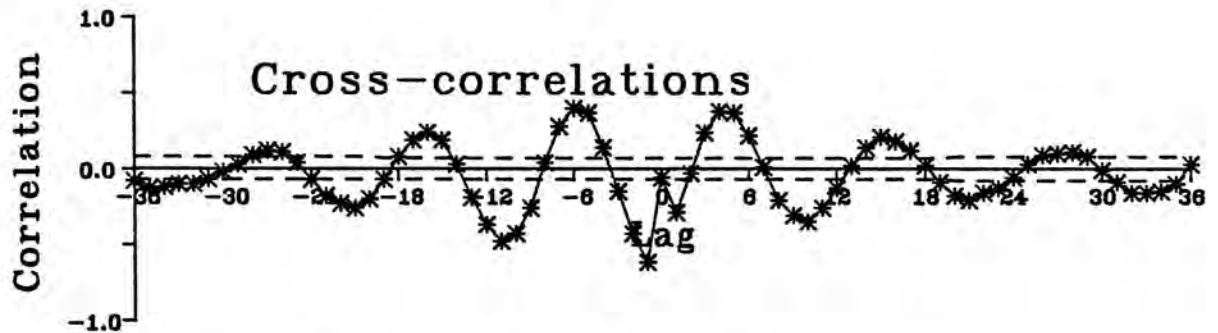Bird 130: Penultimate Session
FIFI 120:120 Autocorrelations

Cross-correlations
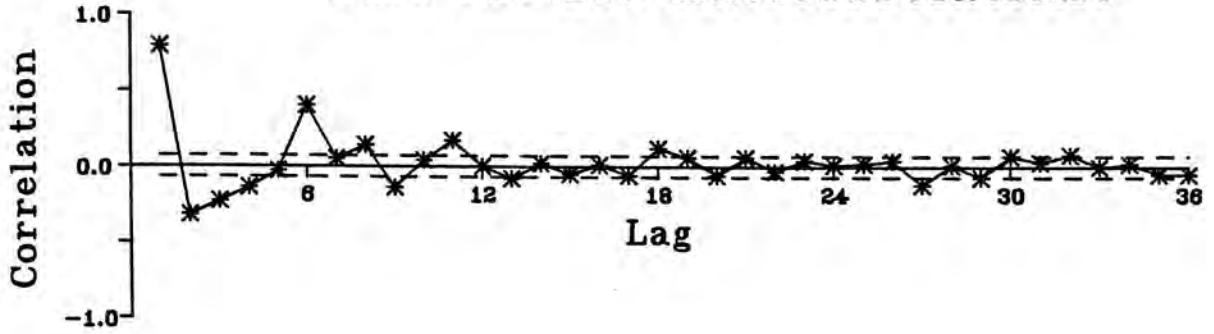
Bird 106: Last Session
VIVIE 120:120 Autocorrelations

Cross-correlations
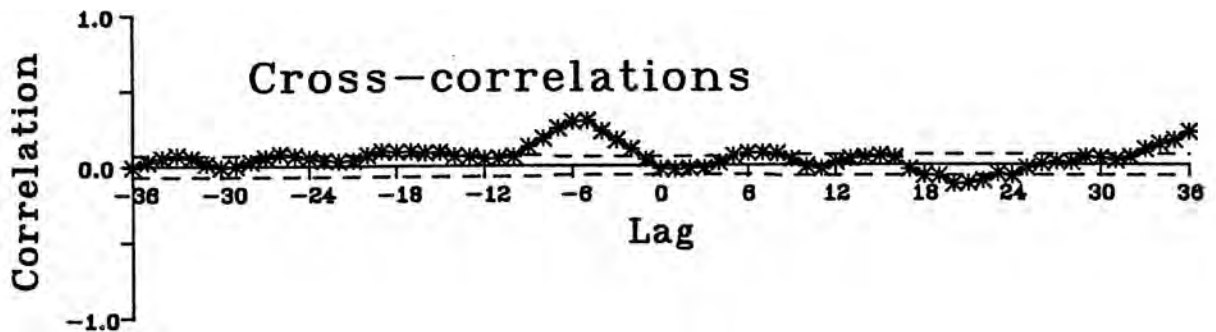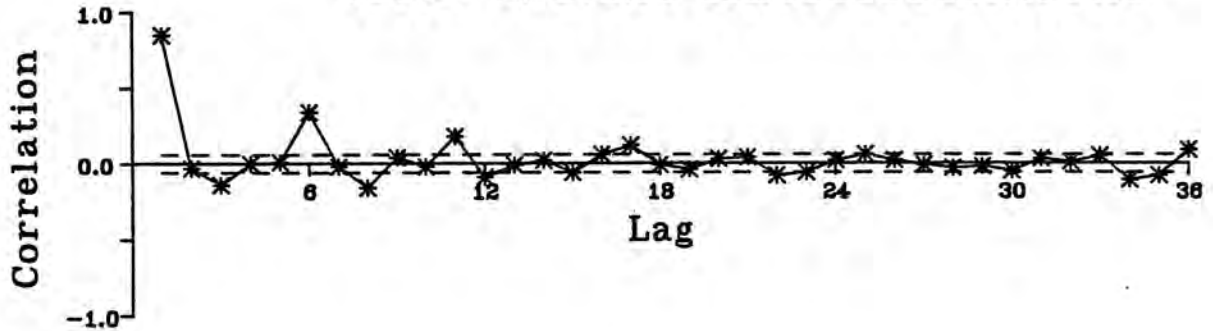
future response proportion, then significant positive cross-correlations at *negative* lags will be observed. In the graphs in figure 8.19 these can be observed, and they imply that a particular value of reinforcement proportion tends to be followed a certain number of cycles later by a similar proportion of responses. However, the existence of significant cross-correlations at *positive* lags means that a particular value of response proportion would tend to be followed by a similar reinforcement proportion at some point in the future. This is simply a result of the schedule payout being to some extent responsive to the subject's behaviour, and does not contradict melioration. A problem arises though, because if, as is often the case (see figure 8.19 upper panel), a timeseries contains significant autocorrelations, and if, as is also common, the series contains significant cross-correlations at *positive* lags (response → reinforcement); then the existence of significant cross-correlations at *negative* lags (reinforcement → response) would be the necessary resultant and need not imply that response proportion is tracking reinforcement proportion. In short melioration is only supported by significant positive cross-correlations at negative lags, in cases where there are not also both significant cross-correlations at positive lags, and significant autocorrelations. In figure 8.19 the upper panels (subject 130) show a typical example of a case where melioration effect could not be discerned because the the significant cross-correlations at negative lags are clearly accounted for by the cross-correlations at positive lags. The case of subject 106, shown in the lower panels of figure 8.19 shows clear evidence of a tendency for a significant cross-correlation at negative lags in the absence of any substantial cross-correlation at positive lags, and therefore this can be scored as a 'hit' for melioration.

Table 8.2:
Number of 5% significant autocorrelations in the peck series
Last 2 sessions of each condition

Concurrent FIFI

Condition	Bird- 144		130		100	
120:120	1	1	6	3	2	4
240:30	0	1	1	1	1	2
60:180	7	3	4	6	1	1

45/23

Concurrent VIVIE

Condition	Bird- 65		106		149	
120:120	9	3	13	9	16	6
240:30	3	3	4	4	4	2
60:180	3	4	12	4	4	6

109/27

Concurrent VIVIA

Condition	Bird- 144		130		100	
120:120	1	1	5	4	4	2
240:30	1	1	13	4	1	3
60:180	4	3	3	4	2	2

62/26

Condition	Bird- 65		106		149	
120:120	1	2	5	3	3	6
240:30	1	3	3	4	4	4
60:180	3	4	3	4	5	8

66/27

Table 8.3: Number of 5% significant cross-correlations at positive lags (Response → Reinforcement): Last 2 sessions

Concurrent FIFI

Condition	Bird- 144		130		100	
120:120	0	3	7	8	0	1
240:30	0	1	1	0	2	1
60:180	0	0	0	7	0	1

32/23

Concurrent VIVIE

Condition	Bird- 65		106		149	
120:120	9	13	0	0	0	13
240:30	0	5	6	13	0	0
60:180	4	5	11	3	1	5

88/27

Concurrent VIVIA

Condition	Bird- 144		130		100	
120:120	0	0	7	3	4	3
240:30	0	0	6	3	0	5
60:180	0	0	7	1	0	0

39/26

Condition	Bird- 65		106		149	
120:120	0	0	2	6	3	16
240:30	1	0	7	0	1	2
60:180	4	0	1	2	0	0

44/27

Table 8.2 gives the number of significant positive partial autocorrelations in the response proportions of each subject on the last 2 sessions under each condition. Table 8.3 shows the number of significant positive cross-correlations at positive lags (response → reinforcement), and table 8.4 shows the same for negative lags (reinforcement → response). Cases of significant negative lags in the absence of both significant positive lag cross-correlations and autocorrelations are underlined.

Table 8.4: Number of 5% significant cross-correlations at negative lags (Reinforcement → Response): Last 2 sessions

Concurrent FIFI

Condition	Bird- 144		130		100	
120:120	0	0	7	9	0	5
240:30	0	5	1	0	0	2
60:180	0	0	3	3	2	1

38/23

Concurrent VIVIE

Condition	Bird- 65		106		149	
120:120	14	6	5	8	0	2
240:30	0	0	6	13	1	0
60:180	5	0	6	5	8	8

80/27

Concurrent VIVIA

Condition	Bird- 144		130		100	
120:120	0	0	2	7	4	3
240:30	0	0	0	8	0	5
60:180	0	1	3	0	0	2

35/26

Condition	Bird- 65		106		149	
120:120	0	0	5	0	0	16
240:30	2	7	3	5	0	1
60:180	0	0	4	1	1	0

45/27

An additional complexity in these data is that 1 in 20 correlations significant at the 5% level, would be expected by chance. Since, where the data were rich enough, correlations were performed up to 30 lags, one or two false positives would be expected in each set. The fraction in the bottom right corner of each panel of tables 8.2 to 8.4 gives the observed number of significant correlations for each group of conditions, over the number to be expected by chance.

The tendency of reinforcement patterning to influence future response choice is not strong. Inspection of table 8.4 shows that there are many cases where there are no more significant cross-correlations at negative lags than would be expected by chance (subject 144 throughout; subjects 100, 130 & 65 under VIVIA). Most of the cases where cross correlations at negative lags do occur correspond to where cross-correlations at positive lags and autocorrelations are also found. We are left with only 8 cases (out of 64) that

The results of these analyses can be summarized as follows. For momentary maximising the only evidence that responses were being placed on the option offering the highest momentary probability of reward was for the first response to a new option. Other responses were not well placed in this regard. Melioration received even less support. The use of cross-correlational analysis did not find any significant tendency for past reinforcement proportion to influence the proportion of responses to each option in the future.

could support a melioration type process. A contingency table was constructed from the number of underlined and non-underlined cells for each subject in table 8.4. A chi-square was then performed to test whether the occurrences of support for melioration were evenly distributed across the subjects. The result for this test was non-significant (χ^2=11.25; df 11; NS); as was a test for differences between conditions (χ^2=1.54; df 2; NS).

8.3 Discussion

Analysing each response for its momentary probability of reward has shown that the vast majority of responses are poorly placed. Any statistic intended to encapsulate these results would yield a negative assessment. Such an approach would overlook however the fact that a certain subset of responses - the first response to a new key, are consistently well placed with high reliability. It might be suggested that this effect is an artefact of the employment of a COD in this study. The COD, as mentioned in the introduction to this chapter, will in some unquantifiable way, add a cost to changing over, because it defines a 2 sec period during which the probability of reward is zero. On a narrow interpretation of momentary maximising this might be predicted to abolish changing over because the momentary probability of reward on the alternative key is zero (Staddon 1983:240); it would certainly tend to produce perseveration on a key. Thus the probability of reward for the new key ignoring COD, when at last the subject chooses to change over, is bound to be high. The constancy of reward probability ratios at changeover shown in figures 8.3 to 8.5 appears to suggest that the subjects have some capability to balance their time since last changeover against time of last response to the same key, to maintain an advantageous momentary probability of reinforcement. The question then arises as to why they do not use this ability to guide all their responding. The analysis of patterning of responding in chapter 9 was partly aimed at this question, but no answer was forthcoming.

Hinson & Staddon's study did not identify any particular subset of responses as being better placed than the rest. One would expect that the effect might be less pronounced because their study did not program a COD. It remains possible, however that in the calculation of their **m** statistic, the effect of many slightly badly placed within-visit pecks was being offset by the

occasional changeover response with markedly favourable reinforcement probability. Though it is not possible to be precise, inspection of Hinson and Staddon's (1983) figures 3 and 4 suggests that even where **m** is over .8 there are still about 30 to 40% of pecks on the wrong side of the switching line – that is to say, responses on the key offering the lower probability of reinforcement. There is, unfortunately, no way of telling from these figures whether these pecks are within-visit or changeover responses.

The analysis of melioration presented here differs from other attempts in the psychological literature (with the exception of Shimp 1982c) to analyse preference controlling mechanisms of this type, by recognizing the two directions of relationship through which matching can be attained. One is the inevitable (on stable schedules) chain of causation by which a certain choice proportion tends to produce a related proportion of reinforcements. The other, and more theoretically interesting direction of relationship, is between the proportion of reinforcements and the future proportion of responding. It is the failure to recognize these two ways in which response and reinforcement proportions can become similar, that has led to the over-estimation, in my opinion, of the matching relation as a description of behaviour. The evidence presented here has suggested that in fact the tendency for reinforcement proportion to influence future responding proportion is very small; although the possibility that this effect occurs over very brief time lags (Horner & Staddon 1986) was not amenable to analysis by the method used here.

The fact that a relationship was found between response proportion and future reinforcement proportions suggests that the breakdown of the session time into cycles, defined by changing over, and aggregates of cycles, was not the factor responsible for the failure to find an effect of reinforcement proportion on response distribution.

In order to test melioration against literal optimisation theory Vaughan (1981) developed a very elaborate form of concurrent schedule. I will call this schedule a 'responsive' schedule, because one of its peculiarities is that it reassess the payout of its two keys according to the response proportion of the subject in the last 4 mins. The other peculiarity was that the VI timer for a key only ran in the 2secs following each response, and not at all during reinforcements. The response proportion was split into eighths, as shown in figure 8.20. In the eighths of greatest preference for each key (relative time on

figure 3.20
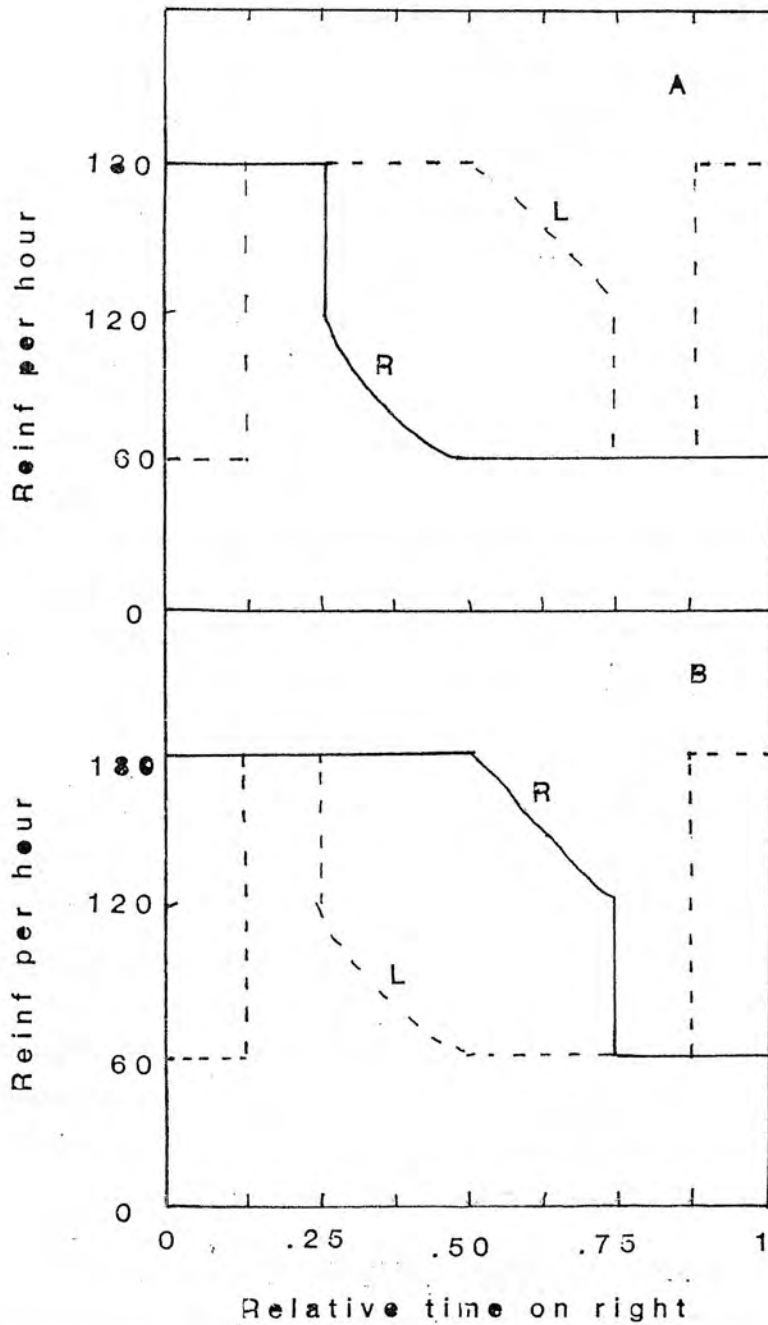
Programmed rates of reinf on left (L) and right (R)

as function relative time on right

Top panel Condition A. Bottom panel Condition B.

from Vaughan 1931

the right 0 → 1/8 and 7/8 → 1) the left key paid off at 60 reinforcements per hour (Rf/h), and the right at 180 Rf/h. Between 1/8 → 2/8 and 6/8 → 7/8 both schedules paid out at the same rate, 180 Rf/h in the former case and 60 Rf/h in the latter. In the range of right key proportion 2/8 → 6/8 the keys moved gradually in their pay out from the 180 Rf/h level to the 60 Rf/h level. Assuming that the inter peck interval is always less than 2 secs so that the programmed rates are actually experienced, this procedure offers two possible response proportion ranges where matching of response and reinforcement proportions is possible (1/8 → 2/8; 6/8 → 7/8); one of which (1/8 → 2/8) provides the highest possible payout. On condition (a) the left key paid off better than the right in the range of response proportion 2/8 → 6/8; for the second condition this relation was reversed so that the right key paid out better in this range. The point of this procedure was that under condition (b) a meliorating subject whose relative time allocation fell within the middle range would be expected to shift preference towards the right and then stabilize in the 6/8 → 7/8 range; which would not produce the level of reinforcement available in the 1/8 → 2/8 range.

Vaughan found that 2 out of 3 subjects produced response proportions between 1/8 → 2/8 for condition (a), and then changed to 6/8 → 7/8 for condition (b), supporting his melioration account (the third subject evinced the same preference under condition (a) but only changed to the 6/8 → 7/8 range on condition (b) after two manipulations of procedure making the 1/8 → 2/8 equilibrium very unstable. Vaughan's argument that this subject was not experiencing the 2/8 → 6/8 preference zone and therefore could not be aware of the changed contingencies is refuted by careful inspection of his figure 2 [consider sessions 35, 37, 47, 48, 49, 50, 51, 52, 53, 54, 55]).

Vaughan's experiment does not distinguish between melioration and any other molecular level account of behaviour (except minimization of deviation from matching, which he acknowledges to have no adherents). Although the observation that two out of three subjects altered their behaviour away from a global maximum as predicted by melioration would contradict a literal optimising account, such a theory has few adherents either (except Rachlin, Battalio, Kagel & Green 1981).

The real interest in this procedure lies in the use to which it was put by Silberberg & Ziriax (1985). They employed a similar responsive form of

scheduling where schedule parameters changed (less elaborately than Vaughan's) according to response proportion. The crucial difference in Silberberg & Ziriax's procedure was that they ran the first four conditions with a timer assessing response proportions and resetting reinforcement timers over just the last 6 secs; the last four sessions used a 4 min timer as in Vaughan (1981). Their rationale for comparing these two timer values was to test the position that the subjects were attempting to momentary maximize but that the 4 min timer was making it impossible for the pigeons to integrate their behaviour over a long enough time period to maximize reinforcement probability effectively. Under Vaughan's condition (b) a subject with relative time to the right in the mid-range (2/8 → 6/8) would, they point out, have to increase time allocation to the left key for four minutes despite its lower reinforcement rate in order to maximize. With a 6 secs timer Silberberg & Ziriax predicted that the birds would be able to maximize because they would be able to integrate their behaviour over time intervals as long as those being used by the experiment controlling computer.

Although melioration in this case was not supported for the 6 secs timer (R_D) varied with response proportion) this is as likely to be due to the small number of sessions run with each of 4 subjects (range 14– 20), and the absence of any attempt to manipulate schedules *post hoc* to produce melioration, as to the change in timer value. Momentary maximization was supported to some extent in 3 out of 4 subjects under the 6secs preference timer, though the data presented are for only 2 out of 8 conditions and do not permit a very clear comparison. Their figure 4 shows individual subject relative time allocation frequencies for condition 1 (one of the 6 secs timer conditions), but summed data for condition 6 (the equivalent condition with a 4 min timer). Inspection of this figure suggests that if the data had been presented summed for both conditions, the difference between them would not have appeared so large. Perhaps if all the results had been presented as separate subject data the difference suggested might also have been less marked.

Though Silberberg & Ziriax interpret their results in a manner consistent with the present study (some support for momentary maximising; none for melioration), I am reticent to give much weight to data from such a peculiar set of scheduled contingencies. No hard and fast distinction can be drawn between 'natural' and 'unnatural' experiments. None the less, conditions where the payout of the schedules is itself continuously reassessed on the basis of

the subjects' responses, produces a disruption in the basic relation between behaviour and environmental effect, that is extremely unlikely to have any even approximate natural analogue. In this account I have followed the original authors in assuming a high and constant rate of response, so that the programmed contingencies are those experienced. However, because the schedules are responsive to the subjects' behaviour in a quite elaborate way, what is really needed is a thorough analysis of the relation between different response patterns and their net returns, and to deduce the optimal strategy. Such an analysis is beyond the competence of the present author.

A preferable method for testing rules of thumb in behaviour is Lester's (1984) employment of a very simple foraging paradigm, placing a small group of goldfish into a tank containing two compartments into which different amounts of food had previously been introduced. This method reduces the impact of behaviour on the environment to the very ecologically plausible process of depletion. Although Lester's study suffers somewhat because it was not possible to test individual subjects, and the level of food remaining in a patch could only be estimated from the time the fish had spent in it, it is none the less a neat paradigm. Lester found cyclings in the preference of the group of fish with the period predicted by melioration ('dynamic matching'). Momentary maximising was not considered, but in this condition would predict a rapid switching of preference as soon as a patch had been depleted below the level of its alternative. Although this pattern is not found in Lester's results on the average, his preference data (presented averaged for the group) do contain small, rapid fluctuations (see his figure 5) which might have been due to individual subjects switching patches more rapidly as their assessment of present patch quality fell below their memory of food availability in the alternative patch.

Kacelnik & Krebs (1985) report some results from a study in press (Kacelnik, Krebs & Ens 1986) utilizing starlings foraging in an aviary containing two operant feeding devices activated by perching on the feeder. They were interested in the starlings' response to a sudden change in the availability of resources. This was done by having each feeder pay out initially at a different probability. There were two conditions, reward probabilities of .75 and .08, and probabilities of .25 and .08. On the second day of testing the payout of the better feeder was abruptly stopped without any signal. Kacelnik & Krebs report that the falloff in preference for the originally better option was too rapid to be

accounted for by any melioration type ('linear operator' in their terminology) model. Unfortunately the data they present (leastways in Kacelnik & Krebs; 1985) are averaged over three birds and ten replications. It would be very interesting to see to what extent the starlings' response to the sudden depletion had changed with increasing experience. None the less this is a study with substantial external validity that raises real problems for a completely melioration type of account. Kacelnik & Krebs suggest a two process model which adds to the melioration (linear operator) theory a process representing the subjects 'confidence' that the reward probability in the patch has not fallen to zero. The authors acknowledge that this amounts to allowing that the birds 'know' that the profitability of a patch may drop to zero. Incorporating a factor assessing the probability that abrupt patch depletion has taken place, brings their model closer to a momentary maximising theory. The difference that remains is that this probability is calculated from the organisms experience of reward or non-reward, whereas the momentary maximising model under discussion here considers the *a priori* probability of reward.

In conclusion, this analysis has found no support for melioration in barbary doves, but some for the momentary maximising found in pigeons in operant experiments by Hinson & Staddon (1981; 1983) and Silberberg & Ziriax (1985). This momentary maximising, however, is only for the first peck to a new schedule, and responses within a visit were found to be unreliably timed. These results contradict Vaughan's (1981) claim that melioration is the choice controlling mechanism on concurrent schedules. It also disagrees with Lester's (1984) study suggesting melioration as the rule of thumb used by goldfish; but may not be inconsistent with Kacelnik & Krebs' (1985) results from starlings.

Melioration is a much simpler process than momentary maximising. To respond on the option having the highest momentary probability of reward, a subject has to have some understanding of the process setting up rewards; whereas to meliorate, all that is required is to shift responding towards the higher local density of prey. It is possible that foragers, where they can (i.e. where the system offers affordances to which they are sensitive), will attempt to forage on the site offering the highest momentary probability of reward. As the environment becomes less stable, less characterized by temporal and spatial regularity, and as the integrative capabilities of the organisms under study become smaller, then all they are left with as a pointer to higher probability of food, are the local rates of reward, and so momentary maximising

reduces to melioration.

Although it would be premature to claim that the process controlling foraging in general is momentary maximising (it is not even clear that momentary maximising [certainly of a priori probabilities], can be defined in all possible contexts); none the less there does appear to be good evidence that some subjects (birds, at least) can integrate information about the environment more elaborately than the simple 'reward tracking' espoused by melioration type models would allow.

CHAPTER 9
KEY PECK EXPERIMENT:
INTER-RESPONSE TIME ANALYSIS

9.1 Introduction

Twice above the need for an analysis of the fine temporal patterning of responding has arisen. In chapter 7 it was mentioned that McNamara & Houston (1983) had shown that optimal responding on variable interval schedules demanded a single inter-response time (IRT), and that any mixture of IRTs was bound to reduce foraging efficiency. In chapter 8 I suggested that a direct analysis of IRTs might shed light on why changeover (CO) responses were placed on the key offering the highest momentary probability of reward, but subsequent pecks were not. Local response patterning is also central to any analysis inspired by the paradigm of 'Variation & Selection' outlined in chapter 3. If the processes of 'blind variation' and 'selective retention' do occur in behaviour we cannot expect to observe them until we have established the units of behaviour that are under the subjects' control.

Molar psychological theories of behaviour have tended to ignore or down-rate the local patterning of responding in operant environments. Shimp (1979) has pointed out that even the gearing of event recorders standard in operant research, has been designed to suppress local irregularities and make the cumulative curve of responding appear smooth. He has further (Shimp 1982c) catalogued a number of cases where molar theorists have ignored interesting response patternings and relationships at more microscopic levels. Rachlin (1978), for example, assumes that the steady state IRT distribution is random, even though he acknowledges that there is evidence this is not so. Nevin (1969; 1979); Heyman (1979); Weiss, Laties, Siegel & Goldstein (1966); Tustin & Davison (1978) and Leslie (1981) have all published data purporting to show that at the local level behaviour on concurrent schedules is indifferent to the pattern of reinforcement. Shimp (1982c) has particular quarrels with Nevin's (1969;1979) analyses, but also points out that in any case, to demonstrate that local effects are random or negligible by any one form of analysis does not prove in general that local patterns of responding are trivial.

There is now a growing body of data, discussed above (chapter 7) showing that the local patterns of one form or another, are legitimate entities for analysis, even on standard concurrent schedules (e.g. Hinson & Staddon 1981; 1983; Real 1983; Real & Dreyfus 1985). Shimp has himself demonstrated in a series of experiments that pigeons can be trained to emit shorter or longer IRTs (1968; 1970); different short patterns of responding across two keys (1982b; 1982c) and can remember the behaviour they emitted a short time before (1982a; 1982b).

The argument from Evolutionary Epistemology (chapter 3) demands more than just that local response patternings should be trainable; they should also arise spontaneously. Lashley (1951) was very concerned about natural rhythms of behaviour, he felt that they would be a consequence of the reverberatory cycles he envisaged operating in the nervous system. He suggested:

> I can best illustrate this conception of nervous action by picturing the brain as the surface of a lake. The prevailing breeze carries small ripples in its direction, the basic polarity of the system. Varying gusts set up crossing systems of waves, which do not destroy the first ripples, but modify their form, a second level in the system of space coordinates. A tossing log with its own period of submersion sends out periodic bursts of ripples, a temporal rhythm.....

> Lashley 1951

Lashley's poetry seems, for the most part, to have fallen on deaf ears – at least so far as comparative psychology is concerned.

Some early authors noticed a tendency for responding on operant schedules to occur in bouts. Mecher (1958) noted that under extinction after responding on a variable ratio schedule, rats tended to respond in distinct bouts or runs of high response rate, rather than continuously responding at an intermediate rate. Blough (1963) developed an interesting method of displaying IRTs on a cathode ray oscilloscope. He found clear non-uniformities in the patterning of a pigeon's IRTs produced by two different schedules. Under VI 4 min two peaks in the IRT record were noted (at approx .4secs and 2→3secs); when the schedule was changed to FR30 these two peaks came close together (.4secs & .7secs). Blough concluded:

> Certain of the responses that contribute to the 'rate' on VI and

DRL [Differential Reinforcement of Low rate] schedules are almost entirely controlled by prior responses; their probability does not vary with extinction, stimulus change, etc., except indirectly through changes in the probability of the responses on which they depend.

In other words Blough found strong evidence for bouting of responses.

More recently the work of Schwartz (1980; 1985) has provided another form of support for Lashley's notion. He gave pigeons a panel of 36 lights that they had to illuminate in any order from top left to bottom right. They did this by pecking on two keys; one of which moved the illumination one row downwards, and the other moved it one column to the right. The birds received reinforcement when the bottom right light was illuminated, independent of the sequence of lights that had been passed through. The only restriction was that they should not attempt, by pecking the two keys, to push the illumination beyond the boundaries of the array. Schwartz observed that although many patterns of movement of the light through the array were possible, only certain patterns were observed; and in most cases a small subset of patternings dominated all the birds' attempts.

In ethology the subjective impression of spontaneous patterning in behaviour is a commonplace. Slater (1973) and Machlis (1977) review the attempts to make these subjective Gestalten (Lorenz 1971/1959) into rigorous mathematical descriptions of behaviour. In the context of very stereotyped behaviours such as pecks, this question resolves itself into the problem of trying to find a rigorous mathematical justification for the natural tendency of ethologists to clump observations of responses into bouts. Early definitions of the bout emphasized a period of pecking uninterrupted by other activity (e.g. Feekes 1972; Wiepkema 1968), or IRTs shorter than some 'bout criterion interval' (BCI) chosen by inspection (e.g. Zeigler, Green & Lehrer 1971; Andrew 1964). As Machlis (1977) points out, these definitions served their authors' purposes.

An early step towards a more rigorous definition of bouts was Nelson's (1964) recognition, analysing the courtship of glandulocaudine fish, that if the intervals between events are statistically independent, then the cumulated sum of intervals greater than a particular value (the survivorship plot) will describe a negative exponential curve; and so if the ordinate is plotted logarithmically (the log survivorship plot) the result will be a straight line of negative slope. When

Nelson did this he found that for some actions of the male fish, a convex curve was found. However, when occurrences of these actions were separated out according to whether a spot on the underside of the fish's body was black or not, it was found that the two new sets of intervals each produced a straight line in the log survivorship graph. The convenience of a chameleon spot on a subject to identify which intervals between events belong to which underlying process is not readily available, and so other methods have been developed to disambiguate mixtures of interval generating processes.

Heiligenberg (1965) analysing mouthing actions of fish compared the distribution of progressively larger intervals between events, until he found the point at which the distribution departed from what would be expected on the assumption of a Poisson process. Most subsequent research on this problem has followed Nelson (1964) in centring on the log survivorship plot and attempting to deduce the optimal BCI from there. If two processes of different rates are generating the intra-bout and inter-bout intervals then the log survivorship function will be convex (Wiepkema 1968; Delius 1969; Duncan, Horne, Hughes & Wood-Gush 1970).

Duncan et al (1970), following Wiepkema (1968), looked for a discontinuity in the survivorship plots as an estimate of the BCI. This method is still common (Clifton 1979; Culshaw & Broom 1980; Petersen 1976; 1978; Slater 1974a; 1974b; 1975).

Machlis (1977) outlined a method intended to make the determination of a BCI more rigorous. She suggested that a computerized procedure be used that worked from the left of the cumulative survivorship plot. Firstly a negative exponential was fitted to the intervals longer than the present point (x), and tested against the data with a chi-square statistic. If the null hypothesis that the intervals greater than x were distributed as a negative exponential could be rejected at the 5% level, then a new value of x further to the right along the survivorship curve was chosen. This process was repeated until the null hypothesis could not be rejected at 5%. The value of x at which a fit occurred was defined as the BCI. In the second part of Machlis' study she subtracted out of the survivorship curve the first negative exponential fitted to the data; and then repeated the procedure to see if a second break point (third distribution) could be fitted to the intervals. In this way she found evidence for 'meta-bouts'; a tendency for the bouts themselves to be grouped into clusters.

The problem with this method, as Machlis (1977) acknowledges, is that the parameters of the negative exponential function fitted to the intervals are dependant on the probability level chosen for the chi-square test (Slater & Lester 1982), and no formal justification is given for taking the point where the curve departs from the negative exponential fitted to the longer intervals as the BCI.

Fagen & Young (1978) suggested firstly, that the parameters of the mixture of exponential distribution fitting the mixture of intervals should be estimated by a maximum likelihood procedure. Then they proposed that the point of intersection of these exponentials, fitting the top and bottom parts of the survivorship curve, should be considered the BCI. The interval at which this occurs, **t′** is given by:

Equation 9.1.1

$$N_w e^{-\lambda_w t'} = N_b e^{-\lambda_b t'}$$

where λ and **n** are the rates and numbers of intervals in the Poisson processes describing intervals falling within (subscript **w**) or between (**b**) bouts. (Note that for a exponential distribution the variance is equal to the mean, which are both 1/rate [e.g. Cox & Lewis 1966]). Slater & Lester (1982) claim (but do not demonstrate) that this is the point where the smallest amount of time is misclassified as between bout when really within, and vice versa. Equation 9.1.1 can be solved for the Fagen & Young BCI **t′$_F$** thus:

Equation 9.1.2

$$t'_F = \frac{1}{\lambda_w - \lambda_b} \ln\left[\frac{N_w}{N_b}\right]$$

Slater & Lester go on to suggest that the point at which the smallest amount of time is misclassified is not actually the best BCI. They propose that the BCI should misclassify a minimal number of intervals, not a minimal amount of time. These are not equivalent, simply because within bout intervals are shorter (by definition) than between bout intervals. Therefore if a minimal amount of time is being misclassified, more within bout intervals than necessary will be misassigned so as to ensure the correct identification of a relatively small number of between bout intervals.

Slater & Lester (1982) demonstrate that the BCI that minimizes the number

of intervals misclassified, t'_S is given by:

Equation 9.2

$$t'_S = \frac{1}{\lambda_w - \lambda_b} \ln\left[\frac{N_w \lambda_w}{N_b \lambda_b}\right]$$

The problem with Slater & Lester's criterion is that by minimizing the number of intervals misclassified it is implicitly weighting the criterion to emphasize the correct classification of within bout intervals, at the expense of between bout intervals. Bearing in mind that within bout intervals are more plentiful than between bout intervals, and that each one of the former occupies less time; Slater & Lester's method pushes the BCI towards the longer end of the spectrum of intervals to ensure the correct classification of as many within bout intervals as possible.

This becomes clearer when we consider the proportion of intervals of different types correctly classified by each criterion. The proportion of intervals classified as within bout is given by:

Equation 9.2.2

$$1 - N_b \, e^{\lambda_w/(\lambda_w - \lambda_b)}$$

for Fagen & Young's BCI: and

Equation 9.2.2

$$1 - \lambda_b/\lambda_w \, N_b e^{(\lambda_w/\lambda_w - \lambda_b)}$$

for Slater & Lester's method. The proportion of intervals classified as between bout is given by Fagen & Young's method as:

Equation 9.3.1

$$N_w e^{-(\lambda_b/\lambda_w - \lambda_b)}$$

and Slater & Lester's BCI gives

Equation 9.3.2

$$(\lambda_b/\lambda_w)N_we^{-(\lambda_b/\lambda_w - \lambda_b)}$$

It can be seen that these are identical except for a weighting factor (λ_b/λ_w) in the formulations based on Slater & Lester's BCI. As the between bout rate is always slower (mean interval longer) than the rate of the within bout process this weighting factor will always be less than one. Therefore the proportion of intervals classified as within bout by Slater & Lester's method will be closer to unity than that from Fagen & Young's method. Similarly Slater & Lester's BCI will classify fewer intervals as between bout than Fagen & Young's. The methods would only assign the same proportions of intervals as within and between bouts in the degenerate case where $\lambda_w = \lambda_b$.

9.2 Method

It has been mentioned above (chapter 6) that this study recorded neither responses that occurred within .3s of each other, nor pecks during the changeover delay (COD). For an IRT analysis, therefore, the COD initiating response (first response to a new key), and also the first peck after the COD, were excluded from calculations. On the basis of a few sessions where complete responding up to the limits of the recording equipment (approximately .05s) was available, it was decided to interpolate the missing very short intervals on the assumption that the peak IRT was at .3s → .4s and that the fall off in the IRT histogram was fairly rapid and linear below that level.

A FORTRAN program was employed here (kindly made available by D. I. Sales of the AFRC Poultry Research Centre, Roslin, Midlothian), to find the maximum likelihood estimates of the parameters of a mixture of exponential distributions (mean [=variance; =1/rate], proportion in mixture). Rather than producing cumulative numbers of intervals longer than a certain duration, and fitting negative exponential distributions to these, it was felt that the scope for error in calculation (both simple rounding error, and programming mistakes) would be reduced if exponential distributions were fitted to the frequencies of intervals falling into a time interval. In other words, fitting distributions to the frequency histogram rather than the survivorship curve. The program enabled the comparison of mixtures of up to four distributions, and provided a residual deviance for each mixture, by which the null hypothesis that adding an additional distribution did not reduce the residual deviance, could be tested (through a chi-square test). The program was based on a maximum likelihood

algorithm by Agha & Ibrahim (1984), that was itself founded on the methods of Hosmer (1973) and Hasselblad (1966,1969). This method is in line with the procedure advocated by Fagen & Young (1978). It was run on data in the form of frequencies of intervals falling into .1s IRT bins. Analyses were performed repeatedly, starting with one distribution and adding in more, one at a time, until the difference in residual deviance between the current model, and a model with one fewer distributions, was not significant at 5%. Although the number of distributions considered present in the mixture was a consequence of the probability level chosen for the chi-square comparison of residual deviances, the parameters of each component distribution were optimal and not affected by the initial estimates supplied by the user.

The BCI was then calculated by the method of Slater & Lester (1982) (equation 9.2 above), and also by the method of Fagen & Young (1978) (equation 9.1). Since the numbers of intervals in each exponential process only enter these equations as their ratio, proportions were substituted for total numbers in both equations.

9.3 Results

Data from the last session of each condition, and the first session of the first condition are presented graphically in figures 9.1 → 9.8. Figures 9.1 and 9.2 show, for bird 100 under FIFI and VIVIA, the cumulated frequency of intervals between pecks longer than **x**, i.e. survivorship plots; figures 9.3 and 9.4 show the same data in the form of semi-logarithmic survivorship plots. The data for subject 65 under VIVIE and VIVIA are shown in the same form in figures 9.5 → 9.8. In the log survivorship plots the BCI if applicable, as defined in equation 9.1 is marked by **F**, and that of the method of equation 9.2 by **S**. The variation within these subjects' data appeared to be as great as that between subjects. The very steep drop off in intervals surviving the first .5s, followed by a lower rate of decay up to about 2s, and then a downward curving, negative exponential falloff in survivorship above that level, are typical features of all subjects' data.

Table 9.1 gives the parameters of the two best fitting exponential distributions to the first two sessions of the first condition, and the last two sessions of all conditions for concurrent FIFI. The columns show condition; key; number of observations (pecks); Number of .1s classes into which the data fall;

figure 9.1   155,

Interpeck Interval Cumulated Frequencies (Survivorships).
Bird 100. Concurrent FI:FI 120:120. First Session

KEY 1   KEY 2
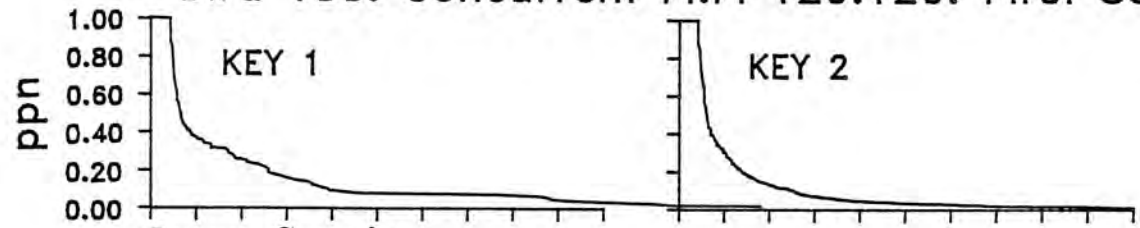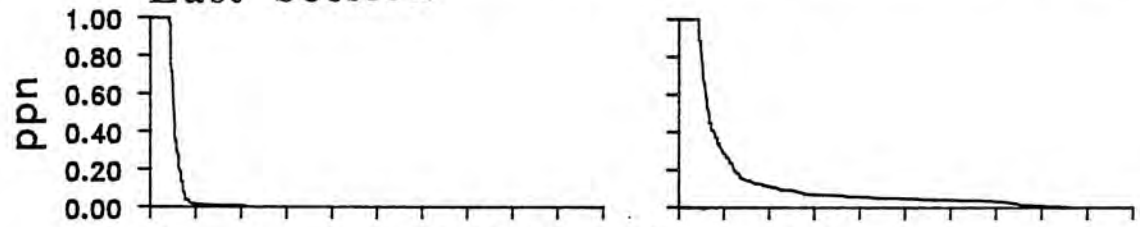
Last Session

Concurrent FI:FI 30:240. Last Session

Concurrent FI:FI 180:60. Last Session
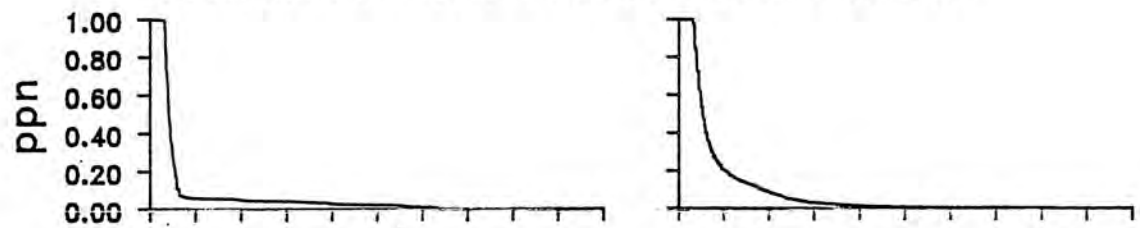
Interpeck Interval (seconds)

figure 9.2   156

Interpeck Interval Cumulated Frequencies (Survivorships)
Bird 100. Arithmetic VI:VI 120:120. First Session

Interpeck Interval (seconds)

figure 9.3

Bird 100 Interpeck Log–survivorships
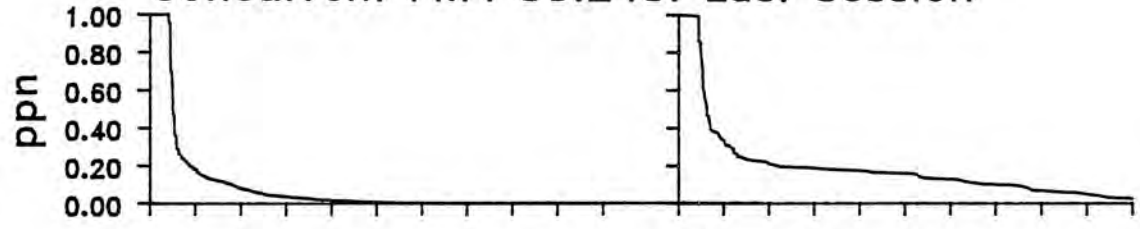## (Interpeck Interval Cumulated Frequencies)
Concurrent FI:FI 120:120. First Session
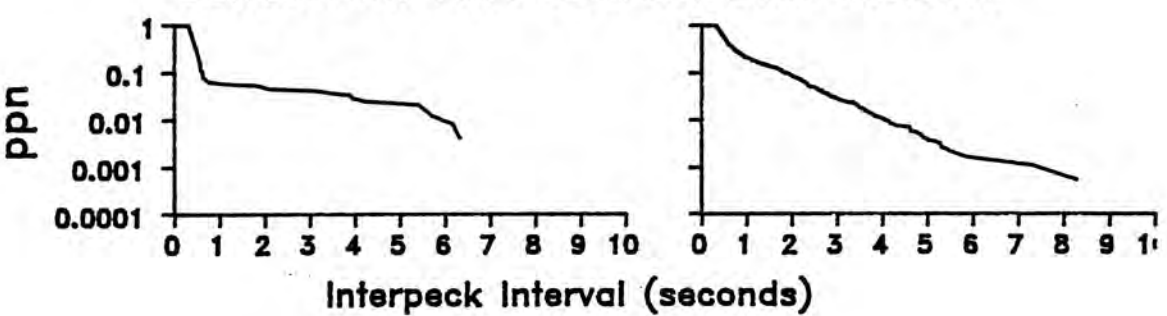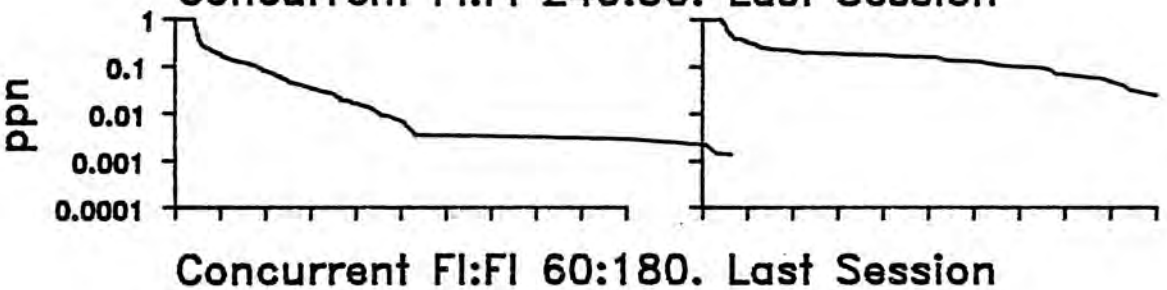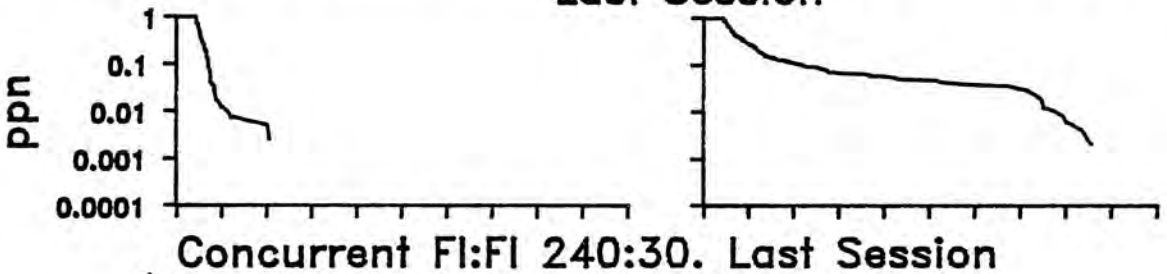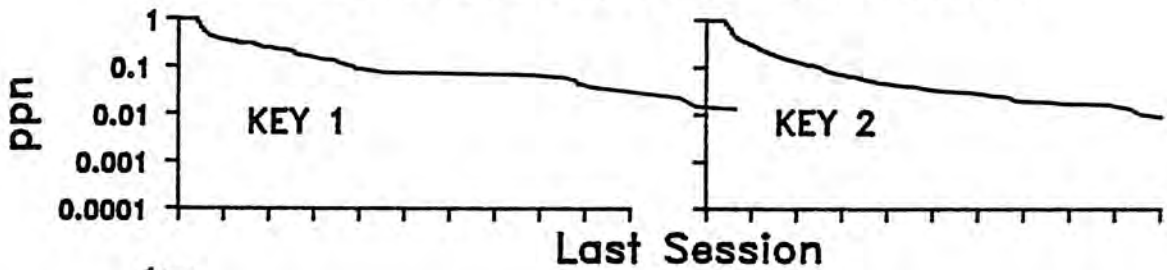
Last Session

Concurrent FI:FI 240:30. Last Session

Concurrent FI:FI 60:180. Last Session

Interpeck Interval (seconds)

figure 9.4   158

Bird 100 Interpeck Log-survivorships
(Interpeck Interval Cumulated Frequencies)
Arithmetic VI:VI 120:120. First Session

Last Session

Arithmetic VI:VI 240:30. Last Session

Arithmetic VI:VI 60:180. Last Session

Interpeck Interval (seconds)

figure 9.5   159
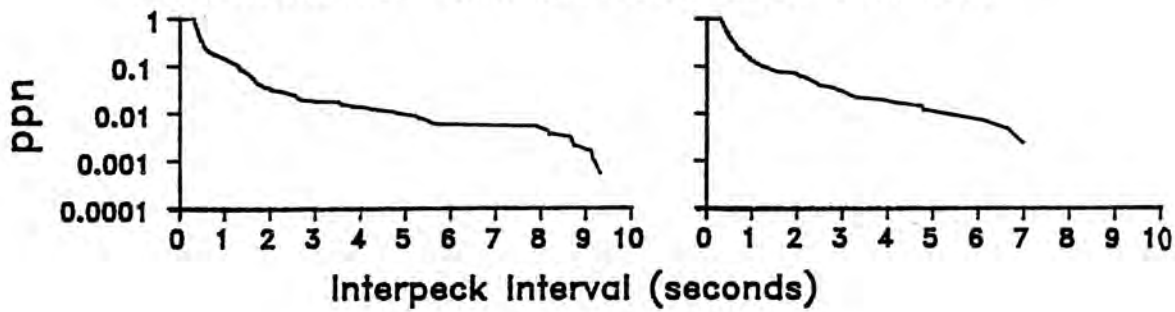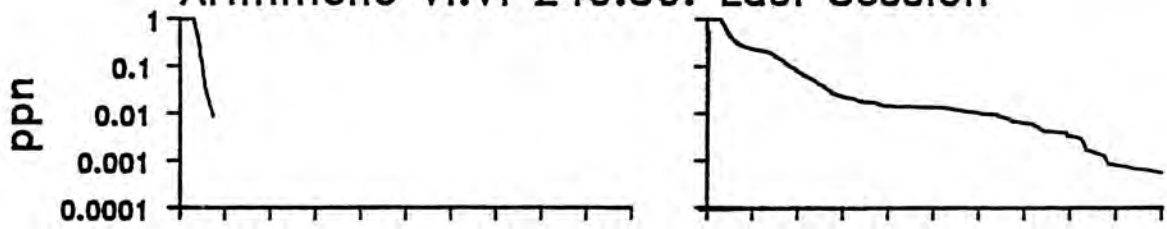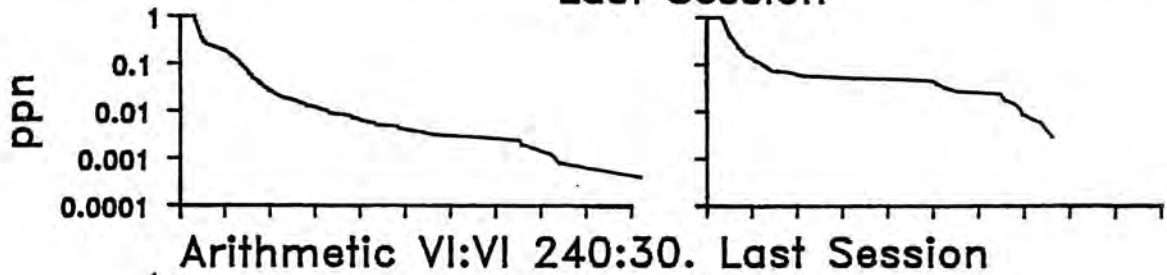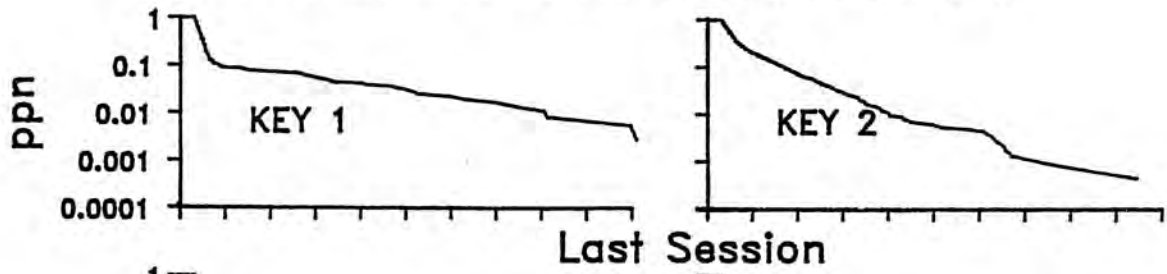
Interpeck Interval Cumulated Frequencies (Survivorships).
Bird 65. Concurrent VI:VI 120:120. First Session

KEY 1   KEY 2
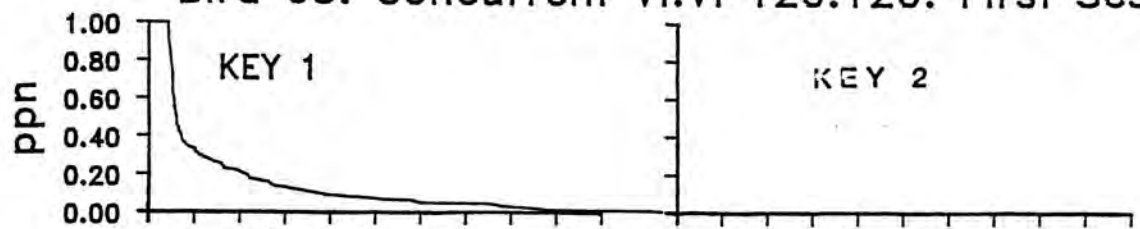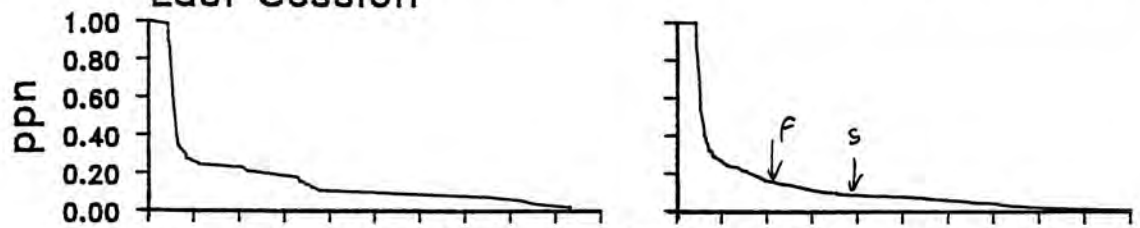
Last Session

Concurrent VI:VI 240:30. Last Session
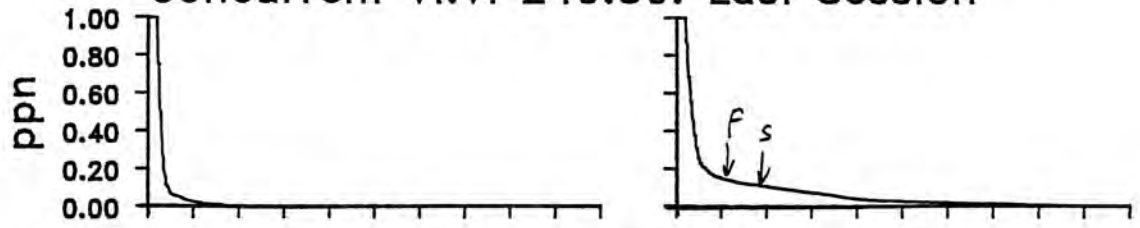
Concurrent VI:VI 60:180. Last Session
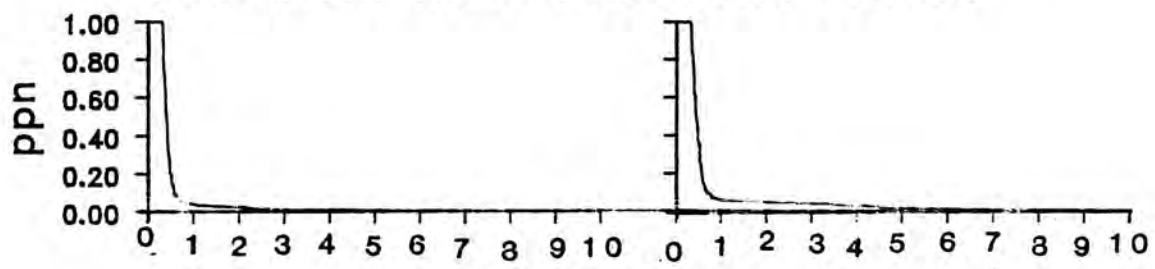
Interpeck Interval (seconds)

figure 9.6   160

Interpeck Interval Cumulated Frequencies (Survivorships).
Bird 65. Arithmetic VI:VI 120:120. First Session

Interpeck Interval (seconds)

figure 9.7   16.1
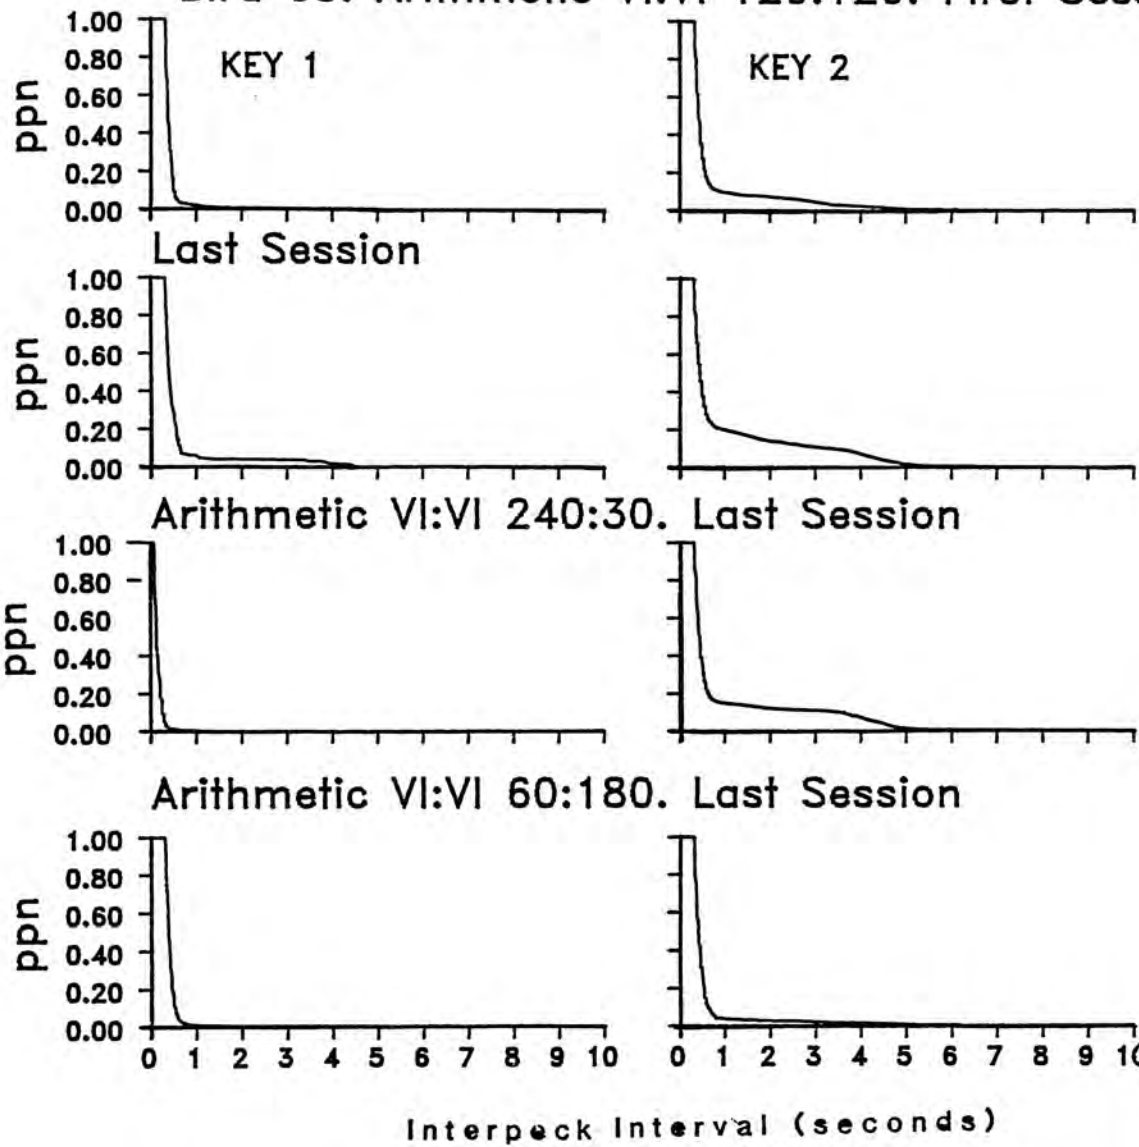
Bird 65 Interpeck Log—survivorships
(Interpeck Interval Cumulated Frequencies)
Exponential VI:VI 120:120. First Session
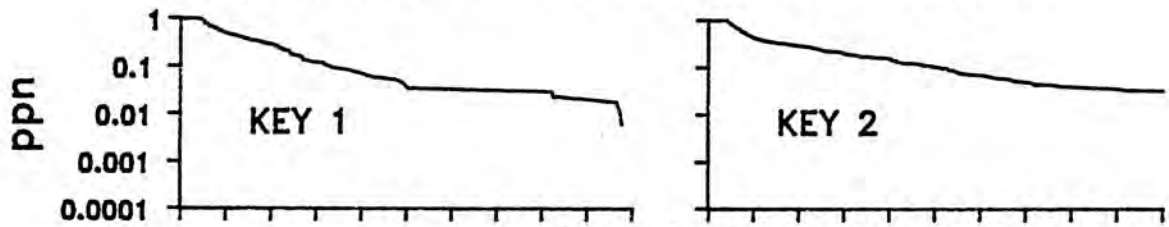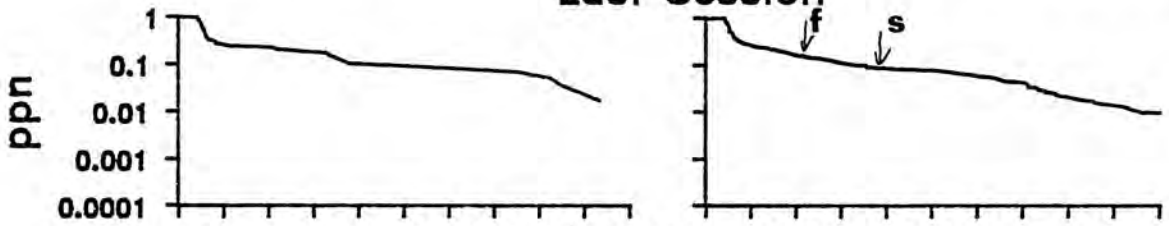
Last Session

Exponential VI:VI 240:30. Last Session
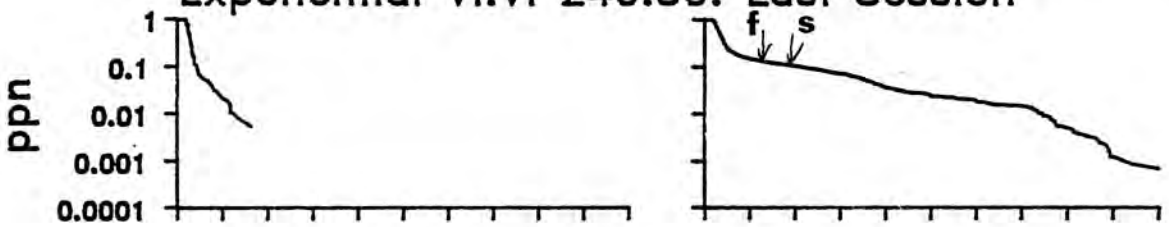
Exponential VI:VI 60:180. Last Session

Interpeck Interval (seconds)

figure 9.8 162

Bird 65 Interpeck Log–survivorships
(Interpeck Interval Cumulated Frequencies)
Arithmetic VI:VI 120:120. First Session

mean intervals (in seconds) and proportion in mixture for the two best fitting distributions; BCI from Slater & Lester 1982; BCI from Fagen & Young 1978. Tables 9.2 and 9.3 give the same data for VIVIE and VIVIA respectively. Where the two distribution fit was not significantly better than that for one distribution, the mean of that one distribution is presented. Also marked in these tables are the BCIs calculated by the method of Slater & Lester (1981; equation 9.2) and from equation 9.1 (Fagen & Young's 1978 method). In no case were more than two distributions required to fit the interval data.

Table 9.1:
Parameters of two best fitting exponential distributions:
Concurrent Fixed Interval

Bird 100

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	122	50	.61					
Begin	2	1147	83	.79					
120:120	1	315	57	.86					
Begin	2	925	88	.85					
120:120	1	36	15	.54					
End	2	158	63	.82					
120:120	1	56	15	.51					
End	2	516	62	.83					
30:240	1	1920	68	.66					
End	2	31	34	.68					
30:240	1	2069	59	.62					
End	2	112	45	.58					
180:60	1	247	29	.43					
End	2	2135	52	.74					
180:60	1	312	24	.41					
End	2	2223	57	.72					

Bird 130

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	43	48	.71					
Begin	2	28	58	.82					
120:120	1	1	1	insufficient data					
Begin	2	58	105	.71					
120:120	1	104	30	.48					
End	2	261	64	1.22					
120:120	1	157	12	.46					
End	2	358	76	1.06					
30:240	1	1885	103	.49					
End	2	136	31	.41					
30:240	1	2318	99	.47					
End	2	72	20	.42					
180:60	1	793	59	.42					
End	2	592	31	.40					
180:60	1	607	76	.41					
End	2	413	30	.40					

Table 9.1: continued

Bird 144

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	0	0	insufficient data					
Begin	2	206	145	.17					
120:120	1	7	17	.60					
Begin	2	273	132	.18					
120:120	1	91	19	.67					
End	2	1726	64	.79					
120:120	1	158	28	.68					
End	2	2884	51	.68					
240:30	1	0	0	insufficient data					
End	2	982	98	.78					
240:30	1	0	0	insufficient data					
End	2	511	74	.65					
60:180	1	218	38	.45					
End	2	1004	44	.62					
60:180	1	160	35	.48					
End	2	1149	49	.45	.79	1.34	.21	1.64	.90

Table 9.2: Parameters of two best fitting exponential distributions: Concurrent Exponential VI

Bird 65

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	134	50	1.03					
Begin	2	415	90	1.12					
120:120	1	239	61	.68					
Begin	2	153	46	1.14					
120:120	1	13	16	.54					
End	2	1573	80	.62					
120:120	1	44	25	.54					
End	2	1004	92	.67	.97	3.50	.04	4.01	2.64
240:30	1	89	5	.28					
End	2	3241	96	.40					
240:30	1	254	15	.29					
End	2	2956	91	.40	.90	2.24	.10	1.91	1.07
60:180	1	1838	44	.40					
End	2	741	33	.44					
60:180	1	1440	40	.42					
End	2	542	34	.44					

Bird 106

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	53	31	.27					
Begin	2	159	108	1.17					
120:120	1	106	54	.57					
Begin	2	150	71	1.01					
120:120	1	32	46	1.21					
End	2	806	58	.46					
120:120	1	43	30	.97					
End	2	840	66	.50					
240:30	1	0	1	insufficient data					
End	2	2057	106	1.12					
240:30	1	0	1	insufficient data					
End	2	2277	97	1.02					
60:180	1	1597	84	.41					
End	2	248	17	.54					
60:180	1	1351	89	.41					
End	2	210	17	.50					

Table 9.2: continued

Bird 149

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	816	101	.38	.85	1.95	.15	1.59	.82
Begin	2	1	1	insufficient data					
120:120	1	366	85	.50	.90	2.22	.10	2.38	1.42
Begin	2	11	27	.66					
120:120	1	476	28	.44					
End	2	903	65	.61					
120:120	1	242	40	.47					
End	2	436	58	.73					
30:240	1	68	97	1.12					
End	2	169	43	.45					
30:240	1	1062	103	1.00					
End	2	140	32	.51					
180:60	1	1501	27	.45					
End	2	1147	30	.50					
180:60	1	1305	26	.45					
End	2	1106	30	.49					

**Table 9.3a: Parameters of two best fitting
exponential distributions: Concurrent Arithmetic VI**

Bird 100

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	3123	58	.53					
End	2	358	38	.50					
120:120	1	3420	56	.54					
End	2	387	39	.49					
240:30	1	202	11	.38					
End	2	3011	65	.69					
240:30	1	159	9	.39					
End	2	2868	74	.65					
60:180	1	2841	58	.51					
End	2	583	38	.51					
60:180	1	2445	62	.51					
End	2	490	42	.52					

Bird 130

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	1081	57	.39					
End	2	204	33	.42					
120:120	1	1410	46	.40					
End	2	423	28	.43					
240:30	1	590	24	.38					
End	2	595	57	.41					
240:30	1	471	18	.39					
End	2	505	48	.43					
60:180	1	2306	83	.40	.82	1.62	.18	1.55	.81
End	2	251	16	.42					
60:180	1	1858	72	.40	.84	1.60	.16	1.62	.88
End	2	356	13	.41					

Bird 144

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	48	23	.53					
End	2	3314	68	.64					
120:120	1	1278	48	.42	.81	1.22	.19	1.61	.93
End	2	116	26	.48					
240:30	1	2	6	insufficient data					
End	2	3165	99	.66					
240:30	1	11	6	insufficient data					
End	2	3051	95	.51	.50	.89	.50	.67	0
60:180	1	87	18	.42					
End	2	1349	49	.42	.81	1.16	.19	1.62	.74
60:180	1	79	14	.43					
End	2	1278	48	.42	.81	1.22	.19	1.62	.74

Table 9.3b: Parameters of two best fitting exponential distributions: Concurrent Arithmetic VI

Bird 65

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	452	28	.41					
End	2	2608	71	.49					
120:120	1	302	21	.41					
End	2	2763	70	.49					
240:30	1	190	10	.40					
End	2	3004	80	.45					
240:30	1	404	11	.38					
End	2	2923	70	.45					
60:180	1	1814	27	.39					
End	2	984	29	.40					
60:180	1	1726	23	.39					
End	2	1068	39	.42					

Bird 106

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	1433	57	.45	.96	1.58	.04	2.69	1.79
End	2	368	42	.58					
120:120	1	1241	51	.43					
End	2	705	44	.59					
240:30	1	317	13	.39					
End	2	2039	75	.69					
240:30	1	327	12	.38					
End	2	1891	79	.65					
60:180	1	1320	51	.39					
End	2	625	34	.52					
60:180	1	1648	55	.40					
End	2	680	30	.52					

Bird 149

Cond	key	No Obs	No Class	Mean 1	ppn 1	mean 2	ppn 2	BCI S	BCI F
120:120	1	2019	50	.49					
End	2	878	13	.45					
120:120	1	2256	47	.52					
End	2	936	26	.45					
240:30	1	909	22	.48					
End	2	2047	44	.47					
240:30	1	1068	27	.50					
End	2	1825	50	.45					
60:180	1	2357	64	.55					
End	2	967	15	.44					
60:180	1	2356	63	.53					
End	2	1065	15	.43					

Tables 9.1 to 9.3 contain data from 168 individual sessions. Twelve cases arose where the program could not fit a mixture of distributions to the data. All of these were simply because there was insufficient data in the session. The

remaining 159 cases can be arranged according to the number of distributions fitted as in table 9.4.

Table 9.4: Number of distributions fitted grouped by bird and session

Condition	bird	Distributions -	1	2
FIFI	100		16	0
	130		15	0
	144		12	1
VIVIE	65		12	2
	106		14	0
	149		13	2
VIVIA(A)	100		12	0
	130		14	0
	144		6	4
VIVIA(B)	65		12	0
	106		11	1
	149		12	0

A χ^2 analysis to test if the differences in numbers of distributions fitted to sessions from different subjects and conditions was greater than would be expected by chance, was significant (χ^2 = 24.43, df=11, p<2.5%). The proportions in each class did not however differ significantly between experimental groups (χ^2= 8.01, df=3, NS), nor between subjects when data from the same subjects on different conditions were grouped together (χ^2= 9.33, df= 5, NS). It can be seen from table 8.4 then, that on average the proportions of sessions most parsimoniously fitted by mixtures of 1 or 2 exponential distributions are 92% and 8% respectively. There is thus a substantial predominance of one distribution fits.

9.3.1 Optimality considerations under VIVIE

In chapter 7 it was shown that optimal responding on VIVIE required a single, constant IRT, and that any mixture of IRTs would reduce foraging efficiency. The data in table 9.2 and figures 9.5 and 9.7 clearly refute this prediction. Not only are the IRTs well described by the most variable of distributions, the exponential; but also, on occasion, more than one distribution is required to fit the data. It was further predicted in chapter 7 that the IRT would be directly proportional to the programmed mean interval between reinforcements.

Table 9.5 gives the results of linear regressions of the mean IRT onto the schedule value for the three subjects run under VIVIE. For the 12 cases where more than one distribution was fitted to the data, this fact was ignored for the purposes of the present analysis, and the overall mean, averaged over the two processes, was used

Table 9.5: Regression of mean IRT onto schedule values

Subject	Intercept	Gradient	p	r^2
65	.72	-.0009	NS	.04
106	.91	-.0017	NS	.06
149	.86	-.0020	<5%	.39
Group	.83	-.0016	<5%	.11

The only significant results in table 9.5 are for bird 149, and for the group data taken as a whole; these show a very slight tendency for mean IRT to decrease with increasing schedule value. This is exactly the opposite result from that predicted from optimality considerations. Although the evidence that mean IRT decreases with increasing schedule value is too weak to deserve any attention in its own right, it is none the less a clear refutation of the optimality prediction.

9.3.2 Bout Criterion Intervals

Tables 9.1 to 9.3 include the BCIs calculated from Fagen & Lester's (1978) method, and also from Slater & Lester (1982). It can be seen that, as predicted in the introduction to this chapter, the BCI from Slater & Lester's formulation is always larger than the result of applying Fagen & Young's equation. This is because it contains an additional weighting factor pushing the BCI towards the longer, between-bout, intervals, and usually producing a criterion that is longer than the mean of either process. Fagen & Young's BCI also occasionally produces a value in excess of the between-bout mean, but is less prone to do so. Fagen & Young's criterion also sometimes produces a figure less than the mean of the within bout process, but only when the proportions of intervals falling into the within-bout and between-bout processes come very close to each other. On balance the Fagen & Young criterion is more consistent with intuition, because it falls in most cases between the means of the two component exponential processes and thereby produces proportions of intervals in each group closer to the true values.

9.3.3 Video Analysis

As noted in chapter 6 occasional video tape records of sessions were analysed for the occurence of a number of behaviour patterns. From these I gained the impression that the birds would lower their heads and peck repeatedly at a key for up to about 2 seconds, then cease pecking and raise their heads slightly. At this point they often looked in the food hopper, apparently to see if their actions had brought any reward; they might walk over to the other key; or, less commonly, stroll around the experimental chamber. Quite often they simply looked around, and then lowered their head again to attack the same key one more. Since these observations are not quantified it would be unwise to attach too much importance to them, but they certainly gave the subjective impression that the pecks could be grouped into bouts, which is not, in general, born out by the statistical analysis above. There are two possible conclusions to draw from this fact. Either subjective impressions should be dismissed in favour of putting all one's trust in mathematics; or the mathematical analysis should be rejected. I am inclined to take the latter course, not because the analysis presented in this chapter was not rigorous, but because it was limited by the fact that very brief IRTs, shorter than .3secs were not recorded. From Machlis' (1977) data we would expect the mean of the within bout process to fall around 1/3 sec, and thus the absence of short intervals (except in so far as I interpolated them) was probably critical to the program's inability to disambiguate a process of within bout intervals from one of between bout intervals in most cases. Where two process have been found at least the problem that Machlis had in trying to attach meaning to her bouts in a simple two-manipulandum experiment, that between bout intervals did not necessarily define crossings over from one manipulandum to the other, could not arise here, simply because the two keys in this experiment were very much further apart than in Machlis' (18cm compared with .2cm and .7cm).

9.4 Discussion

The analysis presented in this chapter adopts a method that has been used extensively in the ethological study of animal behaviour, shows its applicability to the traditional operant paradigm of key pecking in a Skinner box, and also attempts to improve its rigour as a tool.

As reviewed in the introduction to this chapter, the analysis of the temporal

patterning of brief stereotyped behaviours has a substantial history to it, and yet the present study is the first to make use of a maximum likelihood procedure to produce optimal estimates of the parameters of the component distributions in the data. The fact that the present study was not initially designed with this form of analysis in mind has meant that the data contain a substantial problem. Events occurring within .3secs of each other could not be recorded, because to do so would have meant reducing the duration of each session (in order to avoid exceeding the memory capacity of the computer), and this was incompatible with the original purpose of the experiment. Ironically, at the time of writing, for the cost of the machines used in this experiment, micro-computers with four times as much memory space are now available. Because of this shortcoming the failure in general to find within-bout and between-bout processes should not be treated as definitive, especially since it was contrary to the impression gained from observation. It is due to this shortcoming that it has not been possible to pursue the question of momentary maximising with respect to bouts.

The absence of this analysis means that the BCIs deduced from both Slater & Lester's (1982) and Fagen & Young's (1978) methods have not, after all, been put to any use. In some ways this is how it should be. Where at all possible analyses should be based on the objectively defined, optimally extracted parameters of the mean and proportion in mixture of each component distribution. The BCI will always depend, not just on the means of the processes , but also on the proportions they happen to be in in the mixture being considered. Although Slater & Lester's BCI is optimal in its partitioning of intervals, in the sense that it minimizes the number of IRTs misclassified, I feel that it misses the point. The purpose of a BCI is not to optimally partition intervals, but to summarize the series of intervals into its component bouts. In this respect Slater & Lester's criterion, as discussed above, would produce very few bouts of great length. I am therefore in favour of Fagen & Young's criterion because it keeps closer to the true proportions of between and within bout intervals, and therefore the true number of bouts in the series of events.

The major pretext for this analysis was McNamara & Houston's (1983) demonstration that anything but a single interresponse time on a VI schedule must be suboptimal. They have subsequently argued (Houston & McNamara 1985) that in general variability in behaviour runs contrary to the predictions of optimal foraging models. These models generally make predictions of a quite

unvarying character (Houston & McNamara 1985 consider the fixed giving up time of a predator feeding on depleting patches according to the marginal value theorem of Charnov 1976), which are then tested against *average* performance of real predators. Houston & McNamara suggest that the observed variability of behaviour be incorporated explicitly as a constraint into the optimality model. They show that the results of an analysis constrained in this way will differ from those of an unconstrained optimal foraging model, and caution that if the variability in behaviour is not taken into account in an optimality analysis then the predictions from that model should not be compared to the averaged observed behaviour.

This post hoc incorporation of constraints to make data fit the theory was rejected as unscientific in chapter two. Is there any good functional reason to expect variability in behaviour? One argument that has often been proposed to account for variability in behaviour is that it keeps the organism's options open to notice and respond to changes in the environment. Given that the costs of variability in behaviour can be estimated (Houston & McNamara 1985), and that, in principle at least, the variability of the environments in which a subject regularly finds itself might be estimated, then an overweening optimality model might be constructed to balance the costs of behavioural variability in the present environment, with the long term benefits of variability in responsiveness to a variety of environments. In practice, in the present context, it is very difficult to see any plausible context related reason for not just one exponential distribution, but the two distributions found on occasion (and, I feel, probably more widespread than could be demonstrated in these data).

If we accept that bouting of pecks is the result to expect here, then I submit that the most parsimonious way to account for this would be to view the bouting of pecks as a form of behavioural structure that is part of the organisms response to some part of the experimental situation. It may perhaps be a somewhat optimal manner to attack hard, food related objects; but I feel it is more likely, given the appearance of behaviour of this type in Machlis' chicks, to be a physiological constraint on the pecking movement, that rapid pecking must be regularly interspersed with longer pauses.

My point is not that subjects are not responsive to schedule values, it was shown in chapters 6, 7 and 8 that in various ways they are. The argument from

evolutionary epistemology is that, if a behaviour is to be modifiable with experience, then it must first be variable; and then it must be placed in the presence of a selection pressure to which it is vulnerable. The investigation presented in this chapter is just a first step towards an analysis in terms of variation and selection as developed in chapter 3. Had it proved possible to state with confidence the temporal structure of the pecking, it would have been interesting to have asked How (if at all) does this structure respond to the present experiment? That is to say, which aspects (if any) of the experimental procedure constitute a selection pressure on the pecking structure? These might include the key as peckable object; the temporal pattern of reinforcement, the confinement of the experimental chamber, the light behind the key, the grain given as reinforcement, etc. etc. Would the pecking structure have been the same at the beginning at the experiment as at the end? These questions await the combination of rigorous methods of data analysis, with more detailed recording of pecks. One thing, at least, is certain from these data, the behaviour is not optimal.

CHAPTER 10
CONCLUSION

10.1 Optimality

One result that has certainly not been found here is optimal behaviour. This is not surprising. The experiments described in this thesis have numerous artificialities to them. Additionally the optimality analysis used was exceptionally rigorous, and was applied critically. Although results from these unnatural experiments might be dismissed just because they were collected in an unnatural environment, I have tried to argue, particularly in chapter 5, that, a priori, it is not at all obvious what makes the present experiments so much more artificial than experiments that have been accepted just because they produced results which were (or, at least, which were presented as being) in line with predictions from optimal foraging theory. There is a danger here of defining the artificiality of an experiment post hoc, according to whether or not subjects foraged optimally in it. In any case, the birds used in these experiments spend their days in boxes of roughly the same dimensions as the experimental chambers, and simply peck up the grain that is put in. I note with interest that of the two experiments reported here, the one giving the better agreement with optimality predictions was the traditional operant psychology key pecking experiment, and the attempt to develop a new foraging environment with a simpler response requirement, only reduced the degree of agreement with prediction.

But, as I have discussed in chapter 2, failure to forage optimally should be the rule, not the exception. Failures to forage optimally do not refute the theory of evolution by natural selection - it is perfect optimisation that would be difficult to account for through natural selection. In the psychological literature there is no shortage of cases of animals failing to behave optimally (a recent case is responding on concurrent VI:VR [variable ratio] schedules, Herrnstein & Heyman 1979), however these have been easy to dismiss since they were collected in situations that paid little explicit attention to the organisms natural environment. Very recently, however, a few examples of sub-optimal behaviour have become apparent under conditions that are less easy to dismiss. The experiment of Shettleworth (1985b) discussed in chapter 4 showed that pigeons have a strong preference for many small items of food

over a few large items, even when the energetic value of the smaller items is substantially less than that of the larger ones. Shettleworth & Jordan (1986) have further shown that rats, given a choice between seeds with husks on, and seeds without husks presented after a delay equal to the time required to remove the husk, strongly prefer to dehusk the seeds themselves. This imposes on them an energetic cost with no compensating energetic gains. These results might be dismissed because they were collected in psychological choice procedures under schedules of reward drawn from the psychological literature. However they appear to have been well controlled studies and possess no obvious artificialities of procedure.

A study using wild born starlings in an aviary is reported by Inglis & Ferguson (1986). They gave each bird a choice between meal worms that were available in a bowl, and meal worms hidden in holes made in a board. These holes were all covered by flaps, but only 31 out of 126 holes contained worms. From an optimal foraging perspective there was no reason why the starlings should waste effort searching the board while identical prey were freely available in the bowl. Inglis & Ferguson found that a substantial proportion of worms were taken from the holes in the board (72% average for undeprived subjects), and that the proportion of worms taken from the holes decreased with increasing deprivation. These results are unaccountable from an optimality perspective, though they merely replicate a number of results in the psychological literature that Inglis & Ferguson cite.

I strongly suspect that more examples of sub-optimal behaviour will soon become apparent. The danger then may be that there could be a backlash against the use of optimality considerations in the study of animal behaviour. The point that has to be emphasized is that, as I argued in chapter 2, optimal foraging theory is not a theory of behaviour – we should not expect to observe globally optimal behaviour – however, that does not mean that optimal foraging theory does not have a role to play. This role includes pointing towards functional factors behind behaviour, and, by comparison with what animals actually do, suggesting the rules of thumb that may be guiding their behaviour.

10.2 Rules of Thumb

As for rules of thumb – the data presented in chapter 8 have not provided a great deal of support for them either. Advocates of both momentary maximising and melioration have been able to find more support for their theories than I have here. Perhaps part of the reason for that is implicit in the previous sentence – advocates are always better at finding supportive evidence than those indifferent to an issue. Melioration, however, has never before been tested on simple concurrent interval schedules as presented here, so the dismissal of this attractively simple theory does not stand in contradiction to any published work. I would not want to claim on the basis of these data that melioration *never* happens – as I suggest in chapter 8, it may well be a 'reserve' strategy for organisms capable of more complex assessments of their environment; and there may be organisms with limited cognitive capacities who use it as a strategy of 'first choice'.

My results from testing momentary maximising do stand in opposition to those of Hinson & Staddon (1983). Aside from the difference in species used here, two differences of procedure stand out as important. One is the use of a COD in the study reported in chapters 6 to 9. Momentary maximising does not cope very clearly with CODs. Any activity that has a zero probability of reward should not be engaged in according to a literal interpretation of momentary maximising. If, as I did, we choose to ignore COD and work round it in our estimation of momentary maximising, then subjects are likely to be discouraged from changing over by the COD in a manner that we have not accounted for in the momentary maximising analysis.

The second difference of procedure was that I did not attempt to calculate an averaging figure (a 'figure of merit'), for momentary maximising. I feel that the use of such devices is unwise because it hides the different ways in which momentary maximising can be achieved, and allows responses that are well placed with regard to momentary maximising to cancel out, to some extent, those that are misplaced.

Even if I were willing to dismiss my own results and concede that momentary maximising is the response controlling rule on these schedules, it could still not be a general process controlling choice in natural environments, simply because prior probabilities of reward cannot, in general, be defined. If

we were to generalize momentary maximising to be a rule demanding that a subject respond on the option possessing the highest subjective likelihood of reward, then we would have a notion possessing a wider area of applicability. The problem in applying such a rule would lie in defining 'subjective likelihood' without reference to the behaviour observed. The dynamic between considering expectancy based rules of thumb, of the momentary maximising type, and past reward 'linear operator' type rules, of the type to which melioration belongs, will I expect, form a constructive dialectic.

10.3 Variation and Selection

A personal disappointment in these experiments was the failure to find any way of assessing, let alone provide support for, the notions of blind variation and selective retention developed in chapter 3. As I mentioned at the end of that chapter, experiments investigating optimal foraging theory are not well suited to researching variation and selection of behaviour because of their excessive duration. However, if it had been possible to identify a rule of thumb confidently as the behaviour controlling rule under these conditions, a foothold would have been gained, such that a variation and selection analysis might have proved possible. In the event it was not even possible to determine with confidence the temporal patterning of responding, and so, as discussed in chapter 9, the many interesting questions that a blind variation and selective retention viewpoint brings to the fore, could not be addressed.

The consolation to these disappointments is that a great many interesting questions remain open for future research.

REFERENCES

Abarca, N & Fantino, E (1982) Choice and foraging *J. Experimental Analysis Behav.* **38** pp117- 123

Agha,M & Ibrahim, M.T. (1984) Algorithm AS203: Maximum likelihood estimation of mixtures of distributions *Applied Statistics* **33** pp327-332

Andrew, R.J. (1964) The development of adult responses from responses given during imprinting by the domestic hen. *Animal Behaviour* **12** pp542-548

Angermeier, W. F. (1984) *The evolution of operant learning* S. Karger: Basel

Barnard, C. J. (1985) Skinner box ecology: rules to forage by *Behavioural & Brain Sciences* **8** pp330-331

Baum, W. M. (1974) On two types of deviation from the matching law: bias and undermatching. *J. Experimental Analysis Behav.* **22** pp 231-242.

Baum, W. M. (1976) Time based and count based measures of preference *J. Experimental Analysis Behav.* **26** pp27-35

Baum, W. M. (1979) Matching, undermatching, and overmatching in studies of choice. *J. Experimental Analysis Behav.* **32** pp 269-281.

Baum, W. M. (1982) Choice, changeover and travel. *J. Experimental Analysis Behav.* **38** pp 35-49.

Baum, W. M. & Rachlin, H.C. (1969) Choice as time allocation *J. Experimental Analysis Behav.* **12** pp861-974

Bitterman, M. E. (1965) Phyletic differences in learning *American Psychologist* **20** pp396- 400

Blough, D. S. (1963) Interresponse time as a function of continuous variables: a new method and some data. *J. Experimental Analysis Behav.* **6** pp237-246

Boakes, R. (1984) *From Darwin to Behaviourism: psychology and the minds of animals* Cambridge U. P.

Bolles, R. C. (1970) Species-specific defense reactions and avoidance learning *Psy Rev* **77** pp32-48

Bolles, R. C. (1971) Species-specific defense reactions In Brush (Ed) *Aversive conditioning and learning*

Bolles, R. C. (1985) The slaying of Goliath; what happened to

reinforcement theory. In T. D. Johnston & A. T. Pietrewicz (Eds) *Issues in the ecological study of learning* Lawrence Erlbaum Associates: Hillsdale, N.J.

Bradshaw, C. M; Szababldi, E. & Bevan, P. (1976) Behavior of humans in variable interval schedules of reinforcement *J. Experimental Analysis Behav.* **26** pp135-41

Branch, M. N. (1985) Preference for a hypothesis: is the case "closed"? *Behav. Brain. Sci.* **8** pp332-3

Breland, K. & Breland, M. (1961) The misbehavior of organisms *Am Psychologist* **16** p661- 4

Breland, K. & Breland, M. (1966) *Animal Behavior* New York: Macmillan

Brown, P.L. & Jenkins, H. M. (1968) Auto-shaping of the pigeon's key peck *J. Experimental Analysis Behav.* **11** pp1- 8

Campbell, D. T. (1974) Evolutionary epistemology In Schilpp (Ed) *The philosophy of Karl Popper* ch. 12 Open Court: Illinois

Catania, A. C. (1966) Concurrent Operants. In W.K. Honig (Ed) *Operant Behavior: areas of research and application* Ch. 6 pp213-270 New York: Appleton-Century-Crofts.

Catania, A. C. & Reynolds, G. S. (1968) A quantitative analysis of the responding maintained by interval schedules of reinforcement *J. Experimental Analysis Behav.* **11** pp27-383.

Changeux, J.-P.; Heidmann, T. & Patte, P. (1984) Learning by selection In P. Marler & H. Terrace (Eds) *The biology of learning* pp115-133. Springer Verlag: Berlin

Charnov, E. L. (1976) Optimal Foraging: the marginal value theorem *Th. Pop. Biol.* **9** pp129-136

Chatfield, C. (1975) *The analysis of time series: an introduction* London: Chapman & Hall.

Church, R. M. (1980) Short term memory for time intervals *Learning & Motivation* **11** pp208-219

Clifton, P. G. (1979) Temporal patterns of feeding in the domestic chick: 1.ad libitum *Animal Behaviour* **27** pp811-820

Collier, G. H. & Rovee-Collier (1981) A comparative analysis of optimal foraging behavior: laboratory simulation In *Foraging behavior: ecological, ethological and psychological approaches* A.C. Kamil & T. Sargent (Eds) Garland Press

Collier, G. H. (1983) An ecological perspective of reinforcement and motivation In *Handbook of behavioral neurobiology* **6** E. Satinoff & P. Teitelbaum (Eds) Plenum Press

Couvillon, P. A. & Bitterman, M. E. (1985) Analysis of choice in honeybees *Animal Learning & Behavior* **13** pp246- 252

Cowie, R.J. (1977) Optimal foraging in great tits (Parus major) *Nature* **268** pp137- 139

Cox, D.R. & Lewis, P. A. W. (1966) *The statistical analysis of series of events* London: Methuen.

Culshaw, A. D. & Broom, D. M. (1980) The imminence of behavioural change and startle responses of chicks *Behaviour* **73** pp64-76

Davison, M. C. (1972) Preference for mixed interval versus fixed interval schedules: Number of component intervals *J. Experimental Analysis Behav.* **17** pp169-176

Davison, M. C. (1985) Foraging for a science of behavior *Behav. Brain Sci.* **8** pp335-336

Dawkins, R. (1981) *The extended phenotype* Oxford University Press: Oxford

Delius, J. D. (1969) A stochastic analysis of the behaviour of skylarks *Behaviour* **33** pp137-178

Dickinson, (1980) *Contemporary animal learning theory* Cambridge U. P.

Dixon,W.J.; Brown,M.B.; Engelman,L; Frane,J.W.; Hill,M.A.; Jennrich, R.I. & Toporek, J.D. (1983) *BMDP Statistical Software: 1983* Univ. of California Press: Berkeley.

Domjan, M. (1983) Biological constraints on instrumental and classical conditioning 10 years later: implications for general process theory. In G. H. Bower (Ed) *The psychology of learning and motivation* **17** New York: Academic Press

Domjan, M. & Galef, B. G. Jr. (1983) Biological constraints on instrumental and classical conditioning: retrospect and prospect *Animal Learning & Behav.* **11** pp151- 62

Duncan, I. J. H; Horne, A. R; Hughes, B. O & Wood-Gush, D. G. M. (1970) The pattern of food intake in female brown leghorn fowls as recorded in a Skinner box. *Animal Behaviour* **18** pp245-255

Eccles, J. C. (1977) An instruction selection theory of learning in cerebellar cortex *Brain Research* **127** pp327- 52

Fagen, R. M. & Young, D. Y. (1978) Temporal patterns of behavior: durations, intervals, latencies, and sequences. In *Quantitative Ethology* P.W. Colgan (Ed) Wiley: N.Y.

Fantino, E. & Abarca, N (1985) Choice, optimal foraging, and the delay-reduction hypothesis *The behavioral and Brain sciences* **8** pp315- 330

Feekes, F. (1972) "Irrelevant" ground pecking in agonistic situations in burmese red junglefowl (Gallus gallus spadiceus) *Behaviour* **43** pp186-326

Ferster, C. B. & Skinner, B. F. (1957) *Schedules of Reinforcement* Appleton-Century-Crofts: New York.

Findley, J. D. (1958) Preference and switching under concurrent scheduling *J. Experimental Analysis Behav.* **1** pp123-144

Flescher, M. & Hoffman, H. S. (1962) A Progression for generating variable interval schedules *J. Experimental Analysis Behav.* **5** pp529-530

Flugel, J. C. (1951) *A hundred years of psychology* 2nd Edition. Duckworth: London

von Frisch, K. (1967) *The dance language and orientation of bees* Trans by L. E. Chadwick. Cambridge Mass: Harvard U. P.

Garcia, J. & Koelling (1966) Relation of cue to consequence in avoidance learning *Psychonomic Science* **4** pp123-4

Garcia, J.; Clarke, J. C. & Hankins, W. G. (1973) Natural responses to scheduled rewards In *Perspectives in ethology* **1** pp1-42 Plenum Press: London

Gass, C. L. (1985) Reaching for an integrated science of behavior *Behav. Brain Sci.* **8** p337

Gibb, J. A. (1962) L. Tinbergen's hypothesis of the role of specific search images *Ibis* **104** p106- 111

Gibson, J. J. (1977) The theory of affordances In R. Shaw & J. Bransford (Eds) *Perceiving, acting and knowing: toward an ecological psychology* Lawrence Erlbaum associates: Hillsdale N. J.

Gibson, J. J. (1979) *The ecological approach to visual perception* Boston: Houghton- Mifflin

Gottlieb, G. (1985) Anagenesis: theoretical basis for the ecological void in comparative psychology In T. D. Johnston & A. T. Pietrewicz (Eds) *Issues in the ecological study of learning* pp59- 72 Lawrence Erlbaum Associates: Hillsdale, N. J.

Gould, S. J. (1980) *The panda's thumb: more reflections in natural history* Penguin Books: Harmondsworth

Gould, S. J. & Lewontin, R. C. (1979) The spandrels of San Marco and the Panglossian paradigm *Proc. Royal Soc. London* **B 205** pp581-98

Green, L.; Rachlin, H. & Hanson, J. (1983) Matching and maximising with concurrent ratio-interval schedules *J. Experimental Analysis Behav.* **40** pp217-224

Green, R. F. (1985) An interdisciplinary approach to foraging

behavior *Behav. Brain Sci.* **8** pp338

Hailman, J. P. (1985) Historical notes on the biology of learning In T. D. Johnston & P. A. Pietrewicz (Eds) *Issues in the ecological study of learning* pp27-57 Lawrence Erlbaum Associates: Hillsdale, N.J.

Harley, C. B (1981) Learning the evolutionary stable strategy. *J. Theor. Biol* **89** pp611-633

Harley, C. B. (1983) When do animals learn the evolutionary stable strategy? *J. Theor. Biol.* **101** pp179- 81

Harley, C. B & Maynard Smith, J. (1983) Learning - an evolutionary approach *Trends Neurosci* **6** pp204

Harlow, H. F. (1959) Learning set and error factor theory In S. Koch (Ed) *Psychology: a study of a science* **2** McGraw Hill: N.Y.

Harzem, P.; Lowe, C. F. & Priddle-Higson, P.J (1978) Inhibitory function of reinforcement: Magnitude effects on variable interval schedules *J. Experimental Analysis Behav.* **30** pp1-10

Hasselblad, V. (1966) Estimation of parameters for a mixture of normal distributions *Technometrics* **8** pp431-444

Hasselblad, V. (1969) Estimation of finite mixtures of distributions from the exponential family *J. American Statistical Association* **64** pp1459-1471

Heiligenberg, W. (1965) A quantitative analysis of digging movements and their relationship to aggressive behaviour in cichlids *Animal Behaviour* **13** pp163-170

Herrnstein, R. J. (1961) Relative and absolute strength of response as a function of frequency of reinforcement *J. Experimental Analysis Behav.* **4** pp267-272

Herrnstein, R. J. (1974) Formal properties of the matching law *J. Experimental Analysis Behav.* **21** pp159- 164

Herrnstein, R. J. (1982) Melioration as behavioral dynamism In Commons, Herrnstein & Rachlin (Eds) *Quantitative analyses of behavior II: matching and maximising accounts*

Herrnstein, R. J. & Heyman, G. M. (1979) Is matching compatable with reinforcement maximisation on concurrent VI:VR ? *J. Experimental Analysis Behav.* pp209-23

Herrnstein, R. J. & Vaughan, W. (1980) Melioration and behavior allocation In J.E.R. Staddon (Ed) *Limits to action* N.Y.: Academic Press

Heyman, G. M. (1979) Matching and maximising in concurrent schedules. *Psy. Rev.* **86** pp496-500

Heyman, G. M. & Luce, R. D. (1979) Operant matching is not a

logical consequence of maximising reinforcement rate *Animal Learning & Behav.* **7** pp133-140

Hinde, R. A (1982) *Ethology* Fontana

Hinson, J. N. & Staddon, J. E. R. (1981) Maximisation on interval schedules <u>In</u> Bradshaw, Szabaldi & Lowe (Eds) *Quantification of steady state operant behaviour* Elsevier: North Holland.

Hinson, J. N. & Staddon, J. E. R. (1983) Hill climbing by pigeons *J. Experimental Analysis Behav.* <u>39</u> pp25-47

Hofstadter, R. (1955) *Social darwinism in american thought* Beacon Press: Boston

Horner, J. M. & Staddon, J.E.R. (1986) Simple dynamic choice. manuscript

Hosmer, D. W. (1973) A comparison of iterative maximum likelihood estimates of the parameters of a mixture of two normal distributions under three different types of sample *Biometrics* **29** pp761-770

Houston, A. I. (1980) Godzilla v. the creature from the black lagoon: Ethology v. psychology <u>In</u> F. M. Toates & T. R. Halliday (Eds) *Analysis of motivational processes* Academic Press: London

Houston, A. I. (1985) Choice and preference - you can't always want what you get *Behav. Brain Sci.* **8** pp339- 40

Houston, A. I. (1986) The control of foraging decisions. <u>to appear in</u> *6th Harvard Symposium on quantitative analyses of behavior*

Houston, A. I. & McNamara, J. (1981) How to maximise reward rate on two variable interval paradigms *J. Experimental Analysis Behav.* <u>35</u> pp367-396

Houston, A. I. & McNamara, J. M. (1985) The variability of behaviour and constrained optimization *J. Theor. Biol.* **112** pp265-273.

Hull, D. L. (1982) The naked meme <u>In</u> H. C. Plotkin (Ed) *Learning, development and culture* ch14 pp273- 327 Wiley: Chichester

Inglis, I.R. & Ferguson, N. J. K. (1986) Starlings search for food rather than eat freely-available, identical food. *An. Behav.* **34** pp614-6

James, W. (1961) *Psychology the briefer course* Harper & Row: New York. First published 1896

James, W. (1979) Great men and their environment <u>In</u> *The will to believe and other essays in popular philosophy* Harvard U. P.: Cambridge Mass. Essay first published as 'Great men, great thoughts, and the environment', 1880.

James, W. (1981) *The principles of psychology* Harvard U. P. :

Cambridge Mass. First published 1890

Johnston, T. D. (1981) Contrasting approaches to a theory of learning *Behav. Brain Sci.* **4** pp125-73

Johnston, T. D. (1985) Introduction: conceptual issues in the ecological study of learning In T. D. Johnston & A. T. Pietrewicz (Eds) *Issues in the ecological study of learning* Lawrence Erlbaum Associates: Hillsdale, N. J.

Kacelnik, A (1979) *Studies of foraging behaviour and time budgeting in Great Tits (Parus major)* Unpublished D.Phil thesis: Oxford University

Kacelnik, A & Krebs, J. R. (1985) Learning to exploit patchily distributed food In R.M. Sibly & R.H. Smith (Eds) *Behavioural Ecology: ecological consequences of adaptive behaviour* Blackwell: Oxford

Kacelnik,A; Krebs, J. R. & Ens, B (1986) Foraging in a changing environment: an experiment with Starlings (Starnus vulgaris). In M.L. Commons; A. Kacelnik & S. Shettleworth (Eds) *The Harvard Symposium on quantitative analyses of behaviour* **6** *Foraging* Laurence Erlbaum Associates (In Press)

Kalat, J. W. (1985) Taste aversion learning in ecological perspective In T. D. Johnston & A. J. Pietrewicz (Eds) *Issues in the ecological study of learning* pp119- 41. Lawrence Erlbaum: Hillsdale, N. J.

Kamil, A. C. & Roitblat, H. L. (1985) The ecology of foraging behavior: implications for animal learning and memory. *Annual Review of Psychology* **36** pp141- 169

Krebs, J. R. (1973) Behavioural aspects of predation In P. P. G. Bateson & P. H. Klopfer (Eds) *Perspectives in ethology* **1** pp73- 107

Krebs, J. R. (1978) Optimal foraging: decision rules for predators In J. R. Krebs & Davies (Eds) *Behavioural ecology: an evolutionary approach* pp23-63

Krebs, J.R., Ryan, C. & Charnov, E.L. (1974) Hunting by expectation or optimal foraging? A study of patch use by chickadees *Animal Behaviour* **22** pp953- 964

Krebs, J.R.; Erichson, J.T.; Webber, M.I. & Charnov, E.L. (1977) Optimal prey selection in the great tit (Parus major) *Animal Behaviour* **25** pp30- 38

Krebs, J.R.; Kacelnik, A. & Taylor, P. (1978) Tests of optimal sampling by foraging great tits *Nature* **275** pp27- 33

Krebs, J. R.; Stevens, D. W. & Sutherland, W. J. (1983) Perspectives in optimal foraging In A. H. Brush & G. A. Clark (Eds) *Perspectives in optimal foraging* Cambridge University Press: Cambridge

Krebs, J. R. & Davies, N. B. (1984) *Behavioural ecology* 2nd Edition. Blackwells: Oxford

Krebs, J. R. & McCleery, R.H. (1984) Optimisation in behavioural ecology In J. R. Krebs & Davies (Eds) *Behavioural Ecology 2nd Edition*

Krechevsky, I. (1932) "Hypothesis" behaviour in rats *Psy Rev* **39** pp516- 32

Krechevsky, I (1967/1932) "Hypothesis" vs "Chance" in the presolution period in sensory discrimination learning *U. of California Publ. in psychology* **6** p27-44 Reprinted in A. J. Riopelle (Ed) *Animal problem solving:selected readings* Penguin books: Harmondsworth

Kruse, J. M. (1985) Alternative approaches to the psychology of foraging *Behav. Brain Sci.* **8** pp342-3

Larkin, S. (1981) *Time and energy in decision making* Unpubd. D.Phil thesis Oxford University; Oxford

Lashley, K. S. (1951) The problem of serial order in behaviour In L.A. Jeffress (Ed) *Cerebral mechanisms in behavior* Wiley: N.Y. pp112-136

Lea, S. E. G. (1979) Foraging and reinforcement scheduling in the pigeon: optimal and non-optimal aspects of choice *Animal Behaviour* **27** pp875- 886

Lea, S. E. G. (1984) Complex general process learning in non-mammalion vertebrates In P. Marler & H. S. Terrace (Eds) *The biology of learning* Berlin: Springer Verlag

Lea, S. E. G. (1985) Optimality: sequences, variability, learning *Behav. Brain Sci.* **8** p343

Lejeune, H. & Richelle, M. (1981) Temporal regulation in the turtle dove: a comparative study In Bradshaw, Szabaldi & Lowe (Eds) *Quantification of steady state operant behaviour*

Leslie, J. C. (1981) Effects of variations in local reinforcement rate on local response rate in variable interval schedules. *J. Experimental Analysis Behav.* **35** pp45-53

Lester, N. P. (1984) The "feed:feed" decision. How goldfish solve the patch depletion problem *Behaviour* **89** pp174-199

Levine, M. (1967/1959) A model of hypothesis behaivor in discrimination learning set *Psy Rev* **66** pp353-66. Reprinted in A. J. Riopelle (Ed) *Animal problem solving: selected readings* Penguin Harmondsworth.

Lewontin, R. C. (1979) Sociobiology as an adaptationist program *Behavioral science* **24** pp5-14

Lorenz, K. (1969) Innate bases of learning In K. H. Pribram (Ed)

On the biology of learning Harcourt, Brace & World, Inc: New York

Lorenz, K. (1971/1959) Gestalt perception as a source of scientific knowledge (originally published 1959) In *Studies in animal and human behaviour* **2** K. Lorenz trans R. Martin. Harvard U.P.:Cambridge Mass.

Lorenz, K. (1975/1941) kant's doctrine of the a priori in the light of contemporary biology In *Konrad Lorenz* Evans (Ed)

Lorenz, K. (1977/1973) *Behind the mirror* Methuen & Co.:London. Published in German 1973

MacArthur, R. H. & Pianka, E. R. (1966) On optimal use of a patchy environment *Am. Nat.* **100** pp603-9

McFarland, D. J. (1977) Decision making in animals *Nature* **269** pp15-21.

McFarland, D. J. (1985) *Animal behavior: psychobiology, ethology and evolution* Benjamin/Cummings Pub. Co. Inc.: Menlo Park, Calif.

McFarland, D. J. & Houston, A. (1981) *Quantitative ethology: the state space approach* Pitman books: London

Mackintosh, N. J. (1974) *The psychology of animal learning* Academic Press

McNair, J. N. (1985) Optimal foraging for operant conditioners *Behav. Brain Sci.* **8** pp343-344

McNamara, J. M. & Houston, A. I. (1983) Optimal responding on variable interval schedules. *Behav. Anal. Let.* **3** pp157-170.

McNamara, J. M. & Houston, A. I. (1985) Optimal foraging and learning *J. theor. Biol.* **117** pp231- 49

McSweeney, F. K. (1975) Matching and contrast on serveral concurrent treadle-press schedules *J. Experimental Analysis Behav.* **23** pp193-8

Machlis, L. (1977) an analysis of temporal patterning of pecking in birds *Behaviour* **63** pp1-70

Marler, P. (1970) A comparative approach to vocal learning *J. Comp. Physiol. Psychol.* **71** p1-25

Marler, P. (1984) Song learning: innate species differences in the learning process In P. Marler & H. Terrace (Eds) (1984) *The biology of learning* Berlin: Springer Verlag

Marler, P. & Terrace, H. S. (Eds) (1984) *The biology of learning: report of the Dahlem workshop on the biology of learning.* Berlin: Springer Verlag

Maynard Smith, J. (1978) Optimization theory in evolution *Ann. Rev. of ecology and systematics* **9** p31-56

Mechner, F. (1958) Probability relations within response sequences under ratio reinforcement *J. Experimental Analysis Behav.* **1** pp109-121

Mellgren, R. L. (1985) Outcome and mechanism in foraging *Behav. Brain Sci.* **8** pp344-5

Miller, D. B. (1985) Methodological issues in the ecological study of learning In T. D. Johnston & A. T. Pietrewicz (Eds) *Issues in the ecology study of learning* Lawrence Erlbaum assoc: Hillsdale, N. J.

Moore, J. (1984) Choice and transformed inter-reinforcement intervals *J. Experimental Analysis Behav.* **42** pp321-335.

Mountcastle, V. B. & Edelman, G. M. (1978) *The mindful brain* MIT press: Cambridge Mass

Myers, D. L. & Myers, L. E. (1977) Undermatching: a reappraisal of performance on concurrent variable interval schedules of reinforcement. *J. Experimental Analysis Behav.* **25** pp203-214.

Nelson, K. (1964) The temporal patterning of courship behaviour in galndulocaudine fishes (Ostariophysi, Characidae) *Behaviour* **24** pp90-145

Nevin, J. A. (1969) Interval reinforcement of choice in discrete trials *J. Experimental Analysis Behav.* **12** pp875-885

Nevin, J. A. (1979) Overall matching versus momentary maximising: Nevin (1969) revisited. *J. Exp. Psy.: An. Behav. Proc.* **5** pp300-306

· Norman, W. D. & McSweeney, F. K. (1978) Matching, contrast and equalising in the concurrent lever press responding of rats *J. Experimental Analysis Behav.* **299** pp453-462

O'Neil, W. M. (1968) *The beginnings of modern psychology* Penguin: Harmondsworth

Oldroyd, D. R. (1980) *Darwinian impacts* Open Univeristy press: Milton Keynes.

Ollason, J. G. (1980) Learning to forage - optimally? *Theor. Pop. Biol.* **18** pp44-56

Ollason, J. G. (1986) Artificial design in natural history: why it's so easy to understand animal behaviour in press *Perspectives in Ethology* **6**

Olton, D. S. (1978) Characteristics of spatial memory In S. H. Hulse; H. Fowler & W. K. Honig (Eds) *Cognitive processes in animal behaviour* Hillsdale, N. J.: Erlbaum

Oster, G. F. & Wilson, E. O. (1978) *Caste and ecology in the social insects* Princeton: Princeton university press

Passmore, J. A. (1957) *A hundred years of philosophy* London

Peden, B. F. & Rohe, M. S. (1984) Effects of search cost on foraging and feeding: a three component chain analysis *J. Experimental Analysis Behav.* **42** pp211- 221

Petersen, S. (1976) The temporal patterns of feeding over the oestrous cycle of the mouse *Animal Behaviour* **24** pp939-955

Petersen, S. (1978) Effects of testosterone upon feeding in male mice *Animal Behaviour* **26** pp945-952

Pierce, G. J. & Ollason, J. G (1986) *Optimal foraging is dead: long live optimal foraging?* Poster presented to 11th Scottish conference on animal behaviour. Obtainable from authors at Culterty Field Station, University of Aberdeen, Newburgh, Ellon, Aberdeenshire, AB4 0AA

Pliskoff, S. S (1971) Effects of symmetrical and asymmetrical changeover delays on concurrent performances *J. Experimental Analysis Behav.* **16** pp249- 256

Pliskoff, S. S; Cicerone, R & Nelson, T. D. (1978) Local response rate constancy on concurrent variable interval schedules of reinforcement *J. Experimental Analysis Behav.* **29** pp431-446

Plotkin, H. & Odling-Smee, J. (1979) Learning, change and evolution *Advances in the study of behaviour* **10** pp1-41

Plotkin, H. & Odling-Smee, J. (1981a) A multiple level model of evolution and its implications for sociobiology *Behav. Brain Sci.* **4** pp225- 35

Plotkin, H. & Odling-Smee, J. (1981b) Possible mechanisms for a multiple-level model of evolution: authors' response *Behav. Brain Sci.* **4** pp257- 68

Popper, K. R. (1972) *Objective knowledge: an evolutionary approach* Oxford University Press: Oxford.

Popper, K. R. (1980) *The logic of scientific discovery* New edition. First published in German 1934

Popper, K. R. (1983) Evolutionary epistemology In D. Miller (Ed) *A pocket Popper* Fontana paperbacks: Oxford. Essay first published 1975.

Pringle, J. W. S. (1951) On the parallel between learning and evolution *Behaviour* **3** pp175- 215

Pyke, G. H.; Pulliam, H. B. & Charnov, E. L. (1977) Optimal foraging: a selective review of theory and tests *Quart. Rev. Biol.* **52** pp137-54

Rachlin, H. (1978) A molar theory of reinforcement models *Journal of the Experimental Analysis of Behavior* **30** pp345-360

Rachlin, H. & Burkhard, B. (1978) The 'temporal triangle': response substitution in instrumental conditioning. *Psychological Review* **85** pp22-44.

Rachlin, H.; Battalio; Kagel & Green (1981) Maximization theory and behavior. *Behav & Brain Sci* **4** pp371- 388

Real, L. A. (1980) On uncertainty and the law of diminishing returns in evolution and behavior In J.E.R. Staddon (Ed) *Limits to action* A.P.: N.Y.

Real, P. G. (1983) A time series analysis of changeover performance on concurrent variable interval schedules. *Animal Learning & Behav.* **11** pp255-265

Real, P. G & Dreyfus, L. R. (1985) Levels of aggregation: Relative time allocation in concurrent schedule performance. *Journal of the Experimental Analysis of Behavior* **43** pp97-114.

Reed, E. S. (1985) An ecological approach to the evolution of behavior In Johnston & Pietrewicz (Eds) *Issues in the ecological study of learning* Lawrence Erlbaum assoc.: Hillsdale, N. J.

Reid, R. G. B. (1985) *Evolutionary theory: the unfinished syntesis* Croom Helm: Beckenham Kent

Rescorla, R. A. & Furrow, D. R. (1977) Stimulus similarity as a determinant of pavlovian conditioning *J. Exp. Psy.:An. Behav. Proc* **3** p203-15

Revusky, S (1977) Learning as a general process with an emphasis on data from feeding experiments In Milgram, Krames & Alloway (Eds) *Food aversion learning* **1** pp1- 51

Revusky, S (1984) Associative predispositions In P. Marler & H. S. Terrace (Eds) *The biology of learning* Berlin: Springer Verlag

Revusky, S (1985) The general process approach to animal learning In Johnston & Pietrewicz (Eds) *Issues in the ecological study of learning* Lawrence Erlbaum associates: Hillsdale, N. J.

Riedl, R. (1984) *Biology of knowledge: the evolutionary basis of reason* John Wiley: Chichester. Translated from the 3rd German edition 1981

Russell, B. (1914) *Our knowledge of the external world as a field for scientific method in philosophy* Open Court: London

Russell, W. M. S. (1962) The analogy between organic and individual behavioral evolution, and the course of intelligence *General systems* **7** pp157- 193

Schneider, B. A. (1969) A two state analysis of fixed-interval

responding in the pigeon *Journal of the Experimental Analysis of Behavior* **12** pp677-687

Schwartz, B. (1980) Development of complex, stereotyped behavior in pigeons *Journal of the Experimental Analysis of Behavior* **33** pp153-160

Schwartz, B. (1985) On the organization of stereotyped response sequences *An. Learning & Behav.* **13** pp261-268

Scoun, J. (1983) *Changeover delay and concurrent schedules* Unpublished PhD. Thesis. University of Waikato: New Zealand.

Seligman, (1970) On the generality of the laws of learning *Psychological Review* **77** pp406-418

Shettleworth, S. J. (1972) Constraints on learning In *Advances in the study of learning* **4**

Shettleworth, S. J. (1984) Natural history and evolution of learning in nonhuman mammals In P. Marler & H. S. Terrace (Eds) *The biology of learning* pp419- 33

Shettleworth, S. J. (1985a) Questions about foraging *Behavioral and Brain Sciences* **8** pp347- 348

Shettleworth, S. J. (1985b) Handling time and choice in pigeons *Journal of the Experimental Analysis of Behavior* **44** pp139- 156

Shettleworth, S. J. & Jordan, V. (1986) Rats prefer handling food to waiting for it *An. Behav.* **34** pp925-7

Shimp, C. P.(1966) Probabilistically reinforced choice behavior in pigeons *Journal of the Experimental Analysis of Behavior* **9** PP443-455

Shimp, C. P. (1968) Magnitude and frequency of reinforcement and frequencies of interresponse times *Journal of the Experimental Analysis of Behavior* **11** pp525-535

Shimp, C. P. (1969) Optimal behavior in free-operant experiments *Psychological Review* **76** pp97-112

Shimp, C. P. (1970) Concurrent reinforcement of two IRTs: absolute rate of reinforcement *Journal of the Experimental Analysis of Behavior* **13** pp1-8

Shimp, C. P. (1979) the local organization of behavior: method and theory In P.Harzem & M.D. Zeiler (Eds) *Advances in behaviour* **1** Wiley: N. Y.

Shimp, C. P. (1981) The local organization of behavior: discrimination of and memory for simple behavioral patterns *Journal of the Experimental Analysis of Behavior* **36** pp303-315

Shimp, C. P (1982a) On metaknowledge in the pigeon: an organisms knowledge about its own behaviour *Animal Learning & Behav.* **10** pp358-364

Shimp, C. P (1982b) Choice and behavioral patterning *Journal of the Experimental Analysis of Behavior* **37** pp157- 169

Shimp, C. P (1982c) Reinforcement and the local organisation of behavior In Commons, Herrnstein & Rachlin (Eds) *Quantitative analyses of behavior II: matching and maximising accounts*

Sibly, R. M. & McFarland, D. J. (1974) A state space approach to motivation In D.J. McFarland (Ed) *Motivational control systems analysis* Academic Press: London & N.Y. pp213- 250

Silberberg, A & Fantino, E (1970) Choice, rate of reinforcement, and the changeover delay *Journal of the Experimental Analysis of Behavior* **13** pp187- 197

Silberberg, A. & Ziriax, J. M. (1985) Molecular maximizing characterizes choice on Vaughan's (1981) procedure. *Journal of the Experimental Analysis of Behavior* **43** pp83-96

Sibly, R. M. & McFarland, D. J. (1974) A state space approach to motivation In *Motivation control systems analysis* D.J.McFarland (Ed) pp213- 250 Academic Press :London & New York

Simon, H. A. (1956) Rational choice and the structure of the environment *Psy Rev* **63** pp129-38

Skinner, B. F. (1938) *The behavior of organisms* Appleton-Century-Crofts: New York

Slater, P. J. B. (1973) Describing sequences of behaviour In P.P.G. Bateson & Peter H. Klopfer (Eds) *Perspectives in ethology* **1** pp131-151

Slater, P. J. B. (1974a) The temporal pattern of feeding in the zebra finch *Animal Behaviour* **22** pp506-515

Slater, P. J. B. (1974b) Bouts and gaps in the behaviour of zebra finches, with special reference to preening *Rev Comp An* **8** pp47-61

Slater, P. J. B. (1975) Temporal patterning and the causation of bird behaviour In P. Wright; P. Caryl & D. Vowles (Eds) *Neural and endocrine aspects of behaviour in birds*

Slater, P. J. B. & Lester, N. P. (1982) Minimizing errors in splitting behaviour into bouts *Behaviour* **79** pp153-161

Spencer, H. (1855) *The principles of psychology* Longman: London. 1st edition

Staddon, J. E. R. (1975) Learning as adaptation In W. K. Estes (Ed) *Handbook of learning and cognitive processes* **2** pp37- 98 Erlbaum: Hillsdale, N. J.

Staddon, J. E. R. (1980) Optimality analyses of operant behavior and their relation to optimal foraging In J.E.R. Staddon (Ed) *Limits*

to action A.P.: N.Y.

Staddon, J. E. R. (1983) *Adaptive behavior and learning* Cambridge U. P. : N. Y.

Staddon, J. E. R. (1986) Learning as inference <u>To</u> <u>appear</u> <u>in</u> R. C. Bolles & M. D. Beecher (Eds) *Evolution and learning* Lawrence Erlbaum associates: Hillsdale, N. J.

Staddon, J. E. R.; Hinson, J. M. & Kram, R. (1981) Optimal choice *Journal of the Experimental Analysis of Behavior* **35** pp397-412

Staddon, J. E. R. & Motheral, S. (1978) On matching and maximising in concurrent choice experiments *Psychological Review* **85** pp436-444.

Taylor, R. & Davison, M. (1983) Sensitivity to reinforcement in concurrent arithmetic and exponential schedules *Journal of the Experimental Analysis of Behavior* **39** pp191-198

Terrace, H. S. (1984) Animal learning, ethology, and biological constraints <u>In</u> P. Marler & H. S. Terrace (Eds) *The biology of learning* Berlin: Springer Verlag

Testa, T. J. & Ternes, J. W. (1977) Specification of conditioning mechanisms in the modification of food preferences <u>In</u> L. M. Barker; M. R. Best & M. Domjan (Eds) *Learning mechanisms in food selection* Waco, Texas: Baylor U. P.

Thompson, D'Arcy W. (1917) *On growth and form* Cambridge U. P.

Thomson, R. (1968) *The pelican history of psychology* Penguin books: Harmondsworth

Thorndike, E. L. (1911) *Animal Intelligence* Macmillan: N. Y.

Thorndike, E. L. (1932) *The fundamentals of learning* Teacher's college: N. Y.

Thorndike, E. L. (1967/1898) Animal Intelligence *Psych. Rev. Monogr. Suppl* **2** pp1- 9. Reprinted in A. J. Riopelle (Ed) *Animal problem solving: selected readings* Penguin books

Tinbergen, N. (1951) *The study of instinct* Oxford U. P.: Oxford

Toal, L. & Leslie, J. C. (1984) Temporal control by signals of interval duration within variable interval schedules *J. Experimental Analysis Behav.* <u>42</u> pp267-77

Tolman & Krechevsky, I. (1933) Means-end-readiness and hypothesis - a contribution to comparative psychology *Psy Rev* **40** pp60-70

Trevett, A. J.; Davison, M. C. & Williams, R. J. (1972) Performance in concurrent interval schedules *Journal of the Experimental Analysis of Behavior* **17** pp369-374

Tustin, R. D. & Davison, M. C. (1978) Technical note: distribution of response ratios in concurrent VI performance *Journal of the Experimental Analysis of Behavior* **29** pp561- 564

von Uexküll, J. (1926) *Theoretical biology* Kegan Paul, Trench, Trubner & Co. : London

Vaughan, W. Jr. (1981) Melioration, matching and maximization *Journal of the Experimental Analysis of Behavior* **36** pp141-149

Vaughan, W. Jr. & Miller, H. L. Jr. (1984) Optimization vs response strength accounts of behavior *Journal of the Experimental Analysis of Behavior* **42** pp337-48

Vollmer, G (1975) *Evolutionäre Erkenntnistheorie* Stuttgart: Hirzel

Warren, J. M. (1974) Learning in vertebrates In Dewsbury & Rethlingshlaffer (Eds) *Comparative psychology: a modern survey* pp471-501 McGraw-Hill: New York

Watson, J. (1924) *Behaviorism* Norton: N. Y.

Wearden, J. H. (1983) Undermatching and overmatching as deviations from the matching law *Journal of the Experimental Analysis of Behavior* **40** pp332- 40

Wearden, J. H. & Burgess, I. S. (1982) Matching since Baum (1979) *Journal of the Experimental Analysis of Behavior* **38** pp339-348.

Weiss, B.; Laties, V. G.; Siegel & Goldstein (1966) A computer analysis of serial interactions in spaced responding. *Journal of the Experimental Analysis of Behavior* **9** pp619-626

Whalen, (1971) The concept of instinct In McGaugh (Ed) *Psychobiology*

Wheately, K. L. & Engberg, L. A. (1978) Choice performance in several concurrent key-peck treadle-press reinforcement schedules *Journal of the Experimental Analysis of Behavior* **29** pp181- 190

Wiepkema, P. R. (1968) Behaviour changes in CBA mice as a result of one goldthioglucose injection *Behaviour* **32** pp179-210

Williams, D. R. & Williams, H. (1969) Auto-maintenance in the pigeon: sustained pecking despite contingent non-reinforcement *Journal of the Experimental Analysis of Behavior* **12** pp511-20

Wuketits, F. M. (1986) Evolution as a cognition process: towards an evolutionary epistemology *Biology and Philosophy* **1** pp191- 206

Ydenberg, R. C. (1984) Great tits and giving up times: decision rules for leaving patches *Behaviour* **90** p1- 24

Young, J. Z. (1976) Memory and its models In H. Harris (ed) *Scientific models and man* Oxford U. P.

Young, J. Z. (1979) Learning as a process of selection and amplification *J. Roy. Soc. of Med.* **72** pp801- 14

Zach, R & Falls, J. B. (1976) Ovenbird (Aves: Paralidae) hunting behavior in a patchy environment: an experimental study *Canadian Journal Zoology* **54** pp1863- 1879

Ziegler,H.P; Green,H.L. & Lehrer, R (1971) Patterns of feeding behaviour in the pigeon *J. Comp & Phys Psy* **76** pp468-477