BODLEIAN LIBRARY OXFORD
DEPOSITED
THESIS
17-4-80

STUDIES OF FORAGING BEHAVIOUR AND TIME BUDGETING IN GREAT TITS (*Parus major*)

by

Alejandro Kacelnik

(Wolfson College)

A thesis submitted for the degree of

Doctor of Philosophy

at the

University of Oxford

Michaelmas Term 1979

a Li,con amor

a Federico,con bronca

a Oliver,con esperanza

Studies of foraging behaviour and time budgeting in Great Tits (*Parus major*).

Alejandro Kacelnik, D.Phil.. University of Oxford, 1979.

Part One is an investigation of the dawn chorus in the Great Tit. The major conclusions are as follows: Foraging and singing are partly incompatible because of the temporal and spatial separation of microhabitat conditions required for each activity. The dawn chorus occurs at a time of day when: a) foraging efficiency is limited by low light intensity and possibly by lack of mobility of the prey due to low temperature; b) acoustic comunication is more efficient than visual displays because of poor visibility and an enhanced sound transmission; c) the rate of territorial intrusion is at its daily peak. Great Tit's readiness to sing and react agonistically towards a territorial intruder is enhanced by experimentally lowering profitability of feeding, but it does not appear to be affected by time of day. Restricting access to food early in the morning leads to higher body weight in the evenings as an anticipatory compensation.

Part Two is a study of sampling strategies based on the statistical decision paradigm known as the Two-Armed Bandit problem. I studied transition behaviour in foraging experiments using concurrent variable ratio schedules. The birds' foraging behaviour approximated the predictions of a dynamic programing algorithm that calculates the optimal balance between exploring and exploiting for finite time horizons. The birds' response to patch diversity and time horizon was closer to molar maximization (total rewards over a certain period) than to molecular maximization (maximising the instantaneous probability of reward). When the environment is not totally stable these two possibilities conflict, and my results supported molar maximization.

Part Three investigates the trade-off between maximising foraging efficiency and efficient territorial defence. Great Tits behaved close to the predictions of the Marginal Value model when there were no territorial intrusions, but modified their feeding behaviour in the predicted way when the probability of intrusion was higher.

ACKNOWLEDGEMENTS

First, I would like to thank my supervisor, John Krebs, for his constant and enthusiastic encouragement and advice. The many Sundays spent discussing chapters of this thesis were only the last stage of the most fruitful professional association that I have enjoyed.

I would also like to thank especially David McFarland who, unwittingly, was responsible for my approach to quantitative ethology and my coming to Oxford in 1974. Most of my compañeros in the Animal Behaviour Research Group were a source of inspiration and learning by osmosis, but the osmotic gradient was specially steep with respect to some of them, who I feel deserve special mention: Alasdair Houston, Robin McCleery, Richard Sibly, Marian Dawkins, Shelly Cohen, Roger Tarpy, Mac Hunter, Richard Dawkins, Jon Erichsen, David Macdonald, Adeline Nuñez will always mean Oxford for me, as will - very especially - Pat Searle.

Alasdair Houston read and criticised the whole of this thesis, and I benefited enormously from his advice and enjoyed thoroughly all aspects of working with him. Carlos Bernstein, Robin McCleery and Richard Sibly also read some chapters and suggested very good ideas. I did my best to take them into account.

John Roberts introduced me to the world of acoustics. I wish he enjoyed as much as I did all the rainy dawns we spent with Mac Hunter in Wytham.

Meyer and Elisa, my parents, endured bravely the worst times of their lives while I was away. I am most grateful to them for too many things to list, and feel deeply sad that Meyer did not live to the joy of my return home.

My research was supported financially by the following sources: The British Council, The Science Research Council, Wolfson College,

Su and Ed Kleinman, and, last but not least, Lidia Rapaport.

Professor J. W. S. Pringle and Dr. D. J. McFarland provided facilities in the Department of Zoology and the Animal Behaviour Research Group respectively.

Nick Styles gave me efficient and enthusiastic technical support and helped me to train the birds and run some experiments.

Dick Cheney, Dave Palmer and Terry Barker from the different workshops of the Department built most of my gadgets.

With respect to the pre-history of this thesis, I am grateful to Enrique Segura for the good neurophysiological years together and his encouragement to come to Oxford, and to Su Kleinman for teaching me English.

Lidia Rapaport helped me to proof read, by making drawings, typing references, and correcting my English. Her support has been invaluable.

CONTENTS

	Page No.
CHAPTER 1	
Procedure and General Methods	1
1.1. Introduction	1
1.2. The Species	1
1.3. The Subject	2
1.4. Acclimitisation to the Laboratory	2
1.5. The holding aviaries	3
1.6. Roosting boxes	3
1.7. Diet	4
1.8. Experimental Aviaries	6
1.9. The feeding discs	7
1.9.1. The perch	7
1.9.2. The logic	8
1.9.3. The discs	10
1.10. Training birds to use the discs	11
1.11. Setting up territorial defense experiments	13
1.12. Data Recording	16
1.13. Data Analysis	16
PART I - Studies on Time Budget	
CHAPTER 2	
2.1. The Dawn Chorus	18
2.2. Hypotheses about the dawn chorus	18
2.2.1. Proximate explanations	18
2.2.1.1. The fixed threshold	19
2.2.1.2. Internal clock hypothesis	20
2.2.2. Functional explanations	23
2.2.2.1. Ecological constraints	23
2.2.2.1.1. Acoustical factors, especially wind	23
2.2.2.1.2. Non-acoustical effects of temperature	25
2.2.2.1.3. Light intensity	26
2.2.2.1.4. Predation	27
2.2.2.2. Intraspecific factors	27
2.2.2.2.1. The behaviour of receivers	27
2.2.2.2.2. The behaviour of other actors	28
2.3. The dawn chorus in the Great Tit	30
CHAPTER 3	
Environmental Factors affecting the Dawn Chorus	34
3.1. Introduction	34
3.2. Light intensity	35
3.2.1. Methods	35
3.2.1.1. Subjects	35
3.2.1.2. Apparatus	35
3.2.1.3. Prey	36
3.2.1.4. Testing	37
3.2.1.5. Units	37
3.2.1.6. Training	38

∠ cont'd

CHAPTER 3 continued

3.2.2. Results	39
3.3. Temperature	41
3.3.1. Method	41
3.3.2. Results	42
3.4. Wind	43
3.4.1. Methods	44
3.4.2. Results	45
3.5. Conclusions	46

CHAPTER 4

Anticipatory control of body weight	49
4.1. Introduction	49
4.2. Method	51
4.3. Results	52
4.4. Discussion	55

CHAPTER 5

On the motivational mechanisms underlying the dawn chorus	61
5.1. Introduction	61
5.2. Methods	62
5.2.1. Birds	62
5.2.2. Aviary arrangements	62
5.2.3. Experimental design	63
5.2.4. Deprivation	64
5.2.5. Intrusion	65
5.2.6. Recording	65
5.2.7. Measures of responses	66
5.3. Results	69
5.3.1. Description of responses	69
5.3.2. Supression of feeding by intrusion	70
5.3.3. Return to baseline feeding rate	72
5.3.4. Intensity of the song response	73
5.3.5. Return to baseline singing rate	73
5.3.6. Number of attacks	74
5.3.7. Approach	75
5.3.8. Displays	75
5.4. Discussion	75
5.4.1. Foraging opportunity	75
5.4.2. Circadian clock	78
5.4.3. Intruders pressure	79
5.5. General conclusions	80

PART II - Studies of foraging behaviour: Optimal sampling

CHAPTER 6

6.1. Optimal foraging and sampling	81
6.2. The Memory Window	82
6.3. Psychological studies of concurrent schedules	84

∠ cont'd

CHAPTER 6 continued

6.4. The Two Armed Bandit 88
6.5. Summary 90

CHAPTER 7

An experiment on optimal sampling 92

7.1. Introduction 92
7.2. Methods 92
7.2.1. Foraging Patches 92
7.2.2. Experimental Sessions 93
7.2.3. Decision criteria 95
7.2.4. Subjects 95
7.2.5. Pilot experiments 95
7.3. The Model 95
7.4. Results 102

CHAPTER 8

Rules of Thumb for the Two Armed Bandit Problem 108

8.1. Introduction 108
8.2. Methods 109
8.2.1. Simulation procedures 109
8.2.2. The Rules 110
8.2.2.1. IMMAX 111
8.2.2.2. MATCH 112
8.2.2.3. PLWIN 112
8.2.2.4. CULLEN d: Cullen Rules 113
8.2.3. The evaluation of the rules 113
8.3. Results 115
8.3.1. Comparison with real birds 115
8.3.2. Effect of average profitability 116
8.3.3. Effect of patch diversity in a rich environment 116
8.3.4. Effect of the time horizon 117
8.4. Discussion 118

CHAPTER 9

Optimal Sampling and Time Horizon 120

9.1. Introduction 120
9.2. Methods 123
9.2.1. Subjects 123
9.2.2. Instruments 123
9.2.3. Types of run 123
9.2.3.1. Treatments 123
9.2.3.1. Neutralisations 124
9.2.4. Measures of sampling 125
9.2.4.1. Decision point 125

/ cont'd

CHAPTER 9 continued

9.2.4.2. Commitment 126
9.2.5. Experimental sequence 126
9.3. Predictions 126
9.3.1. Molar Predictions 126
9.3.2. Molecular Predictions 129
9.4. Results 129
9.4.1. General Behaviour 129
9.4.2. Test of predictions 130
9.4.2.1. Global amount of sampling 130
9.4.2.2. Difference in sampling between treatments 131
9.4.2.2.1. Graphical analysis 131
9.4.2.2.2. Statistical analysis 133
9.4.2.3. Pattern of switching between patches 133
9.5. Discussion 134

PART III - The trade-off between foraging and territorial defence

CHAPTER 10

Optimal foraging and territorial defence 136

10.1. Introduction 136
10.2. Methods 139
10.2.1. The birds 139
10.2.1. The experimental set up 140
10.2.3. Foraging schedules 140
10.2.4. Experimental sessions 141
10.3. Predictions 142
10.3.1. The behaviour of control sessions 142
10.3.1. The behaviour in intrusion sessions 143
10.4. Results 143
10.4.1. Do undisturbed birds forage optimally? 143
10.4.2. The effect of intrusions 146
10.4.3. Inverse optimality 147
10.5. Discussion 149

REFERENCES

APPENDICES

1. Algol program to solve the finite horizon two armed bandit problem with equal sampling before decision.

2. Kacelnik, A. 1979
The foraging efficiency of great tits (Parus major L.) in relation to light intensity.
Animal Behaviour, 27, 237-241.

/ cont'd

3. Roberts, J., Kacelnik, A. and Hunter, M. 1979
Sound interference patterns in animal acoustic communication. Technical memorandum 66. Polytechnic of the South Bank, London.

4. Roberts, J., Kacelnik, A. and Hunter, M. 1979
Acoustic properties of the environment as determinants of birds' song patterns. Animal Behaviour(in press).

5. Kacelnik, A. and Tarpy, R. 1979
The effect of hunger and vigilance during prey handling in the great tit (Parus major). Manuscript.

6. Krebs, J.R., Kacelnik, A. and Taylor, P. 1978
Test of optimal sampling by foraging great tits. Nature, 275 27-31.

Note: Chapter 7 was done in collaboration with J. Krebs and P. Taylor and Chapter 10 with A. Houston.

CHAPTER 1

GENERAL PROCEDURES AND METHODS

1.1. Introduction

During the course of the experimental work the methods and instrumentation evolved, due in part to the different needs of the various experiments and in part to the accumulation of experience.

In order to make the path less tortuous for future researchers, most of the methods are described in some detail in this chapter. This information needs to be complemented with that under the Methods heading of the pertinent chapters.

1.2. The species

The Great Tit (*Parus major*) already has an extensive history of experimental work in captivity.

Various studies, like that of Dunnet and Hinde (1953) on the roosting behaviour of captive great tits, Blurton-Jones (1968) on aggressive displays, Smith and Dawkins (1971), Smith and Sweatman (1974), Partridge (1976) and Krebs et al. (1972) on foraging behaviour had already established the suitability of the species for ethological experimentation in the laboratory by the time this work started. In addition, more than twenty years of intensive field studies (since Hinde, 1952) provide a great deal of information on the behavioural ecology of the species. In addition to Hinde's work the field work of authors like Krebs, Gibb and more recently, Hunter had special influence on my work and will be mentioned repeatedly in this thesis.

1.3. The subjects

All the experimental subjects were wild birds in adult plumage, captured by mistnetting in or near Oxford.

Some animals were used in more than one experiment. When this is the case, it will be mentioned in the appropriate Methods section.

1.4. Acclimatisation to the laboratory

On arrival in the laboratory the birds were placed indoors in individual cells (0.33 m x 0.33 m x 0.33 m) forming part of sets of 4 called collectively "Parus motels". Figure 1.1 shows a diagram of a motel. Each cell had 5 solid walls and a wire mesh front, a door at the back and two perches. Food was given in a bowl and water in a fountain stuck to the back of the motel. The birds had free access to a roosting box hanging from hooks at the front of the cells.

The aim of this holding method was to facilitate acclimatisation of the birds to laboratory conditions, and to the standard balanced diet. Releasing newly caught birds in big aviaries is not satisfactory because they are easily frightened and tend to spend a long time flying near the roof of the aviary without settling near the food dishes. In contrast, the birds in the motels calmed down very rapidly, and normally they were taking food from the dishes in a few minutes.

The motels were kept facing a window with translucent glass, so that the birds received natural light but could not see people passing near their cages.

New birds were kept for a minimum of two weeks in the motels, and then transferred to permanent holding aviaries. During those two weeks the diet was gradually changed from mostly worms and sunflower seeds to the standard balanced diet.

1.5. The holding aviaries

Birds which were not participating in an experiment were kept in holding aviaries. Initially these aviaries were big communal aviaries where up to 8 subjects lived together, but this was later replaced by individual cages measuring approximately 1 m x 1 m x 0.80 m. The birds seem to keep in better shape when kept individually, and also it was easier to inspect them and to choose a particular animal when it was required for an experiment. All the individual holding cages had roosting boxes hanging from their front, and the birds normally used them at night and when frightened. The holding room was not temperature controlled and had a light cycle of 12 hours of light.

1.6. Roosting boxes

These were aluminium boxes measuring 0.15 m x 0.08 m x 0.08 m, with open front and back, and slots for sliding doors (Figure 1.2). They proved to be a very convenient method of catching birds for either transporting or weighing, avoiding the disturbance caused by handling the animals directly. These boxes were always available to the birds in the holding cages, and most of them used the boxes as a refuge whenever a person approached the cage. They could be trapped simply by lowering the sliding door in the front (the back door was of

course kept in place all the time). Since the boxes were simply hanging loosely they could be detached from the cage and the birds could be weighed or transported to an experimental place without seeing the experimenter.

1.7. Diet

This consisted of the standard balanced diet given ad libitum when the birds were not in an experiment, plus the prey given during the foraging tests.

The balanced diet used initially was described by Krebs et al. (1972), but it was later replaced by a proprietary mixture called Aleckwa Mischung I imported specially from Germany. This mixture was enriched with raw carrots, hard boiled eggs and vitamins (Abidec) as well as with maggots. Peanuts and sunflower seeds were added in winter only.

Three different types of prey were used, segments of mealworm, "ravioli", and fly pupae, the last of these in two different sizes.

The first type of prey used consisted of mealworm (larvae of Tenebrio sp) pieces which were highly palatable to the tits and had the advantage of not requiring any special training. Their main inconvenience is that once they are sectioned, the segments deteriorate very rapidly, and after only 2 or 3 hours in moderately warm weather they become black and the birds reject them. It would be interesting to ask a researcher with similar taste to the Great Tit (Warren and Vince, 1963) if the segments change markedly their taste when they get black.

The "ravioli" were pastry cocoons containing a segment of a mealworm. They were aimed at standardising handling time and at making a prey with a controllable degree of crypsis.

The pastry was made with flour and water, and could be coloured by mixing different proportions of food dyes with the water before preparing the dough. No fat was used because I did not want the birds to eat the pastry.

The ravioli were made by placing portions of mealworms between two thin layers of dough, and cutting small squares by pressing with a wire mesh (Figure 1.3).

While the ravioli are very convenient prey for experiments where handling time and crypsis are important, they have the inconvenience that their preparation is time consuming and that they can not be stored for long, because the texture and colour of the pastry change as it dries.

The final and most successful prey type used were fly pupae. Fly pupae have several features that make them almost ideal for foraging experiments in the laboratory. They are naturally occuring prey, fairly homogeneous in size and shape, the tits take them without training and have a fairly constant handling time.

The birds normally eat only the soft contents and discard the cocoon, so that the number of prey taken can be counted at the end of the experiment and the weight of food eaten can be calculated by weighing the empty husks and comparing with the live pupae. They are immobile live prey, so that prey that are not eaten in a test can be left in place for the next experiment and will keep in good condition.

I used pupae of two sizes, from the species Lucilia sencata and Calliphora vicina.* In both cases they were bought as larvae (maggots) from fish tackling shops and kept in a cold room (3°C) until about 3 days before an experiment, when they were placed at room temperature. After about 48 hours at room temperature the larvae pupated and could be kept alive in a refrigerator for at least two weeks without detectable changes in their state or their palatability for the tits.

The mean weights of the pupae were 26 ± S.D. 1.6 mg (N=150) for Lucilia sencata and 61 ± S.D. 3.2 mg (N=150) for Calliphora vicina.

1.8. Experimental aviaries

The foraging experiments were done in the experimental aviaries of the Animal Behaviour Research Group, which have been described in detail in several sources (recently in Cowie, 1979). Figure 1.4 presents a general plan of the aviaries. Basically they consisted of two indoor experimental areas each one connected by trapdoors to two outdoor individual flight cages. Thus, up to four birds could be kept in residence in the outer compartments, being allowed to enter the experimental area by remotely opening the appropriate trapdoor.

The dimensions of the experimental areas were approximately 4.5 m x 3.5 m with a height of 2.4 m. The arrangement of apparatus inside varied between experiments and is described under the Methods section in each chapter.

---

*I am grateful to Steven Malcolm for identifying the flies.

All the observations were made through one-way perspex windows. All experiments were carried out with the lights in the aviary switched on and the outside lights switched off.

The birds residing in the individual outdoor compartments were normally trained to perform different foraging tasks in the experimental area, and entered immediately after the trapdoor was opened at the beginning of a test. The trapdoors were kept closed during tests and opened again at the end.

1.9. The feeding discs

This apparatus was developed on the basis of an idea suggested by Mike Cullen and John Krebs, and was built in the workshop of the Department of Zoology.

The discs are operant food dispensers analogous in their principles to the Skinner Boxes used in Experimental Psychology. The full apparatus has three components, a perch, a logic unit and a feeding disc (Figure 1.5).

1.9.1. The Perch

This is the portion of the instrument which the experimental subject can operate. Figure 1.5a shows a picture of the perch together with the feeding disc, described below. The perch itself is a cylindrical piece of wood (diameter 0.01 m, length 0.08 m) sticking out of a block of aluminium at 0.08 m above the surface where the whole apparatus is placed.

Each perch was calibrated so that a weight of 15 g halfway between the tip and the base would operate a microswitch fitted inside the block. The closure of the microswitch

was the input signal to the logic that controls the process of delivering rewards.

Since a Great Tit weighs between 16 g and 21 g the birds could operate the perch by simply hopping onto it. Because this is what the birds normally did, the term 'hop' is used in the experiments involving this apparatus as equivalent to the term 'response'.

1.9.2. The Logic

Signals from the perch were sent to a programmable unit consisting of a set of modules made by Behavioural Research Developments Limited (Figure 1.5d). This logic was programmed so as to standardize the length and amplitude of the pulses coming from the perch, and count them individually as units of response. The logic delivered rewards in the feeding disc according to a chosen schedule. In most experiments the schedules used were Variable Ratios (VRn) (Mackintosh, 1974). Under a VRn schedule a reward is delivered on average every n responses. The actual number of responses required in each occasion varies pseudo-randomly within the range 1 to 2n. Thus, under a VR5 the actual number of responses required to obtain a reward ranges between 1 and 10, with a mean of 5. The inverse of the number that characterises the VR schedule (1/n) is called the 'reward probability' associated with that particular schedule. VR schedules resemble field situations where the number of prey obtained is roughly proportional to the amount of searching effort and where no depletion occurs within the period of observation. This type of schedule was chosen

because it may be considered similar to the situation faced by a Great Tit in Winter when they forage mainly for immobile prey like seeds or pupae (Gibb, 1954; Hartley, 1953; Perrins, 1979, Lack, 1971, Betts, 1955).
Under these conditions prey depletion is probably a slow process taking place through the whole season and the number of prey captured is roughly proportional to the effort put by the searching bird. This point about the proportionality of number of captures and hunting effort is emphasised here partly in contrast to the claim made by Staddon (1979) and others about the interest of studies based on variable ratio schedules in relation to Variable Interval schedules (Mackintosh, 1974). In the latter, the number of captures does not depend linearly on the intensity of foraging effort. Variable interval schedules with limited hold are similar to processes where the predator waits for prey to emerge naturally and be available for a short period, as during the eclosion of aquatic insects. The classical variable interval schedule with umlimited hold resembles field situations with behavioural depression of the prey like when wagtails hunt for flies in cowdung, and the flies hide after the first captures and re-emerge when the predator has gone (Davies, 1977).
It is clear that any reward schedule used in the laboratory will face great limitations for the comparison with real life situations (Collier, 1980), but my point is that when assessing the appropriateness of using one type of schedule

or another it is more important to take into account the hunting behaviour under the natural conditions of the species that is under study than the researcher's inclination for systems with many 'degrees of freedom'. Staddon 1979 and Myerson and Miezin, 1980, hold conflicting views about this question.

1.9.3. The discs

Whenever the logic unit delivered the signal for a reward, one food item was made available to the bird in a feeding disc (Figure 1.5b). These were flat metal boxes (0.40 m x 0.40 m , height 0.11 m) covered by a wooden lid. The lid had a window through which the subject could pick up one food item at a time from an individual food cup. Underneath the lid there was a horizontal rotating acrylic disc with the food cups (holes 0.7 cm deep and 0.5 cm in diameter) arranged in circles. Only one food cup was accessible through the window, and a new one appeared every time the disc rotated 5 degrees. The first version of the disc had a single line of holes along its perimeter, and in consequence could be used to deliver up to 72 food items before completing a full circle and being refilled ($72 \times 5^{o} = 360^{o}$). In a later version the disc had 5 concentric lines of holes, or food cups. An additional piece of aluminium with five holes at various radii made available only one of the concentric rows at a time. This aluminium piece was advanced one step every full circle of the acrylic disc, so that the subject had successive access to each of the rows of food cups. This development enlarged

the capacity of the feeding disc from 72 to 360 food items (5 rows of 72 holes each) and made it possible to do much longer tests. With reasonably large prey (say as big as the pupae of a Calliphora sp fly) the disc may be loaded with all the food required for a full day of foraging.

The rotating disc was advanced in steps by means of a solenoid-driven toothed flywheel placed inside the box. The system of the flywheel and a solenoid were obtained from the dialling pieces of old telephones, and the number of holes in each row was determined by the fact that the telephone pieces rotated 5 degrees for each tooth of the flywheel.

The discovery and adaptation of the telephone pieces is due to Mr. Dick Cheney, in the workshop of the Department of Zoology.

1.10. Training birds to use the discs

The training technique improved with experience, but for brevity I shall describe the final and most successful procedure. Previous training procedures achieved the same final performance but were much slower and more laborious.

The training followed several steps:

(a) The subject was given ad libitum food through the window in the wooden lid for a day. This was simply done by placing a food dish in place of the acrylic disc. The bird soon got used to approaching the window looking for food.

(b) The logic was programmed so as to deliver automatically one reward each 90 seconds. The acrylic discs were replaced,

loaded with one pupa in each hole. After a short time the birds learned to associate the noise produced when the disc was advanced with the presence of food through the window. They usually flew around the room until the disc moved, and then approached the window to pick up the pupa. It took normally between one and two hours for the birds to learn this response, but all subjects were left for a full morning with automatic delivery to consolidate it.

(c) The automatic schedule was left running, but the perch was now connected to the logic so as to deliver one additional reward each time that the perch was operated. Several pupae were stuck under a piece of sellotape to the perch block, in a position that the animal could only reach by hopping onto the perch. What normally happened was that while waiting for the next reward to come automatically, the subject would hop onto the perch trying to reach the pupae, and as a consequence the disc would advance one step. The noise of the advancing disc was by now a powerful second order conditioned stimulus (see Tarpy and Mayer, 1978), and the bird responded to it by leaving the perch at once and approached the window to pick up the reward.
Within a few hours the birds learned to operate the perch without any pupa stuck to the perch block.

(d) The automatic schedule was stopped and the logic programmed with a Variable Ratio schedule, normally starting with a VR2 and increasing the ratio progressively every hour towards a VR10. Once the bird was working consistently with a VR10 schedule it was considered to be disc-trained

and the specific training for the particular experiment started. If nothing went wrong a bird could be fully disc trained by this method within two days. Things could (and often did) go wrong in one main way.

If the perch was connected too early, before the noise of the rotating disc was a positive second order reinforcer, the subject reacted with what appeared to be strong signs of fear in relation to the perch. I had birds that after a single hop to the perch developed such aversion to it that they would not touch it again for weeks, even if the perch was loaded with pupae or mealworms and they were obviously hungry. A loud noise resulting from their own behaviour seems to be interpreted as a sign of danger by great tits, or in more descriptive language, it seems to constitute a very powerful negative unconditioned stimulus, with extremely slow habituation.

Special training procedures required for the different experiments are described in the respective Methods sections.

1.11. Setting up territorial defense experiments

When an experiment involved some form of territorial defense the main problem was how to implement the intrusions of supposed 'invaders'. As for the experimental birds, males simply became territorial after a period of several weeks of living individually in the aviaries. Since normally I wanted more than one bird to be territorial in the same area (the experimental indoor aviary)Igave access to this compartment on alternate days to each of the two resident birds in the flight cages leading

to that indoor aviary. This technique limited the number of experimental subjects to 4, the number of flight cages with access to the experimental aviaries available.

With respect to the presentation of the 'intruder', two different techniques were employed.

The two techniques were based on the presentation of a male great tit behind a transparent partition while a loop tape with great tit song was played on a tape recorder. A set of stereo speakers was placed at both sides of the intruder's box. I used a live animal as an intruder, because pilot trials failed to obtain any consistent response using mirrors. The first method of presentation was by means of a 'puppet theatre'. This was a box (0.30 m x 0.66 m x 1.00 m) with a transparent perspex front (Figure 16) and a 'curtain' operated by a D.C. motor. When an 'intrusion' was required, a stimulus bird was placed inside the box before the test and the curtain was lifted at the appropriate moment by turning on the motor with a switch placed outside the aviary. The same switch operated the tape recorder with the loop tape of a song. The intruderbird was transported to and from the theatre by means of a roosting box attached to the back of the theatre. The curtain took approximately 7 seconds to complete the trip up or down, and because of this the minimum time used for intrusions was of 30 seconds, to ensure that during a minimum of 15 seconds the intruder bird was visible to the resident.

This method was very efficient in producing strong responses from the territorial holder, but it was pointed out to me that the 'intruder' might respond in a very submissive manner due to

its unusual situation, and there could be a mismatch between the visual displays of a submissive bird and the clear and loud song generated by the speakers. I did not observe any sign of this in my experiments, but decided anyway to incorporate the suggestion in later experiments by using a territory holder as an intruder. This was done using wooden boxes (0.66 m x 0.66 m x 0.66 m) with sliding doors in the front and back. These boxes were attached to the place of the original trapdoors.
They stayed in place more or less permanently and the birds got used to passing through them to get in and out of the indoor aviary. The sliding door leading outside was a wooden one.
One of the two sliding doors leading to the experimental aviary was solid and the other transparent, made of acrylic. The two resident birds were used alternately as intruder or holder in different tests. When a bird acted as territory holder, the three doors were open, and the box was just a tunnel through which the bird gained access to the indoor aviary. The bird acting as an intruder was trapped in the box by closing the back door while the bird was inside the box. The two front doors were kept closed until the time of the intrusion.

Intrusions were simulated by opening the solid door in the front of the box with the trapped bird. The transparent door remained in place and the two birds could see each other in the same way as in the puppet theatre, but now the intruder was watching its 'own' territory being occupied by another male and was expected to display back to the territory holder. The two outdoor flight cages leading to each experimental sector were

visually isolated to avoid habituation of each bird to the presence of the other. Individual identification might still have occurred by song. Since the two methods of intrusions were used in different conditions I have no means of comparing the responses between them, but both techniques were satisfactory for their purposes and all that can be said is that they resemble different natural situations, with intrusions by strangers (theatre) or neighbours (intruder boxes).

1.12. Data Recording

Data were normally recorded on magnetic tapes using a computer-compatible event recorder (Dawkins, 1971). The magnetic tapes were later played back into a small digital computer, which output paper tape with the times of occurrence of each type of event.

In one of the experiments (Chapter 7) I used an event recorder designed by P. Chatwell (unpublished) which differed from the Dawkins' event recorder in that simultaneous activities could be recorded in the same channel and that the output was sent to a magnetic storing device instead of paper tape. In later experiments, though, I returned to using the Dawkins' event recorder because of its reliability.

Operation of the feeding discs was recorded automatically and observational data (displays, songs, etc.) by watching the birds though the observer's window in the aviaries.

1.13. Data Analysis

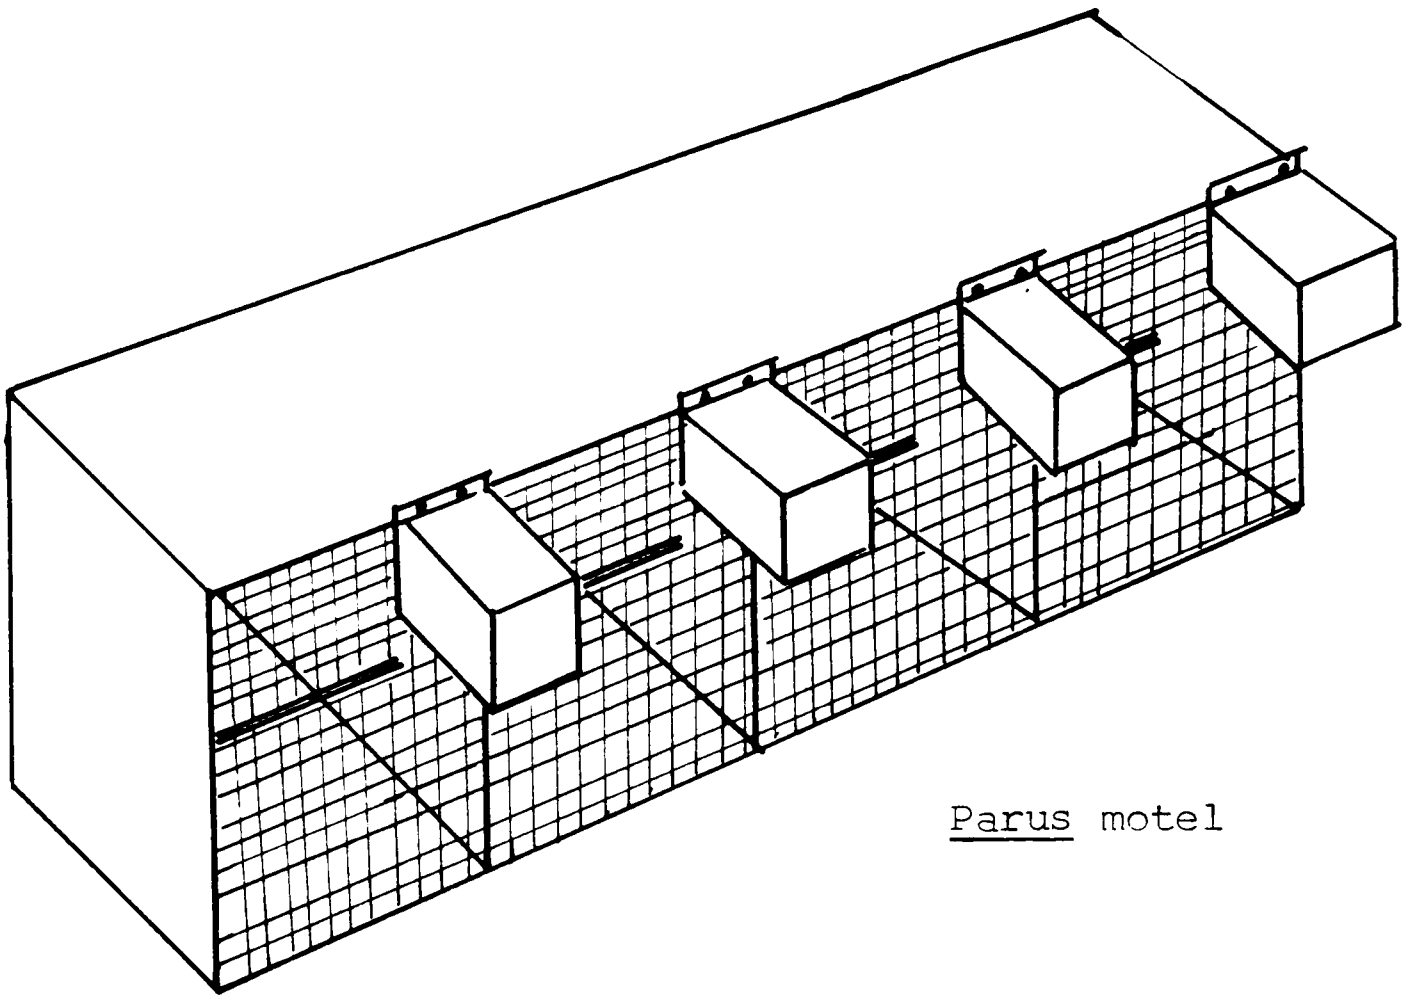
Most of the data analysis was done by specially written programs in Algol-60 language, in a dialect for PDP 8 computers developed by R. Abbot.

Some special statistical programs were used including a multi-factor analysis of variance program and a program for analysis of variance of latin-square designs written by C. Dawkins and a curve fitting program by R. McCleery.

Some graphs were made in a Hewlett-Packard x-y plotter attached to a PDP 8 computer, controlled by specially written programs, using machine code software packages developed by R. McCleery.

Simulations requiring vast storage capabilities, like those in Part 2, were run in the 1906 I.C.L. computer of the Oxford University Computer Service.

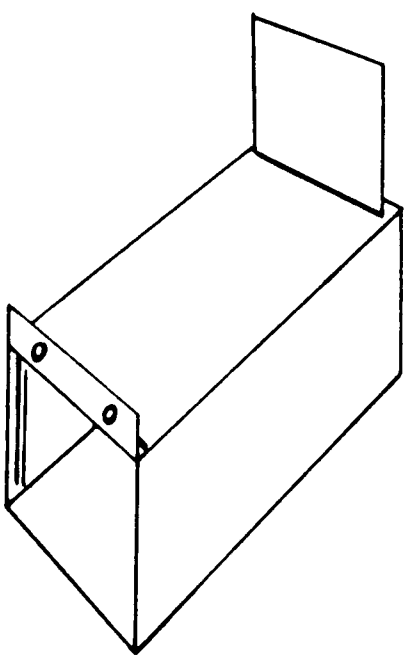
Figure 1.1



Parus motel

On arrival to the laboratory, new tits were kept for two weeks in a motel.

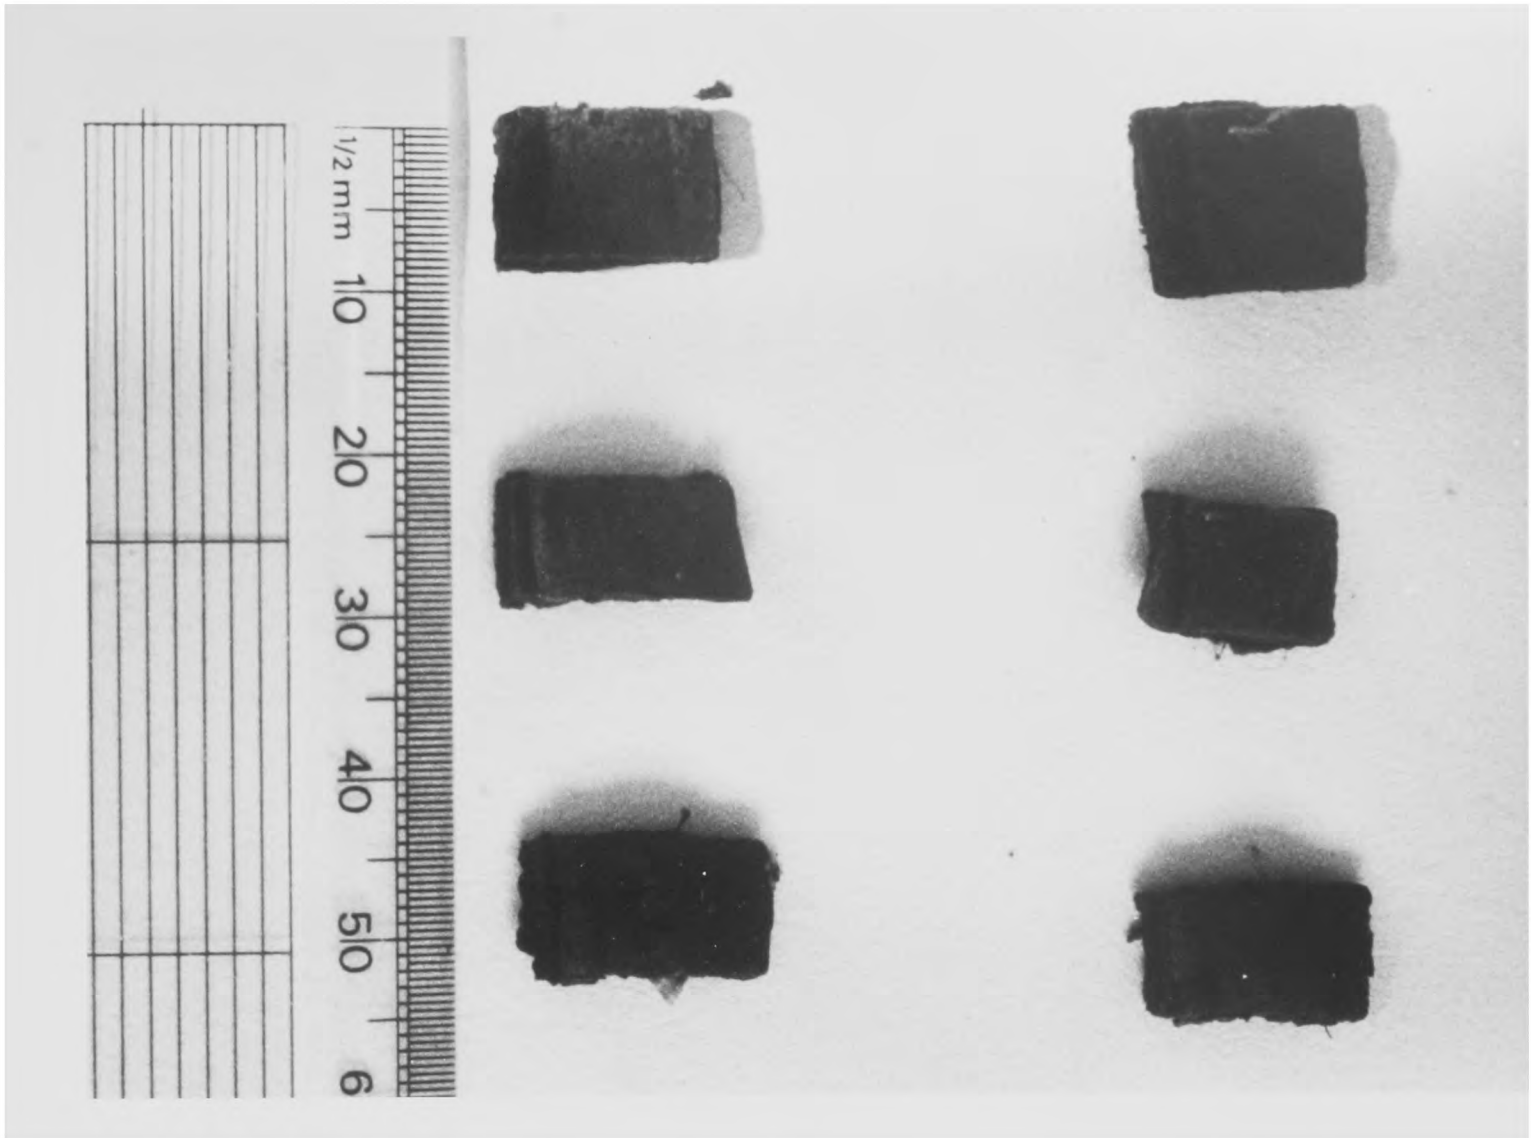
Figure 1.2



'K-box'

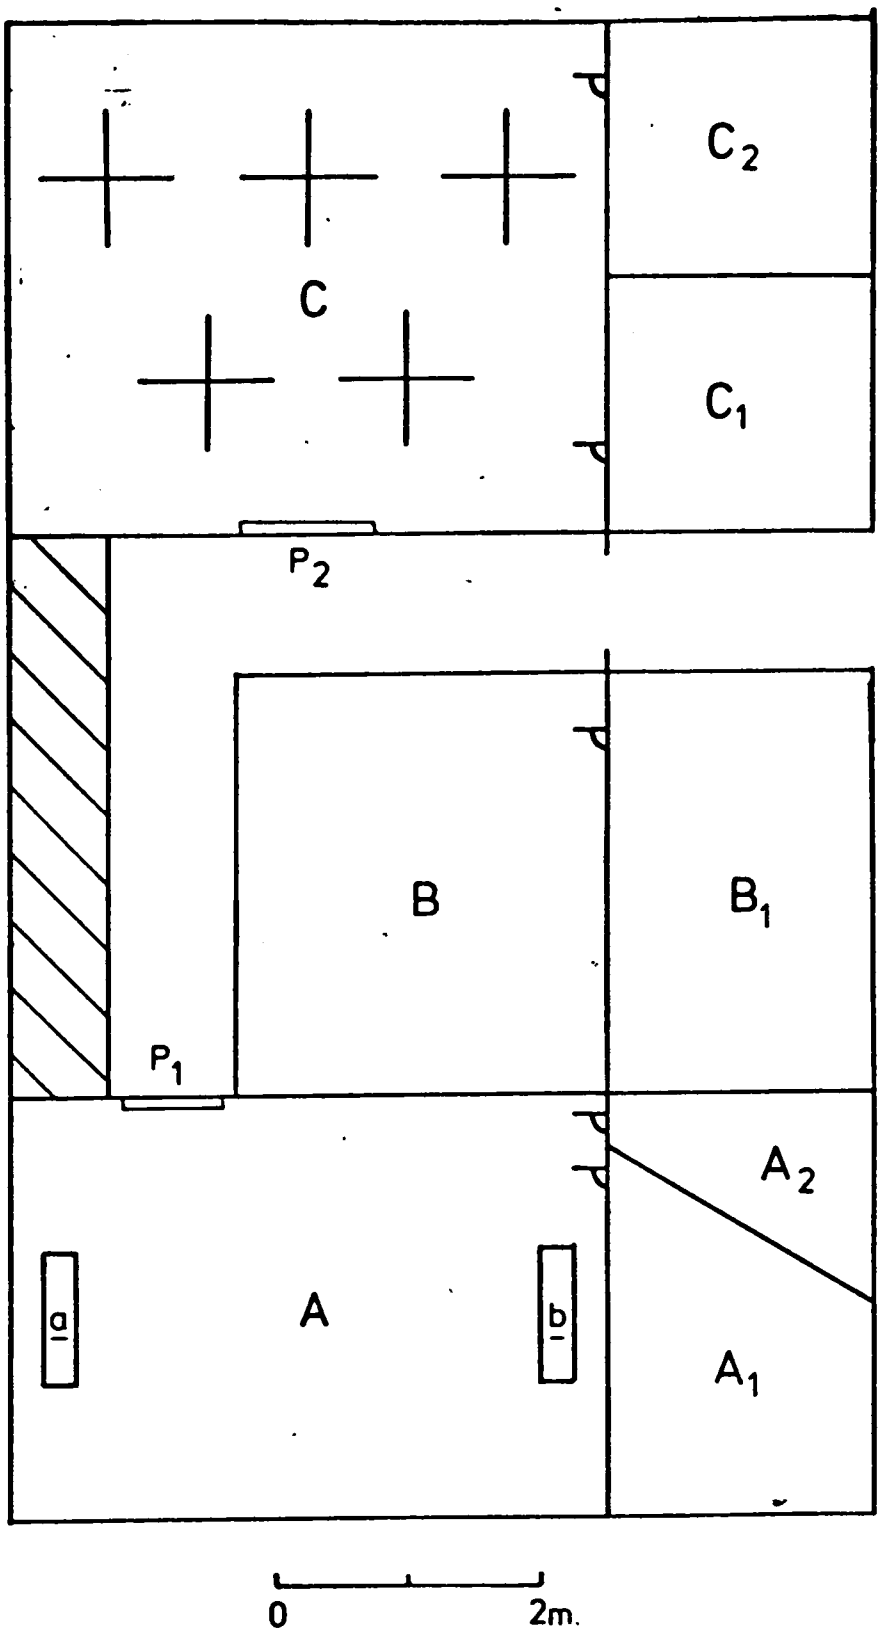
The roosting boxes were used for transporting and weighing the birds without handling them directly.

Figure 1.3



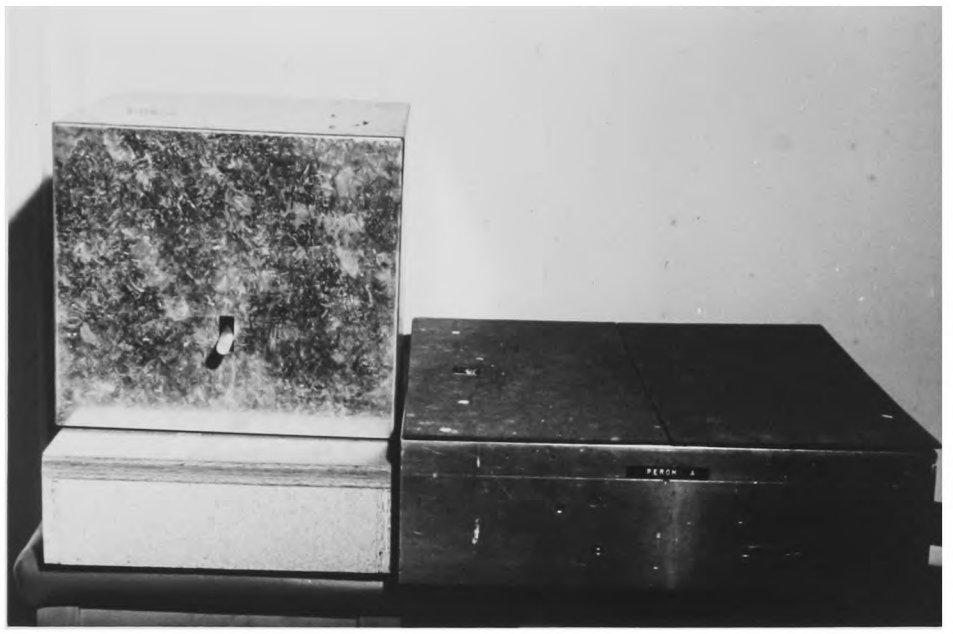
The photograph shows the general appearance and dimensions of the artificial prey used in the experiment described in Chapter 3.

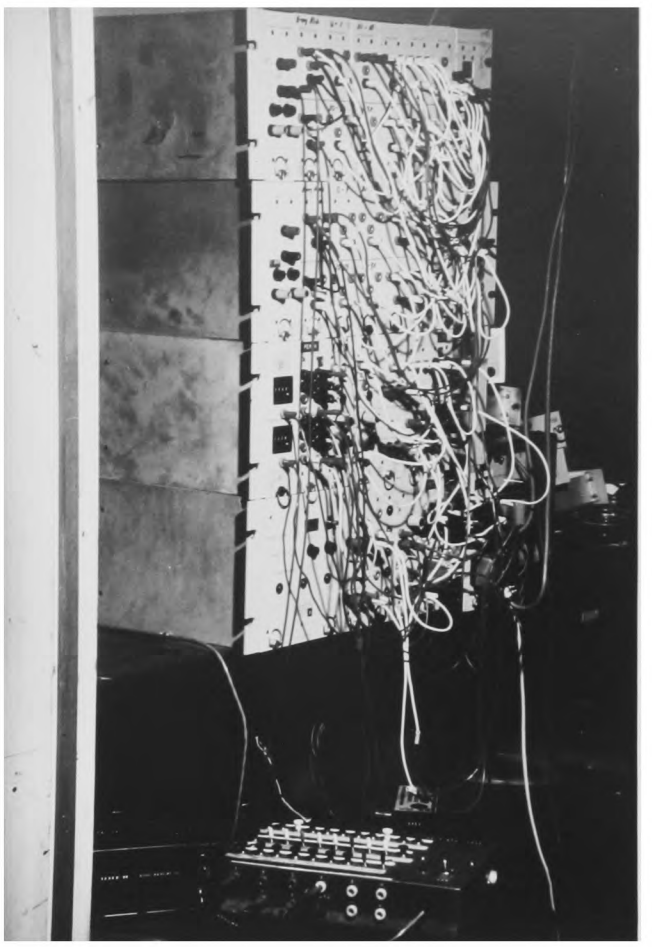
Figure 1.4

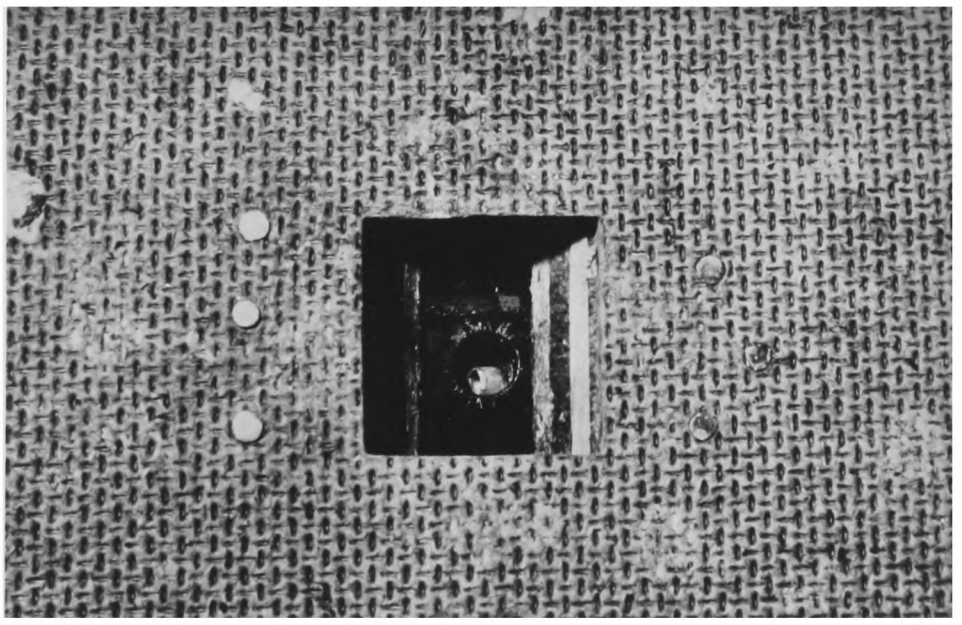


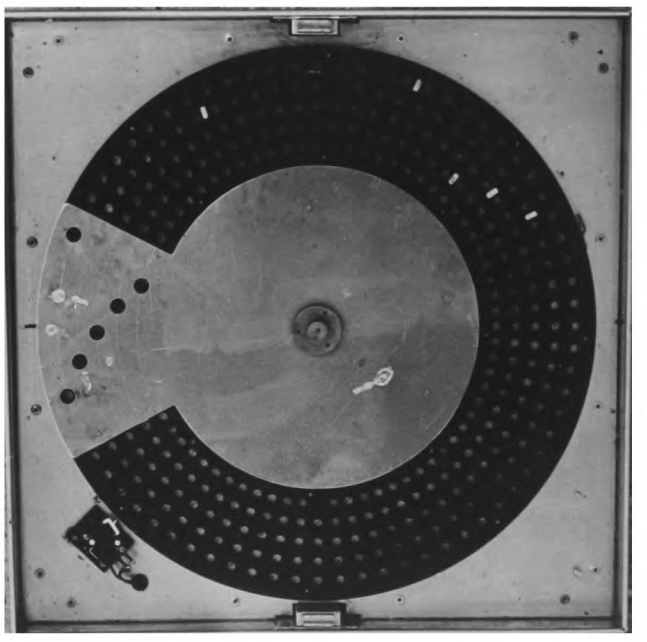
A plan of the aviaries. A and C are the main experimental sections and each are connected by trap doors to two outside flight cages (A_1, A_2; C_1, C_2).
The crosses in Aviary C and the rectangles in Aviary A represent the 'experimental trees' and the 'disc apparatus' respectively. During an experiment the birds were observed through one way mirrors (P_1 and P_2). B and B_1 were used as a holding aviary.
Reproduced from Cowie (1979).

Figure 1.5



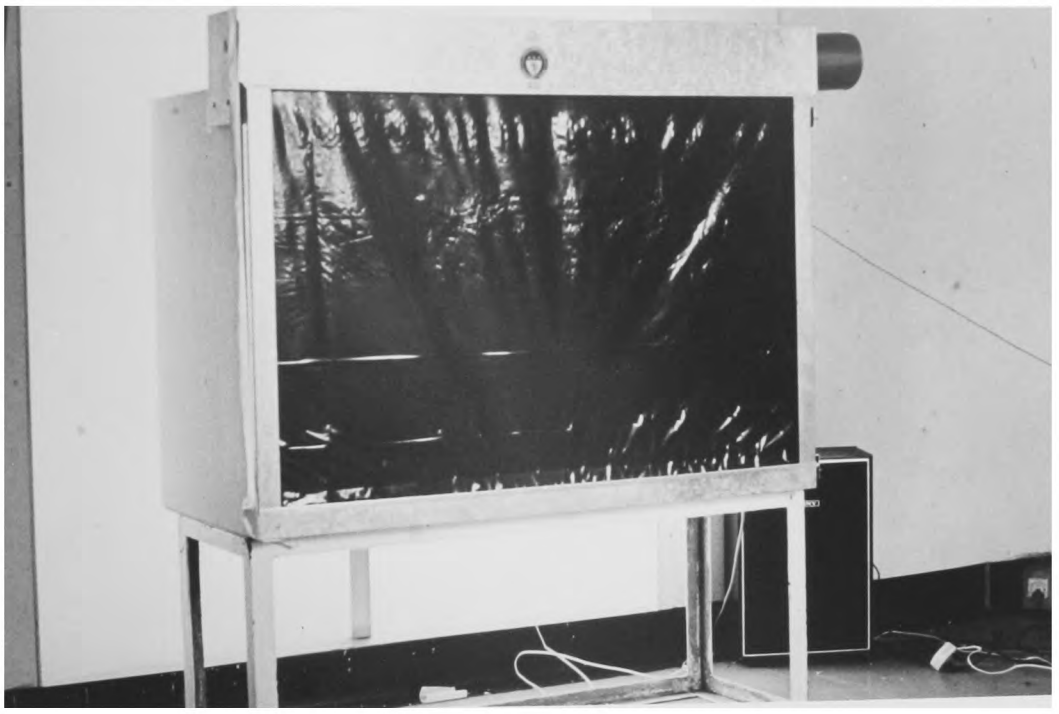


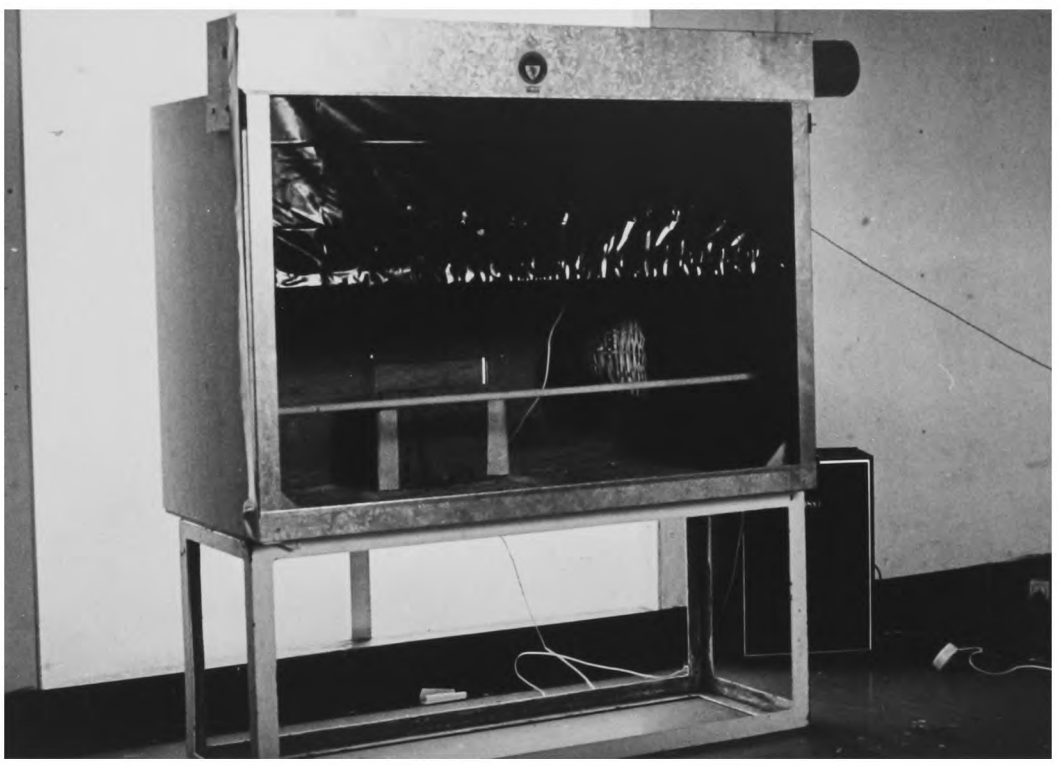




The operant feeding apparatus. Top left: perch plus disc. Bottom left: a bird's view of the window in the feeding disc. Inside, a segment of a mealworm. Bottom right: the inside of a feeding disc, Mark II, showing 360 food cups and the aluminium piece to make only one row available at a time. Top right: the logic.

Figure 1.6 A "puppet theatre", used to simulate intrusions. Top: the theatre closed. Middle: beginning of an intrusion. Bottom: the theatre is open and the 'intruder' bird is fully visible.





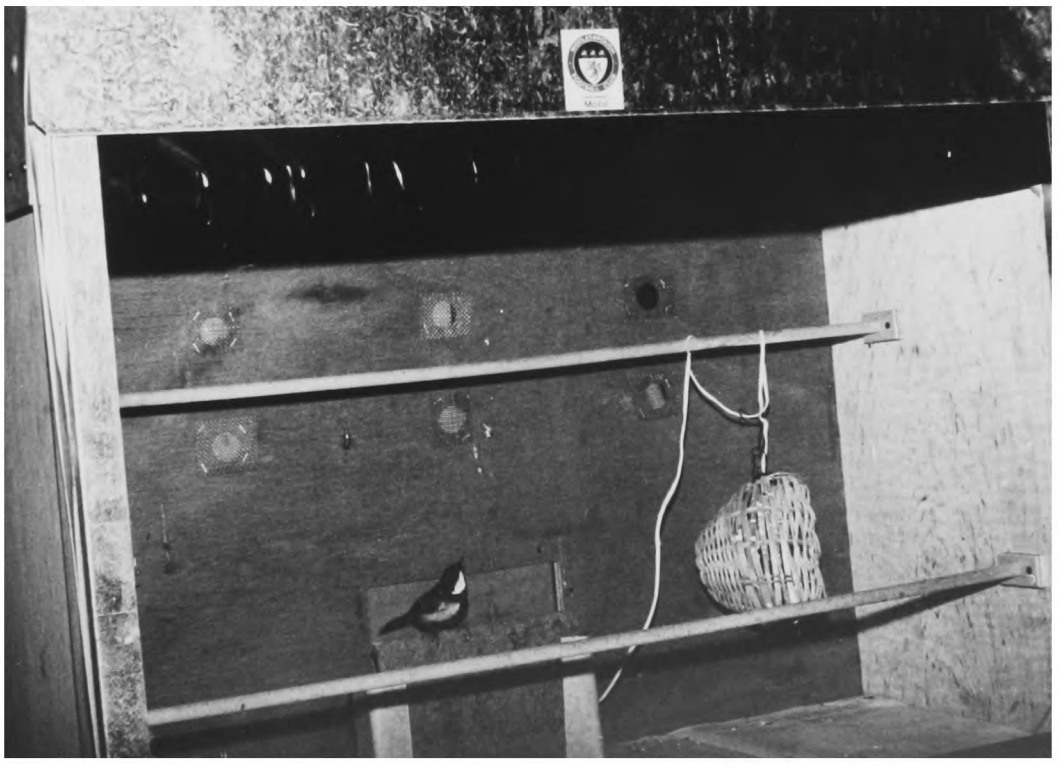


Plate 1 A great tit foraging for 'ravioli' in an artificial canopy (Chapter 3). Top left: the bird inspects a possible hiding place. Top right: having found a prey, it takes it to a suitable handling perch. Bottom left: the bird extracts the content of the raviole (a mealworm segment) while holds the rest with its feet. Bottom right: it drops the empty cocoon and swallows the mealworm segment.









PART I

Studies on Time Budget

CHAPTER 2

2.1. The Dawn Chorus

Many terrestrial animals that communicate by sound show a peak in vocal activity early in the morning. This phenomenon has been termed the dawn chorus by many authors (e.g., Wynne-Edwards, 1962; Morton, 1975; Tenaza, 1976; Henwood and Fabrick, 1979). The dawn chorus is not a geographically restricted phenomenon, it occurs in widely different habitats such as Sonoran Desert in Arizona, tropical rain forest in Borneo, Arctic Tundra, and mixed woodlands in Great Britain (Henwood and Fabrick, 1979; Astrom, 1976; Wynne-Edwards, 1962).

In spite of its widespread occurrence, there is no concrete evidence about the ultimate (functional) significance or proximate causes of the dawn chorus. The purpose of the first part of this thesis is to contribute to the understanding of both functional and causal aspects of the dawn chorus.

2.2. Hypotheses about the dawn chorus

Different authors have advanced tentative explanations of both the causal and functional aspects of the dawn chorus. It is not always easy to classify the hypotheses in these two categories since some of them involve both functional and proximate factors, but the summary below is an attempt to set up a framework for the discussion of both types of problem. My discussion will focus primarily on the dawn chorus of song-birds.

2.2.1. Proximate explanations

Most of the literature regarding the proximal causes of

daily routines refers to activity in general and to the mechanism underlying the onset and termination of daily activity, rather than to the sequence of different activities once that the animals are awake. I attempt to summarise in this section those ideas within the vast literature on daily rhythms that are relevant to the discussion of the dawn chorus.

2.2.1.1. The fixed threshold hypothesis

In the simplest hypothesis about the mechanism of the start of daily activity, including bird song, each species was said to have a characteristic fixed value of light intensity that would operate as a threshold for the onset or termination of general activity. This threshold was called 'Arthelligkeit' by Schwan (1920 and 1921-2, quoted in Astrom, 1976) and given related names by later German authors. In all cases the implication was that bird species start singing using a mechanism more similar to a photometer than to a clock. The obvious prediction from this hypothesis is a strong correlation between the time of the onset of activity and illumination level. This was found to be true for some species (Blackbirds Turdus merula (Haecker, 1916; Zimmer, 1919; Doring, 1920, all quoted in Astrom, 1976)) untrue in others (Reed bunting Emberiza scheonicus (Astrom, 1976), Wren Troglodytes musculus (Lutz, 1931)), and inconclusive for a few other cases (Chaffinch Fringilla coelebs, Greenfinch Carduelis chloris and Great tit in Scheer (1952) and Schmitz and Middel (1966), quoted in Astrom, 1976). The main objections to this hypothesis are the great variations that most species show between the illumination at which they become active at different times of

the year and the difference between awakening and roosting light intensity (Kluijver, 1950; Palmgren, 1949; Hinde, 1952; Astrom, 1976). Also, species that roost in caves and holes in trees still wake up at regular times, in spite of very little variations in illumination in their roosting places. Subbaraj (1977) for instance, found that the emergence time of the bat Taphozus malanopogon, which roosts in caves, varied less than 16 minutes in the whole year, while sunset time varied 41 minutes. The illumination at the emergence time varied from 0.1 lux to 750 lux, or almost four orders of magnitude.

In summary, the fixed threshold hypothesis predicts that there should be a strong correlation between light intensity and onset of activity, that the illumination at awakening time ought to be similar to illumination at roosting time, and that these values ought to be constant for each species throughout the year. Only the first prediction is supported for some species, while it is definitely ruled out for others. The second and third predictions do not seem to hold for any studied species. More complicate versions of this hypothesis, with adjustable thresholds, cannot be totally rejected, but are unlikely to be a common mechanism.

2.2.1.2. Internal clock hypothesis

Most of the failures of the fixed thresholds hypothesis can be overcome by postulating the presence of some internal timing process synchronised with the environment by a suitable cue or Zeitgeber (see Aschoff, 1960; Daan and Aschoff, 1975; Daan, 1976, for reviews). Extensive literature on laboratory experiments

show conclusively that animals are able to measure time quite accurately, and that the phase of the daily rhythms of activity can be adjusted by environmental cues, especially light cycles.

The involvement of clock mechanisms in the timing of the onset of activity in free living animals is also beyond doubt (Astrom, 1976; Subbaraj, 1977; Hinde, 1952; and many others). But the role of circadian clocks in the control of the daily routine of wild animals as well as the nature of the Zeitgeber mechanism operating is less well documented.

A general problem in interpreting the relevance of laboratory information of circadian pacemakers and their entrainment mechanisms to field conditions is that rhythms in the laboratory are usually studied in artificially undisturbed circumstances. While in those situations internal factors are bound to play a crucial role in the timing of activity, in real life animals are in the presence of a whole constellation of factors which may override the internal rhythms. Daan and Aschoff (1975) are aware of this problem, which under the name of 'Masking' (Aschoff, 1960) constitutes a real nuisance for the study of biological oscillators in the laboratory. The ability of certain animals to 'choose' a Zeitgeber has been demonstrated in squirrel monkeys Saimiri sciureus by Sulzman et al. (1977). These animals are 'normally' entrained by the light-dark cycle, but when food is made available for a limited, regularly occurring interval, the presence of food becomes the Zeitgeber for the circadian rhythms of drinking, body temperature, urinary potassium and water excretion, indicating the flexibility of the timing mechanism.

It has recently become apparent that environmental factors often have overriding importance in determining the temporal sequence of behaviours. In particular this can be illustrated with reference to current developments in the study of the control of food and water intake. A large family of models of the control of drinking and feeding patterns was for years developed on the basis of the behaviour of pigeons and rats with water and food ad-libitum (McFarland and McFarland, 1968; Toates and Oatley, 1970; Toates and Booth, 1974). These models predicted accurately the meal patterns and drinking bout structure as the outcome of the rhythmic behaviour of negative feedback systems with thresholds or delays. The driving variable was, in all cases, some measure of internal deficit, and the animals were supposed to drink or eat whenever those deficits were above some critical value. These homeostatic models of intake face difficulties in accounting for the behaviour of free ranging animals. In contrast to laboratory experiments, in the wild there is rarely a situation where resources are freely available and no time has to be invested in doing other activities. Awareness of these complications has led to the proposal of mechanisms that emphasise the role of environmental supply or 'opportunity' in the timing of intake (Wirtshafter and Davis, 1977; Collier, 1980; Lea, 1980, 1979). When more naturalistic experiments were designed, simple deficit-driven models proved to be unable to predict the experimental results (Collier, 1980; Sherry, Mrosovsky and Hogan, 1979, Mrosovsky, 1979).

My approach to the problem of the mechanisms involved in the dawn chorus will emphasise the effect of environmental opportunity. In Chapter 4 I study some features of the mechanism controlling food intake in relation to the daily pattern of food availability, and in Chapter 5 I explore the effect that foraging opportunity has on readiness to sing and defend the territory.

2.2.2. Functional explanations

Two main types of argument can be put forward to explain the dawn chorus in functional terms. First, physical constraints in the environment, such as light intensity, wind speed, temperature and humidity change through the day, and may influence the opportunity to perform different activities. Secondly, there is the possibility that the behaviour of the conspecifics at whom the vocalisations are directed may alter through the day, so as to make it particularly advantageous to vocalise at dawn. It is likely that the second type of effect is a consequence of the first, but I will discuss the two types of argument separately.

2.2.2.1. Ecological constraints

2.2.2.1.1. Acoustical factors, especially wind

Acoustic attenuation (i.e. the difference between the power of a sound signal at different distances away from a source) depends on certain properties of the atmosphere such as wind, relative humidity, temperature, and the presence of fog, which vary regularly with time of the day (Kinsler and Frey, 1962; Delany, 1977). If the benefit derived by an animal from producing an acoustical signal such as territorial song depends on the distance over which the sound is detectable, then there would be

an advantage in vocalising at those times of the day when acoustic attenuation and background noise are at a minimum. Henwood and Fabrick (1979) proposed a model that incorporates the effects of wind, temperature, humidity and background noise, to predict audibility of an acoustic signal at different times of the day. Wind, temperature and humidity are responsible for 'excess attenuation ' - the attenuation of sound in excess of that predicted from the spherical expansion of the sound wave, 6 dB per doubling of distance (Wiley and Richards, 1978) . They guessed the values of background noise and considered that the signal would be audible at any point in space where its power was above the level of background noise. By using average climatological data, their model predicts that an acoustic signal produced at dawn in places as different as the rain forest in Borneo or the Sonora desert in Arizona would be audible over an area 20 times bigger than the same signal produced at noon and twice as big as if the signal was produced at dusk (Figure 2.1).

Henwood and Fabrick conclude that this obvious advantage of dawn vocalisations explains what is being optimised by the dawn chorus. Actually, their model has not yet been thoroughly tested, but assuming that it is correct there may still be other important ecological constraints which are necessary to explain the dawn chorus in all animals and species. Definitive testing of the model requires the comparisons of its predictions with experimental recording of sound attenuation and background noise in the field, when the climatological variables can be measured simultaneously, as opposed to relying on average figures. It woul

be interesting to see if the model can explain the difference in timing of the songs of different species on the basis of the main frequencies of their song and the way in which these frequencies are affected by the predictions. I will briefly describe some field measurements on the effect of wind (the most important cause of excess attenuation) in Chapter 3.

2.2.2.1.2. Non-acoustical effects of temperature

In species where other time consuming activities (especially foraging) cannot be performed at the same time as vocalising, the overall benefit derived from concentrating vocal communication in the early morning will depend not just on the opportunity for acoustic communication, but also on the advantages of performing other activities at that time. For example, if an animal's feeding rate depends on the activity of its prey, and this in turn is strongly dependent on temperature, it may be inefficient to feed at dawn when the temperature is low. This is true for insectivorous birds that hunt flying insects (Davies, 1977). It may therefore be advantageous to concentrate on vocalising early in the morning because the opportunity for foraging is very poor. I will discuss below how this argument may relate to the Great Tit, and in Chapter 3 the problem is discussed further.

An additional effect of temperature may be to modify microhabitat conditions differentially. On a cold morning an animal could gain considerable thermal benefit from sunbathing (Lustick, 1969), so that if feeding is normally done in shaded places

(e.g. a forest undergrowth) and singing in exposed places (such as tree tops) the relative advantage of dawn song may be increased.

2.2.2.1.3. Light intensity

This factor may be involved in the dawn chorus in several different ways. First, the same rationale developed for temperature can be applied to light. In a species where foraging is normally incompatible with singing, any effect of light on foraging opportunity would indirectly alter the overall benefit which would accure from singing at particular times. For example, low light intensity in the early morning may decrease the success of a bird foraging for cryptic prey (see Chapter 3) and therefore increase the relative advantage of singing at dawn. Light may also favour early morning vocalisations by reducing the efficiency of visual communication. It has been shown (Smith, 1979) that temporarily muted redwing blackbirds (Agelaius phoenicus) replaced vocal signals with other forms of territorial defence, including visual displays. If this substitution can operate in both ways, song would be a favourable form of communication when light limits visibility for a particular species.

In nocturnal migrants such as the prairie warbler (Dendroica discolor) and sedge warbler (Acrocephalus scirpaeus) new settlers arriving during the night compete to establish a territory or to attract females in the first hours of the following morning (Catchpole, 1973; Nolan, 1978) and since establishing a territory and attracting females commonly involves vocal interactions, the dawn chorus may be a secondary result of the pattern of migration imposed by light.

2.2.2.1.4. Predation

Singing may be an activity that makes the individual particularly exposed to predation and this may lead animals to sing at times when they are less likely to be detected by predators, just as great tits choose to sing from places sheltered from the view of predators (Hunter, 1978). On the other hand it could be argued that the twilight periods are the most dangerous times of the day because both nocturnal and diurnal predators may be active (Hinde, 1952), so that animals should tend to avoid activities that expose them to predation during this time. These hypotheses could be tested by recording the intensity of predation at different times of the day and during different activities.

2.2.2.2. Intraspecific factors

In this section I refer to factors that result from the daily pattern of behaviour of conspecifics, both the receivers of the vocalisations and other vocalising animals (actors).

2.2.2.2.1. The behaviour of receivers

When the purpose of the vocalisation is territorial advertisement or mate attraction, the vocalising animal would benefit from calling when there are more intrusions from floating individuals or when more females are looking for males, rather than just vocalising at the times favoured by the ecological constraints.

However, receivers may also be influenced by the ecological constraints discussed above. Thus, for example, non-territorial males may tend to seek territories early in the morning when foraging is unprofitable, or when they cannot easily be detected

because it is too dark. Territorial males may then sing most intensely early in the morning just because that is when intruders pressure is strongest.

The tendency of intruders to be active early in the morning may also relate to the mortality of territory owners. Assuming that instantaneous risk of death of each bird is constant throughout the day, there would be a greater number of empty territories 'on the market' at the onset of daily activity than at other times. Floating males would benefit from seeking territories at that time instead of during the middle of the day, when territories are vacated at scattered times. Krebs and Perrins (1977) described an experiment in which the rate of arrival of invaders to empty territories of great tits was measured as a function of time of the day. Their results are shown in Figure 2.2. The rate of arrivals was clearly higher in the first hours of the morning and dropped later on. These results may be a product of either of the two mechanisms proposed here. The Great Tit in Britain is non-migratory (Perrins, 1979), so that the higher rate of instrusions early in the morning cannot be the result of arrival of nocturnal migrants as discussed earlier (Section 2.2.1.3.).

2.2.2.2.2. <u>The behaviour of other actors</u>

The timing of the territorial song may be affected by the need to sing when everybody else is singing. It has been shown that the rate of invasions to empty great tit territories is reduced by broadcasting great tit song (Krebs, 1977). A corollary of this result is that if a male remains silent when other males

are singing it would have to defend its territory against a disproportionately high number of invasions. Thus, once that a tendency to vocalise early in the morning has developed, perhaps because of ecological constraints, competition between males might reinforce the tendency for synchronous dawn song. (For a parallel discussion of synchronous calls to attract females see Partridge and Krebs, 1978.) While competition and interference may produce intraspecific synchrony of dawn songs, it may also lead to interspecific asynchrony and account in part for the fact that species in the same habitat show a dawn peak of song at slightly different times.

Another factor which might influence the dawn chorus is the 'forgetting' of territory boundaries that may take place during the night. Fordham (1978) has suggested that low intensity of feeding activity at dawn in the moorhen Gallinula chloropus is due to "more territorial disputes and increased social activities as residents and prospective immigrant breeders reassert their social positions after the long dark period of reduced or nil contact with neighbours". The functional interpretation of Fordham's observation may be similar to that proposed for intruder ecological constraints or the possibility of overnight death of territory holders may favour early morning attempts to expand boundaries.

Concluding this review of possible functional factors related to the timing of the dawn chorus, it seems likely that any single factor hypothesis will be unable to account either functionally or causally for the diversity of species and ecosystems

where the dawn chorus has been observed. The most promising approach to this problem may be to consider the dawn chorus as a special case of trade off between competing activities. Careful study of the ecology of each population will tell which are the competing pressures operating on the time budget of individual animals, and by assessing the costs and benefits of performing different activities as a function of time of the day it will be possible to derive daily routines as rigorous predictions of optimality models (Sibly and McFarland, 1976; McFarland, 1977). In Part I of this thesis I attempt to advance in that direction by reducing the problem of the dawn chorus in the great tit to the analysis of the costs and benefits associated with foraging and territorial defence during the first hours of the morning.

2.3. The Dawn Chorus in the Great Tit

The Great Tit (Parus major) in Whytham Wood, Oxford, is territorial from mid-January until late May (Hinde, 1952; Krebs, 1971 ; Webber, 1975). During this period the birds wake up at about the time of civil twilight (when the sun is 6^{o} below the horizon) and after a short period of calling and preening the males usually engage in a long burst of song. As the sun rises, song becomes more intermittent, and the birds start to move about in the canopy, presumably feeding. About one hour after sunrise they start to spend more time on the ground, where they do most of their foraging. From then onwards, great tits spend much of their time foraging, occasionally singing from high branches when they detect the presence of an intruder within the territory or when they hear another male singing. The main features of this

daily routine are well documented for the Great Tit (Hinde, 1952; Kluijver, 1950; Hunter, 1978 and unpublished data, Krebs pers comm, personal observations) but they are not understood in either causal or functional terms.

Hinde (1952) and Hunter (unpublished data) recorded the number and length of great tit's songs at different times of the day, and found in both cases that the total amount of song produced decreased during the course of the morning (Figure 2.3). In both cases the pattern is very similar, with a very obvious decline of the amount of song within the first couple of hours and a more or less stable rate of singing later on. Slagsvold (1977) also observed a morning peak of singing in the genus Parus and claims that this peak is particularly pronounced in early spring.

In addition to these quantitative data which refer to average song rates for all the birds in an area, observations of individual males by both Hinde (1952) and Hunter (pers comm.), as well as my own unpublished observations showed that some males spend periods of up to one hour singing from the same perch before they start to feed. While no detailed information on time budgets of Great Tits in the wild is available, it is likely that the time allocated to singing is not a time of active feeding in this species. This can be concluded from microhabitat studies by Hunter (1978), who showed that preferred places for the two activities are separated (Figure 2.4). This separation of feeding and singing places can be related to functional considerations:

Great tits are well adapted to foraging on the ground (Lack, 1971) where they do not have to compete with other related species (Morse, 1978), but sound attenuation is much more severe near the ground than in the higher branches (Wiley and Richards, 1978; Marten and Marler, 1977; Roberts, Kacelnik and Hunter, 1979a,b).

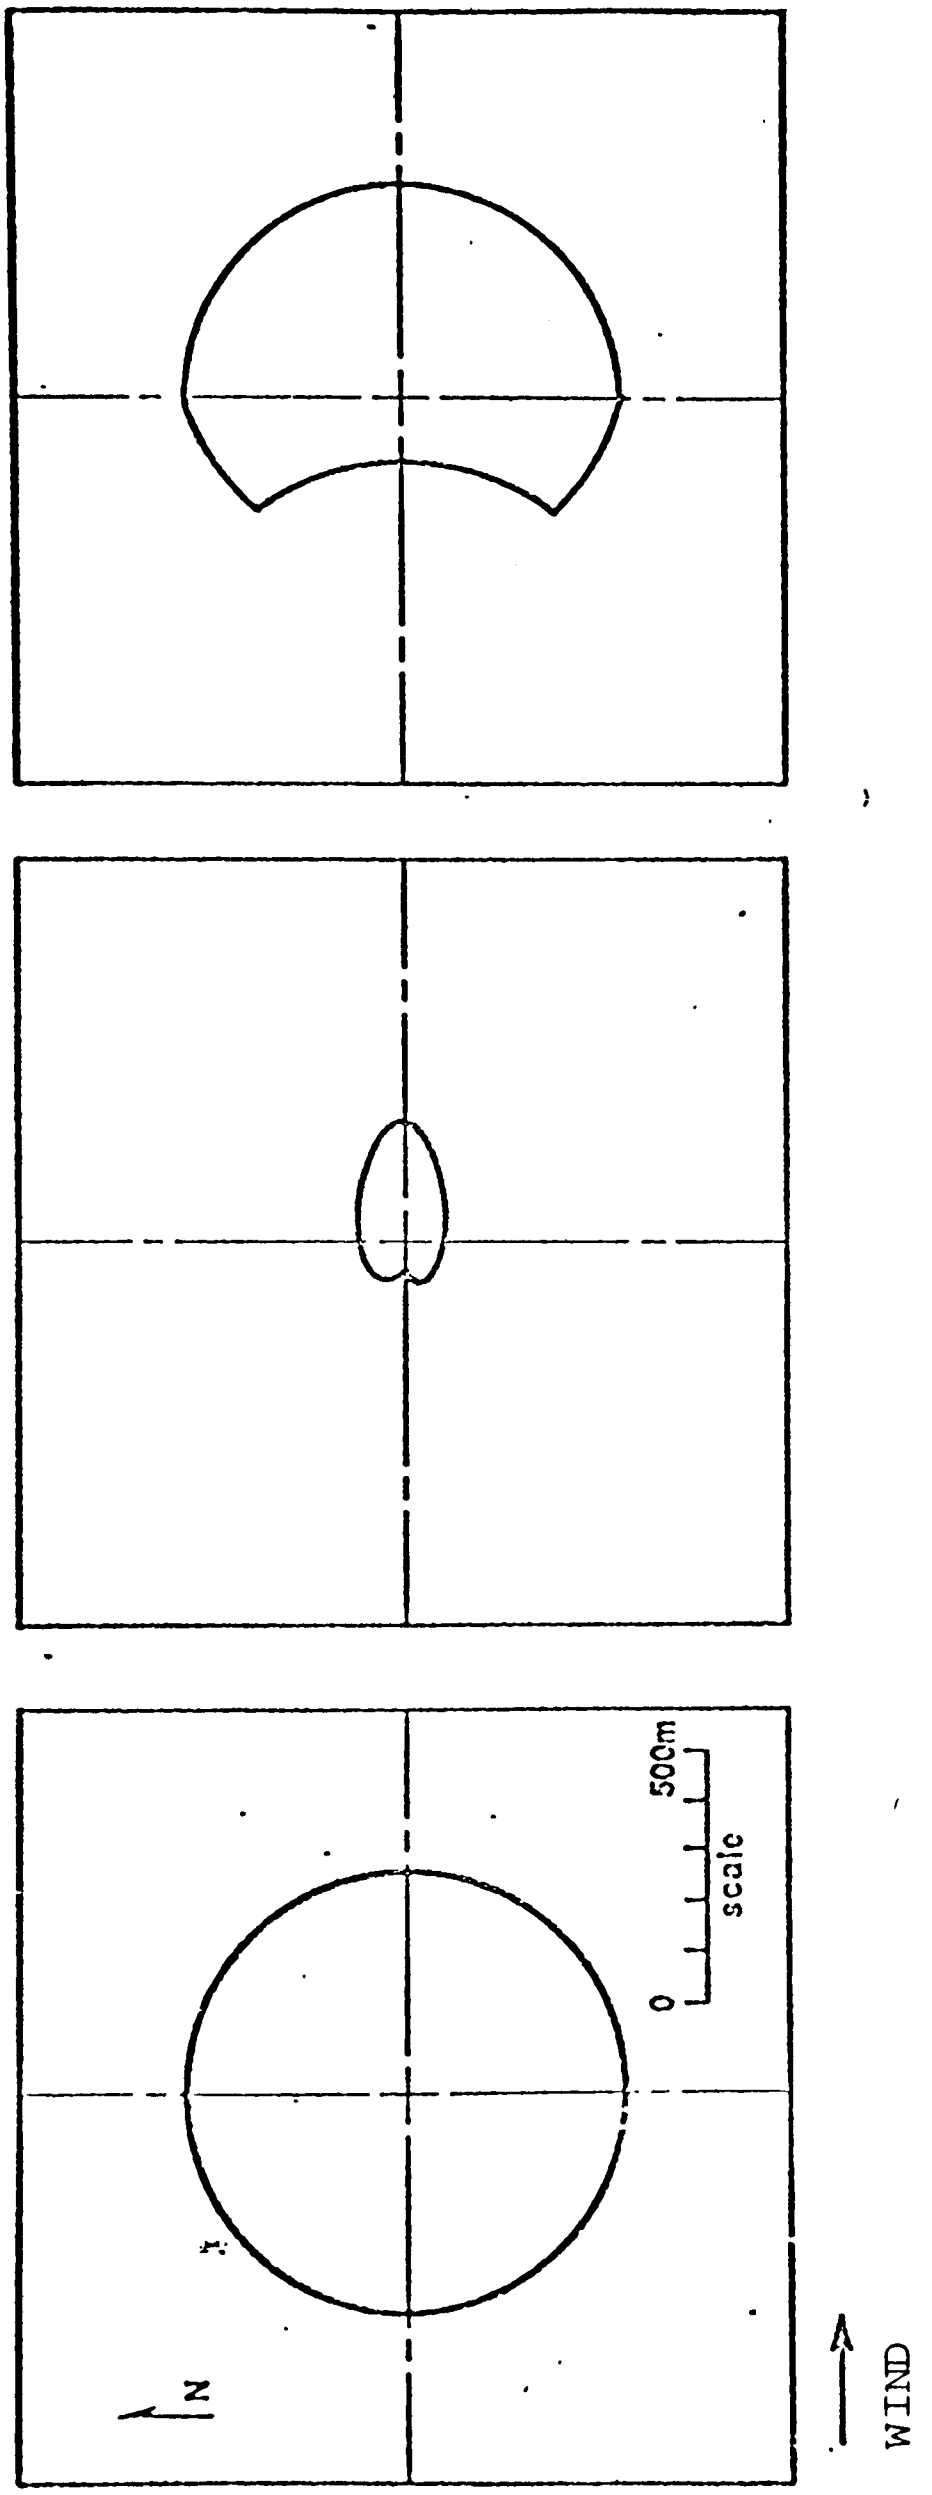
Since feeding and singing are partly incompatible,by allocating most of their time to singing, territorial great tits are paying the cost of not foraging. This cost appears at first sight to be very high because the energy reserves are at their lowest daily point just after the overnight period of fasting. In the related species Parus atricapillus, Chaplin (1974) found that the fat stores in midwinter evenings provide only slightly more energy than expended overnight. In spite of their probable ability to decrease nocturnal metabolic rate by hypothermia (Steen, 1958; Hissa and Palokangas, 1970, but see Grossman and West, 1977) great tits are 10% lighter in the morning than in the evenings (Owen, 1954, Chapter 4 of this thesis).

Song in the great tit seems to be largely associated to territorial maintenance, and territory defence does not appear to increase short term feeding rate, since when food is very scarce (e.g. in snowy weather) the birds abandon their territories. The function of territory seems in the great tit more related to event: occurring considerable time after the season when song is more intense, like reducing predation by spacing out the nests (Krebs, 1970, 1971 ; Dunn, 1977) or protecting the area most heavily used to collect food for the chicks (Krebs, 1971). Therefore, the

question to be discussed in the first part of this thesis becomes: Why is it that great tits in Spring sing instead of foraging early in the morning?

I will look at some of the ecological constraints of both activities, at the mechanisms of responding to overnight weight loss and at how foraging opportunity influences territorial responses to intruders.

Figure 2 . 1



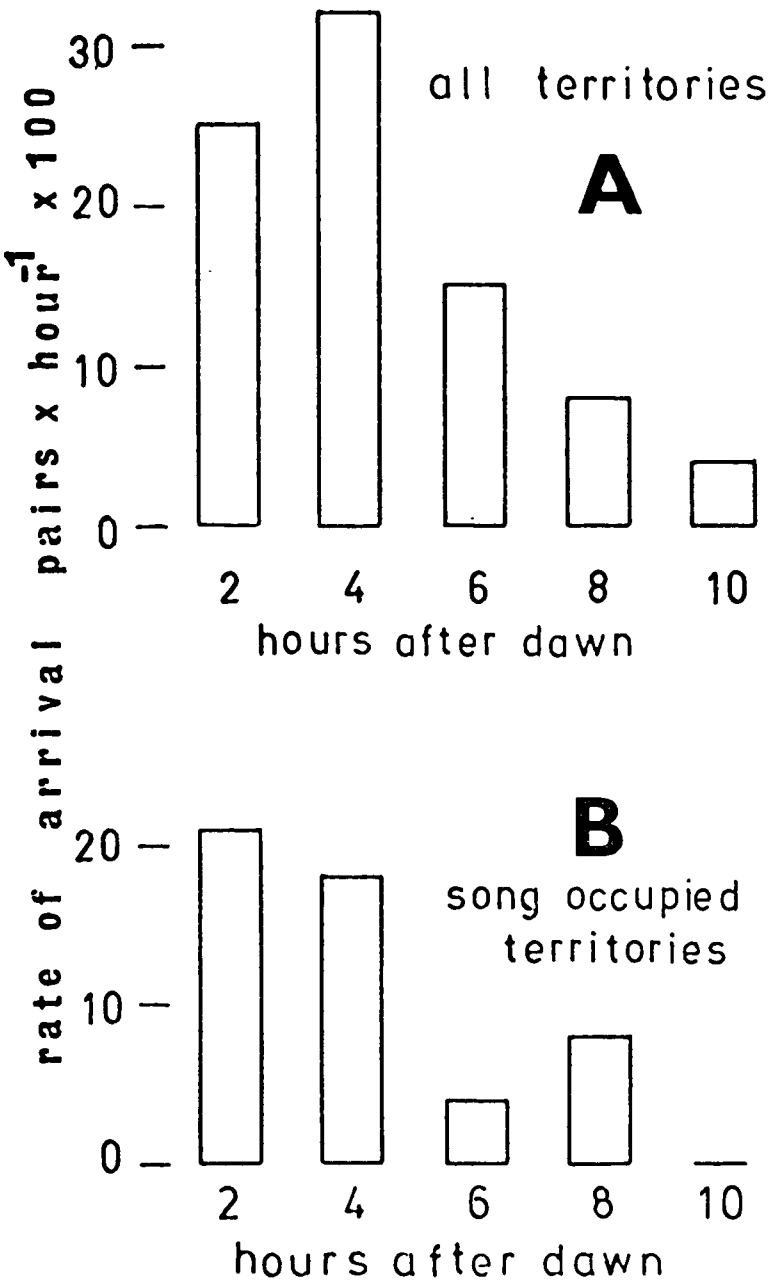
0500 hrs
Area = 7.1 x 10^5 m^2

1200 hrs
Area 3.0 x 10^4 m^2

2000 hrs
Area = 3.2 x 10^5 m^2

Predicted broadcast area in a desert micrcclimate at three different times of the day. Source frequency = 3000 Hz and wind from the west. From Henwood and Fabrick (1979).

Figure 2.2



Rate of arrival of territorial invaders at different time of day. Krebs and Perrins (1977) removed great tits from their territories and recorded the number of new claimants at different times of day. The high rate of arrivals early in the morning was not due to the absence of song in the empty territories, since territories where the owner's song was broadcast (B) showed the same effect as all the territories taken together.

Figure 2.3: Data from Hunter (A) and Hinde (B) show that there is a peak in great tit singing activity at dawn in Wytham Woods. From Hunter (unpublished) and Hinde (1952).

Figure 2.3

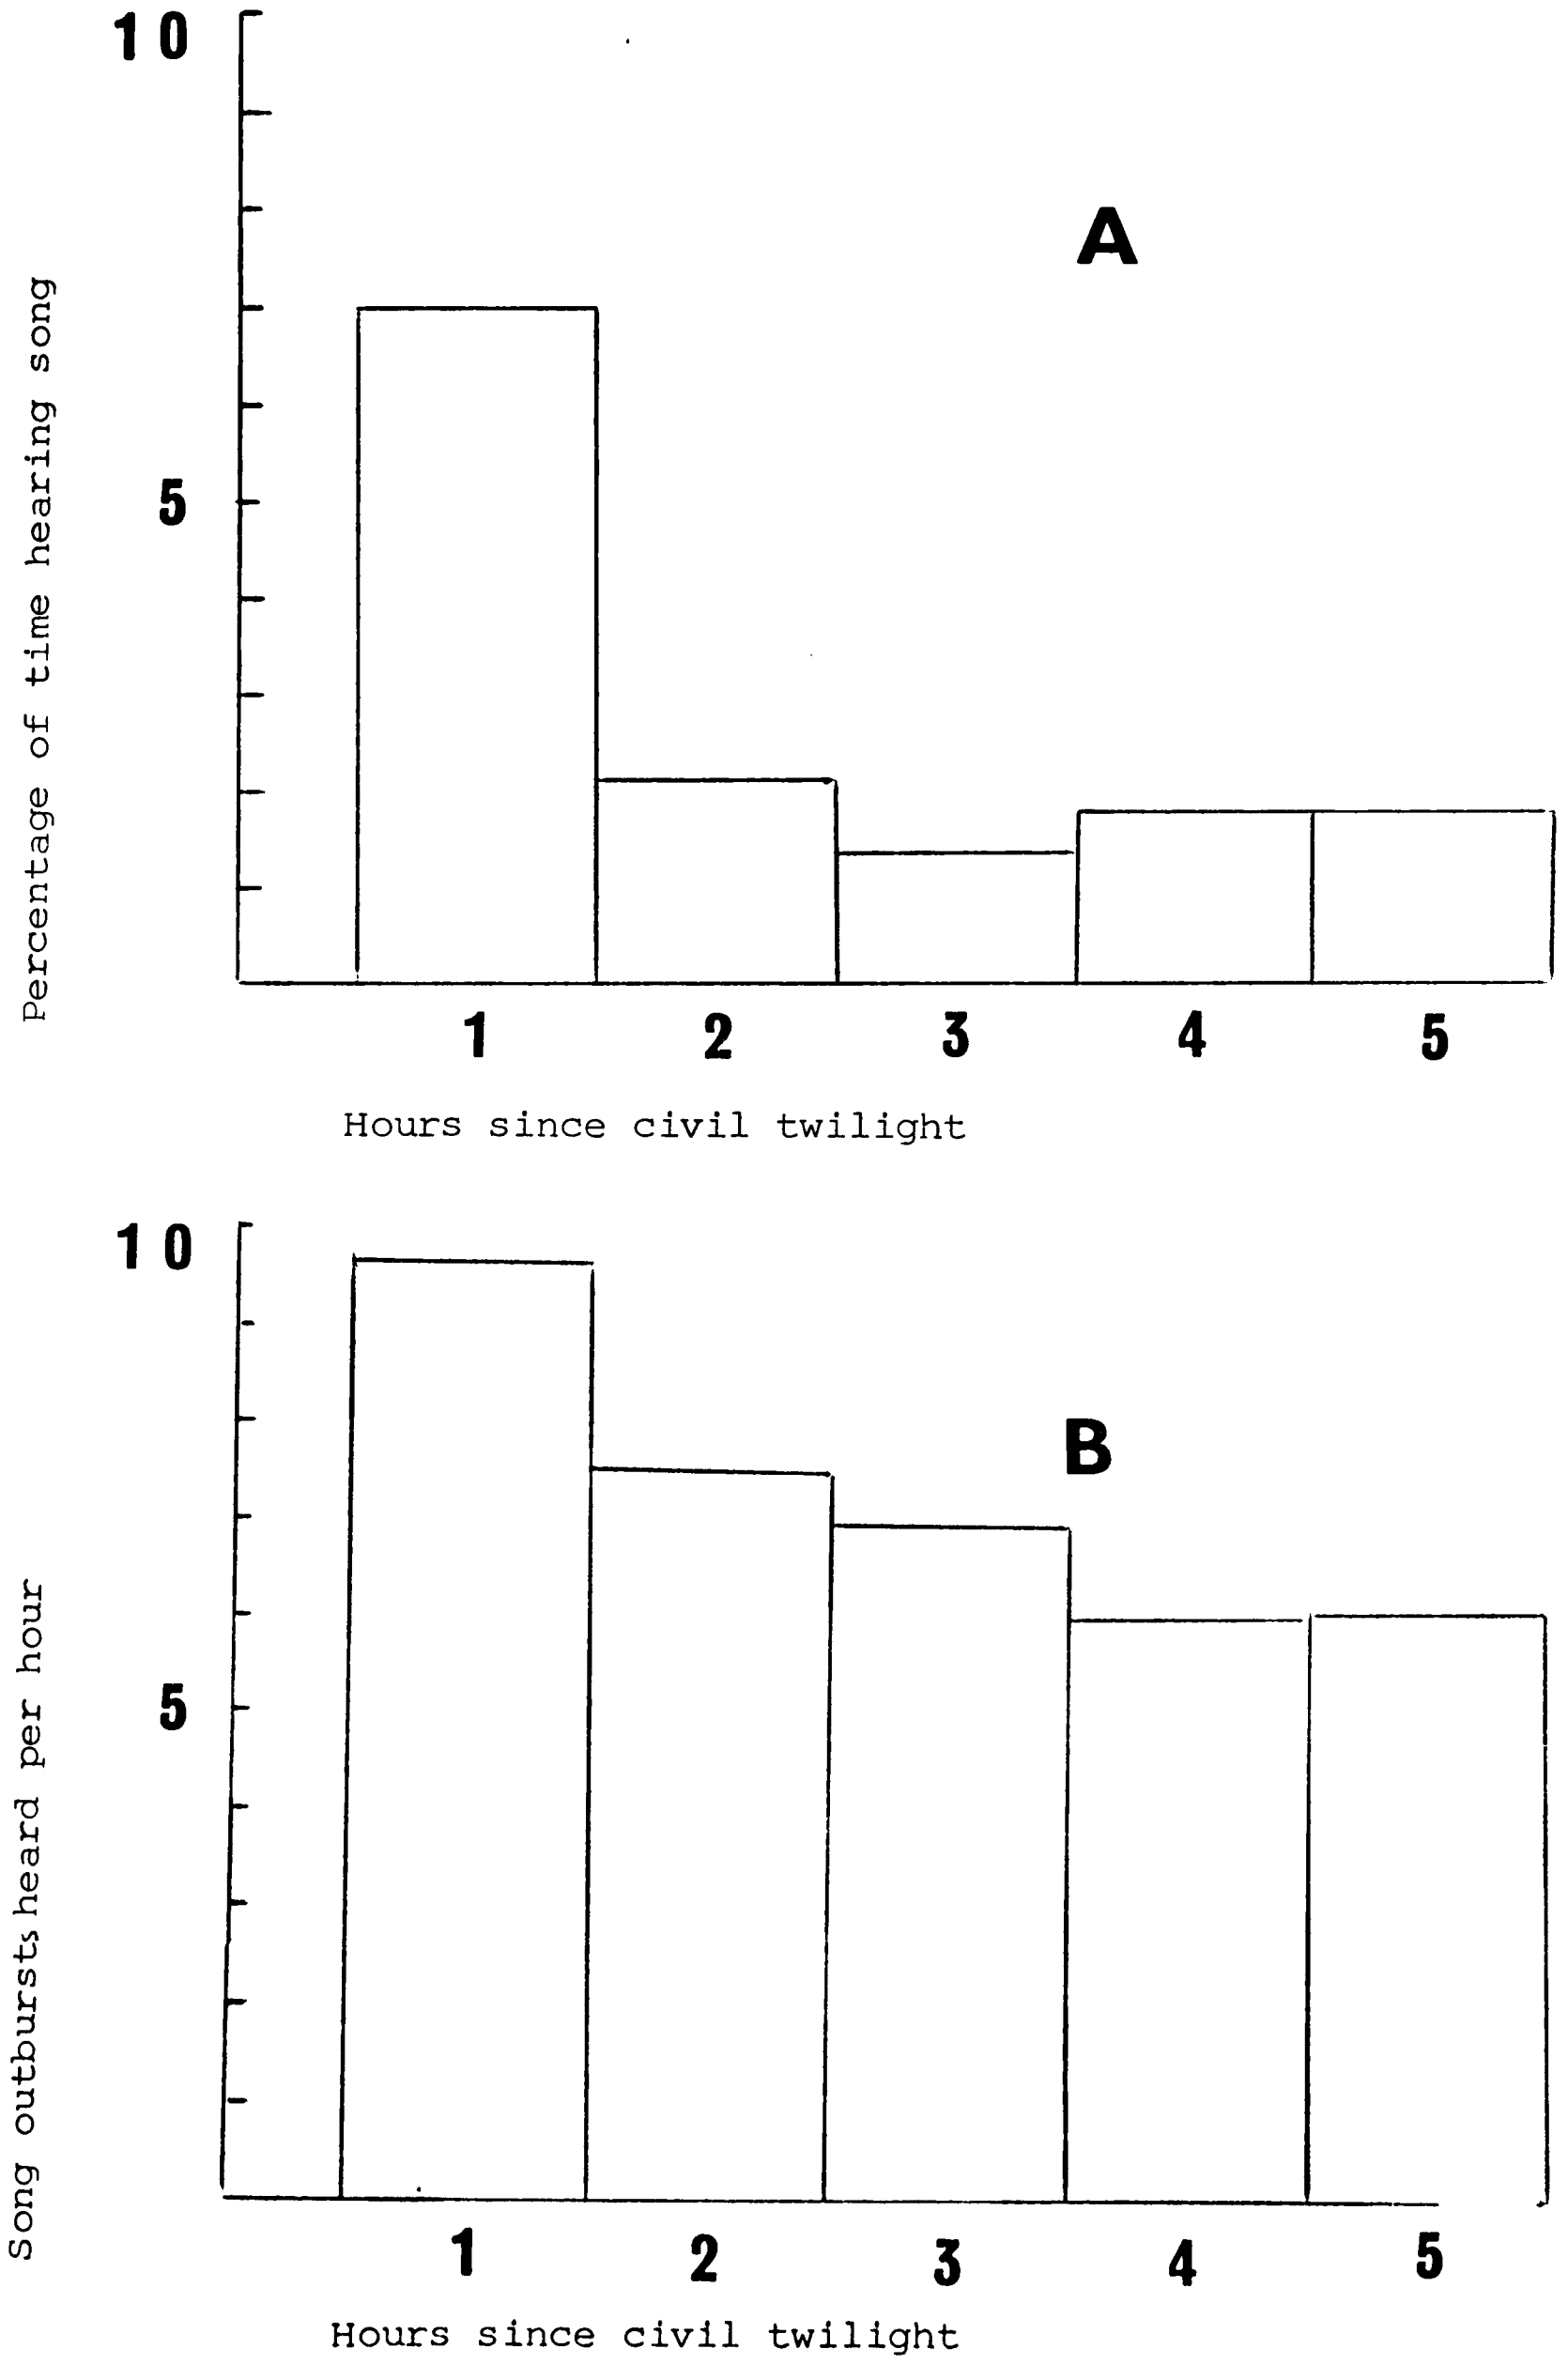
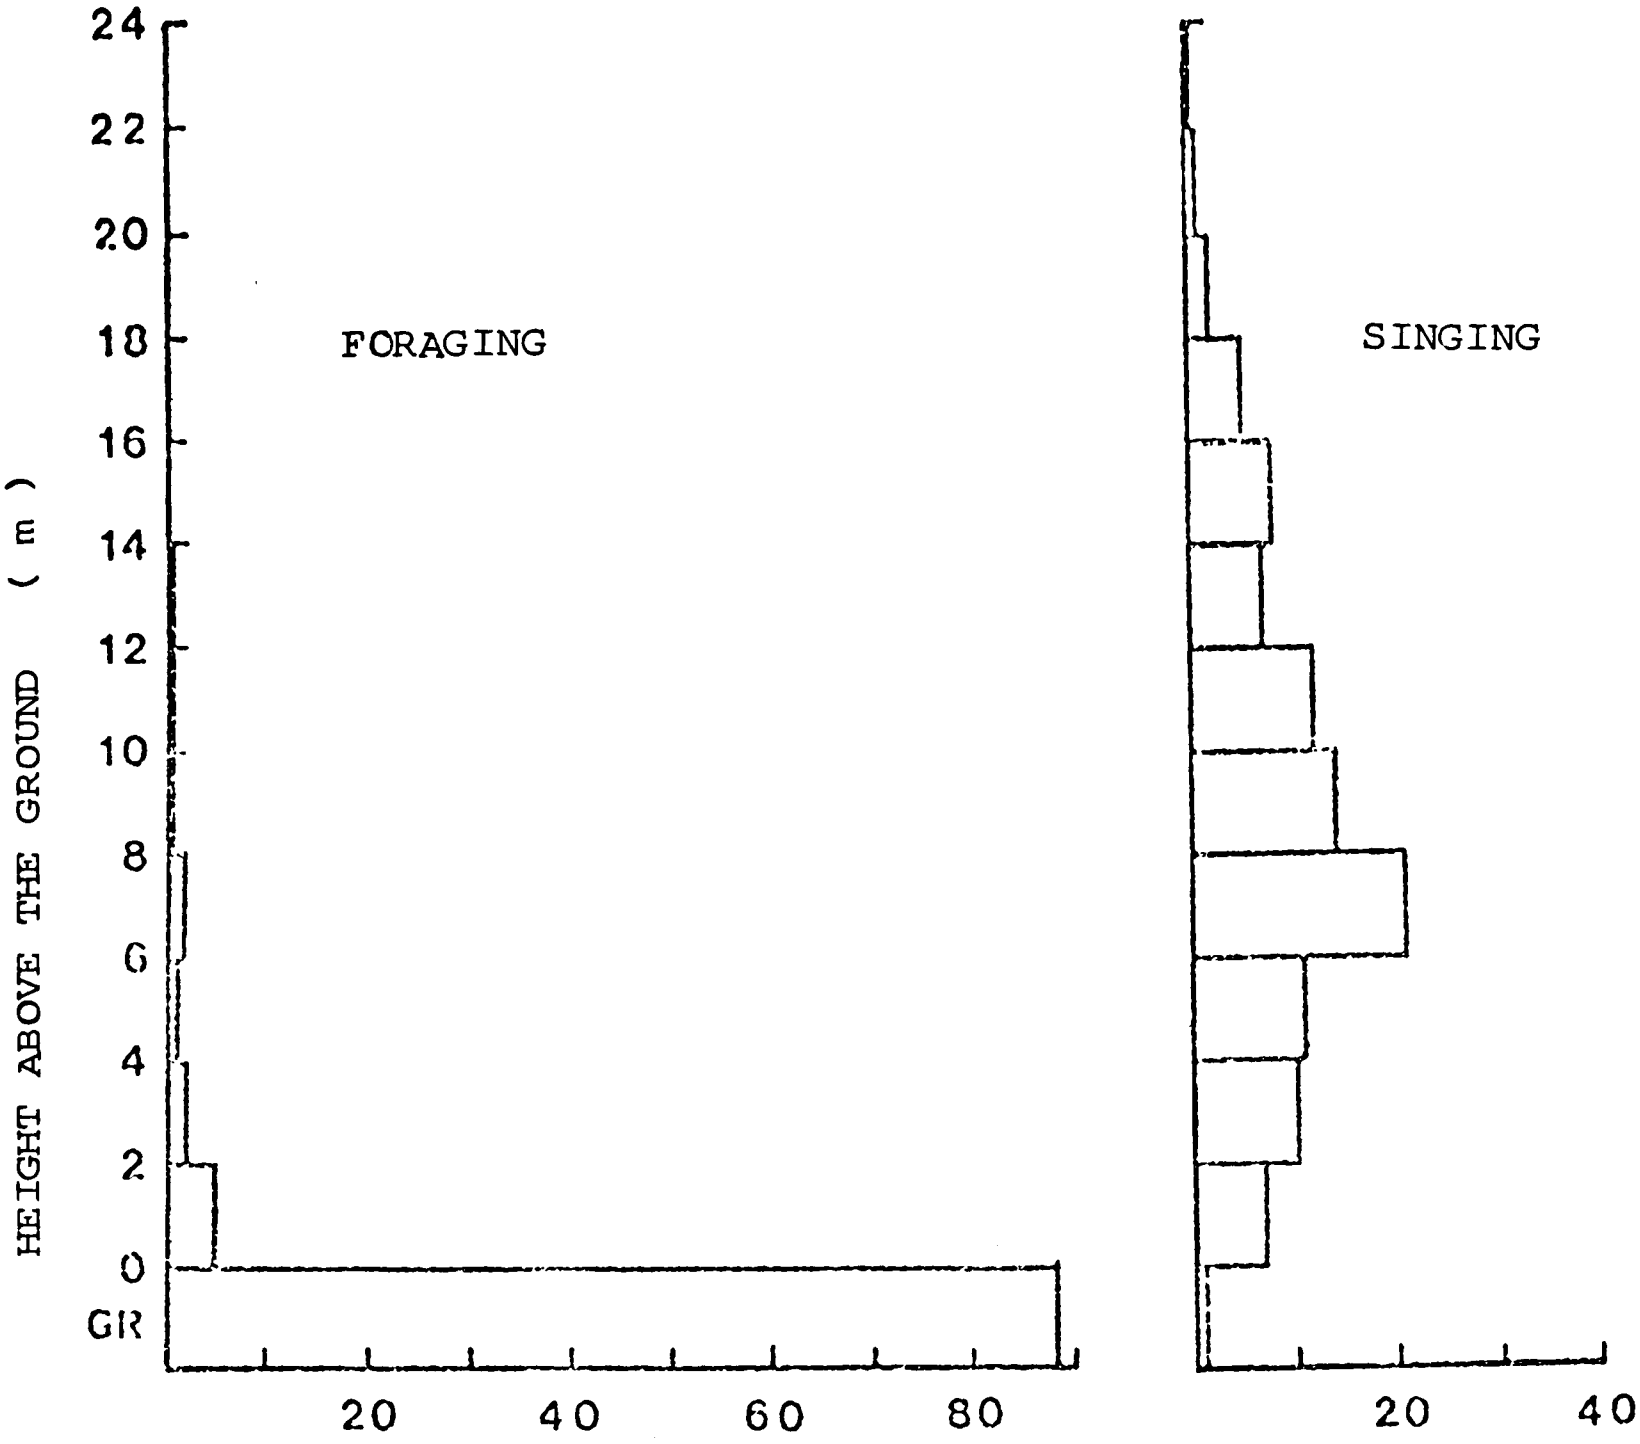

10
5
A
Percentage of time hearing song
1
2
3
4
5
Hours since civil twilight
10
5
B
Song outbursts heard per hour
1
2
3
4
5
Hours since civil twilight

FIGURE 2 . 4



Frequency distribution of foraging and singing at various heights above the ground(GR) upward at 2 m intervals.From Hunter(1978).

CHAPTER 3

Environmental Factors affecting the Dawn Chorus

3.1. Introduction

In the previous chapter I discussed several hypotheses of how physical constraints in the environment may favour vocalising early in the morning. The most important factors discussed were light intensity, temperature and wind. The aim of this chapter is to examine how these variables might influence the dawn chorus of the great tit in Oxfordshire. The chapter is divided into three sections. In the first I shall look at the influence of light on foraging efficiency. If light could be shown to limit foraging efficiency most severely at the time of the dawn chorus an advantage could be suggested for concentrating on activities not affected by light such as singing or preening until illumination increases. The effect of temperature is discussed in the second section. Two hypotheses relating temperature to the dawn chorus were presented in Chapter 2. Temperature may affect foraging opportunity either through its effect on prey mobility or through differential effects on microhabitat conditions. I studied the first possibility by measuring the locomotory activity of ground dwelling invertebrates. I compare both hypotheses with data on average temperatures in Oxfordshire.

Section three concerns the possible acoustic advantages for early morning vocalisations. I shall present some preliminary measurements of the effect of wind (the main physical factor affecting opportunity for acoustic communication) on transmission of sound, and relate these measurements to the dawn chorus by means

of average records of wind as a function of time of the day in Southern England.

3.2. Light Intensity

In this section I describe an experiment designed to test if light intensity might lower feeding efficiency at the time of the dawn chorus. I measured foraging efficiency of great tits in the laboratory under different light intensities, covering the range of prey luminance measured in the field during the first two hours after twilight, when the dawn chorus takes place.

3.2.1. Methods

3.2.1.1. Subjects

I used 12 adult great tits. All were experimentally naive and had spent at least 45 days in captivity. They lived individually in metre cube cages provided with detachable aluminium roosting boxes in which the birds roosted and hid when mildly frightened (see GRAL methods). These cages were kept in a room at approximately 20^{o}C, with light from 08:00 to 19:00 hours. The standard diet was provided ad libitum outside the experimental period and only after 13:00 hours during this. All tests took place between 09:00 and 13:00 hours in a room at 10^{o}C. Each bird was tested only once each morning, and the order was randomised to avoid circadian effects and consistent difference in deprivation between subjects. The mean body weight loss between evening and the beginning of a test was 8% which is within the normal range of body weight loss under natural conditions.

3.2.1.2. Apparatus

The experimental set up is shown in Figure 3.1. It consisted

of two similar aviaries measuring 90 x 90 x 60 cm and separated by a wooden wall. The birds could cross from one to the other through a passage which could be opened or closed by a remotely controlled sliding door. Illumination was the same in both sides, provided by a pair of 150 watt reflector lamps on each side. The four lamps were connected in parallel to a single dimmer which was used to vary light intensity between treatments. The waiting compartment had three perches, and the searching arena was filled up with leafless branches and twigs in which the prey were hidden. The branches were re-arranged every day so that the birds could not remember details of the searching arena. When a test was due to start the subject was caught in its roosting box and this was attached to the waiting compartment (start box in Figure 3.1). The bird's body weight was determined by subtraction before every test.

3.2.1.3. The prey

Fifty artificial prey were distributed in the searching arena before each test. A prey item consisted of a piece of a mealworm (mean weight 14 mg) enclosed in a pastry cocoon made of flour and a water solution of food dyes in a proportion of two to one in volume (see "ravioli" in Gral methods). The water solution was made of 8 ml of green, 4 ml of red, and 2 ml of yellow "Boots food colouring" with water up to 100 ml. The birds handled the prey by pecking open the cocoon and eating the mealworm piece. They did not eat the cocoon, the purpose of which was to make the prey match the background.

3.2.1.4. Testing

Each test consisted of two 15 min blocks. At the beginning of a test, the birds were released in the waiting compartment and allowed to accommodate visually to the current level of illuminatio for 15 min. The passage was then opened and the birds crossed to the searching arena where they hunted for prey for the next 15 min. Searching, handling and non-foraging times were recorded on a computer compatible event recorder (Dawkins, 1971). After finding a prey, a bird usually took it to an horizontal branch, where it held the cocoon with one or both feet while hammering with its bill. Once a hole was opened, it cleanly pulled out the piece of mealworm and swallowed it. After this, birds sometimes discarded the cocoon while others continued with its searching for remaining bits of food (Plate 3.1). I will discuss this point further when analysing the effect of light on handling times.

Each bird was subject to ten treatments, each one consisting of a test at a given luminance. Figure 3.2 shows the luminance values used in the ten treatments. The experimental range stretche from 0.01 cd/sq m to 10.00 cd/sq m. This covers the full range of luminance as measured on the same prey in the field during the first two hours after twilight (see Figure 3.7).

3.2.1.5. Units

References to light are not always clear in behavioural literature. The expression light intensity refers sometimes to luminous intensity (candela), luminous flux (lumen), illumination (lux) or luminance (cd/sq m), and it is difficult to compare between different authors. I shall follow the Royal Society (1968) units and the definitions of Walsh (1958) and Hailman (1978). For

an animal hunting for prey against a visual background the relevant variable affecting detectability is the subjective brightness of the visual target, which is both observer and background dependent, and can not be directly measured. However, subjective brightness varies in the same way as luminance (Walsh, 1958) and in consequence this latter is the best measure of light conditions for the type of problem that I am discussing here. Luminance (units: candela/ square meter, or cd/sq m) is the luminous intensity emitted by a surface. Its relation to illumination (the light intensity falling on a surface) depends on the reflection factor of the surface in question. All experimental conditions are referred to in accordanc to the luminance at the prey surface as measured with an S.E.I. spc photometer. To facilitate comparisons with other experimental situations Figure 3.2 also shows the luminance of a sheet of white matt paper as an approximation to a standard white surface of magnesium whose reflection factor is defined as unity.

3.2.1.6. Training

Before the experiment all subjects received 12 tests, one per day, at intermediate luminance (0.10 cd/sq m).

Figure 3.3 shows the median and range of the birds' performance during the training phase. Figure 3.3a presents the latency from the opening of the passage to the beginning of searching, and Figure 3.3b the number of captures achieved during the first three minutes after this. These two variables were chosen because they indicate the birds' familiarity with the task and readiness to work (latency) and their skill (capture rate).

Although I did not test statistically for asymptotic performance, it is apparent from Figure 3.3 that within the second half of the training period there was no major trend in any of the two measures. Since the order of treatments was randomised, any subsequent improvement would have been cancelled between conditions.

3.2.2. Results

Figure 3.4 shows that two measures of foraging success increased up to an asymptote with increasing light intensity. The first measure, profitability, is defined as prey caught per second foraging (searching and handling). The second measure, searching efficiency, is defined as prey captured per second of searching time. This second measure rises more steeply with light because it ignores handling time: at high light intensity the birds catch more prey and therefore spend more time handling (Holling, 1965), and this sets an upper limit to profitability. This effect is complicated by the fact that handling time is inversely related to light intensity (Figure 3.5), which reduces apparent profitability at low light intensity. The reason for the decrease in handling time with increasing light is that the birds extracted food from the cocoons more thoroughly when light intensity was low, spending additional time gleaning the open cocoon for extra pieces of mealworm. Most of the cocoons recovered after experiments with high light intensity had a clean hole through which the piece of mealworm had been extracted, while cocoons recovered after low light tests were fully opened as a consequence of this extra attention. Table 3.1 shows that the birds actually extracted more

food when they had fully opened the cocoons. It is interesting to note that the increase in handling time with increase in inter-catch interval is what would be expected from optimal foraging theory (Cook and Cockrell, 1978).

Figure 3.6 shows that the proportion of experimental time spent foraging increased with light intensity up to an asymptote of 100%. At low light intensity the birds spend as much as 50% of their time preening, resting or bathing.

In order to relate the laboratory results to the levels of light intensity in the field, I combined the laboratory results with measures of luminance at the surface of the artificial prey in the field (Figure 3.7). The luminance at which searching efficiency reaches 95% of its highest value (pointed by S in Figure 3.4) corresponds in Figure 3.7 to the light intensity about 60 minutes after the onset of the morning song on sunny days and more than two hours after in overcast days, more common in Oxfordshire. Profitability (indicated by P in Figure 3.4), reached nearly asymptotic values when luminance was close to 1 cd/sqm. This corresponded to 40 minutes after twilight in sunny days and over one hour in overcast days. The time at which the birds started to spend more than 95% of their time foraging (T in Figure 3.6) matched very closely the time for asymptotic profitability (see proximity of T and P in Figure 3.7). This is what one would expect if the bird's motivational tendency to forage matches the expected payoff under the circumstances, but it must be remembered that profitability depends on the prey's characteristic handling time and its density, as well as on the capture rate. It would be very interesting to

see if replicating the experiment at different prey densities or with different prey types the proportion of time devoted to forage is readjusted. The predictions derived from this experiment are that the shape of the curve for searching efficiency would be rather resilient to changes in prey type, but that since profitability would be affected according to the disc equation (Holling, 1965), proportion of time foraging ought to follow the same readjustment.

3.3. Temperature

In this section the possible effect of temperature on foraging opportunity at the time of the dawn chorus is assessed indirectly by studying the variations on locomotory activity of several flightless invertebrates. Prey activity is related to daily fluctuations in temperature in the field.

3.3.1. Method

The locomotory activity of four different potential prey types was measured as a function of temperature.

They prey types were spiders (*Zygiella x-notata*), adults and larvae of the flour beetle (*Tenebrio mollitor*), and larvae of the fly *Calliphora vicina* . All of these were highly palatable prey for the great tit and were similar to prey normally captured by great tits in the wild (Betts, 1955).

A total of 100 measurements of activity were taken for each prey type, ten at each of the following temperatures (in centigrade -10, -7, -4, -1, 0, 1, 4, 7, 10 and 20.

Each individual prey item was used only once. They were kept in a room at 20 degrees centigrade with a light cycle of 12L:

12D. Each day a sample of ten individuals of each type was collected and placed in individual containers in a climatic chamber at 95% relative humidity and one of the temperatures listed above. The light cycle was the same as before.

The next morning, starting at 09:00 hours (one hour after lights on) each of the 40 specimens was placed individually in an arena 0.15 m wide and 0.25 m long with a 0.025 m grid marked on the floor. The arena was kept within the climatic chamber.

The subject was left alone for 2 minutes and the number of squares of the grid touched during the second of the two minutes was recorded. Due to their slow movements, three recording minutes were allowed for the larvae of T molitor (mealworms). The order in which subjects belonging to each of the four prey types were tested was randomised. After testing all 40 individuals the temperature of the chamber was set to a new value (the order of the test temperatures was also randomised) and 40 new specimens were put in the chamber after the new temperature had been reached

3.3.2. Results

Figure 3.8 presents the mean activity records of the four prey types as a function of temperature. Locomotory activity of the different prey seemed to be affected by temperature in similar ways. The insect prey showed an "activity threshold" at around 7°C, while the spiders maintained some level of activity below that temperature. From 7°C upwards there was a steep increase in activity in all four prey types, which started to level off between 10°C and 20°C. These results agree in general with those of

Gunn and Hopf (1941), who showed that the flighless beetle Ptinus tectus increase its level of activity steeply between 1° and 10°C. The only difference is the lower activity threshold reported for P tectus than found in this study.

In order to assess if prey activity is likely to vary during the time of the dawn chorus these results can be compared with Figure 3.9a, which shows average temperatures in Oxford as a function of time of the day for the months of February, March and April, 1979. In February temperatures during the whole day are below the activity threshold for the insect prey in this study, and low enough for the spiders to be fairly inactive, but in the other two months it looks as if an increase in prey activity could be expected to occur some time after the period of the dawn chorus.

In summary, prey activity fluctuations due to daily variation of temperature may affect foraging opportunity at the time of the dawn chorus in March and April if prey movements influence foraging efficiency (Robinson, 1966), but its effect is probably negligible before that. Ground frost though is very common in Wytham in February and temperature may have an effect on the dawn chorus during this month by increasing the thermal cost of foraging on the ground early in the morning.

3.4. Wind

This section concerns the possible acoustic advantages of early morning vocalisations. The effect of wind on attenuation of sound of various frequencies covering the main frequencies in great tits' song and on background noise generation was measured

in a representative area of Marley Wood, Wytham.

3.4.1. Methods

A tone generator in series with a Nagra DH amplifier was used to generate pure tones (0.5, 1.0, 2.0, 4.0 and 8.0 kHz). The tones were recorded on a loop tape in a Stellavox SP7 tape recorder, and played back in the field using the same tape recorder and amplifier in combination with an E.R.A. Acoustic Corporation polyplanar speaker (dimensions: 4 inches by 8 inches). The speaker was mounted at the top of a 5 m high aluminium pole. Sound Pressure Level (SPL) was recorded using a Bruel and Kjaer type 2203 Precision Sound Pressure meter with associated type 1613 octave band filter and a Uher 2300 tape recorder. The tape was analysed in the laboratory using a Bruel and Kjaer type 2112 Audio Frequency spectrometer and type 2305 Level recorder. The receiving microphones were placed at the same height as the speakers. The frequency filters of the SPL meter were set for each tone to filter background noise. Background noise was measured in silences between pure tone broadcasts.

The loop tape lasted 5 minutes, with 50 seconds of each tone separated by 10 seconds of silence. Broadcasting started 15 minutes before civil twilight and were repeated at 15 minute intervals until noon, unless climatic conditions (rain, stormy weather) forced an earlier interruption. This happened on several weekends, so that I was left with full records for only two consecutive days (8th and 9th April, 1978). The tests had to be interrupted by bad weather because the recording equipment,

specially the microphones, could not be sheltered without affecting their acoustic performance (see Miller, 1978 for new techniques for recording sounds under rain).

Wind speed was recorded before and after every broadcast with an anemometer fixed 5 m above the ground. The anemometer was reliable for winds stronger than 0.5 meters per second.

3.4.2. Results

In the two occasions when records for the whole morning could be completed, the pattern was very similar. Wind within the wood was nearly nil (0-0.5 m/s) until about one hour after sunrise, when a slight breeze (approximately 0.5 m/s) started to blow from the North-North-East sector. Since the speaker - receiver axis was N-NE to S-SW the wind direction was within 45^{o} of the direction of the sound. As the morning advanced, the wind became gusty, reaching peaks of 2-3 m/s for periods of a few seconds. Because of the small sample size the results will be described qualitatively Figure 3.10 presents typical records obtained before and after sunrise. The effect of the slow breeze was in general to affect the stability of the record, without changing the baseline sound pressure level. Gusts of 2 m/s produced drops in sound pressure level of approximately 10dB, or 5 dB m/s tending to be more noticeable at intermediate frequencies. Gusts approximately 2 m/s produced drops in signal SPL of 9.6 dB (SD = 3.6, n = 13). Assuming for simplicity that the effect of wind speed on SPL is linear, this means a drop of about 5 dB for every increase of one meter per second of wind velocity. Wide band background noise, resulting mainly from the movement of leaves in the surrounding trees, in-

creased by an average of 18.7 dB for gusts of the same speed (SD = 5.8, n = 17). Making again the assumption of linear effect of wind velocity, this gives an increase in background noise of about 9 dB for every increase of one meter per second of wind speed. If these figures are confirmed, they imply a loss of 14 dB in signal to noise ratio per m/s of wind speed.

These results can be related to the dawn chorus using Figure 3.9b, which presents average wind speeds as a function of time of the day for the months of February to April of 1979. The figure shows that in March, for instance, mean wind speed is close to 6.5 m/s at the time of the dawn chorus and 8 m/s at noon, so that communicatory efficiency at noon would be 21.0 dB lower than at dawn.

While my results must be taken with caution because of the small sample size, they tend to confirm the predictions of Henwood and Fabrick's model (1979), in that communicatory advantages are likely to be gained from using acoustic signals at dawn instead of later on in the day.

3.5. Conclusions

The results presented in this chapter show that several ecological factors favour singing at dawn in the great tit.

First, the opportunity for performing what I hypothesised to be the main competing activity - foraging - is very low at dawn because light intensity limits capture rate. Temperature may also contribute to decrease foraging opportunity at dawn in March and April, because the average temperature variation during the

morning at that time of the year is within the range when locomotory activity of different ground invertebrates is very sensitive to temperature. The effect of temperature on foraging is not known because I did not measure how prey activity influences great tits' foraging efficiency. When live spiders were released after testing for their activity in the aviaries where captive tits lived, more active spiders - those just tested at higher temperatures - were detected and captured quicker, but also took longer to handle. These effects were not quantified, so that further experiments are needed before concluding which of them is the dominant. In addition to the fact that profitability is limited during the time of the dawn chorus, two other observations made in the experiment on the effect of light are worth pointing out. The birds exploited more thoroughly individual prey and spent a lower proportion of their time foraging at low rather than at high luminance. In both cases the effect may be a response to light intensity or to rate of prey capture, but those two factors would probably vary together in the field as they did in my experiment. The increase in individual prey exploitation means that great tits can to some extent compensate the decrease in capture rate by adapting other components of their foraging behaviour. The drop in proportion of foraging time observed at low light intensities supports my assumption that more time would be available for singing when foraging opportunity is poor, and argues against a purely homeostatic control of foraging behaviour (see Chapter 4). If the main factor controlling the occurrence

of foraging was internal deficit as opposed to environmental supply, one would predict that the lower the feeding rate the higher the deficit would be and in consequence the harder the animal would work for food. In the experimental situation the birds used the non-feeding time mainly for preening and bathing, and in some cases they just sat immobile on their perches. The effect of feeding opportunity on territorial defence is investigated in Chapter 5. The analysis of the effect of wind on acoustic transmission supports predictions of Henwood and Fabrick's model that acoustic attenuation is lower at dawn than at other times of the day.

In conclusion, it is clear that no single factor hypothesis can account for the ultimate causes of the dawn chorus.

Table 3.1

Sixty fully open and sixty perforated cocoons were classified according to the presence or absence of traces of food. More fully open cocoons were totally empty ($x^2 = 5.40$ $p < 0.05$).

	Some food left	No food left
Fully open	10	50
Perforated	21	39

Figure 3.1

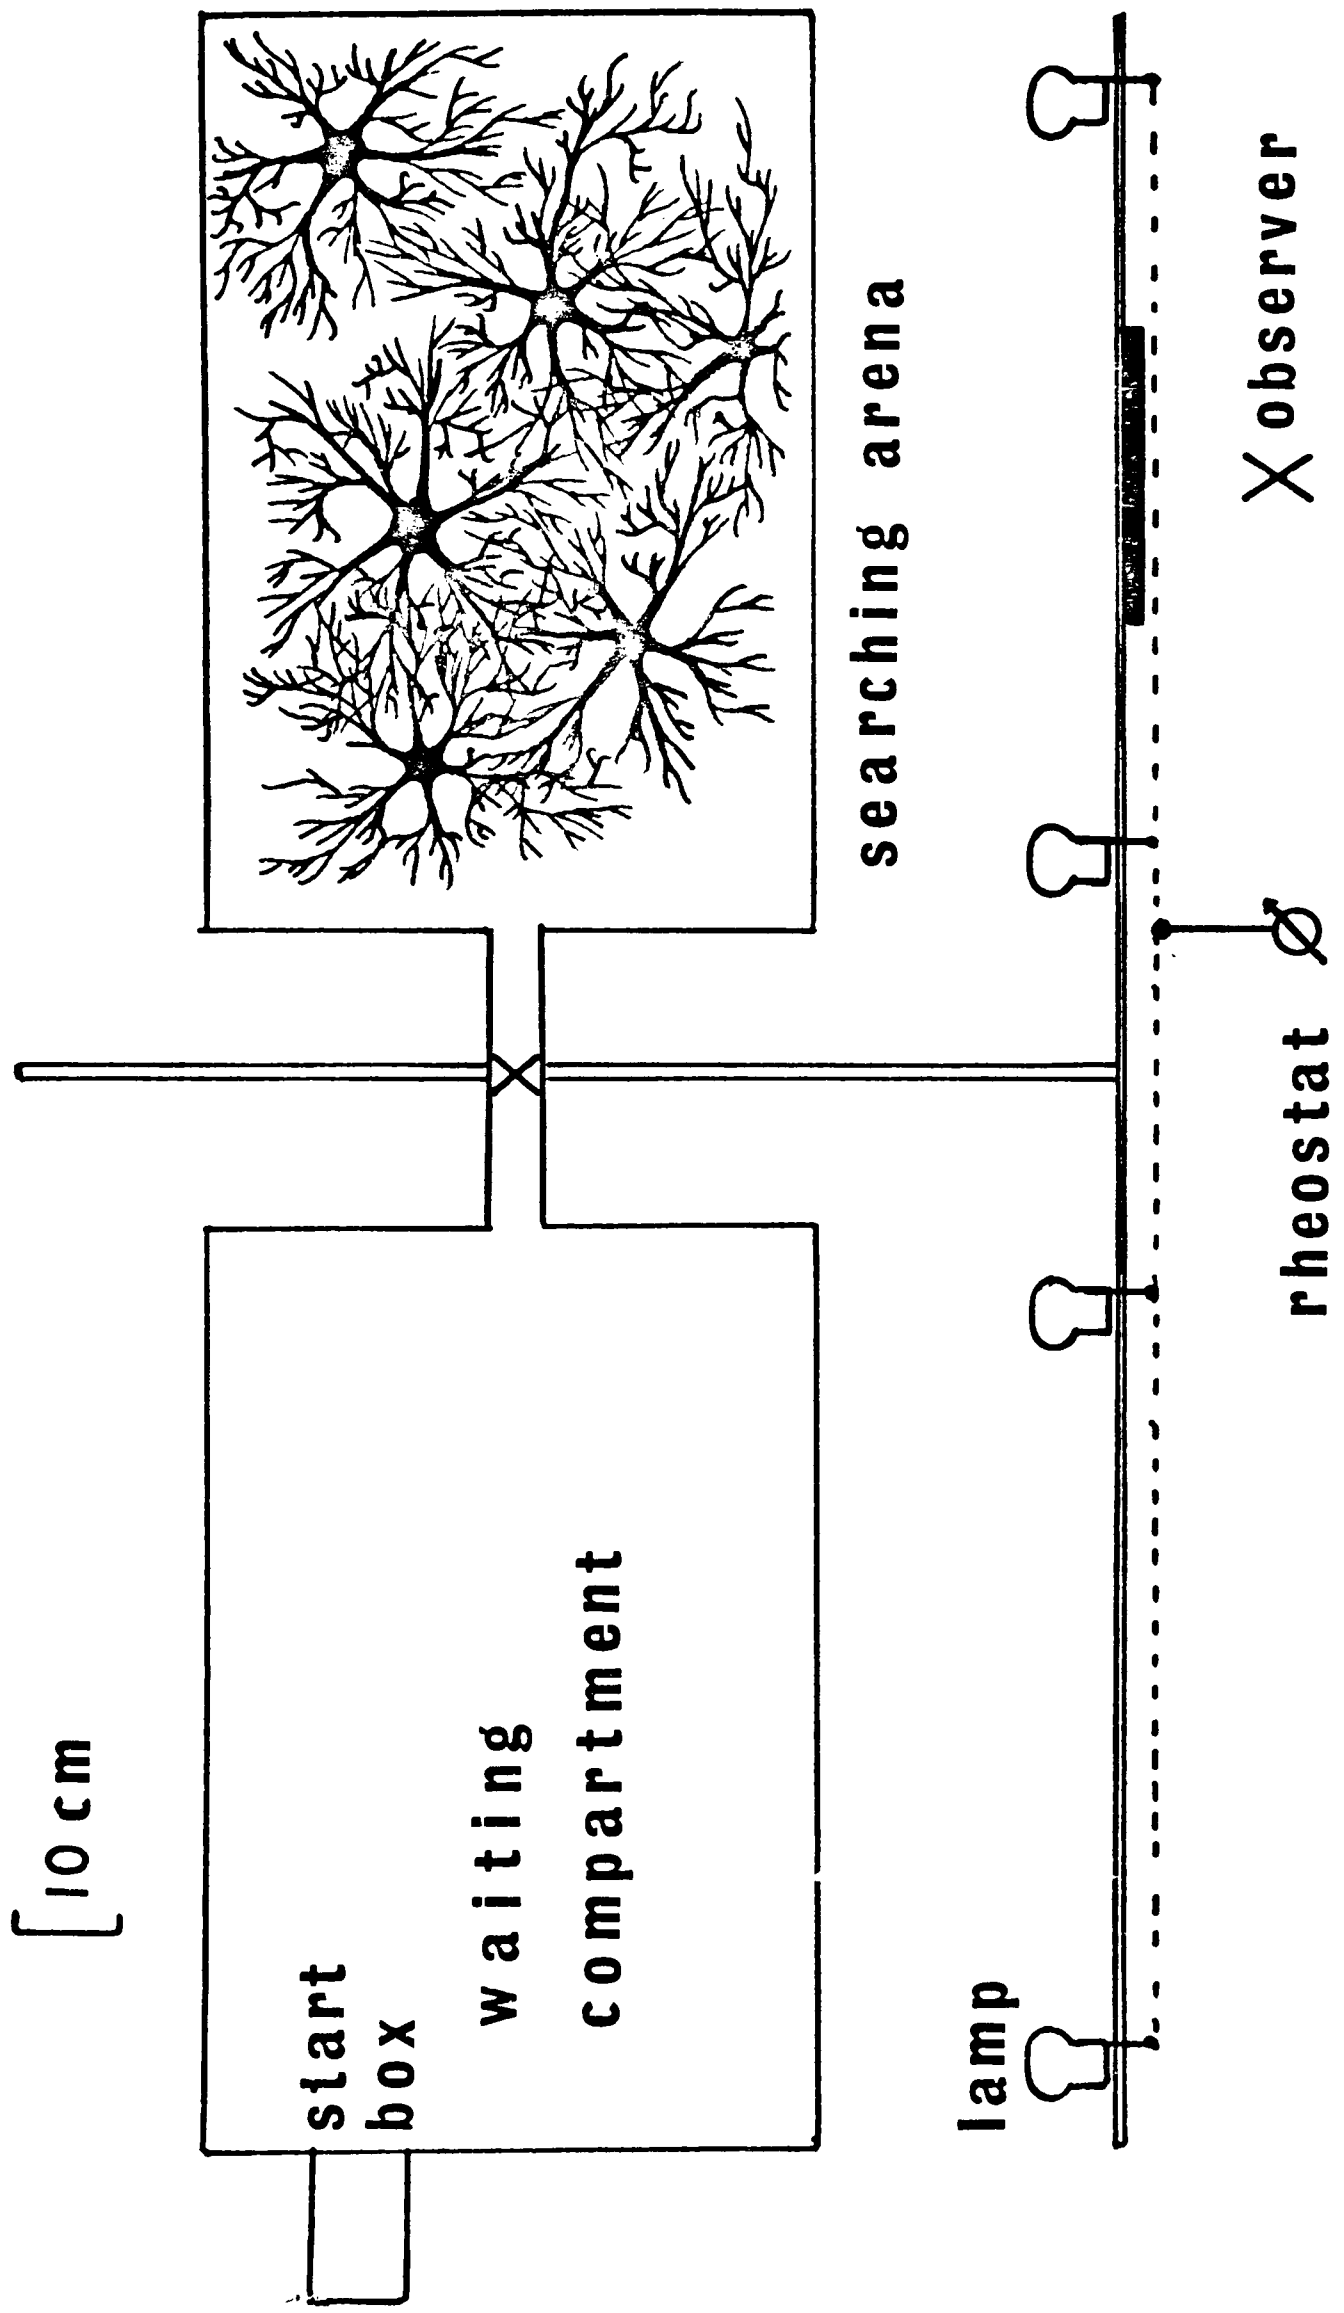
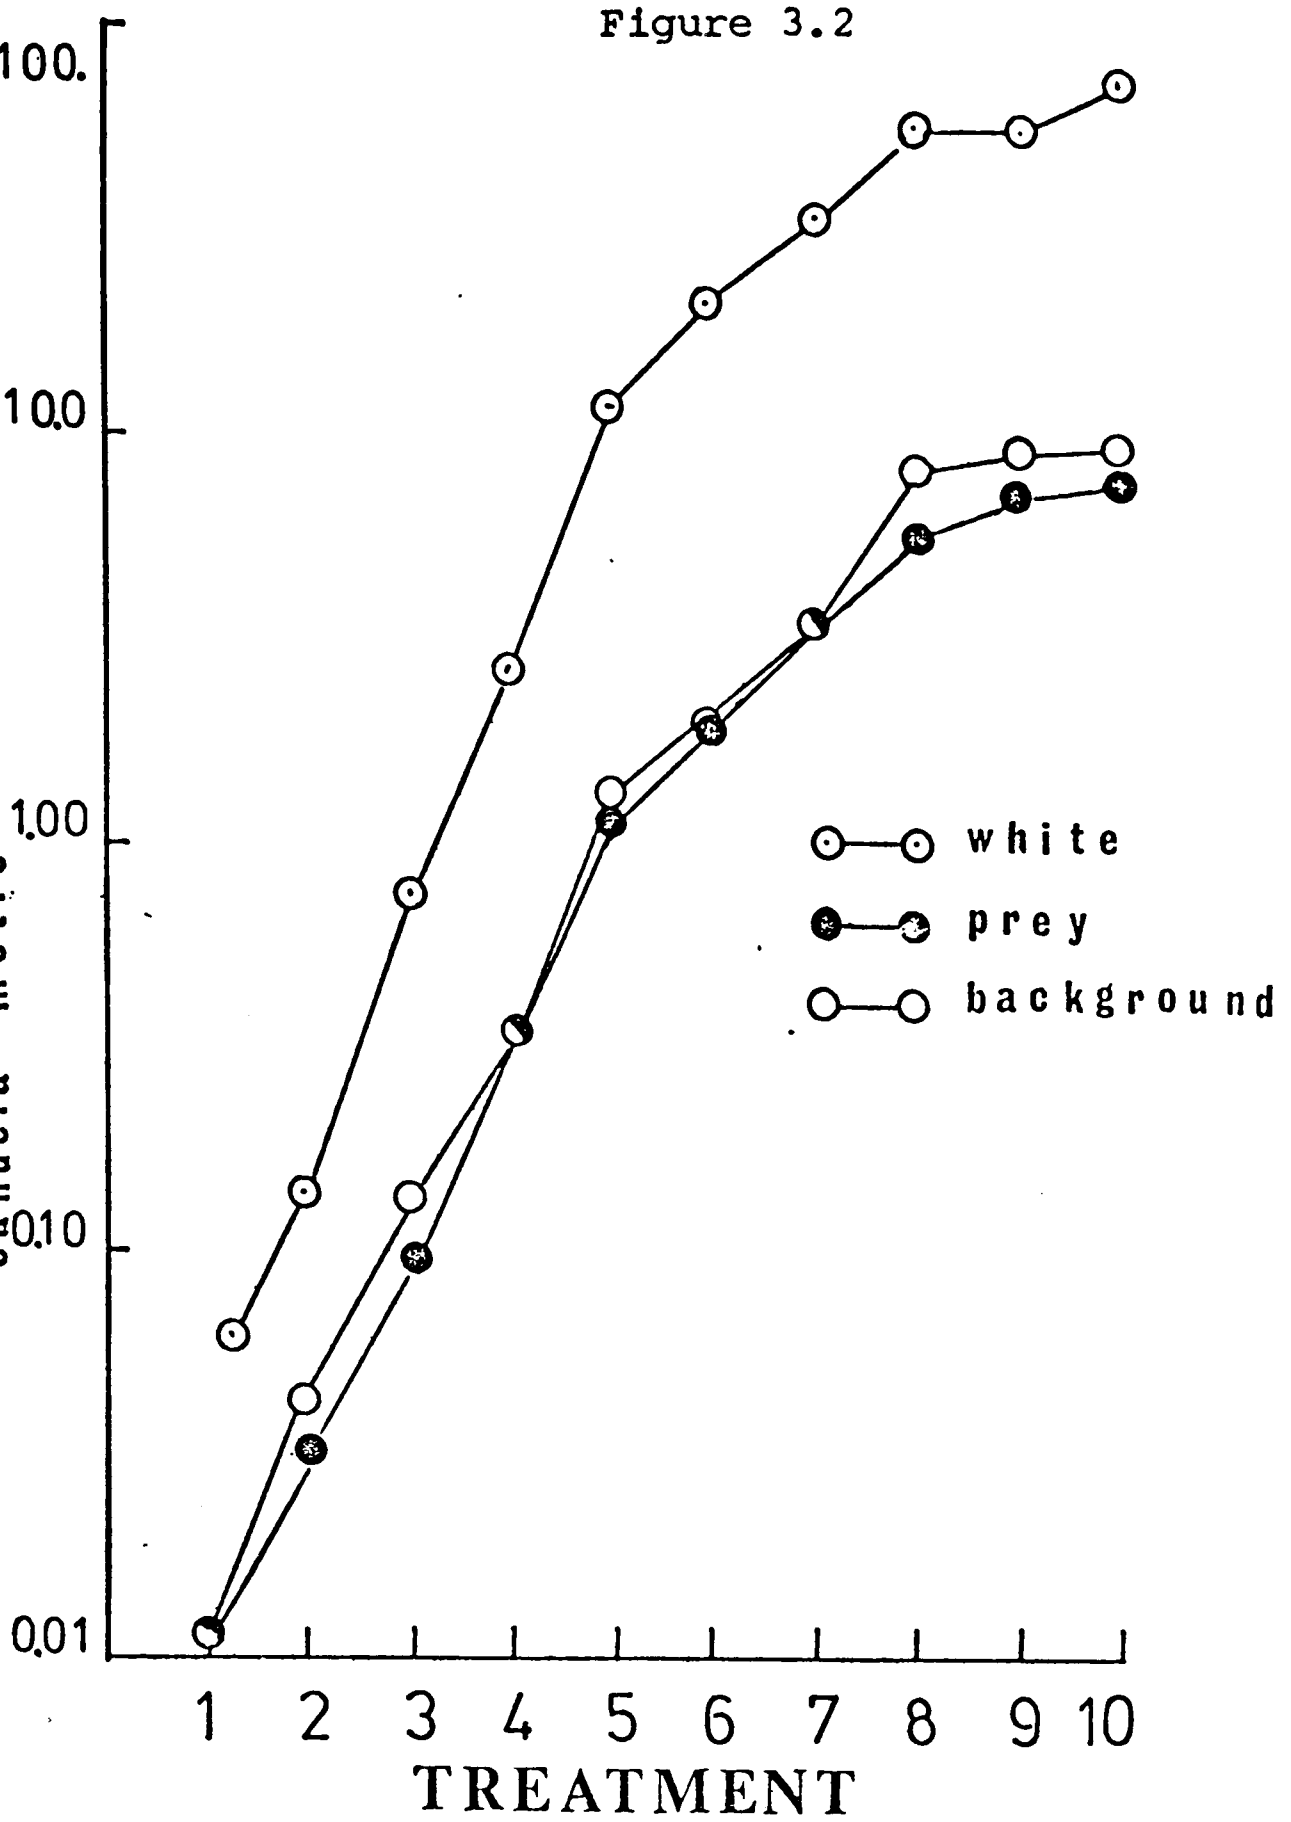


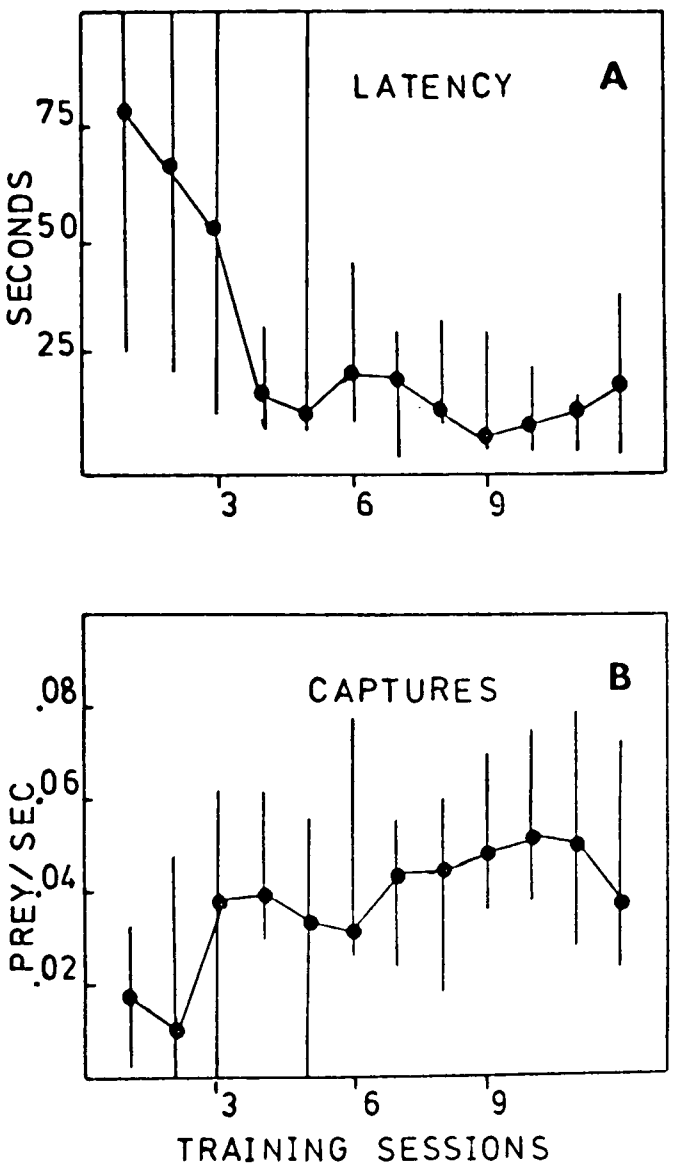
Diagram of the experimental set up.

Figure 3.2



Luminance at prey, background and matt white paper surfaces for the ten different treatments. The values for the white surface are shown to facilitate comparison with other experimental situations.

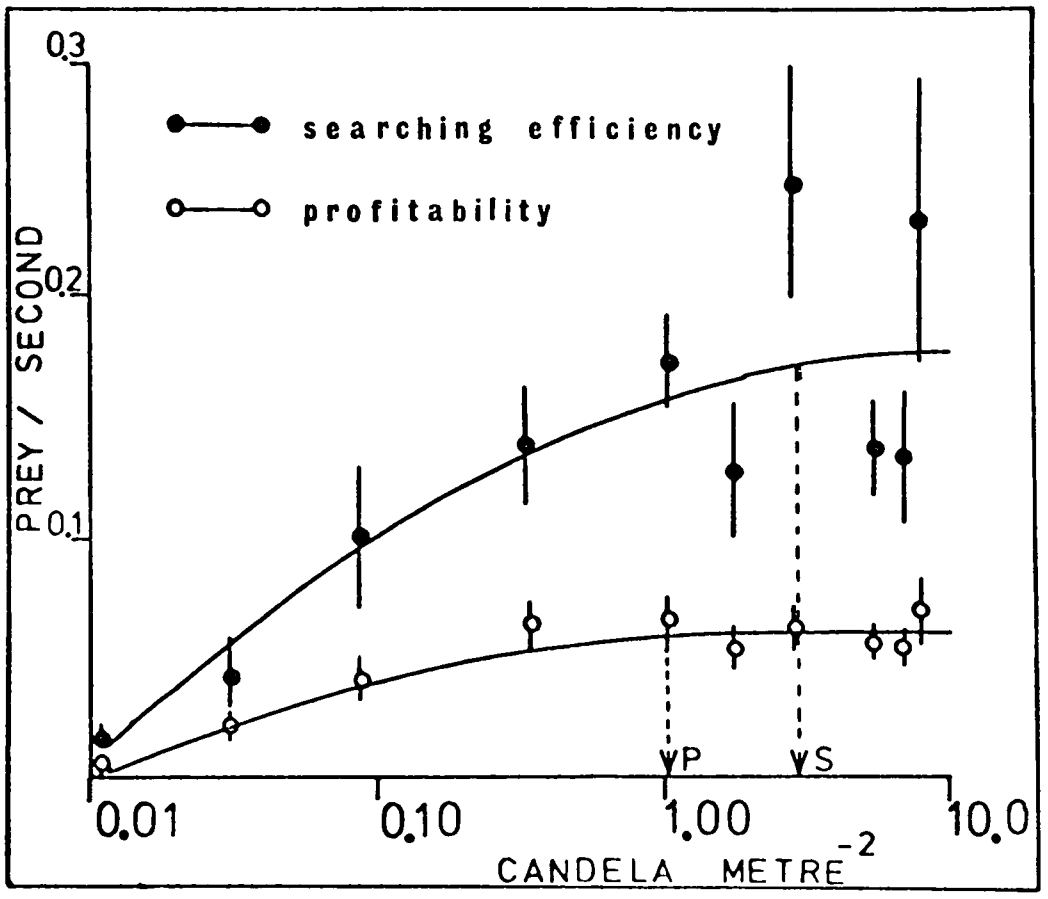
Figure 3.3



A. Latency (seconds) from opening the entrance to the searching arena until the bird began to search for food reached an asymptote after about 6 days of training.

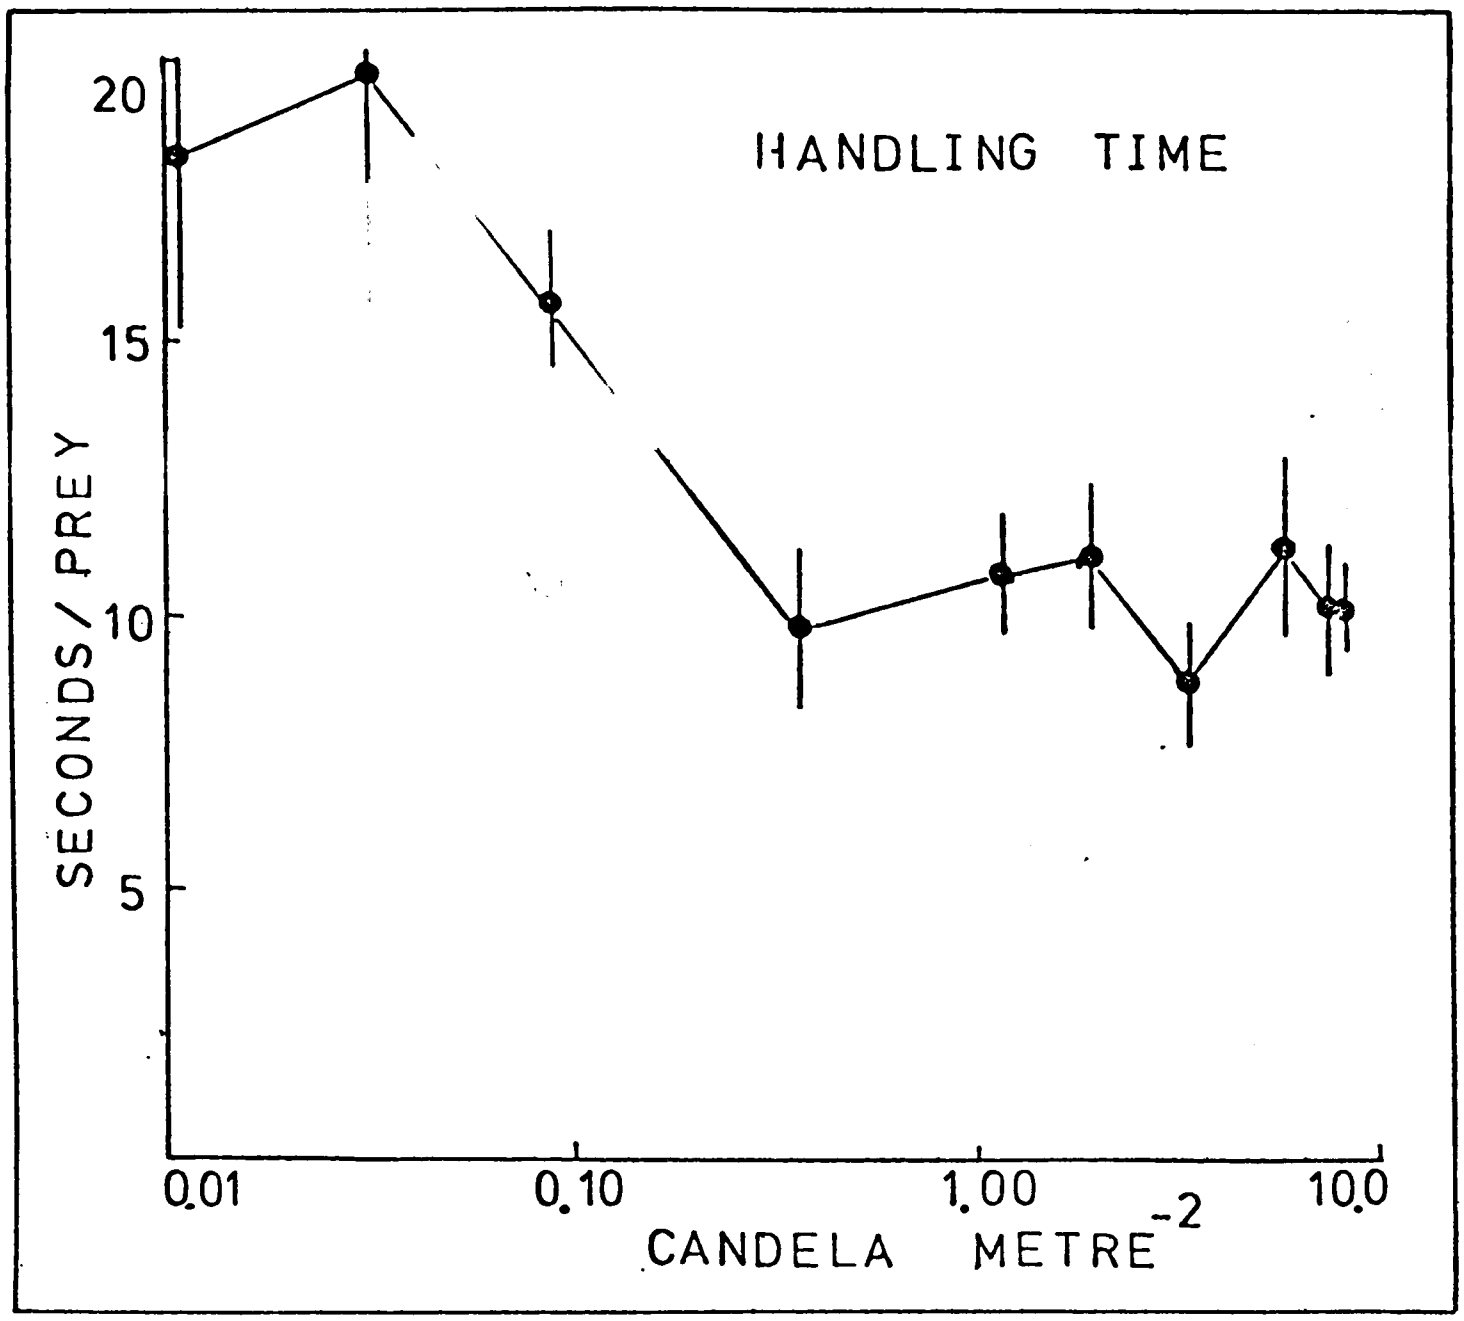
B. Capture rate (prey/second) during the first 3 minutes of hunting also reached an asymptote. Medians and ranges are shown for consecutive training sessions.

Figure 3.4



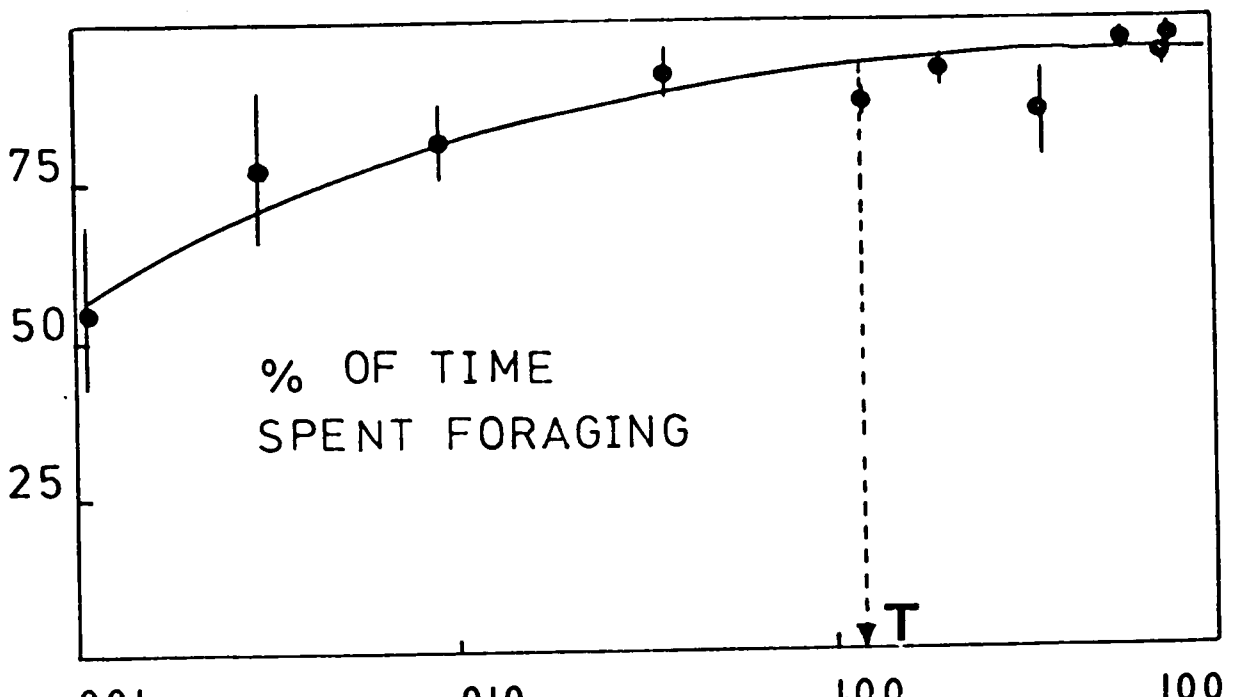
Profitability and searching efficiency for different light intensities (for definitions see text). The curves are second order polynomials fitted to the original data, but only means and standard errors are plotted for clarity of the Figure. In both cases, second order polynomials with negative quadratic terms were better fits ($p<0.05$) than linear ones, indicating that values were reaching a plateau. P and S indicate the luminance values at which profitability and searching efficiency reach 95% of their maximum value.

Figure 3.5



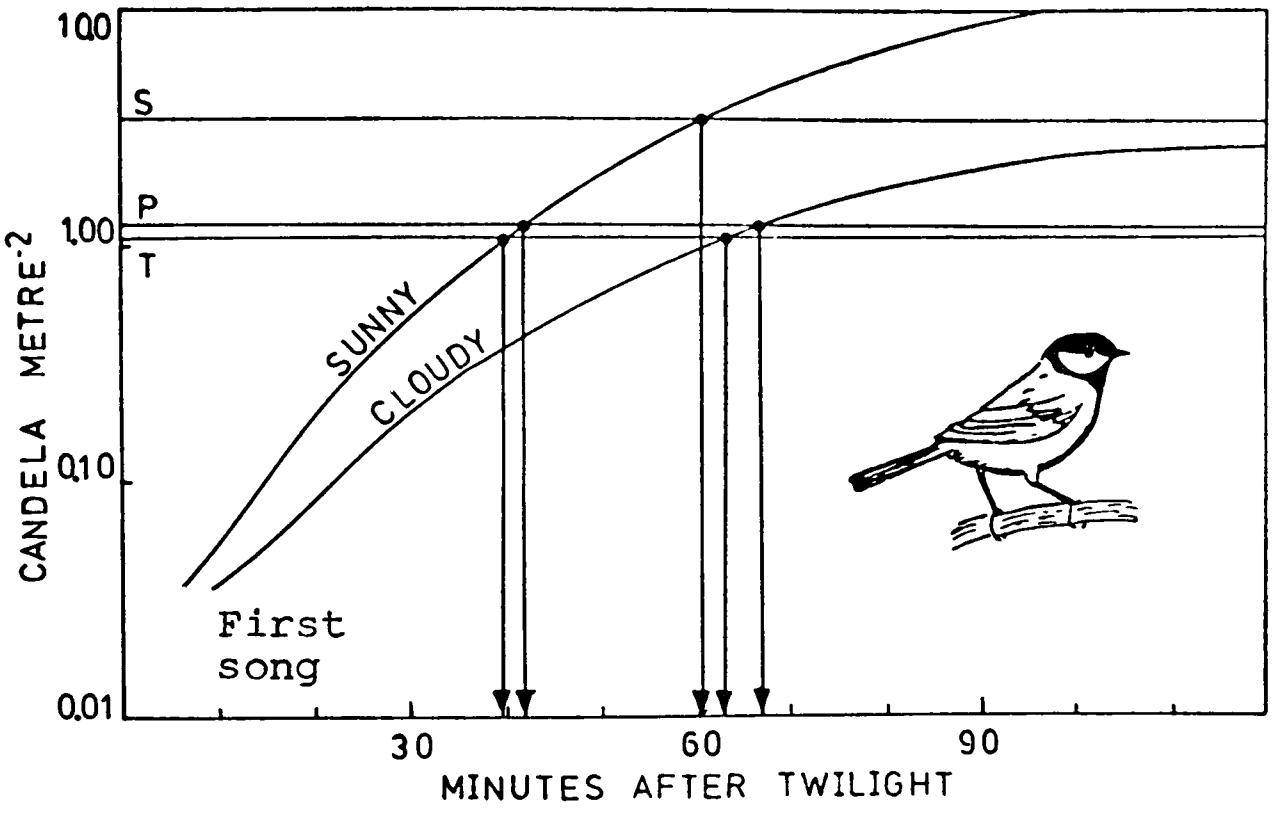
Handling time decreased with increasing luminance. Individual prey items were more thorough exploited when capture rate was low. Means $\pm$ s.e.

Figure 3.6



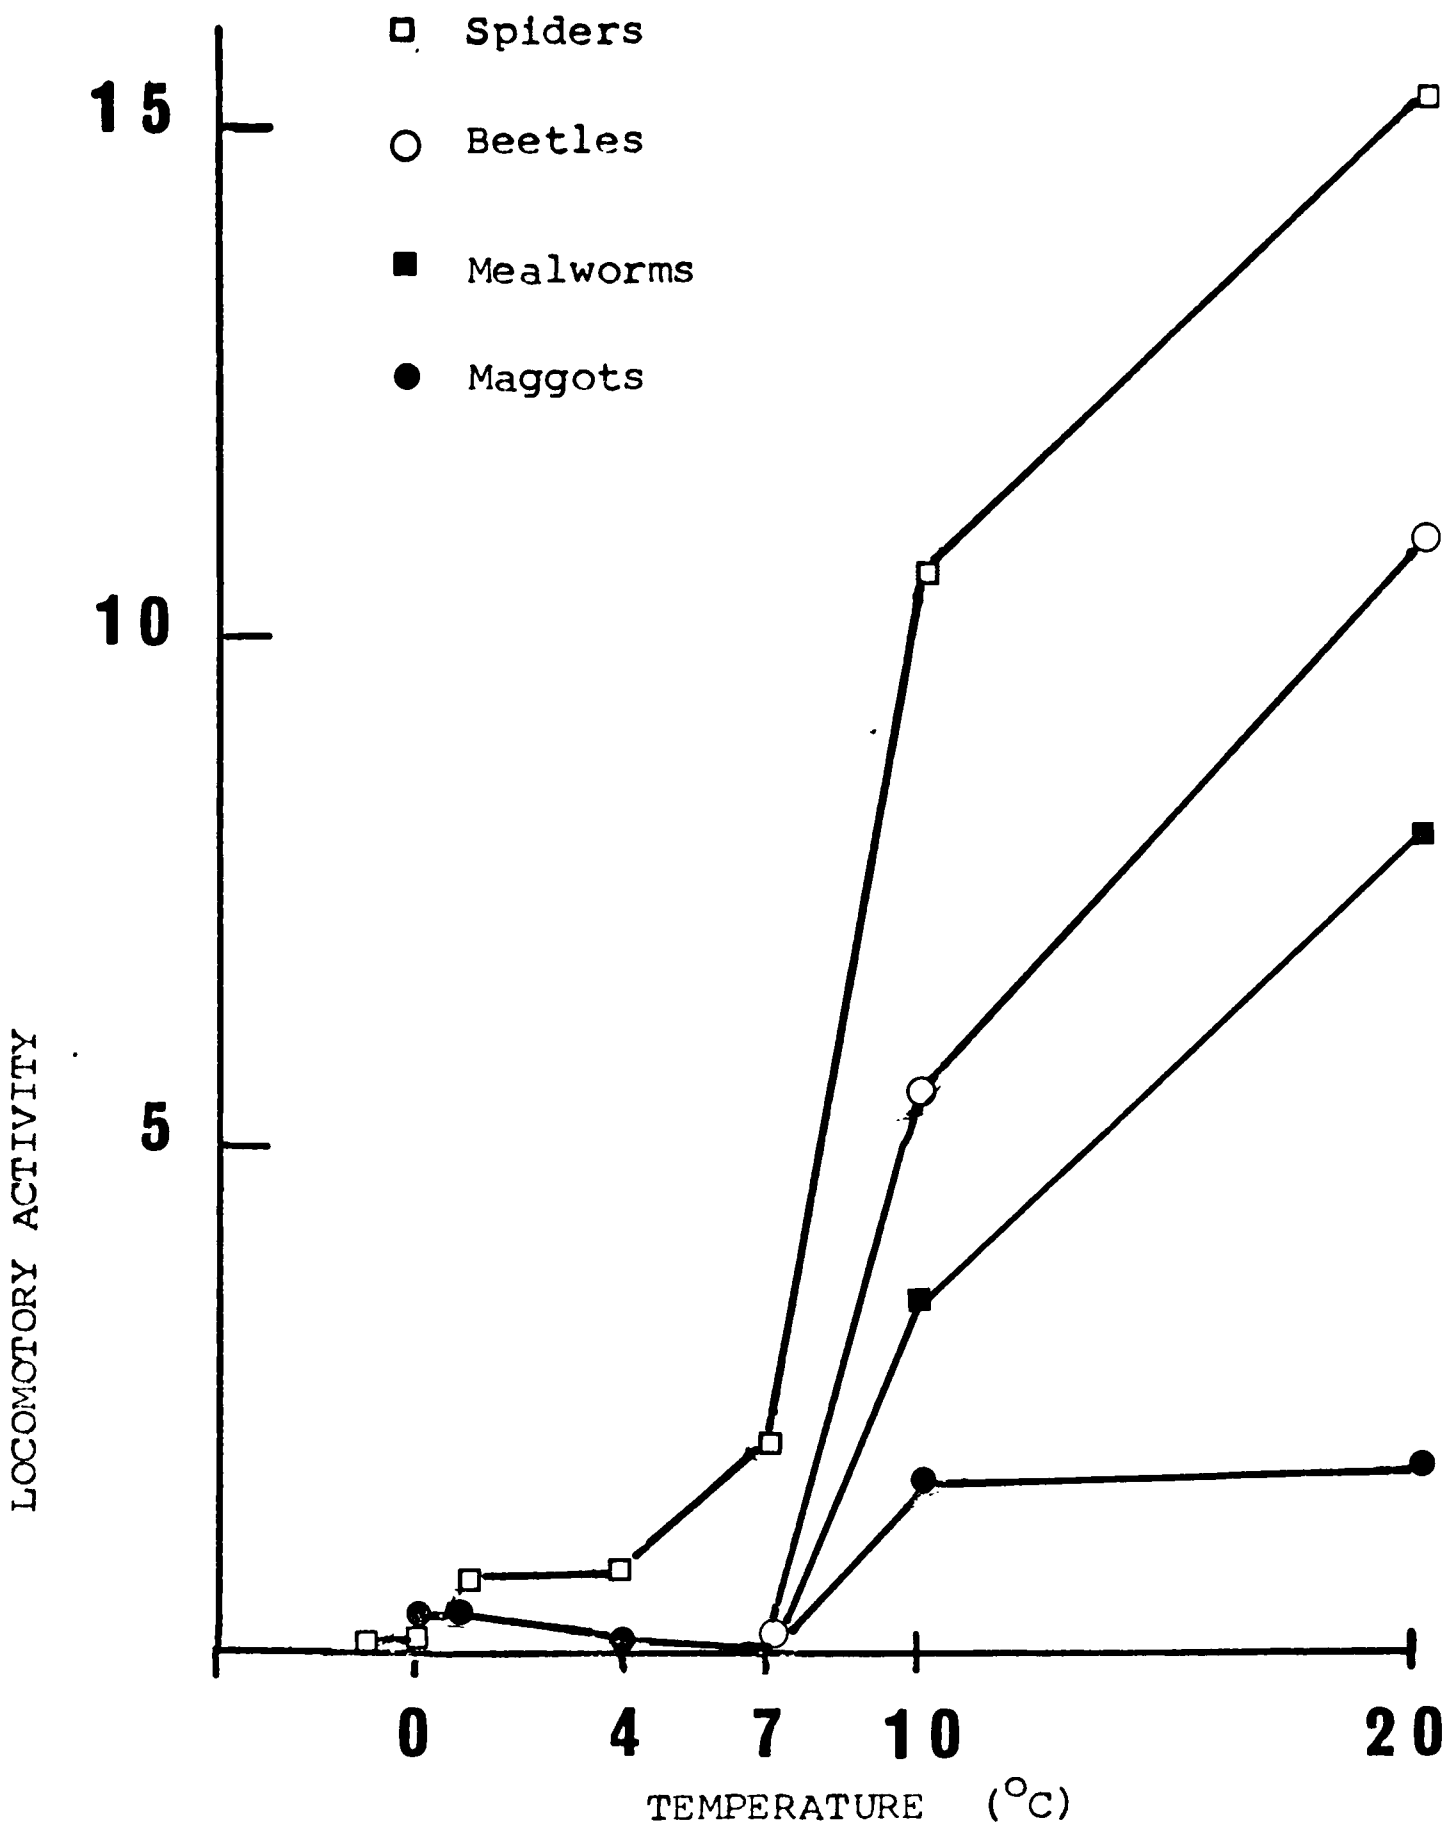
The fraction of time spent foraging increased with luminance up to 100%. The curve is a second order polynomial fitted to the original data, but means $\pm$ s.e. are plotted for clarity. T shows the luminance at which the fitted polynomial reaches 95% of its maximum value.

Figure 3.7



When luminance at great tit's feeding places is plotted against time of the day for sunny and overcast weather, the results from Figures 4 and 6 can be used to estimate at what times light limits foraging efficiency in the wild. S, P and T are from Figure 4 and 6. The time of the first song is approximate from references in the text. Asymptotic searching efficiency (s) is reached between 60 minutes and more than two hours after the first song, depending on the wheather. Asymptotic profitability (P) and full dedication to foraging (T) were reached at similar luminance, corresponding to between 45 and 65 minutes after the first song.

Figure 3.8



Locomotory activity of some potential prey of the great tit as a function of temperature. The ordinate represents the mean number of squares in a grid touched during the 2nd minute (spiders, maggots and beetles) or second to fourth minutes (mealworms) after being released in an open arena.

Figure 3.9 Daily variation in temperature (a) and wind (b) in Oxford. The abscissa is time with respect to civil twilight. The temperature dots are means for the first 14 days of the months of February to April 1979, and the wind dots are means for the whole months during the same period. The vertical bars indicate 12:00 hours. Source: records of the Meteorological Office, Benson, Oxfordshire. The approximate time of the dawn chorus is shown by the shaded zone.

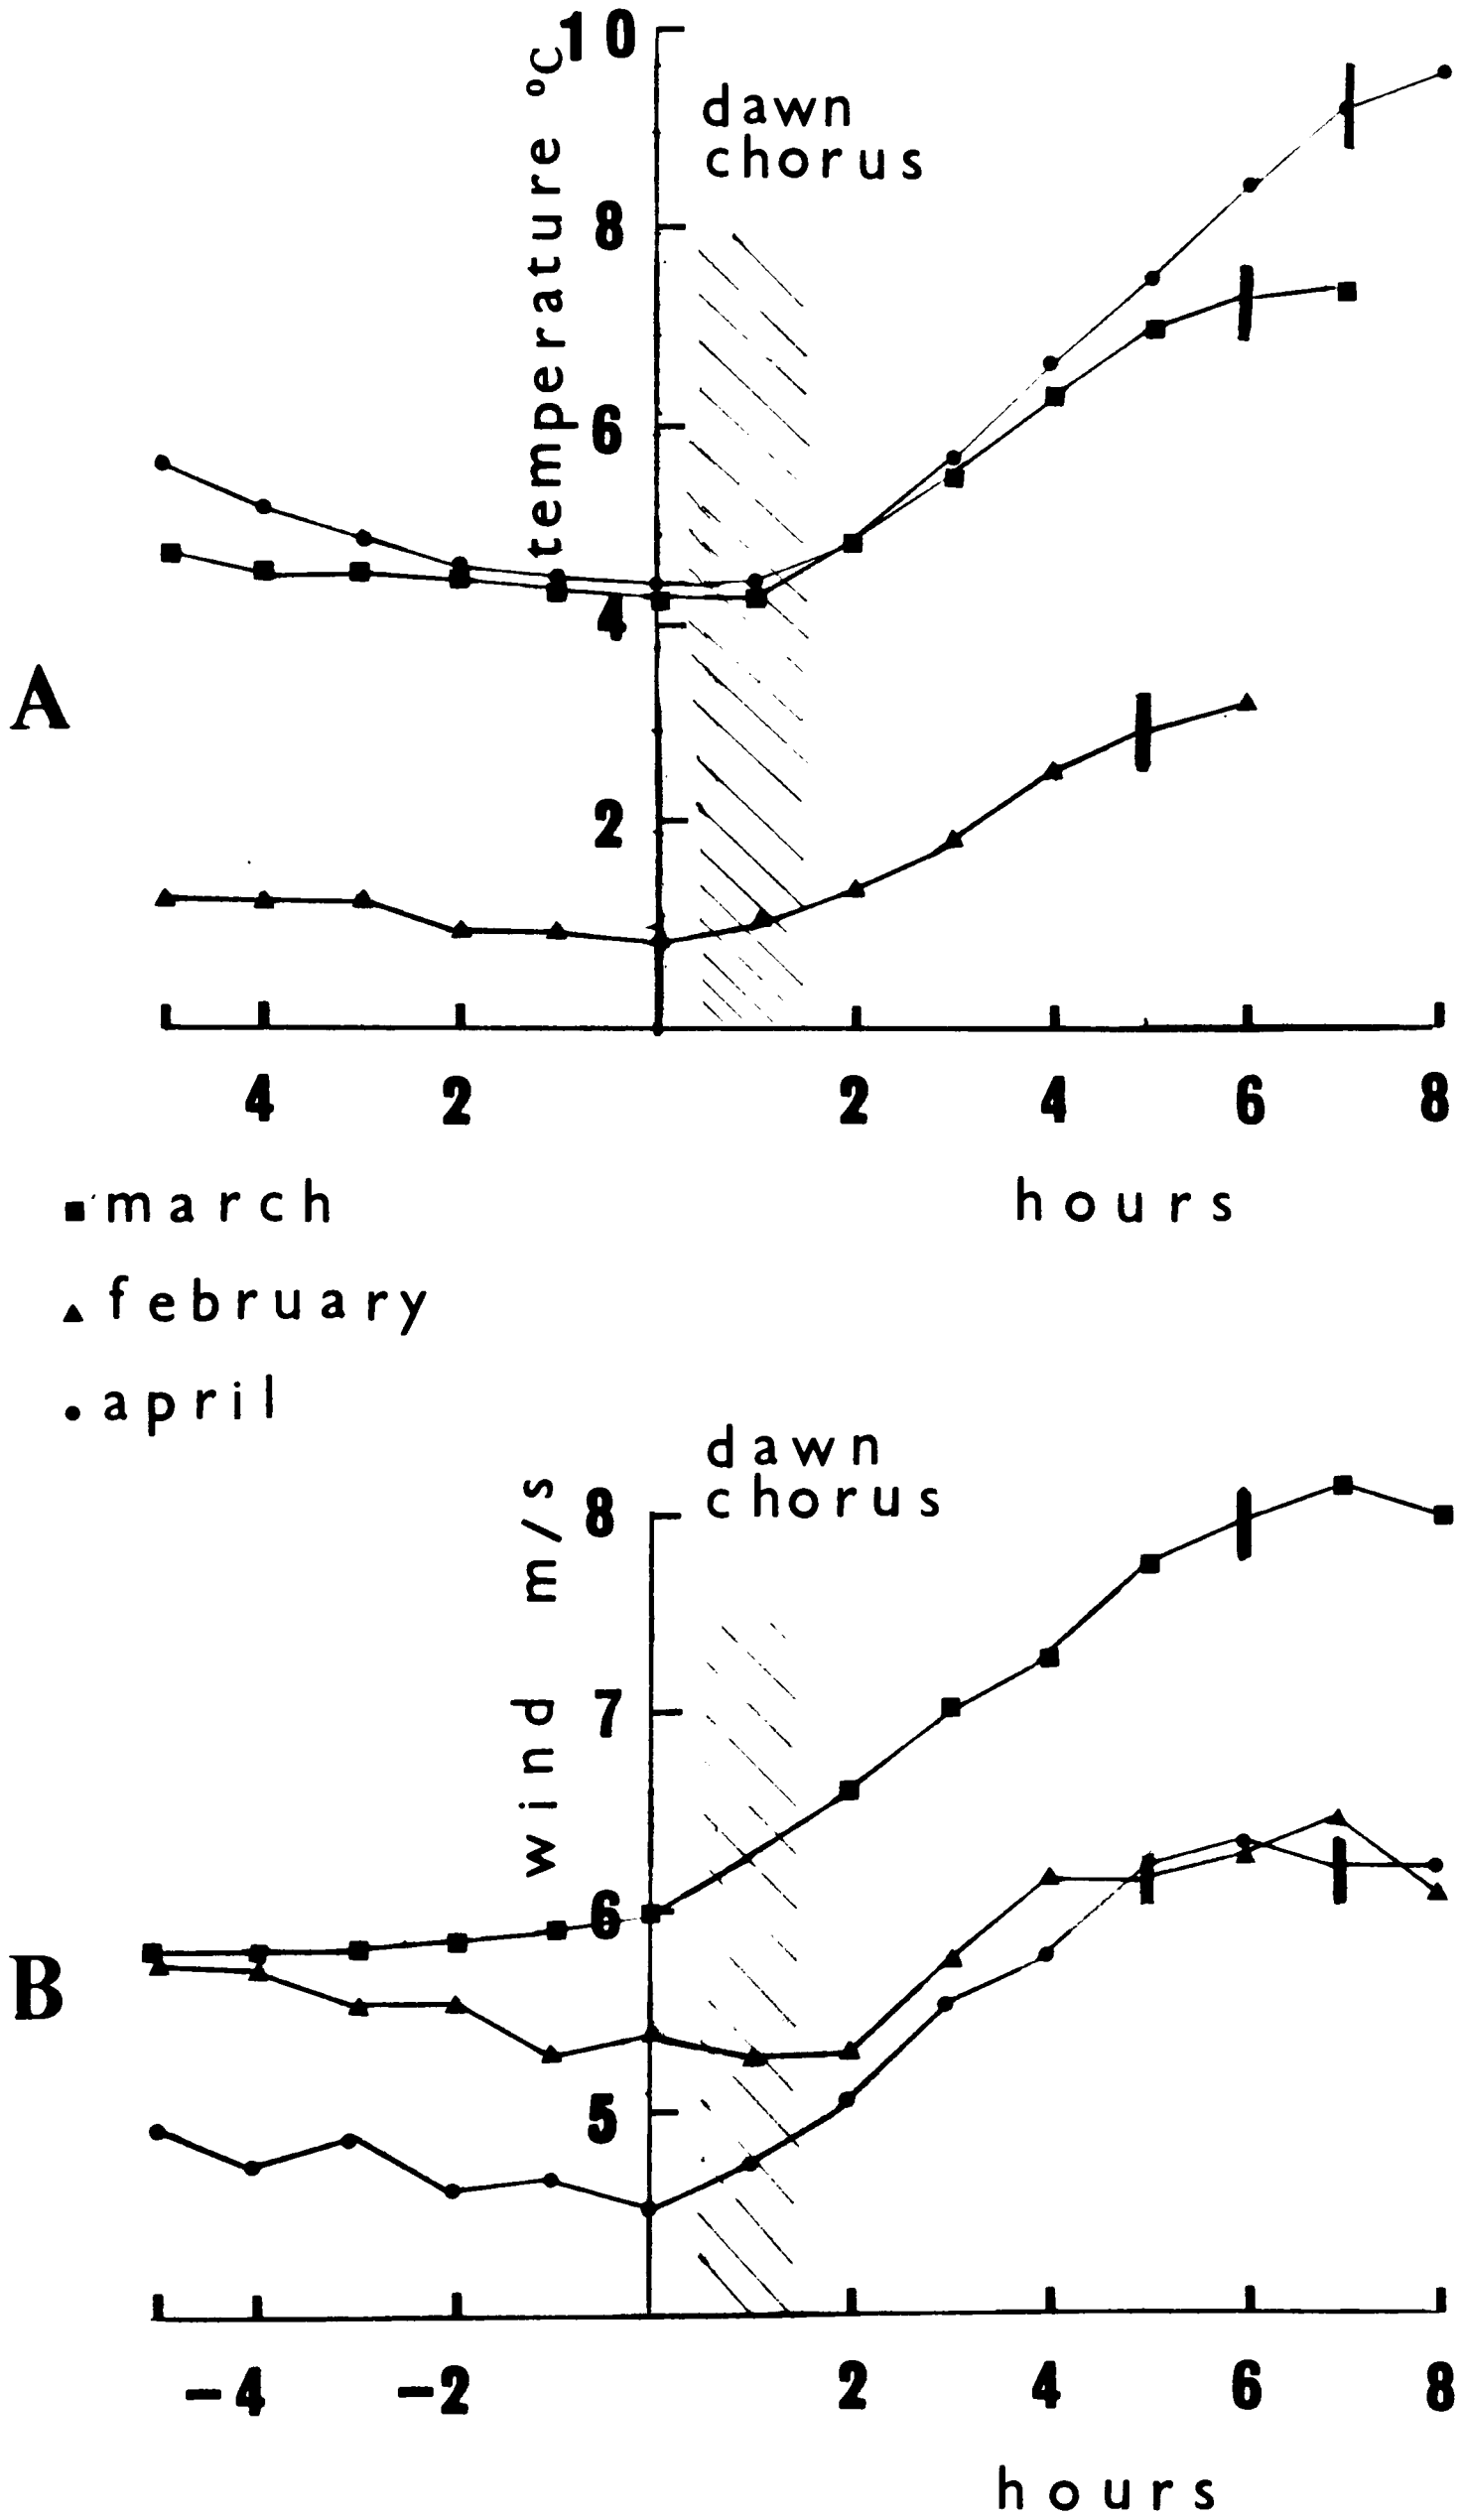
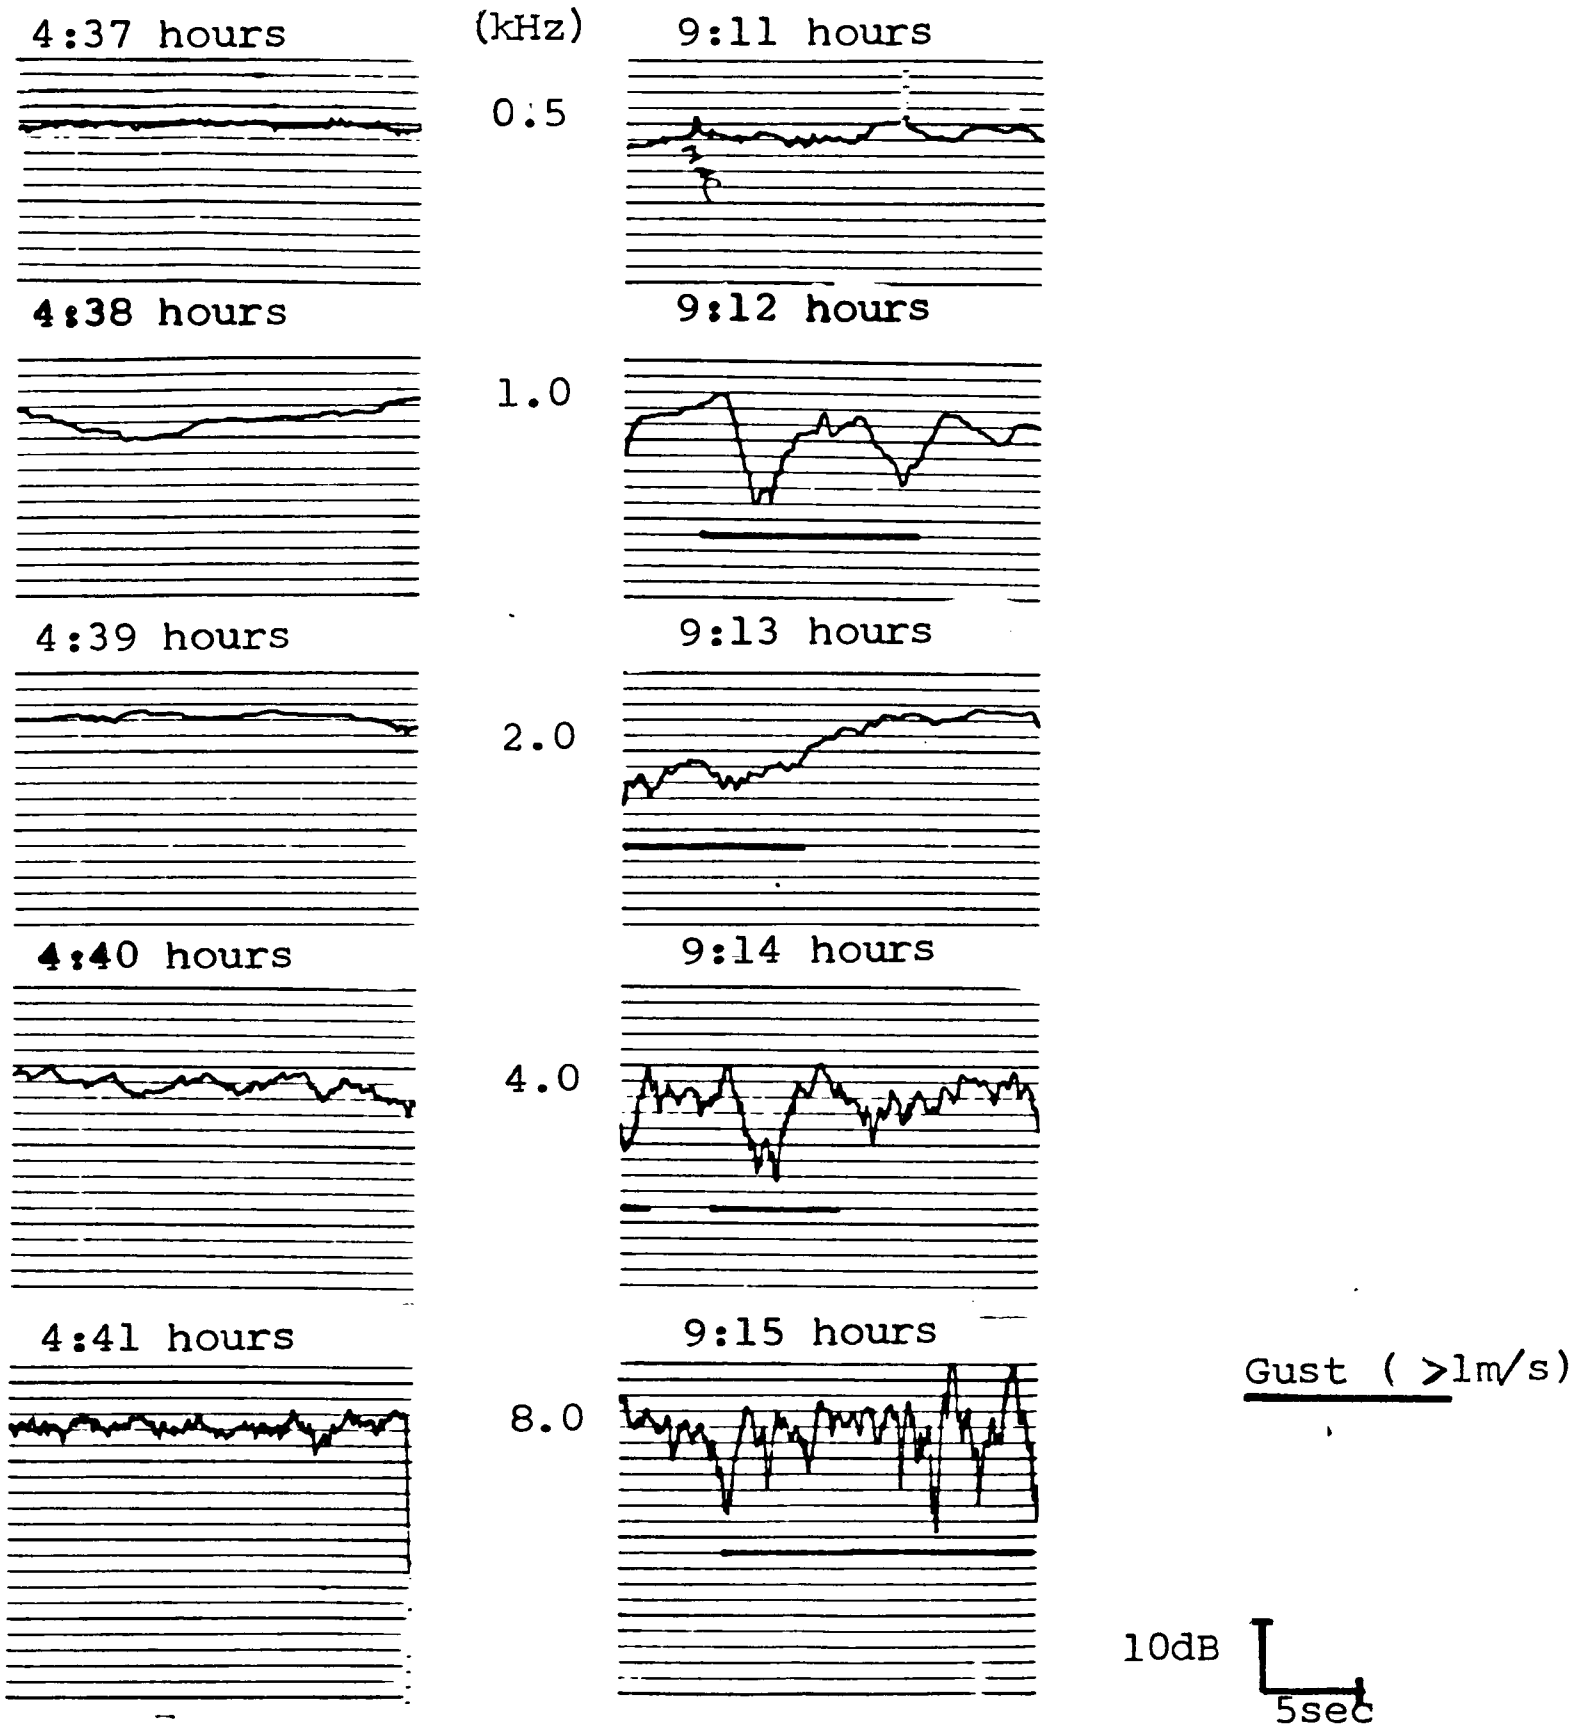

A
temperature °C
10
8
6
4
2
dawn
chorus
4
2
2
4
6
8
hours
■ march
▲ february
• april
B
wind m/s
8
7
6
5
dawn
chorus
−4
−2
2
4
6
8
hours

FIGURE 3•9

Figure 3.10



Effect of wind on acoustic attenuation. The traces are examples of records obtained in Marley Wood on April 8th 1978 (for details see text). Left: records obtained about 10 minutes before civil twilight (04:49 h) in complete calm weather. Right: records four and a half hours after civil twilight, with a breeze of approximately 0.5m/s. The horizontal bars are gusts of stronger wind (c. 2m/s).

CHAPTER 4

Anticipatory control of body weight

4.1. Introduction

This chapter is concerned with the way in which the mechanisms controlling food intake in the great tit might respond to the requirements posed by the dawn chorus. I have mentioned before (Chapter 2) the apparent paradox that in spite of a long overnight period without feeding, great tits sing or preen instead of feeding when they first get up at dawn. In Chapter 3 I discussed possible functional reasons for the dawn chorus, concentrating on light, wind and temperature as possible ecological constraints. In this chapter I will describe an experiment related to the causal mechanisms underlying control of body weight.

If the great tit postpones early morning feeding in the expectation that foraging conditions will improve, the mechanism controlling food intake must involve factors other than immediate food deficit and availability. My discussion is concerned with one aspect of this long-term control, involving anticipation of future conditions. In particular, I wanted to find out whether great tits can anticipate a period of food deficit after waking up in the morning, by increasing their body weight before going to roost. This response, if shown to occur, could be important in allowing the bird to adjust to recently experienced conditions at the time of waking up. For example, in a period of several successive days of intense territorial activity, when the bird had to spend more time than usual in early morning song and ter-

ritorial defence, it could put on extra weight in the afternoon and so decrease its need for early morning foraging. The effect of early morning deprivation is similar to what happens during cold spells. Cold spells and periods of specially high rate of territorial intrusions cannot be accurately predicted from the time of the year (as opposed to what happens with longer term seasonal fluctuations in resource availability (King, 1974; Katz, 1974, Craig et al., 1979)), and because of this one would expect the presence of an adaptive control system, able to adjust certain parameters to recent experience by the addition of further control loops (Milsum, 1966).

Previous workers have shown that small passerines are able to respond to cold spells by becoming heavier (Baldwin and Kendeigh, 1938) and this was corroborated in the great tit (Kluijver, 1952; Owen, 1954; Haftorn, 1952 (quoted in Owen, 1954)). Kendeigh et al. (1969) found that house sparrows Passer domesticus responded to a cold environment by putting on weight when the nights (but not the days) were cold, but not when the days (but not the nights) were cold. This response makes sense in functional terms, since low diurnal temperatures can be compensated for by adjusting the rate of foraging during the day, but cold nights have to be endured with reserves accumulated beforehand if excessive body weight loss is to be avoided.

In this chapter I investigate if the response to morning deprivation is similar to the response to cold nights and present a computer simulation of a possible mechanism underlying the control of body weight.

4.2. Method

I used eight adult wild caught great tits (7♂and 1♀). All the birds had been captured in the vicinity of Oxford more than three months before the experiment started, and some of them had been subjects in the experiment described in the previous chapter. They were accustomed to the laboratory standard diet (see General Methods) which was the only food provided throughout the experiment.

Ten days before the experiment the birds were moved from the holding cages to motels (see Chapter 1) and placed in a room at 13°C and with a daily light cycle of 15D:9L. The birds usually spent their active periods hopping back and forth from the perches to the mesh. While they had enought space to fly, the distances involved were short enough for them to move within the cage by hops rather than by flying. The experiment consisted of three blocks, of three, seven and seven days respectively. During the first and last periods food was available for the whole of the light period. The cages were cleaned every evening just before the lights went off, and food was replaced at 9:00 when the lights came on again. The first block is named hereafter the baseline period, and the last one the reversal control period. The intermediate seven day block differed from the other two in that food was replaced at 11:00 instead of 9:00. This block is named the deprivation period. Water was permanently available. The birds were weighed by difference using the roosting boxes (see General Methods) at 9:00, 11:00, 12:00, 14:00 and 18:00 hours during the

deprivation period. The 11:00 and 14:00 hours weights were not taken during baseline and reversal control periods. The experiment took place during April, 1977.

4.3. Results

Body weight oscillated daily with an amplitude of approximately 1.5 g, with the highest body weight just before the lights were going to be turned off and the lowest in the morning. During the baseline and reversal control periods the lowest body weight was observed just after dawn, and during the deprivation period the lowest point was at the end of the deprivation, which was at 11:00 hours.

Figure 4.1 shows the mean value of the highest and lowest weights for each day of the baseline, deprivation and reversal control periods. In order to emphasize the changes in body weight between periods, a baseline value was obtained for each animal for the morning and evening by taking the mean of the score on the three baseline days. All remaining data were expressed as a ratio to this, so that differences in size between the animals would not affect their weight in the score for a particular day. Figure 4.2 presents the results in this way. Only the fourteen days of deprivation and reversal control periods are shown since the baseline days have an average of one by definition. Lowest body weight outside the deprivation period is identical with morning body weight.

After the start of the deprivation period, evening and morning body weights seemed to take some time to reach new values

but eventually increased by 4% with respect to the baseline. Evening body weight peaked two days earlier than morning weight, producing a temporary increase in the amplitude of the daily fluctuation. The lowest body weight dropped by about 2% with respect to the morning weight during the baseline period, but recovered gradually following the curve of early morning weight but lagging behind it. By the end of the deprivation period lowest body weight was similar to the original weight at dawn.

After reversal to control conditions, evening body weight started to drop gradually, but stayed on average between 1% and 2% above baseline until the end of the experiment. Morning weights instead returned within two days to the original value. On the final experimental day the weight in the morning was exceptionally low. It seems obvious to me that something must have happened that night (possibly someone switched the lights of the room on by accident), but since I was not aware of any special circumstance I left the point as it is. The specially low value does not result from a single individual having a dramatic loss of weight, but instead it is a true drop in the whole sample. The changes in weight at different times of the day after a transition give an interesting indication of the speed at which the birds adjust to food deprivation in the morning. I was also interested in analysing the situation after the transient, in order to establish if the final body weights were significantly different in the deprived versus the control condition. In order to do this, I conducted an analysis of variance eliminating the data obtained in the first

three days after the transition from baseline to deprived and from deprived to reversal control. Table 4.1 shows the results of the ANOVA for evening, morning and lowest body weight of the last four days of deprivation and the last four days of reversal control. Evening body weight and morning body weight were significantly higher during the period of morning deprivation, and lowest body weight during deprivation was not significantly different between conditions. The analysis of the lowest daily weight may have been affected by the last exceptional point, but intuitively it can be seen that the trend in the lowest weight is clearly to increase and that by the end of the deprivation period it was reaching the baseline values in the morning (Figure 4.2).

Figure 43 is a plot of the mean weights at various times of day excluding the first three days after a transition. Each point in this figure results from the average of 32 data points: four days and eight birds. The aim of this figure is to describe in detail the rate of change of body weight during the course of the day. The main features are: (1) At dawn the weight of the deprived animal is higher than those with free access to food. (2) Two hours later the deprived birds reached their lowest value in the day, which was only slightly lower than the morning weight of the control period. (3) By the evening the situation is reversed, and the animals under deprivation are heavier than when food is continuously supplied.

It is remarkable that the rate of change in body weight after the deprived birds have had access to food for a period of 3 to 5

hours after dawn is almost identical with the rate of change of the non-deprived birds during the period between dawn and three hours later. I think that this probably indicates that this rate of change in weight is the fastest that the birds can sustain once that they have filled their crops, or in other words, because of the unlimited food availability in captivity the rate of intake is limited by a digestive bottleneck rather than by external factors. This is similar to the situation of wild pigeons feeding on a superabundant but poor quality food (Kenward ans Sibly, 1977).

4.4. Discussion

The results described so far call for some comments on the possible motivational mechanisms involved in the control of food intake. Feeding motivation has often been discussed in terms of negative feedback models. While all proposed models agree on the basic property that feeding in the long term decreases rate of further feeding, the actual structure of different models differs widely (Geertsema and Reddingius, 1974; Toates and Booth, 1974; Wirthshafter and Davis, 1977). Sibly and McCleery (1975), Sibly (1975), and McCleery (1977) have reformulated the long standing rule that feeding tendency is proportional to deficit (Cicala, 1961; Hinde, 1970; Bolles, 1962) to incorporate the interaction with other drives, but they continue to assume that feeding is driven by immediate factors.

Deficit is normally assumed to be the difference between the current state of reserves (or energy supply in Toates and Booth's model) and some ideal (optimum) state of satiation. In the wild this ideal state cannot be measured independently of the

animal's behaviour and in consequence the definition has a tautological element in it. For certain purposes, particularly for laboratory experiments in simple environments, satiation state can be operationally defined as the state an animal finds itself in after prolonged access to unlimited food, or as that condition beyond which feeding rate equals rate of energy utilisation. In practice authors working with operant paradigms tend to use a certain arbitrary interval between feeding responses as a satiation criterion. Since my experimental animals (and for this matter, great tits in the wild (Owen, 1954)), do not seem to have a period of the day when food intake equals energy consumption (body weight increases through the whole day (Figure 4.3)) all one can say is that if satiation occurs, it can only be the state reached at the end of the day, when body weight is highest. Figure 4.3 also indicates that this value is not the maximum possible body weight under the experimental conditions, since the deprived animals reach a higher evening body weight even after a lower start. The dotted lines in Figure 4.3 point to my estimate of the maximum attainable body weight, obtained by extrapolation of the fastest rate of increase of body weight that the animals can maintain for a long period.

Given that evening body weight is the maximum value reached during the daily cycle and that it does not result from an external or physiological constraint but from the birds behaviour it is tempting to model the motivational control system for feeding as a negative feedback system where evening body weight is the set

point (or at least the 'settling point' (Wirtshafter and Davis, 1977))-

Figure 4.4 shows a diagram of such a model. This particular model is 'hardware independent' (McFarland, 1971) in that it does not require the set point or any of the other variables to be located in particular anatomical structures. There is also a vast physiological literature on neuroanatomical (or 'hardware dependent') models, see for example Keesey et al. (1970), Nisbett (1972), Panksepp (1974), and for a critique, Wirtshafter and Davis (1977). My model has four parameters; Incentive, Maximum rate of intake, Standard (or 'resting') rate of consumption, and Active rate of consumption.

Incentive is only included to show where would one expect that the model would be affected by food palatability, feeding opportunity or other environmental variations. Since in the experiment food was always the same and its availability was all-or-none this parameter was unity in the simulations. The remaining parameters were estimated from the data as follows.

Maximum rate of intake was computed from the rate of body weight gain immediately after the end of the deprivation plus the rate of body weight drop when food is not available during daytime. This latter value was obtained from the morning deprivation period. The rate of consumption during rest was obtained by subtracting evening and morning body weights, assuming for simplicity that overnight weight drops linearly with time. This assumption could prove to be inappropriate if the model is used for very different

lengths of the night period. Since the parameters are obtained from the data, the model is a hypothesis about their relationship. What I want to test is if this type of model, taking these parameters as constraints, would mimic the response of the animals to the experimental treatment.

The model was implemented as an Algol 60 program and run on a PDP-8 computer modifying the simulation techniques of Toates (1974). The integrations were done by iterating the feedback loop of the system ten times every time unit (equivalent to one hour of real time). Figures 4.5 and 4.6 show the result of this simulation. The response of the model to the morning deprivation was qualitatively different from the birds' response. There was no compensation for the extra drop in morning body weight other than by increasing the length of the period of high intake rate until original evening body weight was reached. This result is not surprising, it is just a way to show how a system driven exclusively by instantaneous deficit from an ideal state of zero deprivation is unable to account for the available data. A very simple modification (Figure 4.7) produces an adaptive control system, that is one that is able to modify its structure or the value of certain parameters according to experience. In the model described in Figure 4.7, an additional feedback loop is added so that the model corrects its set point using a learning rule aimed at coping with long term environmental changes. From my own data and those of Kendeigh et al. (1969) it seems that the birds try to avoid reaching extreme low weights without becoming too heavy when it is not necessary. Since the birds can only use available

information to predict future events, I tested the following learning rule: Assuming that the model has a 'natural' or 'relaxed' value of lowest body weight and this equals the value during baseline conditions, the model computes a correcting term that is added to its natural set point. I tested a correcting term based on a weighted average of the deviations from the natural lowest weight in the previous three days. The weighing consists of 'devaluating' the importance of deviations as time goes by. The correcting term is, then:

$$CT = NLW - (0.6 \times LW(d-1) + 0.3 \times LW(d-2) + 0.1 \times LW(d-3))$$

CT = correcting term

NLW = 'natural' lowest body weight

LWd = Lowest body weight on day d.

Figures 4.8 and 4.9 show the output of the model when treated in the same way as the experimental animals. The results are in good qualitative agreement with the data. In Figure 4.2 the dotted line shows another expression of the model's behaviour, and Table 4.1 compares the data with both models. This model does not discriminate between the cause of the excessive loss of weight being due to lower overnight temperature or to a time constraint such as the need to defend the territory. In consequence it is a very economic way of explaining both my results and those of Kendeigh et al. (1969). Kendeigh et al's observation that diurnal temperature was not effective in modifying evening body weight but night temperature was, lead these authors to propose that the stimulus for the change in weight was night temperature. This model shows an alternative hypothesis. Diurnal

low temperatures do not normally affect the set point because the model can compensate them through the short-term feedback loop, by increasing the period of feeding at the maximum rate of intake. Only when the disturbance cannot be compensated through the short-term feedback loop does the next loop up in the hierarchy take action. Many different learning rules based on deviations from preferred daily lowest body weight would produce similar adaptive behaviour. Nevertheless, different learning rules may make very different predictions which could be tested experimentally. The particular rule used above oscillates with a characteristic period and amplitude, and this can be seen by running the model for longer periods than the seven day blocks used in this experiment. These oscillations as well as the speed of adaptation to changes of different intensities than the ones used here can be used to test specific features of the model.

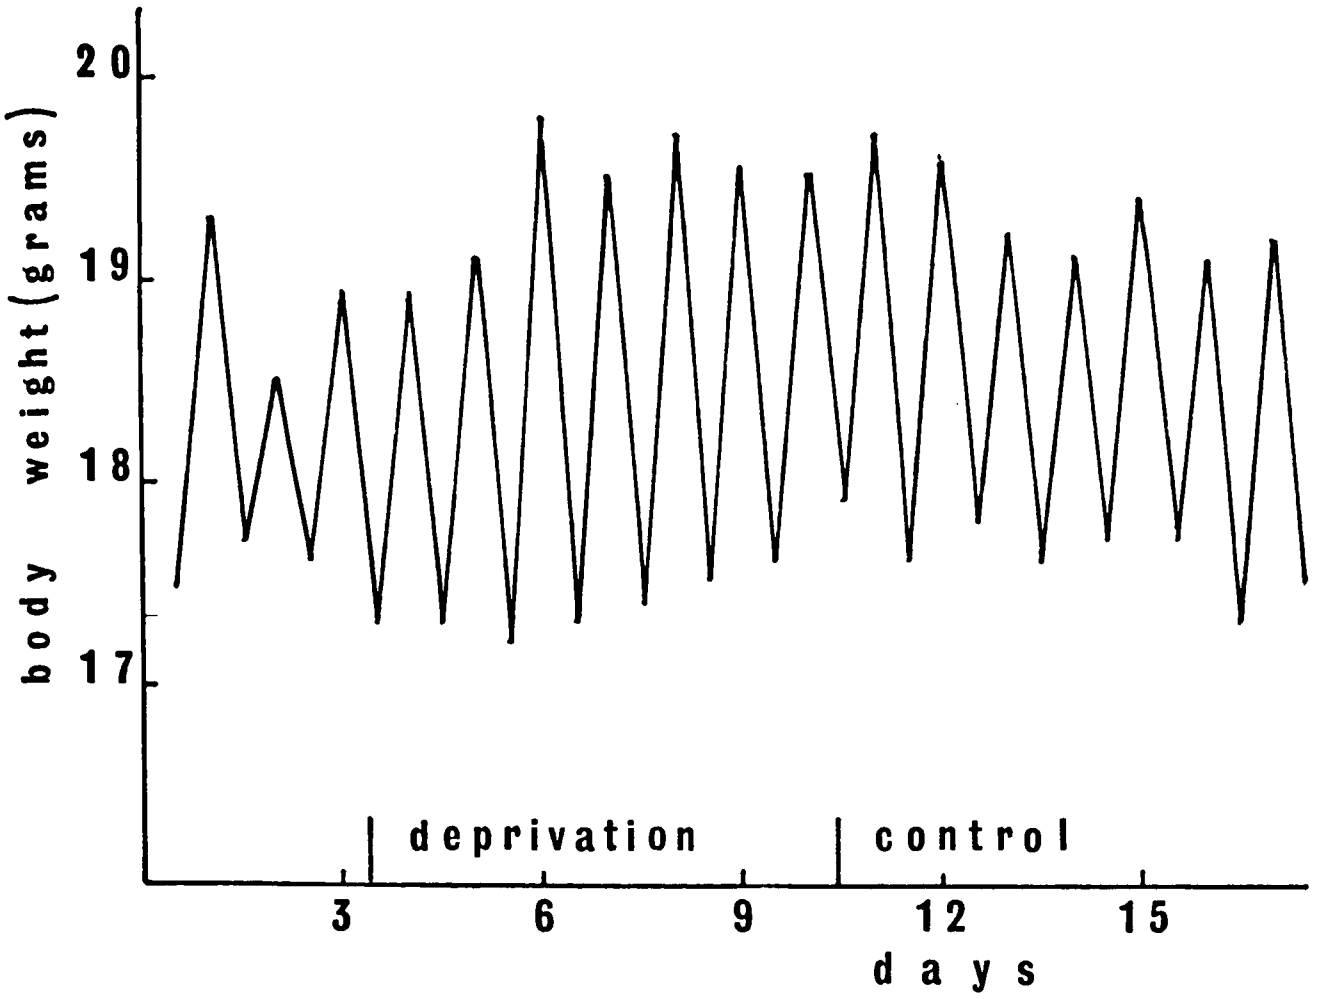
In conclusion, great tits are able to compensate for changes in the daily pattern of food availability by overfeeding in anticipation, and a simple adaptive regulatory mechanism with two hierarchically organised feedback loops can account for this ability.

Table 4.1

		Deprivation	Reversal Control	DIFF	F	p
Dawn Body Weight (g)	Great tits	17.79	17.29 *	+0.52	32.4	0.001
	Fixed set point	16.70	17.30	-0.60	-	-
	Adjustable set point	17.74	17.24	+0.50	-	-
Lowest Body Weight (g)	Great tits	17.15	17.29	-0.14	2.7	Ns
	Fixed set point	16.70	17.30	-0.50	-	-
	Adjustable set point	17.14	17.24	-0.10	-	-
Evening Body Weight (g)	Great tits	19.25	18.90 *	+0.35	12.0	0.01
	Fixed set point	18.80	18.80	0	-	-
	Adjustable set point	19.20	18.70	+0.50	-	-

ANOVA of experimental results and results of computer simulations
Deprivation and Control refer to the last 4 of the 7 day blocks when food was not available until 2 hours after dawn (deprivation) and the last 4 of the 7 day block when ad libitum access to food was restored (control).
Cells marked with an asterisk point at trivial identities resulting from the chosen set point in the simulations.
The column of differences between Deprivation and Control periods result from the dynamics of the models, and are the most relevant comparisons.

Figure 4.1



Mean body weight throughout the whole experiment. The peaks correspond to evening body weight and the troughs to the lowest daily weight.
During the period labelled deprivation (days 4 to 10) food was not presented until 2 hours after dawn. During the initial and control period food was provided ad libitum from dawn to dusk.

Figure 4.2: A different presentation of the data shown in Figure 1. For each animal I computed the evening and morning body weight during the initial period (days 1-3), and this was used as a reference for Morning, Lowest and Evening body weight for the remaining 14 days. Each individual weight was expressed as a fraction of that bird's reference value and all the scores were then averaged. Only deprivation and control days were plotted (days 4-17).
Solid lines: experimental data
Broken lines: simulation results (see text)

Figure 4.2

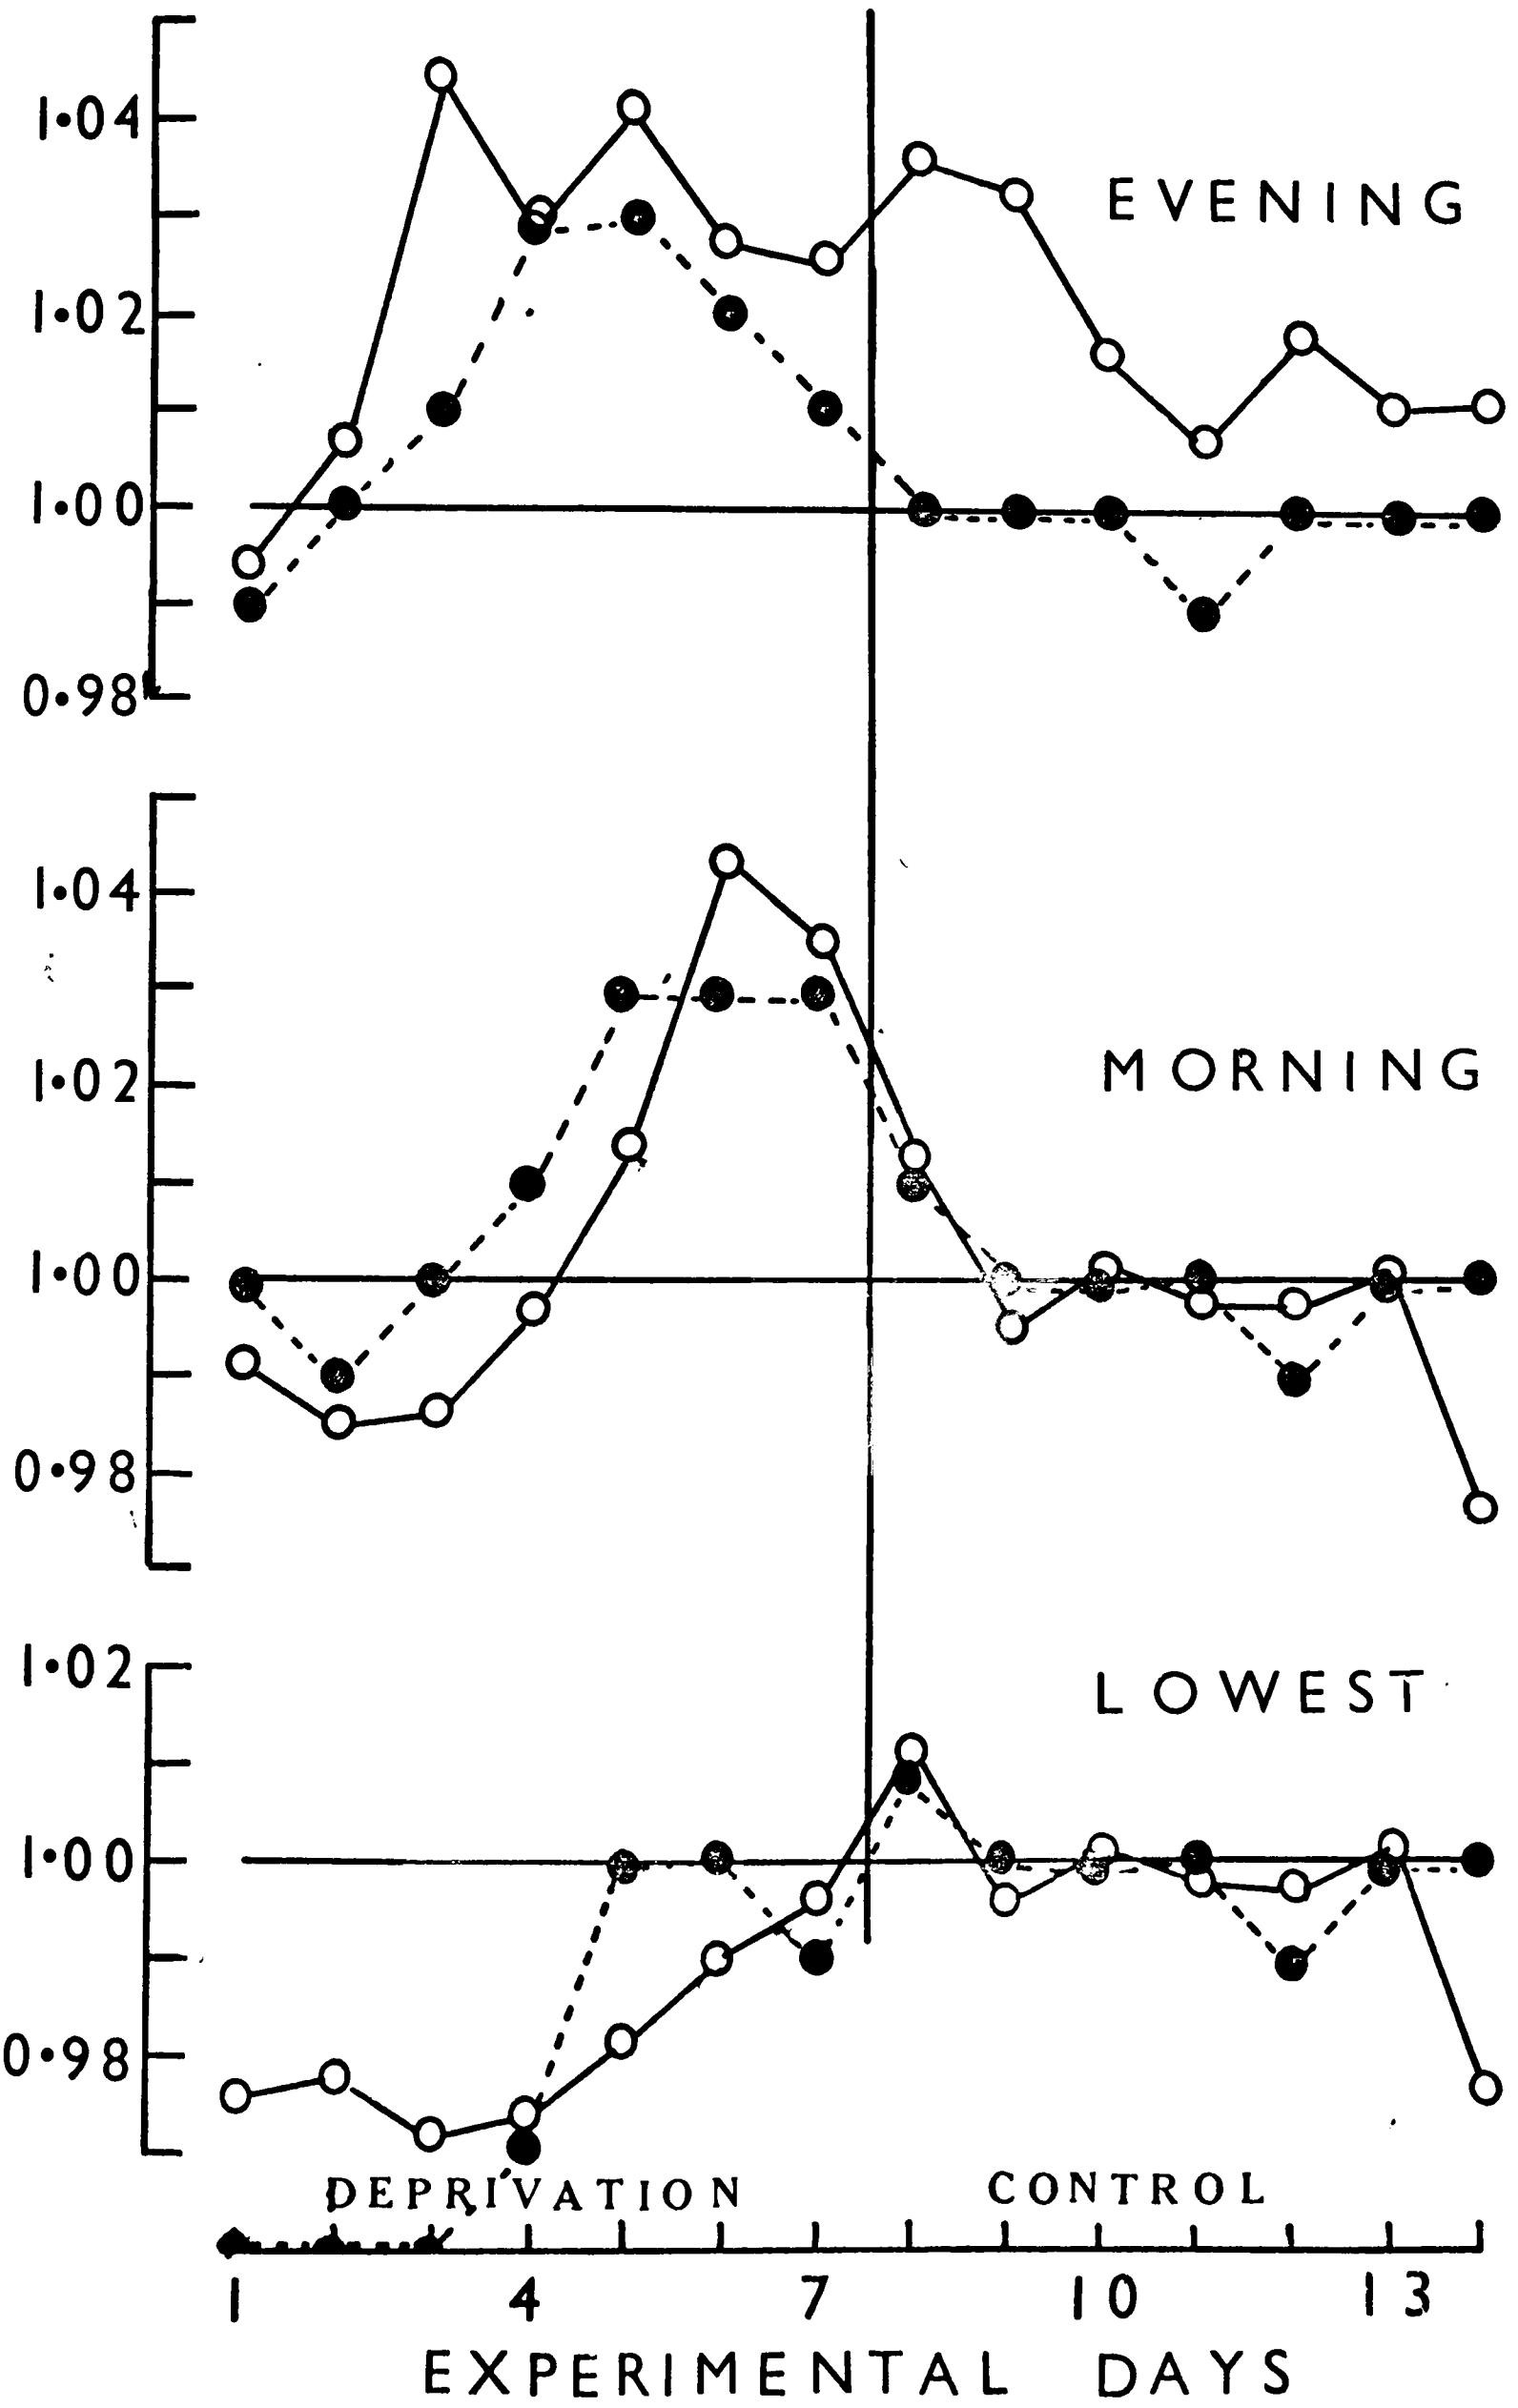
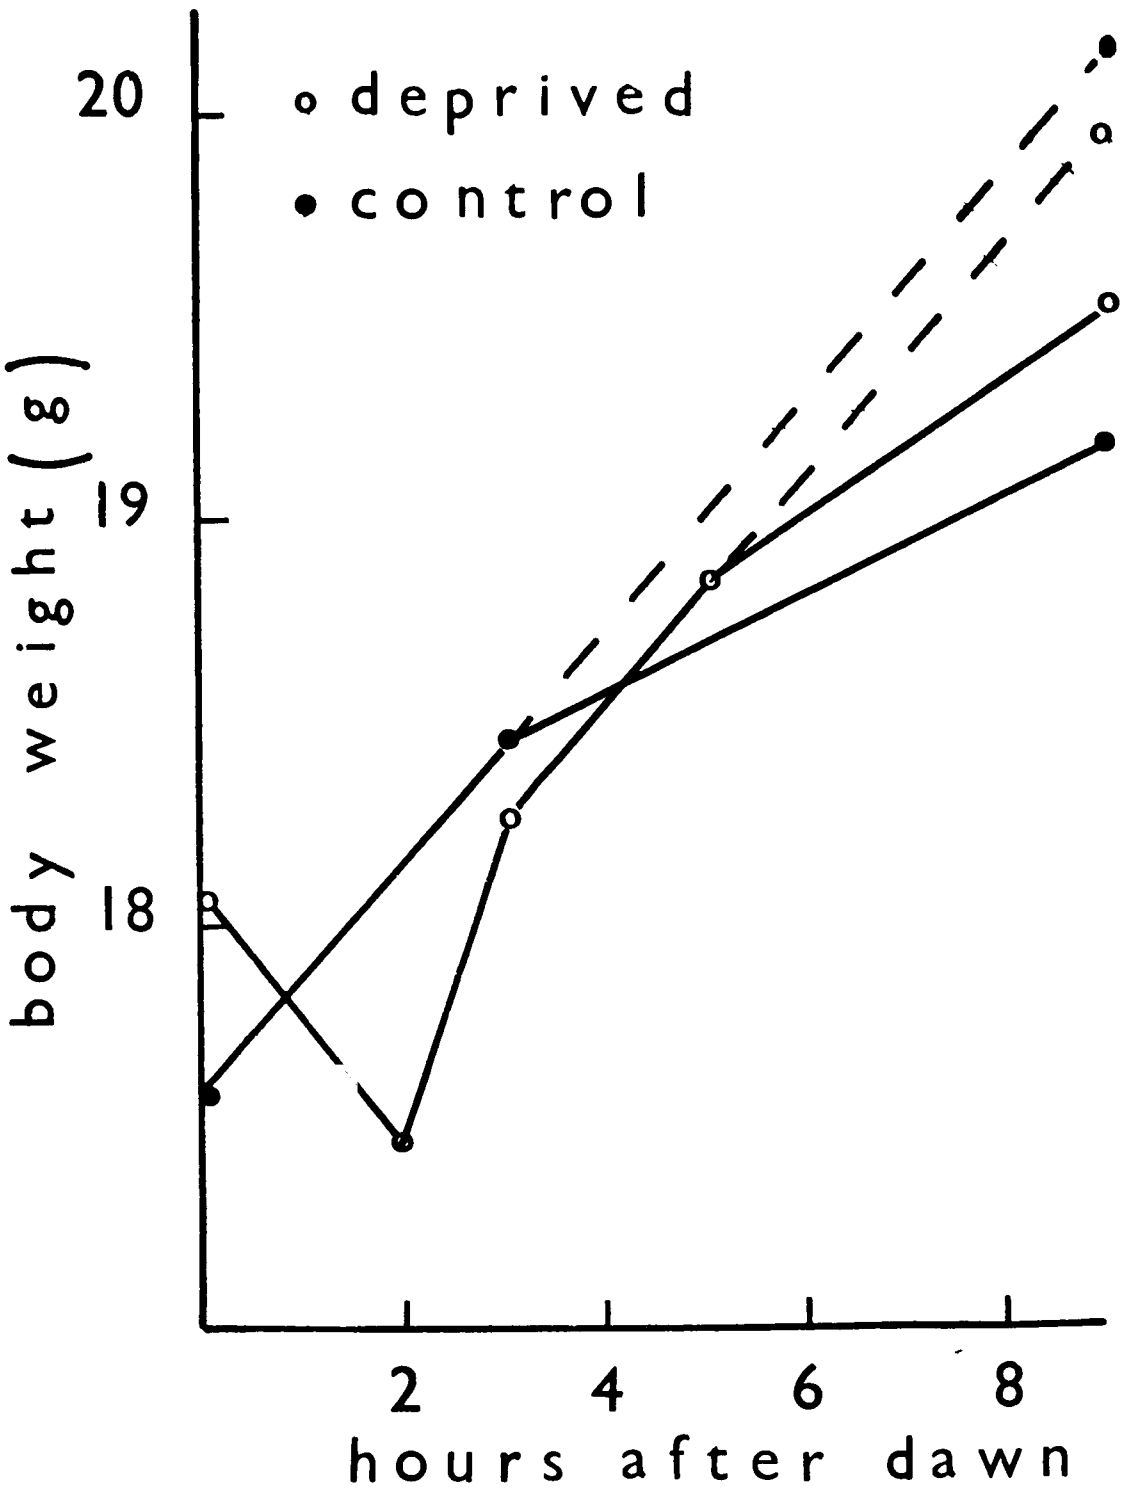

1·04
1·02
1·00
0·98
EVENING
1·04
1·02
1·00
0·98
MORNING
1·02
1·00
0·98
LOWEST
DEPRIVATION
CONTROL
1
4
7
10
13
EXPERIMENTAL DAYS

Figure 4.3



Body weight at different times of day. The values are means of all individuals over the last 4 days of the deprivation and reversal control period.
The dotted lines indicate, by extrapolation, the maximum attainable evening body weight under control (●) and morning deprivation (o) conditions.

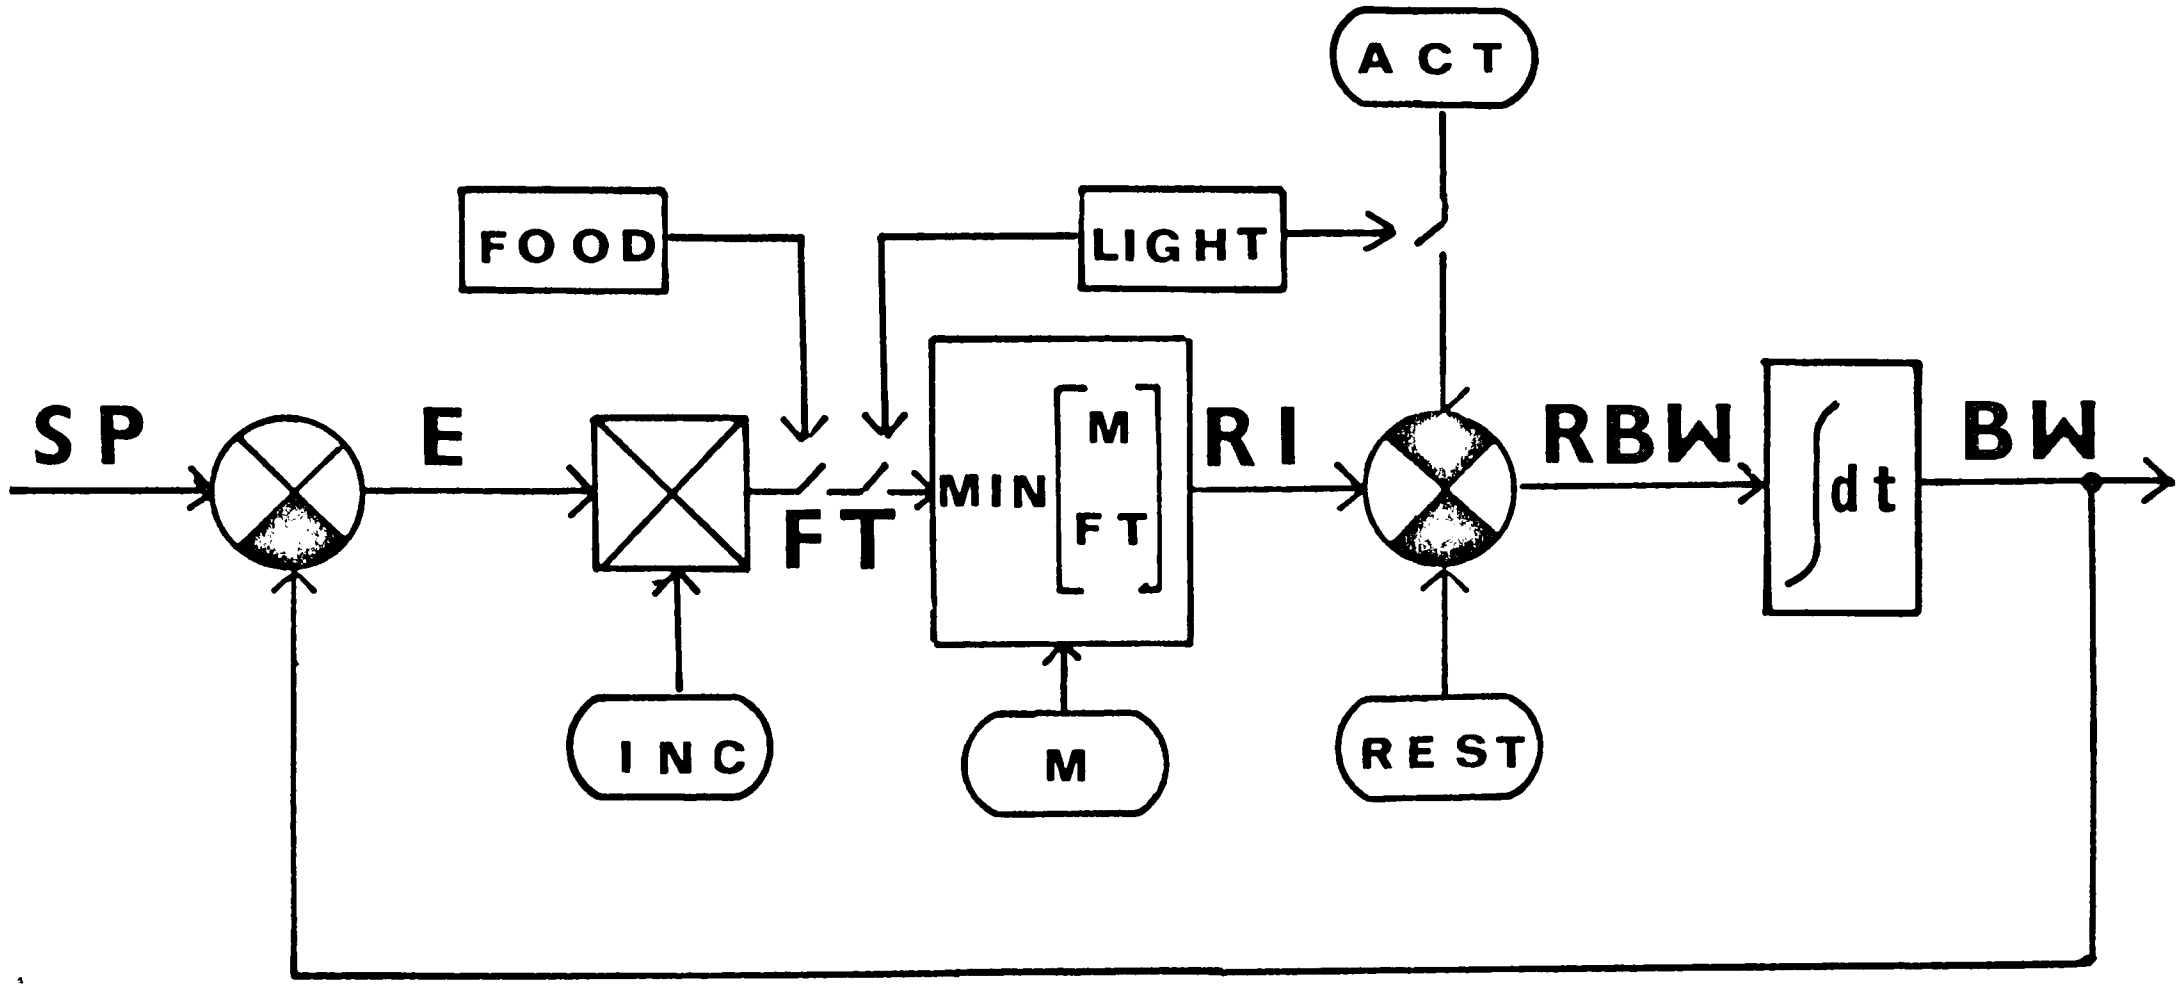
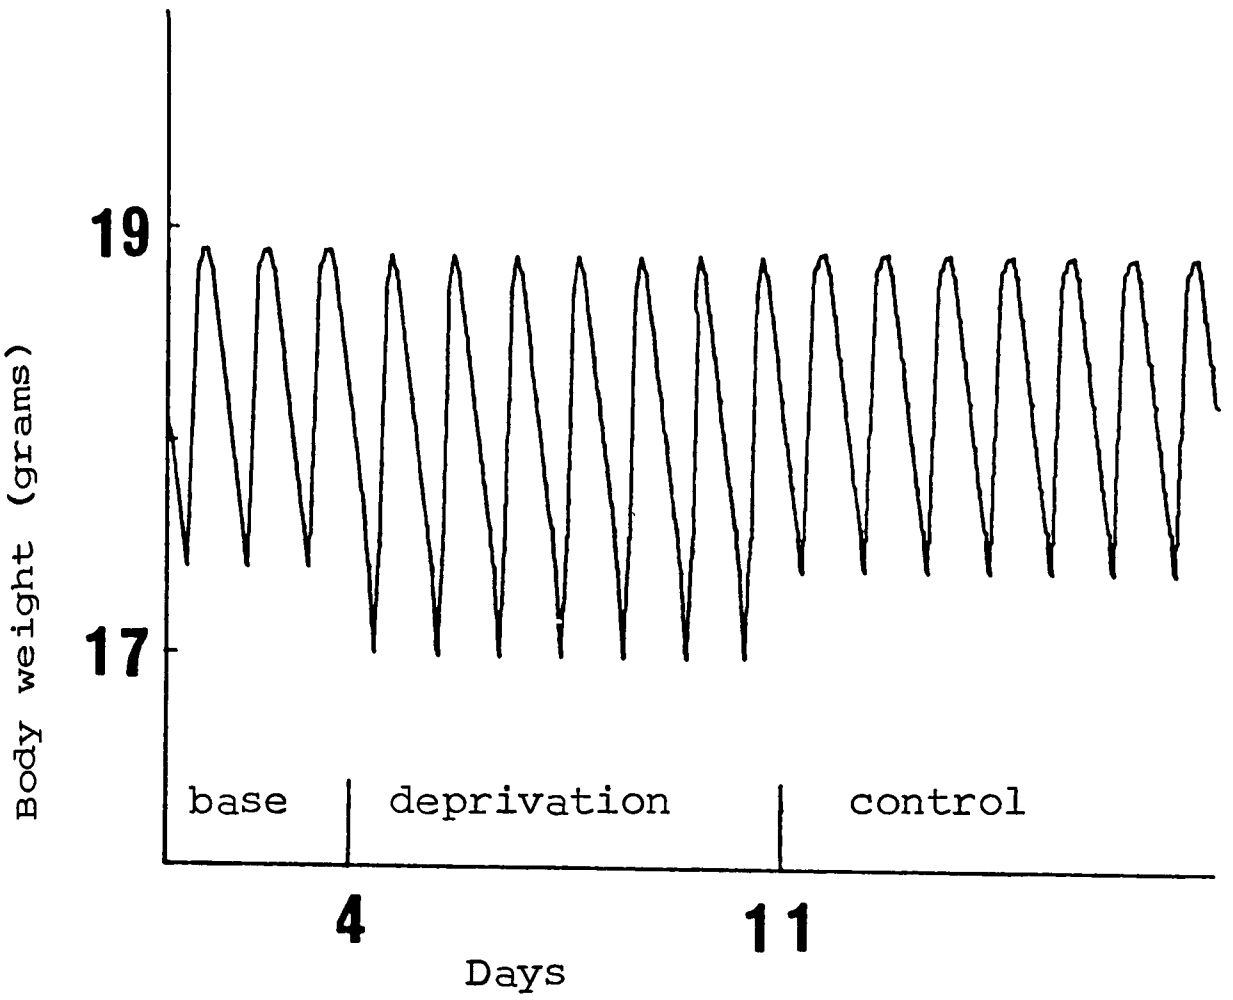


Figure 4.4

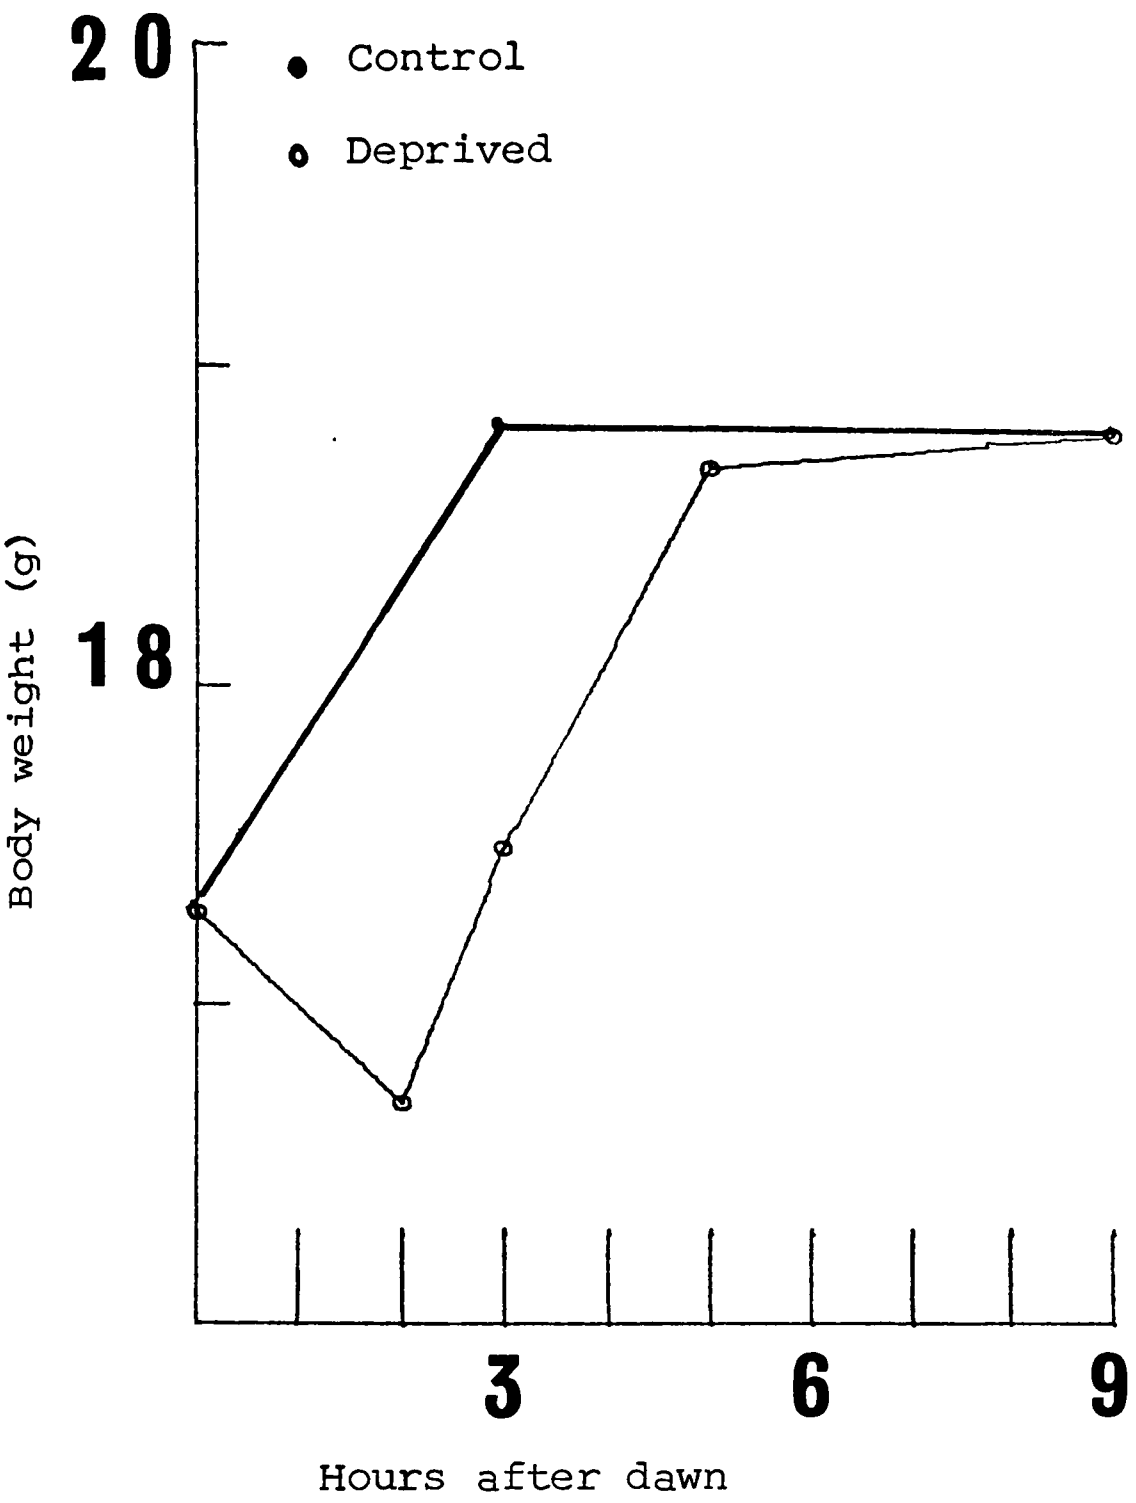
Diagram of an homeostatic model of body weight regulation (SIMWGT 1). Variables: SP, set point; E, error; FT, feeding tendency; RI, rate of food intake; RBW, rate of change in body weight; BW, body weight. Parameters: Act, the increase in the rate of BW loss due to the animal being active; Incentive, a measure of the food palatability or "strength of the stimulus"; M, maximum possible rate of intake due to anatomical constraints. Boxes represent operations and Arrows are variables. The simbols utilised are: ⊗ sum (black means change of sign of input); ⊠ product; MIN[M, FT] minimum (output = lowest of the two inputs); ∫ integral; ↓ switch; (REST) parameter.

Figure 4.5



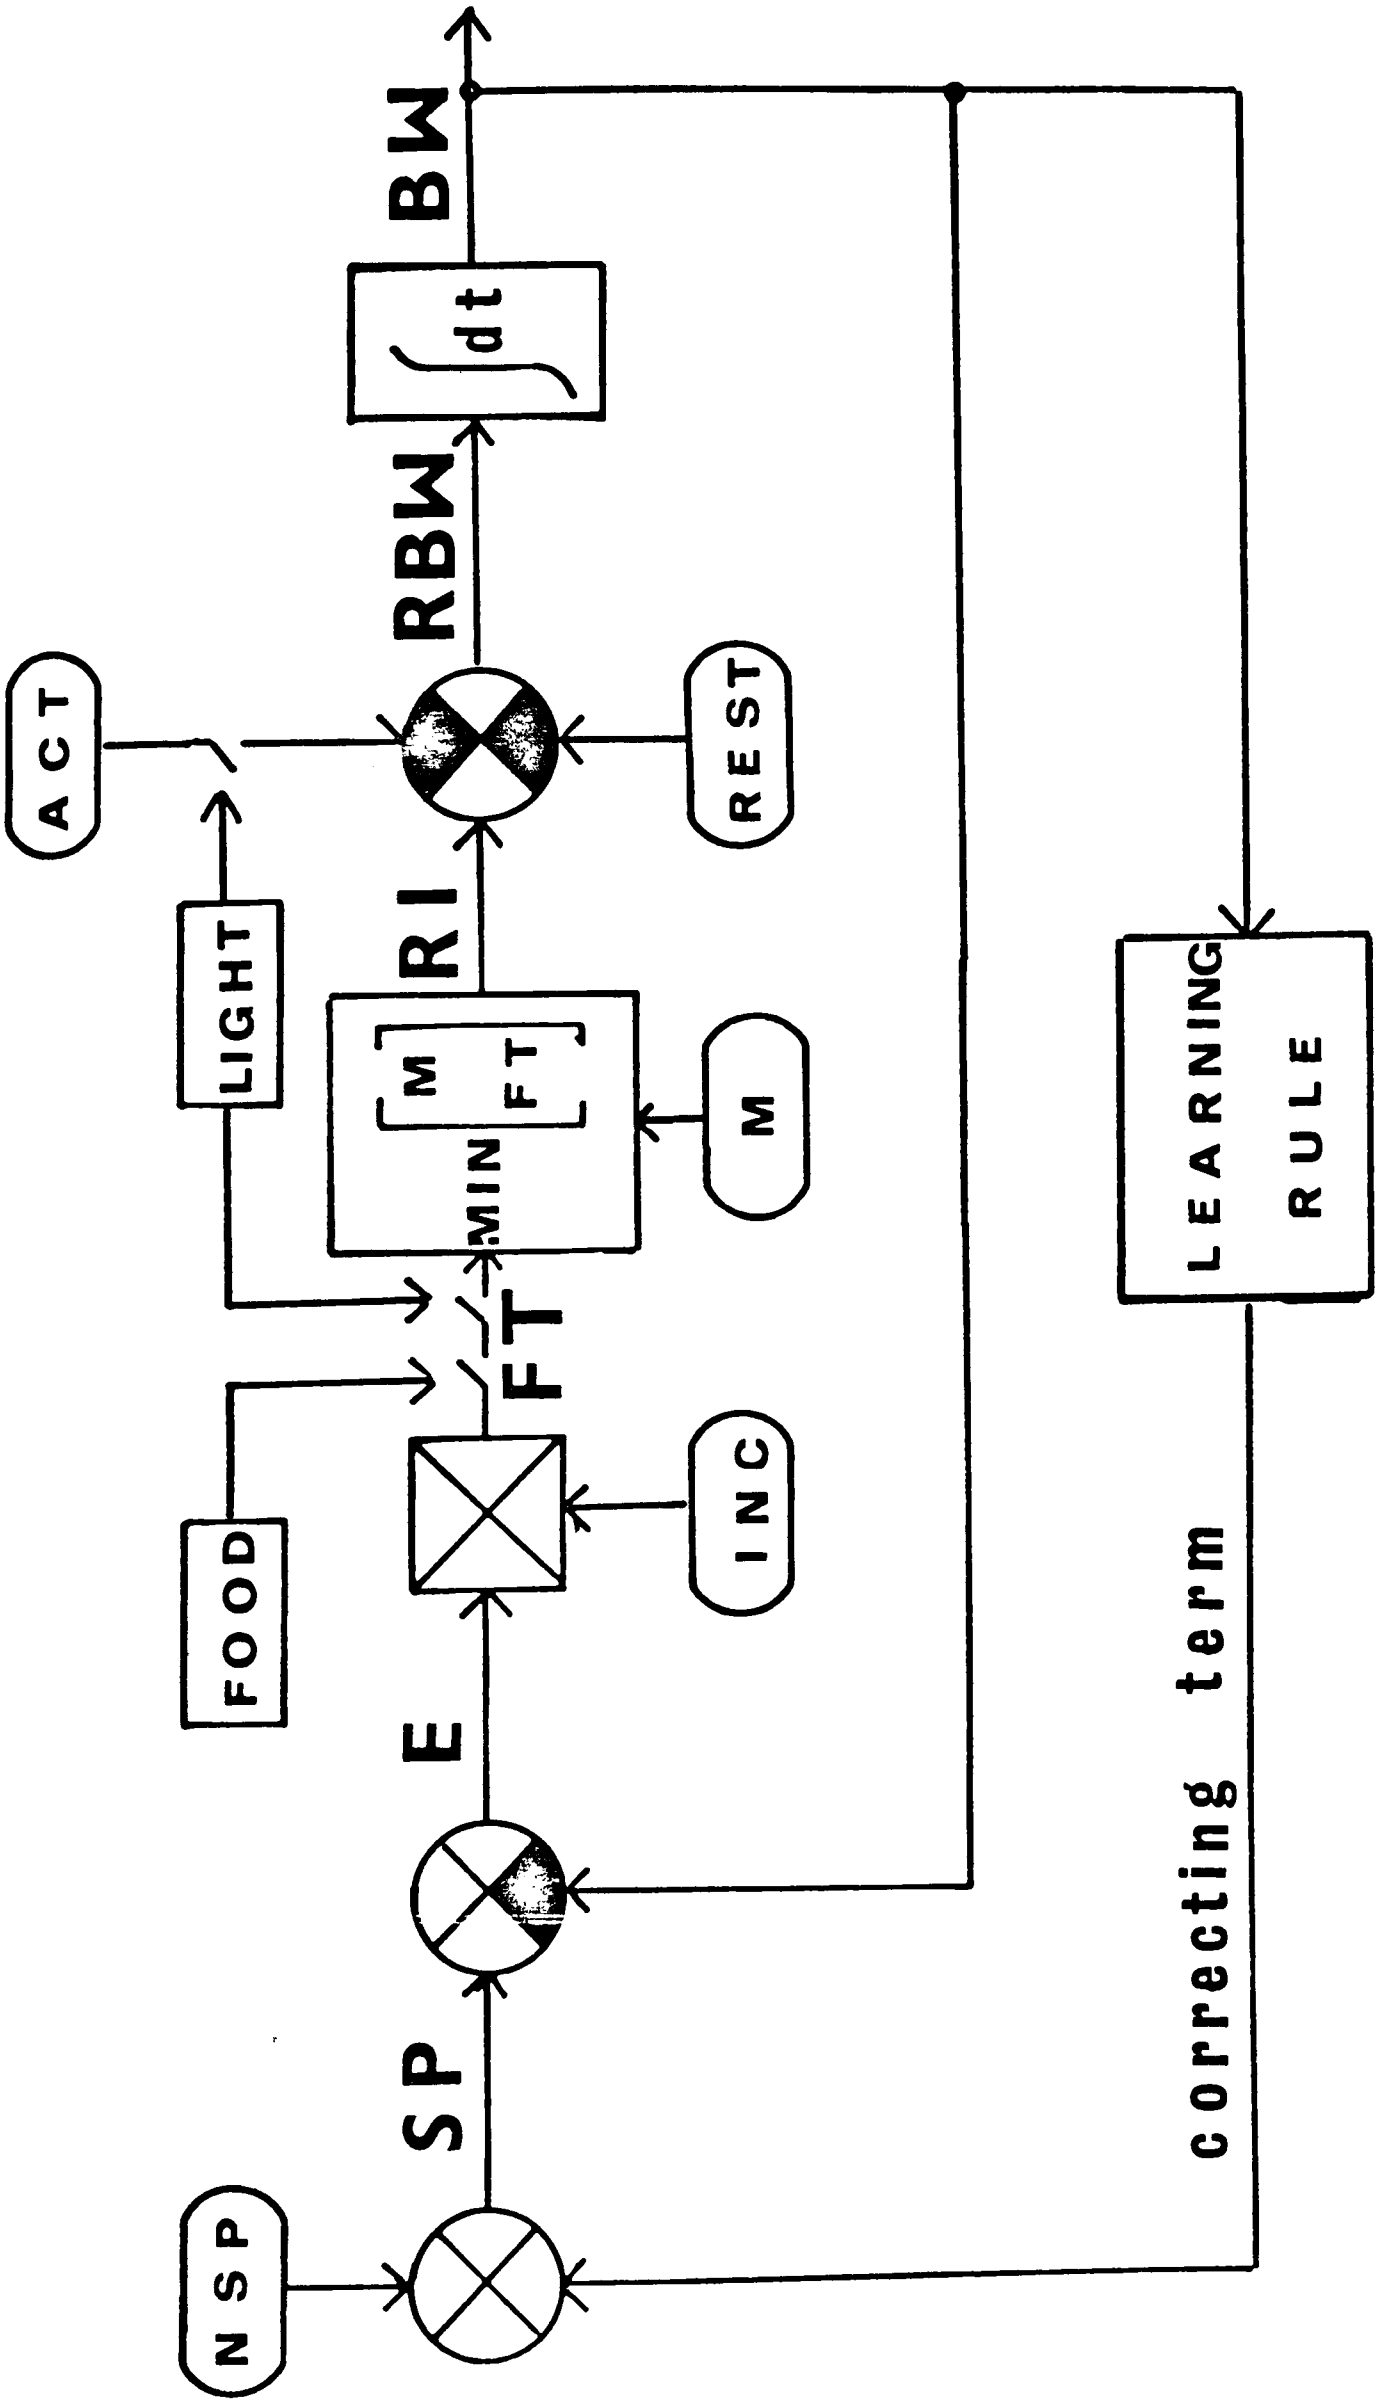
Computer simulation based on homeostatic model depicted in Figure 4. Notation identical to Figure 1.

Figure 4.6



Computer simulation based on homeostatic model depicted in Figure 4. Notation as in Figure 3.

Figure 4.7

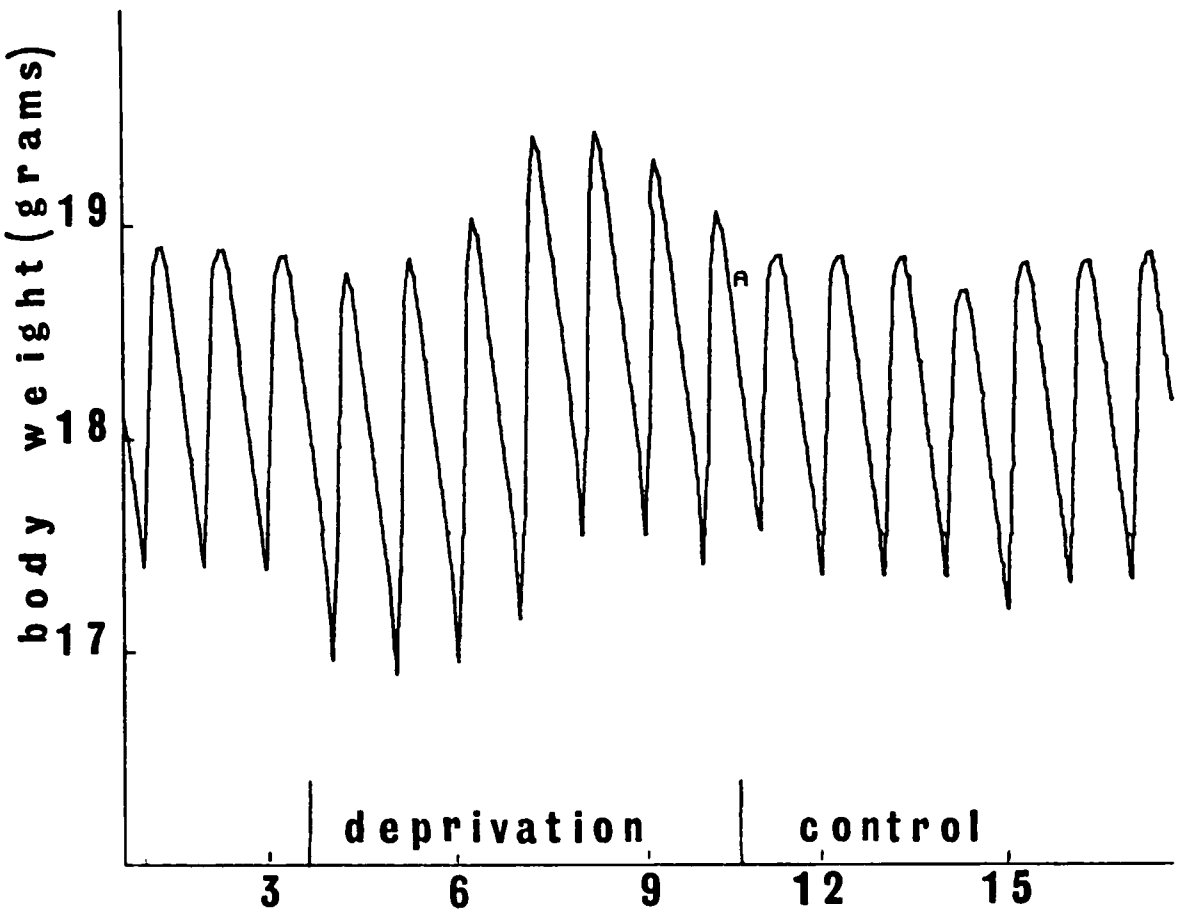


Homeostatic model with adjustable set point (SIMWGT 11). Notation as in Figure 4.4

NSP: Natural set point.

The learning rule is discussed in text and in Figures 8 and 9.

Figure 4.8

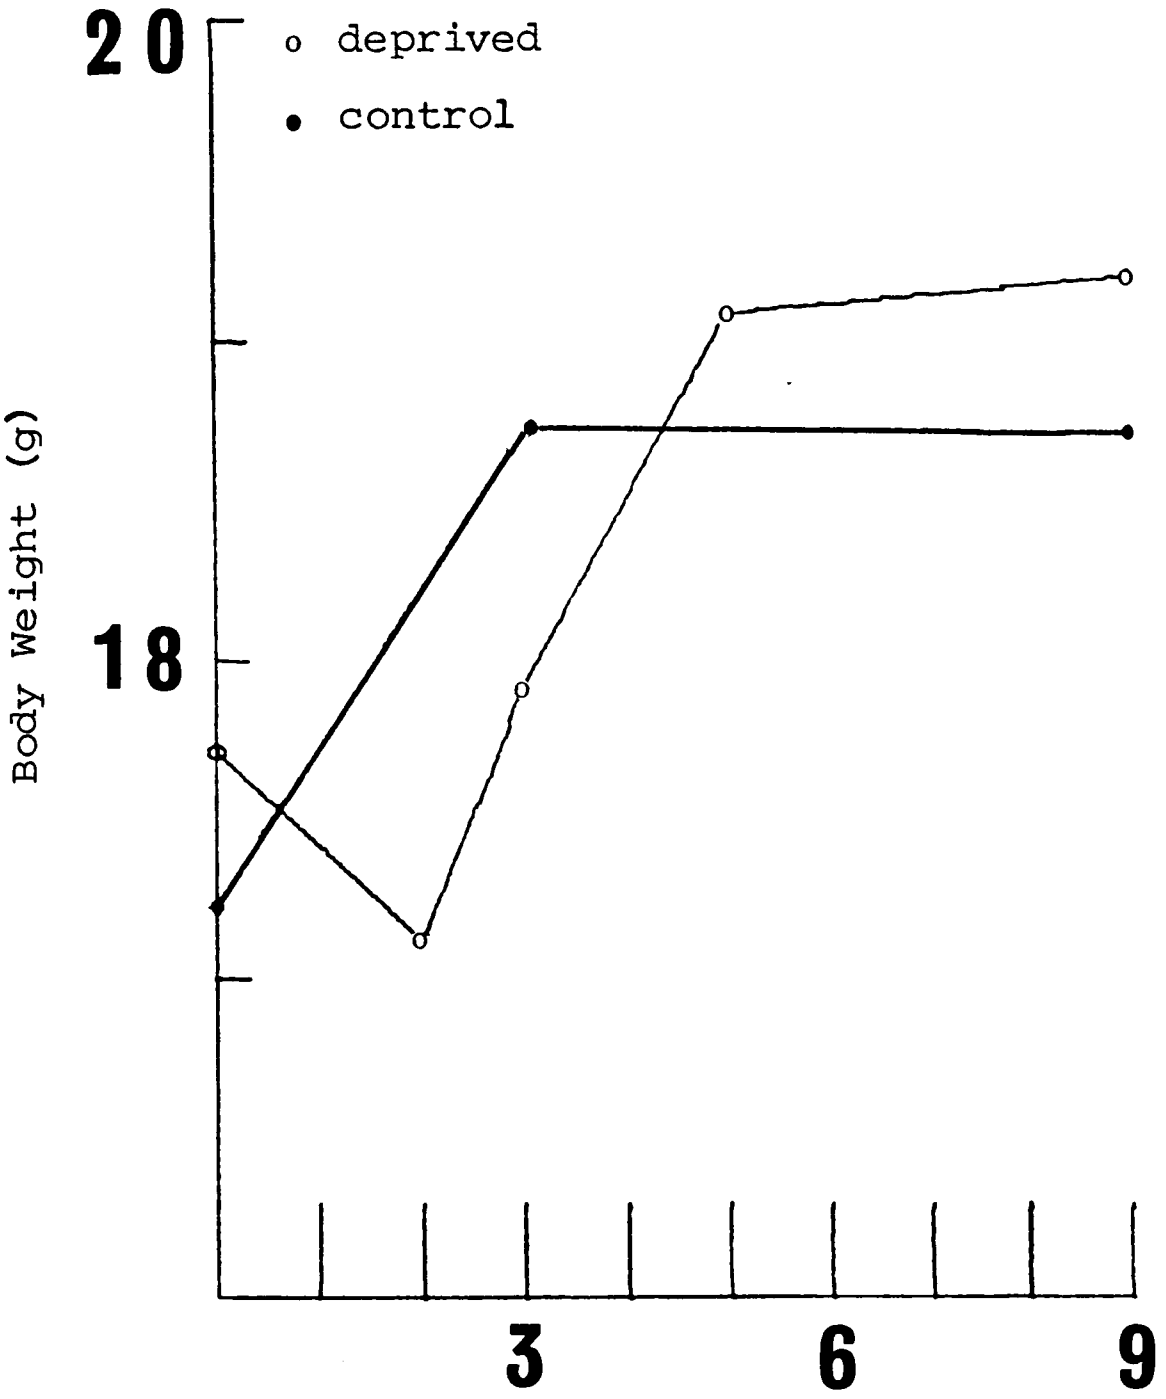


Computer simulation of SIMWGT 11. Notation as in Fig 1. The learning rule was:

$$\text{Correcting Term} = 0.6\ LBW_{(d-1)} + 0.3\ LBW_{(d-2)} + 0.1\ LBW_{(d-3)}$$

LBW: Lowest daily body weight; d: day, with reference to current one.

Figure 4.9



Computer simulation of SIMWGT 11
Notation as in Figure 3.
Learning rule as in Figure 8.

CHAPTER 5

On the motivational mechanisms underlying the dawn chorus

5.1. Introduction

I have argued that a number of environmental variables make the morning twilight period the most advantageous time for territorial song in the great tit (Chapters 2 and 3). My aim in this chapter is to investigate the nature of the mechanisms that produce the peak in morning song.

My study investigates the effects of foraging opportunity, time of the day and intensity of intrusions on the readiness of birds to interrupt foraging to sing or chase a rival. Because foraging opportunity and rate of intrusions are negatively correlated with time of day under normal field conditions the three factors could in principle cause a behavioural sequence of the type observed in the wild. These three factors are not mutually exclusive, because while each of them might be sufficient to account for the peak in morning song, there is no reason to expect that responding to one of the cues would indicate a lack of response to others. Also, they are not unique hypotheses because at least three other factors of functional significance have been shown to correlate with the dawn chorus: light, temperature and wind. The effects of these three factors as cues where not investigated here because the facilities available made them too hard to control.

In my experiments I manipulated foraging opportunity and presence of intruders independently at different times of the day, measuring the pattern of disruption of the ongoing behaviour.

I assumed that if any of the three factors could be shown to have an effect of readiness to sing under laboratory conditions it can also be assumed to contribute to the behavioural sequence in the wild.

5.2. Methods

5.2.1. Birds

I used 4 experimentally naive male great tits. They had been captured in November, 1977, kept for 2 weeks in motels (see Section 1.4) and then released into the outdoor compartments of the set-up shown in Figure 1.4, where they lived individually. They were allowed to become territorial in the indoor experimental aviaries by the procedure described in Section 1.11, which consisted of giving the birds access to the indoor area on alternate days for several months before the experiment took place in May, 1978. All birds were trained to operate the feeding discs (see Sections 1.9 and 1.10) during February and given extensive practice with different variable ratios schedules during March and April.

All the birds were singing by mid-February, both spontaneously as well as in response to playbacks of great tit song. They also sang and displayed when the local wild great tits approached the outdoor aviaries, and to a lesser degree responded to each other.

5.2.2. Aviary arrangements

Figure 5.1 shows the arrangement of the indoor aviaries where the experiment took place. One feeding disc was placed close to one of the walls on a platform 1.5 m high. On the other side of the room, about 3.0 m away, there was a 'Puppet theatre',

a box used to present apparent territorial intruders (Section 1.11). Two artificial 'trees' were placed in the aviary to give the birds resting perches. The 'approach' tree, was placed 0.5 m from the front of the theatre, and the other tree was placed in the centre of the room, equidistant from the feeding disc and the puppet theatre.

At both sides of the theatre there were loudspeakers connected to a stereo tape recorder (Sony TC 270) used to play back songs.

5.2.3. Experimental design

There were 5 different types of test, or treatments (Table 5.1). Four of them resulted from the factorial combination of 2 levels of profitability, and two lengths of intrusion. The 2 levels of profitability corresponded to reward probabilities of 0.25 and 0.05, which I call below high and low feeding opportunities (see Section 1.9.2) in the feeding disc. The 2 lengths of intrusion corresponded to presentations of the intruder bird for either 60 seconds or 30 seconds. All presentations of the intruder were accompanied by playback of great tit song.

The fifth type of test was a control for the effect of the operation of the theatre. It consisted of a test with low feeding opportunity and an operation of the theatre (opening the front and closing it again) for the time of a long intrusion, but with neither intruder bird or song.

Each test lasted 20 minutes, with a single intrusion during the fifth minute. The experiment was run as a latin square replicated 4 times, once for each subject. Each subject was tested

5 times a day on five consecutive days, receiving every day the five types of treatment in a different order. It was not possible to test birds simultaneously, and in consequence each subject was tested on a different (consecutive) week, the whole experiment lasting from May 2nd to May 30th, 1978.

The data collected on the last experimental day were lost due to an accident in the retrieval of the recorded tapes, so that the latin square was not complete for one of the four subjects.

Testing began each day 15 minutes after civil twilight, and from then onwards one test started every 60 minutes.

5.2.4. Deprivation

The food was removed and the aviaries swept every evening after the subject had gone to roost. Hunger during the tests was standardised using the following procedure.

Results obtained in pilot tests showed that the total number of rewards obtained by the birds in a 20 minute period at high feeding opportunity was normally under 30. Using this information, in the interval between tests each animal received the number of rewards necessary to reach a total of 35, when added to the rewards obtained during the previous test. In this way the cumulative intake at the beginning of each test was independent of the type of treatment received before, and was only a function of the number of tests received that morning.

Food ad libitum was provided one hour after the end of the last test of the day. By then, each bird had eaten 175 rewards, which compares with approximately 300 eaten during the same

period of continuous free access to the feeding disc.

5.2.5. Intrusion

The technique used to simulate intrusions was the first intrusion procedure described in Section 1.11. The bird used as the intruder was a male great tit, captured in the same month and place as the experimental subjects, and kept isolated from them since its capture. It was housed in a cube metre cage provided with a roosting box used for transporting the animal to the theatre. The bird was placed in the theatre before each test, and the intrusion consisted of opening the theatre for a given time (one or half minute) while playing back the song of a great tit through the loudspeakers at both sides of the theatre.

During the week before the experiment began the intruder bird was placed in the theatre every morning, so as to familiarise it with the procedure and avoid excessive stress when the tests took place.

5.2.6. Recording

The following events were recorded on magnetic tape using a computer compatible event recorder (see Section 1.12):

(1) operation of the feeding perch, (recorded automatically)

(2) delivery of a reward, (recorded automatically)

(3) Song. One event was recorded for every burst of song

(4) Aggressive displays. I followed Blurton-Jones'(1968) classification and recorded the following displays: (a) head-up, (b) head-forward (horizontal), (c) tail-fanning, and (d) wings-out. I recorded occurrences of

crest-raising and head-down displays as well, but they were difficult to identify reliably and I did not use them in my analyses. Figure 5.2, modified from Blurton-Jones' (1968) shows diagrams of these displays.

(5) Approach. One event was scored for every ten seconds that the experimental bird spent on the approach perches.

(6) Attack. Attacks were normally discrete events lasting for less than ten seconds, but when they lasted for longer, one event was scored for every ten seconds or fraction.

Durations of Approach and Attack were estimates subjectively. All these behaviours are described in further detail in Section 5.3.

5.2.7. Measures of responses

Because I was interested in the effectiveness of the intruder as a modifier of the ongoing activity, all the analyses were based on comparisons between pre-, during- and post-intruder levels of activity.

There were two groups of response pattern while the intruder was present: (1) a drop in the rate of feeding activity, and (2) an increase in the rate of performing all activities oriented towards the intruder (singing, displaying, approaching, attack). On many occasions foraging was completely interrupted during the intrusion. When this happened, the difference in score between the pre-intrusion and the during-intrusion period depended exclusively on the pre-intrusion rate of responding. Because total suppression was a maximum response I did not consider it appropriate to assign different scores to tests that differed in pre-intrusion rate but that showed total suppression. In order to correct for this distortion, I used an index that expressed the

drop in foraging activity in the presence of the intruder in relation to pre-intruder rate of response. The index employed was:

$$m(F) = \frac{Fi - Fb}{Fb}$$

where m(F) is the measure of feeding suppression, Fi is the rate of feeding during the intrusion, and Fb is the pre-intrusion (or baseline) rate of feeding.

m(F) was normally between -1 (total suppression) and 0 (no effect), but in a few cases a small positive ratio of m(F) was recorded.

For the other group of responses this complication did not arise, and the score used was simply the increase in the rate of behaviour during the intrusion with respect to pre-intrusion rate:

$$m(a) = Ai - Ab$$

where m(a) is the measure of response in activity 'a', Ai is the rate of activity during intrusion and Ab is the rate in the pre-intrusion period.

Figure 5.3 is one example of each type of response to illustrate the meaning of the scores.

The number of occurrences of each behaviour was computed for each half-minute interval of the test. The data collected during the first minute were not used to avoid the transient at the beginning of the test when birds were accomodating to the feeding schedule.

The score for each behaviour during the following four minutes was averaged and this mean is termed in the rest of this chapter as the baseline or pre-intrusion level for each particular behaviour.

In addition to the already described indices of the intensity of the response during the presence of the intruder, I was interested in testing the after-effects of the intrusion. I did this by tracing regression lines through the log transformed data for the five minutes after the intruder disappeared. The regression coefficients (positive in the case of feeding suppression and negative in all other behaviours) were compared in order to get an idea of the readiness to return to the previous level of activity under the different conditions. The logarithmic transformation was used because the recovery curve approximated an exponential decay curve, similar to those found in the recovery of baseline readiness to attack smaller fish after a presentation of a dummy to a resident territorial male Haplochromis burtoni (Heiligenberg, 1974). The slopes of the regression lines were used to measure the speed of return to baseline activity. An additional measure of the duration of the disturbance produced by the intrusion was obtained by computing for each test the time after the intrusion when the recovery regression line equalled the average between the pre- and during-intrusion rate of response.

The score obtained is called the response half-life, and differs from the analysis of the slopes in being sensitive to the value of the baseline and the magnitude of the response. When no signs of return are observed during five minutes, the regressior line diverges from the baseline, and the response half-life equals infinite. Because infinite scores did occur, the half-life of the responses could not be analysed parametrically.

5.3. Results

5.3.1. Description of responses

During the period preceding the intruder presentation the birds spent most of their time close to the feeding disc, almost continuously involved in either hopping on the perch or handling food. Short calls were rather common, both as responses to vocalisations from birds in the outer aviaries and spontaneously. The birds also occasionally sang, but in contrast to the short calls, the birds usually flew to the highest perch in the central 'tree' to sing, or at least interrupted their operation of the perch while singing. Sometimes they responded to a distant song from a bird outside by approaching the puppet theatre. The presentation of the intruder evoked dramatic changes in the behaviour of three of the subjects and minor changes in the fourth one who never sang or displayed to the intruder during the experiment. This less responsive animal was the last one to be tested, in the week between 22 and 27 of May. This week roughly coincided with the end of the Spring singing season for wild great tits in Wytham Woods during that year (J. R. Krebs, personal com.), and in consequence it seems likely that its failure to respond aggressively was due to a seasonal effect. While it is not possible to discard completely the possibility of a genuine individual difference between this animal and the rest (as opposed to a seasonal effect), this explanation seems unlikely because this bird had behaved similarly to the others during the training, counter-singing loudly to the sporadic songs coming from the outdoor aviaries. I considered these arguments strong enough to

exclude this subject from the statistical analysis. When the intruder was present the other birds usually flew from the foraging area to either the central or the 'approach' tree, where they sang, displayed and in some cases attacked the intruder. On occasions the interruption of feeding did not last all the time that the intruder was there; the resident bird returned to the feeding disc and resumed feeding after an initial interruption. Also, some birds returned immediately to forage after the intruder disappeared, while others continued to sing and stayed around the intrusion area for several minutes after the intrusion had ended. Singing towards the intruder took place mostly from the highest perch in the central tree. When perching on the 'approach' tree the predominant response was displaying visually, though some song occurred in the approach tree and some displays occurred in the central one.

Tables 5.2 to 5.4 present a general synthesis of the effect of profitability, length or intrusion and time of the day, respectively. The significance levels are extracted from the statistical analysis presented below for each behaviour separately.

5.3.2. Suppression of feeding by intrusion

In 60% of the tests the birds stopped feeding during the presence of the intruder. The index of feeding suppression m(F) (see Section 5.2.7. for definition) was -1.0 for all those cases, resulting in an extreme skewness of the data with 60% of the data points concentrated on one boundary of the distribution. A parametric analysis of variance is not valid in these circumstances, and the use of any other test requiring independence of the data

was also precluded by the fact that my data came from several replicates in a few individuals. Intuitively it can be seen that these difficulties result from the fact that the presence of the intruder in this particular experimental set up is an extremely powerful stimulus, and its effect is so dramatic that it obscures differences between intrusion conditions. In spite of the impossibility of assessing significance levels, the difference in responsiveness shows trends which are worth discussing. Tables 5.2 to 5.4 show the differences in the average effect. Low profitability was associated with stronger suppression than high profitability, while short intrusions were more effective than long ones. No trend was apparent with time of day. While the effect of profitability goes in the direction predicted by the hypothesis, potential differences between the means are hidden by the boundary effect discussed above. In order to describe the data, removing as much as possible the boundary effect, I used the following procedure. Tests belonging to the same day (hence to the same bird) and having the same length of intrusion were arranged in pairs. Two pairs resulted for each day: high profitability-long intrusion vs. low profitability-long intrusion and high profitability-short intrusion vs. low profitability-short intrusion. From then onwards I followed the procedure for performing a sign test, classifying the pairs according to whether the value of $m(F)$ was bigger in the high profitabilit test or vice versa. In this way, pairs where suppression was total in both high and low profitability do not influence the result.

The same procedure was employed to describe the effect of length of intrusion on feeding suppression, and the results of both analyses are presented in Table 5.5. In more than twice as many pairs feeding suppression was more effective when profitability was low ($m(F)$ is more negative) and the same applies to short vs. long intrusions.

5.3.3. Return to baseline feeding rate

In the example shown in Figure 5.3 a few minutes after the intruder disappeared the subject was already feeding at its original, pre-intrusion rate. There was a great deal of variation in the pattern of recovery to the baseline feeding rate. In some tests the subjects switched back to their original rate of activity as soon as the theatre was closed, while in others they remained close to the intruder area for several minutes, singing or calling frequently. But in most of the occasions the pattern was one of gradual return, with the animals going back and forth between the feeding and the intruder area and progressively returning to the baseline pattern.

Tables 5.2 to 5.4 show that the slope of the recovery line was steeper when profitability was poor than when it was high, and when the intrusion was short than when it was long. No clear trend is visible in the time of the day effect. The analysis of variance of the rate of recovery is shown in Table 5.6. The only significant factors were subjects and feeding opportunity. This last effect was in the opposite direction to the expectation from the general hypothesis: I predicted that the rate of recovery should be faster for high opportunity tests than for those with

low foraging opportunity. The effect of profitability upon the median values of the feeding response half life (Table 5.2) show the same trend as the analysis of the slopes, while the length of the intrusion does not seem to have any effect (Table 5.3) nor does time of the day produce any clear trend (Table 5.4).

5.3.4. Intensity of the song response

The effect of feeding opportunity, length of intrusion and time of the day on the average increase in rate of song is shown in Tables 5.2 to 5.4. The song response was almost twice as intense for low than for high foraging opportunity. Neither length of the intrusion nor time of day showed any strong trend. The analysis of variance of these results is shown in Table 5.7. The only significant factor in this analysis was the experimental manipulation of foraging opportunity and length of intrusion. Table 5.8 presents the average response for each of the five treatments.

5.3.5. Return to baseline singing rate

I used the same type of analysis as in the return to feeding baseline. The average responses appear in Tables 5.2 to 5.4 and the analysis of variance of the slopes of the recovery lines appears in Table 5.9. The results are similar to those shown for the return to feeding rate baseline. Once again the return to the baseline was faster in the non-profitable than in the profitable tests. In addition there was a significant result in the contrast between the empty theatre and the real intruder,

which was not observed in the case of the return to feeding baseline. This is due to the absence of singing response when the theatre was empty (Table 5.8). Similar to the feeding response, the median half life of the increase in song rate is much longer in high than low profitability tests (Table 5.2), but no difference is apparent for different intrusion lengths (Table 5.3) or time of day (Table 5.4).

5.3.6. Number of attacks

In 9 experiments out of 75 the resident bird physically attacked the intruder, by launching itself against the transparent front of the theatre and pecking at it. Each attack did not usually last for long. Typically the bird would attempt an attack and after ten seconds would fly back to the 'approach' perch, where it sometimes displayed, sang or returned to another attack. These episodes were recorded as individual attacks, so that I could obtain an index of the intensity of attacking in different tests.

All the attacks took place in long intrusion tests. This is probably due to the fact that it takes some time to reach the level of confrontation required for overt fighting, and short intrusions were simply too short for this level to be reached. All tests in which at least one attack had been observed were matched with the other long intrusion test the same day, in order to compare the score of high and low feeding opportunity. Seven pairs resulted from this procedure, all of them with higher number of attacks in low than in high feeding opportunity. The lack of the data precludes the assessment of statistical significanc

5.3.7. Approach

I recorded one 'approach' for every ten seconds which the subject spent in the approach tree. Tables 5.2 to 5.4 present the mean rate of approach. Table 5.10 shows the results of the analysis of variance of the approach scores, expressed in all cases as the difference between number of approaches during intrusion minus number of approaches during baseline period. In addition to significant individual differences, the only significant factor was the contrast between empty theatre control and real intrusions. The level of approach during tests with an empty theatre was significantly higher than during tests with the intruder.

5.3.8. Displays

All visual displays were combined by addition into a single score of displays, and this value was treated in the way previously described for variables that increased with the presentation of the intruder (presented as 'Displays' in Tables 5.2 to 5.4). Table 5.11 shows the result of the analysis of variance on these combined scores.

As with most previous variables, there was a significant individual effect. Treatment was also significant, mainly due to the stronger response in long than in short intrusions and to the contrast between empty theatre control and all the rest of the tests.

5.4. Discussion

5.4.1. Foraging opportunity

The results described in this chapter support the view

that feeding opportunity may be one of the cues involved in the motivational mechanism of the dawn chorus.

The feeding behaviour of territorial great tits was more disturbed when profitability was low than when it was high, and the intensity of the song response was also greater in low than in high profitability tests.

It is worth pointing out that these two results are independent. Birds in this experimental situation did not need to stop operating the feeding perch while singing or displaying to the intruder. The fact that they did interrupt foraging while singing indicates that the ecological constraints probably operating in the wild to make it advantageous to sing and forage in different places has lead to the evolution of motivational mechanisms that make the two activities partly incompatible on motivational grounds. The possibility of such constraints ought to be taken into account when discussing optimal decision-making by animals. Presumably, if changes in the environment made it advantageous for great tits to spend a higher proportion of their time foraging in the tree tops, selection would operate to eliminate the motivational incompatibility of singing and foraging in the same way in which it would favour morphological adaptations to the new type of behaviour. Until more information is obtained about the genetic variability of morphological or behavioural patterns, there is no reason to believe that morphological features are more genuine constraints to optimality problems than motivational rules.

While the results obtained for responses in the presence of the intruder support the view that feeding opportunity may be

causally linked to the dawn chorus, the after-effects of the intrusion directly oppose this hypothesis. The birds modified their ongoing behaviour for longer when feeding opportunity was high than when it was poor. I have no option but to seek for a post-hoc explanation for this observation, which I find puzzling. It could be argued that differences in hunger are confusing the issue. If birds experiencing high profitability eat more before the intrusion than animals under low profitability, the former may be less hungry when the intruder appears, and in consequence take more time off from feeding than when profitability is low. I tried to control for this effect by setting the timing of the intrusion early in the test so that even for the high profitability tests, the onset of the intrusion took place before any obvious satiation could be observed. Moreover, if birds in profitable tests were less hungry when the intruder appeared, one would expect greater responsiveness during the intrusion in high profitability tests, which is the opposite of the obtained results.

Another possible argument could be that my interpretation that the aggression observed is an expression of territorial defence is unwarranted. If the birds were fighting for the defence of the food resource with a short term perspective, a better source would be more intensely defended that a poor one. Morris (1968) and Kikkawa (1968) have shown that intensity of aggression and dominance relationships in flocks of captive silver eyes Zoosterops lateralis fighting for access to food varied with the amount of food eaten by the most dominant animals. Also,

Carpenter and McMillen (1976) found that Hawaiian honeycreepers Vestiaria coccinea defended territories between two levels of resource abundance, as predicted by an optimality model. I have given reasons why it does not seem likely that the territorial behaviour of the great tit is aimed at maximising feeding rate in the short term (see Chapter 2) but even if this assumption was wrong, the contradiction in the results is not solved. The same as the argument based on hunger level, this hypothesis predicts higher aggressivity in tests with higher profitability since the value of the food source would then be greater. This is the opposite of the results for behaviour during the intrusion.

5.4.2. Circadian clock

In contrast to the effect of profitability the time of day at which the tests took place did not have any significant effect on any of the measures of responsiveness. This result is unexpected at the light of many previous reports of circadian control of activity levels in birds (Daan and Aschoff, 1975), but does not conflict with them directly. Previous authors found that what I called baseline activity was modulated by a circadian pacemaker, but my interest in this experiment was oriented towards the readiness to modify ongoing behaviour rather than to the rate of spontaneous activity. It could be argued that in my experimenta design, time of day is identified with order of the test, and this could confuse the interpretation. Indeed, in spite of the control for standard food intake at the end of each treatment, there is a cumulative satiation problem. The amount of food that a bird had ingested prior to any one test was independent of the type of test that it had experienced before, but it was not independent

of the number of tests experienced. Thus, at the beginning of the test corresponding to four hours after twilight each subject had eaten twice as much as at the beginning of the test corresponding to two hours after twilight.

If the cumulative intake is less (or more) than the animal's target of that time of day, an increase (or decrease) in feeding tendency may counter-balance the commands from any pacemaker, masking circadian variations in responsiveness. This is a real possibility, and it applies to my study as much as to the studies of circadian modulation of baseline activity. Hunger increases activity levels, (see review in Hinde, 1970) and supplying food ad-libitum (as in Daan, 1976) does not seem preferable to my technique of standardising the amount of intake ingested every hour. Diurnal animals will be hungrier at the end of the dark period, and if food is freely available hunger will decrease rapidly after lights are switched on, confusing the interpretation of the peak in baseline activity found at dawn.

5.4.3. Intruders pressure

In addition to the changes in readiness to respond , the behaviour of a territorial bird in the wild may show daily fluctuations that result directly from the stimulation from intruders. The intense disruption produced in the ongoing feeding activity by the appearance of the intruder indicates that the allocation of time to feeding and singing is strongly affected by the arrival of intruders. Given that intruders arrive more often early in the morning than at any other time of the day (Krebs and Perrins, 1978) a foraging animal would engage in territorial activity more often at dawn.

5.5. General conclusions

Of three possible mechanisms of the dawn chorus investigated (foraging opportunity, circadian clock and intrusion rate), two were found to be effective in changing the allocation of time and effort to feeding and singing, while one of them did not show significant effects.

Foraging opportunity was effective in that great tits showed a more complete cessation of feeding activity and reacted more severely towards intruders when foraging opportunity was low than when it was high. The data on the effect of foraging opportunity were paradoxical, however, in that the birds took significantly longer to return to baseline activity when profitability was high.

There was not effect of time of day in any of the variables studied but it is possible that circadian variations were obscured by hunger or other factors. The effect of intrusions was so dramatic that it is most likely that rate of intrusions in the wild would affect the allocation of time to singing and foraging.

TABLE 5.1

Experimental Treatments

High and low profitability refer to reward probabilities of 0.25 and 0.05 respectively. Long and short intrusion lengths were 60 or 30 seconds respectively.

Treatment	Profitability	Intrusion length	Intruder present in the Theatre
1	high	long	yes
2	high	short	yes
3	low	long	yes
4	low	short	yes
5	low	long	no

TABLE 5.2

Effect of Profitability

Profitability	High	Low	Significance
Feeding suppression	-0.76	-0.87	NT
Song response	2.82	4.73	0.01 (ME)
Feeding recovery slope	0.12	0.75	0.01
Song recovery slope	-0.20	-0.44	0.01
Feeding response half life	∞	3.8	NT
Song response half life	6.2	3.4	NT
Approach response	1.84	2.29	Ns
Displays response	2.06	2.66	Ns

Notes:

Ns - Not significant

NT - Not tested

ME - Contrast of high vs low profitability was not tested, but main treatment effect is significant.

Feeding and song half lives are medians. For description of other responses see text.

TABLE 5.3

Effect of Intrusion Length

Intrusion Length	Long	Short	Significance
Feeding suppression	-0.74	-0.89	NT
Song response	3.49	4.05	0.01(ME)
Feeding recovery slope	0.34	0.53	Ns
Song recovery slope	-0.34	-0.30	Ns
Feeding response half life	7.6	8.8	NT
Song response half life	4.15	5.5	NT
Approach response	2.60	1.54	Ns
Displays response	3.14	1.59	0.05

Notes:

Ns - Not significant

NT - Not tested

ME - Contrast of high vs low intrusion was not tested, but main treatment effect is significant.

Feeding and song half lives are medians. For description of other responses see text.

TABLE 5.4

Effect of Time of Day

Hours since civil twilight	1	2	3	4	5	Significance
Feeding suppression	-0.85	-0.83	-0.66	-0.79	-0.90	NT
Song response	2.98	3.11	3.26	2.81	2.76	Ns
Feeding recovery slope	0.70	0.16	0.57	0.40	0.35	Ns
Song recovery slope	-0.28	-0.30	-0.28	-0.24	-0.20	Ns
Feeding response half life	6.5	∞	11.1	4.2	7.8	NT
Song response half life	7.5	3.2	4.2	5.4	4.5	NT
Approach response	2.05	1.19	2.82	3.27	4.45	Ns
Displays response	3.27	0.88	2.34	1.29	1.76	Ns

Notes:

Ns - Not significant

NT - Not tested

ME - Contrast of high vs low profitability was not tested, but main treatment effect is significant.

Feeding and song half lives are medians. For description of other responses see text.

TABLE 5.5

EFFECT OF PROFITABILITY AND INTRUSION LENGTH ON FEEDING SUPPRESSION

a) Effect of profitability

The four intruder tests each day are arranged in two pairs of equal intrusion length.

	number of pairs
m(F) profitable > m(F) non-profitable	15
m(F) profitable < m(F) non-profitable	6
Total suppression on both conditions	17
TOTAL	38

b) Effect of length of intrusion

The four intruder tests each day are arranged in two pairs of equal profitability.

	number of pairs
m(F) short intrusion < m(F) long intrusion	15
m(F) short intrusion > m(F) long intrusion	5
Total suppression under both lengths of intrusion	18
TOTAL	38

$$m(F) = \frac{\text{feeding rate during intrusion - feeding rate baseline}}{\text{feeding rate baseline}}$$

TABLE 5.6

RATE OF RETURN TO BASELINE FEEDING

Analysis of variance of regression coefficients of recovery lines.

Source of Variation	DF	SS	MS	F	p
Individual	2	6.30	3.15	4.0	0.05
Days	4	6.58	1.64	2.1	Ns
Time of day	4	2.59	0.65	0.8	Ns
Treatment main effect	4	6.53	1.63	2.1	Ns
High vs low profitability	1	5.91	5.91	7.5	0.01
Long vs short intrusion	1	0.57	0.57	0.7	Ns
Interaction	1	0.05	0.05	0.1	Ns
Control vs rest	1	0.004	0.004	0.0	Ns
Residual	60	47.45	0.79		
Grand Total	74	69.45			

TABLE 5.7

ANALYSIS OF VARIANCE OF INCREASE IN SINGING RATE AS A RESPONSE TO TERRITORIAL INTRUSION.

Source of Variation	DF	SS	MS	F	p
Individual	2	10.17	5.08	0.5	Ns
Days	4	55.76	1.39	1.3	Ns
Time of day	4	2.57	0.64	0.1	Ns
Treatment	4	247.61	61.90	5.6	0.001
Residual	60	667.02	11.12		
Grand Total	74	983.12			

Note:

Ns - Not significant

TABLE 5.8

INTENSITY OF SONG RESPONSE

Mean increase in song bursts per minute as a response to territorial intrusion.

Treatment	m(Song)
Non profitable-short	5.10
Non profitable-long	4.36
Profitable-short	3.01
Profitable-long	2.63
Control	-0.18

Note:

m(Song)= song rate during intrusion - baseline song rate

TABLE 5.9

RATE OF RETURN TO BASELINE SINGING

Analysis of variance of regression coefficients of recovery lines.

Source of variation	DF	SS	MS	F	p
Individual	2	1.94	0.97	9.9	0.001
Days	4	0.33	0.08	0.8	Ns
Time of day	4	0.11	0.03	0.3	Ns
Treatment main effect	4	2.04	0.51	5.2	0.01
High vs low profitability	1	0.87	0.87	8.9	0.01
Long vs short intrusion	1	0.03	0.03	0.3	Ns
Interaction	1	0.00	0.00	0.0	Ns
Control vs rest	1	1.14	1.14	11.7	0.01
Residual	60	5.88	0.10		
Grand Total	74	10.30			

Note:

Ns - Not significant

TABLE 5.10

APPROACH TO THE INTRUDER

Analysis of variance of increase in rate of approach to intruder area as a response to territorial intrusion

Source of variation	DF	SS	MS	F	p
Individual	2	14.59	7.30	8.9	0.001
Days	4	2.40	0.60	0.7	Ns
Time of day	4	3.21	0.80	1.0	Ns
Treatment main effect	4	7.45	1.86	2.3	Ns
Long vs short intrusion	1	1.11	1.11	1.3	Ns
High vs low profitability	1	0.09	0.09	0.1	Ns
Interaction	1	0.52	0.52	0.6	Ns
Control vs rest	1	5.74	5.74	7.0	0.05
Residual	60	49.18	0.86		
Grand Total	74	76.84			

Note:

Ns - Not significant

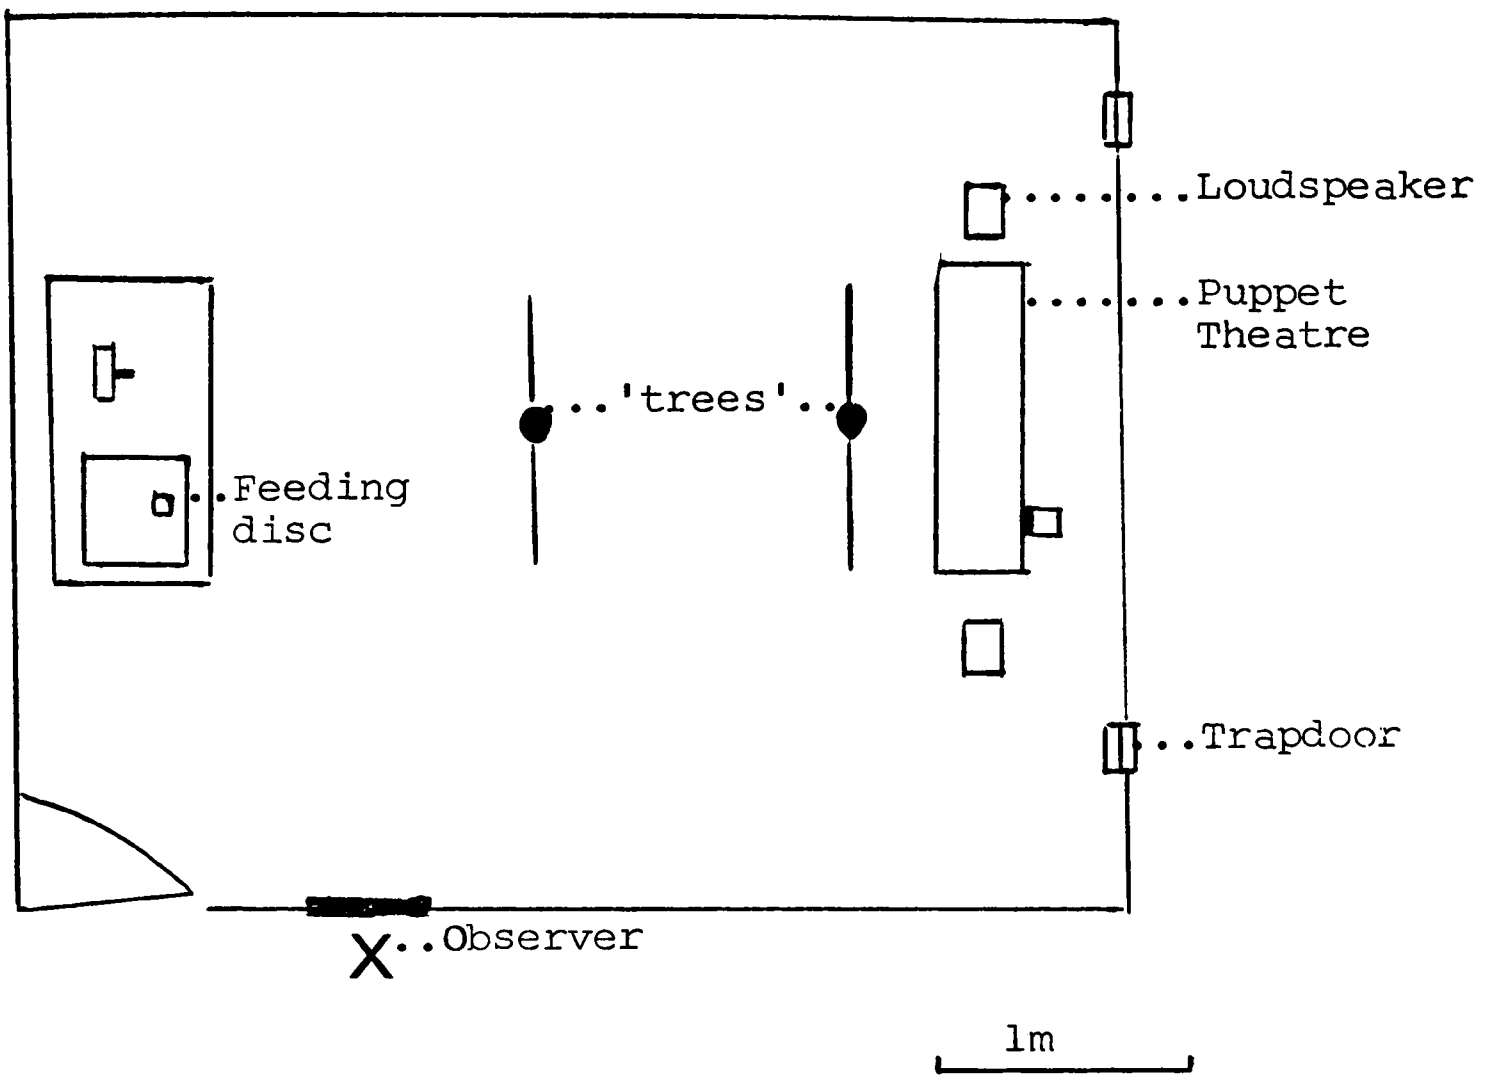
TABLE 5.11

COMBINED DISPLAYS

Analysis of variance of increase in visual displays as a response to territorial intrusion.

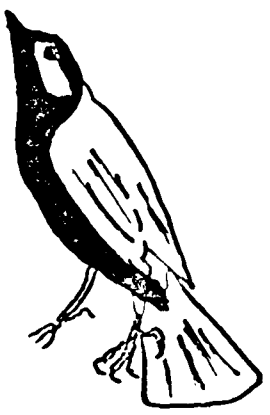
Source of Variation	DF	SS	MS	F	p
Individual	2	6.47	3.23	4.9	0.05
Days	4	1.32	0.33	0.5	Ns
Time of day	4	2.14	0.54	0.31	Ns
Treatment main effect	4	12.47	3.12	4.7	0.05
Long vs short intrusion	1	3.39	3.39	5.1	0.05
High vs low profitability	1	0.65	0.65	1.0	Ns
Interaction	1	2.78	2.78	4.2	0.05
Control vs rest	1	5.65	5.65	8.5	0.01
Residual	60	39.74	0.66		
Grand Total	74	62.15			

Figure 5.1



Aviary arrangements. The central tree had pairs of perches at 3 heights (1.0, 1.5 and 2.0 m) while the approach tree had a single pair of perches (0.8m high). Descriptions of the apparattus under General Methods.

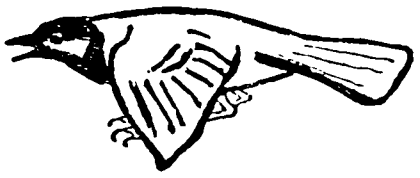
Figure 5.2



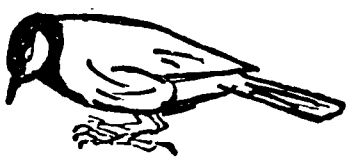
HEAD - UP



HEAD - FORWARD



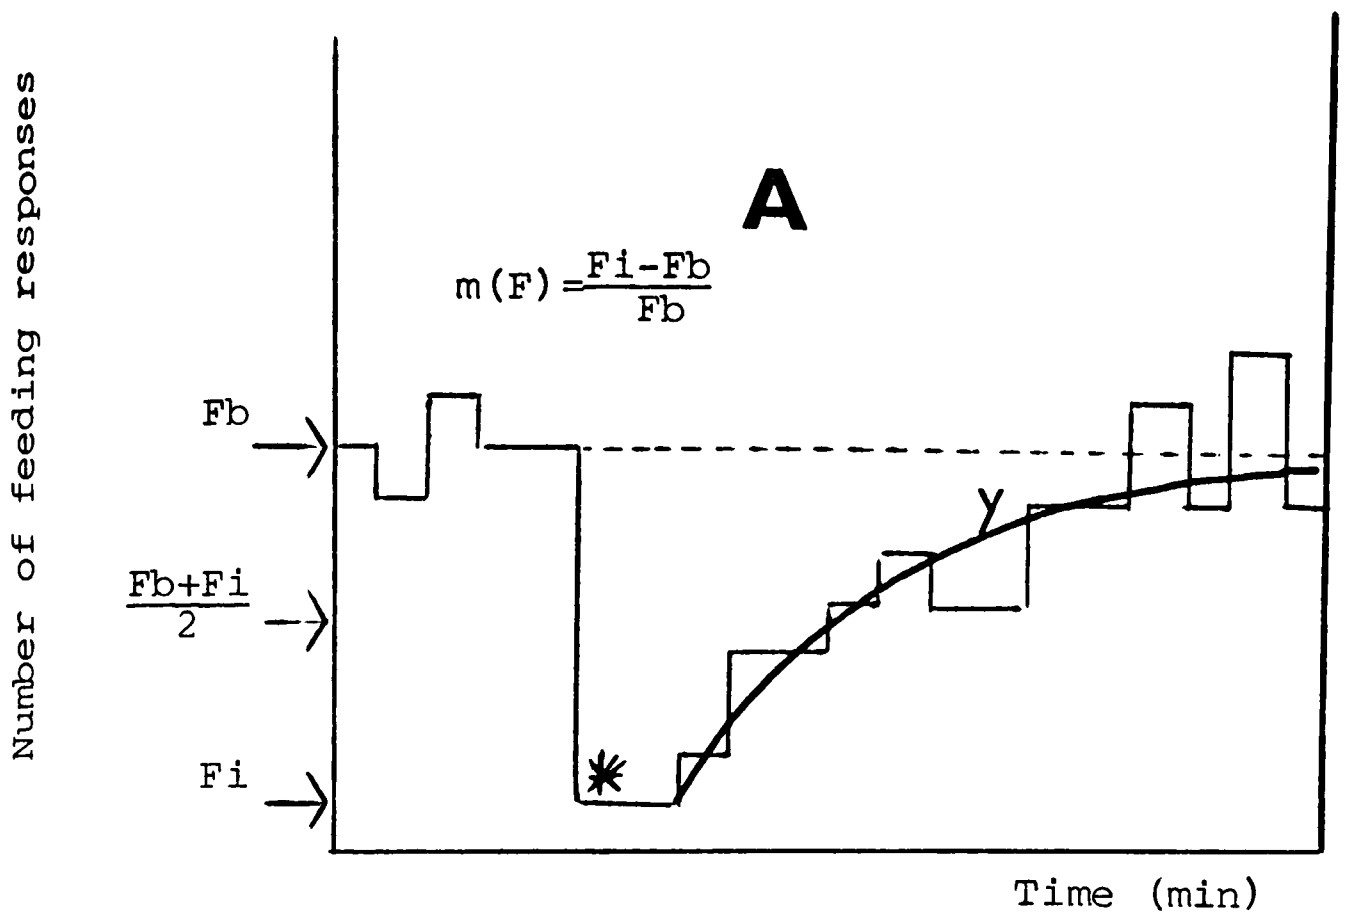
WINGS - OUT

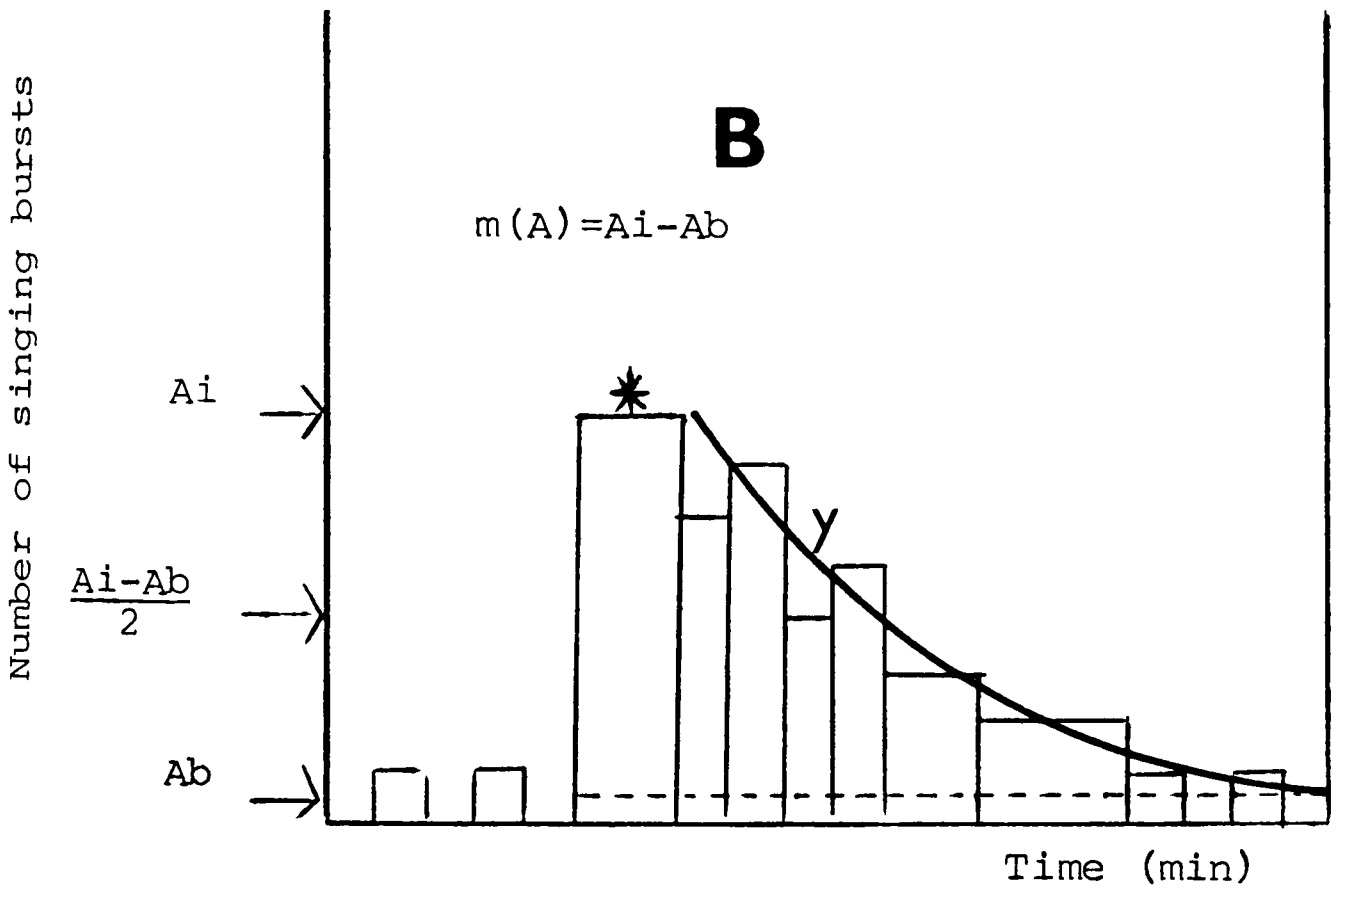


HEAD - DOWN

The different displays of the great tit referred to in this chapter. Tail-fanning can occur simulatenously with any of the other displays(see Head-up and wings-out in this illustration). From Blurton-Jones(1968).

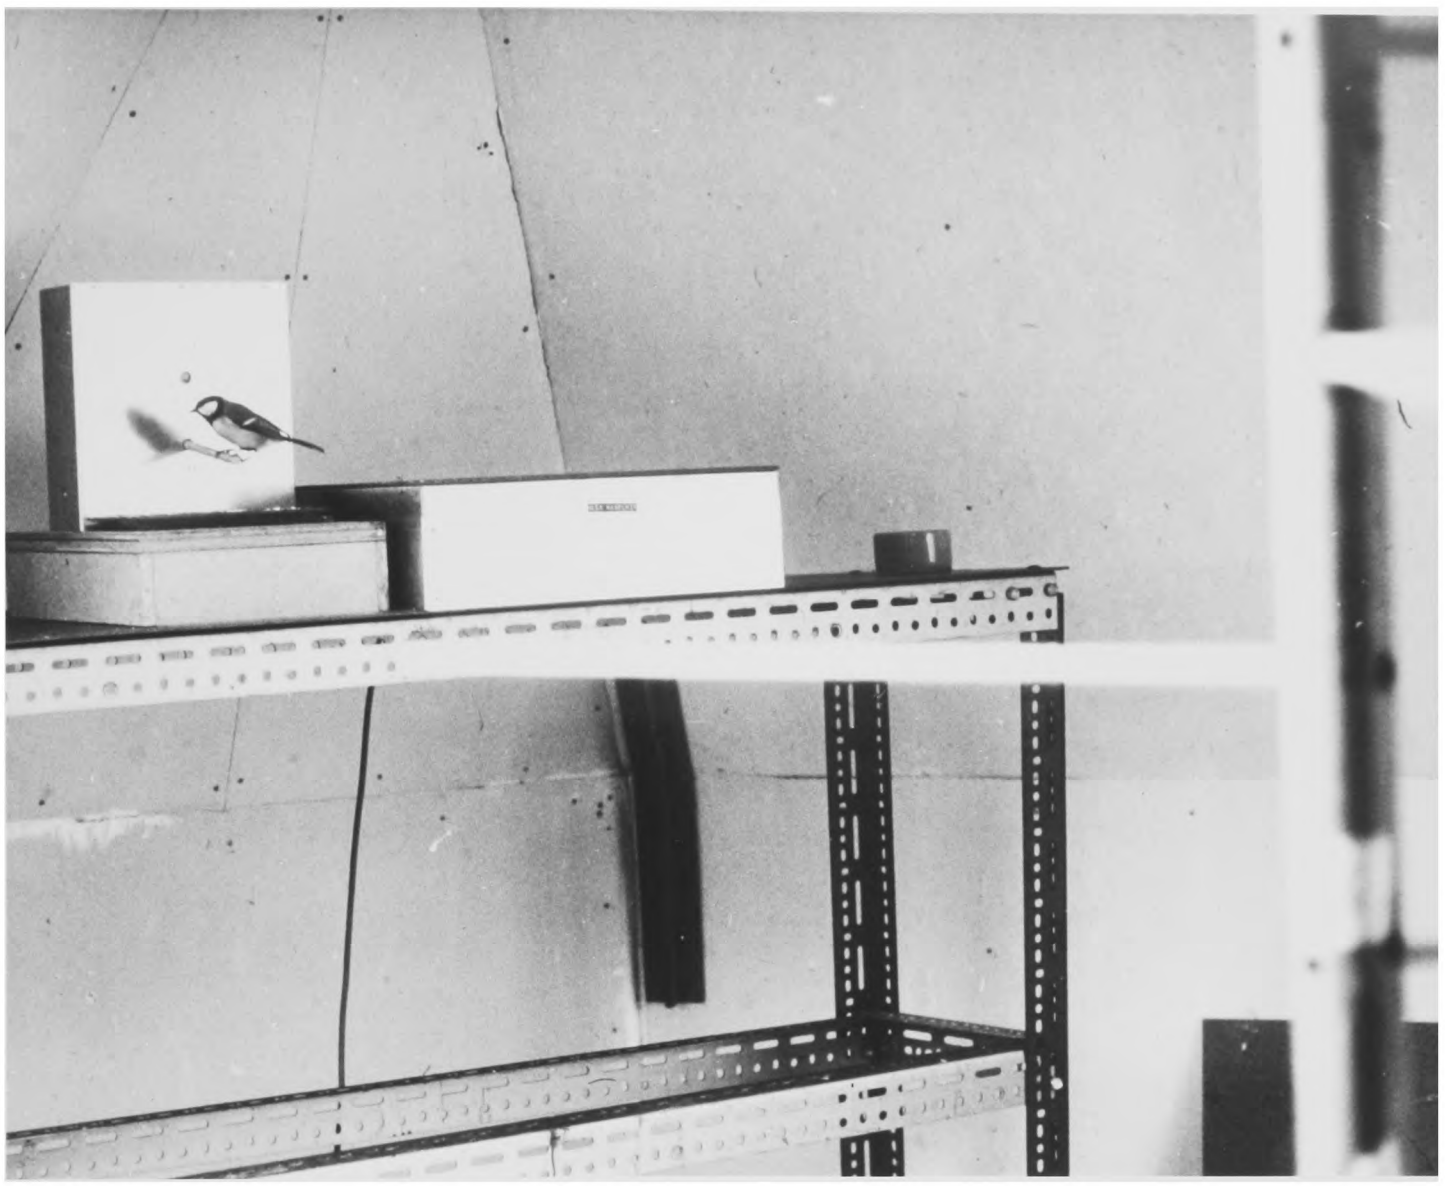
Figure 5.3





Examples of the analysis of feeding (a) and singing (b) responses The asterisk indicates the presentation of the intruder. Fi or Ai: response rate during the intrusion; Fb or Ab: response rate before the intrusion (baseline); Y: recovery curve. m(F)=(RI-RB)/RB is the measure of the feeding response. m(A) is a measure of all agonistic responses.

Plate 2



A great tit operating a feeding disc (Chapters 7 to 10). In the foreground, a 'Krebs tree'.

PART II

Studies on foraging behaviour: Optimal sampling

CHAPTER 6

6.1. Optimal foraging and sampling

Optimal foraging theory is an attempt to formulate rules about decision making by predators, based on the premise that animals are selected during evolution to harvest food in a maximally efficient way. Several extensive reviews have been published recently on both experimental and theoretical work on optimal foraging (Pyke, Pulliam and Charnov, 1977; Krebs and Cowie, 1976; Krebs, 1978 ; Staddon, 1979a,b), and the rate of additions to the literature is such as to make it difficult to keep track of it. But with the exception of few isolated examples (Pulliam, 1975; Milinsky and Heller, 1978) most studies are very uniform in their basic assumptions. The goal is usually assumed to be maximising the rate of energy (or food) intake, and models using this assumption have been remarkably successful in predicting the decision rules of avian and other predators both in the laboratory and in the field (Cowie, 1977; Krebs et al., 1977 ; Smith and Dawkins, 1971; Smith and Sweatman, 1974, to quote only work on the great tit).

The models have dealt with various aspects of predator behaviour: choice of prey items (Werner and Hall, 1974; Charnov, 1976a; Krebs et al., 1977); allocation of foraging effort to different areas or patches (Smith and Dawkins, 1971; Smith and Sweatman, 1974); group size (Barnard, 1979; Pulliam, 1976) and search strategy (Smith, 1974 a,b).

One of the limitations of existing models is that they deal with steady state behaviour in stable environments, in which the predator behaves as if it has estimated the availability of different prey types or patches, in order to calculate its average net rate of intake. As several authors have pointed out (Oaten, 1977; Green, 1979; Oster and Heinrich, 1976) this is an unrealistic feature of the models. More often than not predators are faced with the problem of sampling the environment in order to determine the relative availability of different types of food and the profitability in different patches.

In many situations there might be a conflict between the most efficient way of exploiting the environment according to the current estimate of food distribution and exploration of new potential sources of food.

6.2. The Memory Window

I shall illustrate the need of sampling with an example based on the Marginal Value Model (Charnov, 1976b; Parker and Stuart, 1976), one of the most successful developments of the optimal foraging approach in recent years. This model predicts the behaviour of a predator hunting in a patchy environment where the food within patches is subject to depletion. From measurements of the travel time between patches and average intake rate as a function of time within a patch it is possible to predict the optimal time that the predator ought to spend in each patch. Cowie (1977) tested the predictions of Charnov's model in an elegant laboratory experiment using great tits, finding good

agreement between the predictions of the model and the behaviour of the birds. Charnov's original derivation, and Cowie's experiments, apply to cases where (a) the depletion process is continuous and deterministic; (b) the quality of the environment as a whole does not vary with time and (c) the predator knows at every instant the rate of capture in the patch and the average quality of the environment. Cowie was aware that these assumptions pose problems in more realistic, unstable and stochastic environments, and proposed the idea of a 'memory window' as a mechanism by which predators could update their assessment of the habitat. The memory window consists of a sliding average over a given period for both travel time between patches and capture rate within patches, the length of the window being a function of the variability of the environment in time, the diversity in patch quality and the variance in inter-capture interval within a patch. The idea is ingenious in that no sampling is postulated and yet the predator's behaviour can track the changing parameters of the environment. But it is a conservative policy.

Imagine that a new type of patch became available, in which the interval until the first capture is twice as long as the current average, but the prey are three times as big as the current average. The inclusion of that type of patch in the accepted set would increase the average rate of energy intake, but an animal which used a "giving up time" (Krebs et al., 1974) based on a memory window mechanism might always leave the new patch type before capturing any prey and might remain 'caught' in its old

diet. It would require a more inquisitive animal, one that sometimes stayed longer in a patch than appropriate to maximise momentary rate of intake as assessed from experience, to discover the advantages of the new type of patch. This animal would be sacrificing expected immediate gain in order to acquire information: it would be 'sampling' the environment. The problem of the trade off between maximising momentary expected rate of intake and exploring the environment constitutes the subject of the second part of this thesis. I will centre my investigation around the question: How should an efficient predator divide its effort between exploring and exploiting? Formulating the question in this way differentiates my approach from that of most previous workers in optimal foraging, since the traditional question has normally been: How should a predator behave to maximise its instantaneous rate of intake?

I will be treating the problem of sampling using a very simplified environment, with only two non-depleting patches with equal sized prey, differing in intercapture interval in a way unknown to the predator.

6.3. Psychological studies of concurrent schedules

While the main aim of these chapters is to study the transient period while behaviour is being tuned to the properties of the environment, some attention will also be paid to the pattern of steady-state behaviour once a stable equilibrium response is reached. The optimal foraging prediction for the steady state in such a stable environment is trivial: if the predator knows

the sign of the difference in food availability between the two patches, the most profitable behaviour is to devote all foraging time to the better patch. The question of equilibrium behaviour is interesting though because it has been studied very extensively by experimental psychologists, with very conflicting outcomes (Bitterman, 1965; Shimp, 1969; Herrnstein and Loveland, 1975; Staddon, 1979a,b). The experimental paradigm used in the next chapters is similar (but not identical) to the paradigm used in studies of choice in concurrent variable ratio schedules in Skinner boxes.

In a concurrent variable ratio (see Section 1.9.2) an experimental subject is allowed to operate 2 instrumental devices, each of them delivering a food reward after a certain number of operations (responses). The number of responses required to obtain a reward is more or less randomly distributed around a mean value called the variable ratio (VR) characteristic of that alternative (Mackintosh, 1974).

Psychologists working with operant paradigms identify a food reward with a "reinforcement" and claim that choice behaviour aims at maximising the rate of reinforcement (Thorndyke's Law of Effect, see Hernnstein, 1970; Mackintosh, 1974). They arrive at similar prediction as those from optimal foraging models, i.e. they predict exclusive operation of the option with smaller VR. Bitterman (1965) triggered a long and heated debate by claiming that animals belonging to certain ("lower") classes of vertebrates (fish, birds) failed to achieve this optimum, but instead reached

a steady-state in which the proportion of responses in each side 'matched' the proportion of the probability of reinforcement in that side. More formally,

$$r1/(r1 + r2) = VR1^{-1}/(VR1^{-1} + VR2^{-1})$$

where r1 and r2 are the rate of response in side 1 and 2 and VR1 and VR2 are the respective variable ratios. In spite of the similarity in nomenclature, this situation is conceptually very different from matching in concurrent variable <u>interval</u> experiments (Herrnstein, 1961; Shimp, 1969; Heyman , 1979; Staddon, 1979) for which the so called Matching Law was formulated (Herrnstein, 1961).

In the variable-interval schedules access to food is determined jointly by time and operation of the manipulandum: a Vlt schedule means that the first response after (on average) t seconds since the preceding food delivery will procure food (Mackintosh, 1974; Staddon, 1979). When an experimental subject faces two independent concurrent VI schedules, it is obvious that after spending some time working exclusively in one alternative a reward is very likely to be available in the other one, so that some switching is expected to occur if the subject is to maximise the number of rewards obtained. Many studies using concurrent VI-VI paradigms found that the proportion of responses delivered within a test to each of the alternatives matches the proportion of reinforcements earned there (Herrnstein, 1961; De Villiers, 1977; Baum, 1974), but there is no agreement as to the mechanism underlying matching. Shimp (1969) claimed on the basis of a simulation study that matching would result from the animals

always choosing the alternative which has at the moment the greatest probability of reinforcement, while Rachlin (1978) offered the view that matching would result from maximising the sum of the reinforcements obtained from each alternative, a conclusion similar to that of Staddon and Motheral (1978).

Heyman and Luce (1980) challenge this conclusion on the basis of a theoretical model which shows that maximising overall reinforcement rate does not necessarily lead to matching. Heyman (1979) showed that a second order Markov model description fitted the behaviour of pigeons working on a concurrent VI-VI. In his model, the probability of delivering a peck to one of the options did not decrease as a function of having obtained a reward in that option before. Since the probability of reinforcement does drop after obtaining a reward (the timer is reset), his model is not compatible with momentary maximisation. It is important to point out the difference between molar matching (matching the proportion of responses to the proportion of rewards in each alternative) and molecular matching (matching the probability of each individual response to the proportion of reward rates in each alternative). While molar matching is an accepted empirical result, most authors see it as the outcome of molecular maximising rather than molecular matching.

Myerson and Miezin (1980) provide the first model of molecular behaviour in concurrent schedules which predicts successfully features of the molar behaviour in both transition and steady-state environments.

The relation between matching and maximising is still an open question for concurrent VI-VI schedules, but there is general consensus in the literature that matching is definitely not equivalent to maximising in concurrent VR-VR schedules (Herrnstein and Loveland, 1975; Lea, 1979; Staddon, 1979). I used variable ratio schedules because they are more realistic to simulate great tits foraging in the wild (Section 1.9.2 and Myerson and Miezin, 1980).

It is interesting that these developments within the psychological literature have the same limitations discussed for optimal foraging models: they assume stable environments which are well known to the subject. If the subject cannot assume that the two alterantives have perfectly stable properties, then it may be optimal in concurrent VR-VR paradigms to depart from total commitment to the better alternative as a safeguard against a change in relative quality of the two sources (Oster and Heinrich, 1976).

6.4. The Two Armed Bandit

The maximisation strategies for situations very similar to concurrent VR-VR have been discussed for some years now in mathematical literature under the name of Two Armed Bandit problems (Thompson, 1933; Bellman, 1965; Jones, 1976; De Groot, 1970; Wahrenberger et al., 1977). The name of this category of problem is taken from the gambling situation where a player puts money in a slot machine known as a two armed bandit. The machine has two slots (arms), and on each play, the gambler must decide whether to put his/her money in one arm or the other. As with the con-

current VR-VR paradigm, if the subject knows the reward probability in both arms, the maximisation strategy simply consists of putting all the available money in the arm with a higher probability. But less trivial solutions can be found when the gambler (a) does not know initially the two probabilities, and (b) has a limited amount of money to play. Two armed bandit problems were first considered by Thompson (1933), and since ther a great variety of related problems have been described. Variations of this theme are problems where the player knows the reward probability of one of the arms but not the other (one-armed bandit), where the magnitude of the reward as well as the probability of getting it is variable, and when there are more than two arms (K-armed bandit). While the most realistic model of the problem that confronts a predator is probably a k-armed bandit with variable magnitude of reward and probability in each arm fluctuating with time, this type of problem is very hard to solve. In order to be able to formulate testable predictions about the optimal strategy I shall confine both experimental and theoretical discussion to a restricted case of two armed bandit problem, which is fully explained in the next chapter. McNamara and Houston (ms) provide a general theoretical discussion of the relevance of statistical decision making to foraging behaviour.

One of the problems that will be recurrent in the next three chapters is the question of which is the variable being maximised. This is a traditional problem in biological optimisation studies (Schoener, 1971; Maynard Smith, 1978) but it is

particularly relevant to problems of optimality under conditions of uncertainty. In a two armed bandit problem where the gambler does not know the reward probability of either of the two alternatives maximising the probability of win for every coin leads to a behaviour which is different from that of maximising the sum of the wins for the total amount of money available. The reason is that a gambler playing the second type of game would do best by investing a certain amount of money in acquiring information, even if this implied an initial sacrifice of some coins by playing them in the arm which is currently believed to be the worst. To put it in the language of experimental psychology: In a two armed bandit problem molecular maximising is <u>not</u> identical with molar maximising.

6.5. <u>Summary</u>

In this section of the thesis, my aim is to discuss the foraging behaviour of predators in the framework provided by, (a) the optimal solutions derived from the statistical decision-making literature and (b) the psychological mechanisms and experimental paradigms known by the psychological authors. In Chapter seven I will expand the description of the two armed bandit problem as applied to my particular experimental situation, present an algorithm aimed at finding the global maximisation policy, and compare the predictions of the algorithm to the behaviour of great tits. Chapter eight is purely theoretical. I will present a series of computer simulations of potential rules-of-thumb in an attempt to relate the difficulties of

reaching an optimal solution to the computing capabilities which is reasonable to expect from an avian predator. Finally, Chapter nine is a return to experimentation. In it I describe an experiment aimed at testing directly if great tits foraging in an unknown environment modify the effort invested in sampling according to the predictions of the optimal algorithm introduced in Chapter seven or if, alternatively, their behaviour is more consistent with molecular maximisation.

CHAPTER 7

An experiment on optimal sampling

7.1. Introduction

In this chapter I describe an experiment aimed at testing if great tits choose an optimal balance between exploration and exploitation when foraging in a new environment.

The experimental set up is extremely simple. The naive bird is faced with a choice of two foraging places (patches) which differ in availability of food. Availabilities remain constant throughout a particular experiment, and the bird can only discover by exploration which of the two patches is better. Once the bird has distinguished between the quality of the two patches, it should concentrate on exploiting the better one, but in this experiment I am concerned with how the bird makes the decision about which patch to exploit. I suggest that there are two simple types of maximising rule. On the one hand, the predator might maximise its rate of food intake at every instant in time by always foraging in the patch with the higher expected reward rate. Alternatively, the predator might attempt to maximise its intake over the total foraging time and sacrifice short term gain in order to acquire more information about the relative quality of the two patches. In this second strategy, an efficient predator has to choose the optimal trade off between exploration and exploitation. My main aim was to test whether or not great tits achieve this optimal trade off.

7.2. Methods

7.2.1. Foraging Patches

Two patches, each one consisting of a feeding disc

(Section 1.9) with its operating perch placed together on top of a platform 1.5 m high, were placed at opposite sides of the indoor aviaries (Section 1.8) 2.5 m apart and with a 'tree' standing exactly half way between them.

The bird could obtain rewards (segments of mealworms weighing 0.07 g each) from either disc by operating the perch next to it. Food items were delivered according to variable ratio schedules with certain probabilities after a perch operation (Section 1.9.2.). Eleven reward probabilities were used, ranging from 0.05 to 0.50 in steps of 0.05, excluding 0.25. An additional probability of 0.075 was used for special purposes (see "neutralisations" below).

The reward probability of any given patch was always constant while a bird was allowed to operate the disc, but I changed it in the intervals between these periods.

7.2.2. Experimental Sessions

Experimental sessions are referred to by means of the following nomenclature:

Runs: General name employed for all periods of uninterrupted access to the feeding patches. Runs started by opening the trap door separating the outdoor from the indoor compartments and letting the bird come in. At the end of a run the trapdoor was opened again and normally the bird would fly back to the outdoor compartment. There were 2 types of runs: tests and neutralisations.

Tests: Runs in which the two patches had different reward probabilities. Reward probabilities were chosen so as to make the sum

of both probabilities add up to 0.50 in all tests. Thus, there were 5 types of tests, with the following combinations of reward probability: 0.00 - 0.50, 0.05 - 0.45, 0.10 - 0.40, 0.15 - 0.35 and 0.20 - 0.30. Tests were stopped when either a treatment was complete or when the subject had taken all (72) rewards from the disc without reaching the criterion for ending a treatment.

Treatments: Series of one or more tests with constant probabilities. Treatments are the basic units of analysis of the results and are named according to the reward probabilities experienced (for example 0.15 - 0.35) or according to the difference in reward probability between the 2 patches, which constitutes the patch diversity (in a 0.15 - 0.35 treatment patch diversity is 0.20). Each bird experienced the 5 treatments twice.

Treatments were finished after the subject delivered more than 90% of a series of 100 consecutive hops in one of the 2 patches.

Neutralisations: Series of 1 or more runs when the reward probability was 0.075 in both patches. Neutralisations were employed between treatments to eliminate the bias in favour of one of the two patches before the new treatment started. A treatment was started when the subject showed a bias of less than 60% for any one patch during a neutralisation run. The reward probability of 0.075 was chosen because it proved to be effective during the pilot period.

7.2.3. Decision criteria

The last hop before the beginning of a block of 100 with a bias of more than 90% was called the decision point. This criterion for determining the decision point was chosen after using pilot tests to compute the reliability of criteria of various lengths to be used as predictors of steady state behaviour. It was found that the 90% - 100 hops criterion was 95% reliable in that less than 5% of the tests that reached that criterion reversed to show preference for the other patch. Reversal was arbitrarily defined as a bout of at least 20 consecutive hops in the opposite patch. Figure 7.1 shows the reliability of the different criteria tested.

7.2.4. Subjects

Nine adult great tits of both sexes were used. All animals were naive with respect to the disc but some had been subjects in other experiments in the same aviaries.

7.2.5. Pilot experiments

There was a long period while the experimental design was refined, reward probabilities were chosen and recording methods set up. Since the birds experienced great number of tests during this period but not in a systematic sequence, these data were used to define the decision criterion and establish the neutralisation technique but were not used in the analysis of the results.

7.3. The Model

The predator's problem of choosing the optimal balance between exploration and exploitation is similar to the classical two-armed bandit problem (see Section 6.4), in which the player

is faced by two machines, each with a characteristic, unknown reward rate on a random schedule, and the goal is to maximise the expected number of rewards in N trials. A numerical solution to this problem can be found with a dynamic programming algorithm which works backwards from the Nth trial, at which point the machine with the higher percentage rewards is chosen. The algorithm involves analysing at each stage, n, the payoff for two possibilities (exploring or exploiting) in terms of the known optimal behaviour at state n + 1.

My model is a modification of the optimal solution of the two-armed bandit problem (Lindley and Barnett, 1965; Jones, 1975). It considers that the bird had a total of N trials (this value N is called the time horizon), and at the beginning of an experimental session it samples each patch alternately. After M trials in each patch the bird commits itself to one patch or the other, so that all strategies have the form: M trials in each patch (exploration) followed by N-2M trials on one patch (exploitation). In my model I introduced the constraint of equal sampling effort in the two patches during exploration because this simplifies the mathematical treatment. The model predicts how the value of M, the point at which the predator decides to exploit, changes with p_1 and p_2, the reward probabilities in the two patches. Intuitively, one would expect M to decrease as $p_1 - p_2$ (patch diversity) becomes larger, and the model predicts the exact shape of the curve.

The dynamic programming algorithm uses a Bayesian approach to describe the information possessed by the bird at each stage.

The essential feature of a Bayesian approach to decision making is the definition of a special probability distribution called the prior distribution, or simply the 'prior'. When a Bayesian subject faces a decision problem whose optimal solution requires the knowledge of a certain parameter g, it starts by formulating an hypothesis about the probable values of g. This hypothesis is previous to any observation, and it is normally based on both previous experience and pragmatic reasons of the sort discussed below. The prior results from assigning a given probability to all possible values of g. In intuitive terms: a prior describes the subject's degree of belief in possible parameter values before any observation is made, it is an initial hypothesis about the 'state of nature' (Silvey, 1970).

The Bayesian subject then takes observations and formulates a posterior distribution, (the 'posterior') which is conditional on both the prior and the result of the observations. The observed results change the subject's degrees of belief in different parameter values by changing a prior into a posterior distribution according to the Bayes theorem (de Groot, 1970). The prior and the posterior, in spite of their special property of describing a subjective state in terms of degrees of belief, can be treated as standard probability density functions, and in consequence it is possible to define confidence intervals. To establish a 95% confidence interval of such a function is equivalent to stating that the subject believes that there is a 95% probability that the value g lies within a specified range. As the number of observations increases, the confidence interval of the posterior becomes

narrower, which is equivalent to stating that the subject becomes more confident about its belief. As said before, the choice of a particular prior is influenced by its 'plausibility' and by convenience, this last factor being normally identified with mathematical tractability. For example, unless there are strong contrary reasons, it is preferable to formulate a continuous than a discontinuous prior, and if the solution of the problem involves some repetitive process, it is convenient to choose a prior such that the incorporation of observations would change its mean and variance but not the nature of the distribution, so that the posterior can be treated as the prior in a new iteration of the process. One continuous distribution with the properties defined above is the Beta Distribution, and for these reasons I chose a beta distribution for the simulation below. While the choice of the prior is clearly not a very rigorous procedure, Silvey (1970) states that with a reasonable amount of sampling, considerable variation of the prior may not affect the posterior beyond the limits of convenience.

In my experimental situation each trial (or hop) can be seen as a binomial experiment, ending in either obtaining a reward or not. The prior then is a hypothesis about the value of the probability of being rewarded, itself a Bernouilli process where the parameter is the patch reward rate. For each new trial,

$$p = 1 - q$$

where p is the probability of reward and q is the probability of failing to obtain a reward in that patch. In my model, at the beginning of a treatment the bird's prior estimate of p is

considered to be β - distributed with parameters (α, β), and therefore with mean $(\alpha + 1)/(\alpha + \beta + 2)$ and variance $(\alpha + 1)(\beta + 1)/(\alpha + \beta + 3)$,* after n trials with r successes, the posterior estimate of p is also β - distributed with new parameters (α + r, β + n - r). Thus it is easy to keep track of the state of the bird's information about the two patches by updating p1 and p2 after each trial. In simulating the model, starting values of $\alpha = 0$, $\beta = 0$ were used for both patches, which means that the bird's initial estimate of the mean reward probability in the two is 0.50 with a uniform distribution between 0 and 1.0. Thus after n trials with a_i successes, the bird's posterior estimate of the reward rate p in patch i is $(a_i + 1)/(n + 2)$. I investigated the effect of varying the prior, by calculating the predictions for the following values of α and β : $\alpha = 0, \beta = 2$; $\alpha = 0, \beta = 0$; and $\alpha = 2, \beta = 0$. The model was rather insensitive to changes in the prior, most of the changes in prediction being between 10% and 20%.

If after n trials in each patch, the bird decides to sample once more, its expected payoff for the remainder of the foraging time is E_S, which can be calculated by adding the probability of getting rewards in the next 2 trials (0 or 1 in each patch, with the given probabilities) plus a term that estimates the chance that the probability assigned to each patch may be modified by including the outcome of the next 2 trials and multiplies the newly assessed probabilities by the number of trials

---

* DeGroot (1970) uses a slightly different but equivalent notation. DeGroot's α and β are equivalent to $(\alpha + 1)$ and $(\beta + 1)$ respectively in my notation.

left minus 2. If it decides to exploit, its payoff is E_D, calculated as the number of trials left multiplied by the reward probability experienced in the better patch. The optimal strategy is to choose the larger (E) of these two values E_S can be calculated in terms of the values of E at stage n + 1 as follows:

Let

$$\pi_i = (a_i + 1)/(n + 2)$$

Then

$$E_S(a_i,a_2,n) = \pi_1 + \pi_2 + \pi_2\pi_1 E(a_1 + 1, a_2 + 1, n + 1) +$$

$$\pi_1(1 - \pi_2)E(a_1 + 1,a_2,n + 1) + (1 - \pi_1)\pi_2$$

$$E(a_1,a_2 + 1, n + 1) + (1 - \pi_1)(1 - \pi_2)$$

$$E(a_1,a_2, n + 1).$$

and

$$E_D(a_1,a_2,n) = (N - 2n) \max (\pi_1,\pi_2)$$

Figure 7.2 presents a flow chart of the simulation. A version of the model implemented in Algol 60 is presented in Appendix A.

The procedure in simulating the model is to start with 2n - N and work backwards, generating E_S, E_D and E recursively. At each stage, the model produces an n x n matrix giving the value of $(E_D - E_S)$ for each combination of values of a_1 and a_2. Figure 7.3 presents an example of one of these matrices and Figure 7.4 shows a simplified version of the matrices at 3 stages of the simulation. When the bird is at the last pair of trials, E_D is clearly higher than E_S, assuming there is a difference in the two

reward probabilities. This is because at this point there is no future left, and the loss of expected pay-off resulting from getting the average probability of reward in the next 2 hops can not be traded against a correction of the assessment to improve later performance. The matrices produced by the model are decision matrices. If a subject who has tried 2n times (n trials in each patch) wants to know if it would pay to sample once more or not, it has to find its row (number of rewards in patch A) and its column (number of rewards in patch B) in the matrix of stage $2n$ and read if the corresponding element of the matrix is negative (sampling) or positive (deciding). In the matrices in Figure 7.4 the model prints (≠≠) when deciding would pay and (?) when the opposite is true. As the simulation works backwards, it becomes more likely that, for any particular combination of a_1 and a_2 the payoff for sampling once more in each patch (and hence acquiring more information about the profitability of the two patches) is higher than the payoff for deciding to exploit. The exact point at which this change from sampling to decision occurs depends on the value of the total number of trials N, which in effect represents the bird's time horizon. N would be infinite in a completely stable environment and close to zero in a highly fluctuating habitat. For given values of a_1 and a_2, the larger the value of N, the longer the bird should sample before deciding.

The model was run with a range of values of N. This model only approximates the optimal solution, because of the constraints in the shape of the behavioural strategy: equal sampling followed

by irreversible commitment to one side. Jones (1975) describes an algorithm for the optimal solution, allowing the subject to sample unequally in the 2 patches and to revert a decision if necessary. In this solution when the model reaches (backwards) the stage where the subject is meant to have done n trials, it is necessary to consider all the possibilities of the distribution of the n trials between the 2 patches, as well as the possible proportion of rewarded trials in each patch, as does the model described above. This unconstrained algorithm naturally operates in 4 dimensions instead of 2 as the restricted solution, and the amount of computer core space and simulation time required is such as to make its simulation impracticable for reasonably long horizons of the order of hundreds of trials. (The algorithm uses matrices with $(\frac{N}{2})^4$ elements.) The quantitative approximation of my restricted solution is likely to be very close to the optimal solution specially in cases where approximately equal sampling in both patches results from external constraints on the subjects behaviour.

Qualitatively, there is no reason to expect any serious difference in the response of both models to patch diversity or time horizon.

7.4. Results

The results of the experiments showed that the birds' followed the general pattern of behaviour assumed by the restricted version of the two-armed bandit model. There was an initial sampling period (before decision) during which the birds performed on average 42% of hops on the more profitable and 58% on the less

profitable perch. This proportion does not differ significantly from 50:50 and in consequence is consistent with the constraint of equal sampling in the model. The slight bias towards the less profitable perch may have resulted from the experimental procedure of always starting with the higher reward rate in whichever perch was used less in the preceding neutralization test. After the decision point the birds spent virtually all their time on the more profitable perch; on average 95% of hops were on this perch, and the range for individual treatment means was 88% to 99% (Figure 7.5).

In terms of the optimal decision model, a hop can be considered as being equivalent to a single trial, so that we can predict, for each combination of reward rates, how many hops the bird should do before making the decision to exploit. Figure 7.6 shows the observed mean decision point in each experimental treatment and the predictions of the model for 3 values of N (total number of hops in the experiment) 250, 150, and 50. The fit to the model with N = 150 is good: not only is the observed curve the same shape as predicted, but the predicted values are all fairly close to the observed means. The other values of N give curves of the same general shape, but with predicted values too high or too low. The value of N which fits the data (150) is the modal value for the number of hops performed in a test (Figure 7.7). This suggests that birds are optimising over the "time horizon" which they experienced more often during the experiments but this experiment does not say if the birds would optimize with a bigger value of N if they were used to longer tests. This possibility is tested in Chapter 9. There is a slight difference

in the modal value of test lengths for different treatments, but since the treatments were presented in a random order there is no reason to suspect that the birds could modify their behaviour according to treatment.

One way in which the birds' behaviour differed from that implied in the model is that they did not sample by doing one hop on each perch at a time, but instead they hopped in bouts of about five. (The overall mean bout length before decision was 5.00, the range was 3.00-16.5.) This makes no difference in predicting the <u>average</u> decision point, but if the birds had hopped in rigidly determined bouts, the predicted decision point for a particular experiment would have to be a multiple of 5, which would involve reading every fifth matrix of the model's output. This tendency to hop in bouts probably explains why the observed points are above the predicted at the right hand end of the graph where the bird should make a decision after about ten hops. The model does not include any cost of travelling. If included, it is likely to give some advantage to greater persistence in a patch before switching.

The alternative goal considered in the introduction was the strategy of "immediate maximising" in which the bird always switches to the perch with the higher current expected reward rate. I used the β - distribution and the same prior values of α and β as above to predict, for any given experiment, at which point the birds should no longer switch back and forth if they were following this strategy. This predicted point is reached when expected pay off in the current perch never drops below that expected in the

other one. The observed point was based on the decision criterion discussed earlier. The expected number of switches depends on the actual sequence of obtaining rewards in the first few moments of the experiment. In calculating the expected values I made the conservative assumption that the birds would continue to switch until the difference between patches in expected reward rate was greater than 5%. The observed mean values for the five treatments are shown in Table 7.1. If we relax the assumption that the birds cannot discriminate differences of less than 5% in our experiments, the difference between observed and predicted becomes even greater. Clearly, the birds switch more than expected. In other words, they are not maximising in the short term (Molecular maximising), but instead they behave as if they acquire information in order to achieve a longer term optimum (Molar maximising).

Although the birds do not seem to switch according to the rule of immediate maximising, there is some evidence that switches between perches are related to short term changes in reward rate. Figure 8 is a plot of the number of unsuccessful hops before a switch during the sampling period, as a function of reward rate. The points lie very close to, but slightly above, the random departure curve. This random curve takes into account the fact that the birds work in bouts rather than single hops. Eight out of the ten means lie just above the random curve (p 0.05). This shows that birds have a slight tendency to switch after a run of bad luck. In other words they have a weak component of "immediate maximising" in their sampling strategy.

The bird's behaviour approximates the prediction of a model based on a constrained optimal tradeoff between exploration and exploitation while they do not fit the model of pure exploitation. The fit to the former model is close in spite of the fact that the model does not directly incorporate a cost of switching between patches. In effect I incorporated a small cost by using the rule that if the payoffs for sampling and exploiting were exactly equal, then the bird should exploit. In the experiment this cost was small both in energetic terms and in time. However in further developments of the model the cost of switching might be an important consideration.

I have presented averaged data for the 9 birds tested, but individual variation was very marked. I considered the possibilit that some individuals are more cautious (do more exploring) but are less likely to make mistakes. Five of the nine individuals tested made one or more mistakes (i.e. locked on to the poorer patch) in the 30:20 or 35:15 tests. These individuals did not make decisions after fewer hops than the other four birds. The average number of hops before a decision for all treatments was 23.9 ± 9.6 and 18.4 ± 9.8 for the fallible and infallible individuals respectively. However, by examining individual treatments in which a bird made an "incorrect decision", it is apparent that these tend to be cases in which the bird did little or no sampling Examining the 18 20:30 treatments (9 birds, 2 treatments each), only 3 out of the 13 in which a correct decision was made involved no sampling, while 2 out of the 5 which the bird made a mistake were tests in which there was no sampling period at the beginning.

Although the birds in the experiment behaved on average as if they were calculating an optimal tradeoff between exploration and exploitation, it is likely that they use some simple proximate rule of thumb to approximate the optimum. This issue will be discussed in the following chapter.

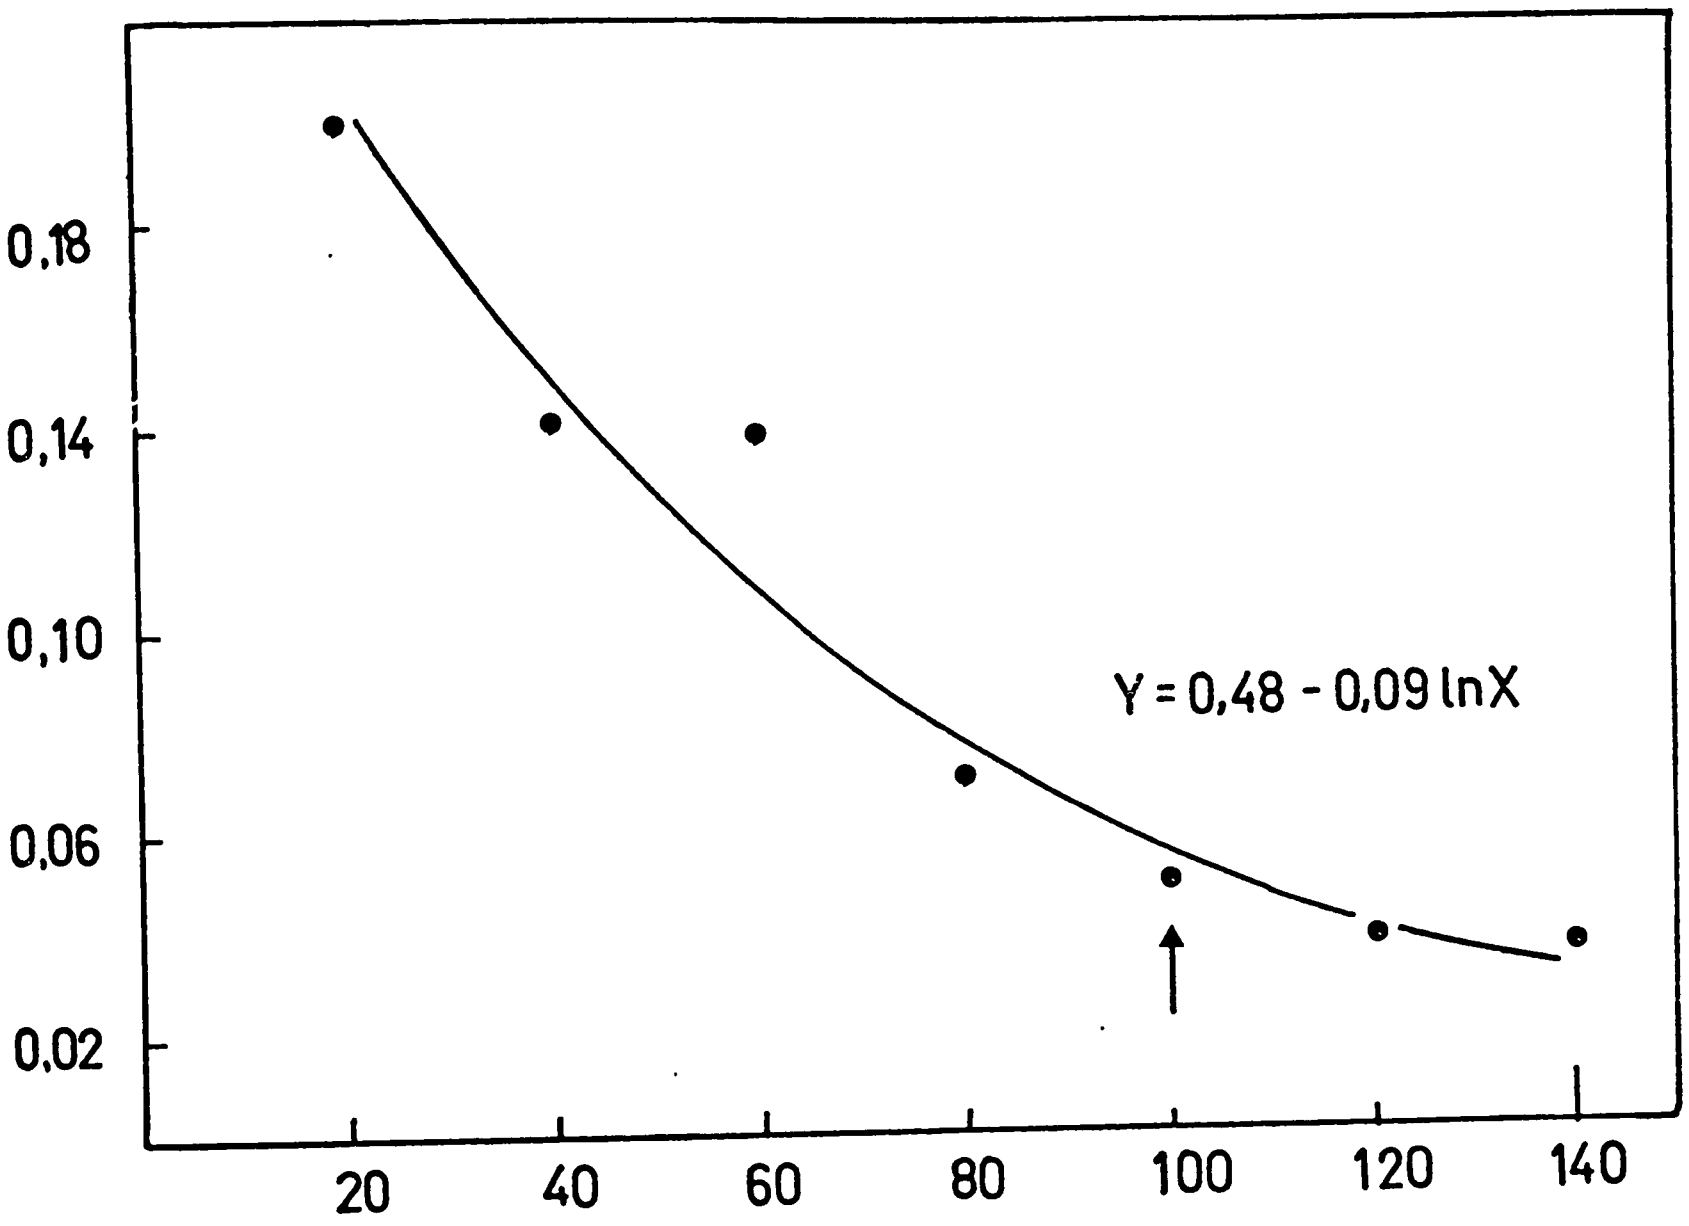
Table 7.1

Observed and predicted number of switches from one patch to another if the bird were operating on the rule of always going to the patch with the higher current expected reward rate.

Test	Predicted ± S.E.	Observed ± S.E.
50:0	0.93 ± 0.2	3.3 ± 0.91
45:5	1.3 ± 0.26	3.6 ± 1.56
40:10	1.36 ± 0.26	3.8 ± 1.1
35:15	1.3 ± 0.34	6.4 ± 2.5
30:20	3.25 ± 1.26	9.6 ± 3.5

Note: In predicting when the birds should have stopped switching, I assumed that they could not discriminate a difference of less than 5% in reward rate.

Figure 7.1



The probability that the bird will reverse its decision to stay in the more profitable perch after performing more than 90% of the last X hops in that perch. A reversal is defined as doing a bout of more than 20 hops in the minority perch.

Figure 7.2

Dynamic programming algorithm to solve the two-armed bandit problem under the constrain of equal sampling in both patches (CONOP).

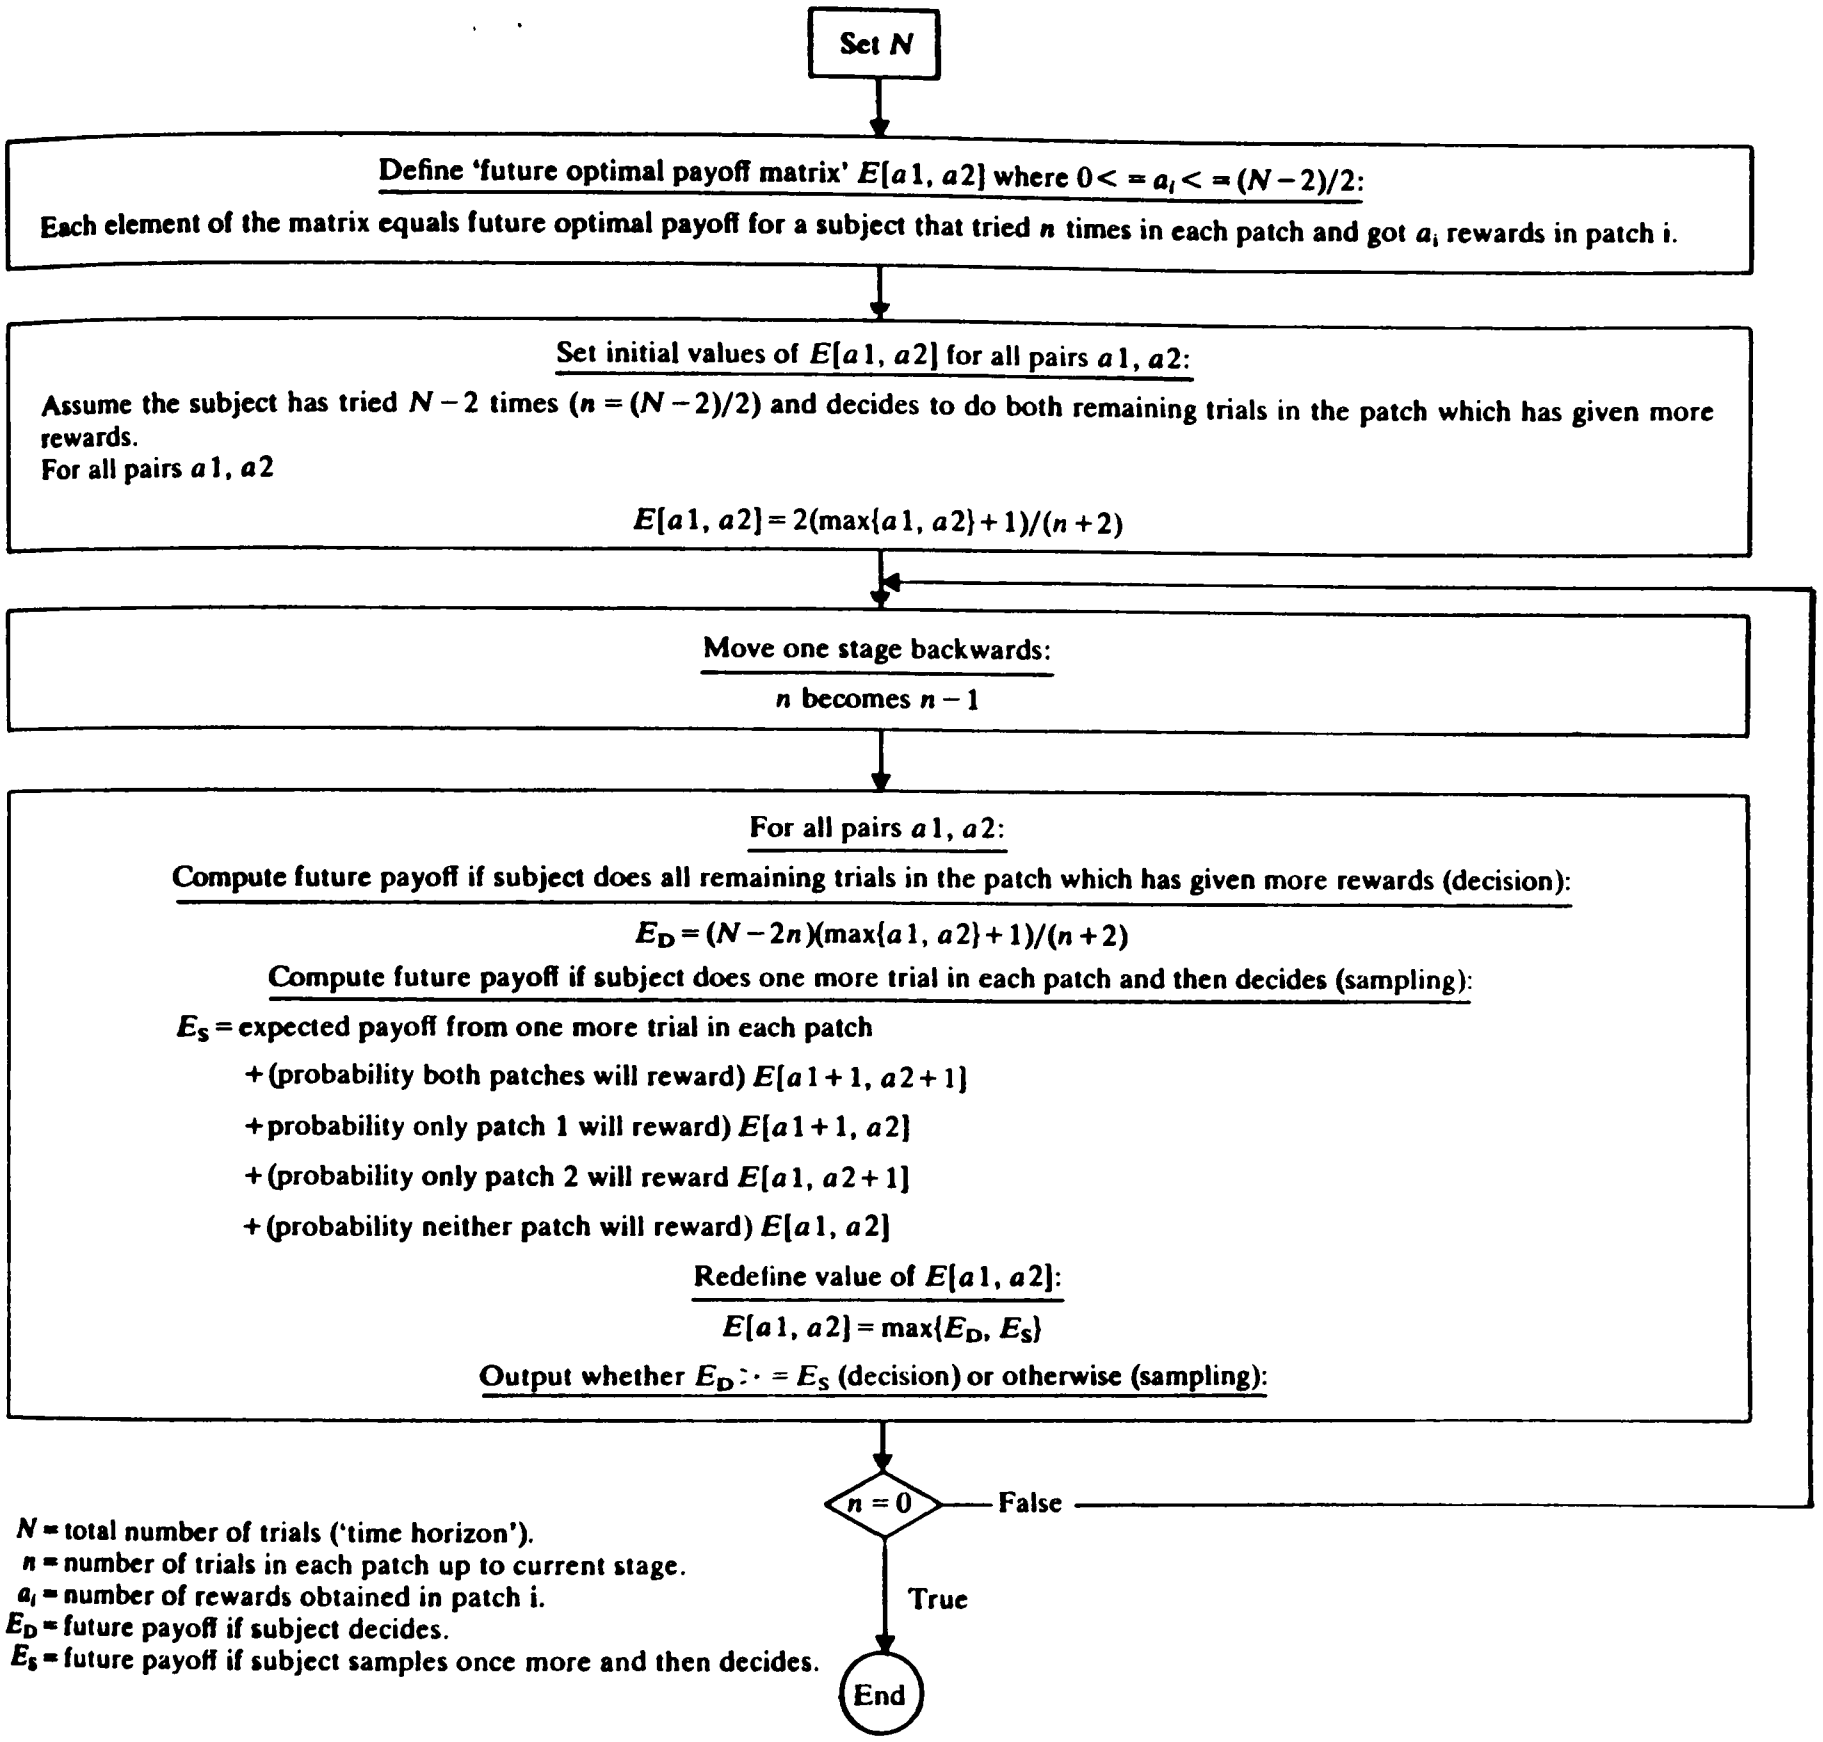
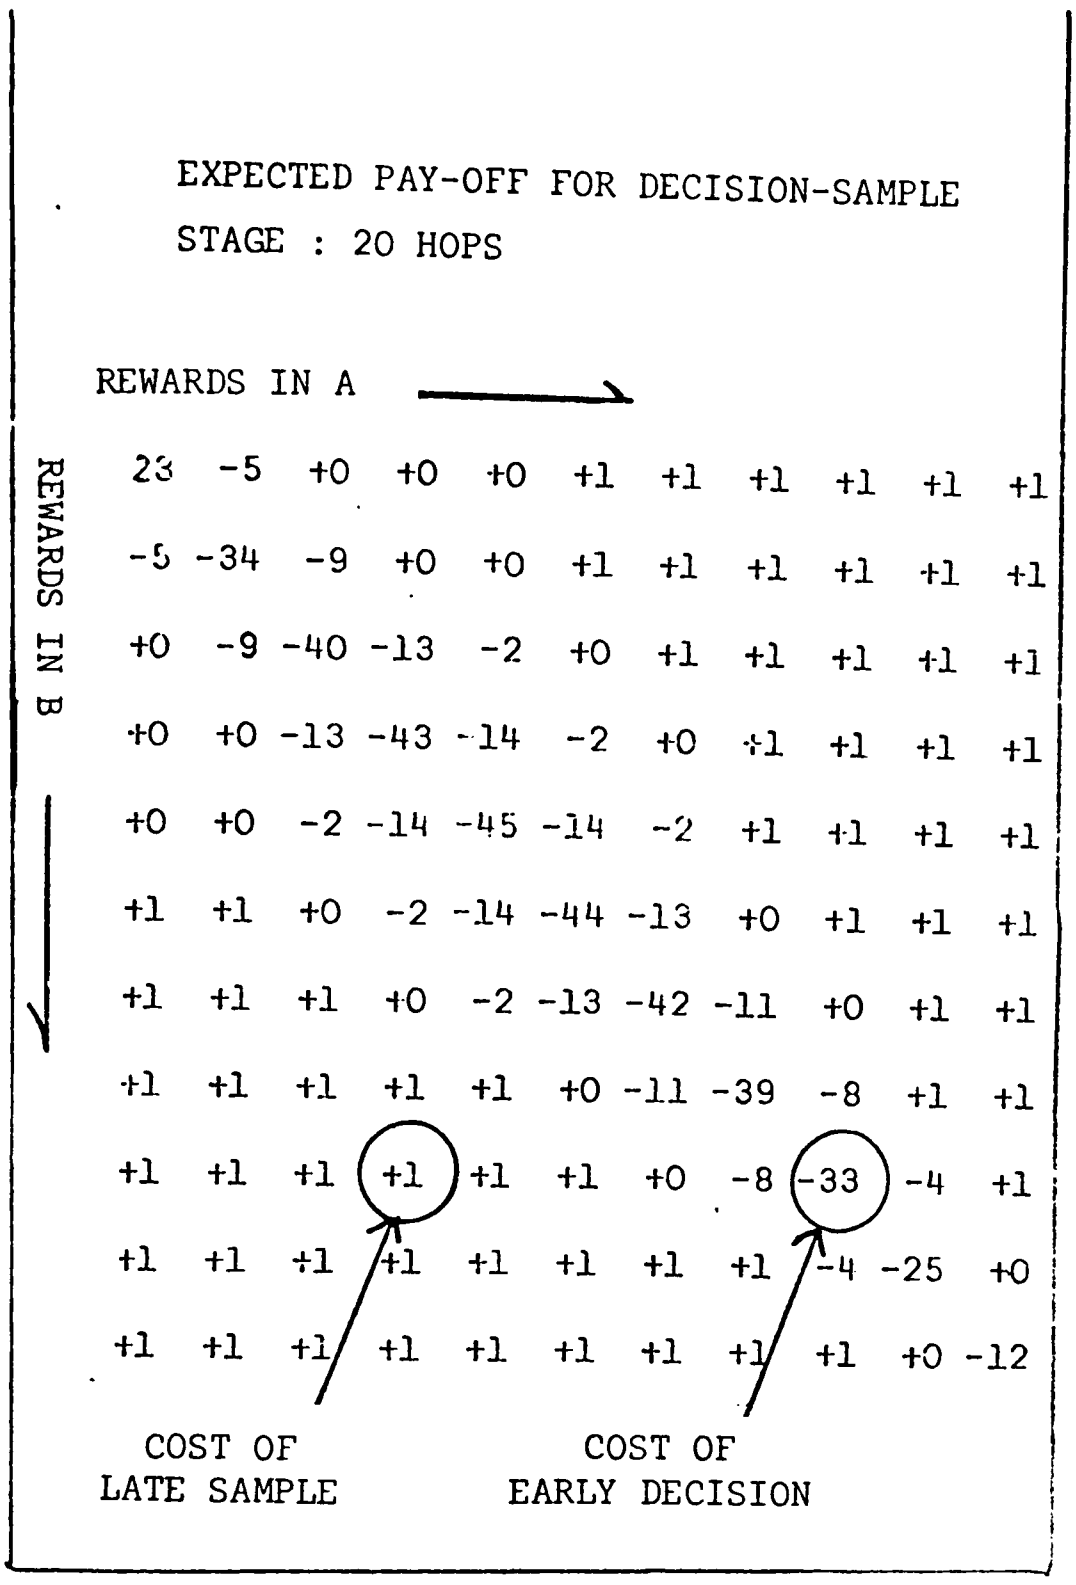
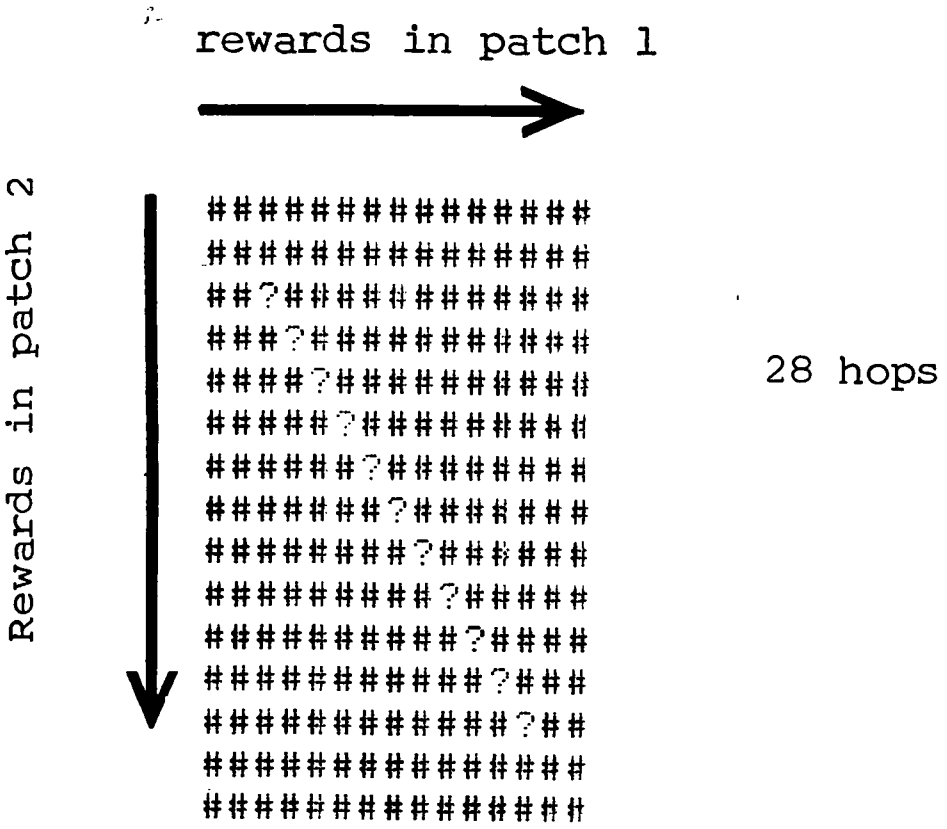


Figure 7.3



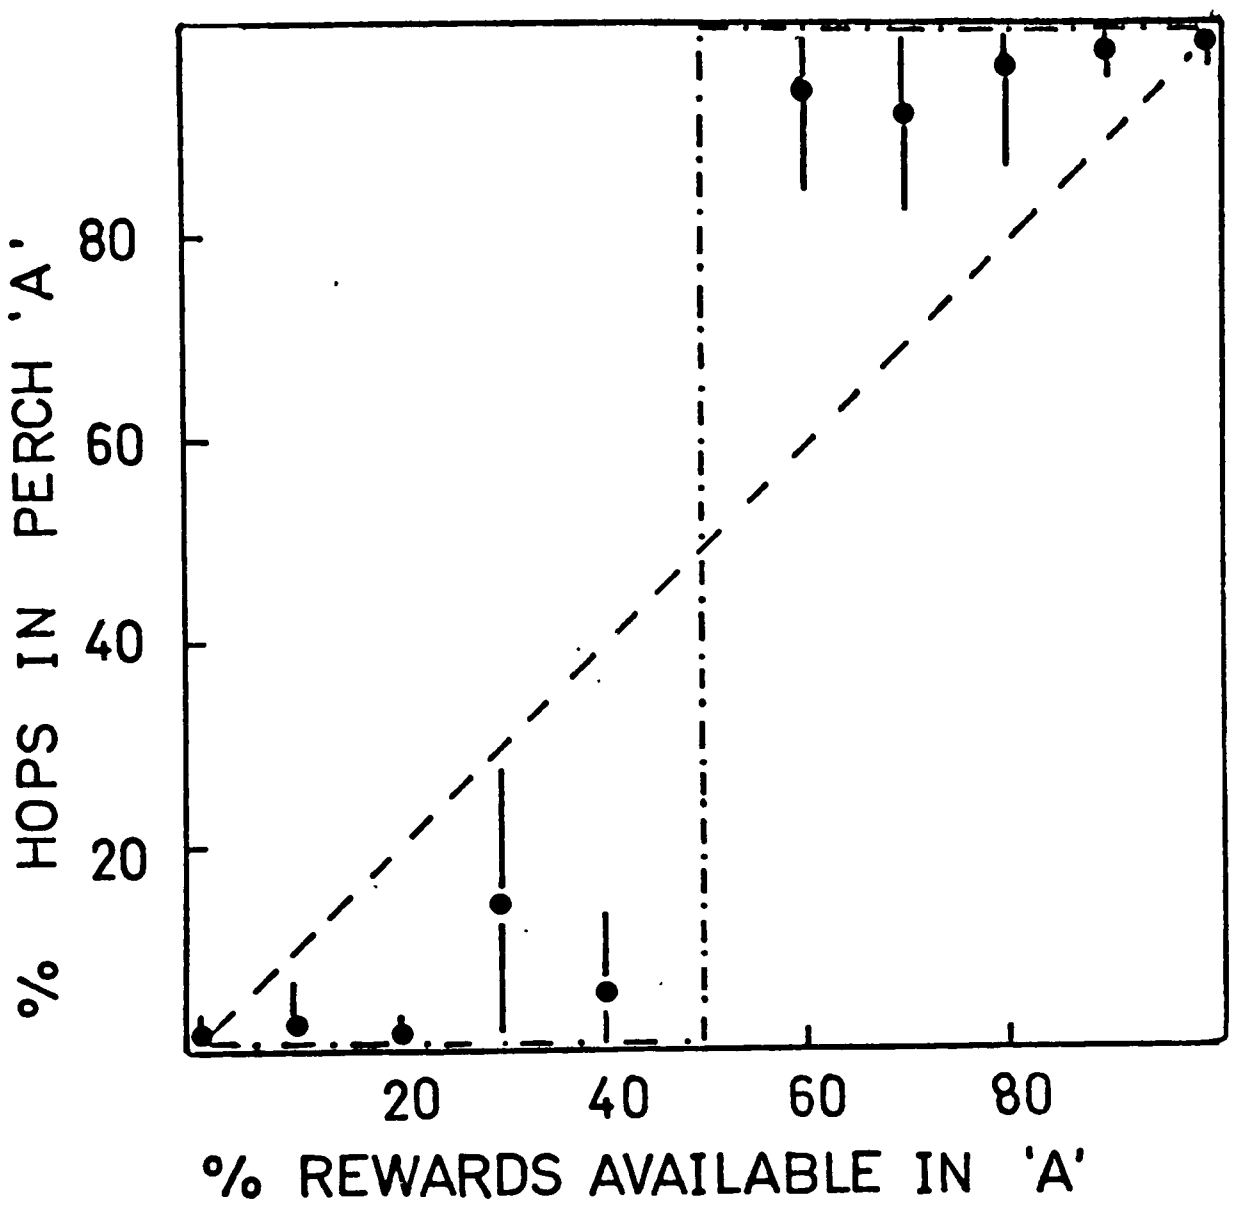
Example of decision matrix generated by the constrained optimal algorithm. Negative numbers are combinations of rewards where it would pay to keep sampling (ES ➢ ED). Positive numbers indicate that the probable pay off for deciding would be higher.

Figure 7.4



Simplified output matrices of the constrained optimal model (CONOP). Question marks indicate reward histories after which the best policy is to sample at least once more. For all other combination of rewards the optimum is to decide in favour of the patch where more rewards were obtained. Question marks are equivalent to negative numbers in Figure 7.3.

Figure 7.5



The proportion of hops in perch A (one perch arbitrarily was designated as A) plotted against the proportion of rewards available in A. The birds do not follow a 45^{o} line, but approximate a step function.

Figure 7.6 The predicted and observed number of hops before making a decision. Three predicted curves are shown, for N=50, 150 and 250. The observed means are close to 150 (continuous line) curve. Also shown are the 95% confidence intervals of the observed geometric means. The x-axis is the difference in reward probability between the two perches (=patch diversity) for the five treatments (e.g.0.50:0 - - patch diversity = 0.50; 0.35; 0.15 - patch diversity = = 0.20).

Figure 7.7 A frequency distribution of lengths of tests. The modal value is close to N = 150, which gave the best fit in Figure 7.6.

Figure 7.6

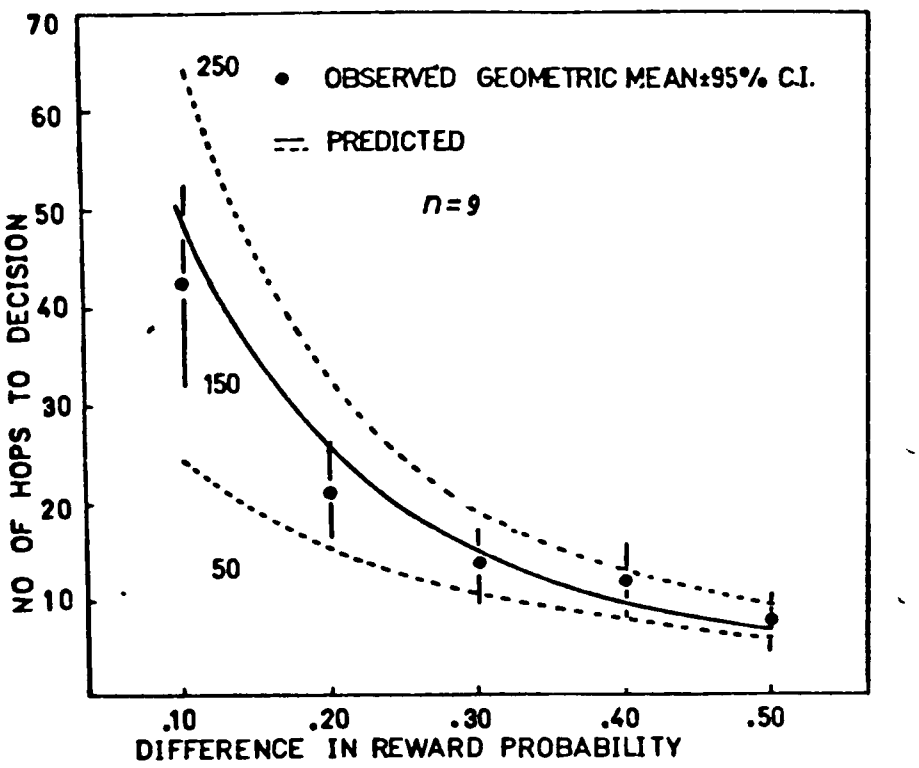


Figure 7.7

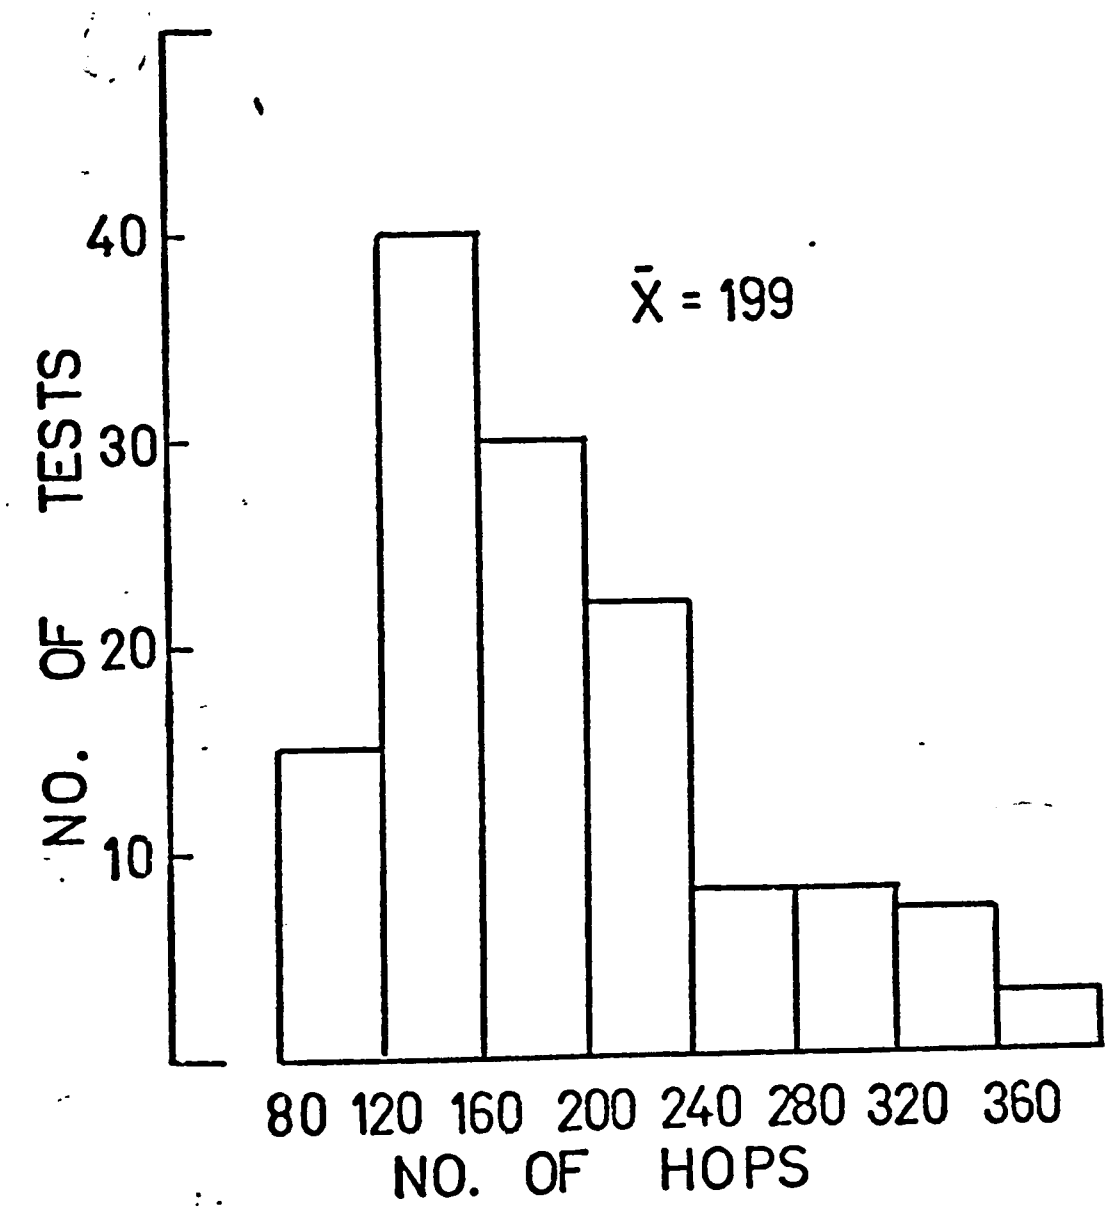


Figure 7.6 The predicted and observed number of hops before making a decision. Three predicted curves are shown, for N=50, 150 and 250. The observed means are close to 150 (continuous line) curve. Also shown are the 95% confidence intervals of the observed geometric means. The x-axis is the difference in reward probability between the two perches (=patch diversity) for the five treatments (e.g.0.50:0 - - patch diversity = 0.50; 0.35; 0.15 - patch diversity = = 0.20).

Figure 7.7 A frequency distribution of lengths of tests. The modal value is close to N = 150, which gave the best fit in Figure 7.6.

Figure 7.6

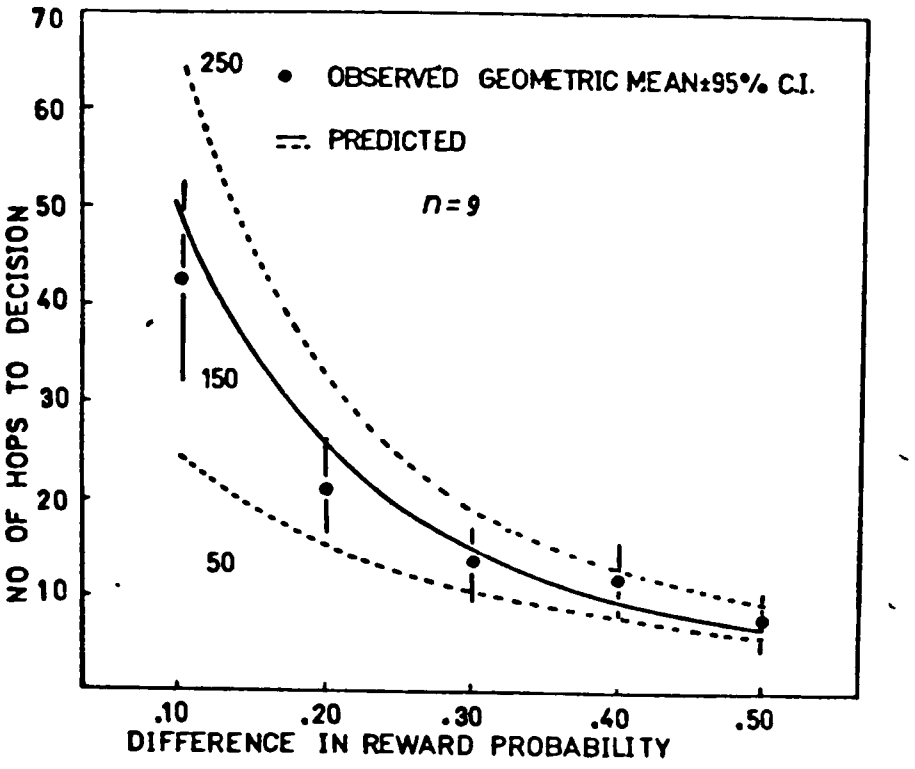
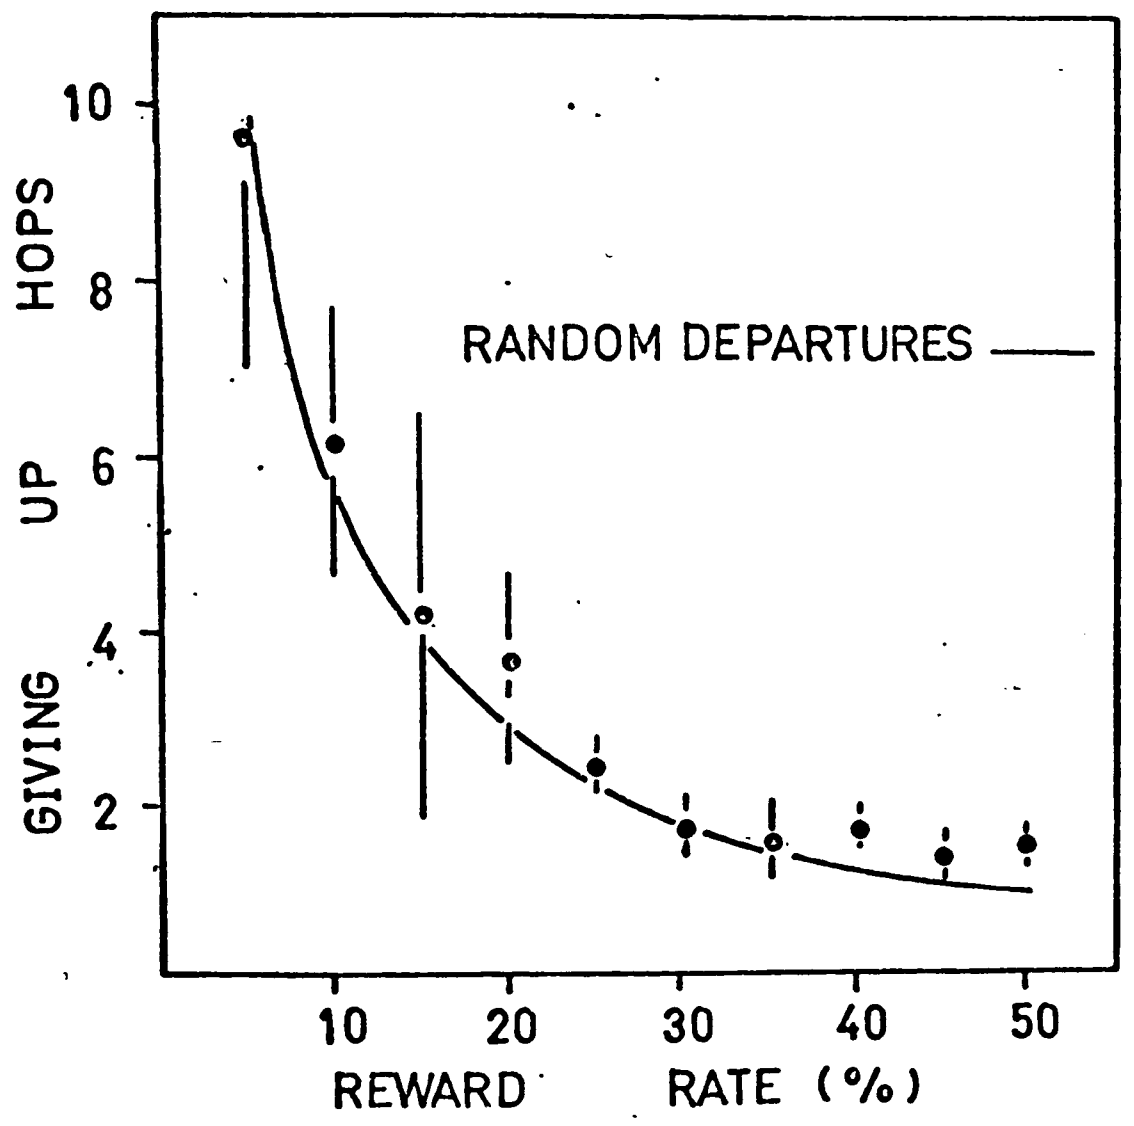


Figure 7.7

Figure 7.8



The observed number of "giving up hops"- (unsuccessful hops before a departure) plotted against reward probability. The random curve is calculated taking into account the fact that the birds tend to hop in short bouts. If p is the probability of getting a reward on a single hop and q = (1-p) then the expected number of giving up hops for bouts of 5 = $qp+2q^3p+3q^3+4q^4+5q^5$. The observed points lie very slightly, but consistently above the random departure line, indicating that the birds tend to leave a perch after a run of bad luck. The observed points are $\pm$ standard error.

CHAPTER 8

Rules of Thumb for the Two Armed Bandit Problem

8.1. Introduction

In Chapter 7 I used a dynamic programming algorithm which required a large amount of computer time and memory space to find a near optimal solution to the two armed bandit problem. The simulation had to make use of a matrix with a number of elements equal to the square of half the time horizon (Figure 7.2) For the time horizon like the one experienced by the experimental animals (150 hops) this meant that 5625 real numbers had to be 'remembered' to decide at every state whether to continue sampling or to switch to exploitation. And yet the birds approximated the optimal decision point very closely.

The complexity of the calculations involved in finding optimal solutions to problems involving stochasticity has been used as an argument agains optimality models. Oaten (1977) showed that the stochastic nature of the world may constrain the optimal solution of foraging strategies, and concluded that the solution of these problems was too difficult for animals to achiev McNamara and Houston (ms) make the point that this pessimistic view ignores the fact that rules-of-thumb may have evolved to produce behaviour that is close to optimal, and Green (1979) in an important development of Oaten's approach proved that a simple Bayesian rule can approximate the optimal behaviour of a predator foraging in a patchy stochastic environment where the predator's problem is when to leave one patch to move to another. My aim here is to show that the same is possible for a predator in the

conditions of my experiment.

The algorithm I used to obtain the near optimal trade-off between exploring and exploiting (referred to as CONOP) defines a balance such that no other strategy can obtain a consistently higher pay off (I used the word 'consistently' because it is possible to define more risky strategies that might do better than CONOP sometimes but would do much worse on other occasions). The procedure to obtain the near optimal solution is not necessarily unique. On the contrary, there may be rules of thumb which approximate the near optimal balance achieved by CONOP. These rules of thumb could never do better than CONOP unless the animal has previous information about the habitat.

Several rules for solving the two armed bandit have been proposed recently, both Bayesian rules (Wahrenberger et al., 1977; Jones, 1975, 1976) and otherwise (Fox, 1974), but none of them for conditions approximating my experiment. In this chapter I tested by simulation the performance of several rules in the same (modelled) situation as the great tits in the experiment. My aim is to investigate if simple rules-of-thumb can approximate the nearly optimal balance between exploration and exploitation calculated by CONOP as closely as did the birds. I shall also initiate the discussion of whether any of the rules tested may be a good description of the rule used by the experimental animals.

8.2. Methods

8.2.1. Simulation procedures

All models were simulated by implementing Algol 60 programs (Rogalgol 40 dialect) and run on a PDP-8F digital computer.

The simulations were based on the Monte Carlo method (described in Snedecor and Cochran , 1967), using a random number generator program to assign failures or successes to individual 'hops' of the model birds. The random number generator produced integers in the range -2047 to 2047 with a rectangular frequency distribution. Dividing the absolute value of these numbers by 2047 transformed the number generator into a random table of 2047 equally frequent numbers between 0 and 1. The probability of a number drawn from the table being less than a given number in the range 0 to 1 was exactly this last number. For example, there was a 0.30 chance of drawing a number smaller than 0.30. This property was used to decide if a hop was rewarded. When a model bird 'hopped' in a patch with nominal reward probability P, a number was drawn from the random table, and the hop was rewarded if the number drawn was less than P.

The models simulated the behaviour of birds in the experimental set up described in the last chapter, which consisted of two nondepleting patches each with reward probability unknown to the predator. In all the rules tested the location of the next hop was based on recent experience, but the rules differed in the way in which they used this experience. Given the stochastic element in the simulations, not all runs of the same model gave the same outcome. Each model was run 100 times and the average result was taken as the characteristic performance of the model. Figure 8.1 presents a flow chart of the general simulation procedure for a single run of a model.

8.2.2. The Rules

Four types of rules were tested. Each rule is identified by an abbreviation and a full name. Both are given below with a

description of the algorithm involved.

8.2.2.1. IMMAX: Immediate Maximising

This is a Bayesian rule. It starts with Beta distributed prior hypotheses about the value of the two probabilities, which are assumed to be equal. The same initial values of the parameters and of the distribution as in the dynamic programming algorithm (Chapter 7) were used. These were: $\alpha = 0$, $\beta = 0$ so that the mean of the priors is updated after every hop by means of the following equation:

$$p_i = \frac{\alpha + r_i + 1}{\alpha + \beta + n_i + 2}$$

where: r_i is the number of rewards obtained in patch i

n_i is the number of 'hops' in patch i

p_i is the estimate of the reward probability of patch i given the information available (subjective probability).

The first 'hop' is assigned at random and later ones to the patch with higher p. This rule is similar to the rules tested by Wahrenberger et al. (1977) except for the choice of prior. The main contrast with CONOP is that the emphasis is on momentary maximisation of reward probability, instead of global maximisation for the whole test. It further differs from CONOP in that it is easier to use. It requires the storage of 4 numbers: the number of rewards and of trials in each of the 2 patches, as opposed to the thousands of numbers required to use the optimal decision matrices produced by CONOP. Jones (1976) tested the efficiency of a very similar rule under the name of One-Step-Ahead, but only for extremely short horizons.

8.2.2.2. MATCH: Probability Matching

This is also a Bayesian rule, and the algorithm used to update the value of the two reward probabilities is the same as for IMMAX, but the allocationof the next hop to a patch is made stochastically instead of deterministically. The place of the next hop is determined using the following equation:

$$prl = \frac{pl}{pl + p2} .$$

Where prl is the probability of choosing patch 1 for the next hop. The next hop is made in patch 1 if a number drawn from the random number generator is lower than prl, and is made in patch 2 otherwise. If the estimate of pl is higher than that of p2, prl will be greater than half, but still there would be a certain probability of choosing the patch with lower expectation of reward This probability was nil in IMMAX.

MATCH seemed worth investigating because it has been found that under certain conditions fish and pigeons seem to match reward probabilites in the two alternatives (Bitterman, 1965; but see review in Chapter 6). This behaviour would clearly not approach the optimum when there is no uncertainty about the values of pl and p2, but it might conceivably be close to the optimum in a fluctuating environment.

8.2.2.3. PLWIN: Play-the-Winner

This rule can be summarised as: Win stay, lose shift. The rule involves no assessment of the environment, it is not a Bayesian rule. An animal starts by trying a patch at random. If it obtains a reward it makes the next hop in the same patch,

otherwise, it tries in the other one. PLWIN has been tested as a sub-optimal strategy for the two armed bandit problem with short horizons (Jones, 1975, 1976). The only storage required is for the outcome and place of the last hop (that is, only 2 bits).

8.2.2.4. CULLEN d: Cullen Rules

This family of rules is named after the celebrated twentieth century ethologist who suggested it to me. The subject starts by making bouts of 2 hops, one in each patch, until a decision is reached. At this point all attention is given to one of the two patches. The decision is reached when the cumulativ number of rewards in both patches differs by more than the parameter d. The parameter defines which member of the family of rules is being used at the time. Note that only the difference in the number of rewards during the 'sampling' period has to be stored not the actual cumulative records. In addition to this value an extra bit is required to identify which of the two patches is ahead at the moment. The models run in this chapter are Cullen 1, 2 and 3.

8.2.3. The evaluation of the rules

The aim of these simulations was to find out if it was possible to generate decision rules that approached the optimal solution at least as closely as the birds in the last chapter, while avoiding the computational complexity of CONOP. The 'goodness' of the different rules will not be judged by comparison with the birds actual behaviour, but with the obtained payoff. The comparison of the obtained payoff is best done by a transformation

of the actual number of rewards obtained in a test. The reason for transforming the data is that the room for improvement over random choice va r ies according to the treatment. For example, when the two reward probabilities are 0.20 and 0.30 a predator behaving at random would obtain on average 0.25 rewards per response, while an omniscient animal, who from the beginning delivered all responses to the best patch would obtain 0.30 rewards per response, an improvement of 0.05 in reward probability When the probabilities are 0.00 and 0.50 instead the room for improvement of omniscient over random behaviour is 0.25. A rule improving performance with respect to random by 0.04 in a 0.20 - 0.30 test would be doing a much more difficult task than one obtaining an increase of 0.08 in a 0.00-0.50 test. Figure 8.2 shows the universe of possible performances in a 150 hop test in the form of a trapezoid (a). The upper margin is the number of rewards that an omniscient animal would obtain in a test of 150 hops with the probabilities of reward marked in the abscissa. The lower boundary is the expected pay off from random behaviour in the same conditions. Two additional polygons are presented (b and c) showing the corresponding range of possible performances under different reward probabilities. Polygon b shows what happens when the difference between patches is held constant as 0.10 but the mean of the two probabilities ranges from 0.50 to 0.85, and polygon c presents the situation when the differences in probability between patches are the same as in the experiment but the absolute value of the two probabilities is raised to 0.65. Similar reasoning applies to comparisons between time horizons. I measured the good n ess of the different rules taking into account

the limit imposed by random and omniscient behaviour respectively. The measure of goodness is referred to as benefit and is defined by the expression:

$$\text{Benefit} = \frac{Ro - Rr}{Rm - Rr} \times 100$$

where Ro is the obtained number of rewards, Rr is the number of rewards expected from random behaviour, and Rm is the maximum number of rewards obtainable, only possible for an omniscient predator.

8.3. Results

8.3.1. Comparison with real birds

The benefit obtained by each of the decision rules in treatments mimicking my experiment is presented in Figure 8.3 together with the results obtained by the birds and by CONOP.

Two of the rules, play-the-winner and probability matching, obtain a lower benefit than the real birds, but the other four (Immediate maximising, CULLEN 1, 2 and 3) approach the constrained optimal solution at least as closely as the real animals.

In particular, the good performance of all CULLEN rules is impressive since these rules are extremely simple. This result demonstrates the possiblity of finding simple rules which do as well as the birds. CULLEN 1, which obtains a consistent advantage over the real birds, only uses two bits, one to compute if the difference in cumulative reward between patches is zero or one and the other to indicate which patch (if any) is ahead.

Figure 8.6a shows the ranking of all strategies including the birds and CONOP, and Table 8.1 presents the average ranking of each strategy over the five experimental treatments. Four

simple strategies did better than the birds in approximating the constrained optimum, while the other two strategies did worse.

8.3.2. Effect of average profitability

I tested the effect of the average value of the probabilitie by running the models with a constant difference between patches (/PA - PB/ = 0.10) but varying the mean probability (PA + PB)/2 = 0.05, 0.25, 0.45, 0.65 and 0.85). The space of possible benefits is shown as polygon b in Figure 8.2. The benefits obtained are plotted in Figure 8.4. The absolute level of probability affects the goodness of all rules, but some, particularly immediate maximising, are specially sensitive to it. When probabilities in both patches are very low, immediate maximising is just as good as CONOP, but for probabilities approaching unity it becomes almost the same as random behaviour. The reverse effect is true for play-the-winner, which is of little value at low probabilities but improves consistently as mean probability increases. In summary, not all rules are equally affected by the average quality of the environment.

8.3.3. Effect of patch diversity in a rich environment

In order to test if the effect of patch diversity was affected by average quality of the environment all rules were faced with patch diversities like those in the experiment but elevating the average probability from 0.25 to 0.65, as polygon c in Figure 8.2.

The results are shown in Figure 8.5. As in the first series of runs, when average reward probability was 0.25, increa-

ing difference in reward probability between patches produced higher benefits from all rules, but the effect of patch diversity was different for the different rules. For example, PLWIN showed a steeper response to patch diversity when average probability was 0.65 (Figure 8.5) than when the average probability was 0.25 (Figure 8.3) while MATCH showed the opposite effect. As a consequence, MATCH obtained a greater average benefit than PLWIN when average reward probability was 0.25 and a lower one when the average reward probability was 0.65. In order to make apparent the effect of different conditions on relative performance, Figure 8.6 shows the ranking of all the rules under the three series of treatments. The changes in rank under varying conditions implies that both average quality of the environment as well as patch diversity modify which of the rules most closely approximates the constrained optimum. It is worth mentioning that in a few points CONOP did not obtain the highest benefit. This result may be due to either the effect of the constraint in the algorithm used here or the fact that some simple rule may be particularly good for a very narrow set of parameters. As expected from the theoretical discussion, through, no strategy did consistently better than CONOP.

Table 8.2 shows the average ranking of each rule across the 15 treatments.

8.3.4. Effect of the time horizon

All previous simulations in this chapter have used horizons of 150 hops. In order to investigate the effect of the horizon the rules that had previously shown good performance (IMMAX and

all Cullen rules) were run with a probability configuration of 0.15-0.35 and time horizons of 50, 150 and 250 hops. The benefit are shown in Figure 8.7 and the ranking of the rules in Figure 8.8. It is clear in Figure 8.7 that all strategies do better for long than for short horizons, but the effect was more marked for Cullen 3. This rule was the worst one for short horizons and moved to the first place below CONOP when horizon was 250 hops (Figure 8.8).

8.4. Discussion

The main aim of the simulations was to explore the possibility of finding simple decision rules which reach the level of performance of the birds in the experiment. This was successfully accomplished.

The question of what type of rule does actually control the birds' behaviour is a different matter, which will be discussed in the next chapter. All the rules make specific predictions about the details of the individual choices, and these predictions can be tested against the real data.

There is one feature of the dynamic programming algorithm that is not shared by any of the rules proposed so far: the response to time horizon. The results of CONOP showed that the optimal amount of sampling increases with the time horizon (Chapter 7). If the birds were shown to increase sampling with length of the test, then it would be necessary to conclude that none of these elementary rules can account for the data.

Regardless of this possibility there is a point that is worth mentioning about the philosophy of the different models.

Of all the rules proposed, only one makes 'sacrifices' in order to gain knowledge, all the others choose the place of the next hop according to the maximum subjective probability of getting a reward. The exception is the family of Cullen rules. The 'sacrifice for the sake of knowledge' arises when the subject playing a Cullen rule has experienced a difference in cumulative number of rewards between the two patches but this difference is smaller than d, the decision parameter. The reason for waiting for d to increase while maintaining equal rate of responding to both patches is to decrease the chance of choosing the wrong patch for the rest of the test. Using a bigger d decreases the number of tests in which a wrong decision is taken, but prolongs the sampling period. The study of the effect of horizon suggests that there must be an optimum value of d for each particular length of test and set of reward probabilities. Waiting longer for the decision would only pay when there is enough test left for the risk of mistakes to outweigh the consequences of sampling for a long time. Thus, while an individual using a Cullen rule with fixed d would not be able to respond to variations in time horizon, a subject which used Cullen and udpated d to the time horizon according to some simple increasing function would still be a possibility if response to horizon is shown in real birds. The question of whether great tits respond to time horizon is the subject of the next chapter.

Table 8.1

Mean rank of decision rules in the conditions like the experiment in Chapter 7.

Rule	Mean rank
Constrained optimum	1.6
Immediate maximising	2.0
Cullen 1	2.9
Cullen 2	3.7
Cullen 3	5.2
Great tits	5.6
Matching	7.0
Play the winner	8.0

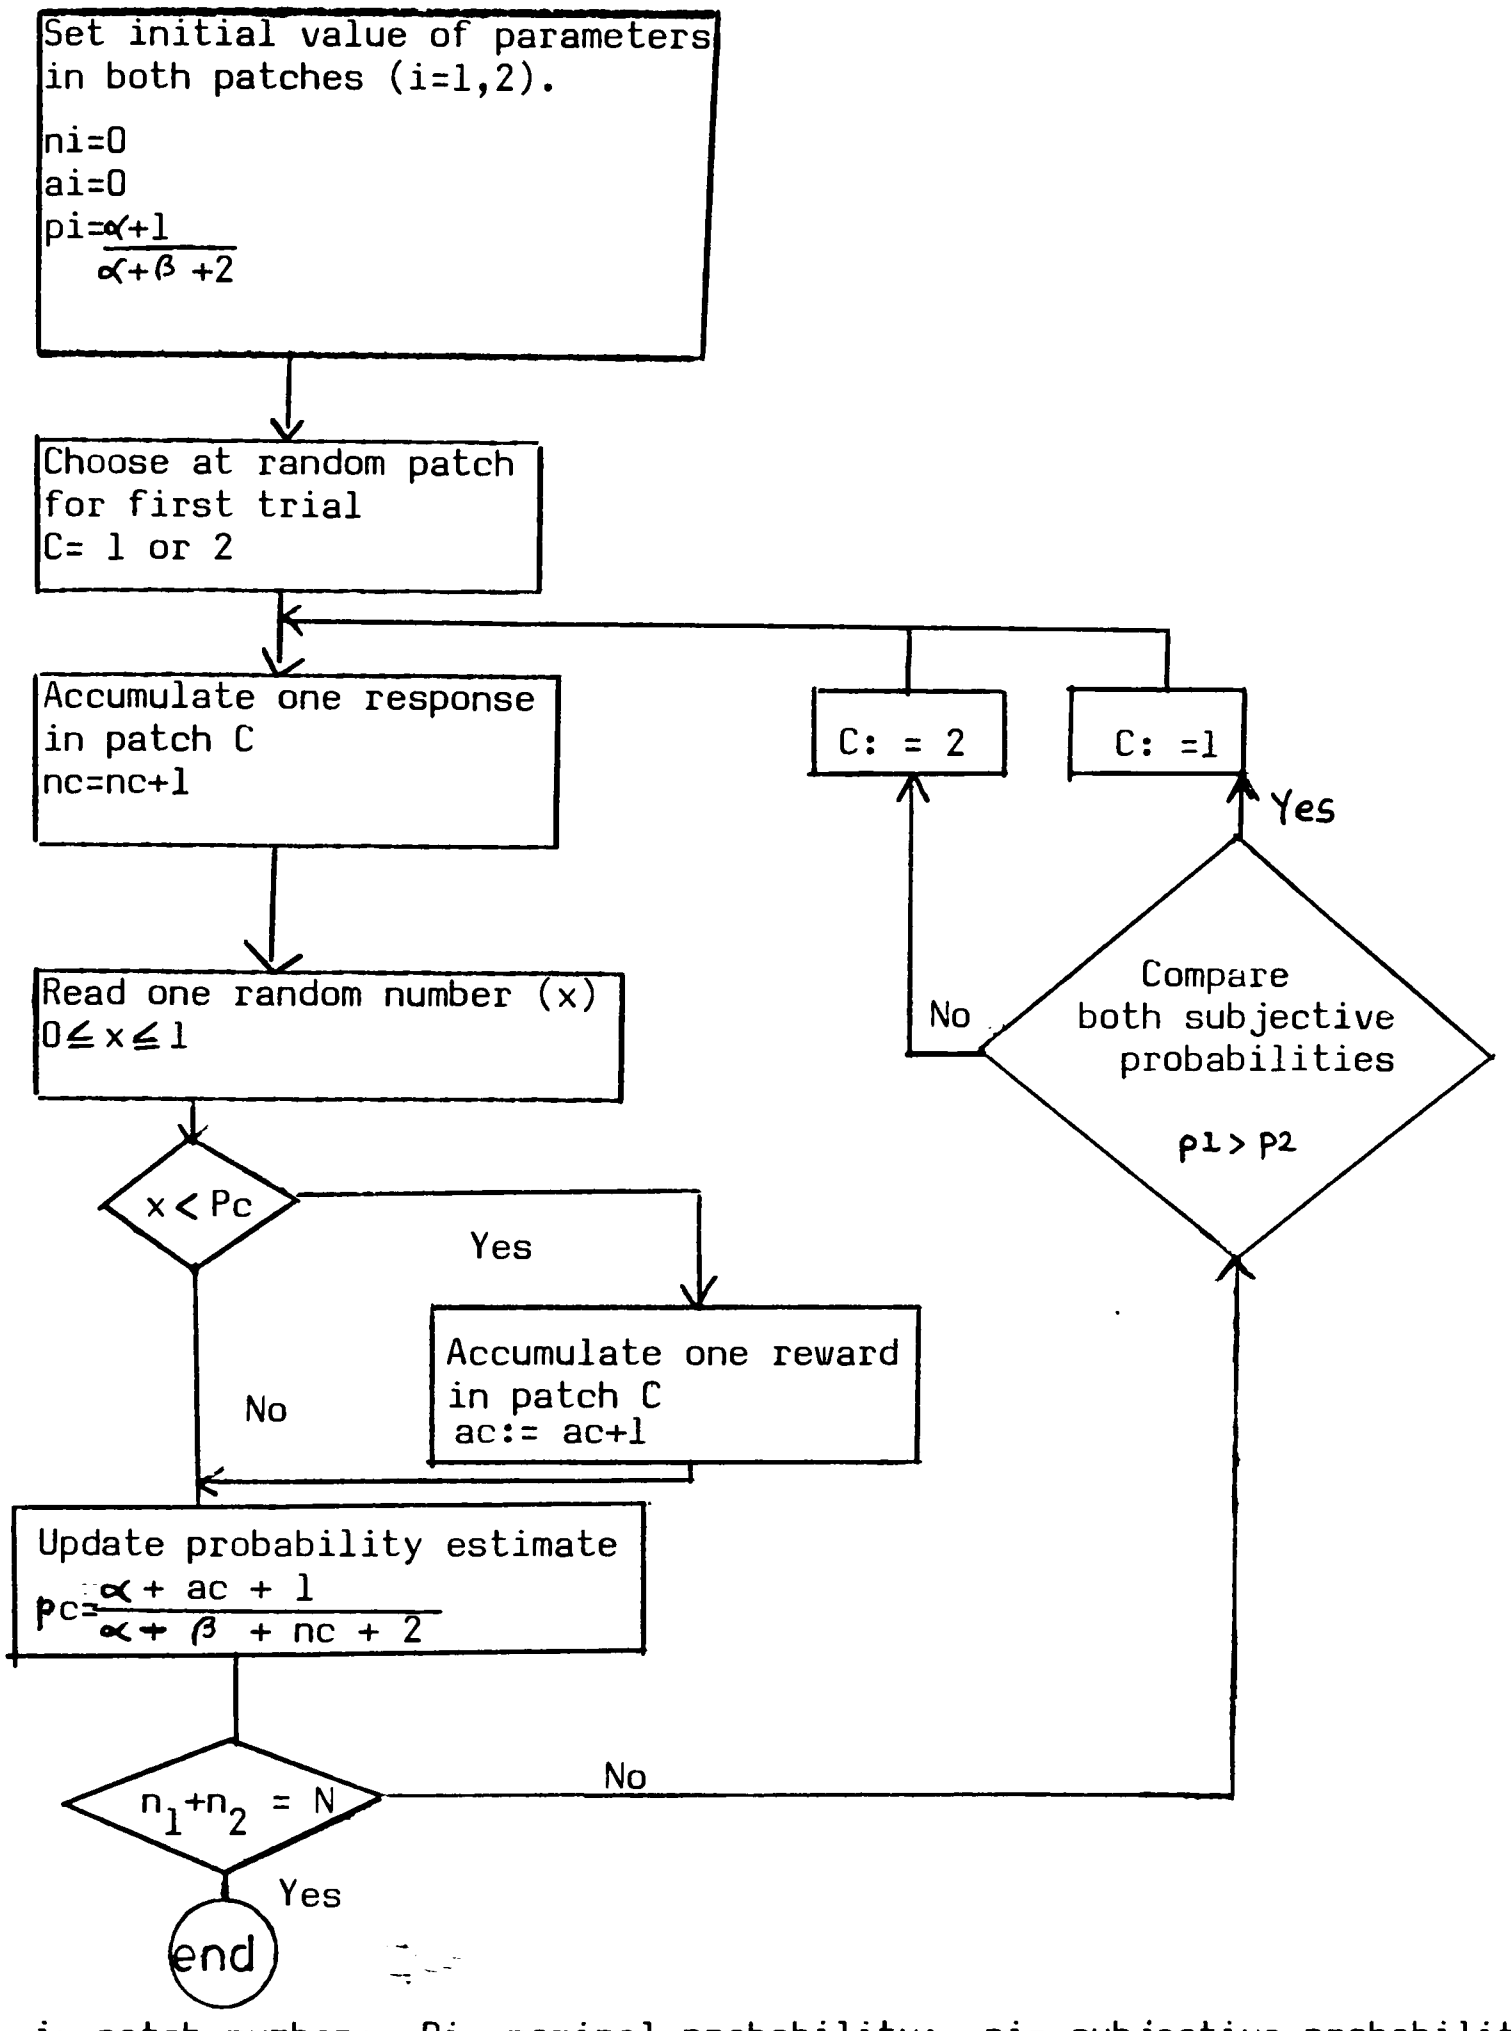
Table 8.2

Mean rank of decision rule over the fifteen conditions tested.

Rule	Mean rank
Constrained optimum	1.5
Cullen 2	3.2
Immediate maximising	3.2
Cullen 1	3.6
Cullen 3	3.9
Matching	6.3
Play the winner	6.3

Figure 8 . 1

Immediate Maximising. Flow chart



i= patch number; Pi= nominal probability; pi= subjective probability;
ni= cumulated responses; ai= cumulated rewards; c= chosen patch;
N= time horizon

Figure 8.2 Limits of possible performance of decision rules in a test lasting 150 hops. Polygons a, b, and c refer to the combinations of reward probabilities indicated in the appropriate scale on the abscissa. In each case the upper boundary is the pay off resulting from exclusive response to the better patch (omniscient bird), and the lower one is the pay off from responding equally in both patches (random bird). Polygon a: conditions as in the experiment described in Chapter 7. Patch diversity ranges from 0.10 to 0.50 but average probability is held constant at 0.25.
Polygon b: patch diversity is constant in the five conditions, but the average reward probability varies from 0.05 to 0.85. Polygon c: the same patch diversities as in polygon a but in a richer environment Average probability is 0.65 instead of 0.25.

Figure 8.2

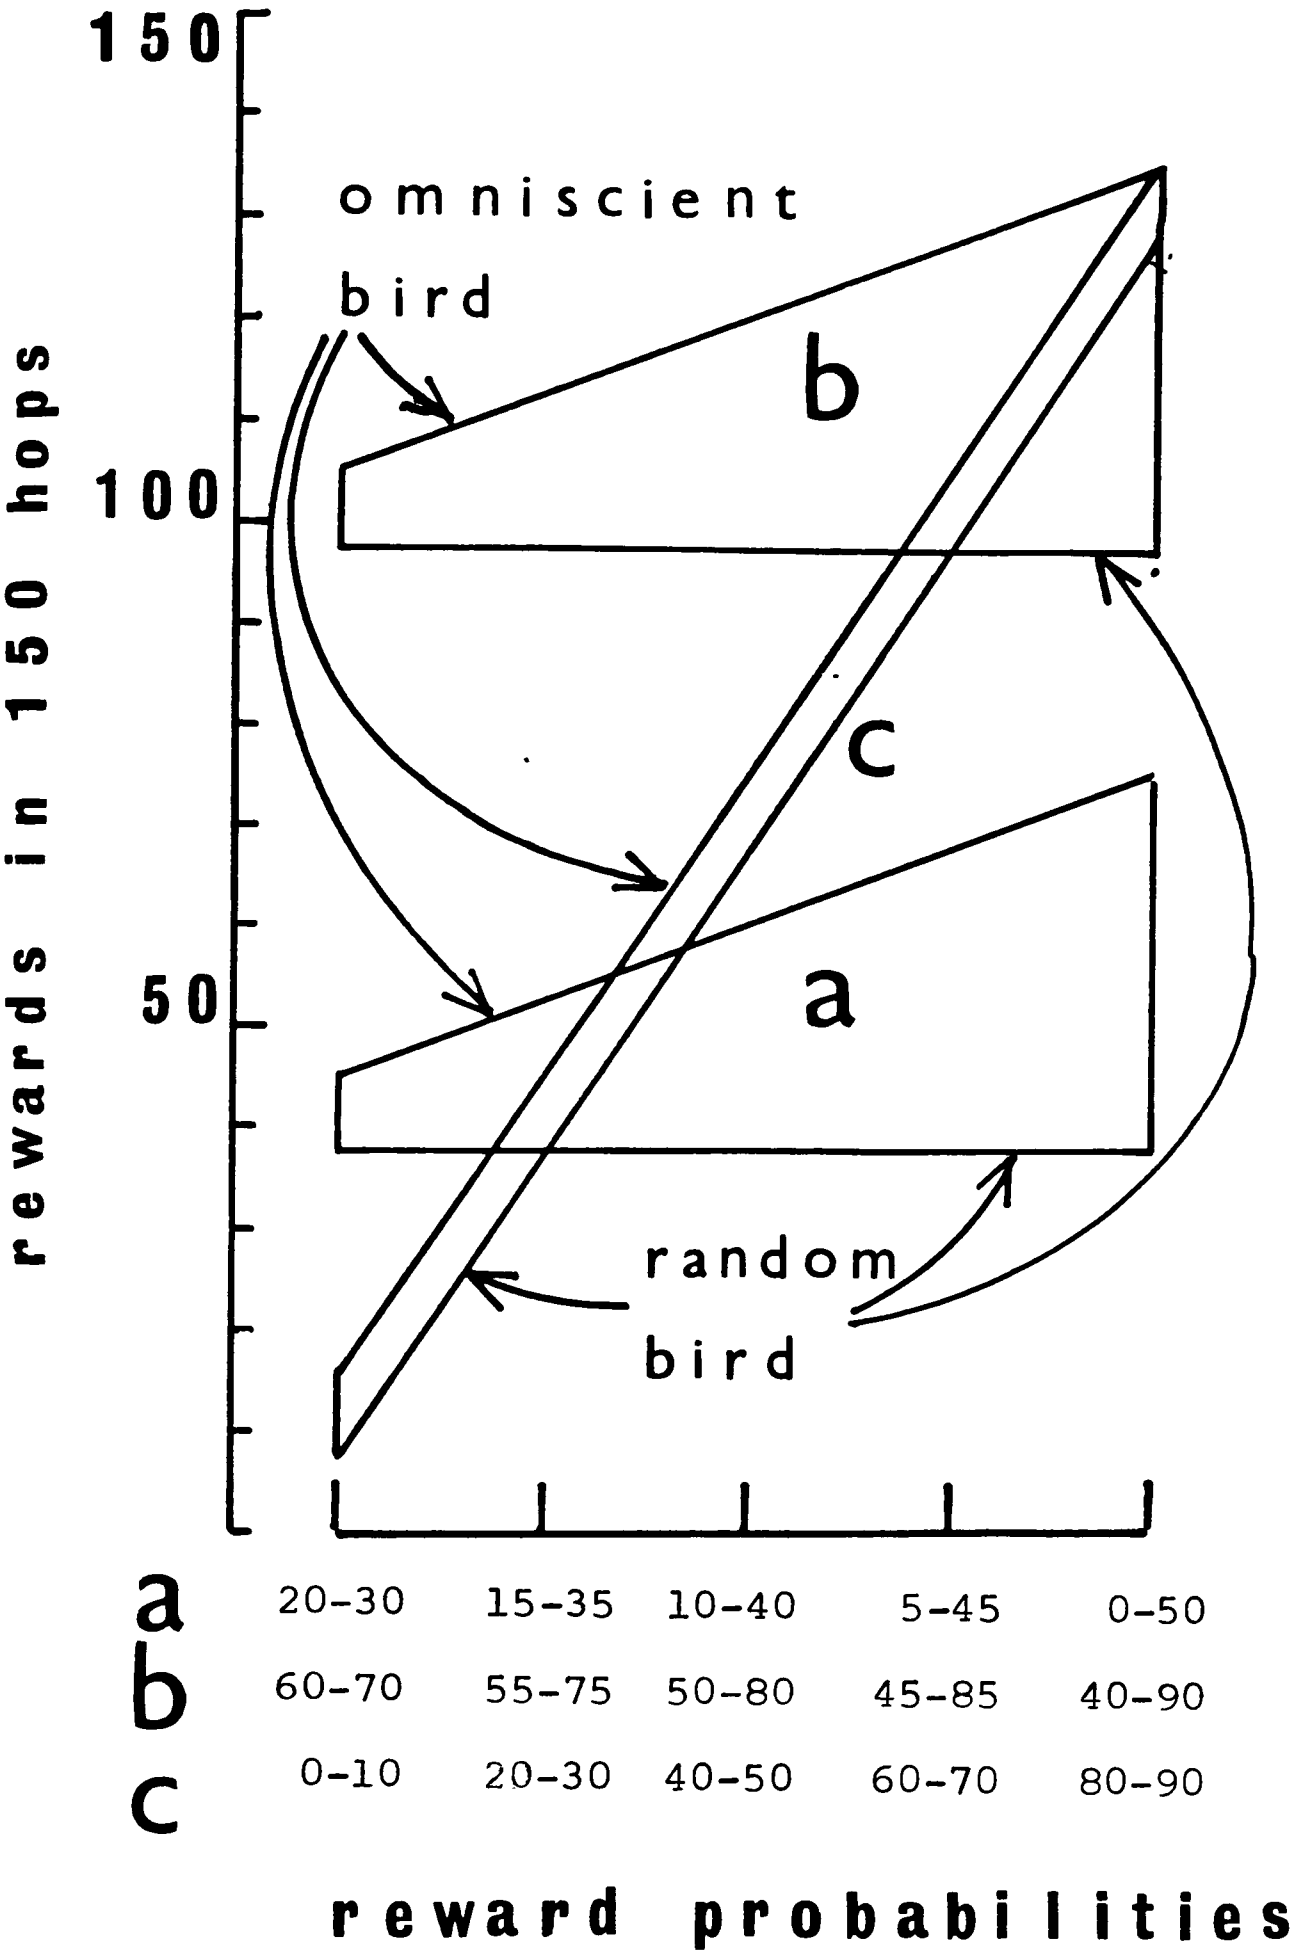
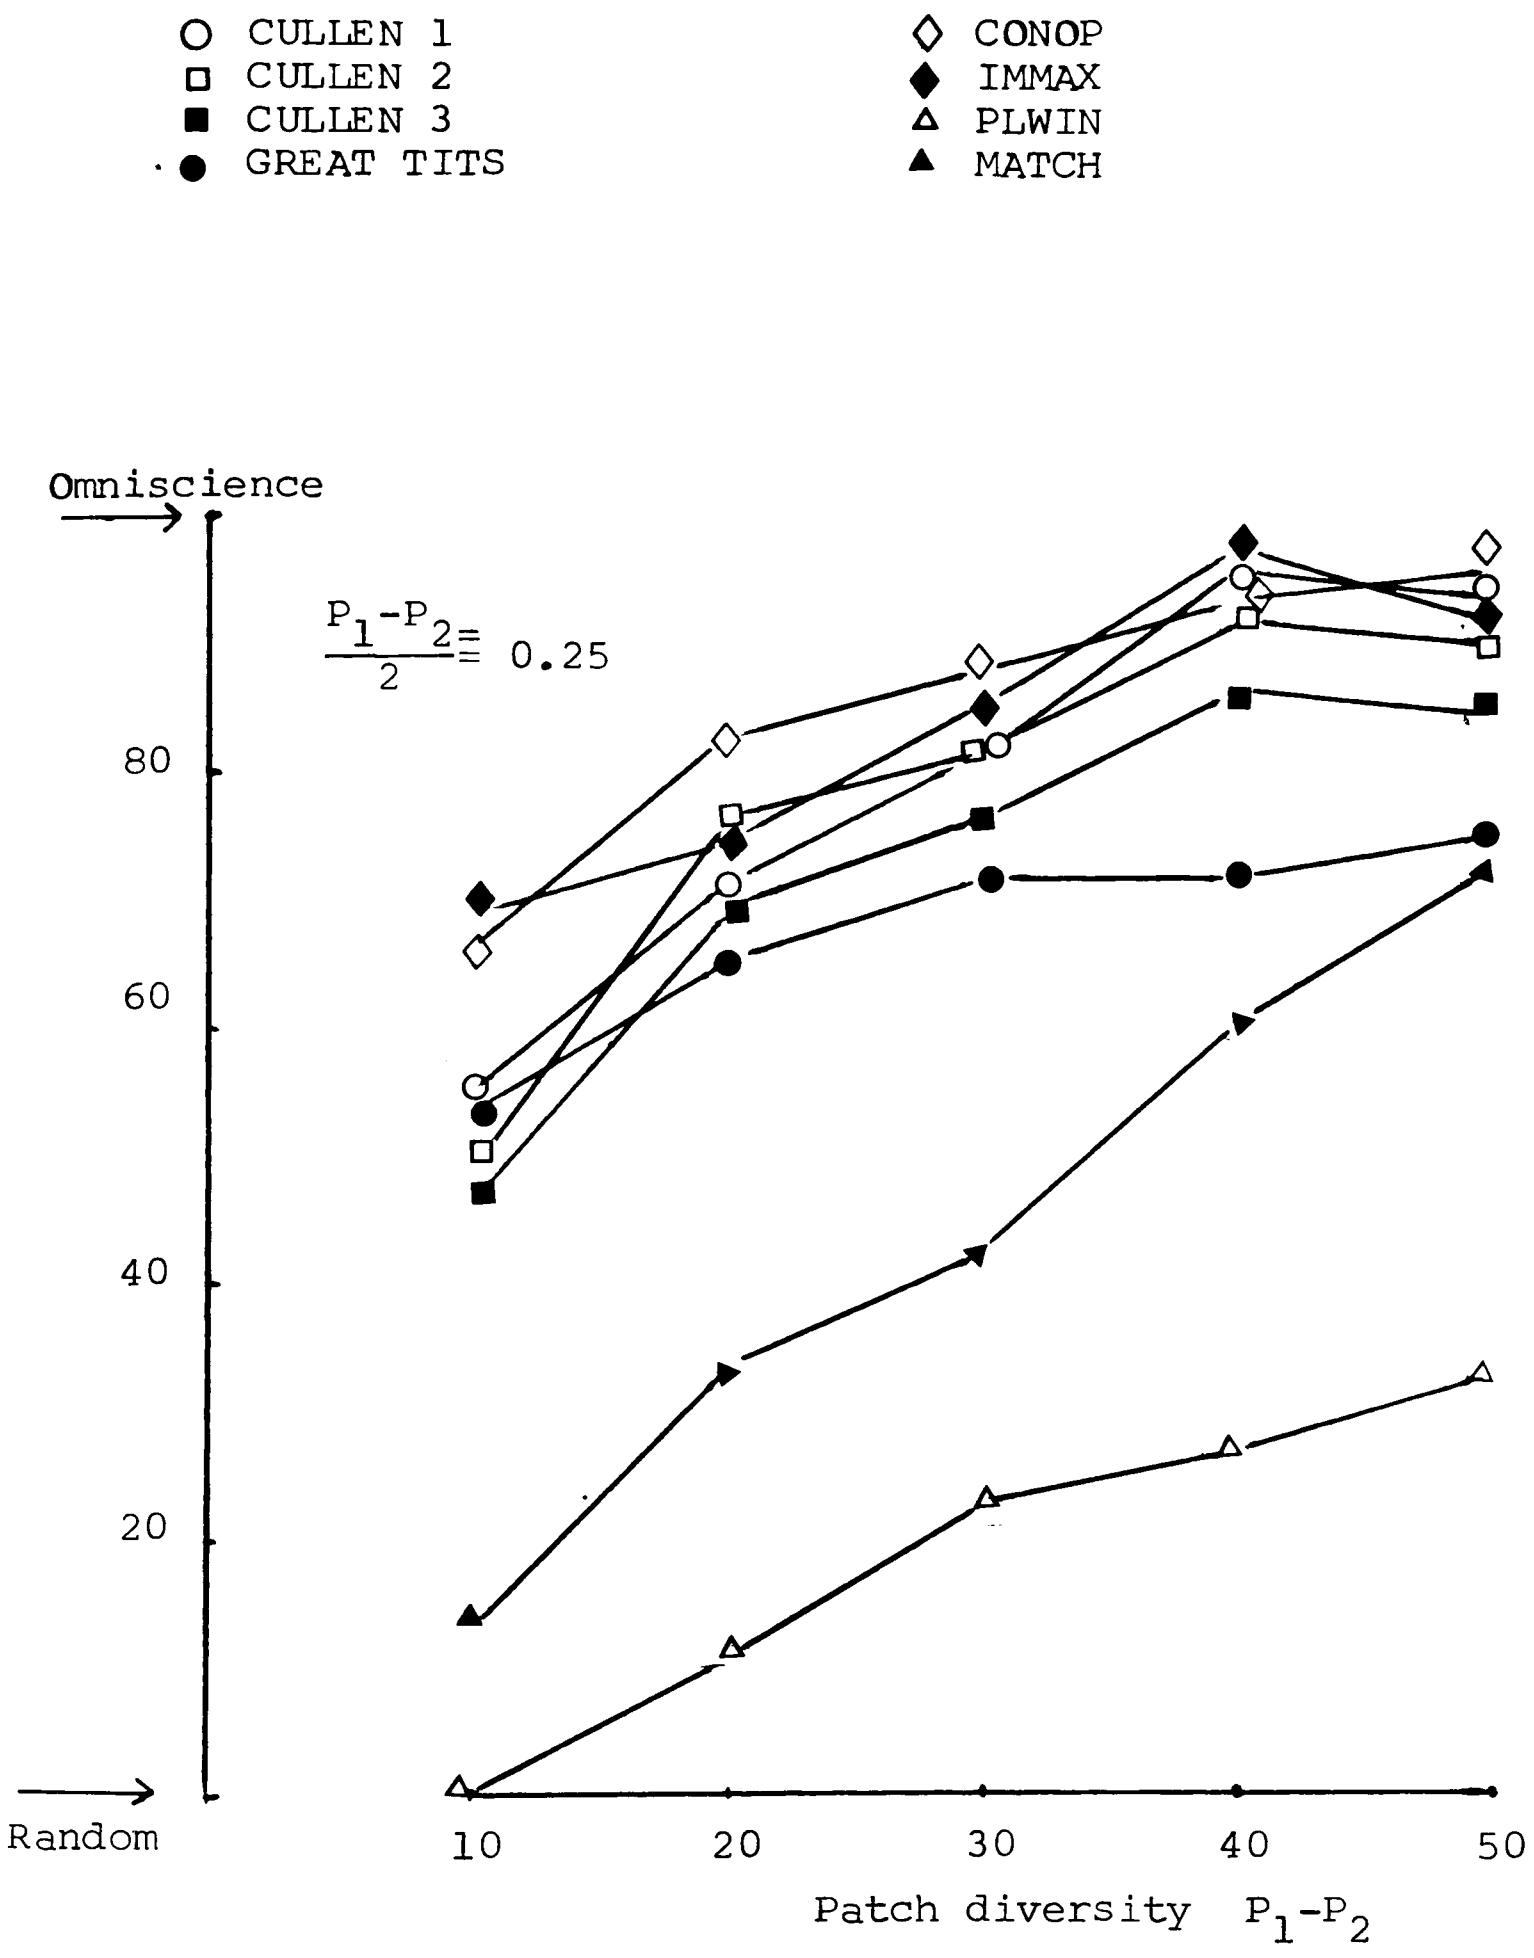
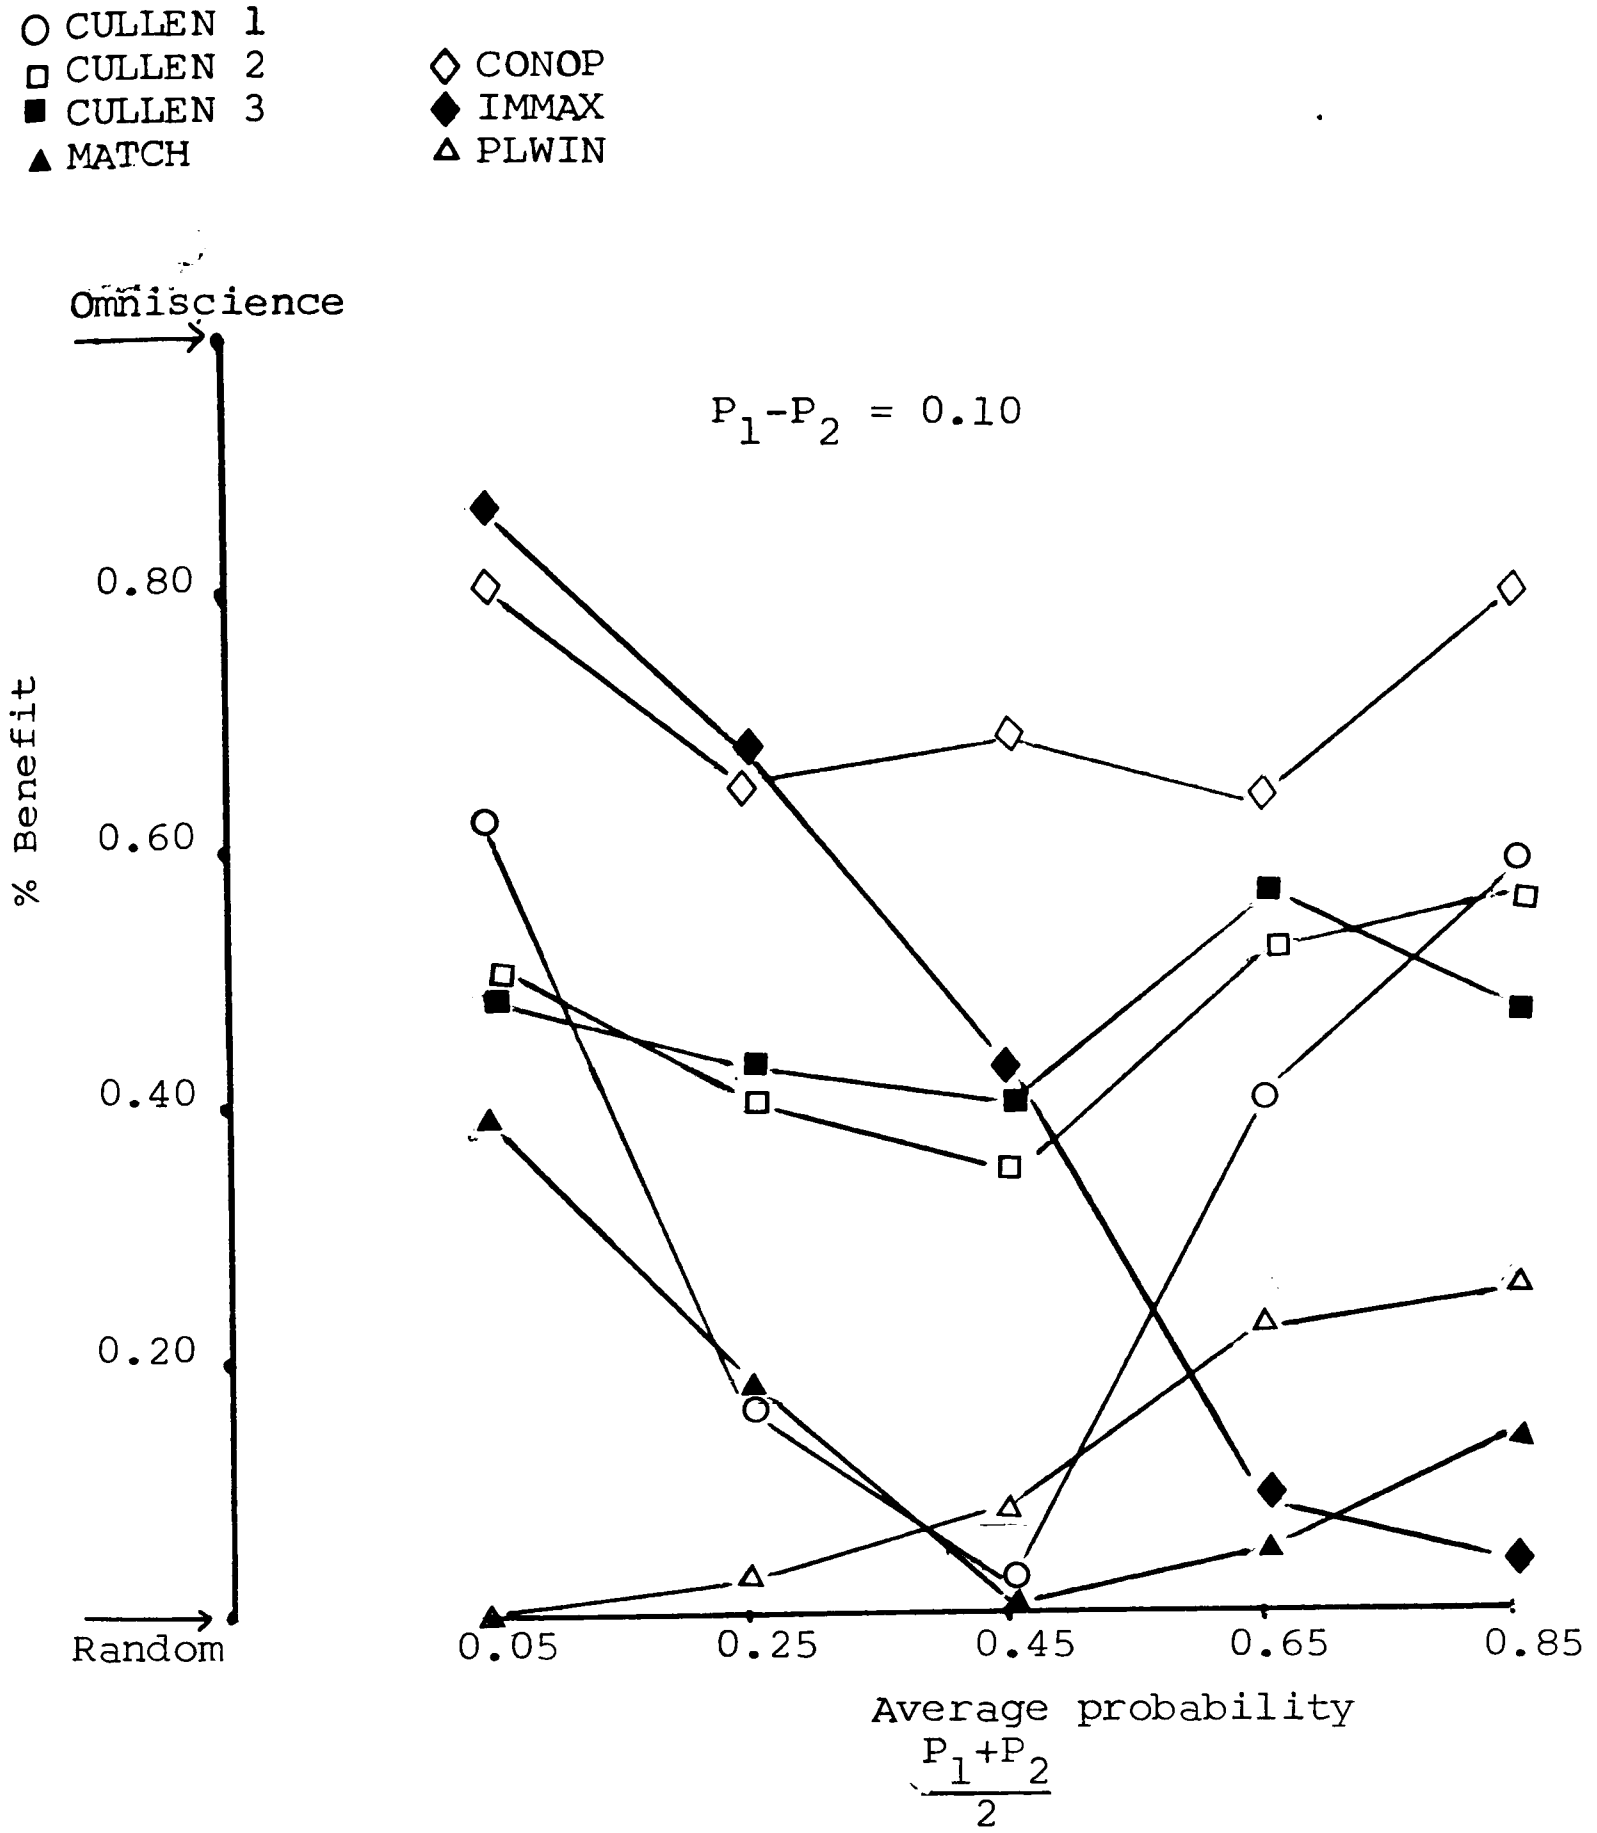

150
100
50
rewards in 150 hops
omniscient
bird
b
c
a
random
bird
a
b
c
20-30
15-35
10-40
5-45
0-50
60-70
55-75
50-80
45-85
40-90
0-10
20-30
40-50
60-70
80-90
reward probabilities

Figure 8.3



Benefit obtained by the decision rules under conditions like those in polygon a of Figure 8.2 Since these condition are like those in the experiment described in Chapter 7, the benefit obtained by the great tits is also plotted. The abbreviations are explained in the text. The measure of benefit takes into account the attainable payoff for an omniscient and for a random bird.

Figure 8.4

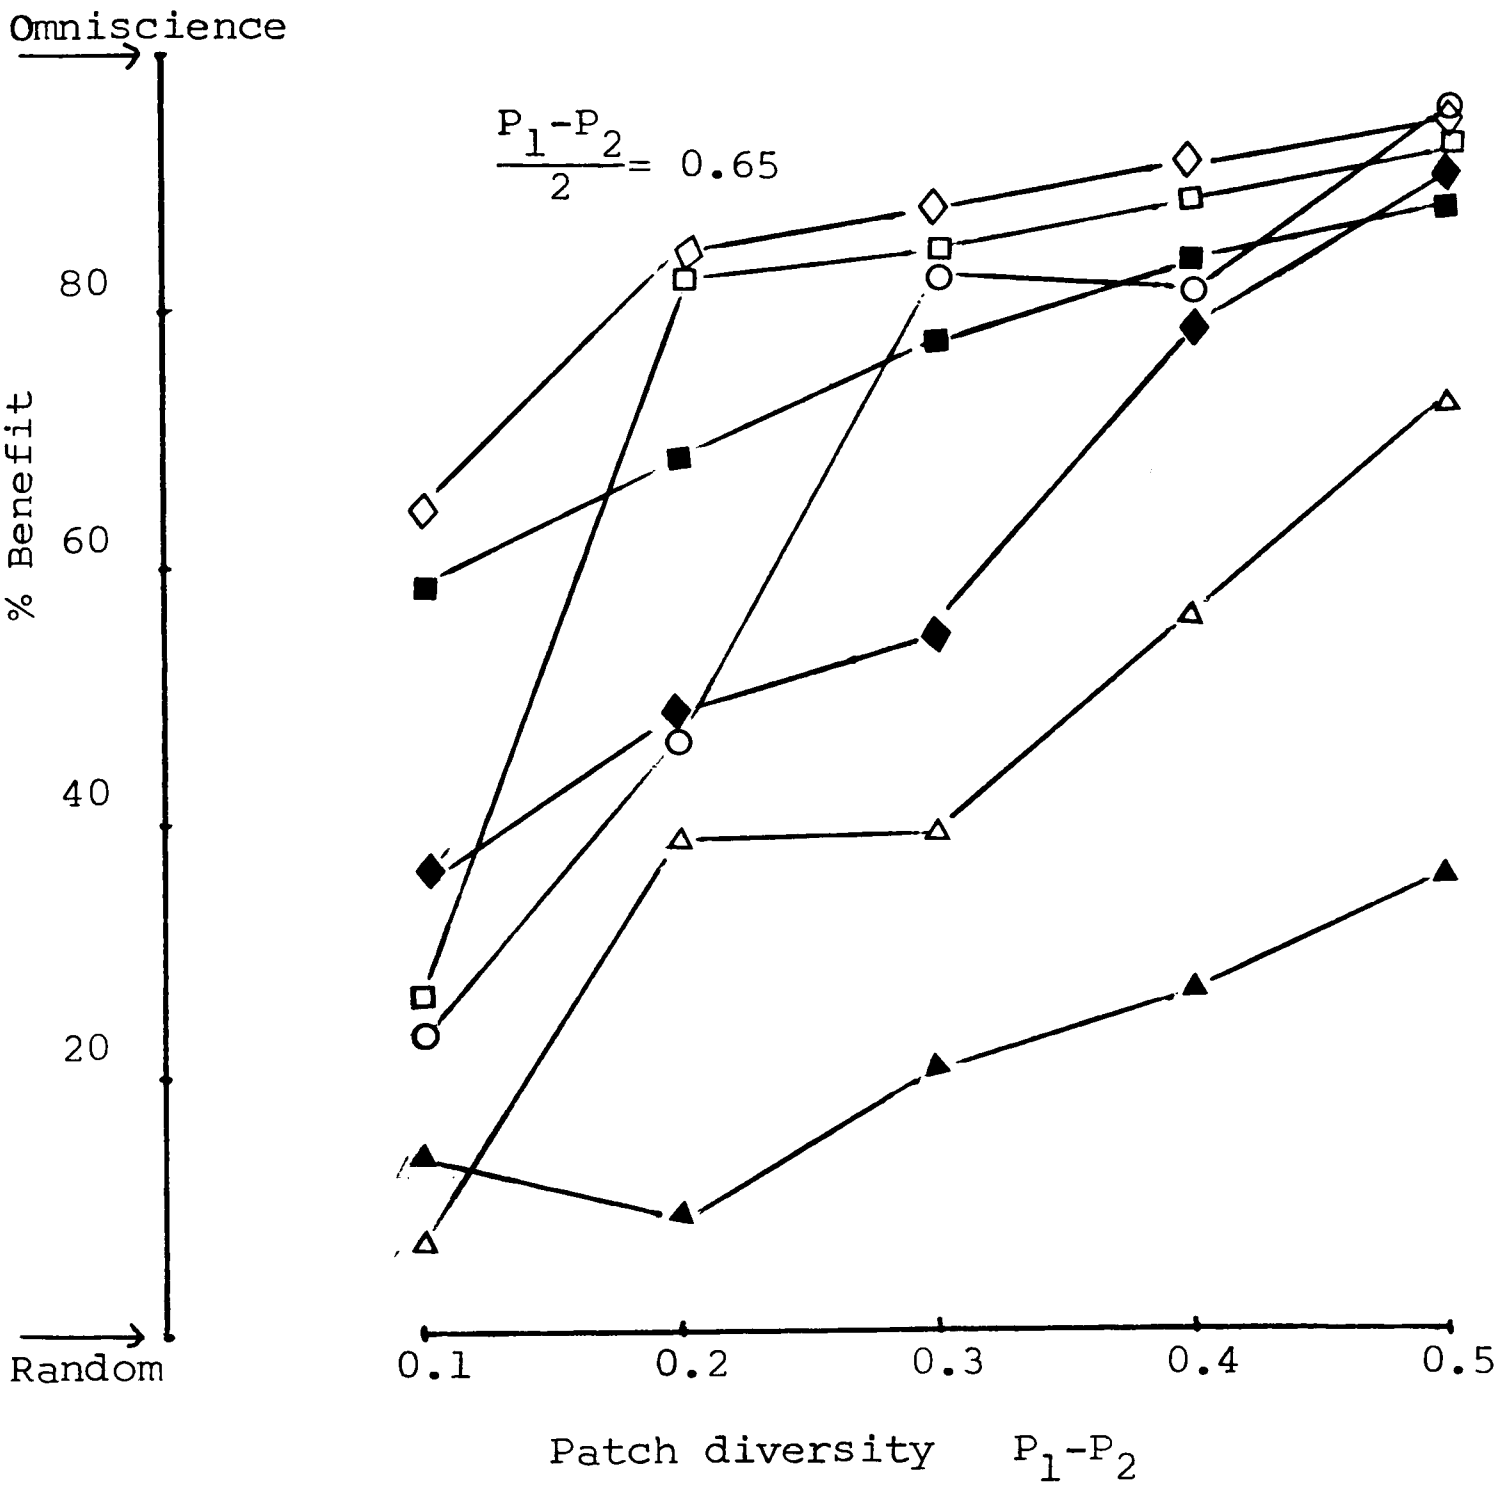


Benefit obtained by decision rules under conditions like in polygon b in Figure 8.2. The abbreviations are explained in the text.

Figure 8.5

◇ CONOP
◆ IMMAX
▲ MATCH
△ PLWIN

○ CULLEN 1
□ CULLEN 2
■ CULLEN 3



Benefit obtained by the decision rules under conditions like in polygon c in figure 8.2. The abbreviations are explained in the text.

Figure 8.6 Ranking of decision rules according to benefit obtained. A, B, and C refer to series of runs with conditions like in the respective polygon in Figure 8.2, (actual benefit obtained is shown in Figures 8.3 to 8.5). Abbreviations are explained in the text.

◇ CONOP

◆ IMMAX

○ CULLEN 1

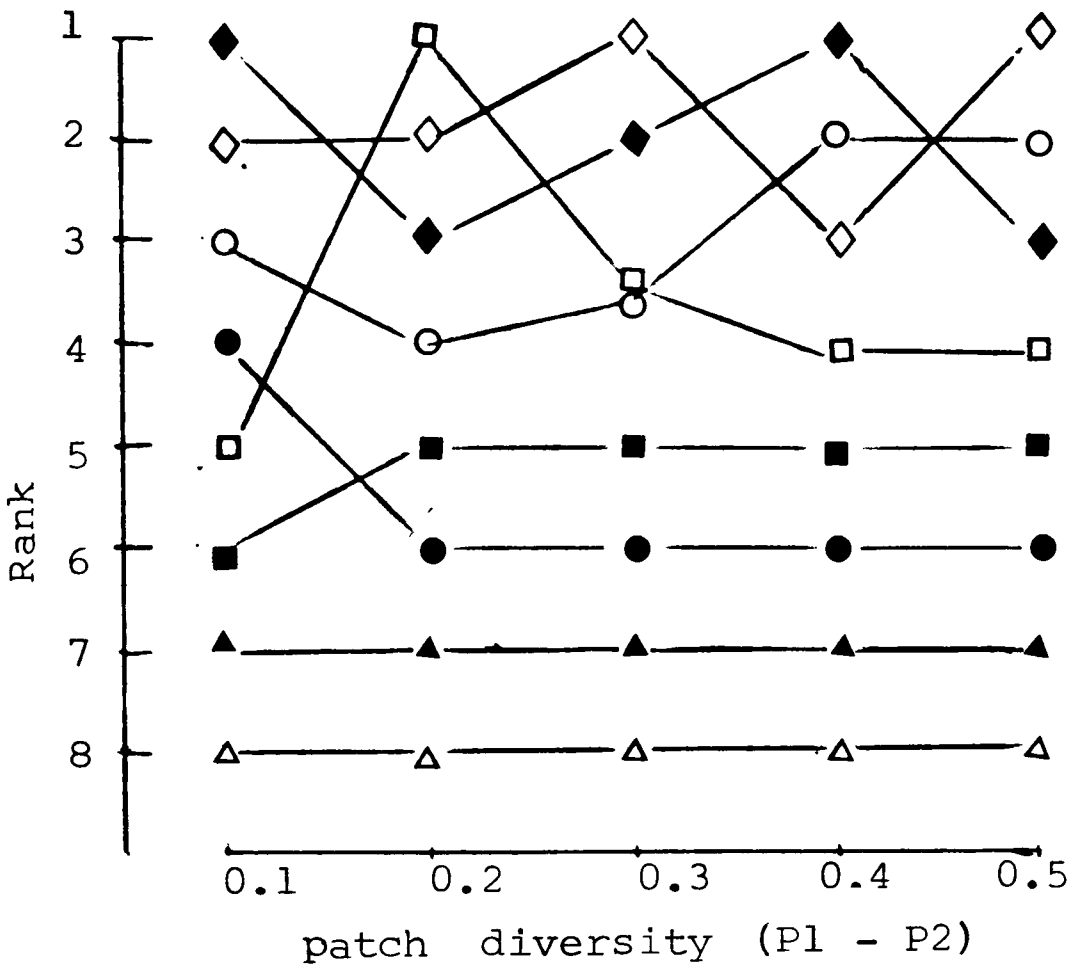
▲ MATCH

■ CULLEN 3

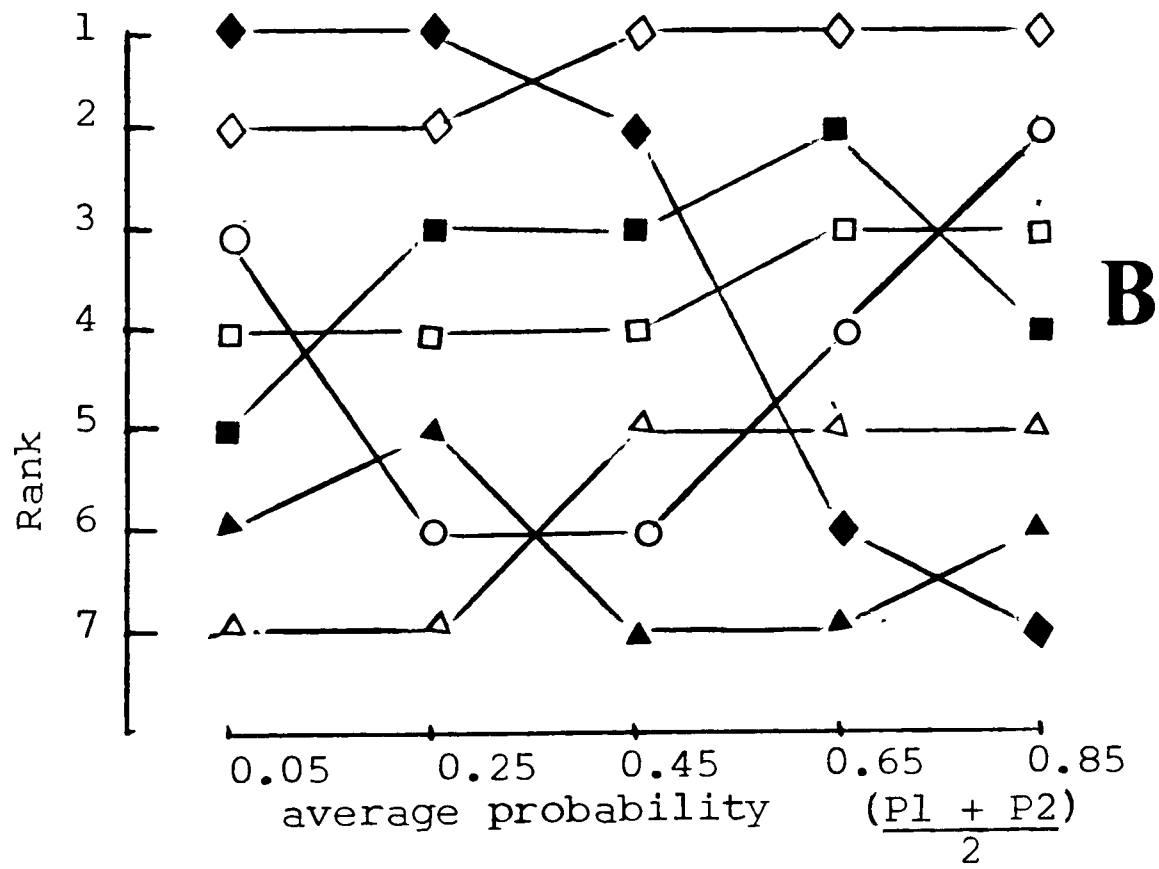
□ CULLEN 2

△ PLWIN

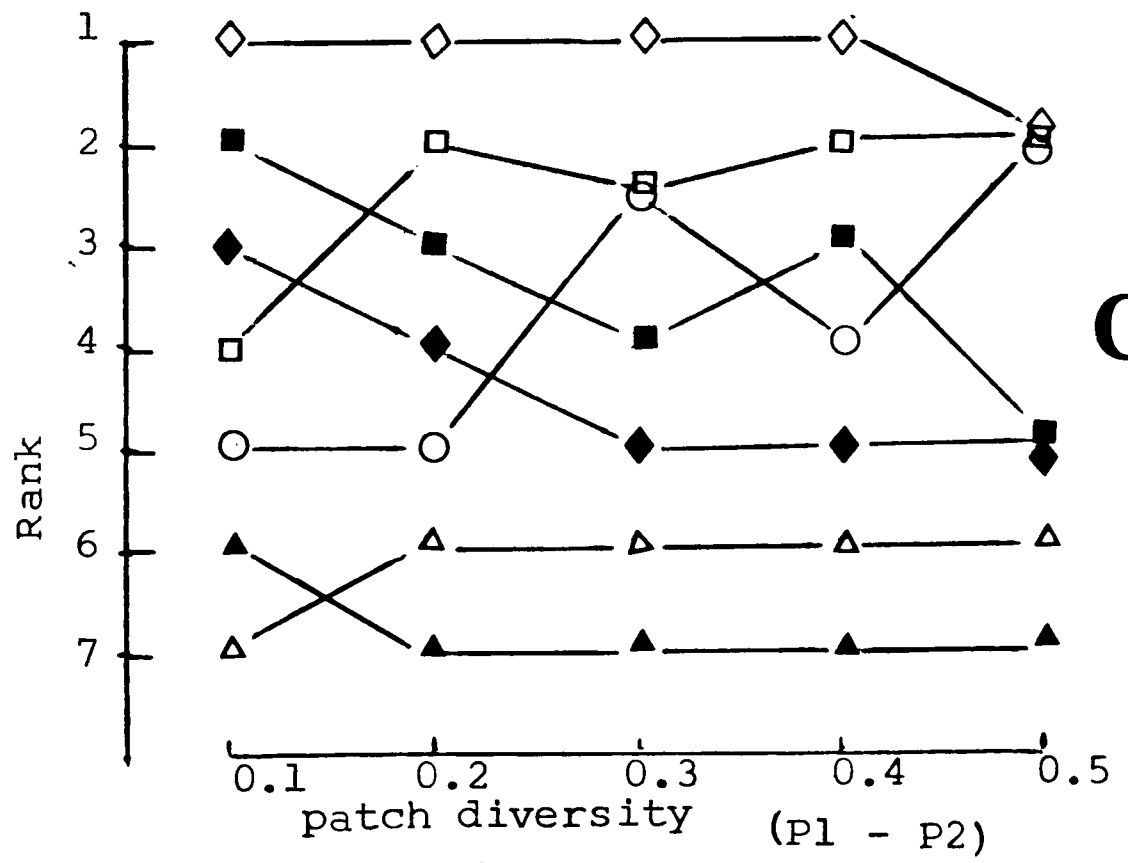
Figure 3 . 6



A $\frac{P1 + P2}{2} = 0.25$



B P1 - P2 = 0.10



C $\frac{P1 + P2}{2} = 0.65$

Figure 8.7 Effect of time horizon.
Benefit obtained by the five rules that obtained higher benefit than the great tits in the experiment are plotted as a function of time horizon. Reward probabilities were held constant at 0.15-0.35. Abbrev. are explained in the text.

Figure 8.8 Ranking of decision rules as a function of time horizon. Reward probabilities are held const. at 0.15-0.35. Notice that Cullen 3 is the worst rule when horizon is 50 hops and shifts to be the best rule below the constrained optimum for a horizon of 250 hops.

Figure 8.7

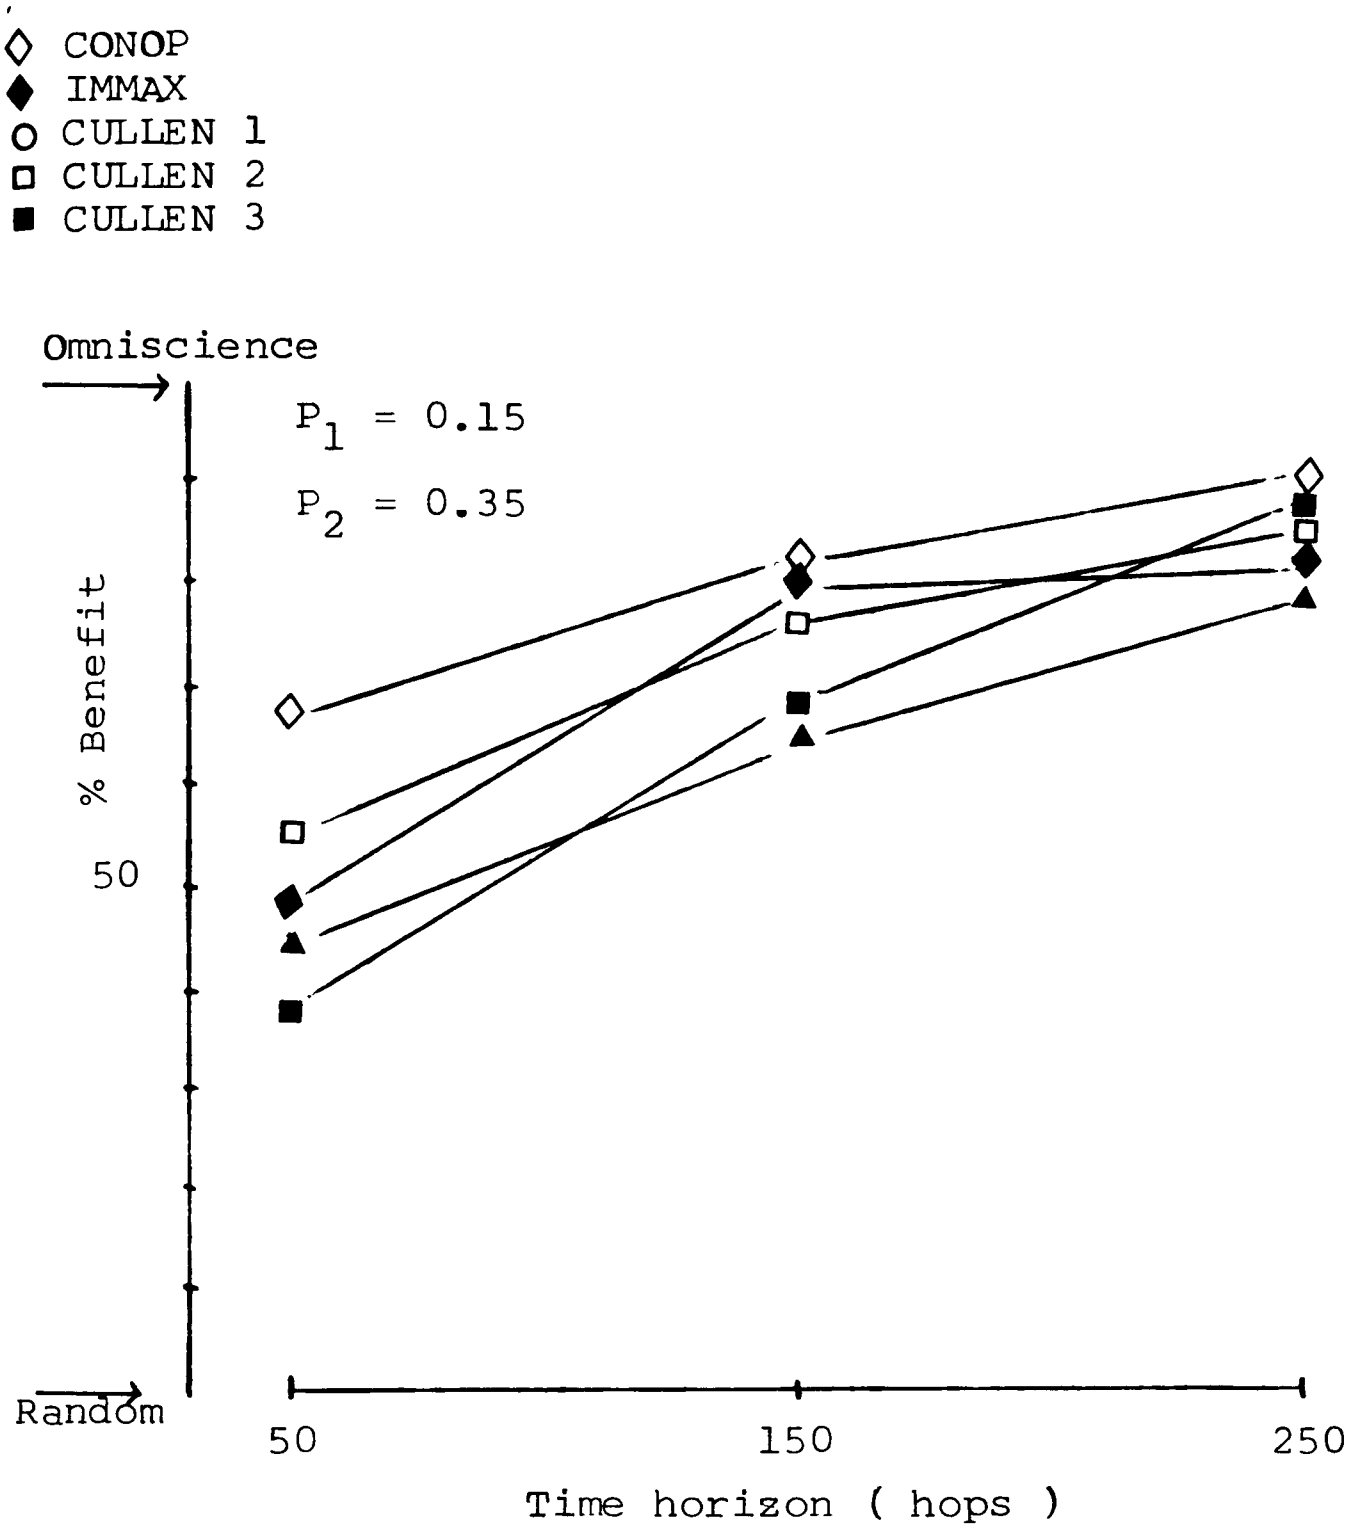
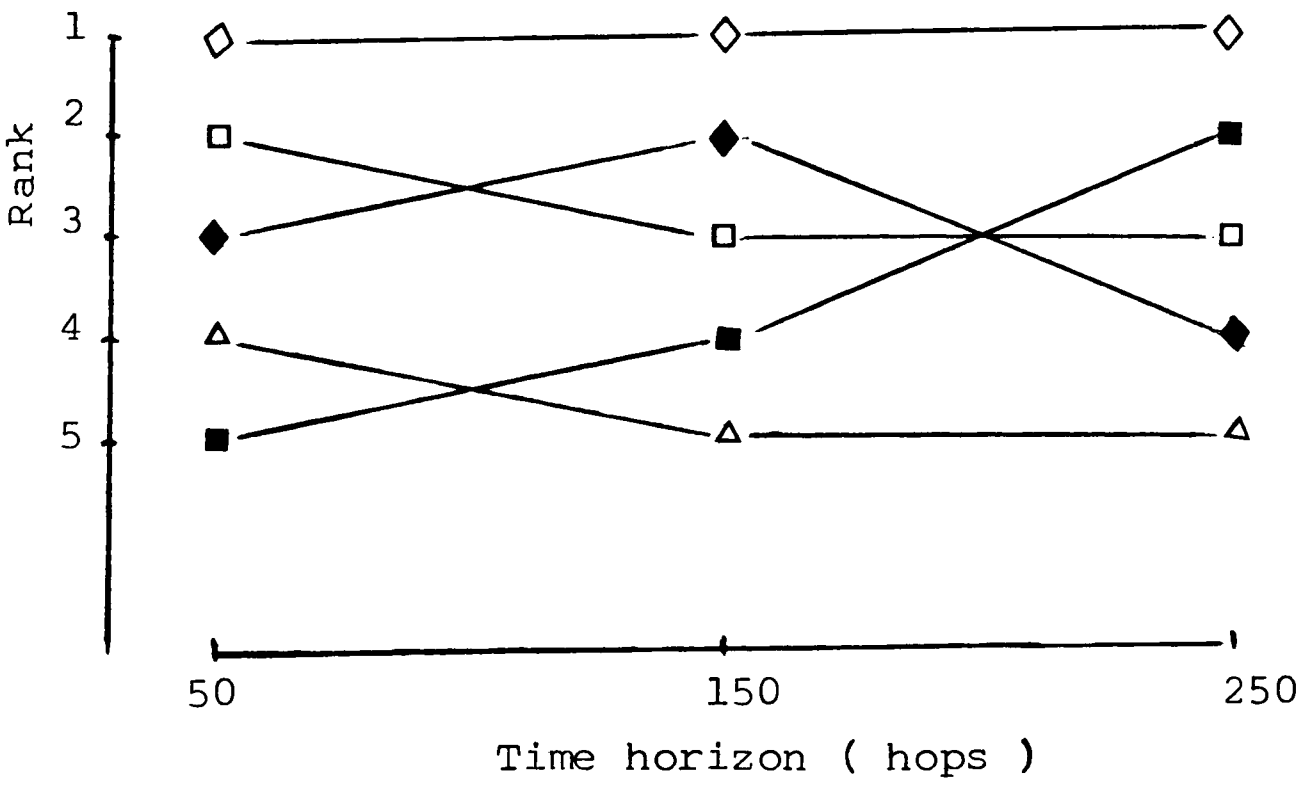


Figure 8.8



CHAPTER 9

Optimal Sampling and Time Horizon

9.1. Introduction

The first experiment in this series, described in Chapter 7, raised at least as many questions as it answered. The experiment was conclusive in showing that great tits learnt to concentrate their foraging effort in the better of two patches after a period of sampling, the length of which depended on patch diversity. The average amount of sampling at different patch diversities approached the predictions from a dynamic programming algorithm (CONOP) that generates a near optimal sampling behaviour subject to some constraints (Chapter 7).

The results of simulations described in Chapter 8 showed that there are simple rules-of-thumb which approximate the constrained optimal algorithm at least as closely as the birds. These simulations investigated the payoff obtained by theoretical birds using different strategies under conditions of varying patch diversity. But the response of CONOP to patch diversity depends on time horizon (as shown in Figure 7.6 and 8.7), a property not shared by any of the simple rules tested in Chapter 8 (but see Chapter 8 for a discussion of possible extensions of Cullen rules). The results obtained from the experiment in Chapter 7, though, showed that the amount of sampling done by the birds was very close to the predictions from CONOP for the length of tests that the birds had experienced most frequently (Figures 7.6 and 7.7), suggesting that there might be a response. Be-

cause of the criterion used to stop the tests there was no way to establish whether the length of the test was responsible for the agreement between data and predictions. Since the treatments were stopped according to a criterion measured over 100 consecutive hops, the minimum length of a treatment was the sampling period plus 100 hops, and at least two possibilities could account for the results:

(a) Great tits might respond to time horizon. The birds may have predicted the modal length of test from their previous experience, and the observed amount of sampling may have reflected their adjustment to this assessment.

(b) Great tits might be totally unable to adjust to different horizons, but may have a 'natural' expectation of patch stability based on evolutionary adaptation (or long-term learning in the wild before capture) to common food patches in the wild. The good fit observed in my experiment may have resulted from a purely accidental coincidence between the length of experimental tests and the birds' natural pre-adaptation.

The problem of the time horizon therefore emerges as one of the most interesting questions raised by the experiment in Chapter 7. The conflict between molecular models, based on momentary maximisation and molar predictions, which emphasise the payoff obtained over a whole session, is most apparent when the adaptation to time horizon is involved. It is clear that molecular maximising does not predict any response to time horizon, since a subject maximising momentary rate of intake (see Discussion

in Chapter 6) would always respond in the alternative with higher expected probability of reward. Molar maximisation instead predicts a balance between short- and long-term benefit which leads to sampling as computed using CONOP. The question is also attractive because of the high predictive information value of the models (Dawkins and Dawkins, 1974), which produce counter-intuitive and unabiguous expectations. For example, CONOP predicts that patch diversity and time horizon should not have an additive effect on the amount of sampling, but a well defined interaction (Figure 7.6). The optimal behaviour involves sampling more when time horizon is longer for all patch diversities, because with longer time horizons the cost of mistakes is bigger, and is more likely to offset the losses due to additional sampling. But the increase in sampling for a given difference in time horizon should be dependent on patch diversity, with bigger increments in sampling expected for the same increment in horizon when the two patches are more similar in reward probability. This is because a given reduction in the probability of a mistake requires more additional sampling when the two patches are similar than when they are very different.

The experiment described in this chapter is aimed at studying the effect of time horizon on sampling strategy. The same experimental set up as in Chapter 7 is used, but the procedures were modified to investigate the interaction between time horizon and patch diversity.

9.2. Methods

Since the experiment described in this chapter is a modification of that in Chapter 7, in this section I will only describe those aspects that were modified, and assume that definitions given in the previous two chapters are still valid.

9.2.1. Subjects

I used four male great tits. All of them had had previous experience with feeding discs.

9.2.2. Instruments

The only modifications withrespect to the experiment in Chapter 7 were that the second version of the feeding discs were used instead of the first one, and small fly pupae (see Section 1.7) were used as prey instead of segments of meal worms. These two changes meant that tests could now be longer than before (up to 260 rewards instead of 72 rewards), and treatments did not need to be interrupted.

9.2.3. Types of run

9.2.3.1. Treatments

Since treatments were now completed in single tests, the terms are used interchangeably, meaning a run with different reward probability in each patch. There were four types of treatments, varying in the reward probabilities and length of test used. The reward probabilities were 0.15-0.35 and 0.10-0.40, (patch diversities of 0.20 and 0.30 respectively). Runs were stopped after a fixed number of hops, which depended on the time horizon being tested. For short horizons the length was 50

hops, and for long horizons it was 250 hops.

9.2.3.2. Neutralisations

Treatments were still intercalated with neutralisations, but the neutralisation technique was different. I used a special neutralisation schedule consisting of rewarding always the first hop in a patch, and never rewarding any other hop. The birds adjusted to this schedule very quickly, by learning to switch between patches after a single hop. Figure 9.1 shows that after four runs of training all the birds were switching after a single hop in more than 50% of the bouts. The main advantage of this technique over the previous one was that it was more effective in suppressing preference for any of the patches. While the old neutralisation established equal use of both patches at global level (similar number of hops in each patch at the end of a run), this schedule operated at molecular level, since it actually drove the birds to alternate every single hop. The change in neutralisation technique leaves one methodological problem unsolved, because none of the two methods is 'clean' with respect to the switching pattern. While the old method penalised switching by having both patches with equal probability (any switch was a waste of time), the new one positively rewarded switching, by giving one reward after every change of patch. This difference might have contributed to any observed differences in bout structure and the amount of sampling of the two experiments. I therefore will not attempt to make direct comparisons between the two sets of results.

9.2.4. Measures of sampling

9.2.4.1. Decision point

In the first experiment the amount of sampling was established by defining for every test a decision point. When a bird made more than 90 of 100 consecutive hops in the same patch it was said to have 'decided', and the period before that was defined as the sampling period. This criterion was very reliable and seemed to be a very appropriate partition of the tests, since proportion of hops in each patch was markedly different before and after reaching criterion. Because this type of criterion could not be used for tests lasting only 50 hops, it was not possible to use it in the present experiment. I attempted to find a decision criterion which could be used in the short horizon tests by testing the reliability of criteria based on five different numbers of consecutive hops and two degrees of bias in the use of the two patches. The problem is that while reliability increases with length of the criterion (i.e., the longer a criterion the less likely it is that the bird will reverse its preference after 'deciding'), the number of short tests that reach criterion decreases as criteria become longer. In order to measure reliability, I counted the proportion of tests that, having reached criterion, maintained the same bias until the end of the test. Since this requires that the run continues for a reasonably large number of hops after 'decision', reliability was established using long tests only. On the other hand, the proportion of tests reaching criterion was only a problem for short tests. A satisfactory criterion would show high reliability as assessed in long tests and a big

proportion of short tests reaching a decision. The results of this search are shown in Figure 9.2. From this figure I concluded that it was not possible to compare tests in this experiment in terms of a decision point.

9.2.4.2. Commitment

As an alternative to looking for an average decision point in each treatment, I used a measure called 'commitment'. Assume that the birds have an initial period of equal sampling and later switch to exploitation. The average proportion of responses in the better patch at a given time from the beginning of a test should reflect the tendency of the birds to have switched to exploitation by that time. If commitment at a particular stage is close to 50% then the bird is considered to be still sampling, any any value above 50% reflects a tendency towards a decision. Treatments were normally divided for analysis in bins of ten consecutive hops, so that short horizon treatments resulted in 5 sequential bins and long horizon tests in 25 sequential bins. Commitment was established for each of these bins and used as the basic measure of sampling. The effect of patch diversity and of time horizon was analysed by comparing the short horizon tests with the first 50 hops of the long horizon tests.

9.2.5. Experimental sequence

Each of the four birds was given a factorial combination of the two time horizons and the two patch diversities, with four replicates. The experimental sequence for one of the birds is presented as an example in Table 9.1.

9.3. Predictions

9.3.1. Molar Predictions: Global measures of sampling. CONOP predicts the appropriate amount of sampling for any combination of patch

diversity and time horizon. These values, referred to as 'static predictions' can be compared with experimental results. However, total amount of sampling does not provide an easy way to distinguish between the predictions of the different models tested in Chapter 8. A more convenient way of formulating the predictions is to describe what the difference in sampling should be between the treatments. I shall call these 'dynamic predictions'. Different models respond in different ways to variation in patch diversity or time horizon, some of them do not respond at all to these variations, and this means that analysing the dynamics of sampling with respect to patch diversity and time horizon may be more instructing than comparing the absolute amount of sampling. I shall list below some of these dynamic predictions, and will mention in each case how they apply to some of the potential decision rules tested in Chapter 8.

Prediction 1. Effect of experience

Commitment to the better patch should increase with cumulative number of hops. This prediction is common to all rules involving some form of assessment of the environment. It would lead to rejection of Play-the-Winner, since this strategy is absolutely dependent upon the outcome of the last hop.

Prediction 2. Effect of patch diversity

Increase in commitment with experience should be faster when patch diversity is bigger. This prediction reflects the birds' difficulty in assessing the differences in quality between the patches, and like the first prediction is a common prediction of all rules involving learning.

Prediction 3. Effect of time horizon

Commitment should increase faster when time horizon is shorter. This prediction arises because for an animal that is trading off immediate gain for certainty of choosing the better patch the cost of mistakes is more important for long than for short time horizons. None of the simple rules discussed in Chapter 8 makes this prediction although Cullen rules can be modified to fit this behaviour (see Chapter 8). In particular, it is in direct conflict with the hypotheses based on momentary maximisation.

Prediction 4. Interaction of patch diversity and time horizon

The effect of patch diversity should be greater for short horizons, and the effect of time horizons should be greater for small than for big patch diversity. This prediction is also in conflict with all the simple rules based on momentary maximisation It is the prediction with highest predictive information value because it arises from the shape of the responses to patch diversity and time horizon (see Figure 7.6).

Prediction 5. Behaviour in steady state

Regardless of the trajectory to equilibrium, CONOP predicts absolute commitment to the better patch after asymptotic behaviour has been reached (i.e., no differences are expected between treatments). This is also the prediction from IMMAX and from all CULLEN models, but differs from the predictions of PLWIN and of MATCH. Table 9.2 presents the predictions quantitatively.

9.3.2. Molecular Predictions: Pattern of switching between patches.

Prediction 6. Switching pattern and subjective reward probability

Switching between patches during the sampling period according to CONOP or CULLEN rules is aimed at accumulating information and should not be related to experienced reward probability in each patch. This lack of relation between experienced reward probability and switching is directly opposed to the predictions of Immediate maximising. If an animal is always maximising the probability of reward for the next hop, it should only switch when the expected reward probability in the other patch is bigger than that in the current patch. The formal statement of this prediction is: if the birds are using IMMAX, nearly 100% of the switches ought to be in the direction of the patch that has given higher number of rewards per hop since the beginning of the treatment. If instead the birds are following CONOP or one of CULLEN models, the proportion of switches during sampling towards higher reward probability would approach 50%. Notice that the prediction refers to subjective probability and not to nominal probability.

9.4. Results

9.4.1. General Behaviour

The birds usually started a run with a series of one-hop bouts made alternately in each patch. In neutralisation tests this lead to obtaining nearly one reward per hop, but in experimental treatments this was not the case, and the birds soon started to make longer bouts, until a bias in favour of one of the patches emerged.

The data are presented graphically in Figures 9.3 to 9.6. Each of these figures shows the results from one of the treatments, and all tests are presented individually in different panels. The ordinate of each panel represents commitment and ranges 0 to 1, with the horizontal axis at 0.5, the commitment when there is no preference. The abscissa represents successive groups of ten hops.

In more than 90% of the cases (58 out of 64) the tests ended with a bias towards the better patch, but the way in which this bias was reached varied widely. In order to make apparent the general pattern, commitment was averaged for each treatment and group of ten hops and these averages are shown in Figures 9.7 to 9.10. The first part of the long tests can be compared directly with the 50 hops of the short tests, as shown in Figure 9.11.

9.4.2. Test of predictions

9.4.2.1. Global amount of sampling

The static prediction about the optimal balance between sampling and exploitation can be tested by comparing the degree of commitment to the better patch for both CONOP and the birds. The birds' commitment was established as defined above, and the commitment of the model over the first 50 hops can be obtained using the following formula:

$$C_{50} = \frac{50 - M}{50}$$

where C_{50} is commitment over the first 50 hops and M is the number of hops in each patch before decision. Table 9.3 summarises these results

The commitment shown by the birds falls consistently below the predicted by CONOP. The same effect can be seen in Figure 9.12, which shows the benefit (see Chapter 8 for definition) obtained by the birds compared with the theoretical benefit from decision rules investigated last chapter. The lower payoff obtained by the birds could result from (a) making more mistakes, i.e. choosing the wrong patch more often than the models; (b) taking longer to make up their minds than the models, or (c) failing to pay 100% attention to the chosen patch after having reached a decision.

The number of mistakes does undoubtedly play some role, since the birds ended up preferring the wrong patch in a greater number of cases than expected from the models. The second possibility cannot be tested because I cannot define a decision point, but it is worth pointing out that some effect of this type may be caused by the need to 'neutralise' the birds' preference between treatments. The prior hypotheses in the simulations were chosen so as to be 'weak' hypotheses, rapidly changeable by experience. The birds, however, experienced tests with equal probability between treatments, and may have taken more sampling to accept a difference as true.

The models' lack of 'memory' between tests is a help in this experiment. The third possibility refers to behaviour under steady state, and has been discussed in Chapter 6 and under Prediction5 in this chapter. I shall return to this problem when discussing Prediction 5 below.

9.4.2.2. Difference in sampling between treatments

9.4.2.2.1. Graphical analysis

The qualitative analysis of the agreement between dynamic predictions and experimental results can be done by examining Figure 9.11. The four predictions are examined below in terms of what they state about the results plotted in this Figure.

Prediction 1: According to this prediction, all graphs in Figure 9.11 ought to show increase in commitment with number of bins. This is clearly true for all cases.

Prediction 2: The two graphs in the left column (big patch diversity) ought to show greater commitment than the two graphs on the right column (small patch diversity). This prediction is also fulfilled.

Prediction 3: The graphs in the bottom row (long term horizon) should present less commitment than the graphs in the top row (short time horizon). This prediction is met in the right column, corresponding to small patch diversity tests, but it is not so obvious in the left column, where patch diversity is big.

Prediction 4: Because of the interaction between patch diversity and time horizon the effect of patch diversity should be more intense in the bottom than in the top row, and the effect of time horizon should be stronger in the right than the left column. This prediction is completely supported by the data.

Prediction 5: This is a prediction about asymptotic behaviour, after sampling is over. According to CONOP, IMMAX and Cullen rules regardless of the trajectory towards decision, commitment after this point should be total. MATCH, instead, predicts different degrees of commitments under steady state for the two patch diversities.

Unfortunately, there are problems in using the available data for testing predictions about steady state behaviour. Since it is not possible to detect the decision point reliably, the steady state must be defined as the period after which bias reaches asymptotic value. This is definitely not possible in the short horizon tests, and this reduces the sample size by a half. By looking at the graphs in Figure 9.9 and 9.10, the situation seems better for the long horizon tests. Average commitment seems to have reached its maximum values after 30 hops in the big patch diversity tests and after about 100 hops in the small patch diversity tests. The asymptotic values appear to be much closer to the prediction of matching than of maximising (traced as broken lines in Figures 9.9 and 9.10).

9.4.2.2.2. Statistical analysis

Predictions 1 to 4 were tested by means of an analysis of variance with quintile of tests, patch diversity and time horizon as factors, on the basis of the design of 4 subjects and 4 replicates in each subject.

Table 9.4 shows the results of this analysis. The main effects of quintiles (fifths) and patch diversity (predictions 1 and 2) were significant at the 0.001 level. The main effect of time horizon approached significance ($p < 0.10$) (prediction 3) and the interaction between patch diversity and time horizon (prediction 4) is significant at the 0.01 level.

9.4.2.3. Pattern of switching between patches

Prediction 6: In order to test if there was a relation between each individual switch between patches and the reward probability

experienced so far in that test, I computed atthe time of each switch the number of rewards and hops in the current patch and in the other one. The subjective reward probability of each patch was established using the beta distribution with $\alpha=\beta = 0$. If the first patch had lower subjective probability of reward than the second, the switch was called a switch 'up', otherwise it was called a switch 'down'. The proportion of switch 'up' with respect to total number of switches did not differ significantly from 50% for any of the birds. The percentage of switches 'up' was established for each test, and the 16 scores for each subject were compared with 50% using a binomial test.

9.5. Discussion

The results of this experiment are conclusive in showing that great tits can adjust their sampling behaviour changes in time horizon. Their performance was poorer than the constrained optimal solution, and at least two factors were responsible for this: (a) the birds preferred the 'wrong' patch on occasions when the dynamic programming algorithm would not have done so and (b) they did not concentrate all their effort in the better patch after reaching steady state behaviour. The experiment showed a strong effect of learning, rejecting the possibility that their behaviour might be controlled by immediate factors as assumed in the play-the-winner strategy. The hypothesis of momentary maximisation (IMMAX) is not the sole mechanism of patch choice. This is because (a) the birds responded to time horizon as predicted by a global maximisation hypothesis, and (b) they did not tend to switch between patches in the direction of in-

creasing reward probability.

It must be pointed out though, that the birds tested in Chapter 7 did show some component of immediate maximisation (Figure 7.8).

Probability matching (MATCH) cannot account for the adaptation to time horizon, but there was some evidence in its favour in that the degree of commitment to the better patch differed between patch diversities when behaviour seemed to have reached an asymptote. This differs with the trend observed in Chapter 7 (Figure 7.3), where most of the birds approached maximisation. However, the rigorous test of steady state behaviour requires longer tests than those used here. Simple Cullen rules (see Chapter 8) could not account for the data given the birds' ability to adjust to different time horizons, but the possibility that some form of Cullen rule with an adjustable parameter might be used by the tits cannot be rejected.

Table 9.1

Experimental design

Block	Aim	Length of runs (hops)	Recording	Runs
1	Acclimatisation to short horizon	50	No	15-35, N, 10-40,N 10-40, N, 15-35,N 10-40, N, 15-35,N 10-40, N, 15-35.
2	Test sampling in short horizon	50	Yes	10-40, N, 15-35, N 10-40, N, 15-35, N 15-35, N, 10-40, N 15-35, N, 10-40.
3	Acclimatisation to long horizon	250	No	10-40, N, 15-35, N 10-40, N, 15-35, N 10-40, N, 15-35, N 15-35, N 10-40.
4	Test sampling in long horizon	250	Yes	10-40, N, 15-35, N 15-35, N, 10-40, N 10-40, N, 15-35, N 15-35, N, 10-40.

Example of the experimental sequence for one of the four birds. Two birds were treated first with short and later with long horizons and the other two in reversal order. The succession of the two patch diversities within a block was randomised, but birds were always "neutralised" between two tests.

Table 9.2

Proportion of hops in the better patch predicted by maximising (CONOP, IMMAX or CULLEN) and by probability matching (MATCH) for the period after decision.

	Maximising	Matching
Reward probabilities		
0.15-0.35	100%	70%
0.10-0.40	100%	80%
Difference	0%	10%

Table 9.3

Proportion of hops in the better patch during the first 50 hops. The top figure in each cell is the experimental result and the bottom one is the model's prediction.

PATCH DIVERSITY	TIME HORIZON 50 hops	250 hops	50-250 difference
0.20	64.9 ± 3.3 S.E 84.0	54.7± 4.2 S.E 74.0	10.2 10.0
0.30	71.6 ± 2.6 S.E 89.0	73.9 ±3.1 S.E 83.0	-2.3 6.0
Difference 0.20-0.30	6.7 5.0	19.2 9.0	

Table 9.4

Analysis of variance

Number of hops in the better patch in ten hop bins up to 50.

Source	DF	SS	MS	F	p
Bird	3	104.5	34.8	10.2	0.001
Patch diversity	1	127.5	127.5	37.5	0.001
Time horizon	1	11.2	11.2	3.3	Ns
Fifth of test	4	160.9	40.2	11.8	0.001
Div x Hor	1	33.8	33.8	9.9	0.01
Div x Fifth	4	31.7	7.9	2.3	Ns
Hor x Fifth	4	3.3	0.8	0.2	Ns
Div x Hor x Fifth	4	1.0	0.2	0.1	Ns
Final Residual	297	1009.5	3.4		
Grand Total	319	1435.5			

Figure 9.1 Proportion of rewarded hops experienced by each subject (letters on the right) during the first 4 runs when only the first hop after a switch was rewarde. All birds scored over 50% by the fourth session. This means that in more than 50% of the feeding visits they switched after a single hop in a patch.

Figure 9.1

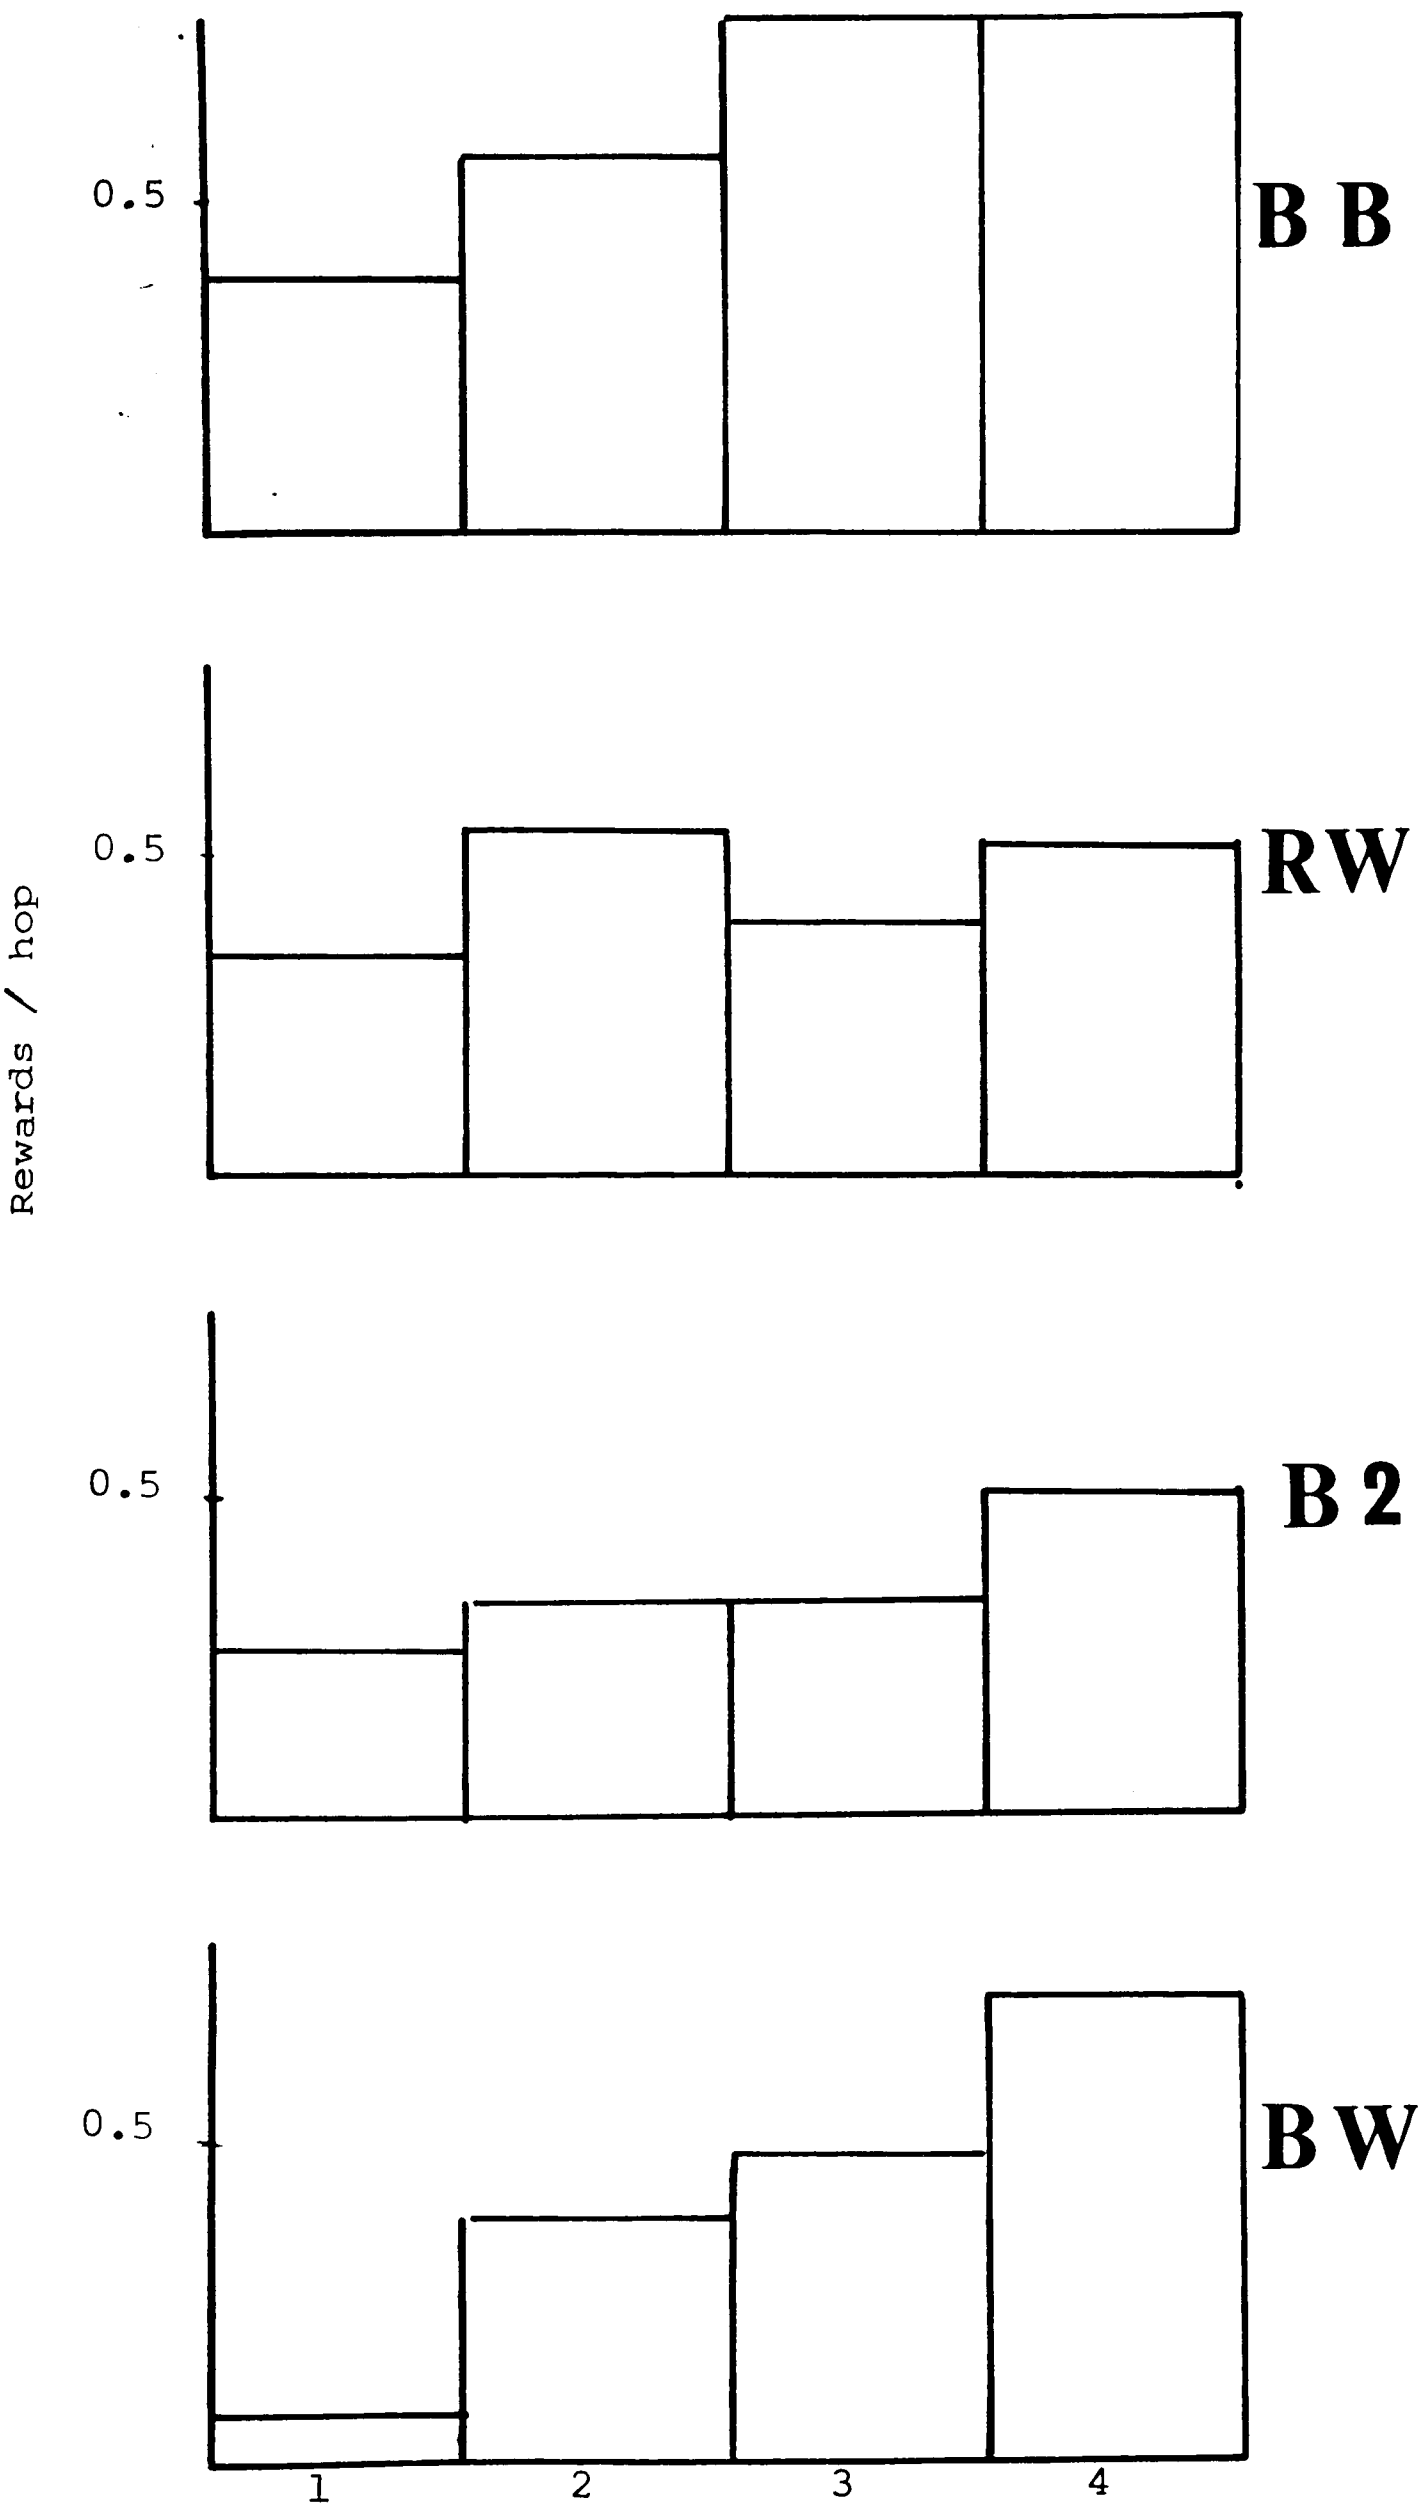

0.5
B B
0.5
RW
Rewards / hop
0.5
B 2
0.5
B W
1
2
3
4
Successive training sessions

Figure 9.2

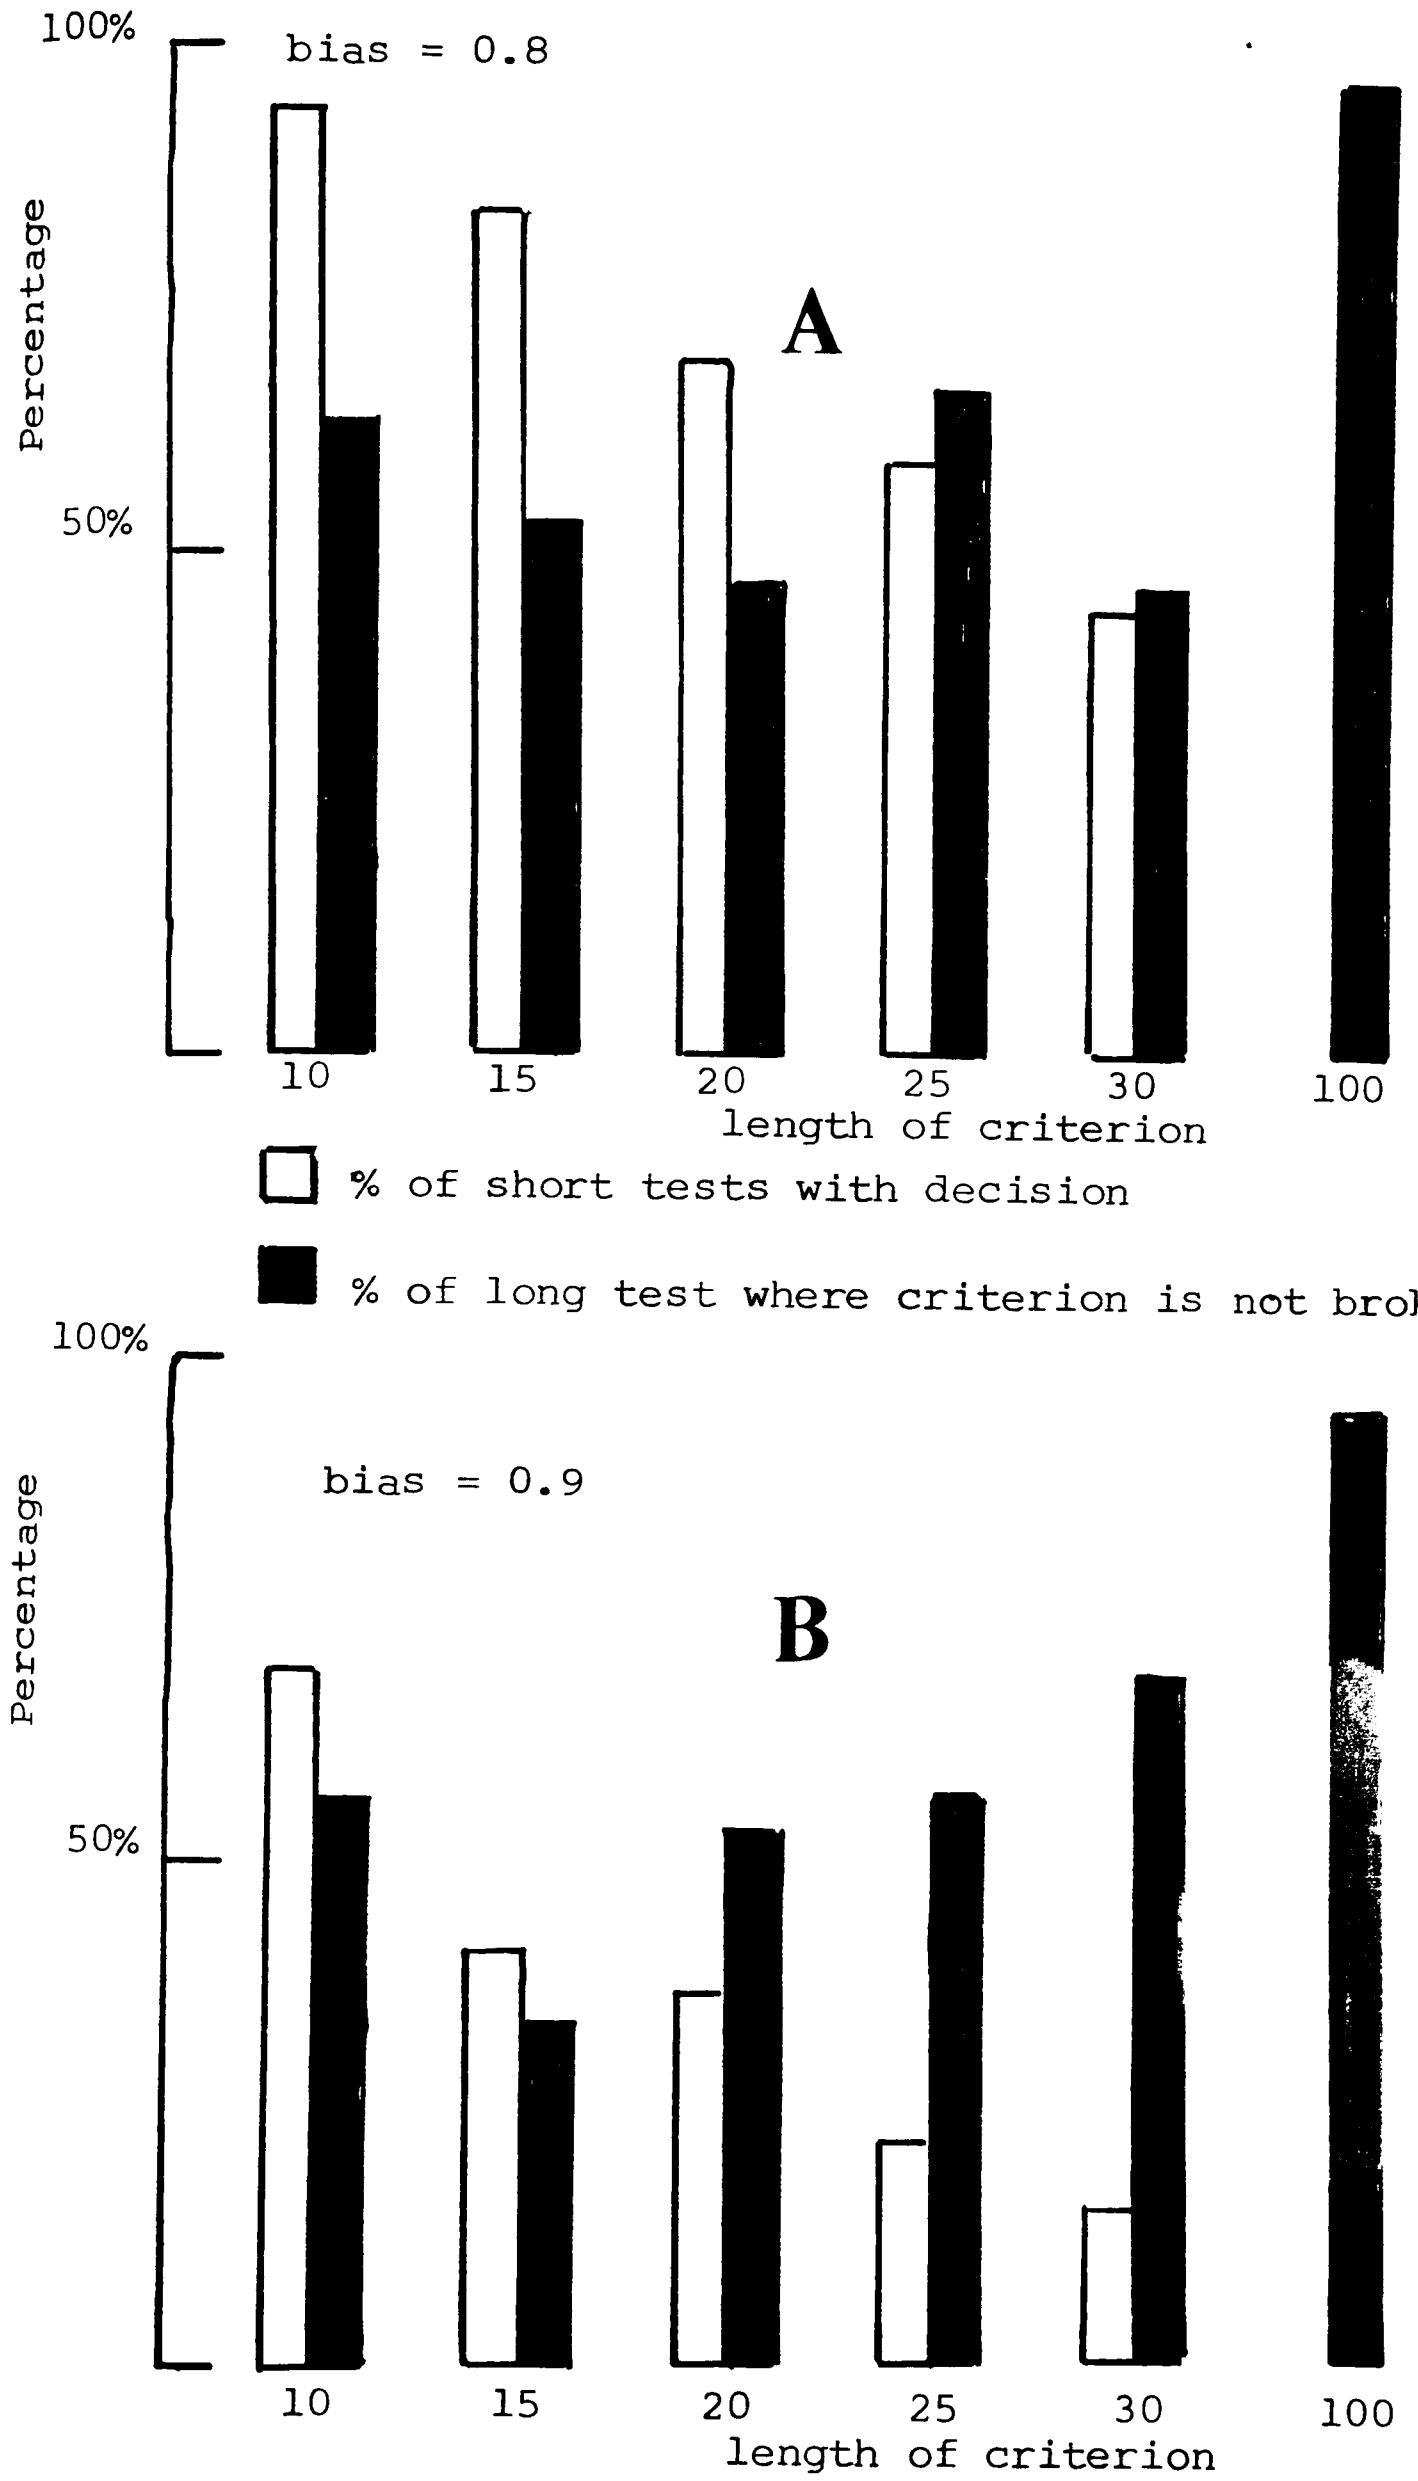

100%
bias = 0.8
Percentage
50%
A
10
15
20
25
30
100
length of criterion
% of short tests with decision
% of long test where criterion is not broken
100%
bias = 0.9
Percentage
50%
B
10
15
20
25
30
100
length of criterion

Figure 9.3 Proportion of hops in the better patch in tests with short horizon and small patch diversity. Each panel presents data from one individual test. The 50 hops of each test were divided into bins of 10 consecutive hops each, and the proportion of hops in the better patch is plotted as the height of the bar for each fifth of test. The ordinate ranges 0 to 1, with the abscissa traced at 0.5, the level of equal hopping in both sides. The number in each panel identify the bird.

Figure 9.4 Proportion of hops in the better patch in tests with short horizon and large patch diversity. Details as in Figure 9.3.

Figure 9.3

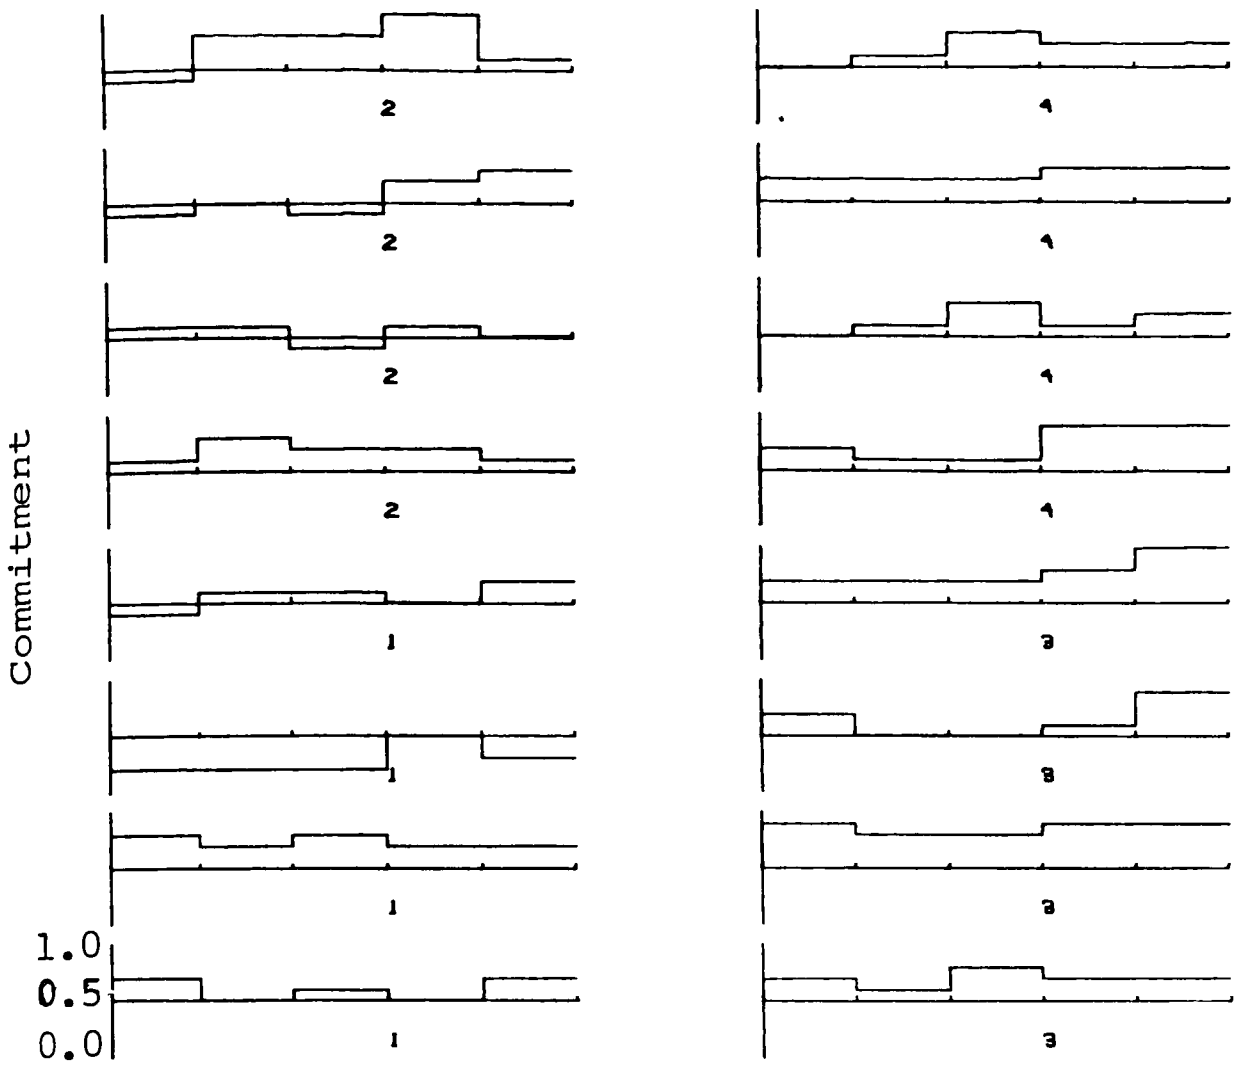


Figure 9.4

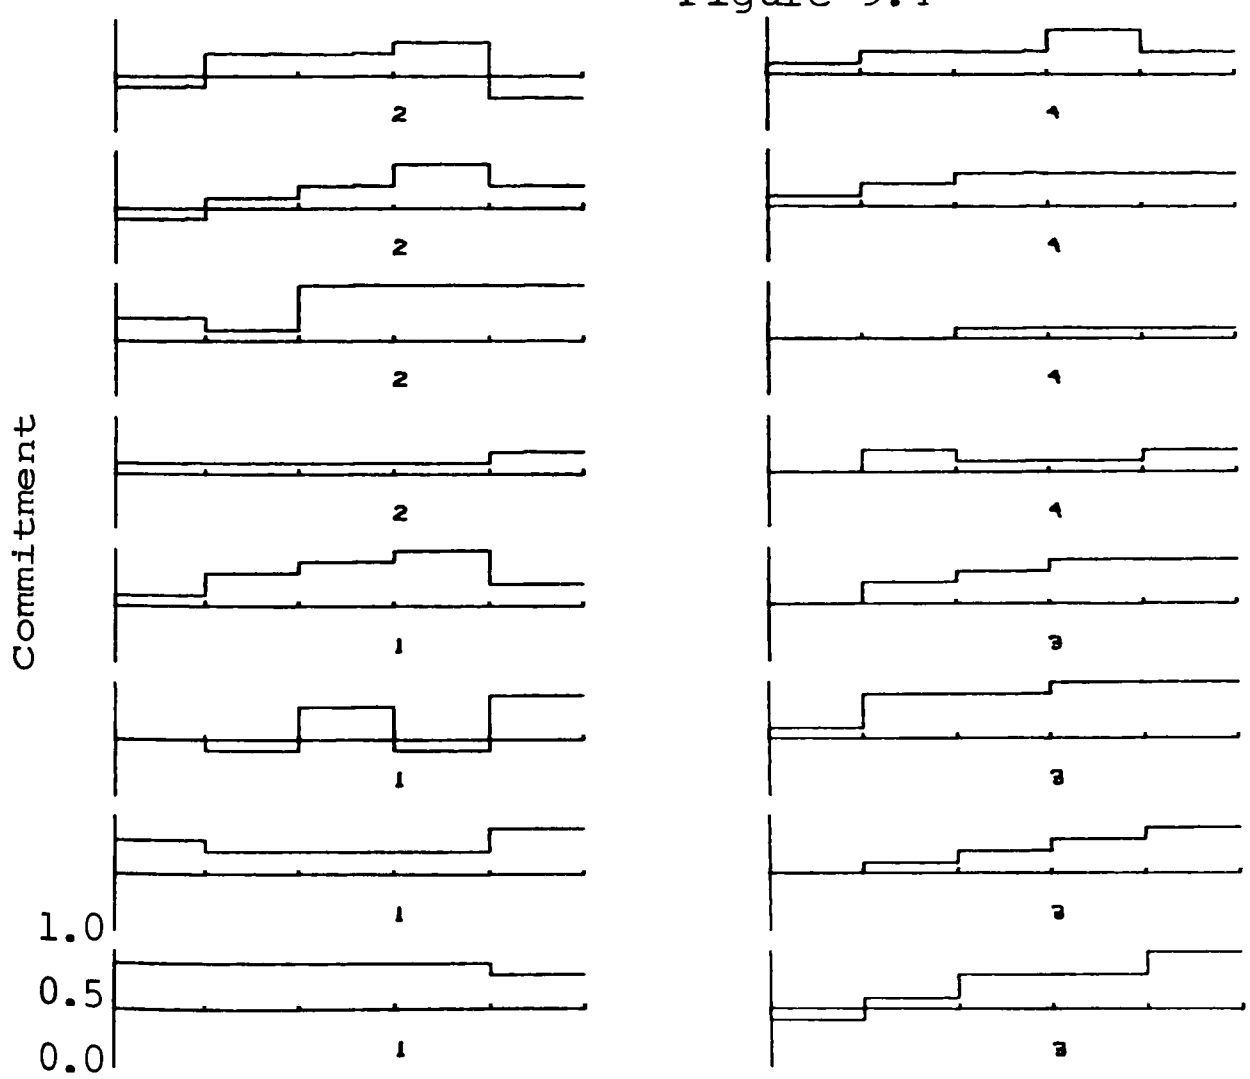


Figure 9.5 Proportion of hops in the better patch in tests with long horizon and small patch diversity. Details as in Figure 9.3 with the exception of the number of bins in the abscissa (25 in this graph as opposed to 5 in the short tests).

Figure 9.6 Proportion of hops in the better patch in tests with long horizon and large patch diversity. Details as in Figure 9.5.

Figure 9.5

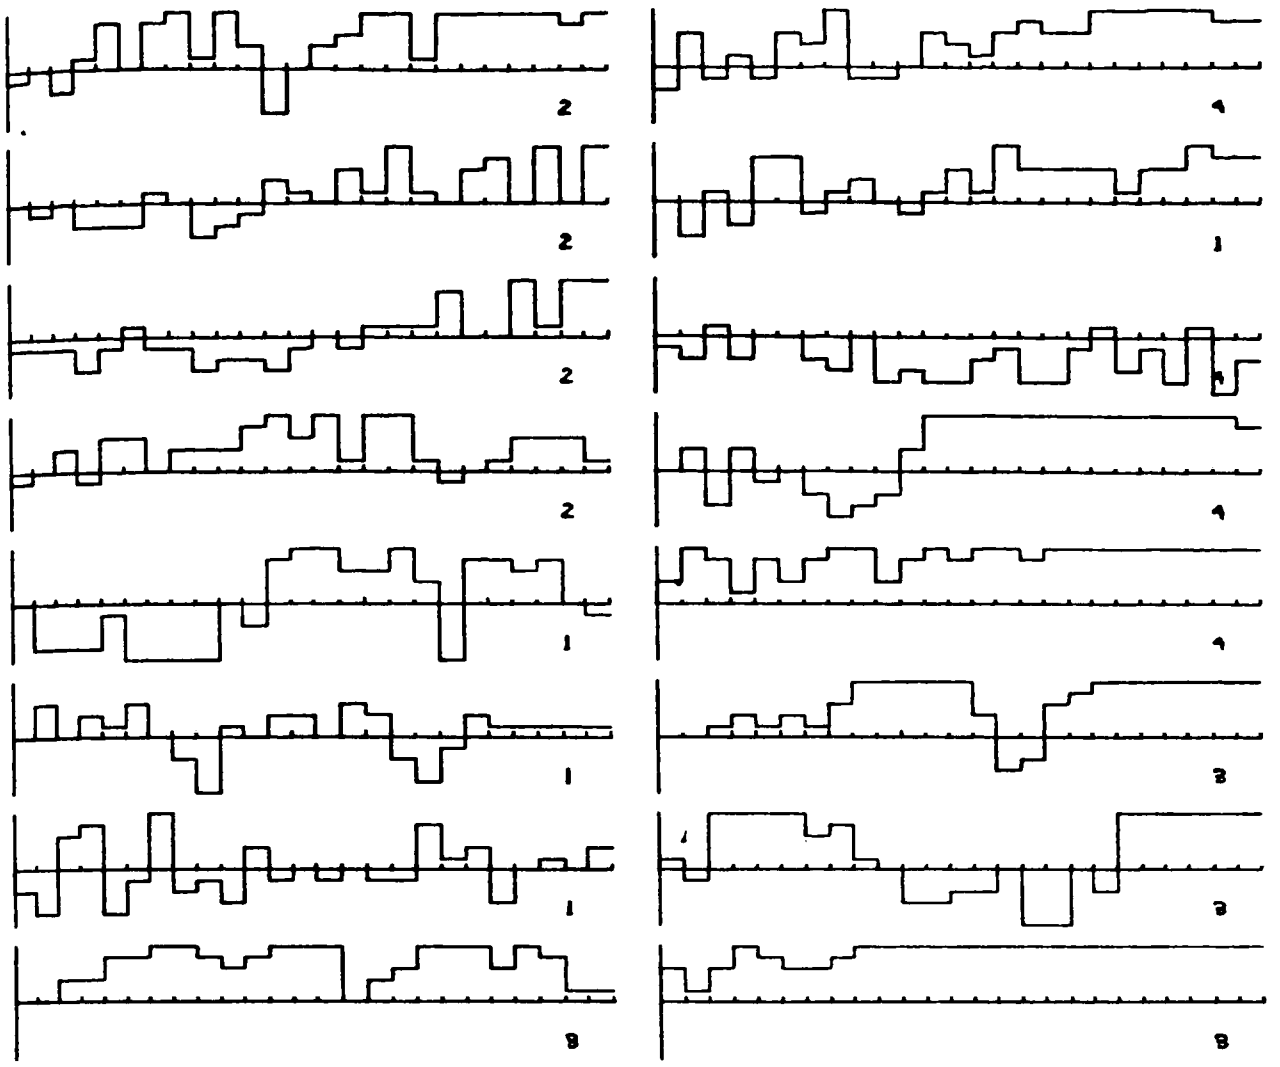


Figure 9.6

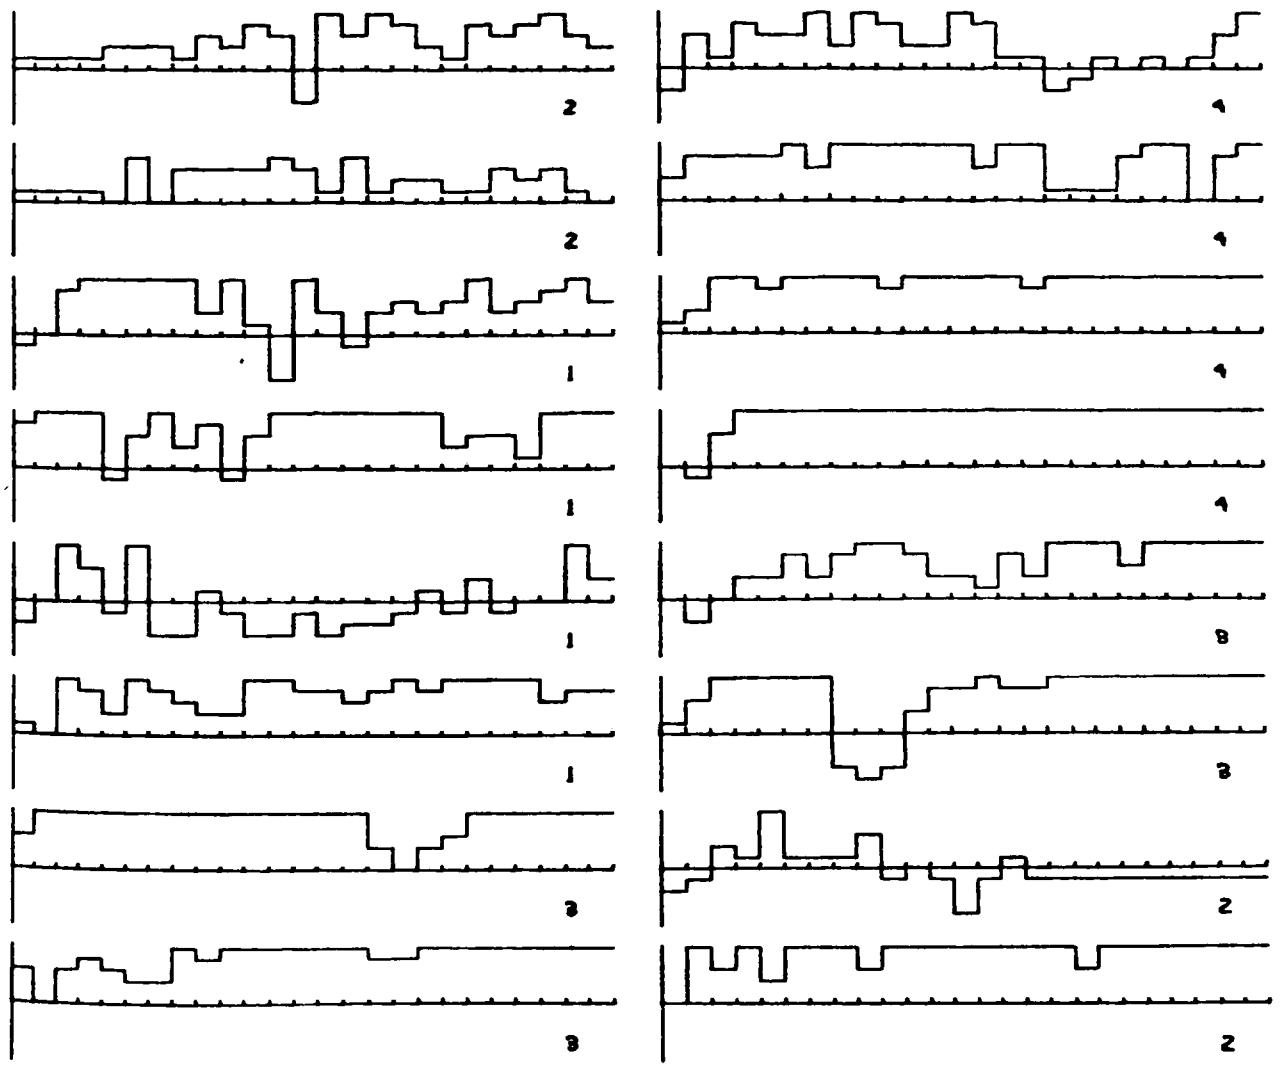


Figure 9.7 Mean proportion of hops in the better patch in tests with short horizon and small patch diversity. The data shown in Figure 9.3 were averaged for each fifth of test and the average is plotted as in previous figures.

Figure 9.8 Mean proportion of hops in the better patch in tests with short horizon and large patch diversity. The data shown in Figure 9.4 were averaged for each fifth of test and presented in the same way as in previous figures.

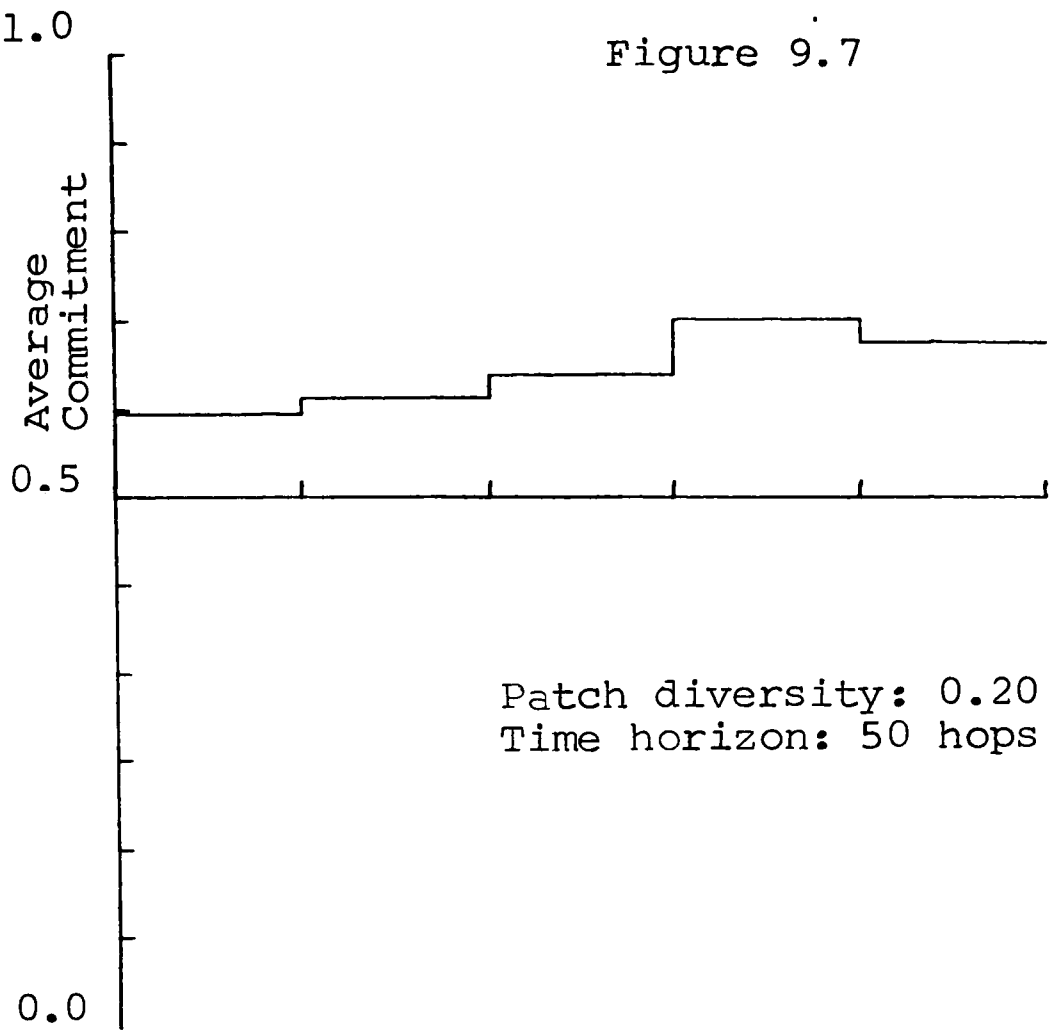


Figure 9.7

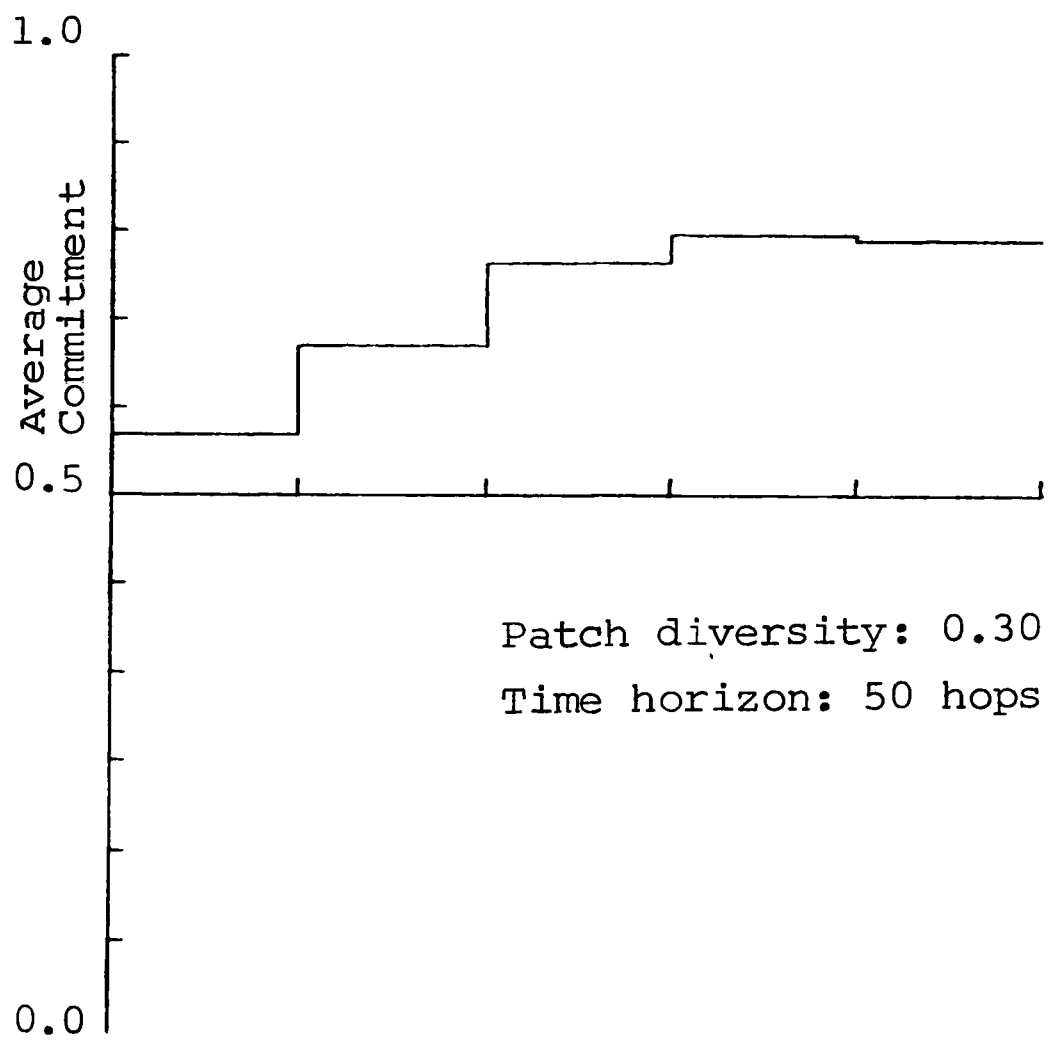


Figure 9.8

Figure 9.9 Mean proportion of hops in the better patch in tests with long horizon and small patch diversity. Each bin results from averaging the values in Figure 9.5

Figure 9.10 Mean proportion of hops in tests with long horizon and large patch diversity. Bins are the mean of all tests presented in Figure 9.6, plotted as in previous figures.

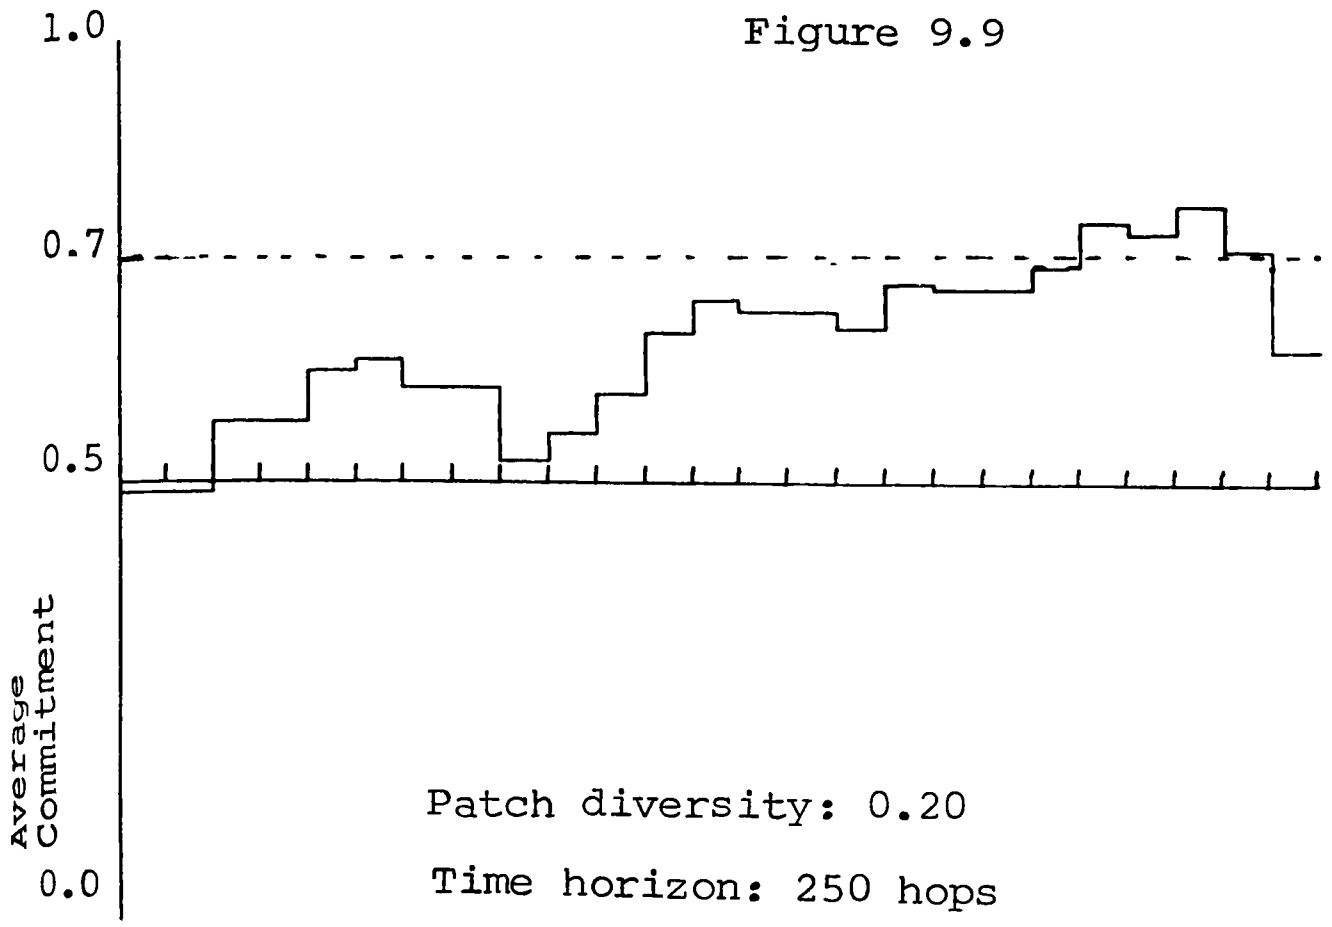


Figure 9.9

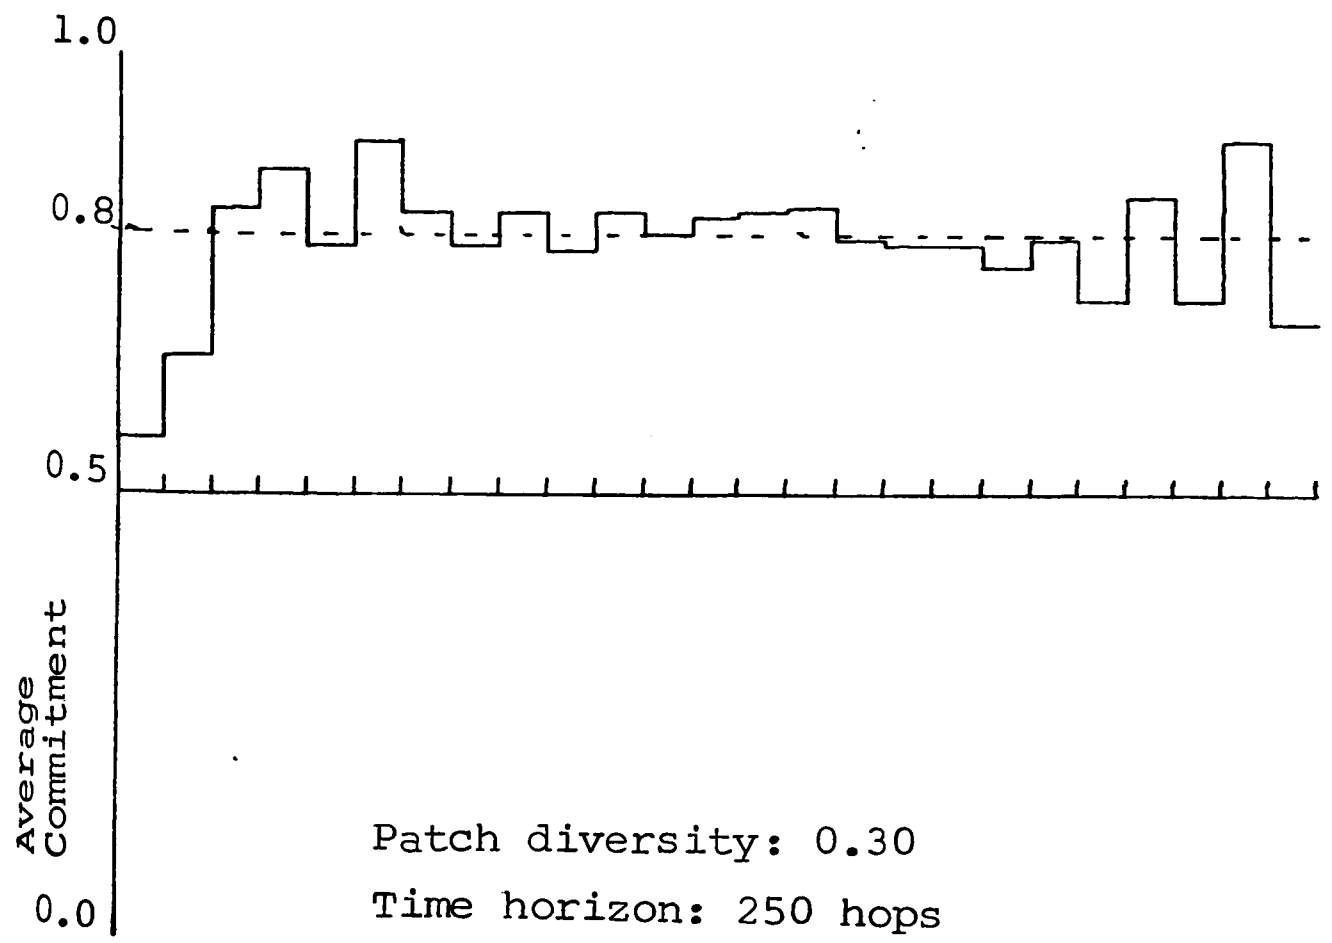


Figure 9.10

Figure 9.11 Mean proportion of hops in the better patch in the first 50 hops of all tests. A and B are identical to Figures 9.5 and 9.6. C and D are the same as in Figures 9.9 and 9.10 but only the first 50 hops were plotted, with the same scale as A and B.

Figure 9.11

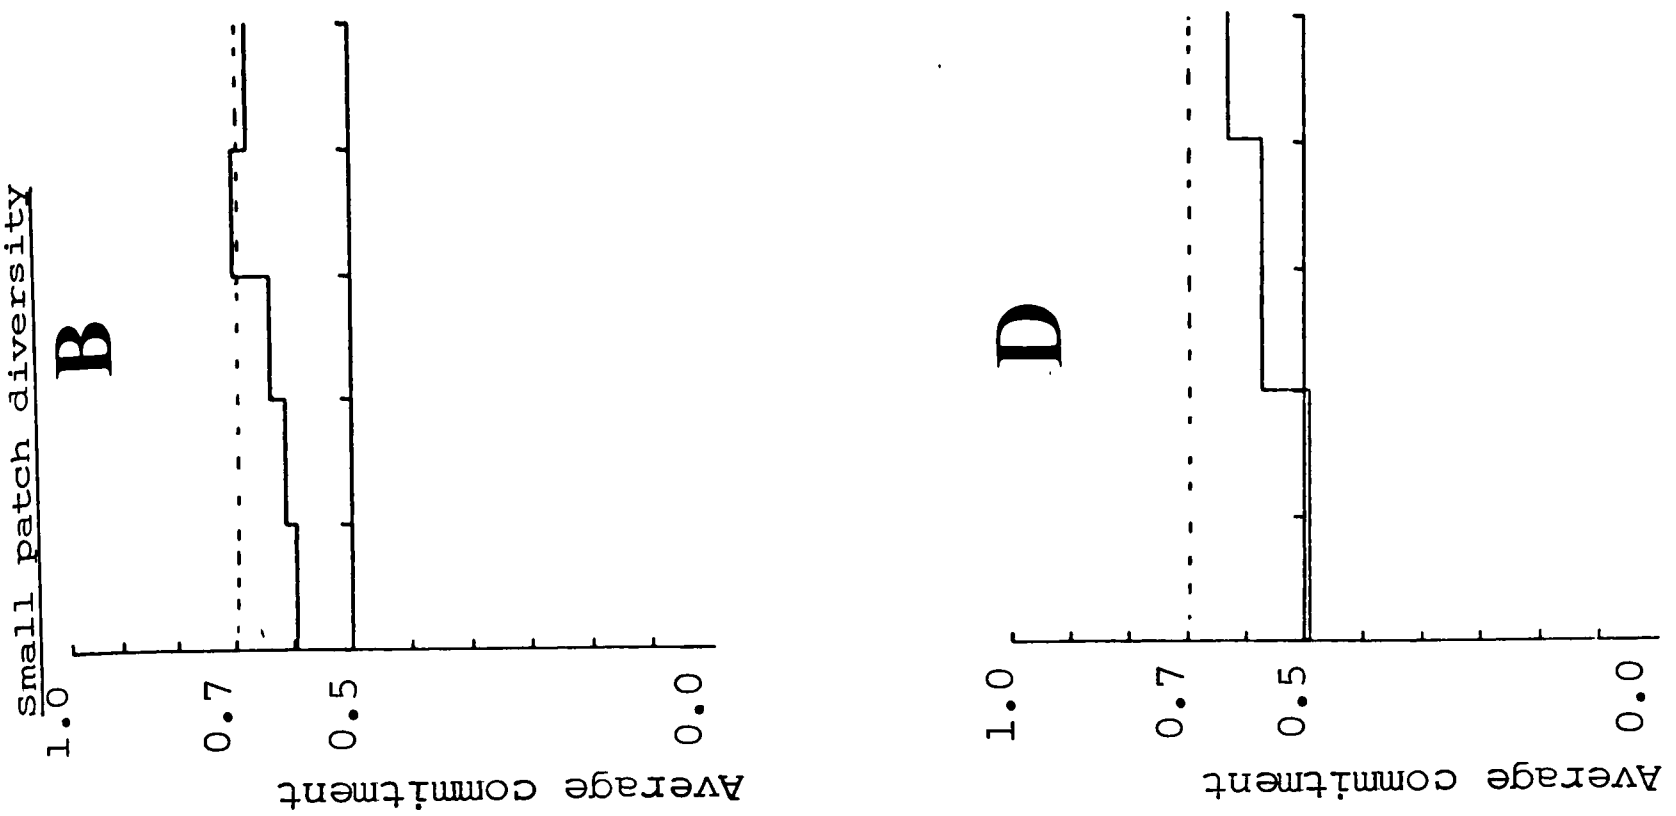
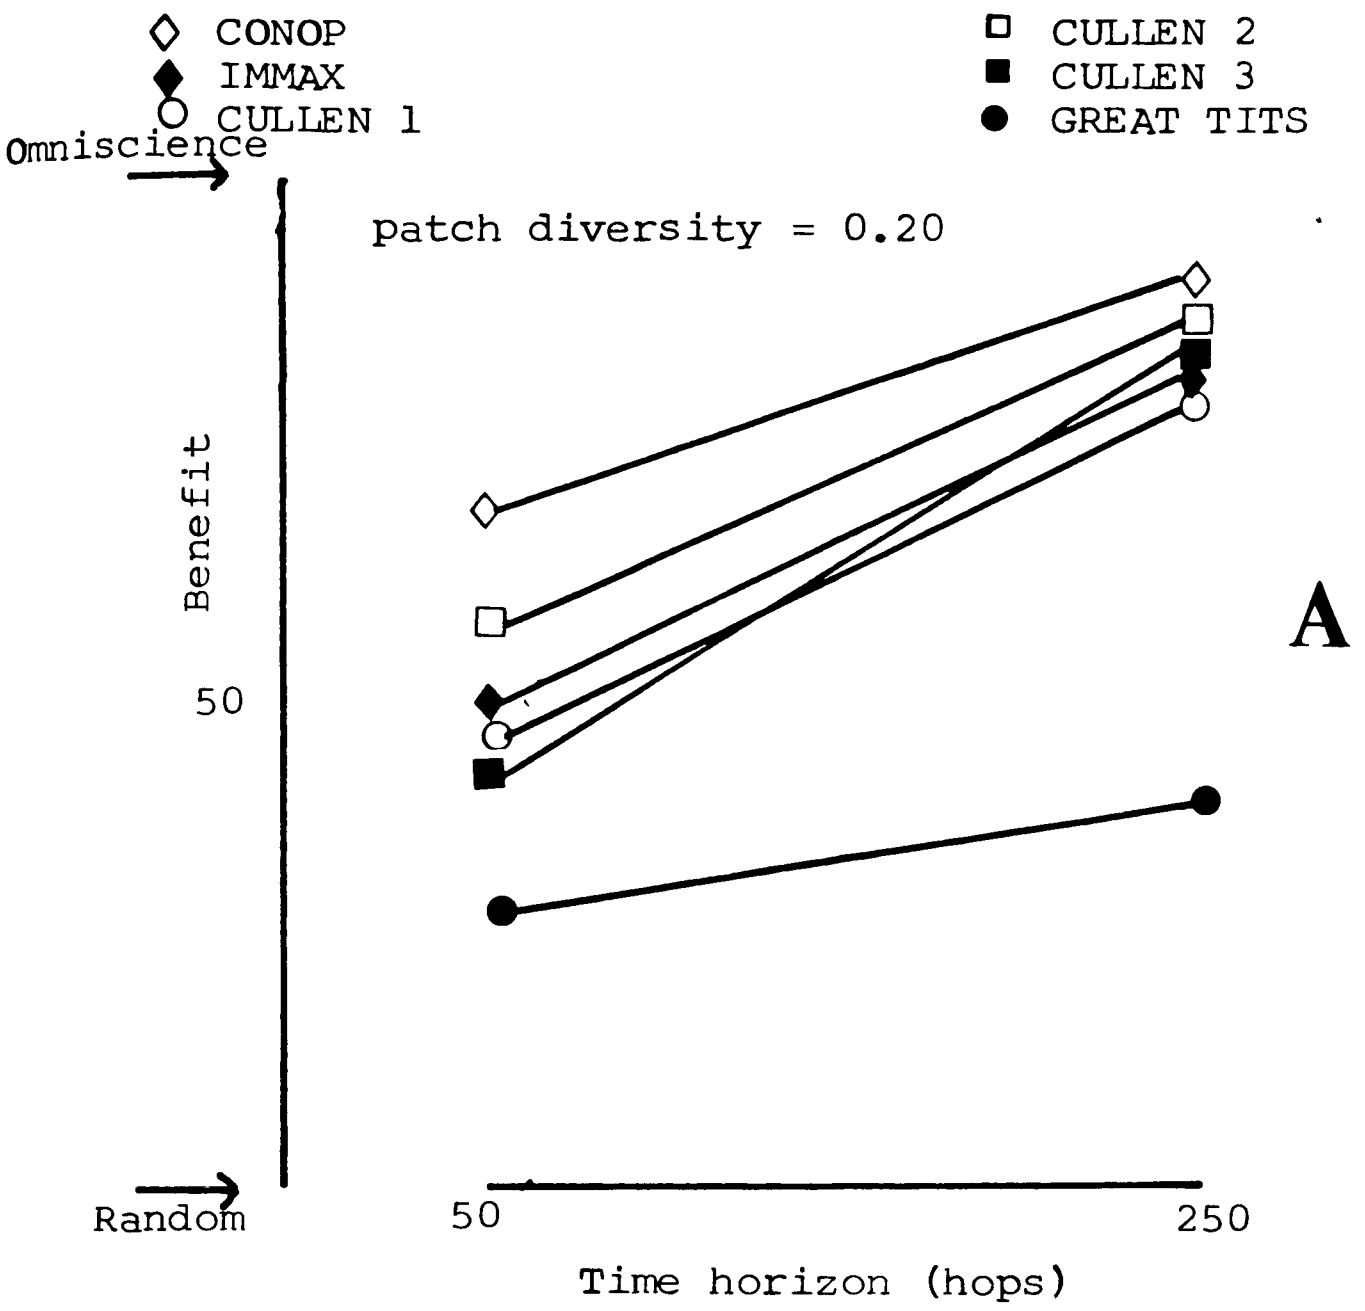

Small patch diversity
B
1.0
0.7
0.5
0.0
Average commitment
D
1.0
0.7
0.5
0.0
Average commitment
Large patch diversity
A
1.0
0.8
0.5
0.0
Average commitment
Short time horizon
C
1.0
0.8
0.5
0.0
Average commitment
Long time horizon

Figure 9.12 Benefit (see definition in chapter 8) obtained by the birds and results from simulations with decision rules investigated in Chapter 8. All the rules obtained higher benefit than the great tits.

Figure 9.12


CONOP
IMMAX
CULLEN 1
CULLEN 2
CULLEN 3
GREAT TITS
Omniscience
patch diversity = 0.20
Benefit
50
A
Random
50
250
Time horizon (hops)
Omniscience
patch diversity = 0.30
Benefit
50
B
Random
50
250
Time horizon (hops)

PART III

The trade-off between foraging and territorial defence

CHAPTER 10

Optimal foraging and territorial defence

10.1. Introduction

In this chapter I describe an experiment designed to investigate the trade-off between foraging behaviour and territorial defence. Most optimal foraging models have been based on the assumption that the foraging decisions maximise the net rate of food intake (see Pyke et al, 1977; Krebs, 1978 ; Staddon, 1979 for reviews of optimal foraging theory) or net rate of intake subject to nutrient constraints (Pulliam, 1975; Belovsky, 1978). However, there may often be a cost associated with feeding, for example because a foraging animal is less vigilant for predators. If this cost is related to the rate of feeding, in such a way that the faster an individual feeds the higher is the cost, it may be advantageous not to maximise the rate of food intake but instead to adopt a policy which achieves an optimal trade-off between foraging and, for example, scanning for predators (Milinski and Heller, 1978; McCleery, 1977, 1978). Milinski and Heller (1978) conducted an experiment in which they tried to measure this trade-off. In their experiment, a stickleback (Gasterosteus aculeatus) was released in a tank containing transparent tubes with various densities of water fleas (Daphnia magna) in them. Normally the fish would direct its initial attacks preferentially towards the denser swarms, but when a model of an avian predator (a silhouette of a kingfisher) was flown above the tank before the test, the fish directed a higher proportion of the first ten attacks to the containers with lower densities of water fleas. Milinski

and Heller assume that foraging on denser swarms leads to higher rate of intake, but because of the confusion due to the swarm density, the fish is less able to detect predators. Thus, by changing the foraging behaviour so as to sacrifice some feeding efficiency in exchange for a gain in ability to detect potential predators, the sticklebacks attain a global optimum in terms of fitness. There are however several difficulties in the interpretation of this experiment. The most serious problem is that the confusion effect due to high density of the prey is assumed to act so as to decrease ability to detect predators but not ability to capture prey. If the second effect of prey density is as important as the first, then the fish preference for less dense swarms may be due to purely feeding factors. In fact, when all the attacks in a test are considered-instead of just the first ten - both distrubed and undisturbed fish prefer to attack the lower density swarms. The effect of the predator on total number of attacks is a general decrease of number, without changing the distribution of attacks between different densities of the prey. The difference between the effect of the predator on the first ten bites and the total nubmer of bites is explained by the authors using an optimality model in which they assume that the fish are less hungry at the end of the test and in consequence are more willing to sacrifice food intake to pay more attention to predators. Their model cannot be directly applied to their data, however, because the sticklebacks were not allowed to eat during the tests, and the authors had to rely on a post-hoc hypothesis (that tendency to bite against the transparent containers decays in a similar

fashion to feeding tendency) to compare the model's predictions with their results. These difficulties could be solved by measuring the changes of feeding rate and latency to react to a predator when sticklebacks forage on swarms of various densities. Until this information is available, the hypothesis of an optimal trade-off between feeding rate and vigilance cannot be considered to have been tested, although it is consistent with Milinski and Heller's results.

Another study of the trade-off between feeding and a competing activity is that of Sibly and McFarland (1976). In their experiment, a Barbary dove Streptopelia risoria was allowed to work for food and water rewards in a Skinner box after being deprived of food and water for various durations. They assumed that the cost associated with a given state of deprivation is proportional to the square of the deficit, and on those grounds they predicted the alternation between feeding and drinking that would minimise cost for every particular combination of hunger and thirst. Their results matched their predictions very convincingly, but of course, the exercise would be much more convincing if the associated costs could have been measured instead of guessed.

In this chapter I shall report an experiment in which great tits were exposed to a conflict between maximising rate of food intake and patrolling their territory for intruders. While the experiment was done entirely indoors, this conflict is likely to be experienced by great tits in the wild because they usually forage on or near the ground, where visibility is poor (Hunter,

1978). I predicted that when the risk of territorial intrusion is low, the birds would behave according to a simple model of maximising intake, but that when intruders had been seen recently the birds would adjust their behaviour to a compromise between foraging and territorial vigilance.

In my experiment I tried to avoid some of the problems of interpretation in previous studies. The intake maximising strategy was predicted independnetly of the birds' behaviour, and the predicted behavioural changes were certain to lead to higher detection of the invaders at the expense of feeding efficiency. It was not possible to measure the risks associated with each behaviour and as a consequence I am unable to formulate a quantitative prediction of exactly how much should the behaviour change. An inverse optimality (McFarland, 1977; Houston, 1977) approach makes it possible to formulate an hypothesis about the possible risks.

10.2. Methods

10.2.1. The birds

I used four territorial male great tits. Three of them had participated in the experiment described in Chapter 9 (they will be named in this chapter Pontryagin, Nahuel and Hamilton) and were in consequence fully territorial and trained to use the feeding discs (see Chapter 1). The fourth bird (named McNewan) was experimentally naive. This bird became territorial in the aviary very soon because in the wild it had held a territory including the roof of the building where the outdoor aviaries were. He was trained to use the feeding discs using the technique des-

cribed in Chapter 1 and after this received the same treatment as the other animals.

10.2.2. The experimental set up

Figure 10.1 presents a diagram of the aviary layout. There were two feeding patches, each consisting of a feeding disc mounted on a platform one meter high. A T-shaped wall was built to divide the aviary as indicated in the diagram. The leg to the partition that separated both patches was taller than the section separating each patch from the central space. The birds had to go over the partition in order to fly from one patch to the other, but they could see the other patch through the wire mesh window in the central leg of the partition.

The trapdoors used in all previous experiments to separate the experimental aviaries from the outdoor compartment were replaced by Intruder Boxes (described in Section 1.11), designed to make it possible to use one of the resident territorial males as an 'intruder' when other birds were in the indoor aviary. The connection between the inside and outside aviaries was controlled by sliding doors.

The intruder boxes could not be seen from the feeding patches but a foraging bird could see them when switching between the patches by flying over the partition (see Figure 10.1).

10.2.3. Foraging schedules

On arrival at a patch, the bird could obtain rewards after the 2nd, 3rd, 5th, 12th, 28th and 60th hop, and the sequence started again after every switch between patches. Figure 10.2 presents this 'progressive ratio' schedule as a cumulative record of rewards

after a switch between patches. This schedule simulated an environment where, (a) the activity of the predator produces resource depression (Charnov, Orians and Hyatt, 1976), either by the removal of the prey or by behavioural responses of the prey and (b) at any given stage of depression of the resource the time to obtain a food item is roughly proportional to the rate of work of the predator.

When expressed as rewards per number of hops, this schedule can give up to a maximum rate of reward of 0.67 reward per hop, if the subject swtiches always after the 3rd hop in a patch. When first exposed to this schedule, the birds tended to stay too long in each patch, but soon learned to switch back and forth to keep the reward rate high. They were trained until they reached asymptotic performance before starting the experiment (the asymptote being defined as when the rate of rewards per hop did not vary more than 0.10 rewards per hop in six consecutive sessions).

10.2.4. Experimental sessions

There were two types of sessions: (a) control or 'undisturbed' and (b) 'intrusion' sessions. In a control session a bird was allowed to enter the experimental aviary and forage for 10 minutes. The door of the intruder box was operated four times but no intruder was visible to the foraging animal. In intrusion sessions an 'intrusion' by another territorial male was simulated (see section 1.11) during the first minute of the test. The subject was left to forage for nine more minutes, with a short bout of song being played back within the fifth minute of the test. I

assume that the effect of this treatment is to make the foraging animal behave as if the probability of intrusions is higher on intrusion than on control sessions. Each bird was tested 16 times: ten undisturbed sessions and six intrusion sessions, in a random order.

10.3. Predictions

10.3.1. The behaviour on control sessions

The foraging behaviour of undisturbed birds can be predicted assuming that they tend to maximise the rate of prey captured per unit of time. The capture rate can be calculated using the following equation:

$$\frac{E}{T} = \frac{Np}{TT + (Np \cdot Ht) + (Nh \cdot IHI)} \qquad (1)$$

where $\frac{E}{T}$ is the rate of captures per second foraging:

TT is the travel time between patches

Ht is the time required to handle one prey

IHI is the inter hop interval, or the time required to make one unrewarded hop

Np is the cumulated number of prey eaten since arrival to the patch

Nh is the number of unrewarded hops since arrival to the patch.

The total number of hops since the last switch, or bout length can be calculated as:

Bout length = Np + Nh

TT, Ht and IHI are assumed to be constraints for each individual, and can be measured empirically. Once they are known, it is easy

to calculate the optimal bout length by evaluating (1) following the depleting schedule described above. In this form, the prediction about the behaviour of the undisturbed animal takes the form of a particular bout length, that is, the bout length that maximises Equation 1.

10.3.2. The behaviour in intrusion sessions

Under a greater risk of intrusion the foraging behaviour of a territorial animal should change so as to increase the chances that an intruder will be detected. In this particular set up this implies switching between the two patches more often than the foraging optimum. The specific prediction is that during intrusion sessions the birds ought to forage in smaller bouts than in undisturbed sessions.

10.4. Results

10.4.1. Do undisturbed birds forage optimally?

Table 10.1 shows the mean value of TT, Ht and IHI for each bird obtained from the control sessions. Using the values in Table 10.1 it is possible to assess the payoff of all possible bout lengths. This is done graphically in Figure 10.3. Bout lengths which end in a reward (2,3,5,12) appear in the graph as local maxima. This is because when the animal is operating in a deterministically depleting schedule like this, the only 'rational' policy is to leave after a reward. This is strictly dependent on the deterministic character of the schedule, and must be considered as an artefact in relation to real patches with stochastic depleting processes.

I have until now referred to bout length in terms of number of hops, but this is translatable into time units using the graphic approach developed by Charnov (1976b) and Parker and Stuart (1976). This method is used in Figure 10.4 which shows the cumulative gain as a function of time spent in a patch. It is clear from Figures 10.3 and 10.4 that the most profitable bout length for all individuals is five hops, that is, they ought to switch between patches after having obtained three rewards since the last switch.

Figures 10.5 to 10.8 present the frequency distribution of bout lengths for each bird in both patches. Firstly, this figure shows an effect that has already been mentioned (Table 10.1, legend): two of the birds (Hamilton and McNewan) used the two patches differently. For these two birds, only the preferred patch was used in the analysis, while for the other two the responses in both patches were combined. The birds did not, as predicted, concentrate on a single bout length. Instead their bout lengths were scattered over a certain range. However, the most common bout length was the most profitable one for three of the four animals (Pontryagin, Hamilton and McNewan). In the case of Hamilton the mode appears after the sixth hop, as predicted taking into account the special technique of this animal of handling the prey sitting on the perch (see Table 10.1, legend). McNewan has a bimodal distribution, with one of the peaks in the most profitable bout length and the other two hops later. The fourth animal, Nahuel, had considerably longer bouts than predicted.

The coincidence between predicted and modal value of bout length may conceivably be due to a preferred switching rate independent of the schedule and travel time, instead of reflecting

an optimal foraging strategy. This hypothesis is unlikely because great tits have been shown to respond to travel time between patches by adjusting the time in each patch (Cowie, 1979). In order to test this in my set up, I modified the partition so as to change the travel distance between the two patches, and tested two of the birds (Hamilton and McNewan) under the new conditions. The travel distance was made first shorter and then longer than the experimental one, and the optimal bout length was recalculated for each bird under the new conditions. The variations in travel time which could be achieved within the limits of the aviary were not enough to modify the optimum bout length but the shape of the curves of profitability as a function of bout length changed. With short travel distance, after the third hop the curve reached almost the same profitability as after the fifth, and profitability dropped abruptly for bouts longer than five hops. With long travel distances the opposite was true, with short bout lengths having a very low profitability, and longer bouts being almost as high as five hop bouts. Both birds responded by keeping the same modal bout length as before, but the frequency distributions of bout lengths changed in the direction predicted, with more short tests in the short travel times and more long bouts in the longer travel times. All these changes were significant for the two birds. Table 10.2 summarises these results.

In summary, the birds do not show exclusive concentration in the optimal bout length, but they switch most often after the optimal number of hops, and the frequency distribution of bout lengths is sensitive to travel distance between patches in the direction of the optimal foraging prediction.

10.4.2. The effect of intrusions

(a) Latency to feed

The presence of the intruder usually provoked the type of response already seen in Chapter 5. In many cases the birds did not feed until the intruder disappeared, spending all their time displaying towards the intruder, singing from the top of the partition or even attacking the intruder. On other occasions they paid some attention to the intruder and then started to feed. Table 10.3 presents the average latency to start feeding in control and intrusion tests.

(b) Rate of hopping

The rate of hopping after the onset of feeding was not affected by the presence of the intruder: None of the birds showed any significant difference (Mann-Whitney test) in the number of hops during the first minute of feeding between control and intrusion sessions (Table 10.3).

(c) Bout length

Figures 10.9 to 10.12 show the frequency distribution of bout length in control and intrusion tests. The three birds that had shown a modal bout length coinciding with the predicted value when undisturbed showed some shift towards shorter bouts. Two of them (Hamilton and McNewan) had modes at the third hop instead of the fifth one, and the third one (Pontryagin) showed no change in mode but increased the proportion of short bouts (1-2 hops). Nahuel did not show any apparent change. A Kolmogorov-Smirnov test between control and intrusion sessions showed a significant shift to shorter bouts in McNewam and Pontryagin and no significant change

for the other two birds. The results of this test are shown in Table 10.4.

10.4.3. Inverse optimality

While the most profitable bout length was predicted directly using some elements of each individual's behaviour as constraints, the same procedure cannot be applied to predict the optimal shift to shorter bouts during intrusion tests. The reason is that it is not possible to estimate the benefit obtained from monitoring the territory in a quantitative manner. It is possible though to use the 'inverse optimality' approach (Houston, 1977; McFarland, 1977) to calculate how much that benefit would have to be to produce a shift of the optimal bout length from five hops to three hops, as observed for three of the birds. These values can be obtained graphically, generalising the method used for calculating optimal time in a patch (Figure 10.3). The method can be extended by adding the benefit B obtained by the bird while travelling between patches, when this travel time is used to check for the presence of potential invaders in the territory. In order to keep the figure format, B is plotted in equivalents of prey units below the level of the abscissa. The optimum bout length for any known B can be obtained from the tangent to the resource depression curve drawn from the new origin P', situated B units below the old origin of the tangent. Alternatively, if the value of B is not known, as in this case, but the shift in the value of the time in patch is obtained from the data and there are convincing reasons to assume that the shift in bout structure obeys to a 'rational' policy, the point P can be found by tracing the tangent to the resource depression curve at the point of the new

bout length. The same results can be obtained if Equation 1 is modified by including the benefit from territorial monitoring together with the benefit in terms of energy intake.

Thus,

$$\text{Benefit/time} = \frac{Np + B}{TT + Ht \,.\, Np + Nh \,.\, IHI} \qquad (2)$$

where B is the Benefit (in prey equivalents) of checking for the presence of the intruder.

For a bout of three hops to score higher than a bout of five, the following inequality must hold:

$$\frac{2 + B}{TT + 2Ht + IHI} > \frac{3 + B}{TT + 3Ht + 2IHI} \qquad (3)$$

Solving for B

$$B > \frac{TT - IHI}{Ht + IHI} \qquad (4)$$

As pointed out above, the term on the left of the inequality is unknown, but the term on the right can be calculated for each individual. This makes it possible to rank the birds according to the likelihood that the conditions for a change to bouts of three hops be met. In other words, it is possible to state how big the benefit of checking the territory ought to be to shift the optimal bout length from five to three hops.

The resulting critical values of B are: Pontryagin: 0.61; McNewan: 0.37; Hamilton: 0.28 and Nahuel: 1.00.

It is interesting to note that the only bird which did not show any shift towards shorter bouts (Nahuel) was also the animal who, given its constraints, would have required the greatest benefit to make the shift appropriate.

10.5. Discussion

The results of this experiment support the hypothesis that territorial great tits, while able to forage approximately in the most profitable way when there is no danger of invasion will feed less efficiently in order to check for the presence of rival males when this danger increases. The ability to forage in the most profitable manner is indicated by the fact that three out of the four subjects when undisturbed organised their feeding in bouts with a modal length equal to the number of hops which maximisesthe rate of energy intake per time unit.

When the cost of changing patch was modified for two of the birds, both changed the frequency distribution of bout lengths as predicted by an optimal foraging model. this result confirms those obtained by Cowie (1979). It is worth pointing out an interpretational problem that this experiment shares with a great number of studies of optimality in animal behaviour. The model used to predict the most profitable bout length is a deterministic one, and predicts for each individual the exact length of bout that the bird ought to make if it is to maximise the rate of energy intake. The data, instead, look very much like most biological results, having a certain degree of scatter around central values. There is nothing in the theory to state which is the critical degree of scattering of the data which can be accepted as sufficient to either reject or accept the agreement of predictions and results, so that the decision is left to the biological common sense of the researcher(and his readers). In the results obtained

in this experiment the degree of scattering around the modal value is considerable, but then, the difference in payoff resulting from bout lengths in the range of 3 to 12 hops is very small (Figures 10.3 and 10.4) and if anything it is quite striking that the birds (or at least three of them) were able to detect this difference and switch most frequently at the appropriate bout lengths.

The spread may be the result of the decision rules used by the birds to implement the optimal behaviour. I discussed in Chapter 8 the difference between the way in which the experimenter and the foraging predator may reach the solution to a particular problem. In this particular case, I calculated the most profitable behaviour by means of Equation 1. but since the birds have neither stop watches nor counters it would not be reasonable to expect that their way of reaching decisions would be based on the same procedure. Simple rules-of-thumb have been shown to achieve acceptable convergence with the theoretical optimum in the non-depleting schedules used in Part 2 of this thesis. For concurrent depleting schedules like the ones used in this chapter a simple rule-of-thumb could be: leave a patch after x seconds since arrival. Cowie (1979) suggests that great tits used a time expectation rule which he argues is less affected by stochastic variations in intercatch interval than a giving-up-time as proposed by Krebs et al. (1974). As opposed to the prediction calculated with an equation, behaviour adjusted using a rule-of-thumb may have some stochastic components in assessment of time and this may explain the scattering. In addition to this explanation based

on the birds' constraints, a functional (and difficult to test) justification for this stochasticity might be that a bird having some variation in bout lengths may be quicker to adjust to changes in the environment than a bird which always leaves a patch after a fixed number of seconds.

Table 10.1

Average travel time, handling time and inter hop interval for each individual

	Travel Time	Handling Time	Int-hop-Int
Pontryagin	6.0	7.2	1.6
Nahuel	5.3	3.1	1.1
McNewan	4.6	6.9	1.5
Hamilton	2.7	5.0	1.0

One of the subjects, Hamilton, has a specially low travel time. This bird achieved this short travel time by making a parabolic flight over the partition, between the two perches, always landing directly on the perch of the other patch. The other birds instead, normally made straight flights from the perch to the partition, from the partition to the platform where the other perch was placed and from there they hopped to the perch. Hamilton seemed to be specially able at manouvering. Another remarkable trick that this bird learnt was to handle the prey sitting on the perch. After obtaining a reward in either of the two patches he would fly without stopping to the perch of patch number 1 where he would sit while handling the prey. By doing this, this subject in effect achieved a 'free' hop in patch 1 after every reward. On the other hand this meant that the animal was effectively hunting only in patch 1 staying in patch 2 only until getting the first reward, at the second hop. Visits to patch 2 were in fact just resettings of its preferred patch. Another bird, McNewan, also developed the stereotyped behaviour of using one of the patches to reset the preferred one, but in his case the preference was reversed with respect to Hamilton. These stereotypes were taken into account when calculating true optimal strategy (see text).

Table 10.2

Effect of travel distance on bout length

a) Modal bout length was not changed by inter patch distance

	Bout length (mode)		
Distance	Short	Medium	Long
Hamilton	7	6	7
McNewan	5	5,7	7

b) Short bouts were more common in shorter travel distances
Summary of the statistical comparison of the frequency distributions of bout length.

	Kolmogorov-Smirnov statistic		
Distance	Short vs Medium	Medium vs Long	Short vs Long
Hamilton	1.49 $p<0.05$	3.18 $p<0.001$	3.02 $p<0.001$
McNewan	2.57 $p<0.001$	1.39 $p<0.025$	3.19 $p<0.001$

p is the significance level of a one-tailed test.

Table 10.3 (a)

Latency to first hop

	Control	Intrusion	p
Pontryagin	< 10	31.25 ±18.63	NT
Nahuel	< 10	37.85 ±22.46	NT
McNewan	8.79 ±5.15	18.02 ±9.95	0.05
Hamilton	3.04 ±1.64	4.57 ±1.23	0.05

Table 10.3 (b)

Hops in first minute of feeding

	Control	Intrusion	p
Pontryagin	13.17 ±0.08	12.43 ±0.90	Ns
Nahuel	26.00 ±3.41	19.00 ±1.57	Ns
McNewan	19.11 ±2.62	17.00 ±2.07	Ns
Hamilton	17.6 ±3.85	15.5 ±2.9	Ns

The onset of feeding was delayed by the intruder (a), but rate of hopping after the first hop was constant for each subject (b). Latency could not be measured accurately for Pontryagin and Nahuel in control sessions, but always it was shorter than 10 seconds.

NT: not tested; Ns: non-significant (Mann-Whitney)

Table 10.4 (a)

Effect of intruder on modal bout length

	Patch used	Bout length (mode)	
		Control	Intrusion
Pontryagin	both	5	5
Hamilton	1	6	4
Nahuel	both	9	9
McNewan	2	5 - 7	3

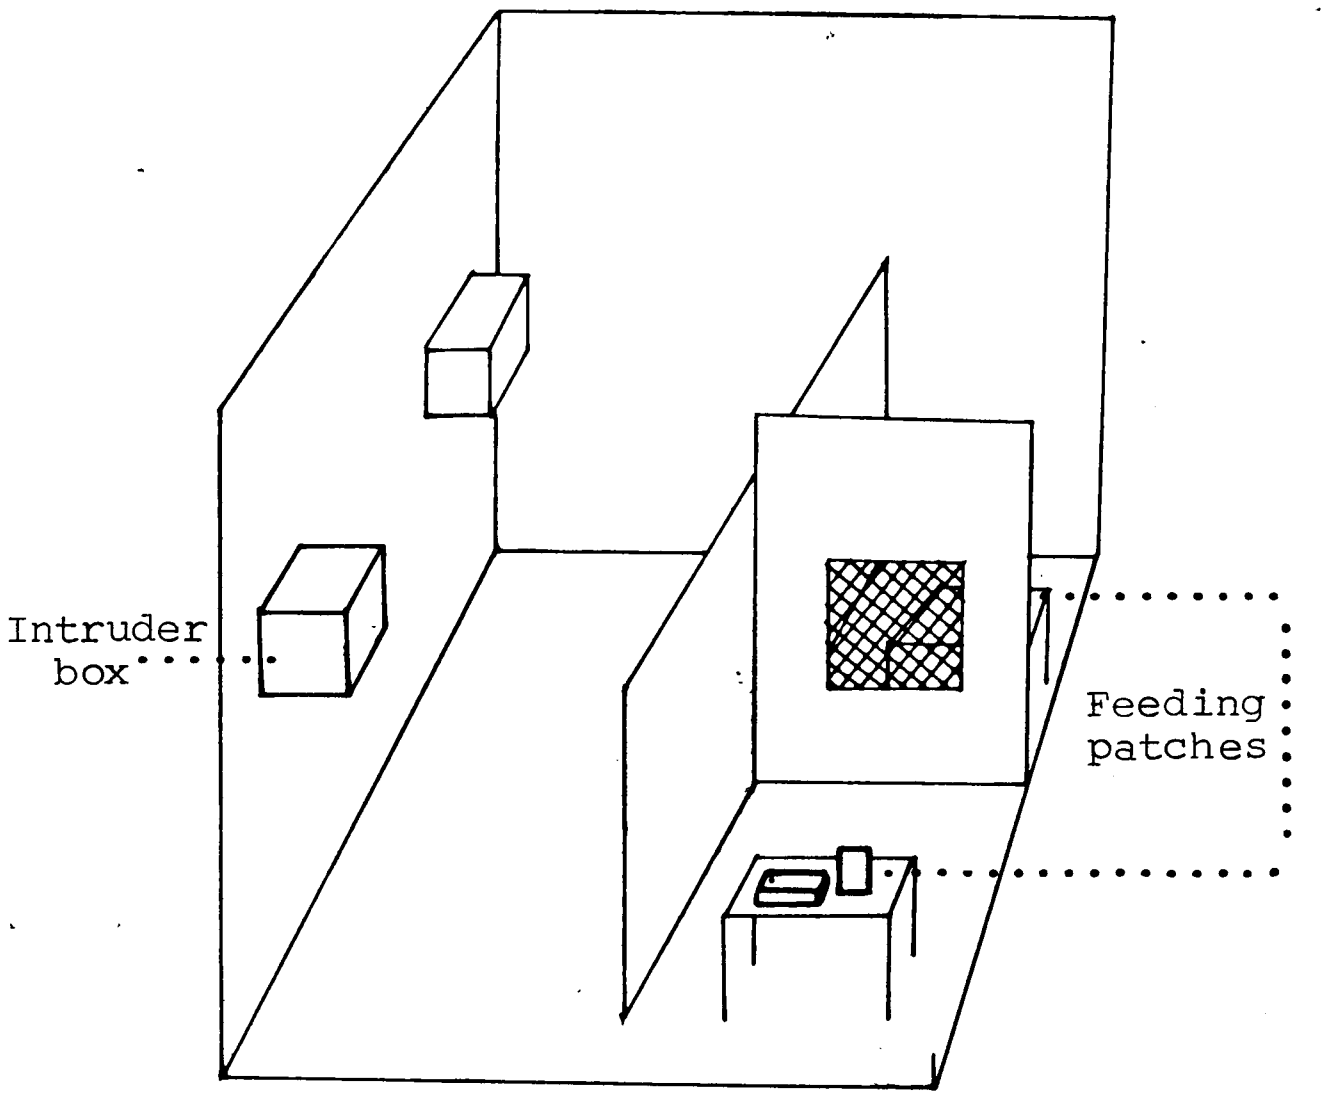
Table 10.4 (b)

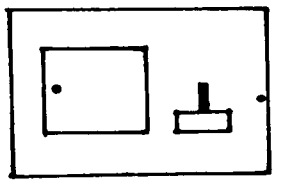
Effect of intruder on frequency distribution of bout lengths
Kolmogorov-Smirnov tests of shift to shorter bouts.

	K	p
Pontryagin	1.76	0.02
Hamilton	1.10	Ns
Nahuel	0.58	Ns
McNewan	1.50	0.05

The presence of an intruder changed the modal bout length towards shorter bouts in Hamilton and McNewan, and shifted significantly the frequency distribution of bout lengths in Pontryagin and McNewan. Nahuel changed neither the mode nor the frequency distribution.

Figure 10.1

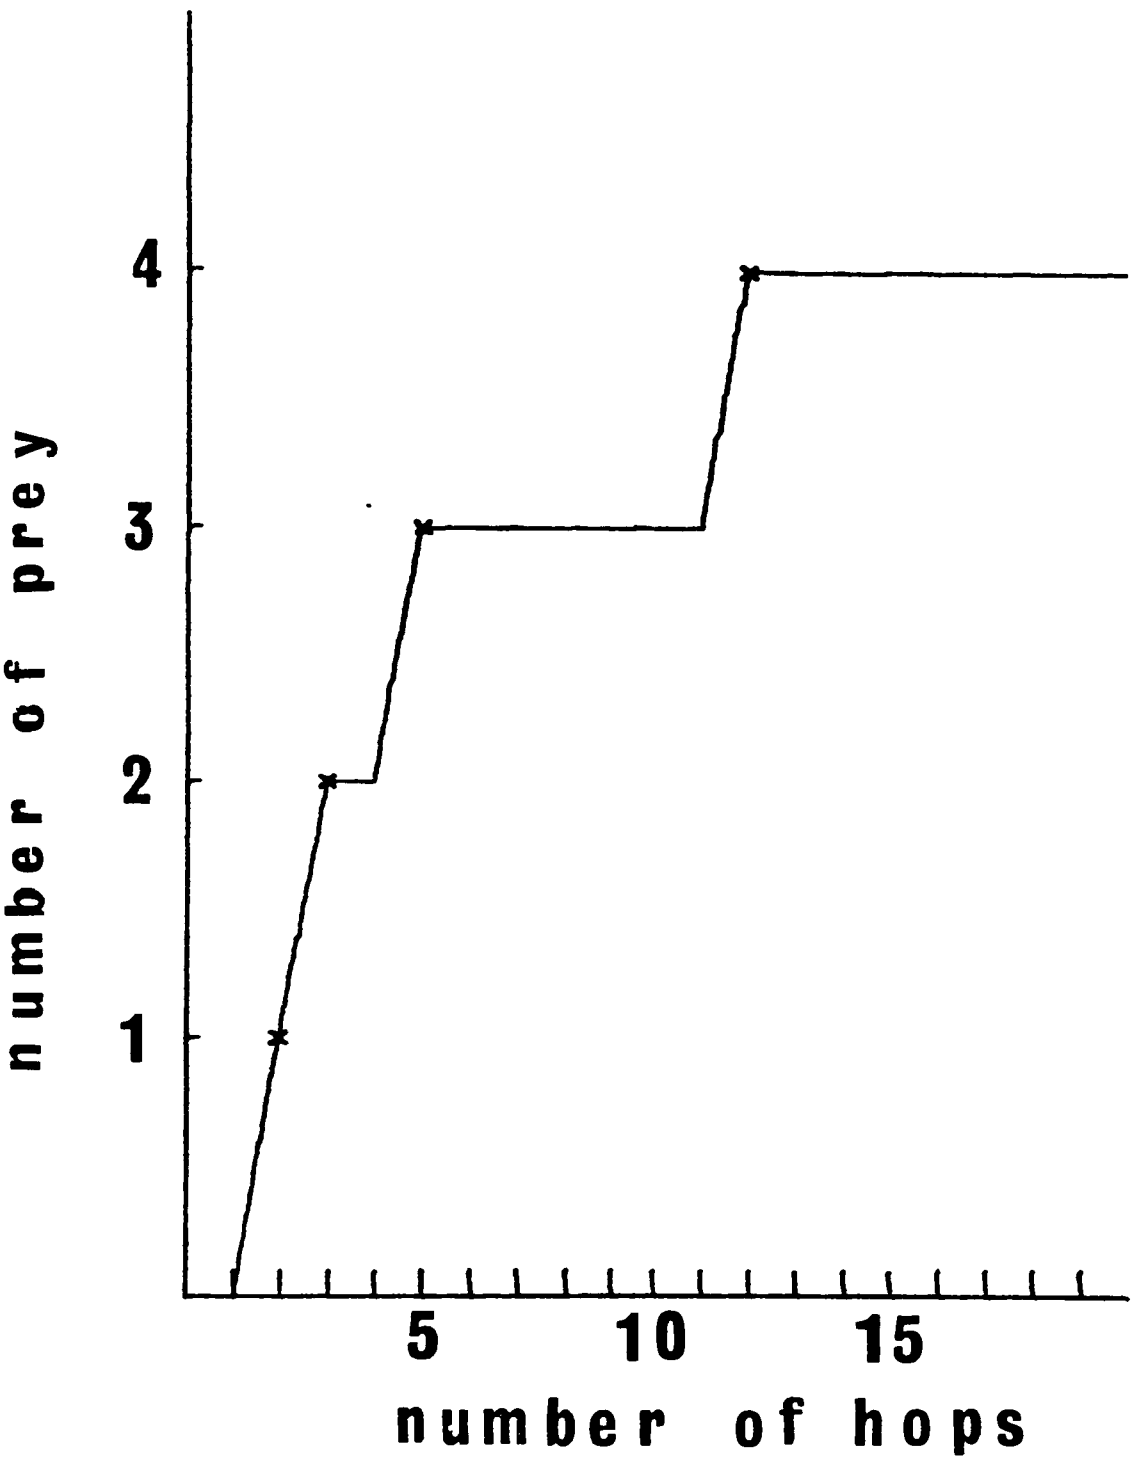




Experimental aviary.

This diagram shows one of the indoor compartments in Figure 1.4, as arranged for this experiment. The central leg of the T-shaped wall separating the feeding patches leaves 0.20m of clearance from the ceiling.

Figure 10.2



Cumulative number of rewards as a function of number of hops. Only the first 20 hops are shown because the birds very rarely stayed longer than that in a patch. This schedule approximates an exponential depletion curve. Number of prey as a function of time depends on the rate of hopping and handling time.

Figure 10.3 (a, b, c, d) Individual rate of intake as a function of bout length. In each panel equation 10.1 (see text) is evaluated for all lengths of bout. The crosses indicate hops where a reward was obtained. The highest capture rate for all the birds is at the fifth hop, indicating that this is the optimal bout length in undisturbed sessions

TT = Travel time; HT = Handling time;

IHI = Inter hop internal

Figure 10.3 (a)

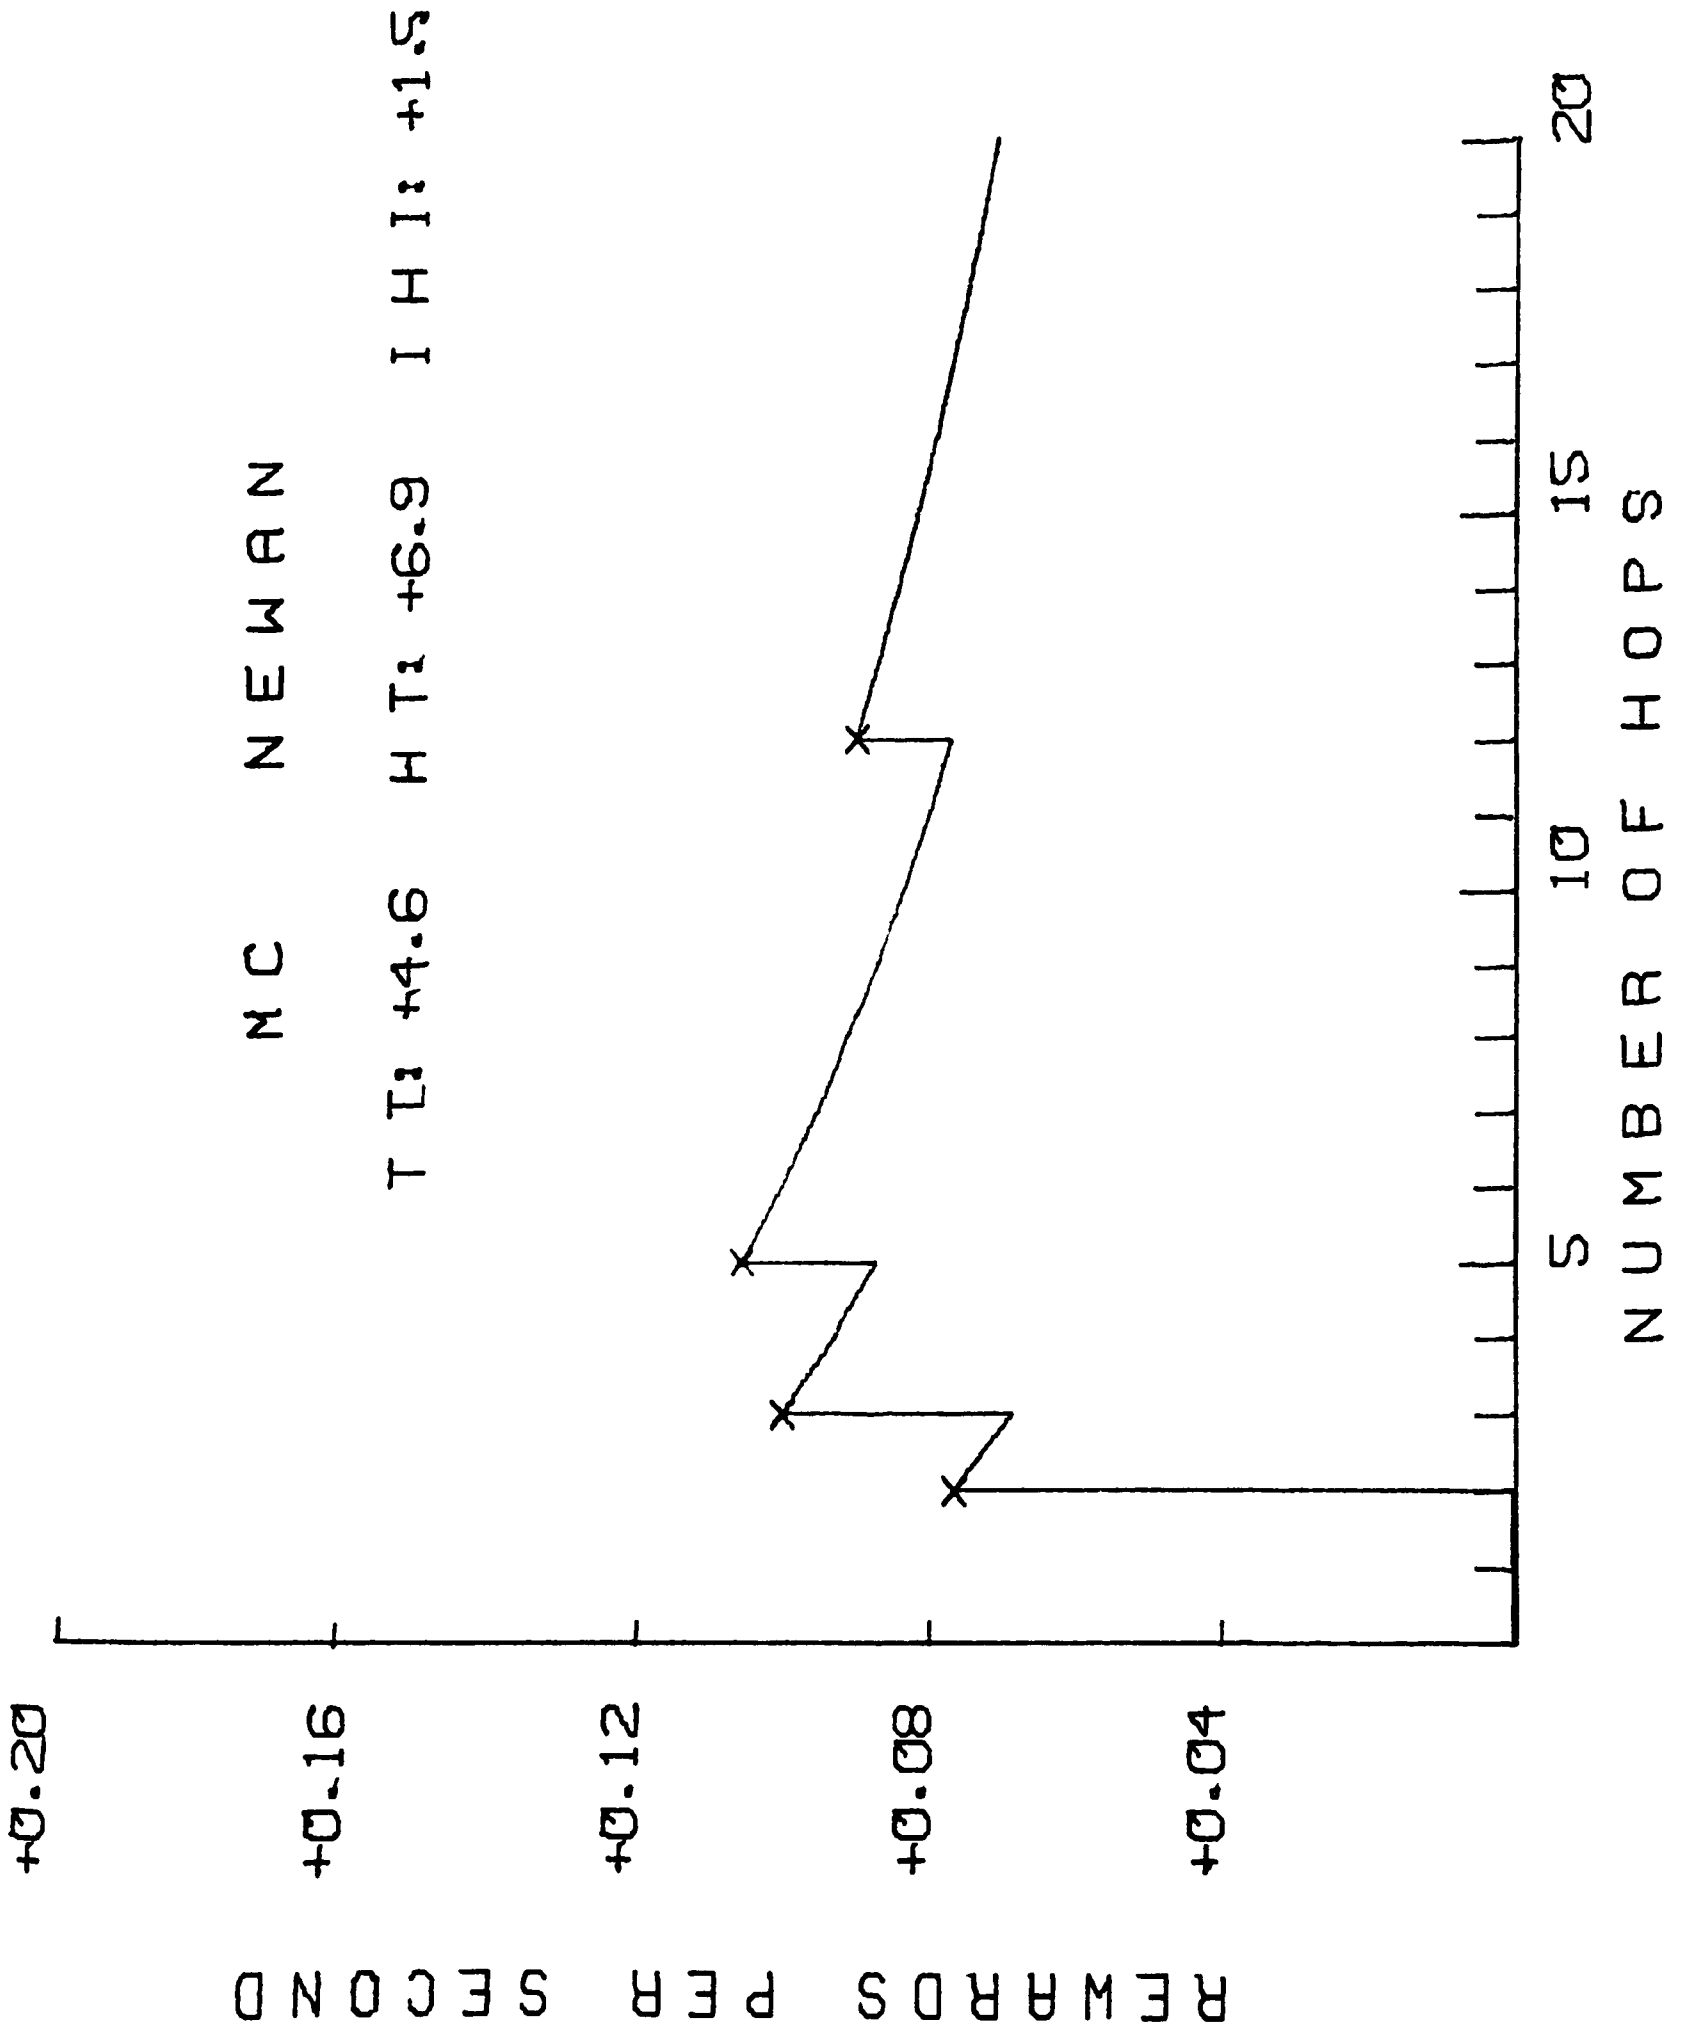

NEWAN
H T1 +6.9 I H I1 +1.5
T T1 +4.6
+0.20
+0.16
+0.12
+0.08
+0.04
REWARDS PER SECOND
5
10
15
20
NUMBER OF HOPS

Figure 10.3 (b)

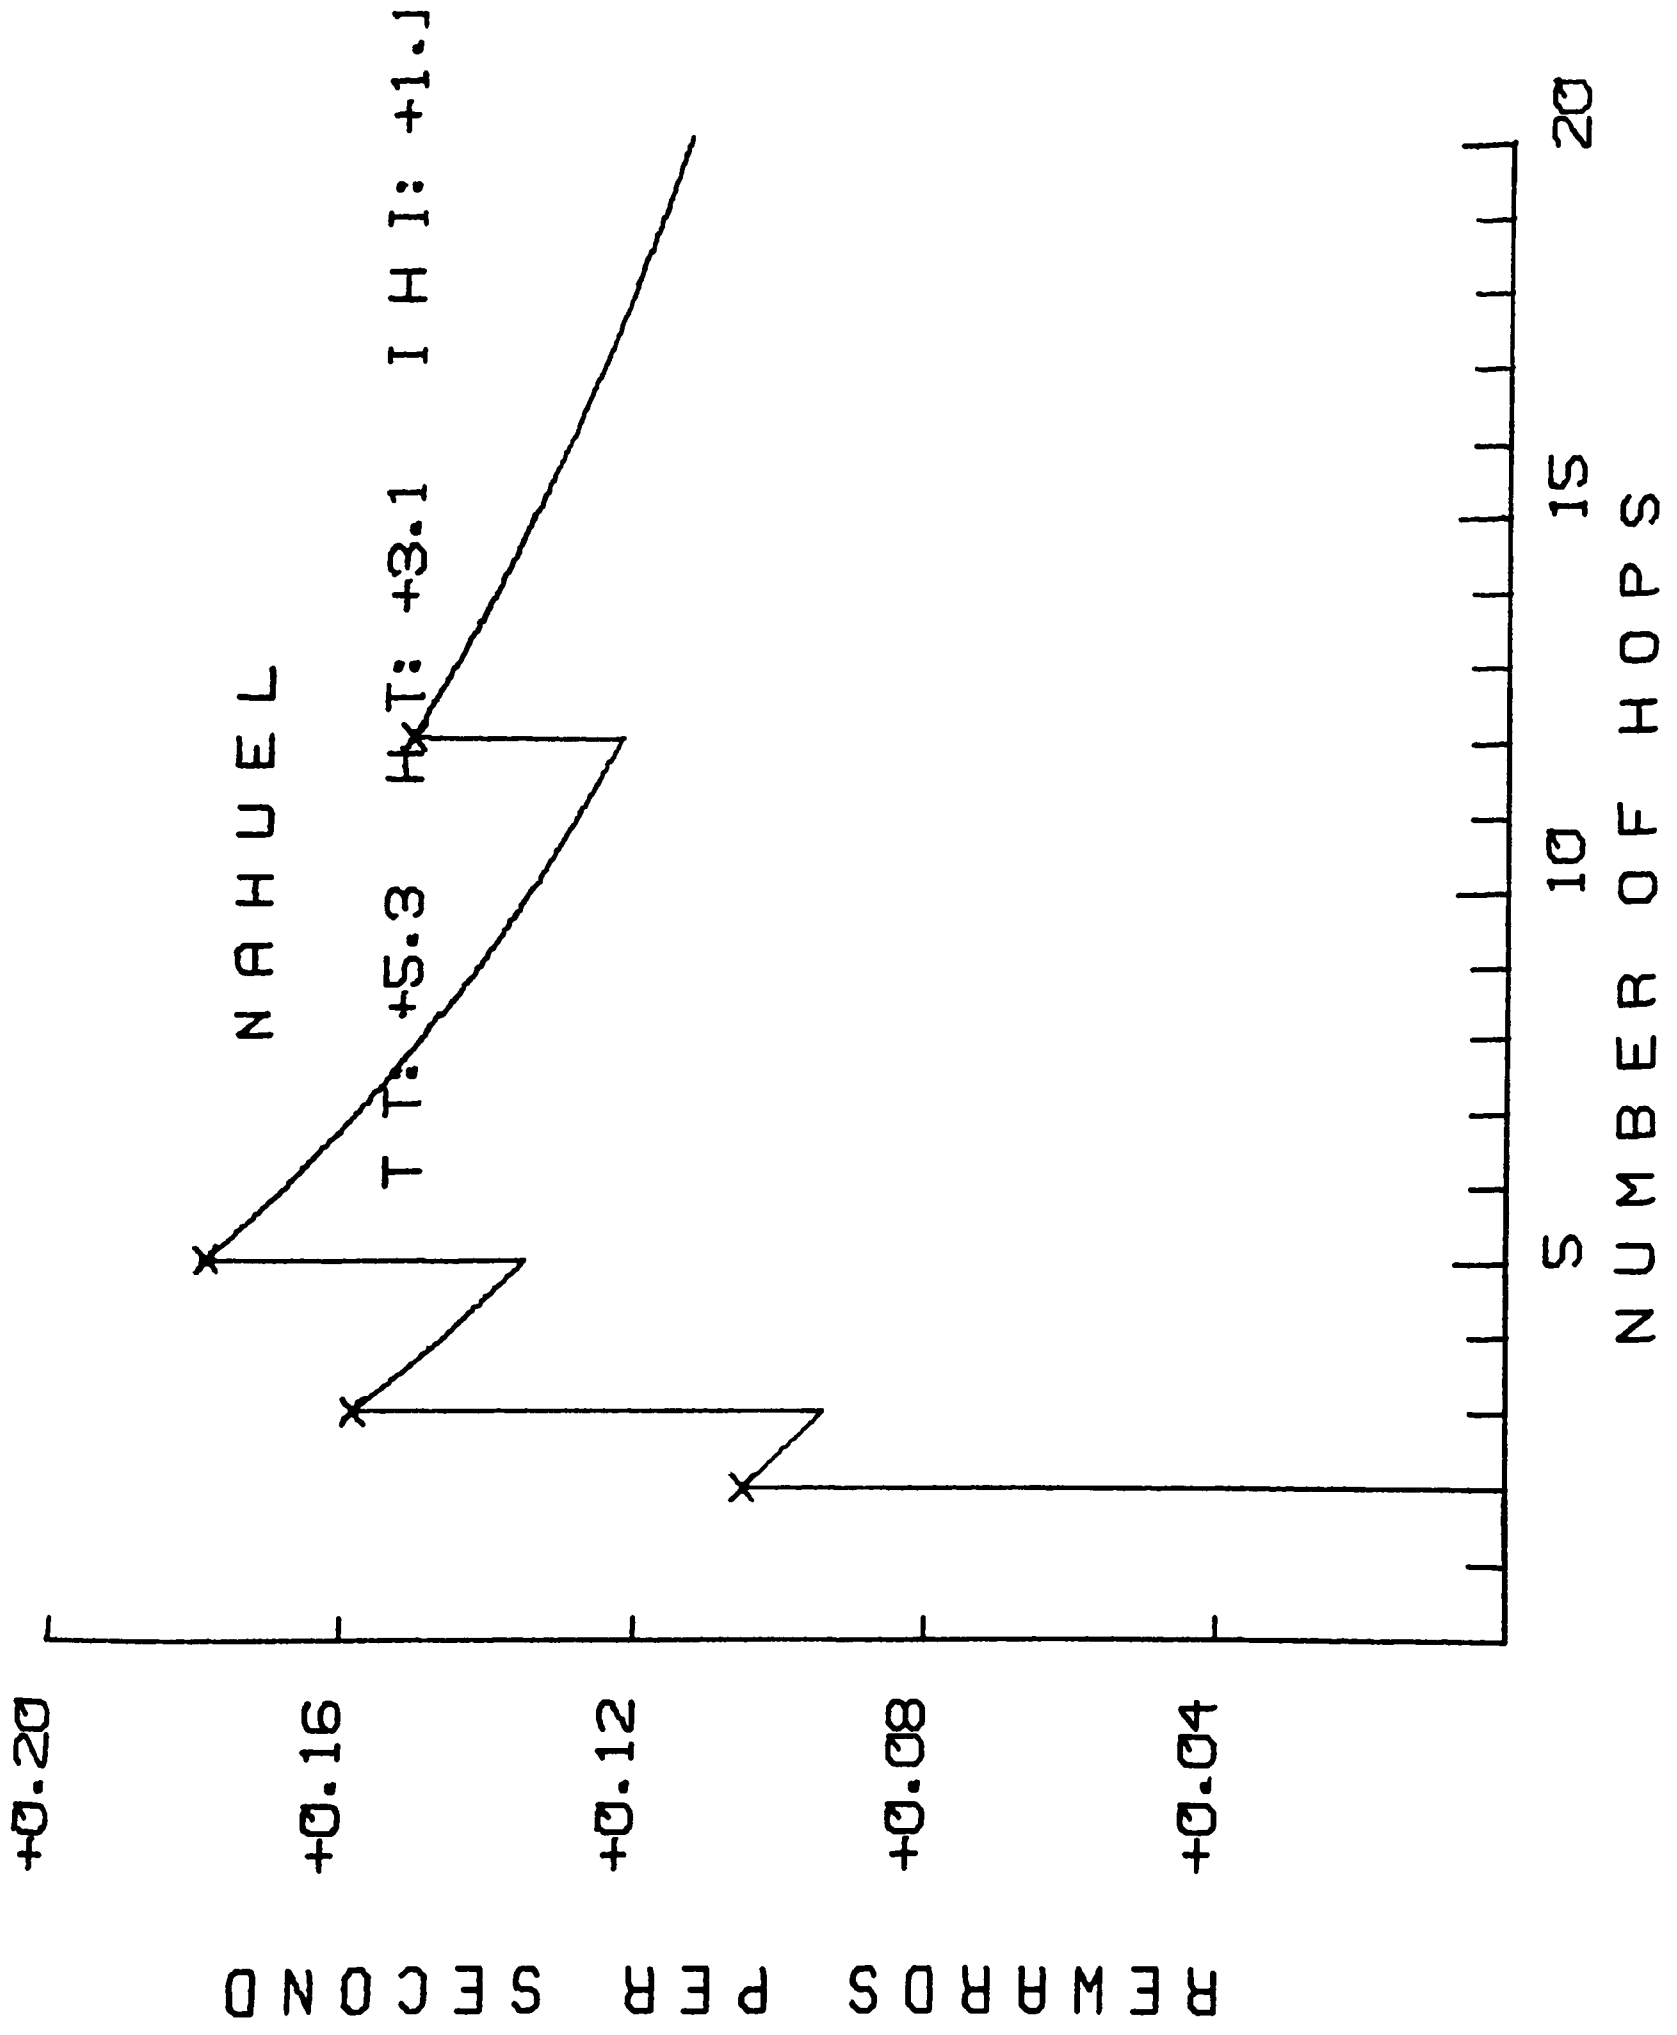

REWARDS PER SECOND
+0.20
+0.16
+0.12
+0.08
+0.04
NUMBER OF HOPS
5
10
15
20
NAHUEL
T T: +5.3
H T: +3.1
I H I: +1.1

Figure 10.3 (c)

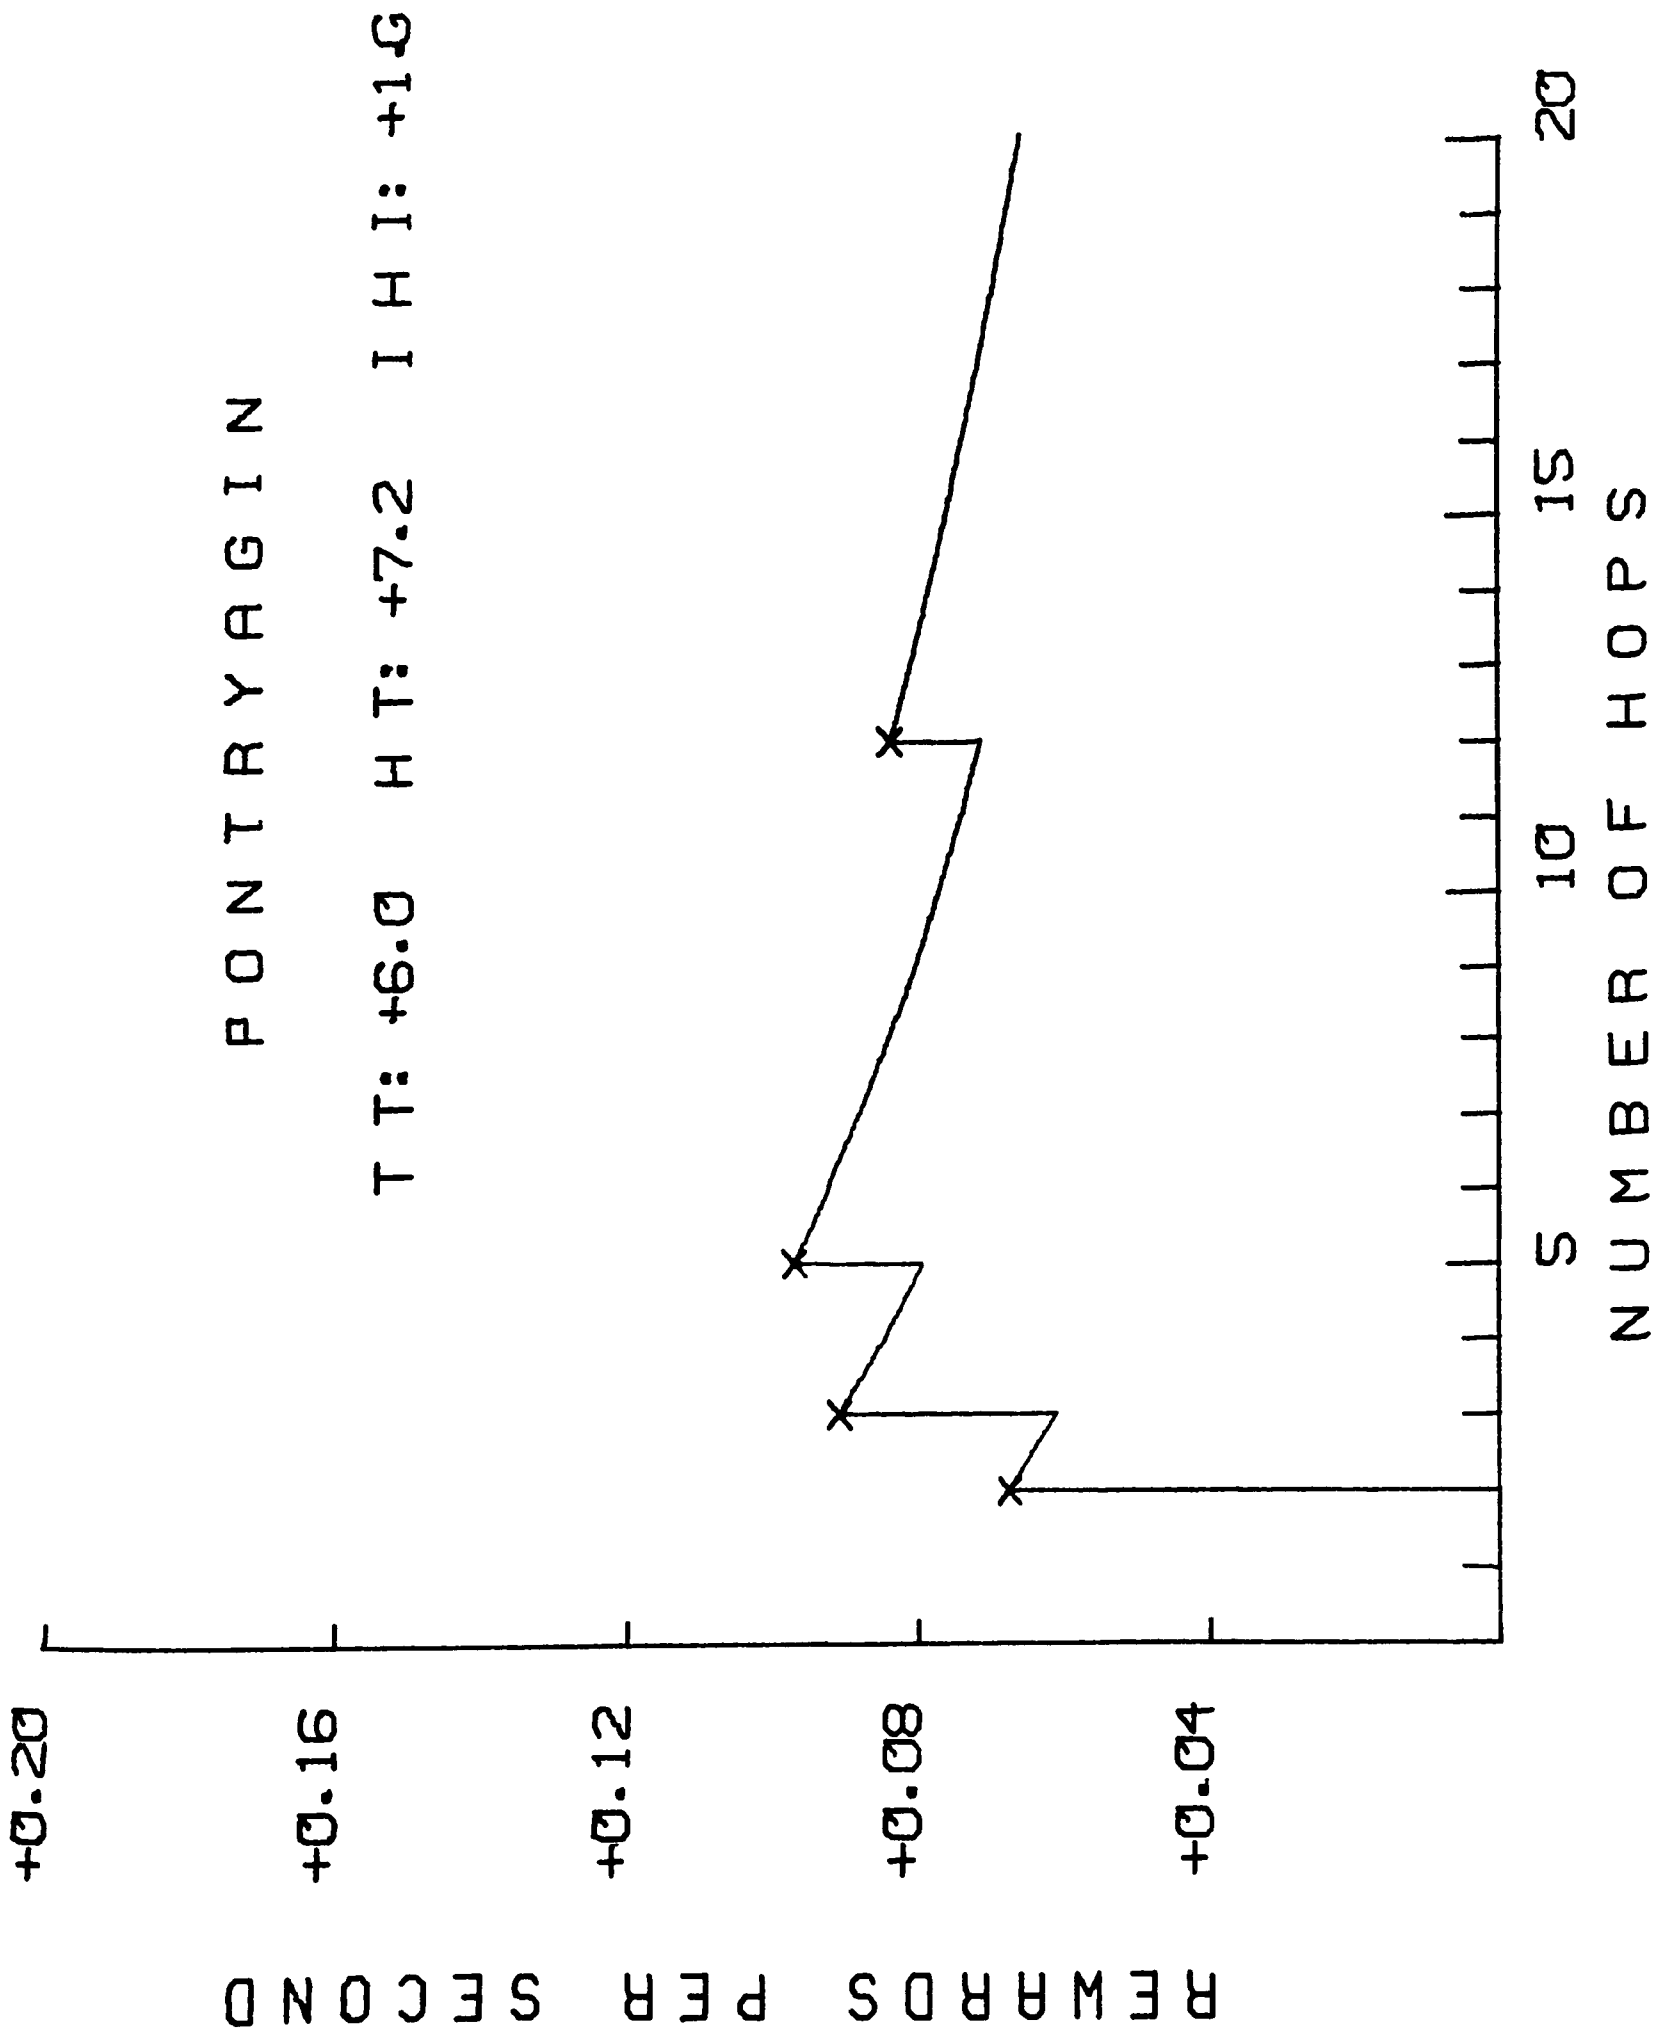

PONTRYAGIN
TT: +6.0 HT: +7.2 IHI: +1G
REWARDS PER SECOND
+0.20
+0.16
+0.12
+0.08
+0.04
5
10
15
20
NUMBER OF HOPS

Figure 10.3 (d)

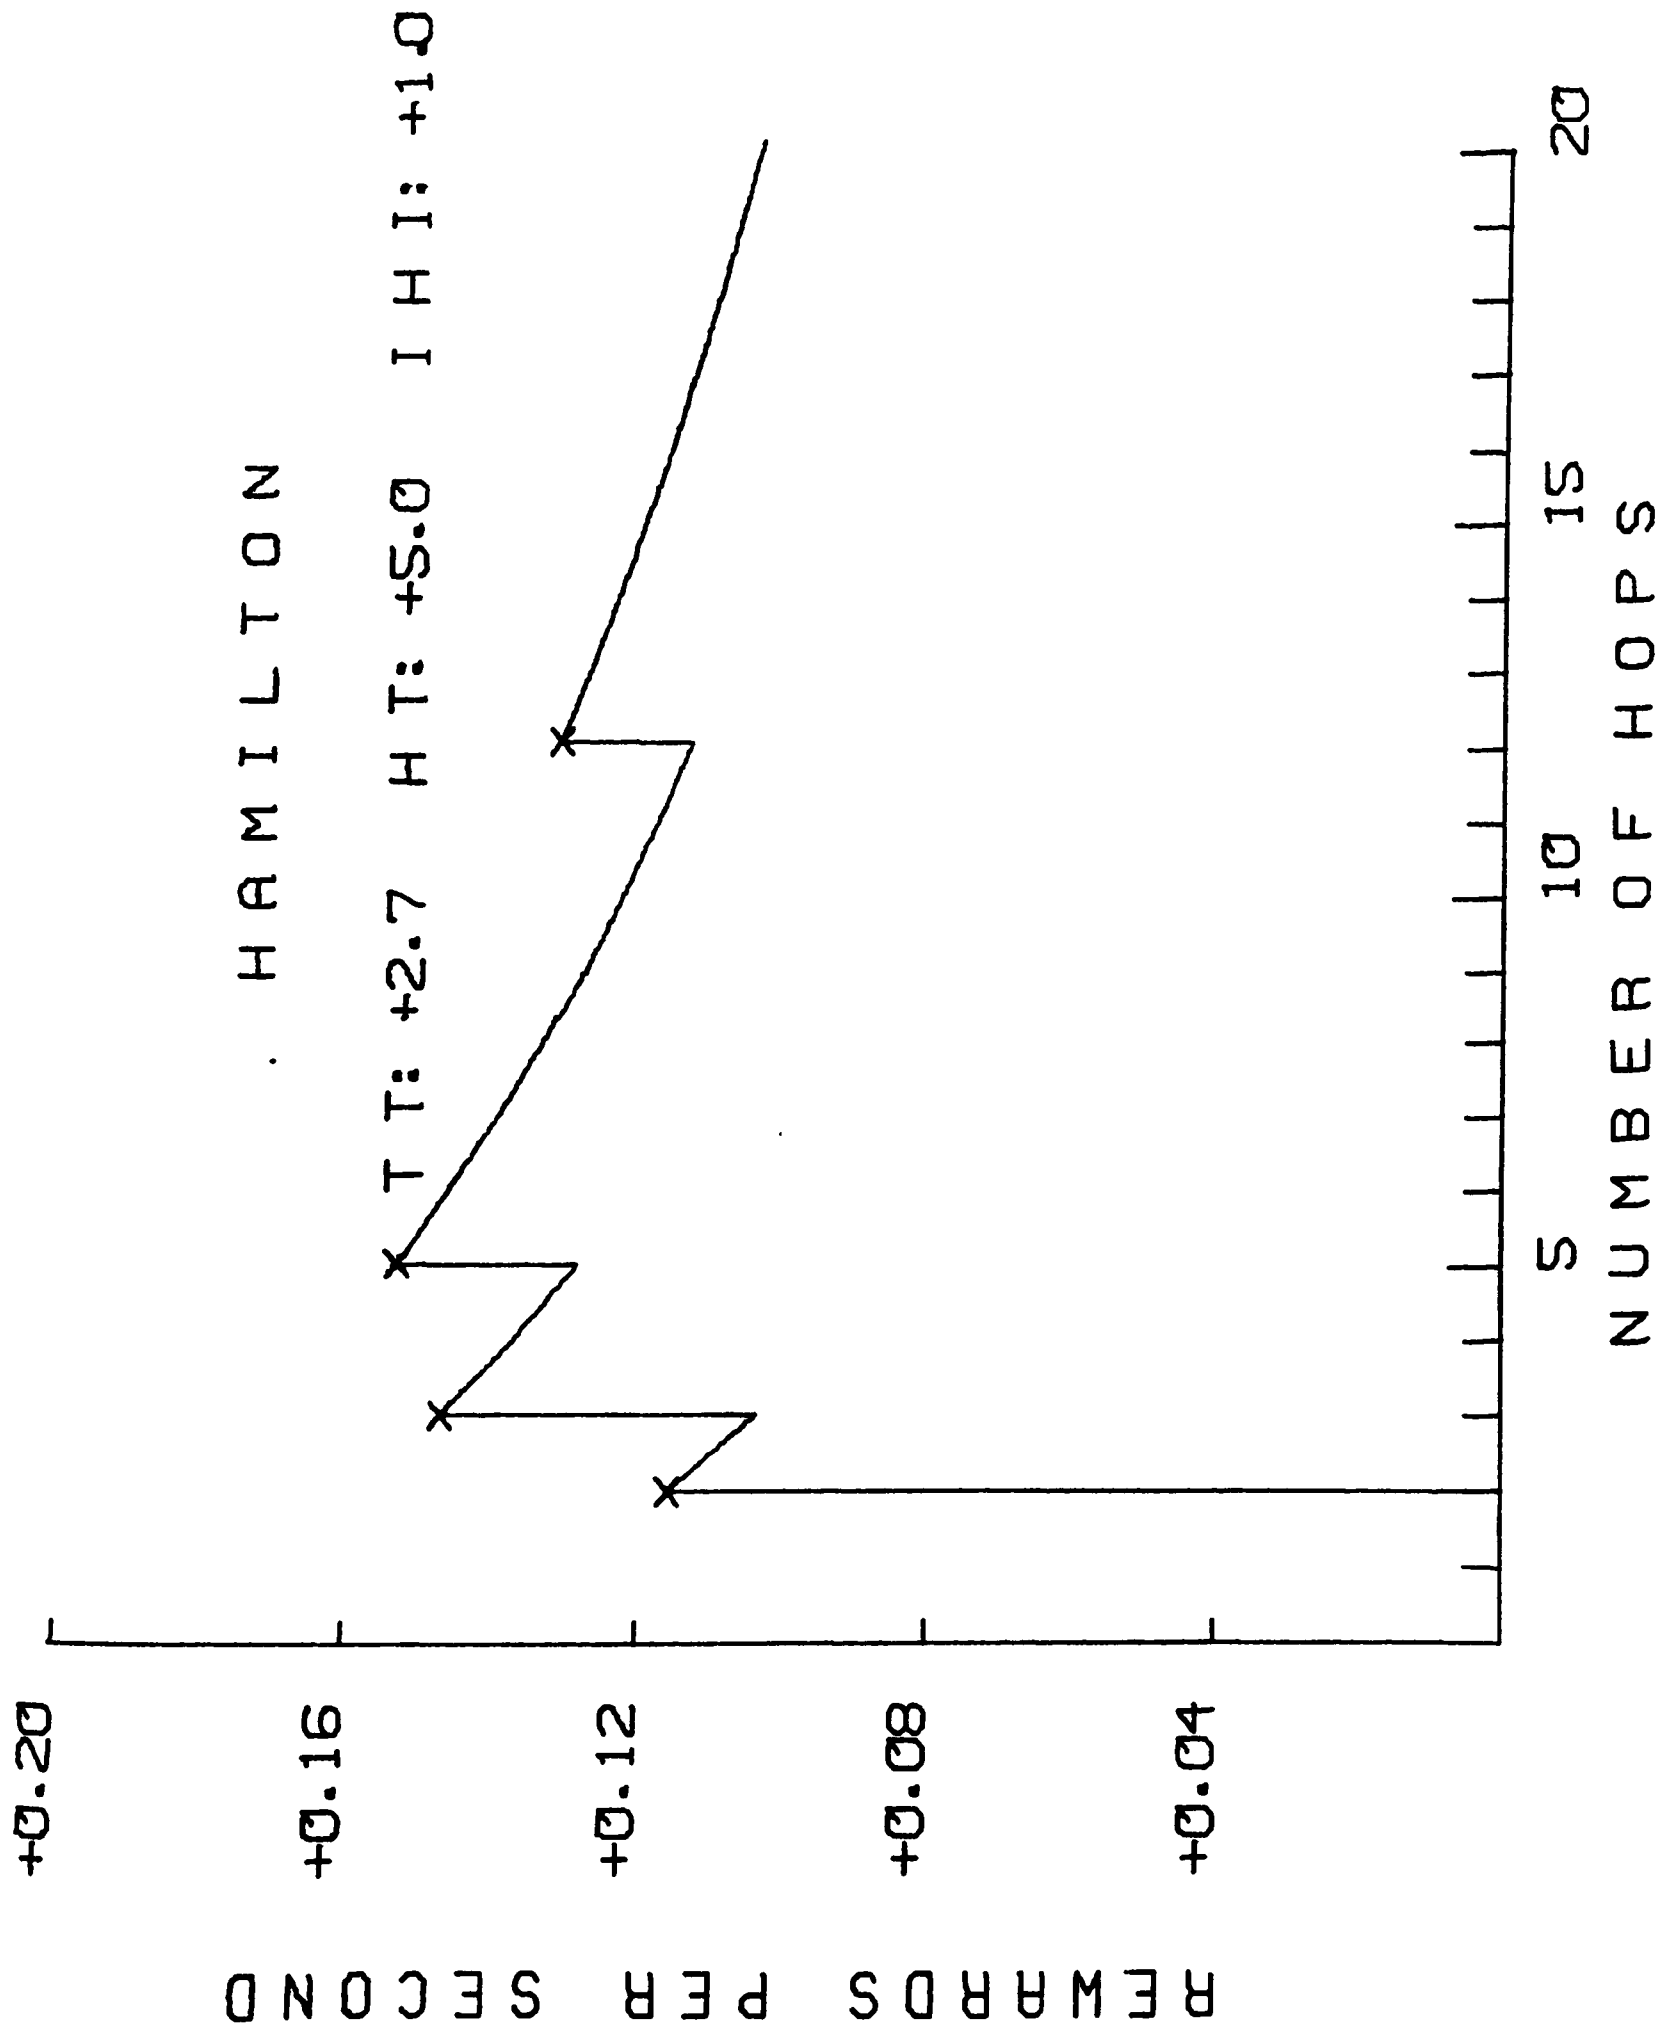
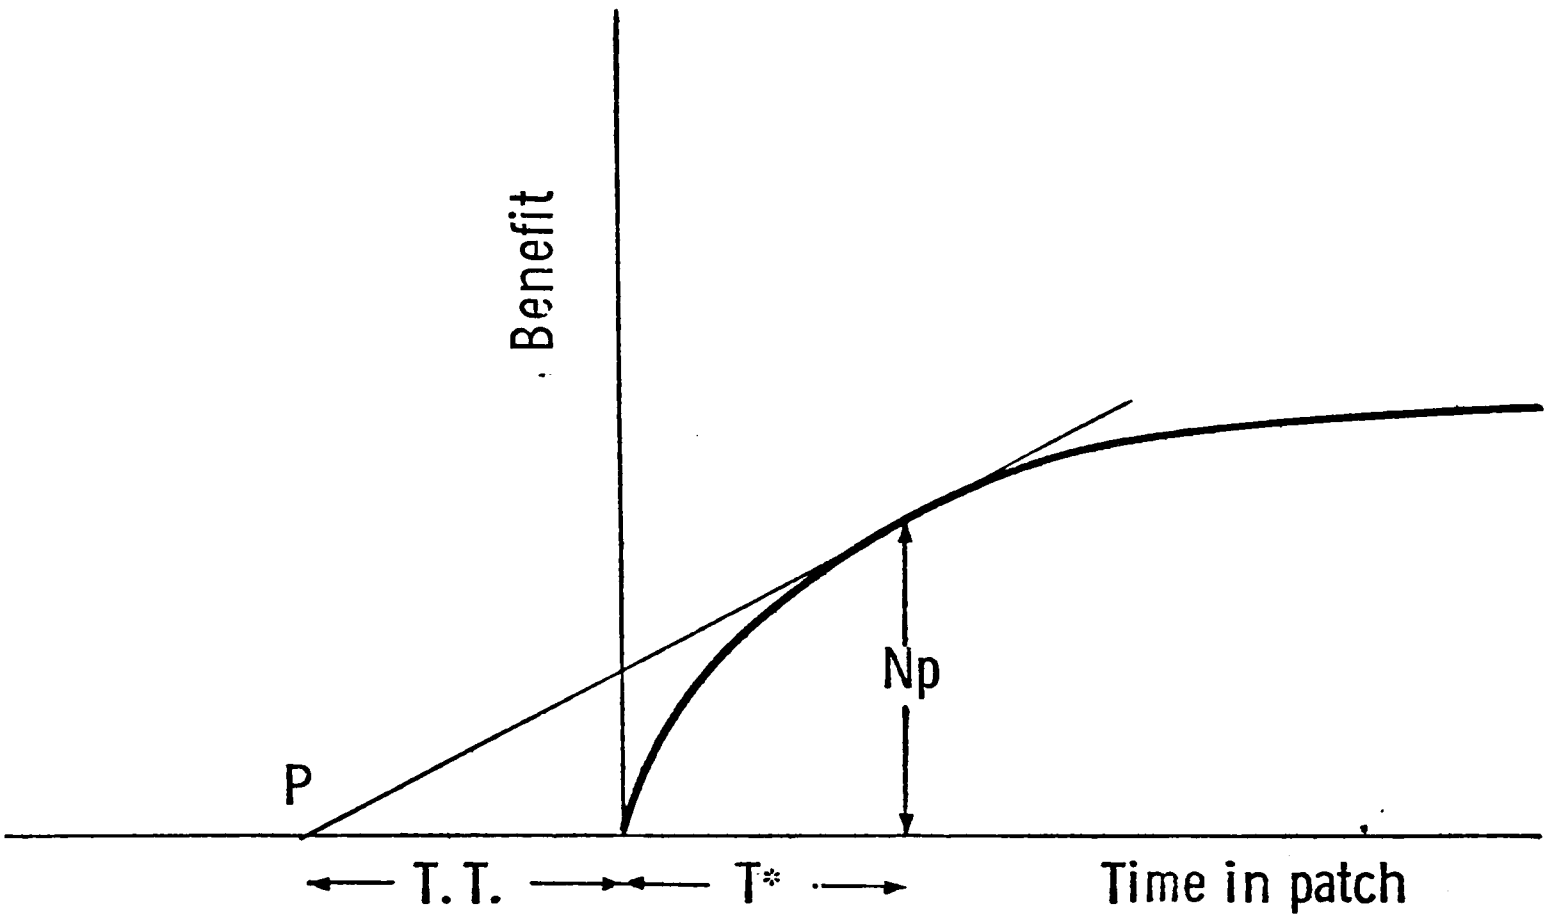
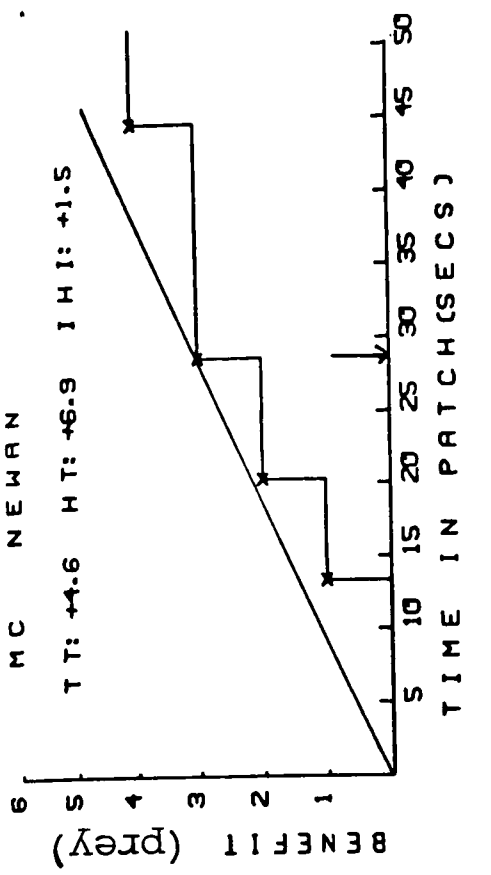

HAMILTON
TT: +2.7 HT: +5.0 IHI: +1Q
REWARDS PER SECOND
+0.20
+0.16
+0.12
+0.08
+0.04
5
10
15
20
NUMBER OF HOPS

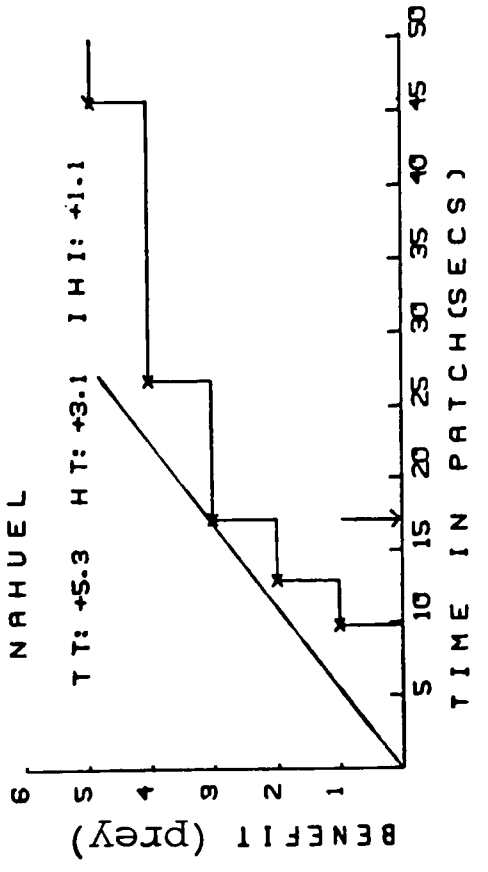
Figure 10.4 (a)

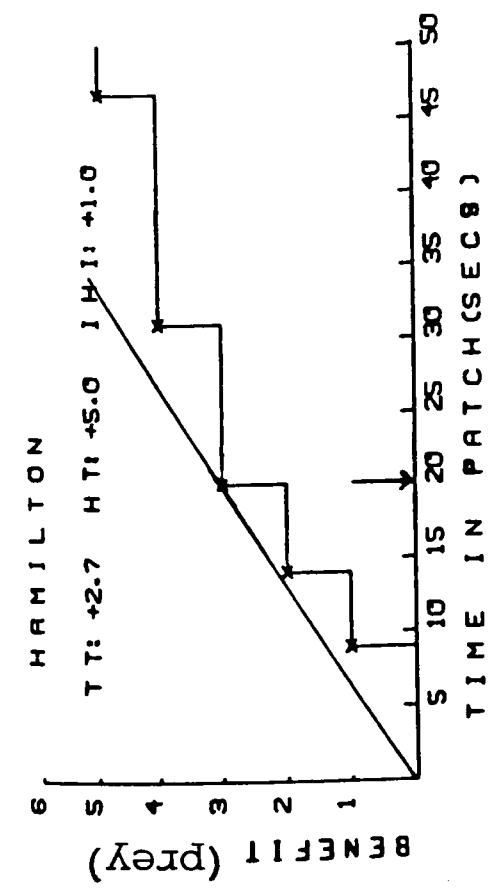


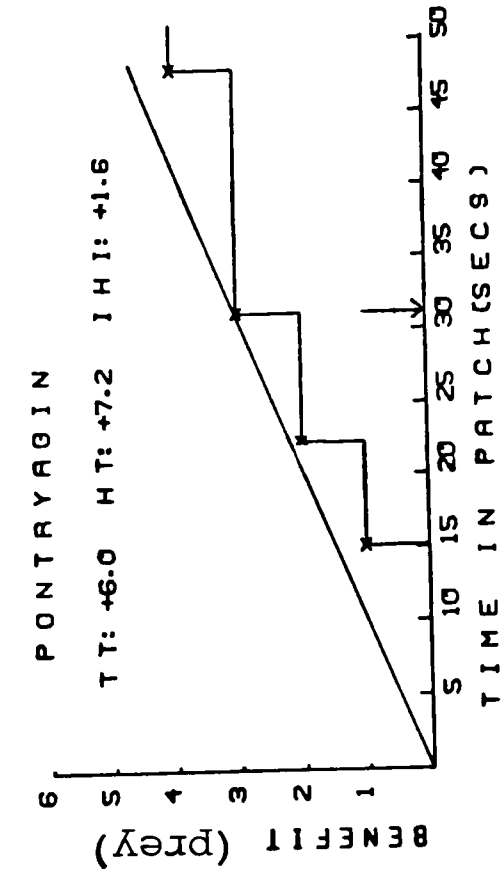
Graphic method to determine optimum time in a depleting patch. Cumulative number of prey is plotted against time since arrival. Mean travel time (TT) is shown to the left of the time origin . Optimum time in a patch (T) is given by the tangent to the cumulative benefit curve through P, which indicates travel time (from Parker, 1978).

Figure 10.4 (b)



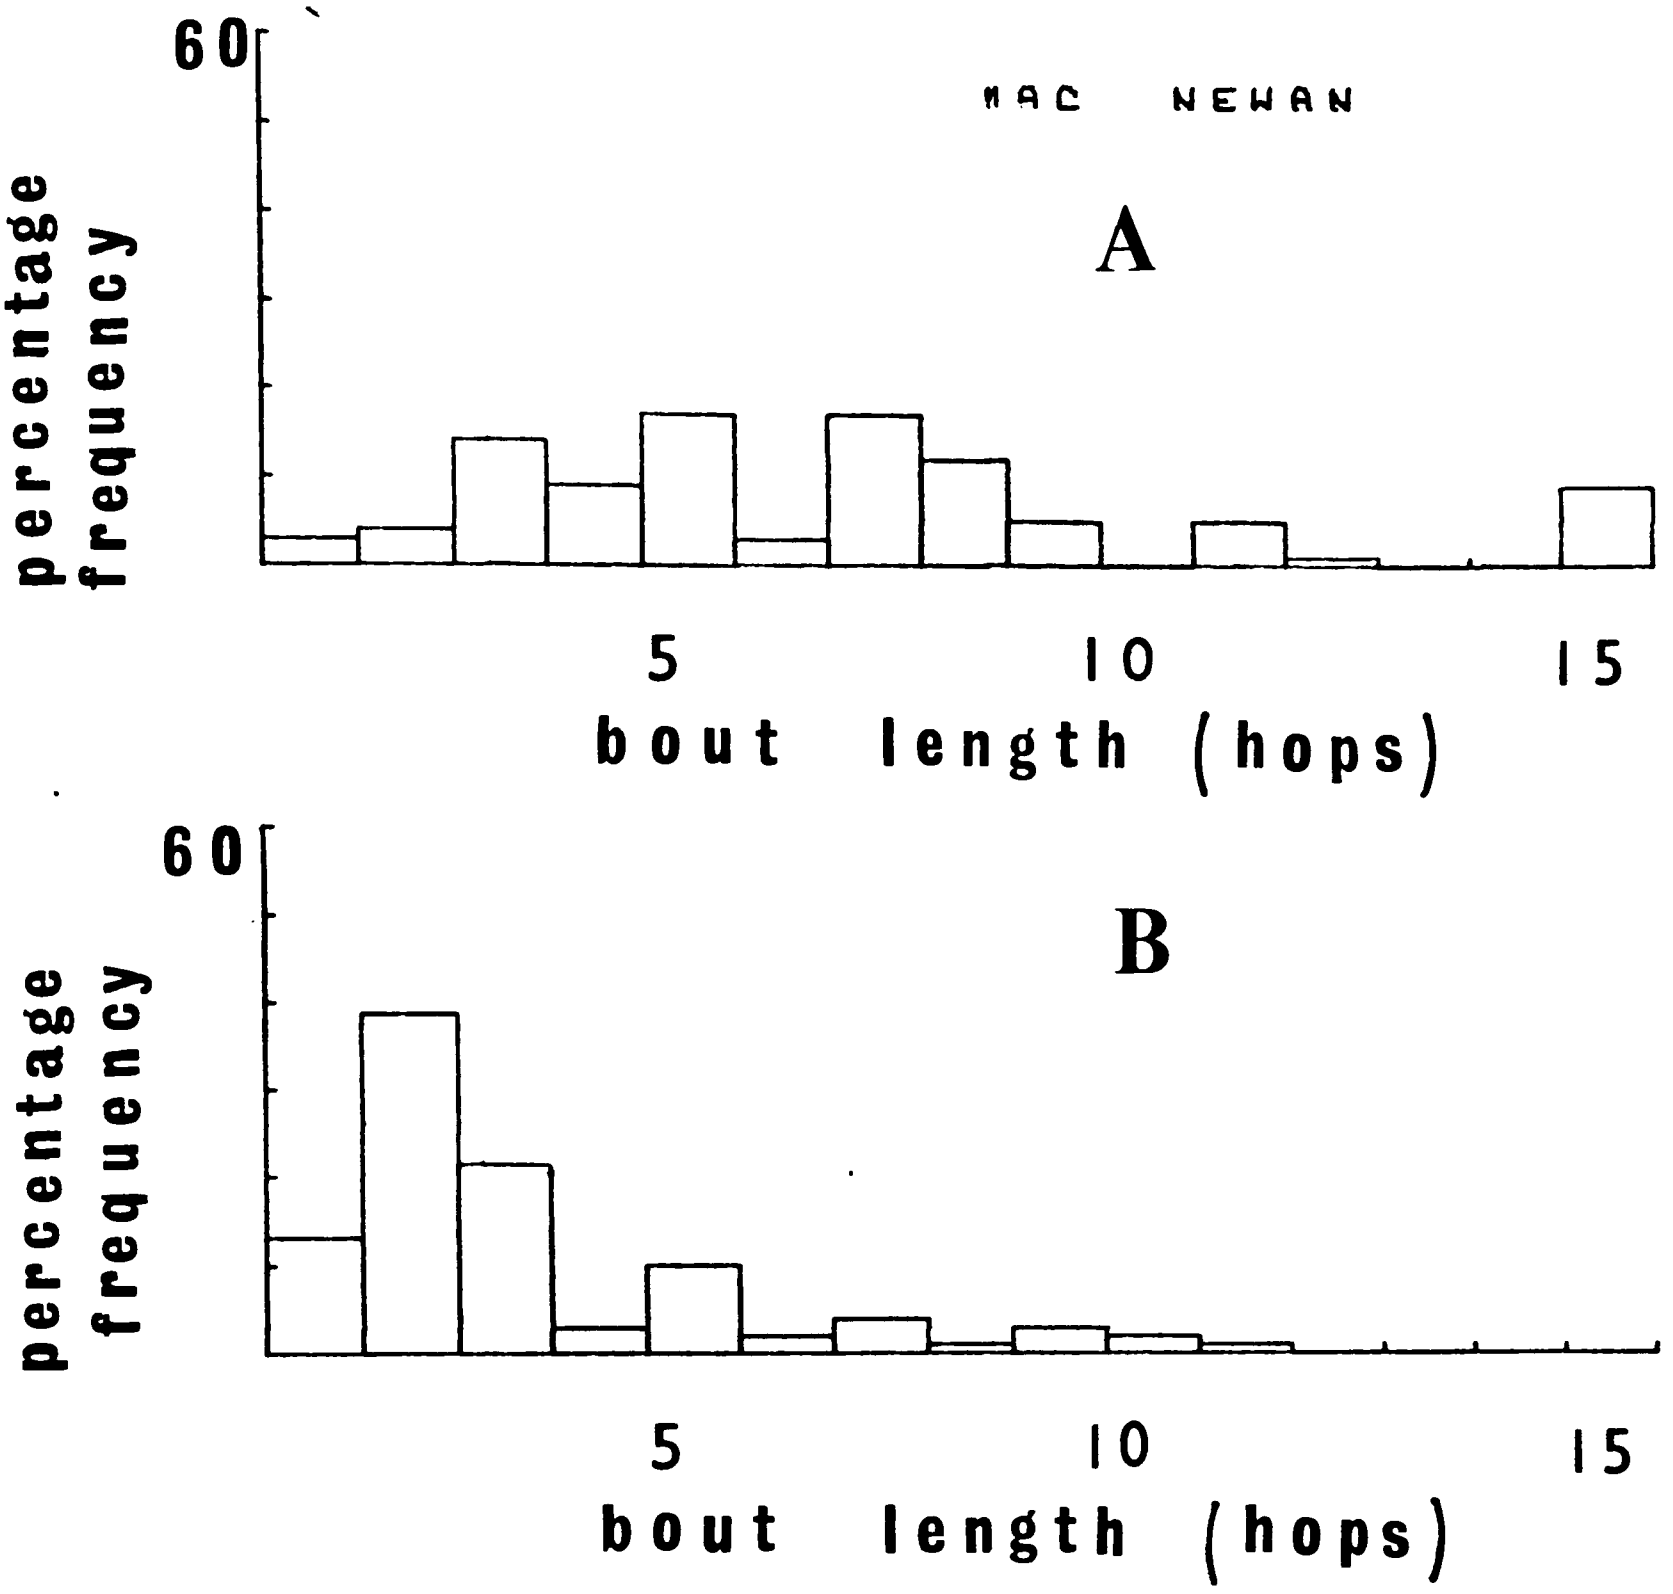






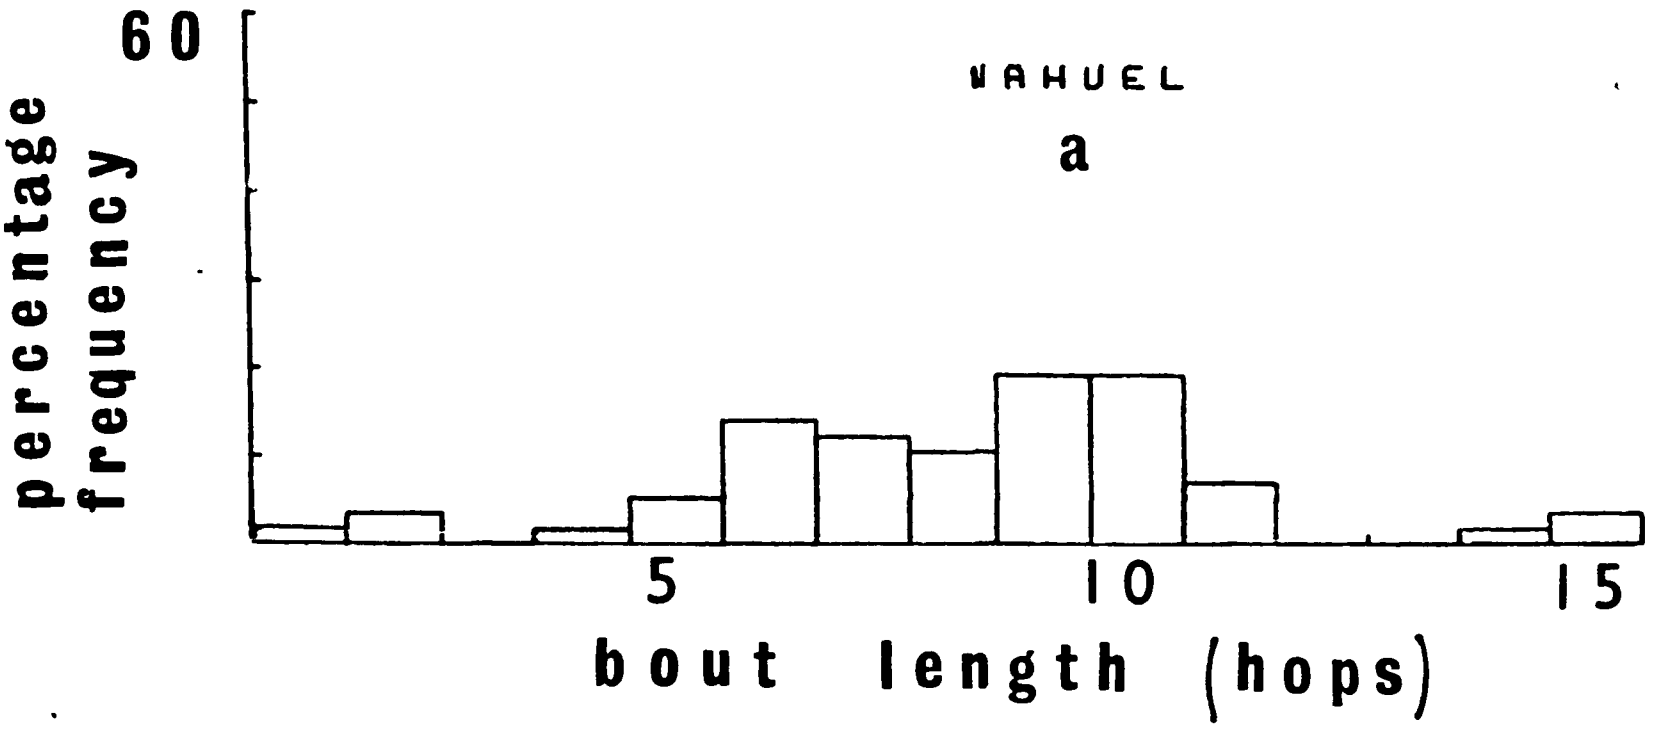
Cumulative rate of intake as a function of time in a patch. The optimum visit length is indicated by an arrow. It results from applying the method described in Figure 10.4(a), but in this figure travel time is included in the abscissa. TT: travel time; HT: handling time; IHI: inter hop internal

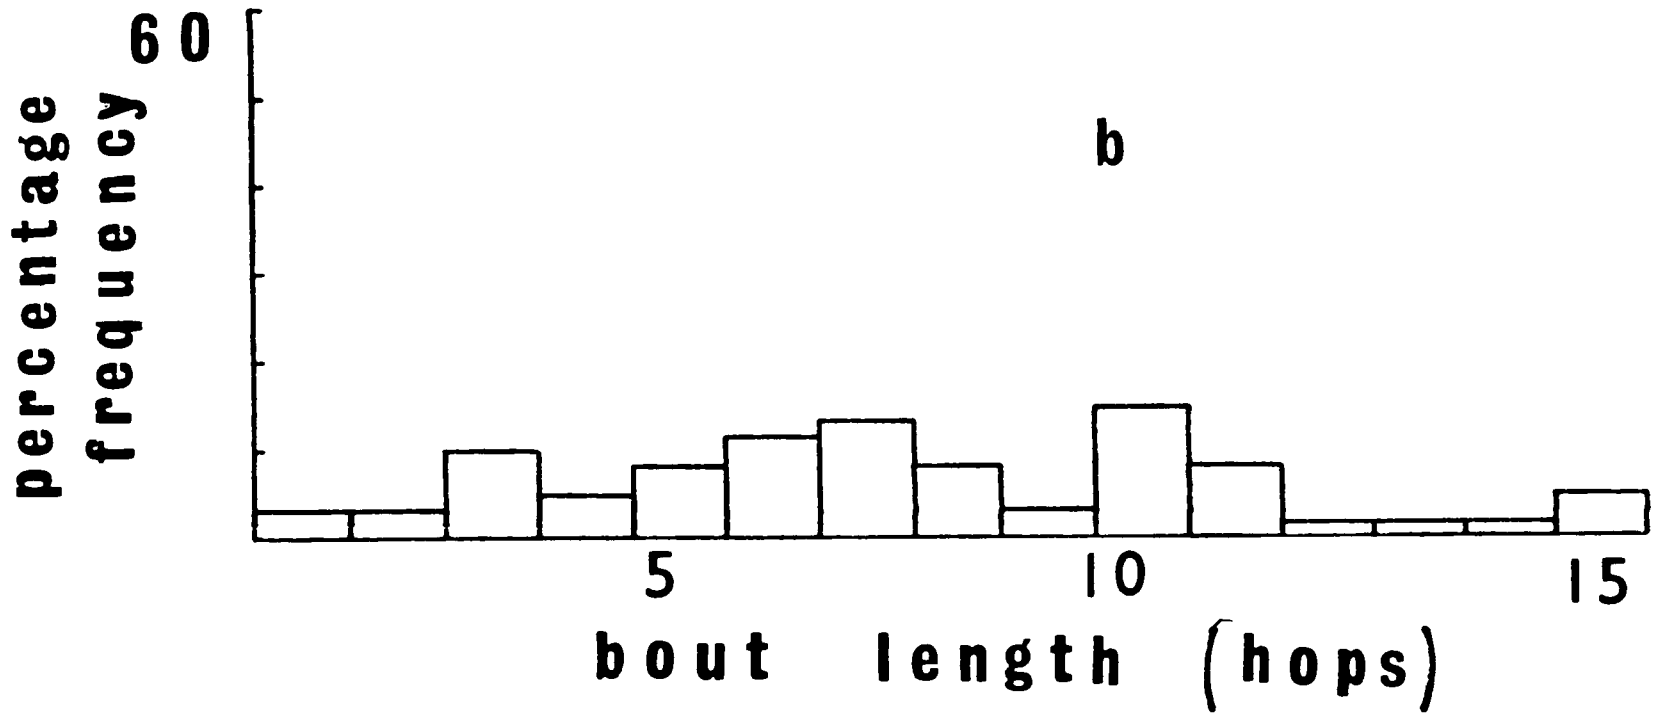
Figure 10.5



Frequency distribution of bout lengths for Mac Newan in control sessions. a: patch 1; b: patch 2. This bird used both patches differentially (see Table 10.1, legend). The number of hops in patch 2 was much smaller than in patch 1, which was used to do all the analysis.

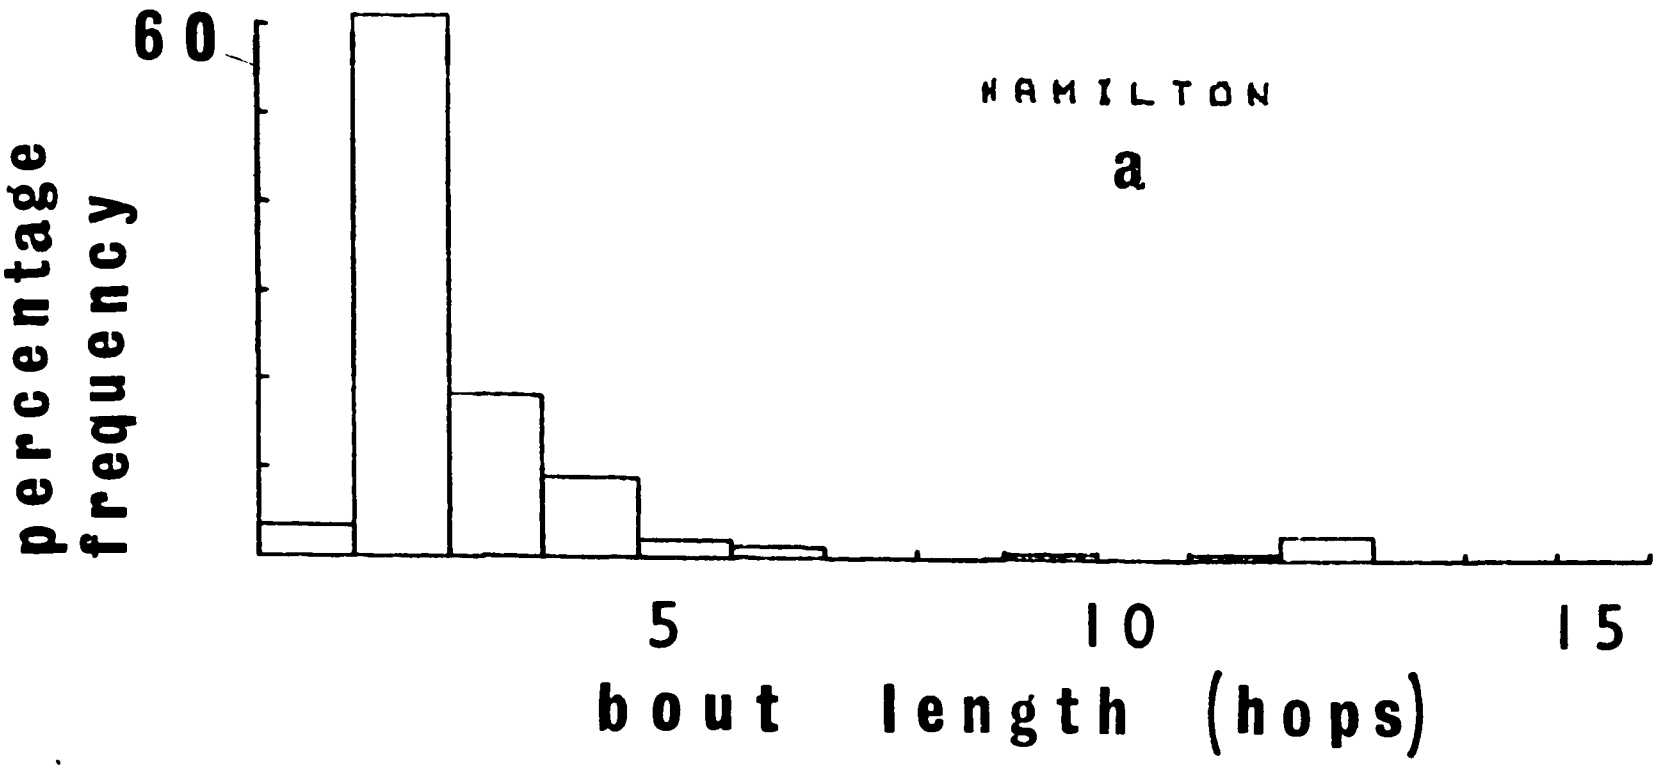
Table 10.6

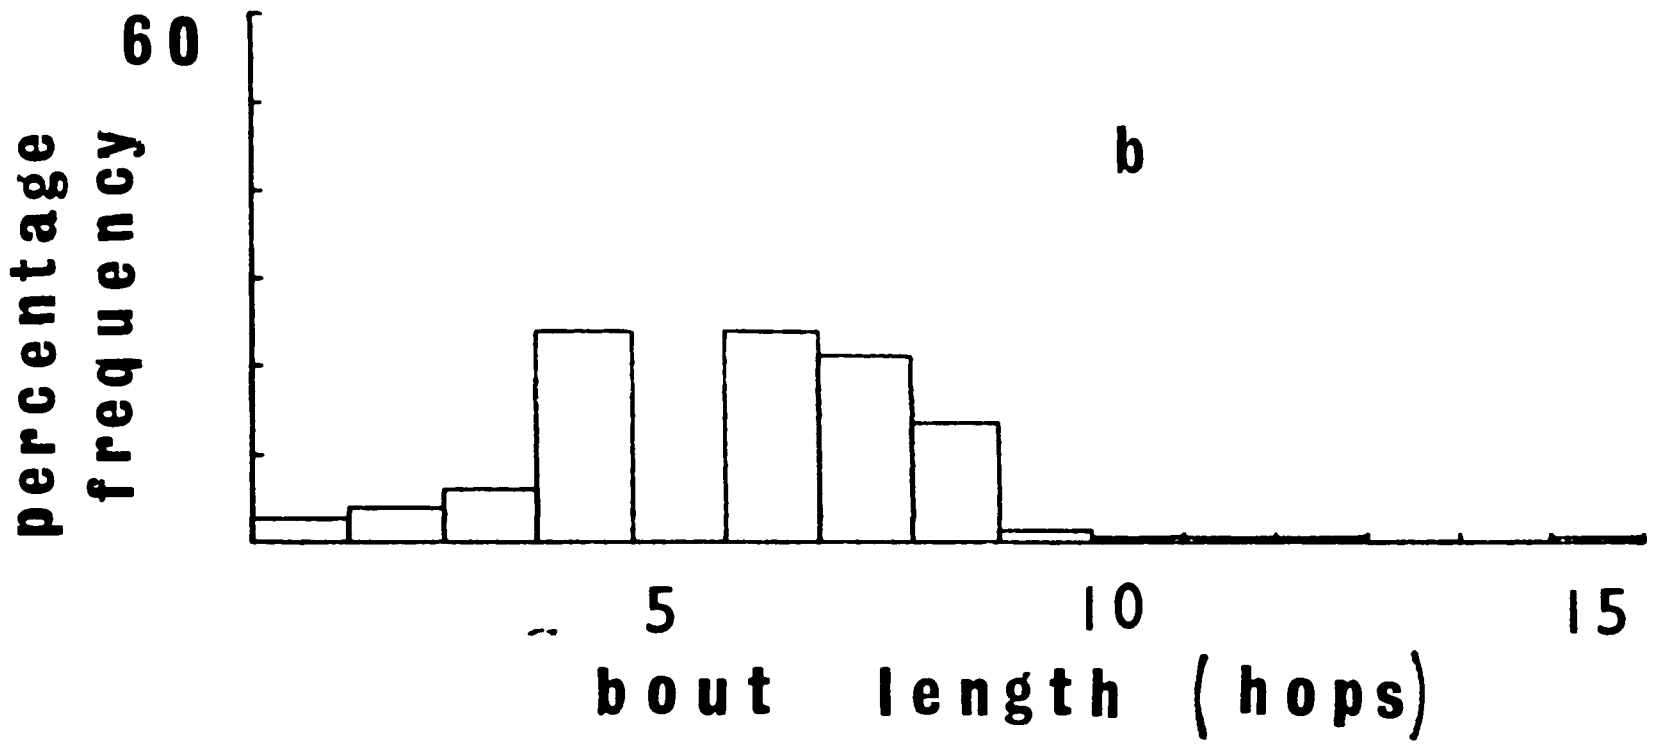




Frequency distribution of bout lengths for Nahuel in control sessions. a: patch 1; b: patch 2.

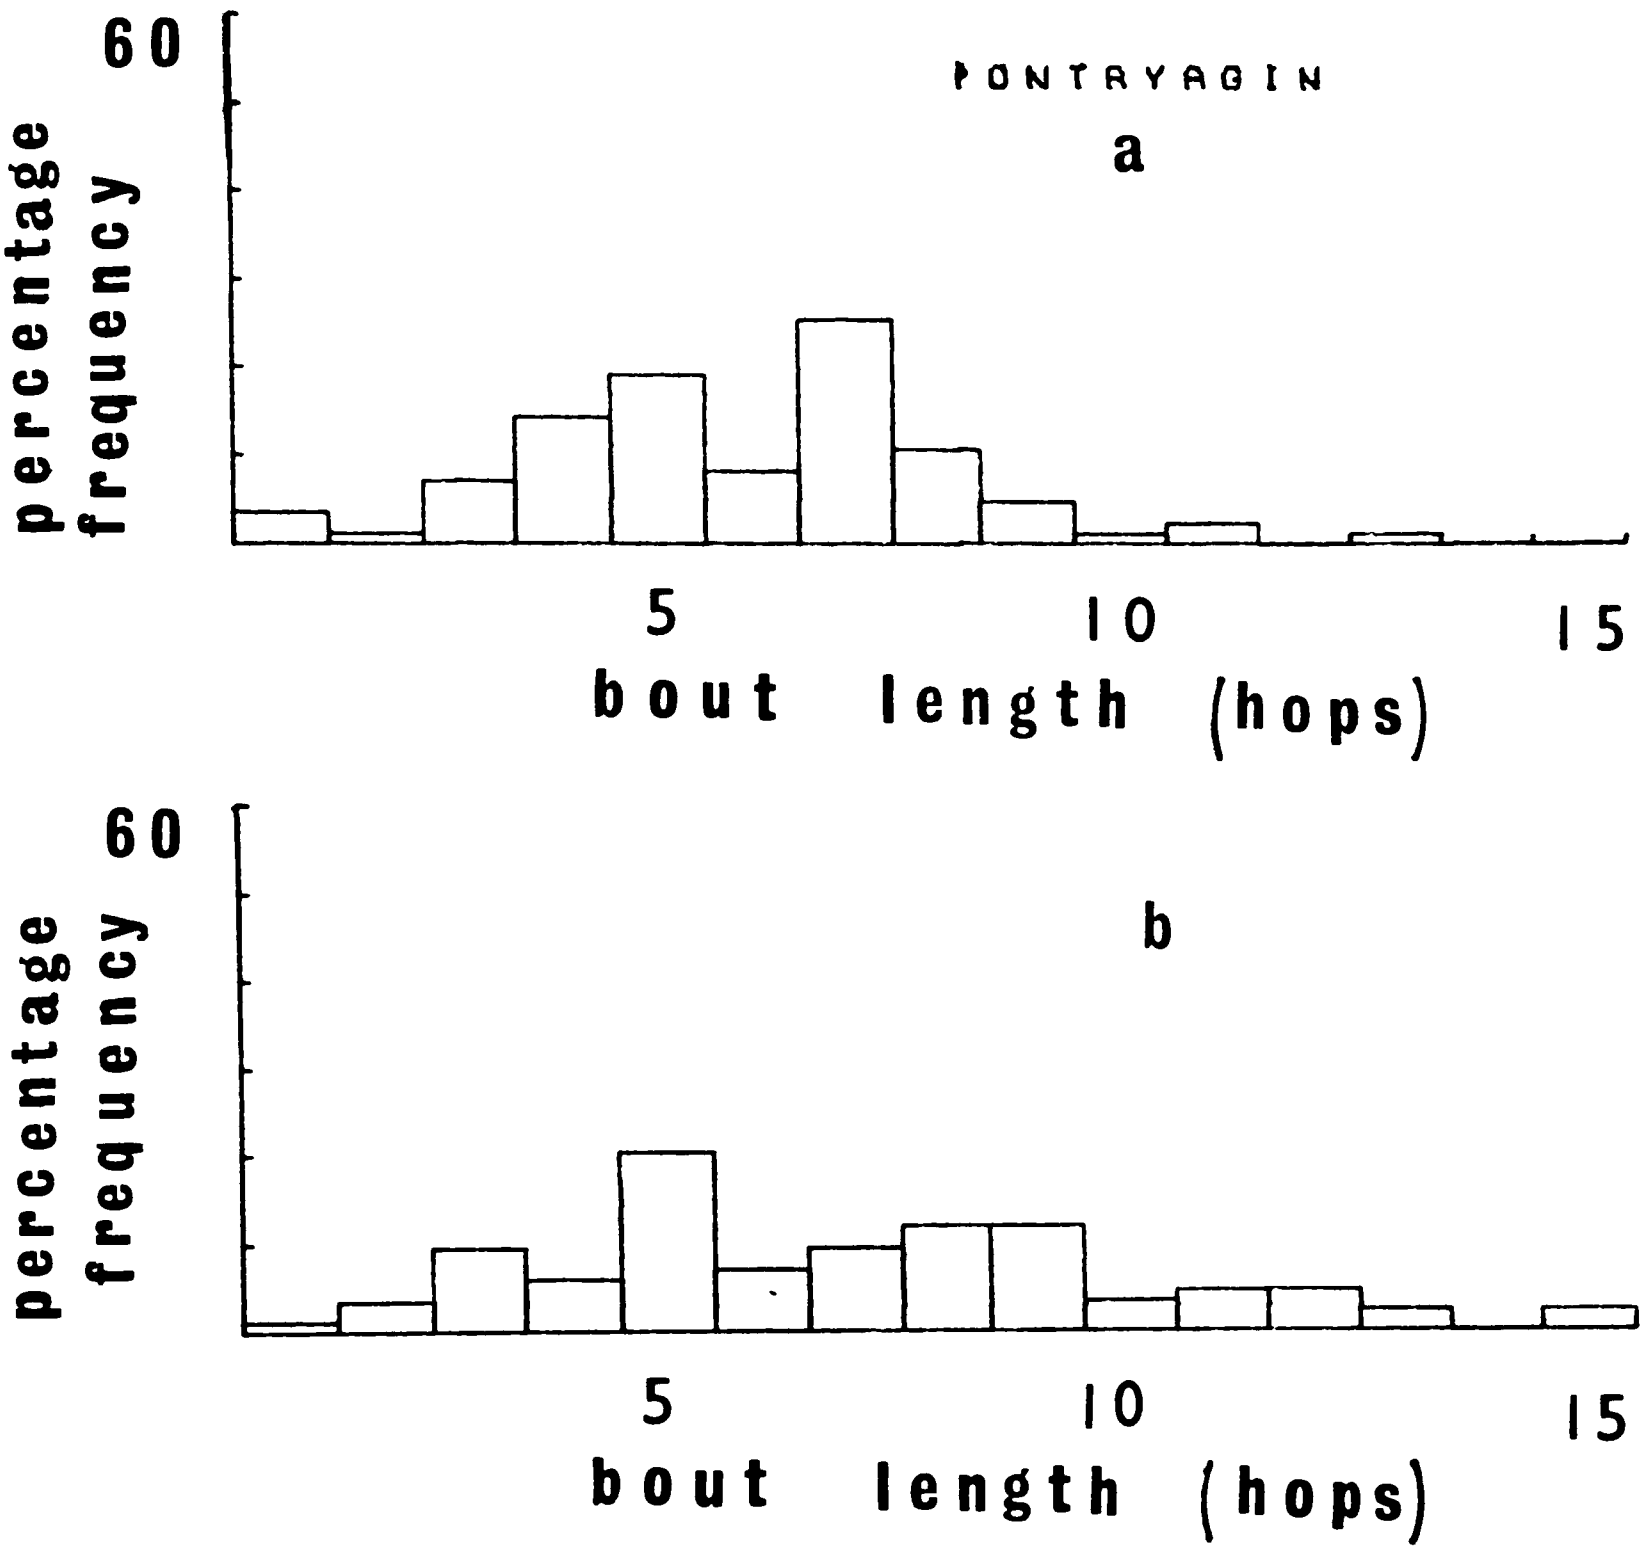
Figure 10.7





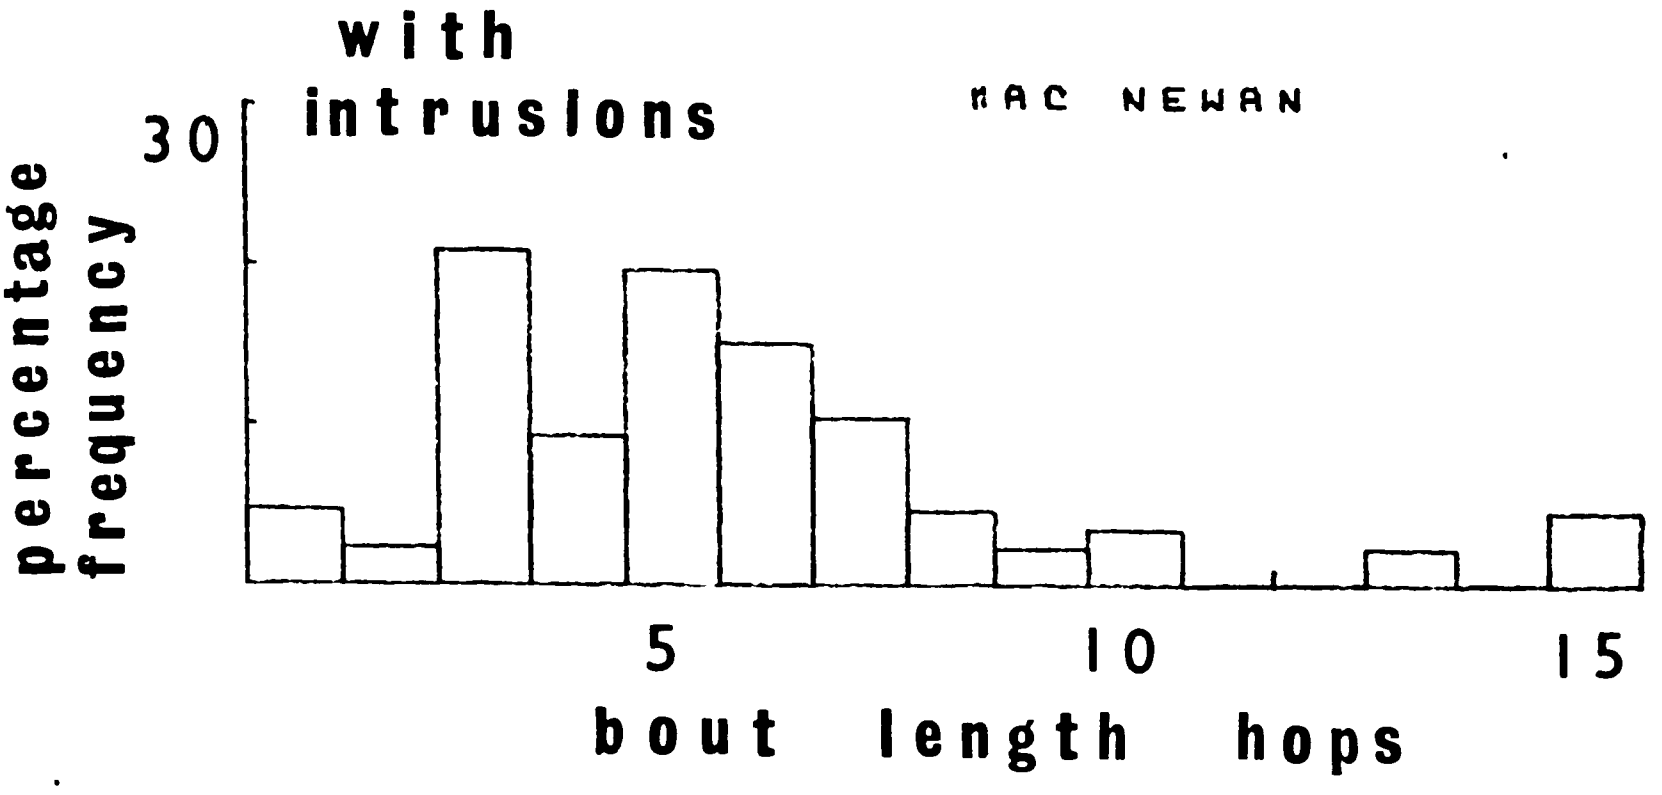
Frequency distribution of bout lengths for Hamilton in control sessions. a: patch 1; b: patch 2. As for Mc Newan (fig. 10.5), this bird used both patches differentially (see Table 10.1 legend), and only patch 2 was used for analysis.

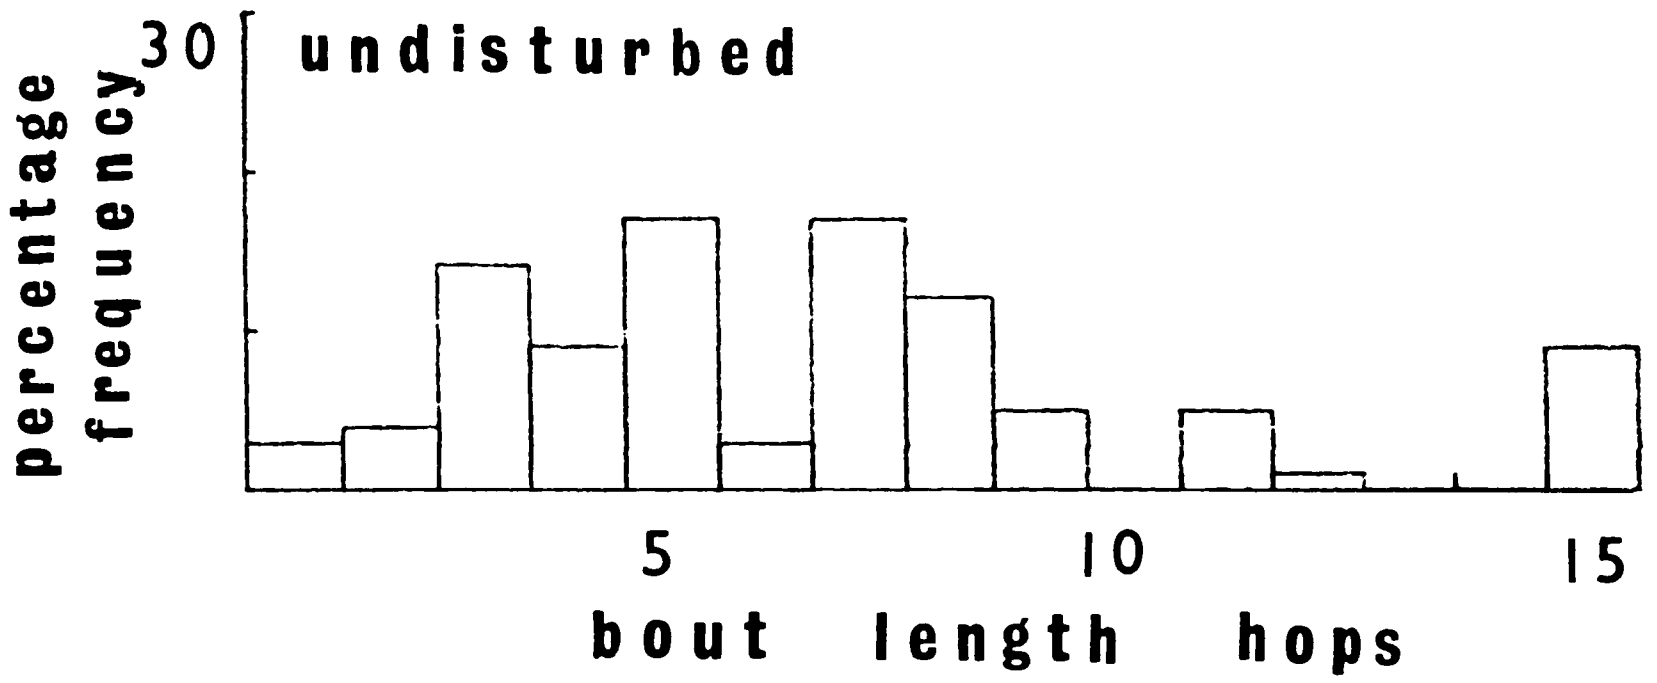
Figure 10.8



Frequency distribution of bout lengths for Pontryagin in control sessions. a: patch 1; b: patch 2.

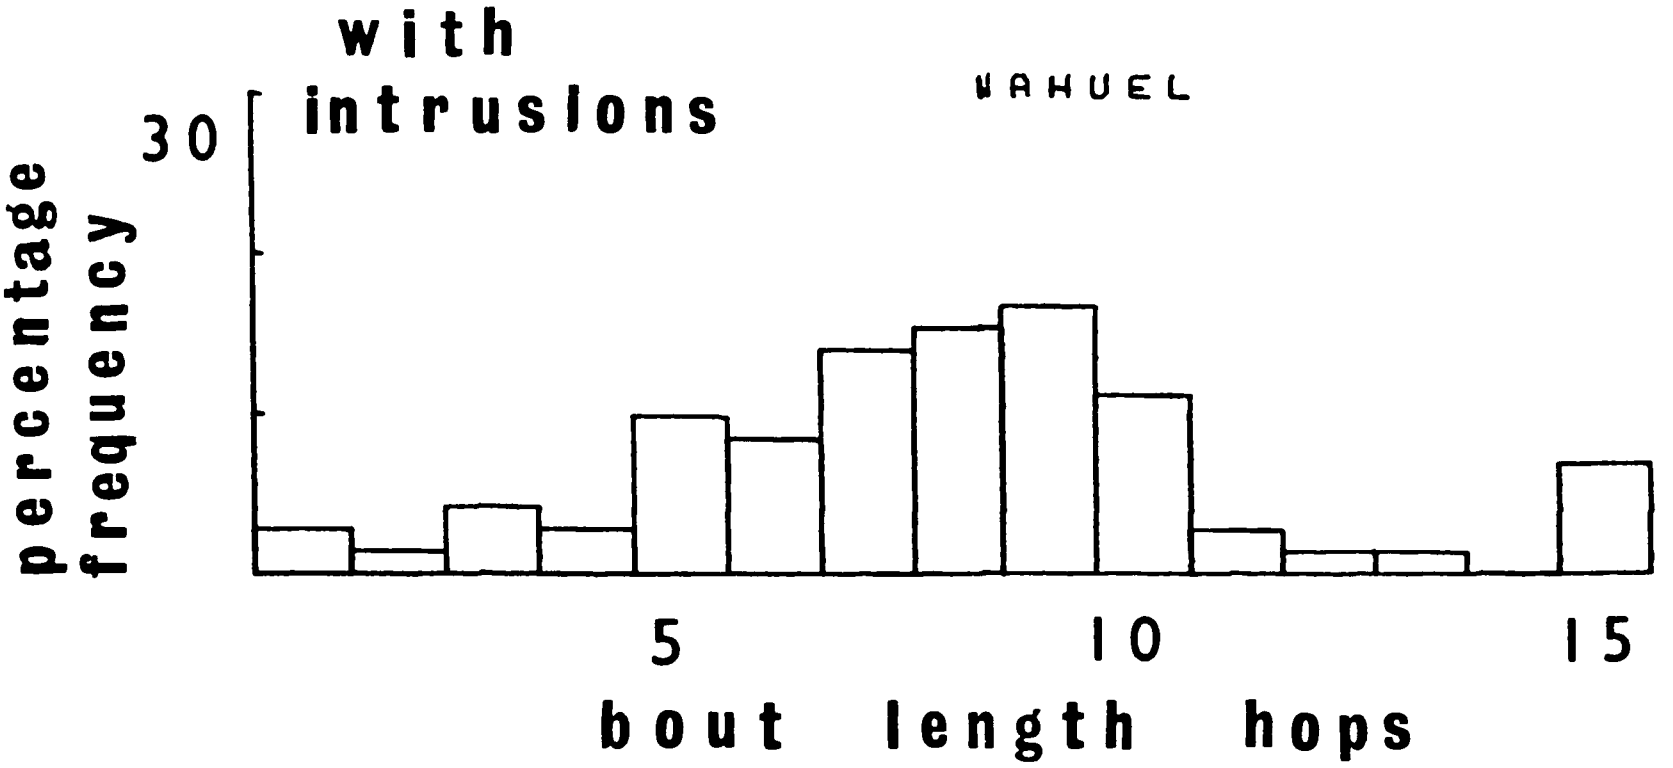
Figure 10.9

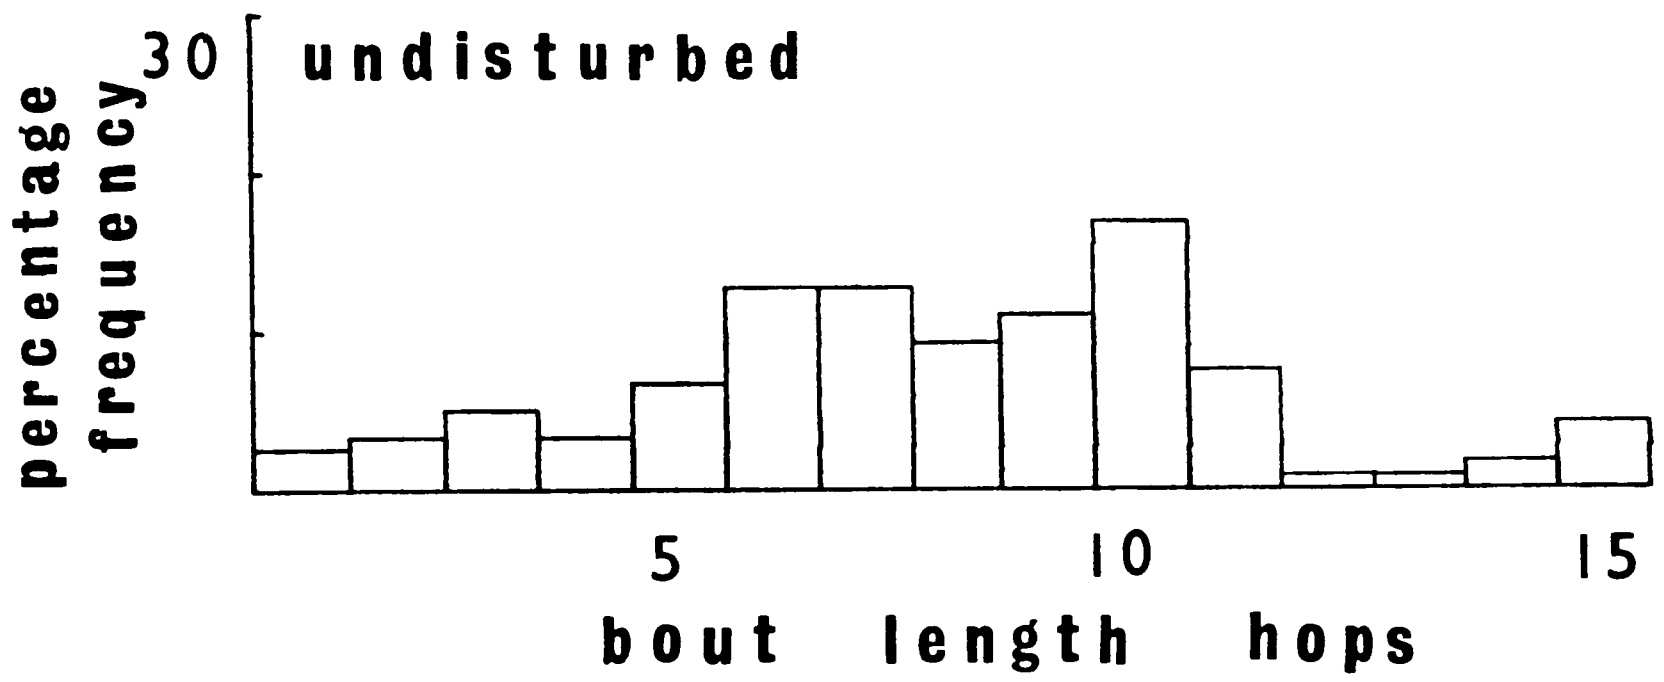




Frequency distribution of bout lengths in the preferred patch for Mac Newan. Intrusion and undisturbed sessions The effect of intrusions is to shift the distribution towards shorter bouts.

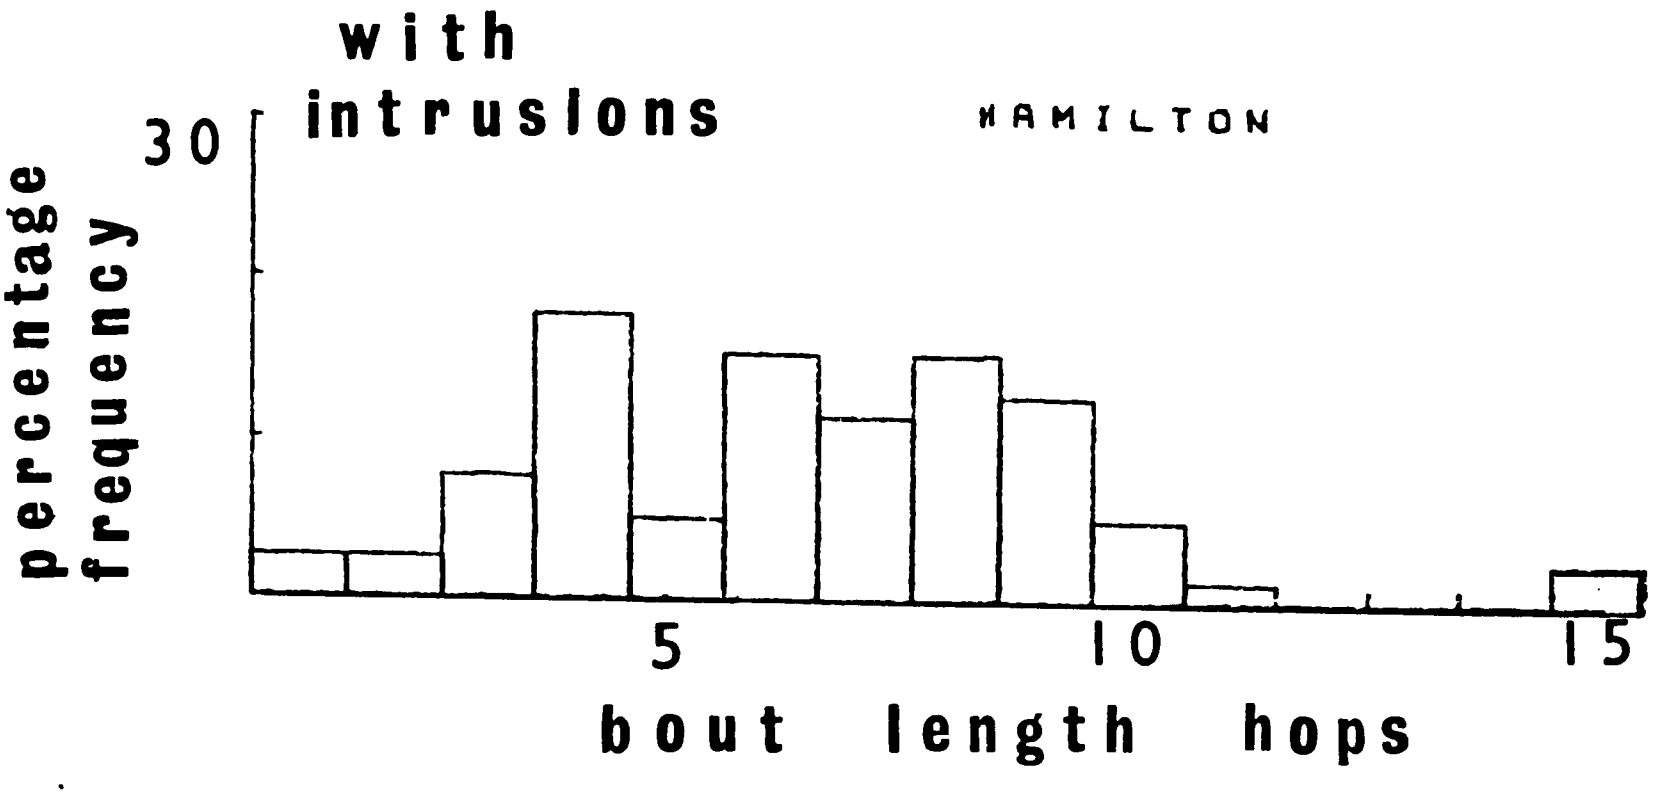
Figure 10.10

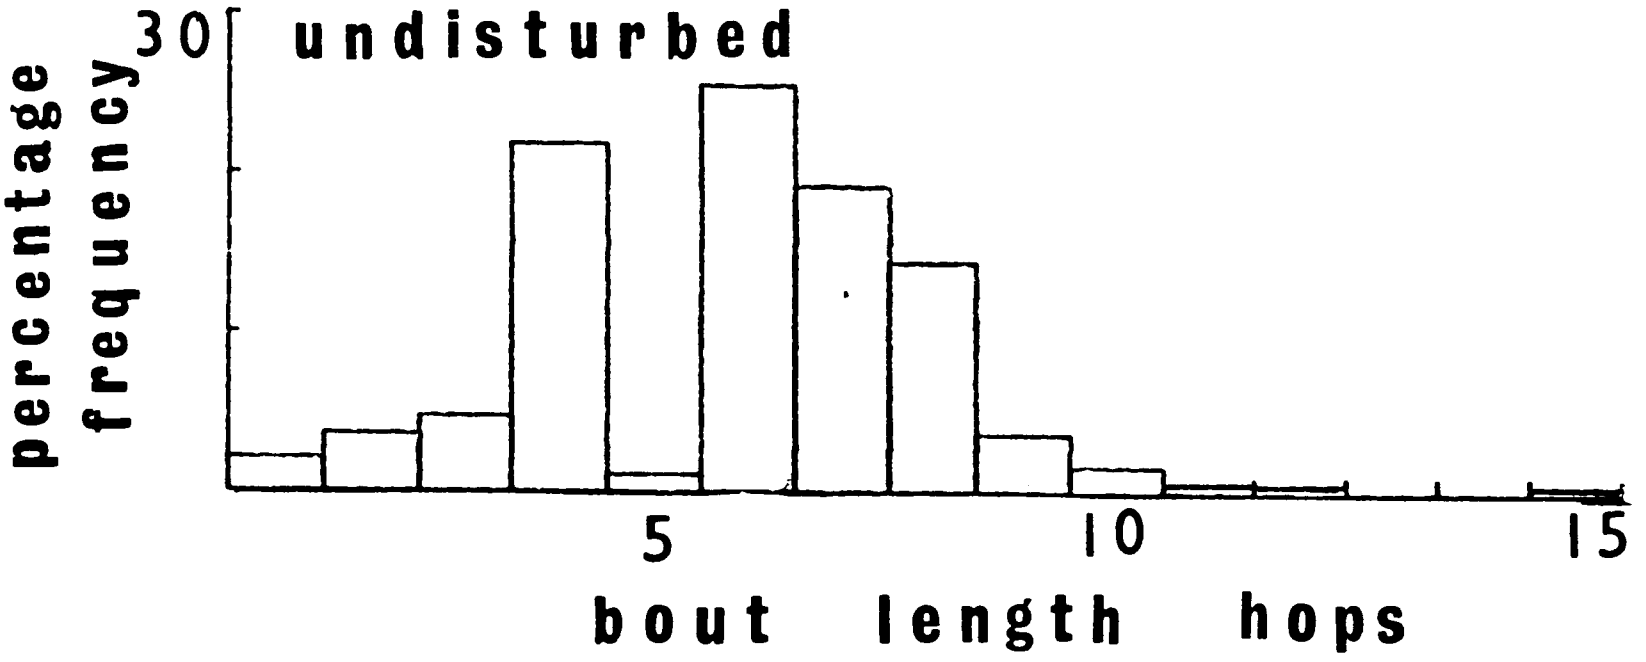




Frequency distribution of bout lengths in both patches for Nahuel. Intrusion and undisturbed sessions. There is no apparent effect of intrusions.

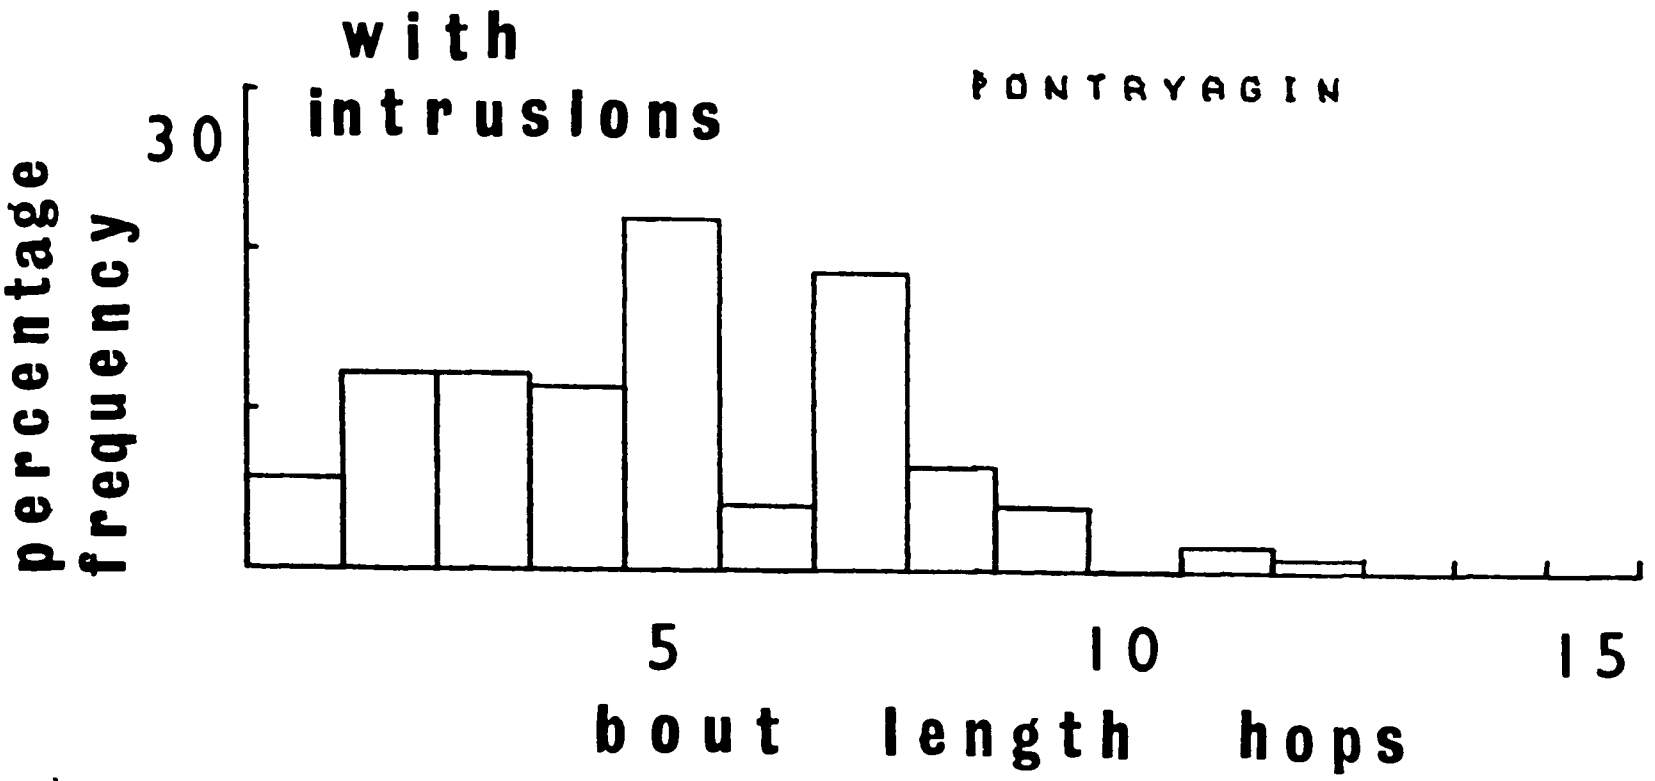
Figure 10.11

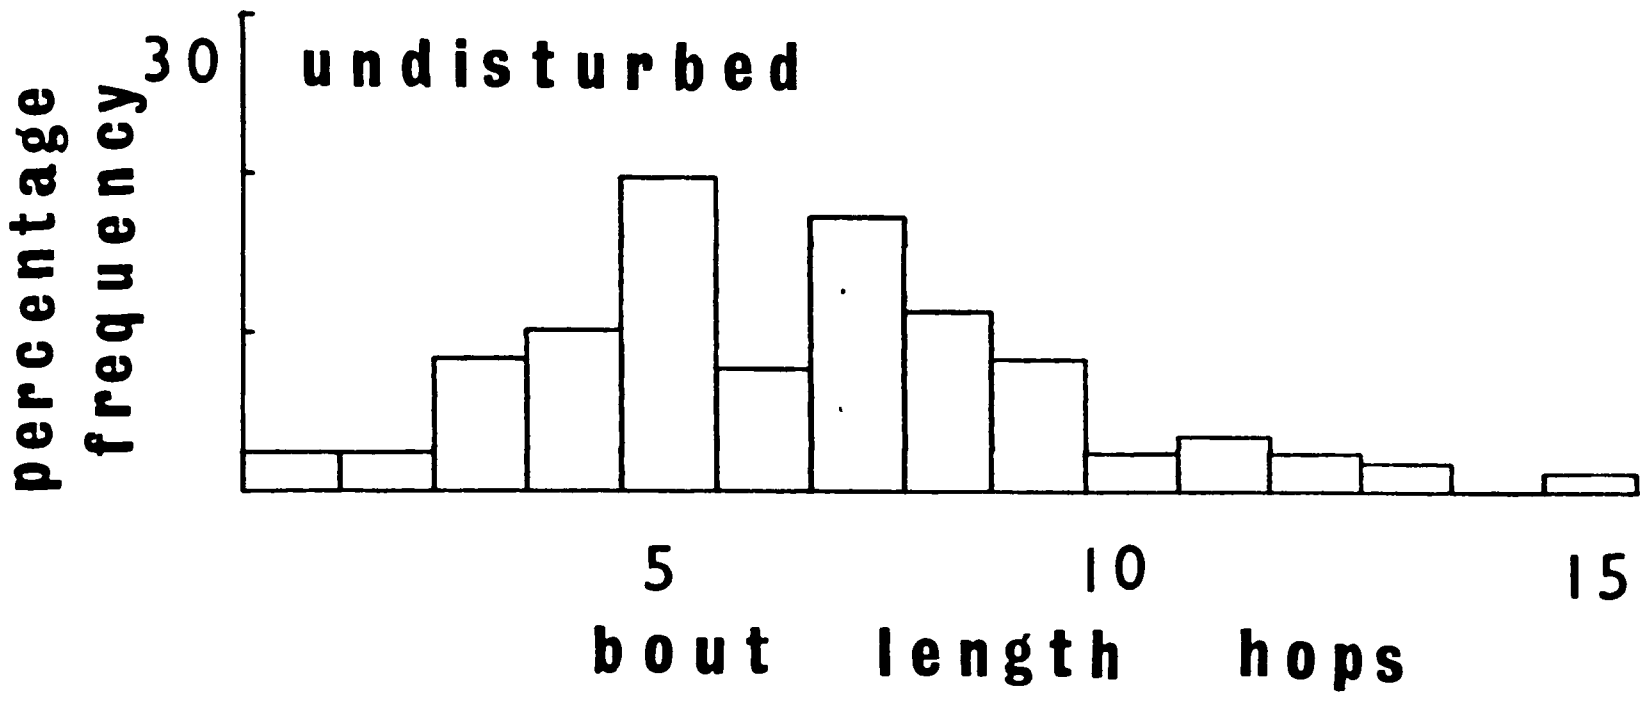




Frequency distribution of bout lengths for Hamilton. Intrusion and undisturbed sessions. Intrusions provoke some (non statistically significant) shift to shorter bouts.

Figure 10.12

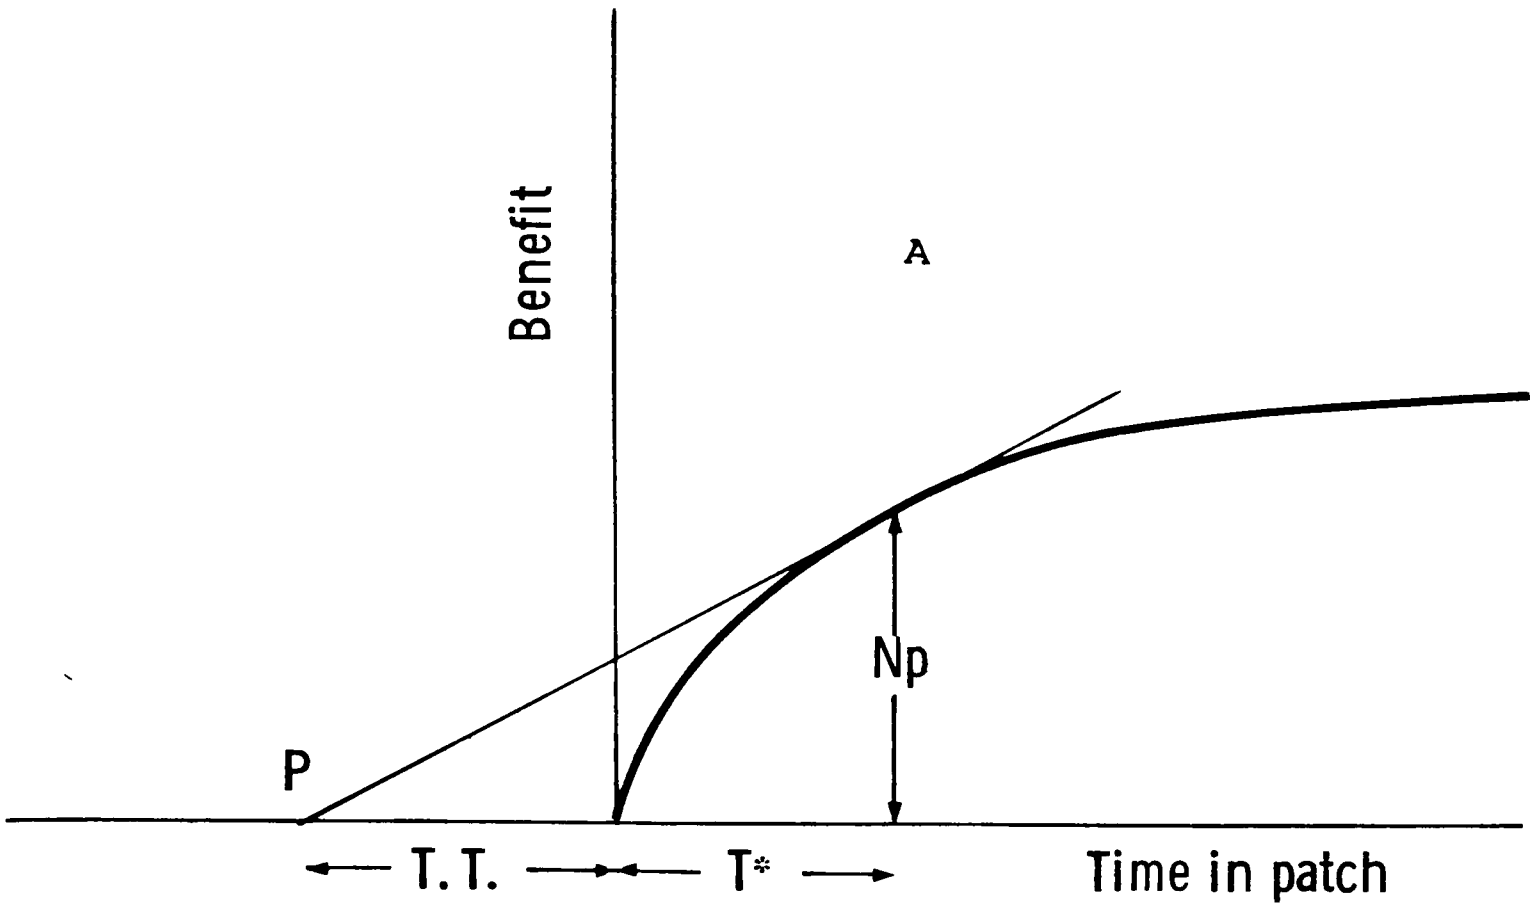


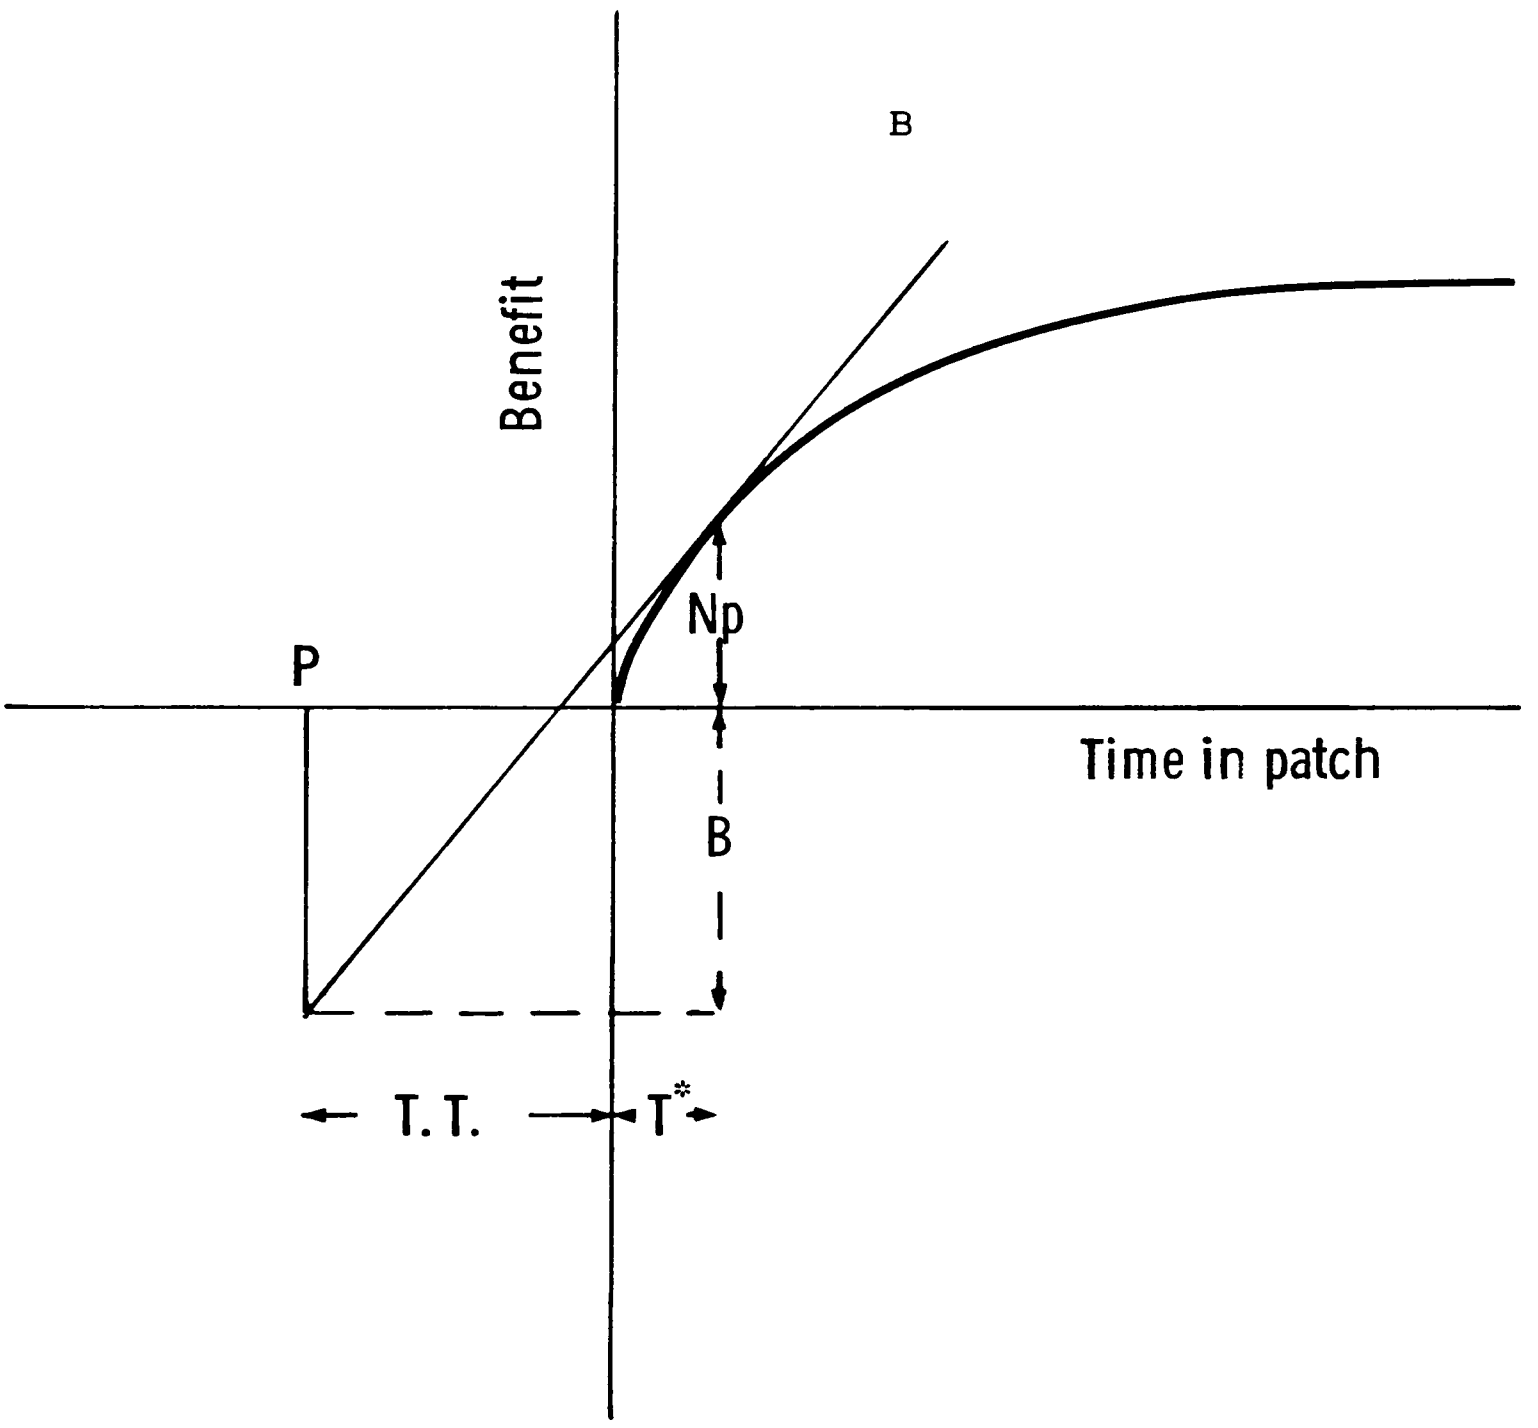


Frequency distributions of bout lengths for Pontryagin. Intrusion and undisturbed sessions. Intrusions provoke shift towards shorter bouts.

Figure 10.13 An inverse optimality method to estimate the territorial benefit required to make a given change in bout length advantageous. A is a replicate of Figure 10.4(a), and the same notation is used. In B, instead of tracing the tangent which contains P, the tangent is traced touching the cumulative benefit curve at the point of the observed time in the patch (or bout length). The ordinate value of the intersection of the new tangent with the projection of P in the x axis is the territorial benefit "B" that would have been required to make the new bout length optimum.

Figure 10.13


A
Benefit
Np
P
T.T.
T*
Time in patch


B
Benefit
Np
P
B
T.T.
T*
Time in patch

R E F E R E N C E S

Astrom, Gustav (1976). Environmental influence on daily song activity of the Reed Bunting Emberiza schoeniclus (L). Zoon, Suppl 2 (82pp)

Baldwin S.P. & Kendeigh S.C. (1938). Variations in the weight of birds. The Auk, 55 416-467

Aschoff, J. (1960). Exogenous and endogenous components in circadian rhythms. Cold Spring Harbor Symposium of Quantitative Biology 25 11-28

Barnard, C. (1979). Aspects of winter flocking and food fighting in the house sparrow. D.Phil. Thesis, University of Oxford

Baum, W.M. (1974). On two types of deviations from the matching law. Journal of Experimental Analysis of Behavior, 22 231-42.

Bellman, R. (1956). A problem in the sequential design of experiments. Sankhya 16 221-229.

Belovsky, G.E. (1978). Diet optimization on a generalist herbivore: the moose. Theoretical Population Biology 14 105-134

Betts, M.M. (1955). The food of titmice in oak woodland. Journal of Animal Ecology, 24 282-323

Bitterman, M.E. (1965). Phyletic differences in learning. American Psychologist, 20 396-410.

Blurton-Jones, N.G. (1968). Observations and experiments on causation of threat displays of the great tit (Parus major). Animal Behaviour Monographs, 1 74-158.

Bolles, R.C. (1962). The readiness to eat and drink: the effect of deprivation conditions.
Journal of Comparative and Physiological Psychology 55 230-234.

Carpenter, F.L. & MacMillen, R.E. (1976). Threshold model of feeding territoriality and test with a Hawaiian Honeycreeper.
Science, 194 639-642.

Catchpole, C.K. (1973). The functions of advertising song in the sedge warbler (Acrocephalus schoenobaenus) and the reed warbler (Acrocephalus scirpaceus).
Behaviour, 46 300-320.

Chaplin, S.B. (1974). Daily energetics of the black capped chickadee Parus atricarpillus in winter.
Journal of Comparative Physiology, 89 321-330.

Charnov, E.L.F. (1976). Optimal foraging: attack strategy of a mantid.
American Naturalist, 110 141-151.

Charnov, E.L.F. (1976). Optimal foraging: the marginal value theorem.
Theoretical Population Biology, 9 129-136.

Charnov, E.L.F.; Orians G.; Hyatt, K.. (1976). Ecological implications of resource depression.
The American Naturalist, 110 247

Cicala, G.A. (1961). Running speed in rats as a function of drive level and presence or absence of competing response trials.
Journal of Experimental Psychology, 62 329-334.

Collier, G. (1980). A comparative analysis of optimal foraging Behaviour, Laboratory Simulations. In press.

Cook, R.M. & Cockrell, B.J. (1978). Predator ingestion rate and its bearing on feeding time and the theory of optimal diets.
Journal of Animal Ecology, 47 451-469.

major).
Nature, 268 137-139.

Cowie, R.J. (1979). Foraging behaviour of the great tit (Parus major)(L.) in a patchy environment.
D.Phil. Thesis. University of Oxford.

Craig, R.B.; De Angelis, D.L.; Dixon, K.R. (1979). Long and short-term dynamic optimization models with application to the feeding strategy of the loggerhead shrike.
The American Naturalist, 113 31-51.

Daan, S. (1976). Light intensity and the timing of daily activity on finches (Fringillidae).
Ibis, 118 223-236.

Daan, S. & Aschoff, J. (1975). Circadian rythms of locomotor activity in captive birds and mammals: their variation with season and latitude.
Oecologia, 18 269-316.

Davies, N. (1977). Prey selection and the search strategy of the spotted fly catcher (Muscicapa striata): a field study on optimal foraging.
Animal Behaviour, 25 1016-1033.

Dawkins, R. (1971). A cheap method of recording behavioural events for direct computer access.
Behaviour, 40 162-173.

Dawkins, M. & Dawkins, R. (1974). Some descriptive and explanatory stochastic models of decision-making.
Motivational Control Systems Analysis
McFarland, D.J. (editor).

De Groot, M.H. (1970). Optimal statistical decisions.
New York, Mc Graw Hill

De Villiers, P.A. (1977). Choice in concurrent schedules and quantitative formulations of the law of effect.
Operant Behaviour II. Konig, W.K. & Staddon, J. (Eds);
New York, Prentice Hall.

Delany, M.E. (1977). Sound propagation in the atmosphere: a historical review.
Acustica, 38 201-223.

Doring, (1920). Das Erwachen der Vogelwelt im hoheren sachsischen Erzgebirge zu den verschiedenen Jahreszeiten. Tharandt, forstliche Jahrbuch, 71 242-263.

Dunn, E. (1977). Predation by weasels (Mustela nivalis) on breeding tits (Parus spp) In relation to the density of tits and rodents
Journal of Animal Ecology, 46 633--52.

Dunnet, G.E. & Hinde, R.A. (1953). The winter roosting and awakening behaviour of captive great tits.
British Journal of Animal Behaviour, 1 91-95.

Fordham, R.A. (1978). Differential intensity of moorhen (Gallinula chloropus) (Rallidae) feeding at dawn and dusk in Spring.
New Zealand Journal of Ecology, 1 109-117.

Geertsema, S. & Redingius, H. (1974). Preliminary considerations in the simulation of behaviour.
Motivational Control Systems Analysis, McFarland, D J. (ed), London, Academic Press.

Gibb, J. (1954). Feeding ecology of tits, with notes on Treecreeper and Goldcrest.
Ibis, 96 513-543.

Green, R.F. (1979). Bayesian birds: a simple example of Oaten's stochastic model of optimal foraging
Technical Report No. 50, Department of Statistics University of California at Riverside.

Grossman, A.F. & West G.C. (1977). Metabolic rate and temperature regulation of Winter acclimatized black-capped Chickadees, Parus atricarpillus of interior Alaska.
Ornis Scandinavica, 8 128-138.

Gunn, D.L. & Hopf H.S. (1941). The biology and behaviour of Ptinus tectus Boie (Coleoptera: Ptinidae) a pest of stored products . II The amount of locomotory activity in relation to experimental and previous temperatures.
Journal of Experimental Biology, 18 278-289.

Haecker, V. (1916). Reizphysiologisches uber den Abendgesang der Vogel.
Pflug Arch, ges. Physiologie, 204 718-725.

Haftorn, S. (1951). En underskelse over vekt-variasjoner i vinterhalvaret hos ktottmeis.
Norsk. Zool. Foren. Tidsskr. Fauna, 2 83-91.

Hailman, J.P. (1978). Optical signals
Animal Communication and Light
Indiana University Press. Bloomington and London.

Hartley, P.H.T. (1953). An ecological study of the feeding habits of the English titmice.
Journal of Animal Ecology, 22 261-288.

Heiligenberg, W. (1974). Processes governing behavioural states of readiness.
Advances in the Study of Behaviour; Volume 5.
Academic Press. New York and London.

Henwood, K & Fabrick, A. (1979). A quantitative analysis of the dawn chorus: temporal selection for communicatory optimization
American Naturalist, In Press.

Herrnstein, R.J. (1961). Relative and absolute strength of response as a function of frequency of reinforcement.
Journal of Experimental Analysis of Behaviour, 7 151-157.

Herrnstein, R.J. (1970). On the law of effect.
Journal of Experimental Analysis of Behaviour, 13 243-266.

Herrnstein, R.J. & Loveland, D.H. (1975). Maximising and matching on concurrent ratio schedules.
Journal of Experimental Analysis of Behaviour, 24 107-116.

Heyman, G.M. (1979). A Markov model description of changeover probabilities on concurrent variable interval schedules.
Journal of Experimental Analysis of Behaviour, 31 41-53

Heyman, G.M. & Luce, R.D. (1980). Operant matching is not a logical consequence of maximising reinforcement rate. Manuscript.

Hinde, R.A (1952). The behaviour of the great tit. Behaviour, Supp II.

Hinde, R.A. (1970). Animal behaviour (2nd Edition) New York, Mc Graw Hill

Hissa, R. & Palokangas, R. (1970). Thermoregulation in the titmouse (Parus major L.) Comparative Biochemistry and Physiology, 33 941-953.

Holling, C.S. (1965). The functional response of predators to prey density and its role in mimicry and population regulation. Memories of the Entomological Society of Canada, 45 1-60

Houston, A.I. (1977). Models of animal motivation. D.Phil. Thesis. University of Oxford

Hunter, M. Jnr. (1978). Some environmental effects on singing behaviour of great tits D.Phil. Thesis. University of Oxford.

Jones, P.W. (1975). The two-armed bandit. Biometrika, 62 523-524.

Jones, P.W. (1976). Some results for the two-armed bandit problem. Mathemalisch Operations forschung und Statistik, 7 471-475.

Katz, P.L. (1974). A long-term approach to foraging optimization The American Naturalist, 108 758-783.

Keesey, R.E.; Boyle, P.C.; Kemnitz, J. W.; Mitchell, J.S. (1970). The role of lateral hypothalamus in determining the body weight set point. Hunger: Basic Mechanisms and Clinical Implications Novin, D.; Wyrwicka, W.; Bray, G.; (editors). New York, Raven Press.

Kendeigh, C.; Kontogiannis, J.; Mazac, A.; Roland, R.; (1969) Environmental regulation of food intake by birds. Journal of Comparative Biochemistry and Physiology, 31 941-957.

Kenward, R.E. & Sibly R.M. (1977). A woodpigeon (Columba palumbus) feeding preference explained by a digestive bottle-neck. Journal of Applied Ecology, 14 815-826.

Kikkawa, J. (1968). Social hierarchy in winter flocks of the grey-breasted silvereye Zosterops lateralis (Latham). Japanese Journal of Ecology 18 235-246.

King, J.R. (1974). Seasonal allocation of time and energy resources in birds.
Avian Energetics. Paynter, R.A. (editor).
Nuttal Ornithological Club
Cambridge, Massachusetts.

Kinsler, L. & Frey A. (1962). Fundamentals of acoustics.
John Wiley & Sons Inc.
New York

Kluijver, H.N. (1950). Daily routines of the great tit.
Ardea, 38 99-135.

Kluijver, H.N. (1952). Roosting habits, sexual dominance and survival in the great tit.
Cold Spring Harbor Symposia on Quantitative Biology.

Krebs, J.R. (1970). Regulation of numbers in the great tit (Aves: Passeriformes).
Journal of Zoology (London), 162 317-333.

Krebs, J.R. (1971). Territory and breeding density in the great tit.
Ecology, 52 2-22.

Krebs, J.R. (1977). Song and territory in the great tit Parus major
Evolutionary Ecology; Stonehouse, B. & Perrins C. (editors)

Krebs, J.R. (1978). Optimal foraging: decision rules for predators.
Krebs, J.R. & Davies N.B. (editors)
Behavioural Ecology
Oxford, Blackwells.

Krebs, J.R. & Cowie R.J. (1976). Foraging strategies in birds. Ardea 64 98-116.

Krebs, J.R.; Erichsen, J.T.; Webber, M.I.; Charnov, E.L. (1977). Optimal prey selection in the great tit (Parus major). Animal Behaviour, 25 30-38.

Krebs, J.R.; McRoberts, M.; Cullen, J.M. (1972). Flocking and feeding in the great tit - an experimental study. Ibis, 114 507-530.

Krebs, J.R. & Perrins C. (1977). Behaviour and population regulation in the great tit (Parus major) Population control by social behaviour. Symposium of the Institute of Biology, London.

Krebs, J.R.; Ryan, J.C.; Charnov, E.L. (1974). Hunting by expectation or optimal foraging? A study of patch use by chickadees. Animal Behaviour, 22 953-964.

Lack, D. (1971). Ecological isolation in birds Blackwell Scientific Publications Oxford and Edinburgh.

Lea, S. (1979). Supply as a factor in motivation. Abstracts: Conference "Analysis of Motivational Processes" Milton Keynes

Leas, S. (1980). Correlation and contiguity in foraging behaviour. Hargen, P. & Zeiler, M.D. (editors) London, Wiley (in press).

Lindley, D.V. & Barnett, V. (1965). Sequential sampling: two decision problems with linear losses in binomial and normal random variables. Biometrika, 52 507-532.

Lustick, S.I. (1969). Bird energetics: effects of artificial radiation
Science, 163 387-390.

Lutz, F.E. (1931). Light as a factor in controlling the start of daily activity of a wren and stingless bees.
American Museum Novitates, 468 1-9.

Mc Cleery, R.H. (1977). Temporal organisation of feeding and drinking behaviour.
D.Phil. Thesis, University of Oxford

Mc Cleery, R.H. Optimal behaviour sequences and decision making.
Behavioural Ecology
Krebs, J.R. & Davies N.B. (editors)
Oxford, Blackwell's.

McFarland, D.J. (1971). Feedback mechanisms in animal behaviour.
Academic Press
London

McFarland, D.J. (1977). Decision making in animals.
Nature, 269 15-21.

McFarland, D.J. & McFarland, F.J. (1968). Dynamic analysis of an avian drinking response
Medical Biological Engineering, 6 659-668.

Mackintosh, N.J. (1974). The psychology of animal learning
Academic
New York.

Mc Namara ,J. and Houston,A.I. (1979). The application of statistical decision theory to animal behaviour.
Manuscript.

Marten, K & Marker, P. (1977). Sound transmission and its significance for animal vocalization. I. Temperate habitats.
Behavioural Ecology and Sociobiology, 2 271-290.

Maynard Smith, J. (1978). Optimization theory in evolution
Annual Review of Ecology and Systematics, 9 31-56.

Milinski, M. & Heller, R. (1978). Influence of a predator on the optimal foraging behaviour of sticklebacks (Gasterosteus aculeatus L)
Nature, 275 642-644

Miller, L.N. (1978). Sound level of rain and of wind in the trees
Noise Control Engineering, II 101-109

Milsum, J.H. (1966). Biological control systems analysis.
New York, Mc Graw-Hill.

Morris, D. (1968). Some ethophysiological relations as shown in captive flocks of silvereyes (Zosterops lateralis, L.)
B. Sc. Thesis, University of Queensland
Australia.

Morse, D.H. (1978). Structure and foraging patterns of flocks of tits and associated species in an English woodland during the Winter.
Ibis, 120 298-312.

Morton, E.S. (1975). Ecological sources of selection on avian sounds
American Naturalist, 109 17-34

Mrosovsky, N. (1980). Animal anorexias
Manuscript.

Myerson, J. & Miezin, F.M. (1980). The kinetics of choice: an operant systems analysis.
Psychological Review
In press.

Nisbett, R.E. (1972). Hunger, obesity and the ventromedial hypothalamus.
Psychological Review, 79 433-453.

Nolan, V. Jr (1978). The ecology and behaviour of the prairie warbler. Dendroica discolor
Ornithological Monographs, 26 596 pages.

Oaten, A. (1977). Optimal forging in patches: a case for stochasticity
Theoretical Population Biology, 12 263-285.

Oster, G. & Heinrich,B. (1976). Why do bumblebees major? A mathematical model.
Ecological Monographs, 46 129-133

Owen, D.F. (1954). The winter weights of titmice.
Ibis, 96 299-309

Palmgren, P. (1949). On the diurnal rhythm of activity and rest in birds.
Ibis, 91 561-576

Panksepp, J. (1974). Hypothalamic regulation of energy balance and feeding behaviour
Federation Proceedings, 33 1150-1165

Parker, G.A. & Stuart, R.A. (1976). Animal behaviour as a strategy optimizer: evolution of resource assessment strategies and optimal emigration thresholds.
American Naturalist, 110 1055-1076

Partridge, B. & Krebs, J.R. (1978). Tree frog choruses: a mixed evolutionarily stable strategy?
Animal Behaviour, 26 959-963

Partridge, L. (1976). Individual differences in feeding efficiencies and feeding preferences of captive great tits.
Animal Behaviour, 24 230-240

Perrins, C. (1979). British tits
London, Collins

Pulliam, H.R. (1975). Diet optimization with nutrient constraints
American Naturalist, 109 765-768

Pulliam, H.R. (1976). The principle of optimal behaviour and the theory of communities.
Perspectives in Ethology, Volume 2 311-332
Bateson, P.P.G. & Klopfer, P.H. (editors)
New York, Plenum Press.

Pyke, G; Pulliam, H.R.; Charnov, E.L. (1977). Optimal foraging: a selective review of theory and tests. Quarterly Review of Biology, 52, 137-154

Rachlin, H. (1978). A molar theory of reinforcement schedules Journal of Experimental Analysis of Behaviour, 30 345-360

Roberts, J.; Kacelnik, A.; Hunter, M. (1979). Sound interference patterns in animal acoustic communication. Technical Memorandum 66
Polytechnic of the South Bank
London

Roberts, J.; Kacelnik, A.; Hunter, M. (1979). Acoustic properties of the environment as determinants of birds' song patterns
Animal Behaviour
In Press

Robinson, M.H. (1966). Anti-predator adaptations in stick-and leaf-mimicking insects.
D.Phil. Thesis. University of Oxford.

Scheer, G. (1952). Beobachtungen und Untersuchungen uber die Abhangigkeit des Fruhgesanges der Vogel von inneren und ausseren Faktoren.
Biologische Abhandlungen, 3-4 1-68

Schoener, T.W. (1971). Theory of feeding strategies
Annual Review of Ecology and Systematics, 2 369-404.

Schwan, A. (1920). Vogelgesang und Wetter, physikalisch-biologisch unter sucht. Vorlanfige Mitteilung.
Pflugers Archiv fur die gesambe Physiologie, 180 341-347.

Schwan, A. (1921-1922). Uber die Abhangigkeit des Vogelesanges von meteorologischen Faktoren, untersucht auf grund physikalischer Methoden.
Verhandlungen der ornitologischen Gesellschaft von Bayern, 15 9-42, 146-186.

Sherry, D.F.; Mrosovsky, N.; Hogan, J.A. (1980). Weight loss and anorexia during incubation in birds.
Journal of Comparative and Physiological Psychology, In Press.

Schmitz, J. & Middel, A. (1966). Die Abhangigkeit des Vogelgesanges von der Helligkeit.
Ornitologische Mitteilungen, 18 111-114.

Tenaza, R.R. (1976). Songs, choruses, and countersinging of kloss' gibbons (Hylobates klossi) in Siberut Islands, Indonesia.
Zeitschrift fur Tierpsychologie, 40 37-52

Shimp, C.P. (1969). Optimum behavior in free-operant experiments.
Psychological Review, 76 97-112.

Sibly, R. (1975). How incentive and deficit determine feeding tendency.
Animal Behaviour, 23 437-446

Sibly, R. & Mc Cleery, R.H. (1976). The dominance boundary method of determining motivational state.
Animal Behaviour, 24 108-124

Sibly, R & McFarland, D.J. (1976). On the fitness of behaviour sequences.
Animal Naturalist, 110 601-617

Silvey, S.D. (1970). Statistical inference.
Chapman and Hall Ltd
London

Slagsvold, T. (1977). Bird song activity in relation to breeding cycle, spring weather, and environmental phenology.
Ornis Scandinavica 8 197-222

Smith, D. (1979). Male singing ability and territory integrity in the redwing blackbird (Agelaius phoeniceus) Behaviour, 68 193-206

Smith, J.N.M. (1974). The food searching behaviour of two European thrushes. I: Description and analysis of search paths. Behaviour, 48 276-302

Smith, J.N.M. (1974) The food searching behaviour of two European thrushes. II: the adaptiveness of the search patterns. Behaviour, 49 1-61.

Smith, J.N.M. & Dawkins, R. (1971). The hunting behaviour of individual great tits in relation to spatial variation in their food density. Animal Behaviour, 19 695-706.

Smith, J.N.M. & Sweatman H.P.A. (1974). Food searching behaviour of titmice in patchy environments. Ecology, 55 1216-1232.

Snedecor, G.W. & Cochran, W.G. (1969). Statistical methods. The Iowa State University Press Ames, Iowa.

Staddon, J.E.R. (1979). Optimality analyses of operant behaviour and their relation to optimal foraging. Adaptation to Constraint Staddon, J.E.R. (editor) New York, Academic Press In Press.

Staddon, J.E.R. (1979). Operant behavior as adaptation to constraints. Journal of Experimental Psychology In press.

Staddon, J.E.R. & Motheral, S. (1978). On matching and maximizing in operant choice experiments. Psychological Review, 85 436-444.

Steen, J. (1958). Climatic adaptation in some small northern birds. Ecology, 39 625-662.

Subbaraj, R. (1977). Rigid internal timing in the circadian rhythm of flight activity in a tropical bat. Oecologia, 29 341-348.

Sulzman, F.M.; Fuller, C.A.; Moore-Ede, M.C.; (1977). Feeding time synchronizes primate circadian rhythms Physiology and Behaviour, 18 775-779.

Tarpy, R. & Mayer R.(1978). Foundations of learning and memory. Scott, Foresman and Co. Glenview Illinois

Thompson, W.R. (1933). On the likelihood that one unknown probability exceeds another in view of the evidence of two samples. Biometrika, 25 285-294.

Toates, F. (1974). Computer simulation and the homeostatic control of behaviour. Motivational Control Systems Analysis D.J. McFarland (editor).

Toates, F. & Booth, D.A. (1974). Control of food intake by energy supply. Nature, 251 710-711

Toates, F. & Oatley, K. (1970). Computer simulation of thirst and water balance. Medical Biological Engineering, 8 71-87.

Wahrenberger, D.L.; Antle, C.E.; Klimko, L.A. (1977).
Bayesian rules for the two-armed bandit.
Biometrika, 64 172-174.

Walsh, J.W.T. (1958). Photometry. Third Edition (revised).
Constable & Co Ltd
London

Warren, R.P. & Vince, M.A. (1963). Taste discrimination
in the great tit (Parus major)
Journal of Comparative and Physiological Psychology,
56, 910-913.

Webber, M.I. (1975). Some aspects of the non-breeding
population dynamics of the great tit (Parus major)
D.Phil. Thesis. University of Oxford.

Werner, E.E. & Hall, D.J. (1974). Optimal foraging and the
size selection of prey by bluegill sunfish. (Lepomis
macrochirus)
Ecology, 55 1216-1232.

Wiley, R.H. & Richards, D.G. (1978). Physical constraints
to acoustic communication in the atmosphere: implications
for the evolution of animal vocalization.
Behavioral Ecology and Sociobiology, 3 69-94.

Wirtshafter, D. & Davis, J.D. (1977). Set points, settling
points and the control of body weight.
Physiology and Behavior, 19 [illegible]-78.

Wynne-Edwards, V.C. (1962). Animal dispersion in relation
to social behaviour
Oliver & Boyd (editors)
Edinburgh

Zimmer, C. (1919). Der Beginn des Vogelgesanges in der
Fruhdammerung.
Verhanlungen der Ornitologischen Geseltschaft
von Bayern.

A P E N D I C E S

THE TWO ARMED-BANDIT ALGORITHM (CONOP)

```
BEGIN
 INTEGER N1, N,PICK,A1,A2,CRUNCH,AA1,AA2;
 REAL ES,B,ED,AL,BE,Y,AL1,AL2;
 AL:=0;
 BE:=0;
  Y:=AL+BE+2;
  WRITETEXT('('-ALPHA,--BETA,--MEAN,VARIANCE')');
  NEWLINE(2);
  PRINT(AL,7,1);
  PRINT(BE,7,1);
  PRINT((AL+1)/Y,7,3);
  PRINT(((AL+1)x(BE+1))/((Y↑2)x(Y+1)),7,4);
  NEWLINE(2);
  FOR N := 50,150,250 DO
  BEGIN REAL ARRAY E[0:(N/2)-1,0:(N/2)-1];
   INTEGER ARRAY X[0: (N/2)-1, 0: (N/2)-1];
   WRITETEXT('(MAXN=')');
   PRINT(N,6,0);
   NEWLINE(1);
   N1 :=(N/2)-1;
    FOR A1 := 0 STEP 1 UNTIL N1 DO
    FOR A2 := 0 STEP 1 UNTIL N1 DO
    E[A1,A2]:= ((N-2xN1)/(N1+Y))x(IF A1>A2 THEN A1+1+AL ELSE A2+1+AL);
    FOR N1 :=(N/2)-2 STEP -1 UNTIL 0 DO
    BEGIN
     FOR A1:= 0 STEP 1 UNTIL N1 DO
     FOR A2:= 0 STEP 1 UNTIL N1 DO
     BEGIN AA1 := A1+1;
      AL1 := AA1+AL;
      AA2 := A2+1;
      AL2 := AA2+AL;
      ED := ((N-2xN1)3(N1+Y))x(IF A1?A2 THEN AL1 ELSE AL2);
      ES := (AL1+AL2)/(N1+Y)+(AL1xAL2xE[AA1,AA2]+AL1x(N1+2-AL2)x
            xE[AA1,A2]+(N1+2-AL1)xAL2xE[A1,AA2]+(N1+2-AL1)x(N1+2-AL2):
            xE[A1,A2])x(N1+Y)↑(-2);
      E[A1,A2] := IF ED GE ES THEN ED ELSE ES;
      X[A1,A2] := ENTIER((ED-ES)X10+0.5);
    END
     PRINT(2xN1,5,0);
     WRITETEXT('('HOPSTOGO:')');
     PRINT(N-N1x2,6,0);
     NEWLINE(1);
     FOR B := 0.00 STEP 0.05 UNTIL 0.5 DO
     BEGIN WRITETEXT('('P:')');
      PRINT(B,1,2);
      WRITETEXT('('=')');
      PRINT(BxN1,2,1);
      IF X[ENTIER(BxN1+0.5),ENTIER((0.50-B0xN1+0.5)]
       GE 0 THEN
       WRITETEXT('('DECIDE')') ELSE
       WRITETEXT('(SAMPLE ONCE MORE')');
     END
    NEWLINE(3);
 END
END
END
```

Appendix 2

REPRINT No. 12 from *Animal Behaviour*, 27, 1, February, 1979

Anim. Behav., 1979, 27, 237–241

THE FORAGING EFFICIENCY OF GREAT TITS (*PARUS MAJOR* L.) IN RELATION TO LIGHT INTENSITY

BY ALEJANDRO KACELNIK

Animal Behaviour Research Group, Department of Zoology, University of Oxford

Abstract. I report an experiment aimed at testing whether foraging efficiency of great tits is limited by light intensity at the time of the dawn chorus. Captive great tits hunting for prey under different luminance conditions were less successful in finding prey when foraging, hunted for a lower proportion of their time, and handled individual prey items for longer when luminance was under approximately 7 cd/m^2. This luminance is not reached in the field until after the time of the dawn chorus, suggesting that in the early morning foraging is limited by light intensity. I suggest that a satisfactory functional explanation of the dawn chorus must take into account the comparatively low foraging opportunity early in the morning, as well as the factors affecting the opportunity for singing and other territorial activities.

The great tit (*Parus major*) in Wytham Wood, Oxford, is territorial from mid January until May (Hinde 1952; Krebs 1971). Song is one of the major forms of territorial defence and advertisement (Krebs 1977), and, in common with many other species, great tits have a major period of song at dawn and a minor peak at dusk (Hinde 1952; Garson & Hunter in preparation) paralleling the bigeminus (large morning and small evening peak) pattern of activity typical of passerine birds (Aschoff 1966). The morning peak in singing activity starts when the sun is between 5° and 3° below the horizon (Kluijver 1950), that is in the middle of the morning twilight period, and lasts about 2 h (Hinde 1952, Table 8). The dawn chorus, as it is usually called, poses a paradox. Great tits sing almost exclusively from perches not less than 5 m above the ground (Hunter, in preparation), while during winter and early spring they forage intensively on nuts and seeds on the ground (Gibb 1954; Lack 1971), so that singing and foraging are mutually exclusive activities. Since they lose between 5 and 10% of their body weight over night in winter (Owen 1954), one might expect feeding motivation to be high in early morning, to ensure that the bird does not lose even more weight. In contrast with this prediction, 'In the early morning, however, the birds (great tits) do not in general seem to feel hungry' (Kluijver 1950), and in fact they allocate a large proportion of time to preening and territorial activity. This apparent paradox can probably be explained by two factors: (1) the advantages of territorial defence just after dawn seem to be particularly high, due to higher intruder pressure (Krebs 1977) and to advantageous acoustic conditions (Henwood & Fabrick in press); (2) the benefits of foraging at this time may be very low, if the rate of prey capture is limited by light intensity.

Foraging tendency is likely to depend both on internal state (food deficit) and incentive (food availability) (Sibly 1975), so that if foraging efficiency is limited by light intensity, foraging tendency may be very low in spite of the great food deficit, and this could 'disinhibit' other behavioural tendencies (McFarland 1974), explaining at least partly the great amount of singing and preening observed at the twilight.

In this paper I describe an experiment designed to test if light intensity does actually lower feeding efficiency at the time of the dawn chorus. I measured foraging efficiency of great tits in the laboratory under different light intensities, covering the range of luminance observed in the field during the first 2 h after twilight, when the dawn chorus takes place.

Methods

Subjects

I used 12 adult great tits. All were wild caught and had spent at least 45 days in captivity. They lived individually in metre cube cages provided with detachable aluminium roosting boxes (8 × 8 × 15 cm) in which the birds roosted and hid when mildly frightened. These cages were kept in a room at approximately 20 C, with light from 08.00 to 19.00 hours. Food was standard diet for insectivorous birds (modified from Krebs et al. 1972) provided ad libitum outside the experimental period, and only after 13.00 hours during this period. All tests took place between 09.00 and 13.00 hours in a room at 10 C.

Each bird was tested only once each morning, and the order was randomized to avoid circadian effects and consistent difference in deprivation between subjects. The mean body weight loss between evening and the beginning of a test was 8% which is within the normal range of body weight loss under natural conditions.

Apparatus

The experimental set-up consisted of two similar aviaries, a waiting compartment and a searching arena, measuring 90 × 90 × 60 cm and separated by a wooden wall. The birds could cross from one to the other through a passage which could be opened or closed by a remotely controlled sliding door. Illumination was the same on both sides, provided by a pair of 150-W reflector lamps on each side. The four lamps were connected in parallel to a single dimmer which was used to vary light intensity between treatments. The waiting compartment had three perches, and the searching arena was filled up with leafless branches and twigs in which the prey were hidden. The branches were rearranged every day so that the birds could not get familiar with the searching environment. When a test was due to start the subject bird was caught in its roosting box which was then attached to the waiting compartment.

The Prey

Fifty artificial prey were distributed in the searching arena before each test. They consisted of a piece of mealworm (mean weight 14 mg) enclosed in a pastry cocoon made of flour and a water solution of food dyes in a proportion of two to one in volume. The water solution was of 8 ml of green, 4 ml of red, and 2 ml of yellow 'Boots food colouring, with water up to 100 ml. The birds handled the prey by pecking open the cocoon and eating the mealworm piece. They did not eat the cocoon, the purpose of which was make the the prey match the background.

Training and Testing

The birds were trained to hunt for prey in the searching arena under intermediate luminance (0·10 cd/m^2) until they reached a constant latency to search and capture rate. Each test consisted of two 15 min blocks. In the first block, the birds were released in the waiting compartment and allowed to accommodate visually to the current level of illumination for 15 min. The passage was then opened and the birds crossed to the searching arena where they hunted for prey for the next 15 min. Searching, handling and non-foraging times were recorded on a computer compatible event recorder. After finding a prey, a bird usually took it to an horizontal branch, where it held the cocoon with one or both feet while hammering with its bill. Once a hole was made, it cleanly pulled out the piece of mealworm and swallowed it. After this, birds sometimes discarded the cocoon while others continued with it, searching for remaining bits of food. Handling time is defined as the period between the first contact with a prey until discarding the empty cocoon.

I will discuss this point further when analysing the effect of light on handling strategy.

Each bird was subject to ten different treatments, each one consisting of a test at a given luminance. Figure 1 shows the actual values of luminance used in the ten treatments. Since both prey and background are secondary luminous sources of low reflection factor, I included in the same figure luminance as measured on a white matt sheet of paper in the same situation

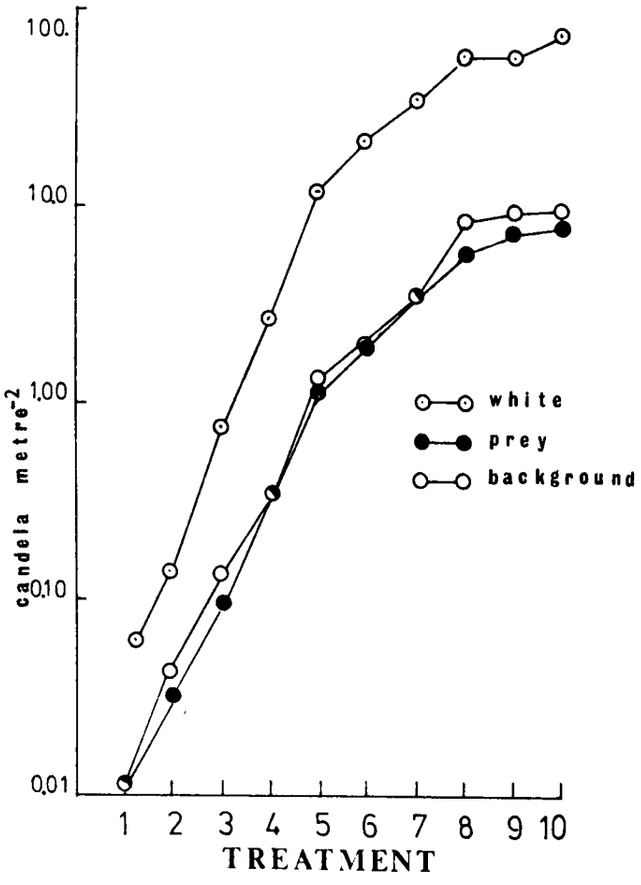


Fig. 1. Luminance at prey, background and matt white paper surfaces for the 10 different treatments. The values for the white surface are shown to facilitate comparison with other experimental situations.

b

to facilitate comparisons with different experimental situations. The experimental range was from 0·01 to 10·0 cd/m^2, which covers the full range of luminance as measured on the same prey in the natural feeding grounds during the first 2 h after twilight (see Fig. 5). I assume that if the prey has the same luminance in the field and laboratory situations, and the two backgrounds are similar, the subjective brightness will be the same and so will be its detectability (Walsh 1958).

All luminance measurements were made with a S.E.I. photometer.

Results

Figure 2 shows the effect of luminance on two measures of foraging success. The first measure, profitability, is defined as prey caught per second foraging (searching + handling). The second measure, searching efficiency, is defined as prey captured per second of searching time. While profitability increases up to a maximum at a luminance of 1 cd/m^2, searching efficiency increased with luminance throughout the whole experimental range. This second measure rises more steeply with light because it ignores handling time: at high light intensity the birds caught more prey and therefore spent more time handling (Holling 1965), and this sets an upper limit to profitability. This effect is complicated by the fact that handling time is inversely related to light intensity (Fig. 3), which reduces apparent profitability at low light intensity. The reason for the decrease in handling time with increasing light is that the birds extracted food from the cocoons more thoroughly when light intensity was low, spending additional time gleaning the open cocoon for extra pieces of mealworm. Most of the cocoons recovered after experiments with high light intensity had a clean hole through which the piece of mealworm had been extracted, while cocoons recovered after low light tests were fully opened as a conse-

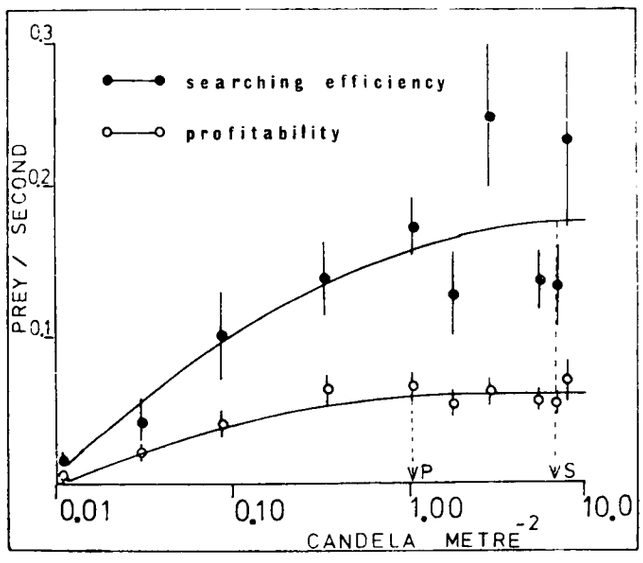


Fig. 2. Profitability and searching efficiency for different light intensities (for definitions, see text). The curves are second order polynomials fitted to the original data, but only means and standard errors are plotted for clarity of the figure. In both cases, second order polynomials with negative quadratic terms were better fits ($P < 0·05$) than linear ones indicating that values were reaching a plateau. P and S indicate the luminance values at which profitability and searching efficiency reach 95% of their maximum value.

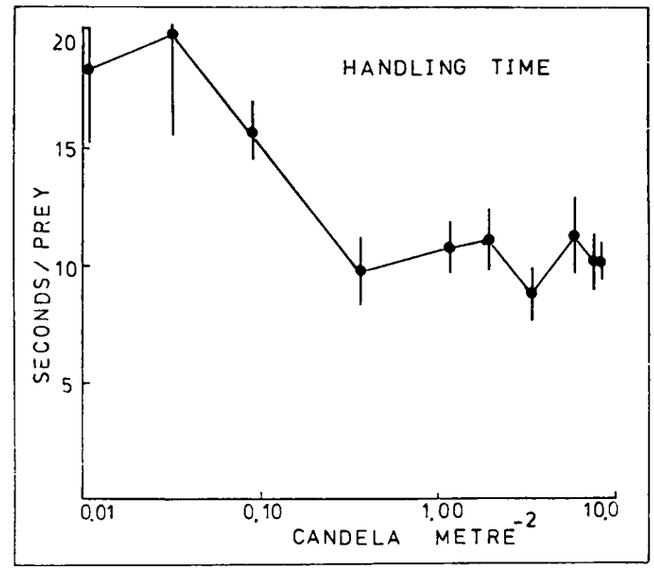


Fig. 3. Handling time decreased with increasing luminance. Individual prey items were more thoroughly exploited when capture rate was low. Mean ± SE.

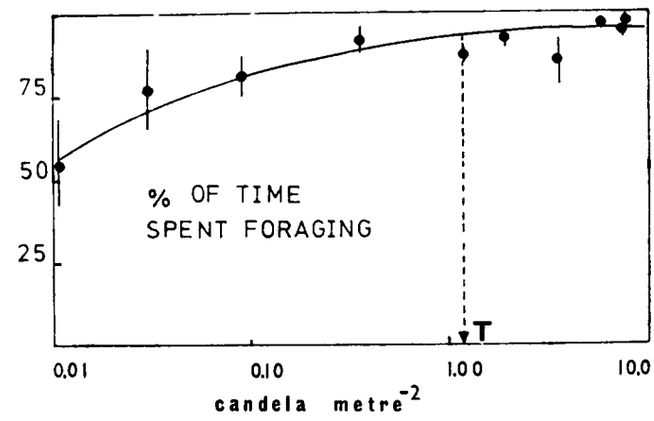


Fig. 4. The fraction of time spent foraging increased with luminance up to 100%. The curve is a second order polynomial fitted to the original data, but means ± SE are plotted for clarity. T shows the luminance at which the fitted polynomial reaches 95% of its maximum value.

quence of this extra attention. Table I shows that the birds actually found more food when they had fully opened the cocoons. It is interesting to note that the decrease in handling time with decrease in inter-catch interval is what would be expected from optimal foraging theory (Cook & Cockrell 1978).

Figure 4 shows that the proportion of experimental time spent foraging increased with light intensity up to an asymptote of 100%. At low light intensity the birds spent as much as 50% of their time preening, resting or bathing.

In order to relate the laboratory results to the levels of light intensity in the field, I combined the laboratory results with measures of luminance at the surface of the artificial prey in the field (Fig. 5).

From Fig. 5 I conclude that:

(a) Searching efficiency is light limited for over 1½ h after the onset of great tits' activity on sunny days and more than 2 h on overcast days, which are more typical in Oxfordshire.

(b) Profitability for birds hunting for prey requiring a handling time similar to the experimental prey would be limited by light for 40 min on sunny days and over 1 h on overcast days.

(c) The proportion of time allocated to foraging matched profitability under the experimental conditions, i.e. the birds foraged nearly all the time when luminance was above 1 cd/m^2, the luminance at which profitability reached its asymptotic value.

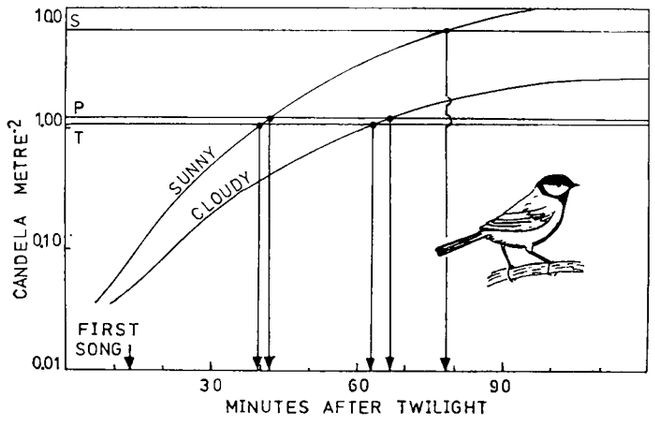


Fig. 5. When luminance at great tits' feeding places is plotted against time of day for sunny and overcast weather, the results from Figs. 2 and 4 can be used to estimate at what times light limits foraging efficiency in the wild. S, P and T are from Figs. 2 and 4. The time of the first song is approximate from references in the text. Asymptotic searching efficiency (s) is reached between 80 min and more than 2 h after the first song, depending on the weather. Asymptotic profitability (P) and full dedication to foraging (T) were reached at similar luminance, corresponding to between 30 and 60 min after the first song.

Discussion

My results show that part of the explanation of the dawn chorus in the great tit may be that light intensity limits foraging efficiency in the early morning. Birds in the wild may, to some extent, compensate for the effect of light intensity by exploiting each item more thoroughly, as they did in my experiments, and by switching to normally less preferred but more conspicuous sorts of prey. Nevertheless, in both cases, the overall pay-off for foraging would be lower than at other times of the day.

The poor profitability of foraging at dawn and dusk would itself favour the concentration of other incompatible activities at these times, provided that the birds expect to have enough profitable foraging time during the rest of the day. Singing and preening are two examples of such activities: they are partly incompatible with foraging because great tits in late winter sing and preen in the canopy (Garson & Hunter, in preparation) while they feed mainly on nuts which they find on the ground (Gibb 1954). The evening peak of song may be smaller than the morning one because of the rate of change of light intensity. While early in the morning profitability is improving, at dusk the bird has to gather enough food to survive overnight even if profitability is low.

In addition to poor profitability there is evidence that there are special advantages for territorial activity at the time of the dawn chorus. In the first place, activity of potential invaders is higher early in the morning than at any other time, making territorial advertisement by song more important as a deterrent at that time. This has been shown by J. Krebs (in preparation, quoted in Krebs & Perrins 1977) who removed territory holders from a piece of mixed wood and replaced some of them by broadcasting their

Table I. Sixty Fully Open and 60 Perforated Cocoons Were Classified According to the Presence or Absence of Traces of Food. More Fully Open Cocoons Were Totally Empty ($\chi^2 = 5{\cdot}40$, $P < 0{\cdot}05$)

	Some food left	No food left
Fully open	10	50
Perforated	21	39

own songs, recording the rate of arrival of new birds at different times. The rate of arrival was higher during the first 4 h of the morning, regardless of whether song was being broadcast or not. This result may be a consequence of the fact that it pays non-territorial birds to seek territories at the time of low foraging efficiency.

Secondly, there seems to be an intrinsic advantage for acoustic communication at dawn. Sounds in the frequency range of the great tit song are likely to be less attenuated (i.e. reach greater distances) when produced around sunrise than later on (Henwood & Fabrick in press; Roberts et al. in preparation).

These are not, of course, alternative explanations but rather different elements of the intricate set of factors that must be governing the temporal organization of great tits' behaviour.

It is possible that the relative importance of the different factors will ultimately be resolved by further quantification of the risks and advantages of performing different activities at certain times of the day and by analysing how these risks and advantages are reflected in the animal's motivational system (Sibly & McFarland 1976). It may develop instead, that further quantification may prove insufficient until a more workable theory of the functional determinants of temporal organization of behaviour is available.

Acknowledgments

I am very grateful to my supervisor, Dr J. Krebs, for his interest and help, and to Dr D. McFarland, Dr R. Sibly and M. Hunter, Jr, for many useful discussions and ideas.

REFERENCES

Aschoff, J. 1966. Circadian activity pattern with two peaks. *Ecology*, **47**, 657–661.

Cook, R. M. & Cockrell, B. 1978. Predator ingestion rate and its bearing on feeding time and the theory of optimal diets. *J. anim. Ecol.*, **47**, 451–469.

Garson, P. J. & Hunter, M. L. In preparation. The singing behaviour of wrens (*Troglodytes troglodytes* L.) and great tits (*Parus major* L.): effects of temperature and time of year.

Gibb, J. 1954. Feeding ecology of tits, with notes on treecreeper and gold-crest. *Ibis*, **96**, 513–543.

Henwood, K. & Fabrick, A. In press. A quantitative analysis of the dawn chorus temporal selection for communicatory optimization. *Am. Nat.*

Hinde, R. 1952. The behaviour of the great tit. *Behaviour Suppl.*, **2**, 1952.

Holling, C. S. 1965. The functional response of predators to prey density and its role in mimicry and population regulation. *Mem. ent. Soc. Can.*, **45**, 1–60.

Kluijver, H. N. 1950. Daily routines of the great tit. *Ardea*, **38**, 99–135.

Krebs, J. 1971. Territory and breeding density in the great tit. *Ecology*, **52**, 2–22.

Krebs, J. 1977. Song and territory in the great tit *Parus major*. In: *Evolutionary Ecology* (Ed. by B. Stonehouse & C. Perrins), pp. 47–62. London: Macmillan.

Krebs, J., MacRoberts, M. H. & Cullen, J. M. 1972. Flocking and feeding in the great tit–an experimental study. *Ibis*, **114**, 507–530.

Krebs, J. & Perrins, C. 1977. Behaviour and population regulation in the great tit (*Parus major*). In: *Population Control by Social Behaviour*. Symposium of the Institute of Biology, London, 20–21 September 1977.

Lack, D. 1971. *Ecological Isolation in Birds*. Oxford and Edinburgh: Blackwell Scientific Editions.

McFarland, D. J. 1974. Time sharing as a behavioral phenomenon. In: *Advances in the Study of Behavior* (Ed. by D. Lehrman, J. Rosenblatt, R. Hinde and E. Shaw), Vol. 5. London and New York: Academic Press.

Owen, D. F. 1954. The winter weights of titmice. *Ibis*, **96**, 299–309.

Roberts, J., Kacelnik A. & Hunter, M. L. In preparation. A model of sound interference for acoustic communication.

Sibly, R. 1975. How incentive and deficit determine feeding tendency. *Anim. Behav.*, **23**, 437–446.

Sibly, R. M. & McFarland, D. 1976. On the fitness of behavior sequences. *Am. Nat.*, **110**, 601–617.

Walsh, J. W. T. 1958. *Photometry*. Third edition (revised). London: Constable.

(*Received* 4 *April* 1978; *revised* 25 *May* 1978; *MS. number:* 1731)

Polytechnic of the South Bank

Borough Road
London SE1 OAA

Institute of Environmental Science and Technology

TECHNICAL MEMORANDUM 66

SOUND INTERFERENCE PATTERNS IN ANIMAL ACOUSTIC COMMUNICATION.

JOHN ROBERTS, ALEJANDRO KACELNIK*
and MALCOLM J. HUNTER. Jr.**

*Animal Behaviour Research Group,
Department of Zoology, Oxford University.

**Formerly of the Department of Zoology Oxford University,
now at School of Forest Resources,
University of Maine at Orono, U.S.A.

MAY 1979.

1.0 SUMMARY

We describe a simple model which calculates the sound pressure level (SPL) some distance away from an acoustic source on the basis of atmospheric attenuation and interference between the direct sound wave and the sound wave reflected by the ground. SPL as measured at points in the sound field where interference is constructive or destructive is respectively above or below that which would be expected from spherical spread of the sound energy. If SPLs at two locations within the sound field are subtracted, the value obtained is a function of the exact location of the two receiving devices as well as of their separation. We claim that the 'sound window' or zone of less attenuated frequencies described by previous authors is a consequence of the methodology employed and has little relevance for the evolution of animal vocalisations.

The patterns of peaks and troughs in the SPL created by interference becomes more pronounced as the height above the ground increases, so that selection against vocalising in pure tones can be expected to be stronger for animals vocalising higher up than for animals vocalising near the ground. The frequency span of acoustic signals is, according to this, expected to be correlated with average height of vocalising stations.

2.0 INTRODUCTION.

Several empirical studies have recently explored the possibility that animal vocalisations which are designed to maximise transmission distances are shaped by acoustical properties of the environment (Morton 1970, 1975; Chappuis 1971; Jilka and Leisler 1974; Linskens et al. 1976; Marten and Marler 1977; Marten et al. 1977; Richard and Wiley 1978.). Morton (1975) and Marten and Marler (1977) found that so-called 'excess attenuation' was particularly low for certain frequencies and referred to these frequencies as the 'sound window'. They considered this phenomenon as an important factor in the evolution of animal acoustic signals. In this paper we present a model which attempts to explain its physical basis.

Our simple model takes into account both atmospheric attenuation and interference between the direct sound wave and the wave reflected from the ground. It has five variables: the frequency of the sound, the distance from sender to receiver, their height above the ground, the acoustic impedance of the ground and the sound absorption by the air as determined by temperature and relative humidity. In this report we describe our model, test its plausibility by comparing its predictions with data sets obtained by various workers, and further explore its implications for animal vocalisation.

3.0 THE MODEL.

The signaller is assumed to be an omnidirectional source vocalising in pure tones above a flat, horizontal, unobstructed, locally reacting plane. In practice these limitations are of little significance provided the ground is reasonably flat and there is little or no wind. The presence of trees and the fact that animals are not omnidirectional sources should be important only at the higher frequencies.

From the point of view of the receiver there are two sound waves from any one signaller. One is the direct sound wave from the source; the other is the wave reflected from the ground, that appears (on plane wave theory) to come from an image source the same distance below the ground as the source is above (see Figure 1), its sound power and phase depending upon the acoustic impedance of the ground.

If the signaller is at height HS, the observer at height HO and the two are a distance D apart horizontally then the sound pressure received will be of the form:-

$$P = \text{constant}[(\exp(jkD))/D + RP(\exp(jkD2))/D2] \quad \ldots\ldots(1)$$

where j is the square root of minus one, k is the wave number of the acoustic signal,

$$D2 = SQRT(D*D + (HS+HO)*(HS+HO))$$

$$RP = (COSPHI - BETA)/(COSPHI + BETA).$$

BETA is a function of the acoustic impedance of the ground and is determined by measurement, and PHI is the angle the direct sound makes with the vertical (see Figure 1).

The above model is generally considered satisfactory if both source and receiver are at least half a wavelength above the ground. Thus, for example, if both source and receiver were 0.15 m. above the ground then equation (1) would have a lower limiting frequency of about 1.5 kHz. For frequencies lower than this a fuller analysis is required.

The error in the above model becomes obvious if both source and receiver are on the ground. Then HS = HO = 0, COSPHI = 0, and RP = -1, cancellation occurs between the direct and reflected waves and P = 0. This is not what happens in practice and the plane wave approximation must be modified to allow for the spherical nature of the real wave.

The difference in sound levels predicted by the plane wave model and those actually heard may be regarded as due to an extra wave;-

$$P = \text{constant}[(\exp(jkD))/D + (RP + (1 - RP)F)(\exp(jkD2))/D2] \quad \ldots(2)$$

The term (1 - RP)F represents the divergence from plane wave conditions. We may thus say immediately that F will tend to zero at high frequencies, large source-receiver separations and large values of the ground impedance (BETA<< 1).

Several authors, for example Chien and Soroka (1975) have expressed F in terms of the other system parameters. We have used their expression for F with values of BETA we have measured for grassland, using a Bruel and Kjaer type 4002 standing wave apparatus, to construct a relatively simple computer program (Appendix 1) with which to predict likely pressure levels in the sound fields of various sources.

Inherent in such a model is the assumption that the value of BETA measured for normal incidence in the standing wave apparatus is valid for all angles of incidence - that is grassland is assumed locally reacting. Bolton (1978) has shown that such an assumption is acceptable save for values of the angle PHI between about 40 degrees to 70 degrees.

The measured value of BETA was obtained using 100 mm. deep samples taken early in the morning (approx. 7.30 am.) after a night without rain and the measurements taken within two hours. The values thus obtained are· (BETA =A - jB)

$$A = 0.00525 + 0.134*f - 0.024*f*f + 0.00129*f*f*f$$

$$B = -0.00129 + 0.125*f + 0.0224*f*f - 0.00121*f*f*f$$

Where f is in kHz. These results are accurate to within about 5%.

The model and the program were tested for sensitivity to BETA by varying the value of BETA in the equations over the range one tenth of the measured value to ten times the measured value. This range of 100:1 in BETA produced no great changes in the predicted values of the sound pressure levels at positions of constructive and destructive interference but did, of course, move the frequencies at which these peaks and troughs occurred. The larger the value of BETA the lower the frequency at which the first maximum and minimum occurred, though the frequency difference between peaks and troughs remained unaltered.

However, as we are concerned with general phenomena the value of BETA obtained for grassland will be used for predicting sound pressure levels above such surfaces as sandy and loamy soil and the results interpreted in the light of the given specific conditions.

The present model does not allow for such effects as attenuation due to wind or turbulence in the atmosphere which may be doubly important at long source-receiver separations as they will not only attenuate the sound but also cause fluctuations in the relative phase of the direct and reflected waves. This latter effect would cause large fluctuations of the SPL in those areas where the interference is destructive and critically dependent upon the phase difference between the two waves. Close to the ground where wind effects are likely to be small the model should be in relatively close agreement with measured data, but when source-receiver heights are such that wind effects will be noticable the model will be most useful in predicting general trends rather than precise values.

4.Ø EXCESS ATTENUATION.

Originally our aim was to compare results predicted by our model with empirical data obtained by Marten and Marler (1977), and hence it is necessary to describe briefly their techniques. Marten and Marler broadcast pure tones from a loadspeaker at heights of Ø.15, 1.Ø, 2.Ø, 5.Ø and 1Ø.Ø m. SPLs were measured at the same heights with microphones 2.5 and 1Ø2.5 m. from the source. They used 25 frequencies between Ø.3 and 11.Ø kHz. in '.... a large hayfield with standing dead and growing grasses to 2Ø cm.'. Marten and Marler subtracted the SPLs measured at 2.5 and 1Ø2.5 m. to determine 'excess attenuation' which they defined as the difference between the measured attenuation and the attenuation expected due to spherical expansion of the sound wave. With equation (2) we calculated the expected SPLs at 2.5 m (SPL1) and at 1Ø2.5 m. (SPL2) for a sound source of arbitrary power. Since we are concerned only with difference in SPLs the sound power of the source is unimportant.

Excess attenuation (E.A) was determined from

$$E.A. = (SPL1 - SPL2) - 32.3 - c$$

where 32.3 dB results from spherical spread (2ØLog1Ø2.5/2.5) and c is the attenuation due to absorption of sound energy by the air (Bazley 1976). Marten and Marler did not publish meteorological data and we have assumed mean values of temperature and humidity for central New York State in July, 24C and 65% respectively (Meteorological Office 1958).

The model was also used to predict actual sound pressure levels in the sound fields of various sources for a series of source-receiver separations and heights at ØC and 9Ø% relative humidity.

Figure 2 shows predicted values of the SPL over grassland for various heights and source-receiver separations. The effect of frequency of the signal on SPL is strongly dependent upon both height and separation. For a given separation (see for example, 1ØØ m. in Figure 2) comparison between heights shows that SPL appears as a smooth curve close to the ground, while a series of peaks and troughs occurs at regular frequency intervals for greater heights. These fluctuations are a consequence of the interference between direct and reflected sound waves (Figure 1), and their general effect is to reduce stability in the reception of any pure tone signal when source and receiver are some metres above the ground. High frequencies are more attenuated than lower ones for all heights, but this is an appreciable effect only at very high frequencies or large source-receiver separations. For example, at a height of 2 m. and a source-receiver separation of 5Ø m.

the SPL due to a source broadcasting at about 2 kHz is less than the SPL due to a source of 5 kHz., but at 100 m. the former source will produce a SPL some 10 dB higher than the latter. Such effects can be discerned at most heights.

Close to the ground all signals of frequency greater than a few hundred Hz are more attenuated than at greater heights. This effect, which can be seen by comparing heights in Figure 2, is non-linear, being very strong below 1 m. but almost negligible for further increases in height. 'Excess attenuation', defined in Methods, is equivalent to the difference between two curves in any of the frames of Figure 2. It can be appreciated that its value is strongly affected by the particular frequency of the signal and the microphones' positions. If, for example, SPL due to a source of about 3 kHz is measured at 100 m. and 200 m. from the source, at 2 m. above the ground, we may find that the SPL is higher at 200 m. than at 100 m., due to the interference effect (Figure 2). If SPL were measured for the same separation between microphones but at different distances from the source, this apparently paradoxical effect would not be observed.

Figure 3 compares our predictions for 'excess attenuation' between points 2.5 m. and 102.5 m. from the source with Marten and Marler's data. On the left are the predicted sound pressure levels in dB re an arbitrary datum at the two microphone positions, and on the right the 'excess attenuation' calculated from the curves is compared with the measured values. Not only do the calculated and measured values follow the same trends but quantitative agreement is also quite close, especially at the higher source-receiver heights. Generally the differences may be attributable to three factors; (i)the presence of wind, (ii)errors in the assumed values of air absorption due to lack of precise data on relative humidity and temperature, and (iii)errors in the assumed value of BETA.

It should be noted that when source and microphone heights are 1 m. and the nearer microphone 2.5 m. from the source the angle PHI will be about 51 degrees and hence in the range for which the assumption of locally reacting grassland is least valid. This is a possible explanation for the relatively poor agreement between measured and predicted values of 'excess attenuation' for this given source-receiver configuration.

'Excess attenuation' as predicted by our model and measured in the field is very strong close to the ground and decreases with height, but above 2 m. the effect of a further increase in height is negligible. An important and counter-intuitive prediction of the model that also agrees with the field data is that 'excess attenuation' does not increase monotonically with frequency; it decreases up to a certain frequency, which is dependent upon source-receiver

height and then increases (Figure 3, 1.Ø and 2.Ø m.). The presence of a minimum value for 'excess attenuation' was first shown empirically by Morton (1975): he called it a 'sound window' and ascribed importance to its presence or absence in different habitats as a factor shaping the songs of different avifaunas. Marten and Marler found a sound window only at ground level and at 1.Ø m. above the ground. Our model instead predicts that these zones of minimum 'excess attenuation' should occur whenever interference between the direct and reflected waves is constructive and air attenuation is not large.

It has also been possible to compare predicted values of 'excess attenuation' with those measured by other researchers. Aylor (1975) measured 'excess attenuation' over weather-slaked sandy loam before and after disking. Source and both receivers were at a height of 1.Ø m. and the receiving microphones were 3 m. and 52 m. from the source. Figure 4 shows Aylor's measured 'excess attenuations' and the SPLs predicted at the two measuring points.

It can be clearly seen that the peak values in the 'excess attenuation' is a result of destructive interference occurring at 52 m. and constructive interference occurring at the nearer measuring position. The precise magnitude of the peak values and the frequency at which they occur will depend upon the actual acoustic impedance of the ground. Here, as expected, the rougher, more porous ground obtained by disking moved the interference pattern to lower frequencies. In fact the change in magnitude and frequency of the peak 'excess attenuation' obtained by disking may be accounted for if the disking increased BETA by a factor of about 3.

Morton (1975) measured 'excess attenuation' through long grasses for source-receiver heights of 1 ft (Ø.3 m.) and microphone receivers 7.62 m. and 91.44 m. from the source. If the various attenuation processes are assumed simply additive it is possible to make a rough approximation for the predicted 'excess attenuation' through the grass using our model together with data from Marten and Marler (1977) on the absorption of acoustic energy by, for example, coniferous forest. Because the data of Marten and Marler is likely to underestimate the actual absorption by the grass the predicted approximation should be lower than the 'excess attenuation' measured by Morton, see Figure 5.

5.Ø SOUND FIELDS.

Comparisons are also possible between model predictions and direct measurement of sound pressure levels for various source-receiver configurations. Embleton et al (1976)

measured SPL over grassland for a source-receiver separation of 15.24 m. and source-receiver height of 1.2 m. Comparison between the predicted and measured values are shown in Figure 6. The general shapes of the curves agree quite closely, though there is some discrepancy in the precise frequency at which the minima occur - probably due to the angle PHI being approximately 80 degrees and thus lying near the range for which assumption of locally reacting ground is least valid. The spikey nature of the practical measurements probably results from such factors as non-uniform ground, ground which is not flat, wind effects, etc.

Embleton et al (1975) also measured the SPL for source and receiver close to the ground. Their results are shown together with the predicted values in Figure 7 for a source-receiver separation of 15.24 m., source height 30 mm., and receiver heights of 30 mm., 305 mm., 610 mm., and 1220 mm. Qualitative agreement appears to be very good, confirming that low frequencies are transmitted better the closer the source and receiver are to the ground. The theory predicts and the measurements confirm that for acoustic waves transmitted in close proximity to the ground the presence of the ground acts as a low pass filter with a very sharp cut-off. Thus there is a great advantage when both source and receiver are close to the ground, to signal at as low a frequency as possible.

Jilka and Leisler (1974) compared the measured SPLs in the sound fields of a loudspeaker with the preferred frequencies and frequency ranges of three reed warbler species in their respective habitats. their measured SPLs and those predicted by our model are shown together in Figure 8 for a source-receiver separation of 16 m. and source and receiver at a height of 0.2 m. The decrease in measured SPL at higher frequencies below that predicted is ascribed to the absortion of acoustic energy due to the presence of reeds and other vegitation.

6.0 DISCUSSION.

'Excess attenuation' and consequent 'sound windows' as predicted by our model are the result of subtracting SPLs at two points in the sound field and depend on both the relative microphone positions and the actual attenuation between source and receiver. Thus they are mainly methodological artefacts and have little relevance for animal vocalisations.

We take the good agreement of our model with the reported field data as a provisional validation of the model's structure and we now proceed to explore its implications for animal vocalisations.

A specific pattern of peaks and troughs of SPLs exists for each source-receiver distance and height; at greater heights and shorter source-receiver distances the pattern is more pronounced (Figure 2). Animals are unlikely to be able to maximise the carrying power of their vocalisation by using solely frequencies which coincide with SPL maxima because they will know neither the exact distance to their receiver nor its height. However, this pattern of SPLs does have potential relevance for animal communication because it becomes more pronounced as height increases. This implies that for animals communicating through a narrow band channel (i.e. using pure tone signals) their signal would fluctuate more as their height above the ground increased.

The general features of a signal which avoids the receiver being in a position of destructive interference and minimum SPL may be determind from Figure 2. Close to the ground it is important to broadcast at as low a frequency as possible and even a pure tone will transmit well if of sufficiently low frequency. As the height of source and receiver above the ground increases to, say 2 m., it can be seen that a pure tone could give destructive interference at the receiver. Thus at this height to ensure good reception the bandwidth of the signal should be sufficiently broad as to cover at least adjacent maxima and minima in the interference pattern. As the height of the source and receiver increases further the interference peaks and troughs come closer together until the frequency bandwidth of the signal may be quite narrow and still span several maxima of the pattern.

Such an analyisis leads to the prediction that, all other things being equal;-

(1)Animals vocalising at some height above the ground would find it advantageous to use signals of wider frequency span than animals vocalising on the ground. This prediction was tested comparing the vocalisation frequency span of the bird species listed by Morton (197Ø; pers. comm.) whose singing perch heights were clearly defined. For 14 species singing between 1 and 5 m. above the ground the average span was 1.96 kHz (SD = 1.11 khz) while for 14 species singing on the ground the average span was Ø.78 kHz (SD = Ø.77 kHz). This difference in average spans was statistically significant at the 1% level ('t' test).

For 1Ø canopy species the mean frequency span was 2.56 kHz (SD = 2.57 kHz). This span was significantly different from that of the ground species at the 2.5% level ('t' test) but, as would be expected, was not significantly different from that of the species singing between 1 and 5 m.

(2)The main vocalising frequency for birds on the ground should be lower than for those vocalising from some height above the ground. This prediction was again tested using the data of Morton. For the 14 species singing on the

ground the average main frequency was 1.83 kHz (SD = 1.27 kHz), while for the 14 species singing from perch heights between 1 and 5 m. the average main frequency was 2.96 kHz (SD = 1.17 kHz). This difference is significant ('t' test) at the 2.5% level. However, for the 10 canopy species the average main frequency was 2.40 kHz (SD = 1.89 kHz) which is not significantly different from that of the ground species or from that of the species singing at the intermediate heights.

There are, of course, factors quite independent of the need for better transmission that effect the frequency composition of animal vocalisations. For example, the lowest frequency that an animal can generate depends upon its size. The larger the animal the lower the pitch generated. In consequence, from size alone one should expect larger animals to have a wider frequency span and a lower mean frequency. Since body size does have an effect on the height at which animals operate within, say, a wood, it is important to bear in mind the possibility that certain trends regarding mean frequency of vocalisation, etc., may be merely side effects of non-acoustic constraints.

In his analysis of the height separation within the genus Parus, David Lack (1971) found that larger species within the same genus tend to spend a higher proportion of their time on the ground than smaller ones. This type of trend effects the predictive value of any model (see Dawkins 1974). The more the departure of the model's predictions from a general trend expected from independent factors the higher the predictive information of that model.

In our case, if smaller size is related to higher stations in the canopy, the agreement between available data and predictions about increased frequency range at greater heights becomes stronger because it goes against the general trend expected independently. However, the fit between data and prediction of lower frequencies close to the ground is less impressive since one would expect such a trend even if birds chose their main vocalising frequency independently of optimising information transmission over the greatest distance.

7.0 ACKNOWLEDGEMENTS.

John Krebs and Henry Bennet-Clark read a previous version of this report and made helpful comments on it.

David Hardy for his 'patience' and help in explaining his many and varied program packages.

8.Ø FIGURE CAPTIONS

Figure 1; Diagram of the sound paths considered in the model. Sound Pressure Levels (SPL1 and SPL2) are the result of interference between the direct and reflected waves.

Figure 2; SPLs as predicted by the model for various source-receiver heights and separations. Each frame corresponds to one height and three distances (D = 5Ø, 1ØØ, and 2ØØ m.). the five source heights are (a) Ø.15 m., (b) 1.Ø m., (c) 2.Ø m., (d) 5.Ø m. and (e) 1Ø.Ø m. In each case HO = HS.

Figure 3; 'Excess attenuation' calculated from equation (4) and measured by Marten and Marler (1977) over grassland. Source-receiver heights are equal in each case with (a) Ø.15 m., (b) 1.Ø m., (c) 2.Ø m., (d) 5.Ø m. and (e) 1Ø.Ø m. The source-receiver separations are 2.5 m and 1Ø2.5 m.

Figure 4; 'Excess attenuation' calculated from equation (4) and measured by Aylor (1975) over weather-slaked sandy loam before and after disking. Source-receiver height 1.Ø m. and the source-receiver separations 3 m. and 52 m.

Figure 5; 'Excess attenuation' calculated from equation (4) modified to allow for absorption by vegitation, and measured by Morton (1977). Source-receiver height Ø.3 m. and source-receiver separation 7.62 and 91.44 m.

Figure 6; SPLs measured by Embleton et al (1976) ,and calculated from equation (2) over grassland for a source-receiver separation of 15.24 m., and source-receiver height 1.2 m.

Figure 7; SPLs measured by Embleton et al (1976) over grassland for a source-receiver separation of 15.24 m., a source height of 3Ø.5 mm., and receiver heights (1) 3Ø.5 mm., (2) 3Ø5 mm., (3) 61Ø mm. and (4) 122Ø mm., and calculated from equation (2).

Figure 8; SPLs measured by Jilka and Leisler (1974) for source-receiver separation 16 m. and source-receiver height Ø.2 m., and calculated from equation (2).

9.0 REFERENCES.

Aylor, E. 1976. Noise reduction by vegetation and ground. JASA 51 197-205

Bazley, E. 1976. Sound absorption in air at frequencies up to 100 kHz. National Physical Laboratory, Acoustics Report Ac 74, Feb. 1976.

Bolton, J. 1978. On the use of local reaction boundary conditions in theories of sound propagation. Spring Meeting Institute of Acoustics

Chappuis, C. 1971. Un example de l'influence du milieu sur les emissions vocales des oiseaux: l'evolution des chants en foret equatoriale. Terre Vie 25, 183-202.

Chien, C. and Soroka, W. 1975. Sound propagation along an impedance plane. JSV 43 9-19

Dawkins, M. and Dawkins, R. 1974. Some descriptive and explanatory stochastic models of decision making. In Motivational Control Systems Analysis., Ed. McFarland, D. J., A.P. New York.

Henwood, K. and Fabrick, A. 1978. A quantitative analysis of the dawn chorus: temporal selection for communicatory optimisation. American Naturalist (in press).

Jilka, A and Leisler, B. 1978. Die Einpassung dreier Rohrsangeraiten (Acrocephalus schoenobaenus, A. scriapaes, A. arundinaceas) in ihre Lebensraume in bezug auf das Frequenzpektrum ihrer Reviersange. J. Ornithol. 115. 192-212.

Lack, d. 1971. Ecological isolation in birds. Blackwell's Scientific Publications, Oxford.

Linskens, H et al. 1976. The acoustic climate of plant communities. Ecologia. 23 165-177.

Marten, K. and Marler, P. 1977. Sound transmission and its significance for animal vocalisation. I. Temperate habitats. Behav. Ecol. Sociobiol. 2, 271-290.

Marten, K., Quine, D. and Marler, P. 1977. Sound transmission and its significance for animal vocalisation. II. Tropical forest habitats. Behav. Ecol. Sociobiol. 2, 291-302.

Meteorological Office 1958. Tables of temperature, relative humidity and precipitation for the world. Part I, North America, Greenland and the North Pacific Ocean. H.M.S.O. 84pp.

Morton, E. 1970. Ecological sources of selection of avian sounds. Ph. D. Thesis, Yale University.

Morton, E. 1975. Ecological sources of selection of avian sounds. Amer. Natur. 109, 17-34.

Richards, D. and Wiley, R. 1978. Reverberation and amplitude fluctuations in the propagation of sound in a forest: implications for animal communication. Behav. Ecol. Sociobiol. 3, 69-94.

Waser, P. and Waser, M. 1977. Experimental studies of primate vocalisation: specialisations for long distance propagation. Z. Tierpsychol. 43, 239-263.

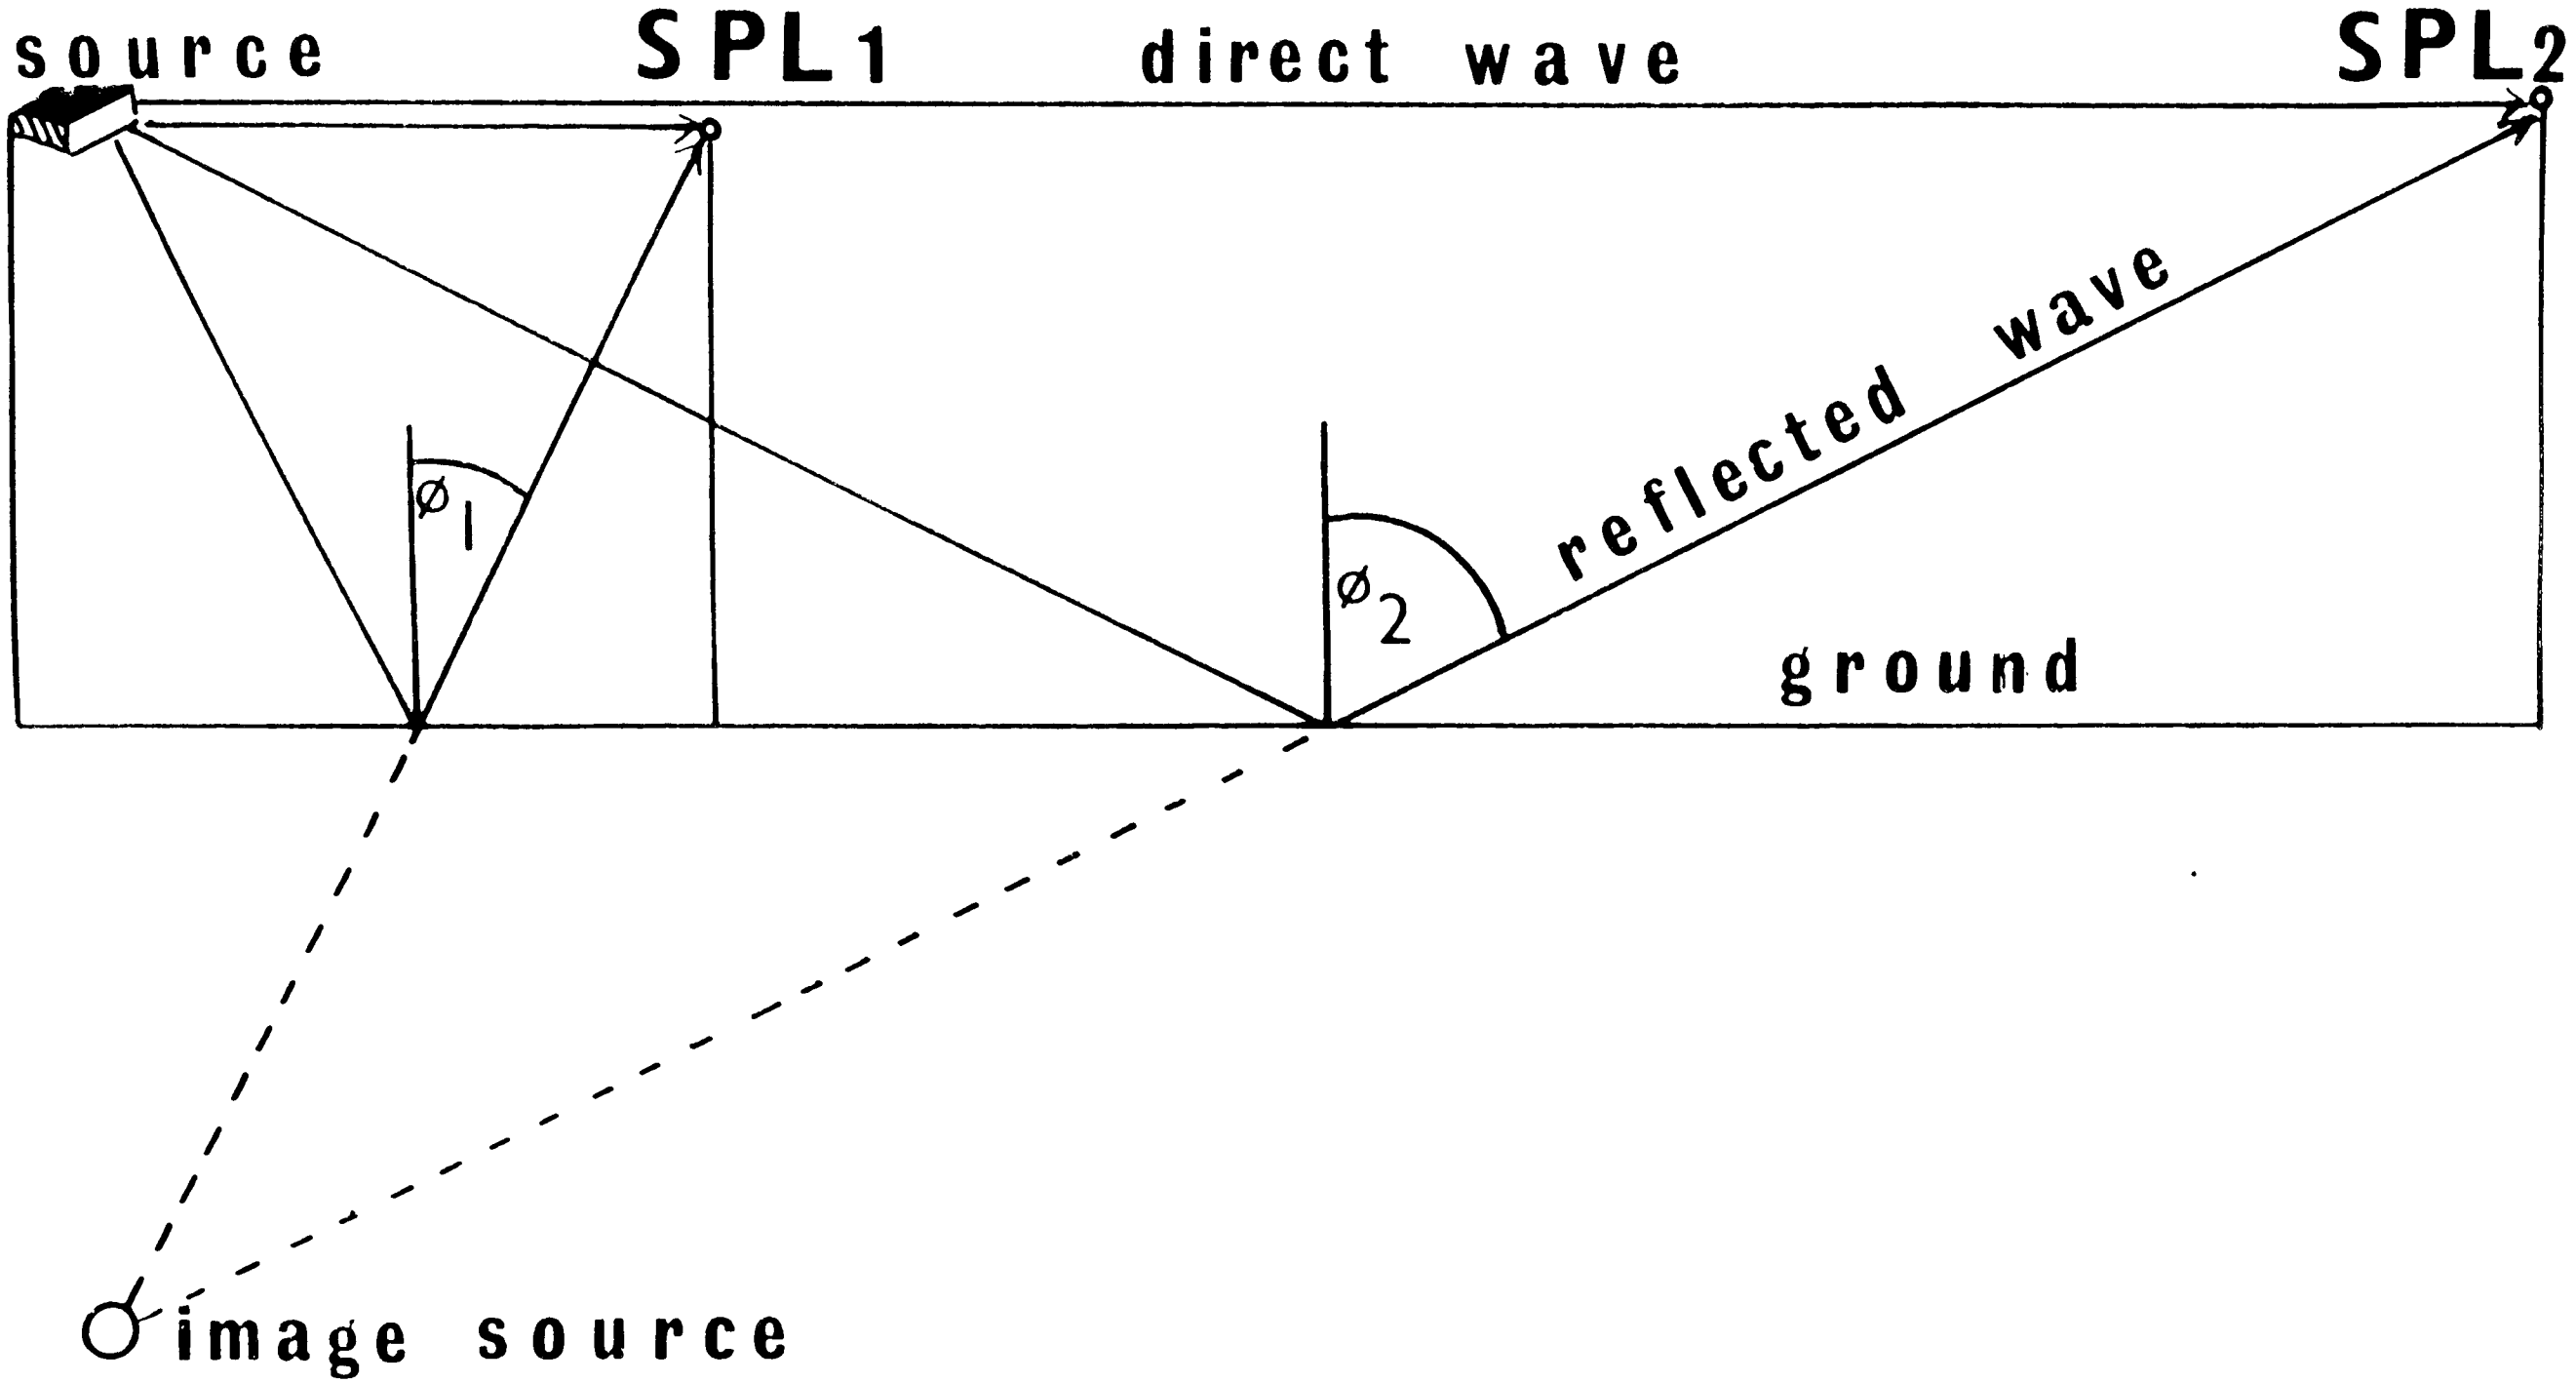
APPENDIX 1.

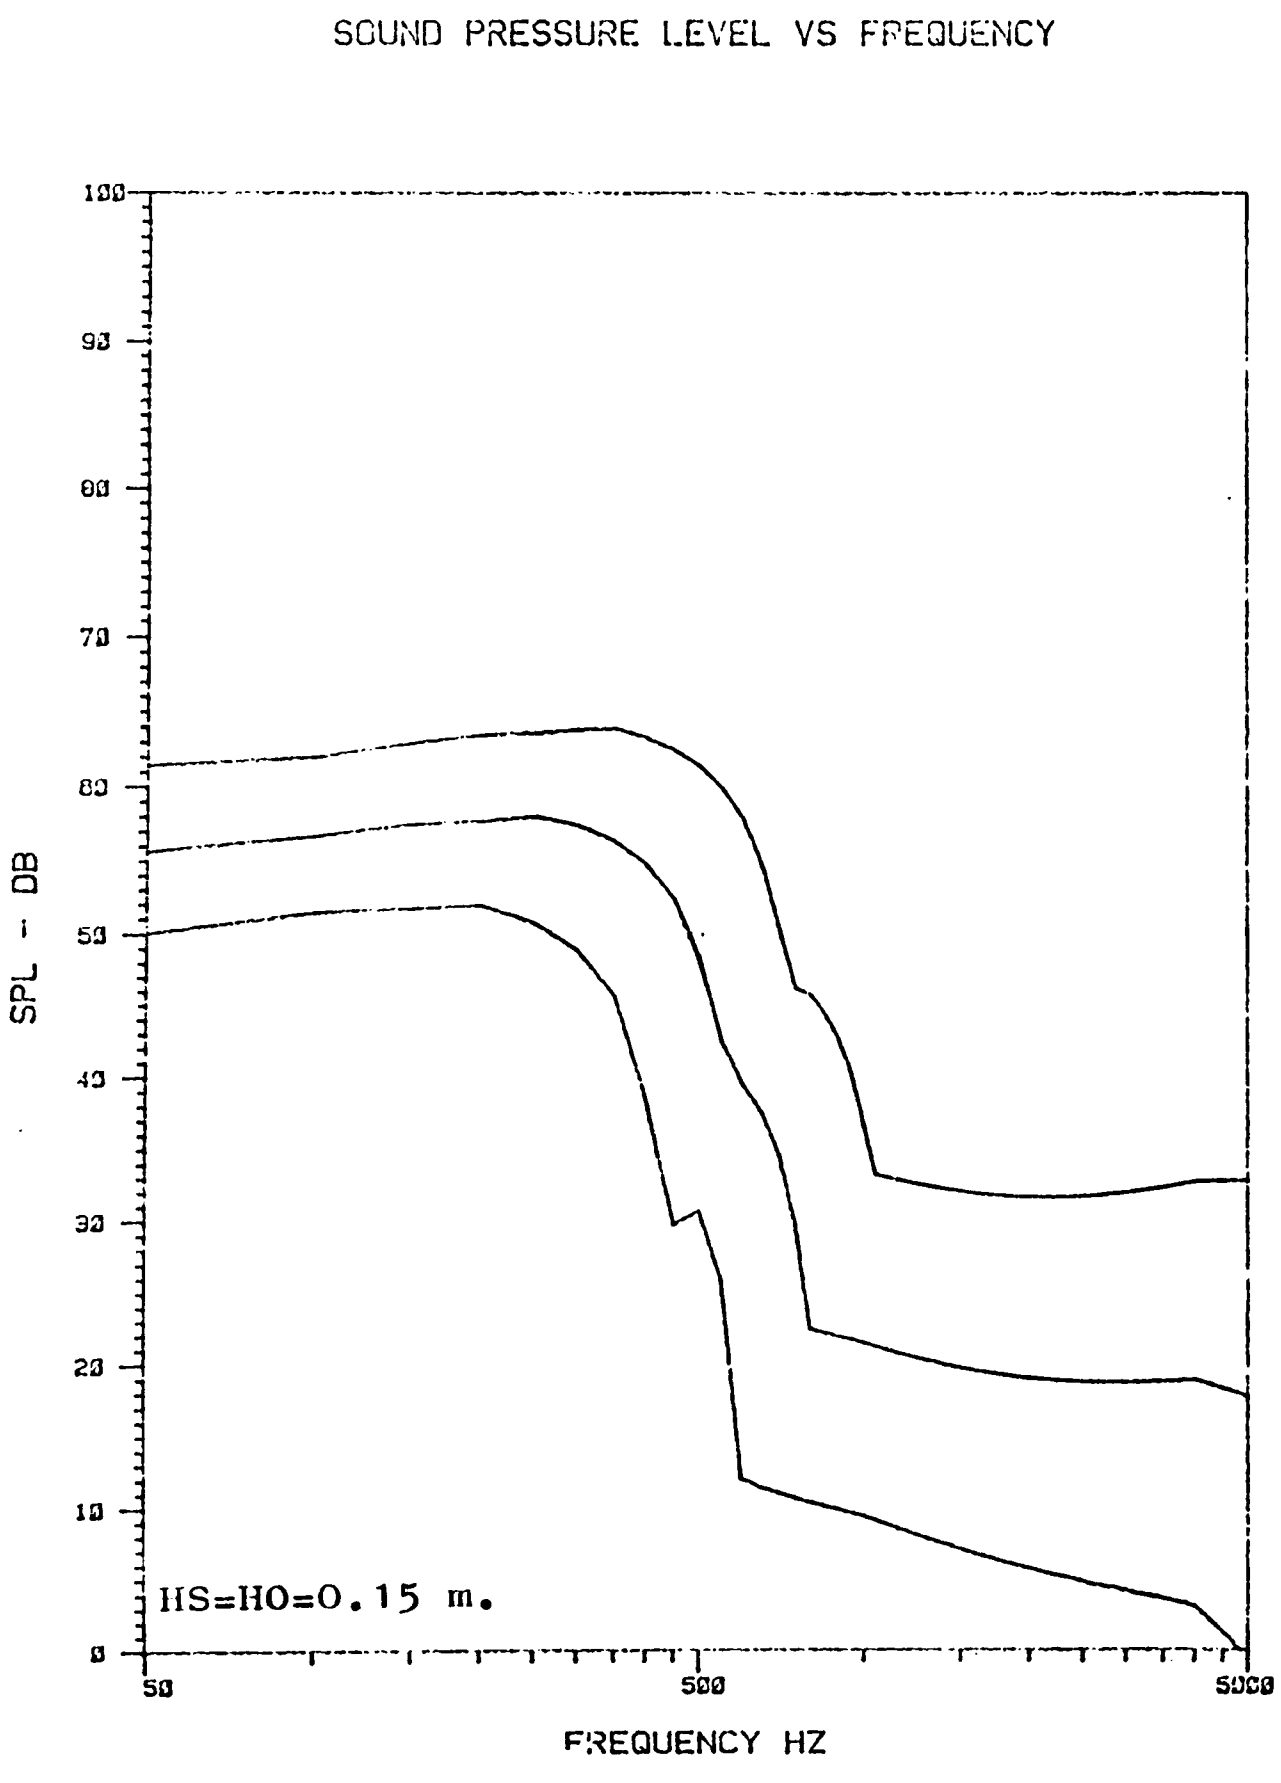
Fortran program for predicting the Sound Pressure Level at an observer of height H, a horizontal distance D from a source of sound power 0.01 W at a height HS above a flat reflecting plane.

```
PAGE 1
00100            INTEGER F,FREQ(1000),I,I1,I2,I3,VELOCITY
00200            REAL COSPHI,C1,C2,DENSITY,DISTANCE(100),H(7),HS(7),K,MOD
,MODLOG
00300            REAL PLOG,PIE,R1,R2,SPOWER,LEVEL,ATTEN,MODW,EF,DSQR
00400            COMPLEX B,BETA,J,P,W,P1,P2,RP
00500            B(FQ)=CMPLX(5.246E-3+1.34E-1*FQ-2.395E-2*FQ*FQ+1.29E-3*F
Q**3,
00600           1  +1.92E-2+1.253E-1*FQ-2.241E-2*FQ*FQ+1.21E-3*FQ**3)
00700            DENSITY=1.2
00800            VELOCITY=345
00900            SPOWER=0.01
01000            PIE=4*ATAN(1.0)
01100            WRITE(3,53)
01200    53      FORMAT(' HEADING SOUND PRESSURE LEVEL VS FREQUENCY',/' F
RAME',/
01300           1  ' YAXIS 0 100 10 10',/' YTITLE SPL - DB',/' XLOG',/
01400           2  ' XAXIS 50 5000 450 9',/' XTITLE FREQUENCY HZ'/
01500           3  ' NUMBER 8',/' PENS 1 1 1 1 1 1 1 1')
01600            C1=SQRT(DENSITY*VELOCITY*SPOWER/(4*PIE))
01700            J=CMPLX(0.0,1.0)
01800            DO 50 I=1,5
01900            DO 45 I1=1,3
02000            WRITE(3,54) I1
02100            DO 35 I3=1,100
02200            AI=I3
02300            F=50*(AI)
02400            BETA=B(F/1000.0)
02500            K=2*PIE*F/VELOCITY
02600            HS(1)=0.15
02700            HS(2)=01.0
02800            HS(3)=2.0
02900            HS(4)=5.0
03000            HS(5)=10.0
03100            I2=I
03200            H(I2)=HS(I)
03300    54      FORMAT(' POINTS', I6)
03400            DISTANCE(1)=50.0
03500            DISTANCE(2)=100.0
03600            DISTANCE(3)=200.0
03700            DSQR=(DISTANCE(I1))*(DISTANCE(I1))
03800            R1=SQRT(DSQR+(HS(I)-H(I2))*(HS(I)-H(I2)))
03900            R2=SQRT(DSQR+(HS(I)+H(I2))*(HS(I)+H(I2)))
04000            C2=(HS(I)+H(I2))*(HS(I)+H(I2))
04100            COSPHI=SQRT(C2/(C2+DSQR))
04200            P1=CEXP(J*K*R1)/R1
04300            P2=CEXP(J*K*R2)/R2
04400            RP=(COSPHI-BETA)/(COSPHI+BETA)
04500            W=0.5*(J*K*R1)*(COSPHI+BETA)*(COSPHI+BETA)
04600            MODW=CABS(W)
04700            IF (MODW-0.2) 19,21,21
```

```
04800    19        EF=1.0+1.5*MODW
04900              GO TO 30
05000    21        IF (MODW-1.0) 22,23,23
05100    22        EF=1.28+0.125*MODW
05200              GO TO 30
05300    23        IF (MODW-5.0) 24,25,25
05400    24        EF=1.65-0.25*MODW
05500              GO TO 30
05600    25        IF (MODW-7.5) 26,27,27
05700    26        EF=0.83-0.1*MODW
05800              GO TO 30
05900    27        IF (MODW-15) 28,29,29
06000    28        EF=0.3-0.02*MODW
06100              GO TO 30
06200    29        EF=0
06300    30        P=-J*C1*(P1+P2*(RP+(1-RP)*EF))
06400              MOD=CABS(P)
06500              MODLOG=20*ALOG10(MOD/2*1.E5)
06600              IF (F-1000) 37,39,39
06700    37        ATTEN=0
06800              GO TO 43
06900    39        IF (F-4000) 41,42,42
07000    41        ATTEN=(11.0*(F-1000)/9000.0)
07100              GO TO 43
07200    42        ATTEN=(11.0*(F-1000)/9000.0+8.2*(F-4000)/6000.0)
07300    43        LEVEL=MODLOG-ATTEN*DISTANCE(I1)/100.0
07400              PLOG=20*ALOG10(C1/(2.0*(DISTANCE(I1))*1.E-5))
07500              WRITE(3,10) F,LEVEL
07600    10        FORMAT(I6,F8.4,$)
07700              IF (I3.LT.100) WRITE(3,55)
07800    55        FORMAT()
07900    20        CONTINUE
08000    35        CONTINUE
08100    40        CONTINUE
08200              WRITE(3,56)
08300    56        FORMAT(1H+,' LAST POINT')
08400    45        CONTINUE
08500    50        CONTINUE
08600              END
*
*
*
*
*


***
```


source
SPL1
direct wave
SPL2
ø1
ø2
reflected wave
ground
image source



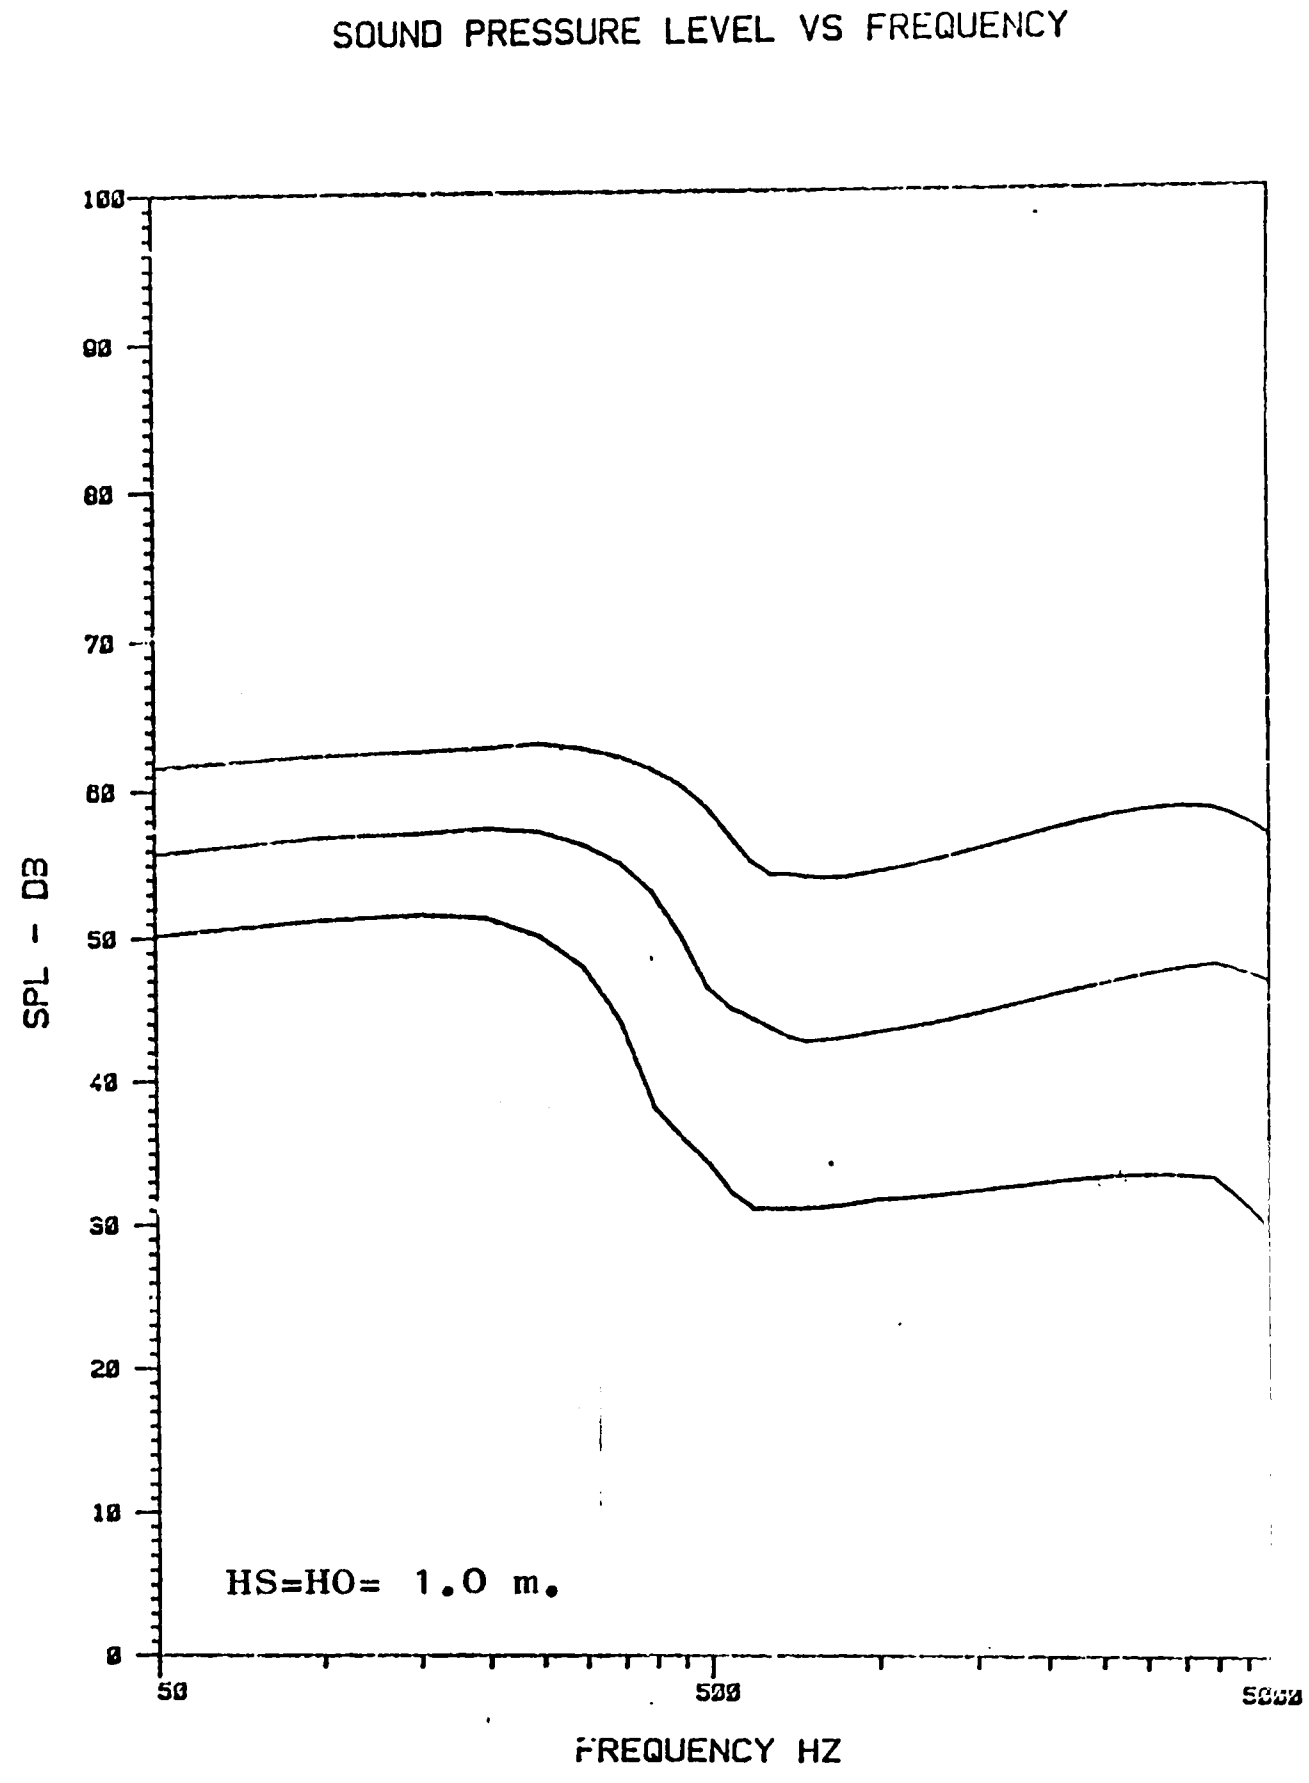
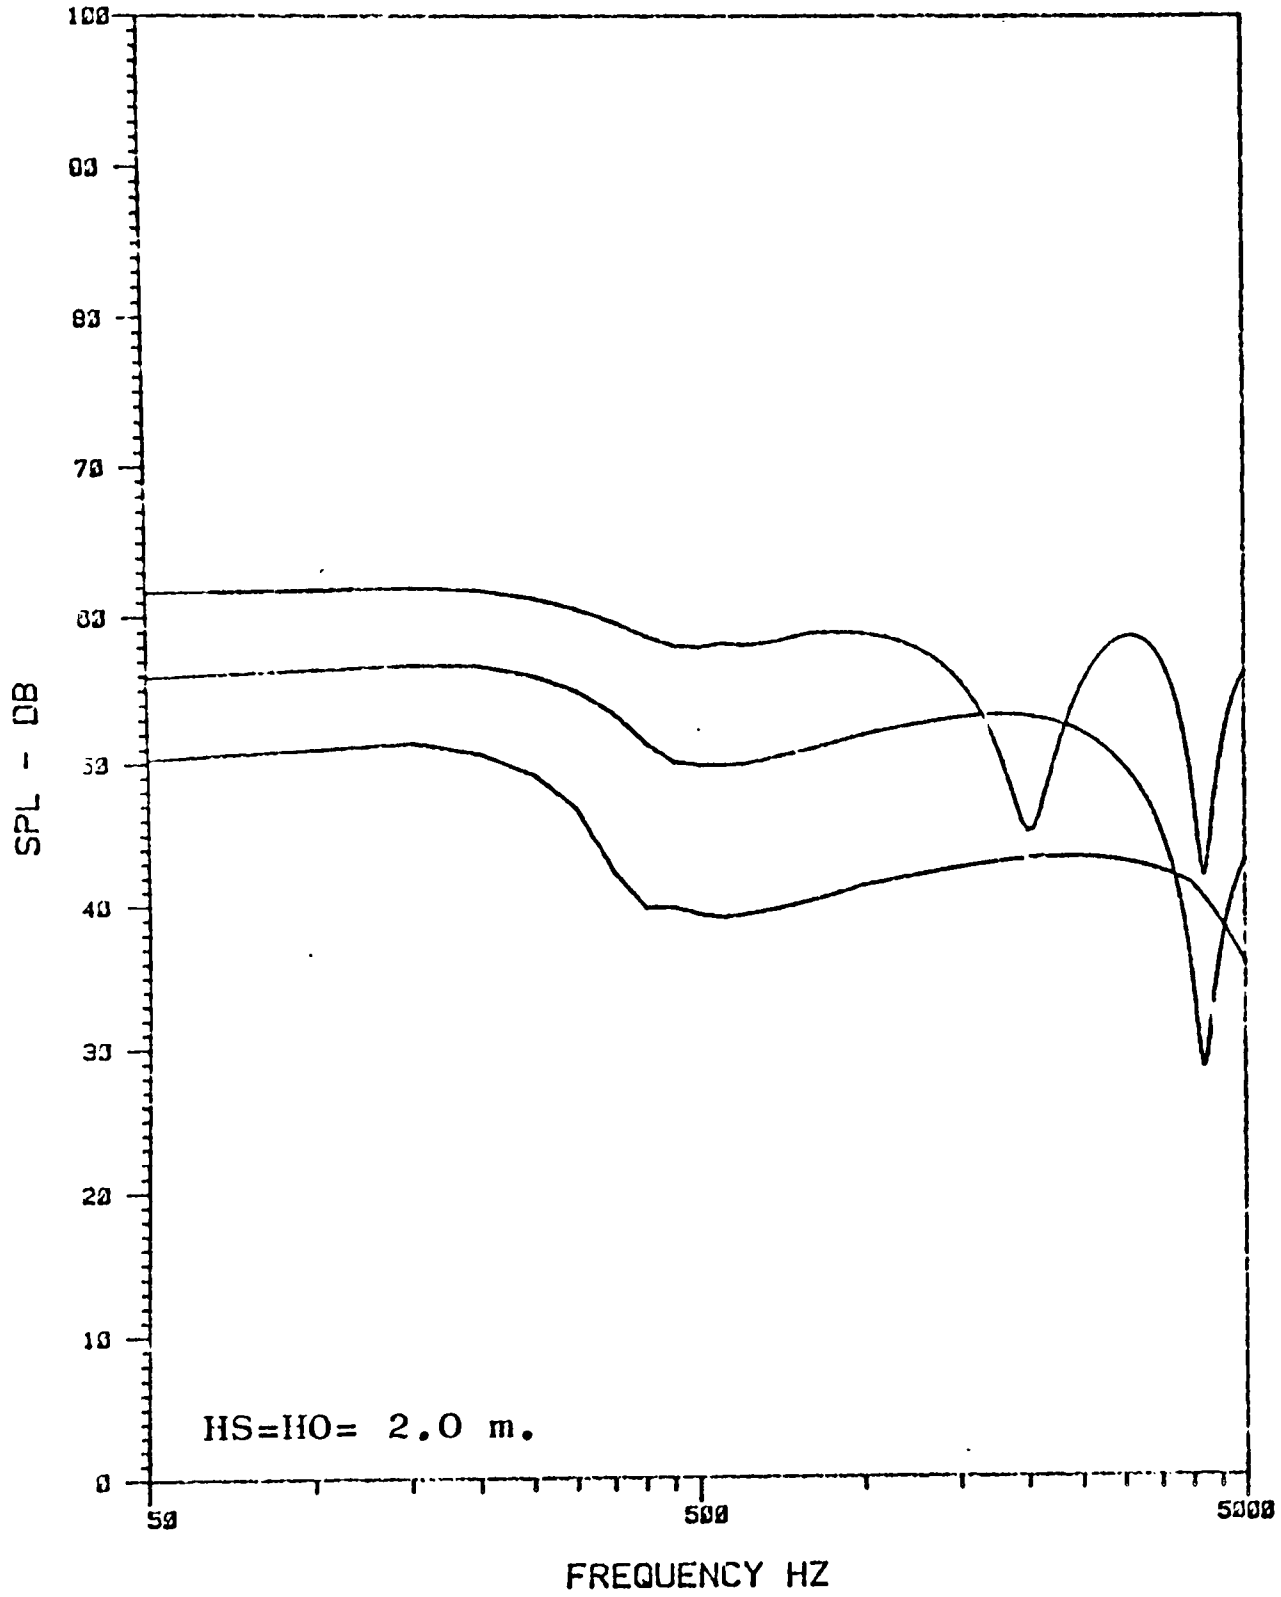


Fig. 2. SPLs predicted by the model, equation (2) for various source-receiver heights and geometries. Each fig. corresponds to one height and receivers 50, 100 and 200 m.


100
90
80
70
60
50
40
30
20
10
0
SPL - DB
HS=HO= 2.0 m.
50
500
5000
FREQUENCY HZ

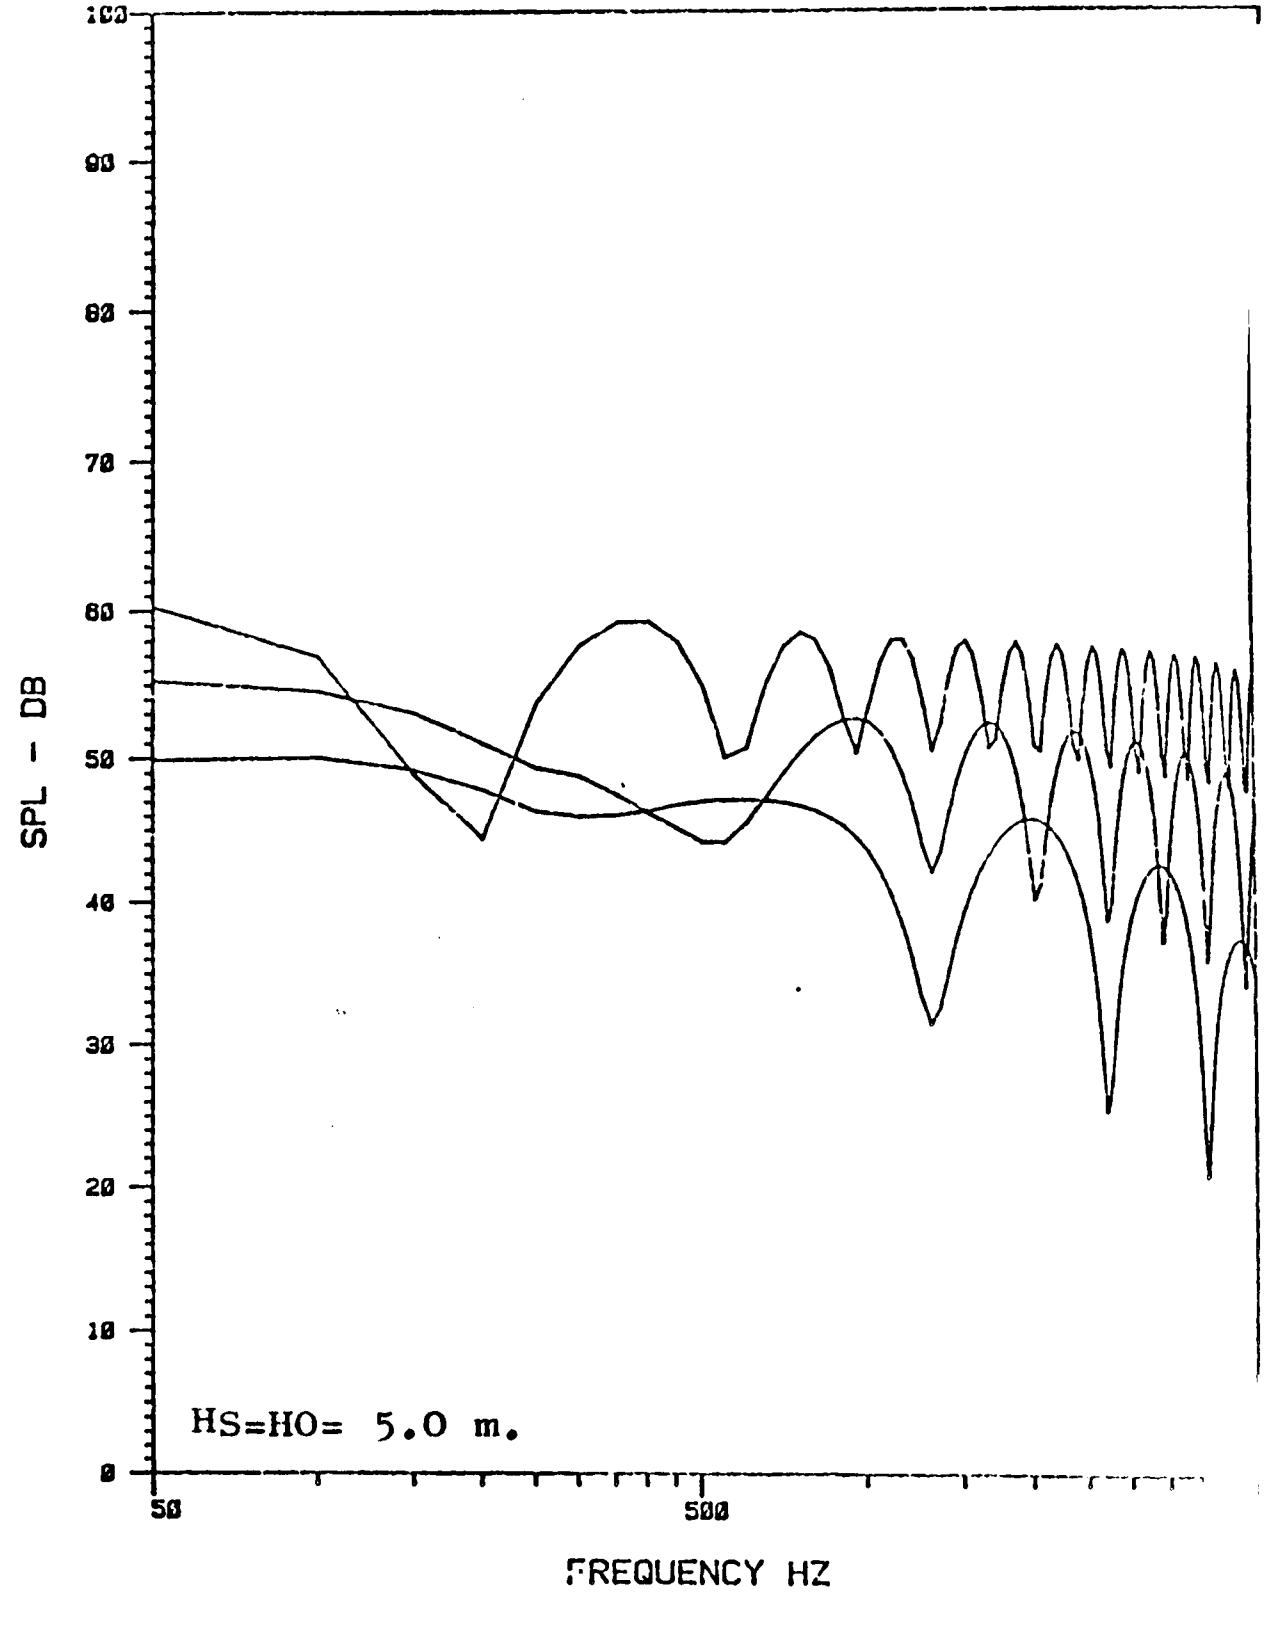
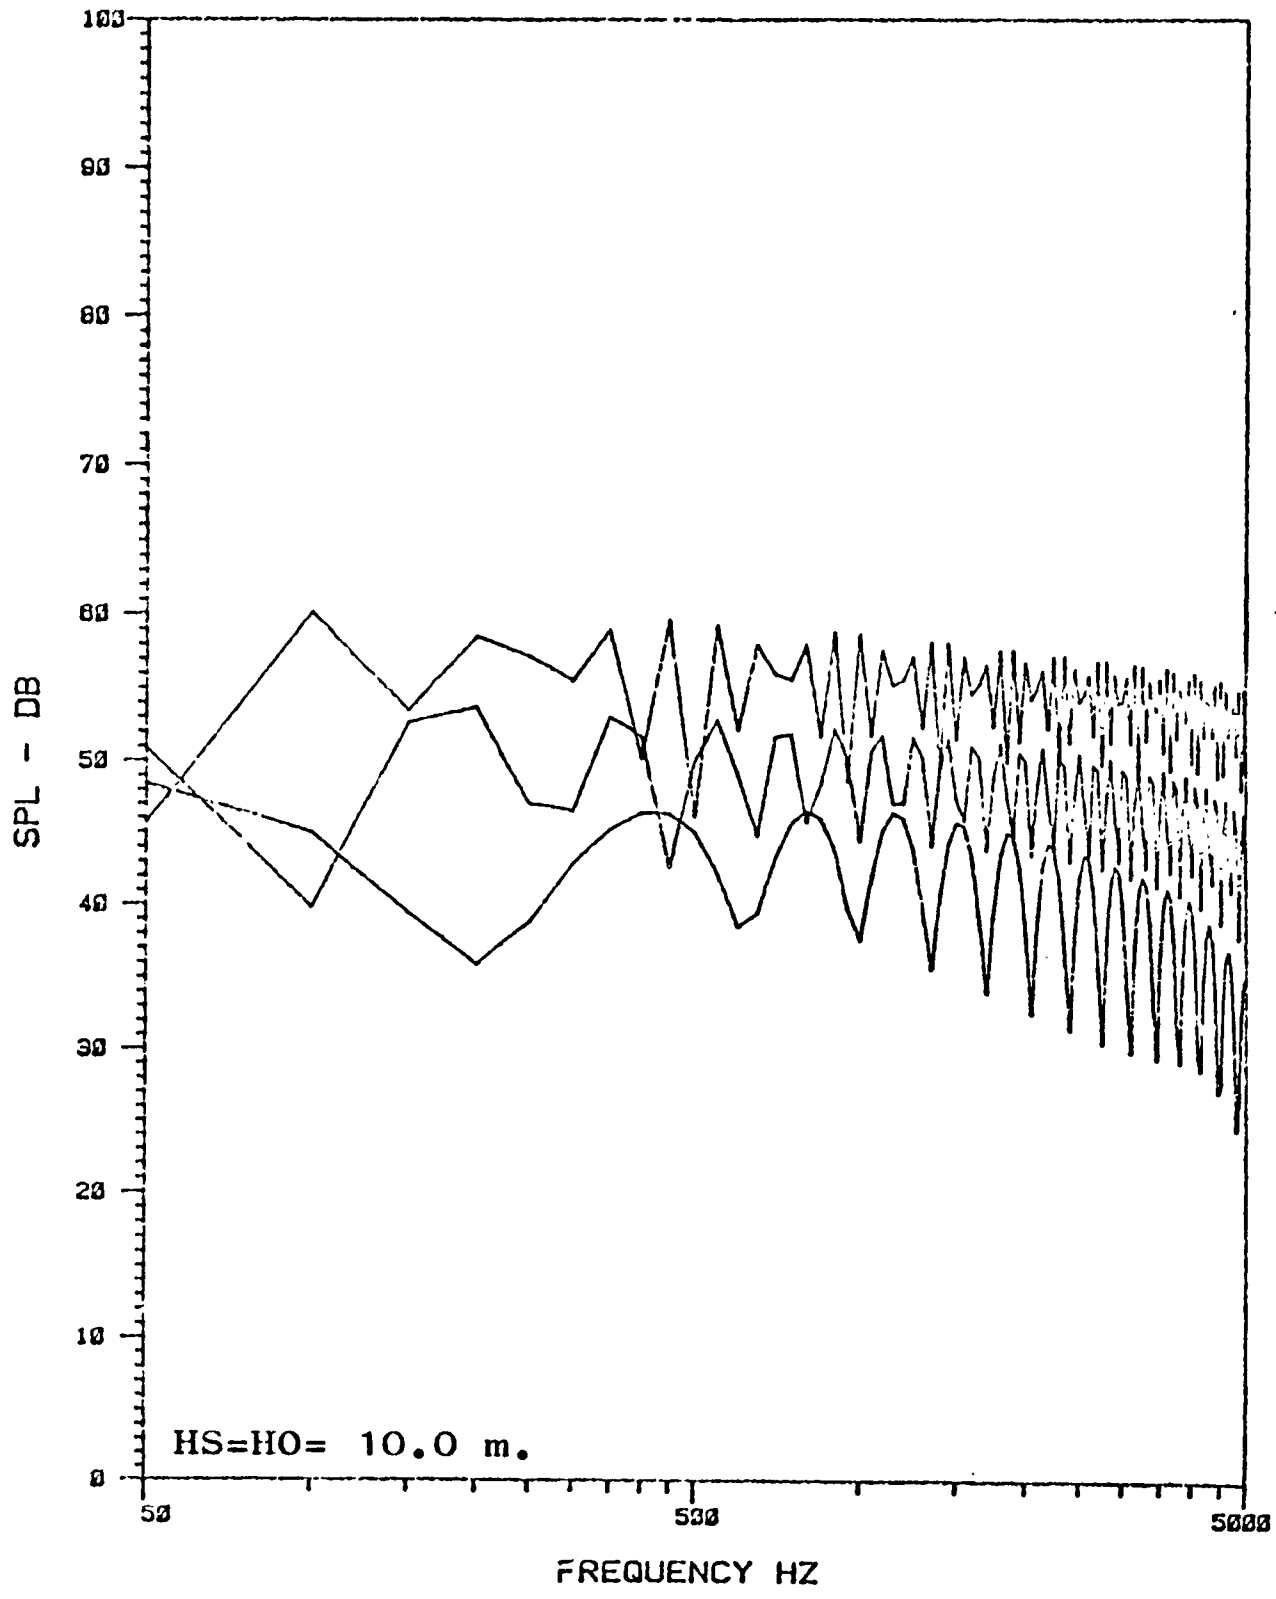
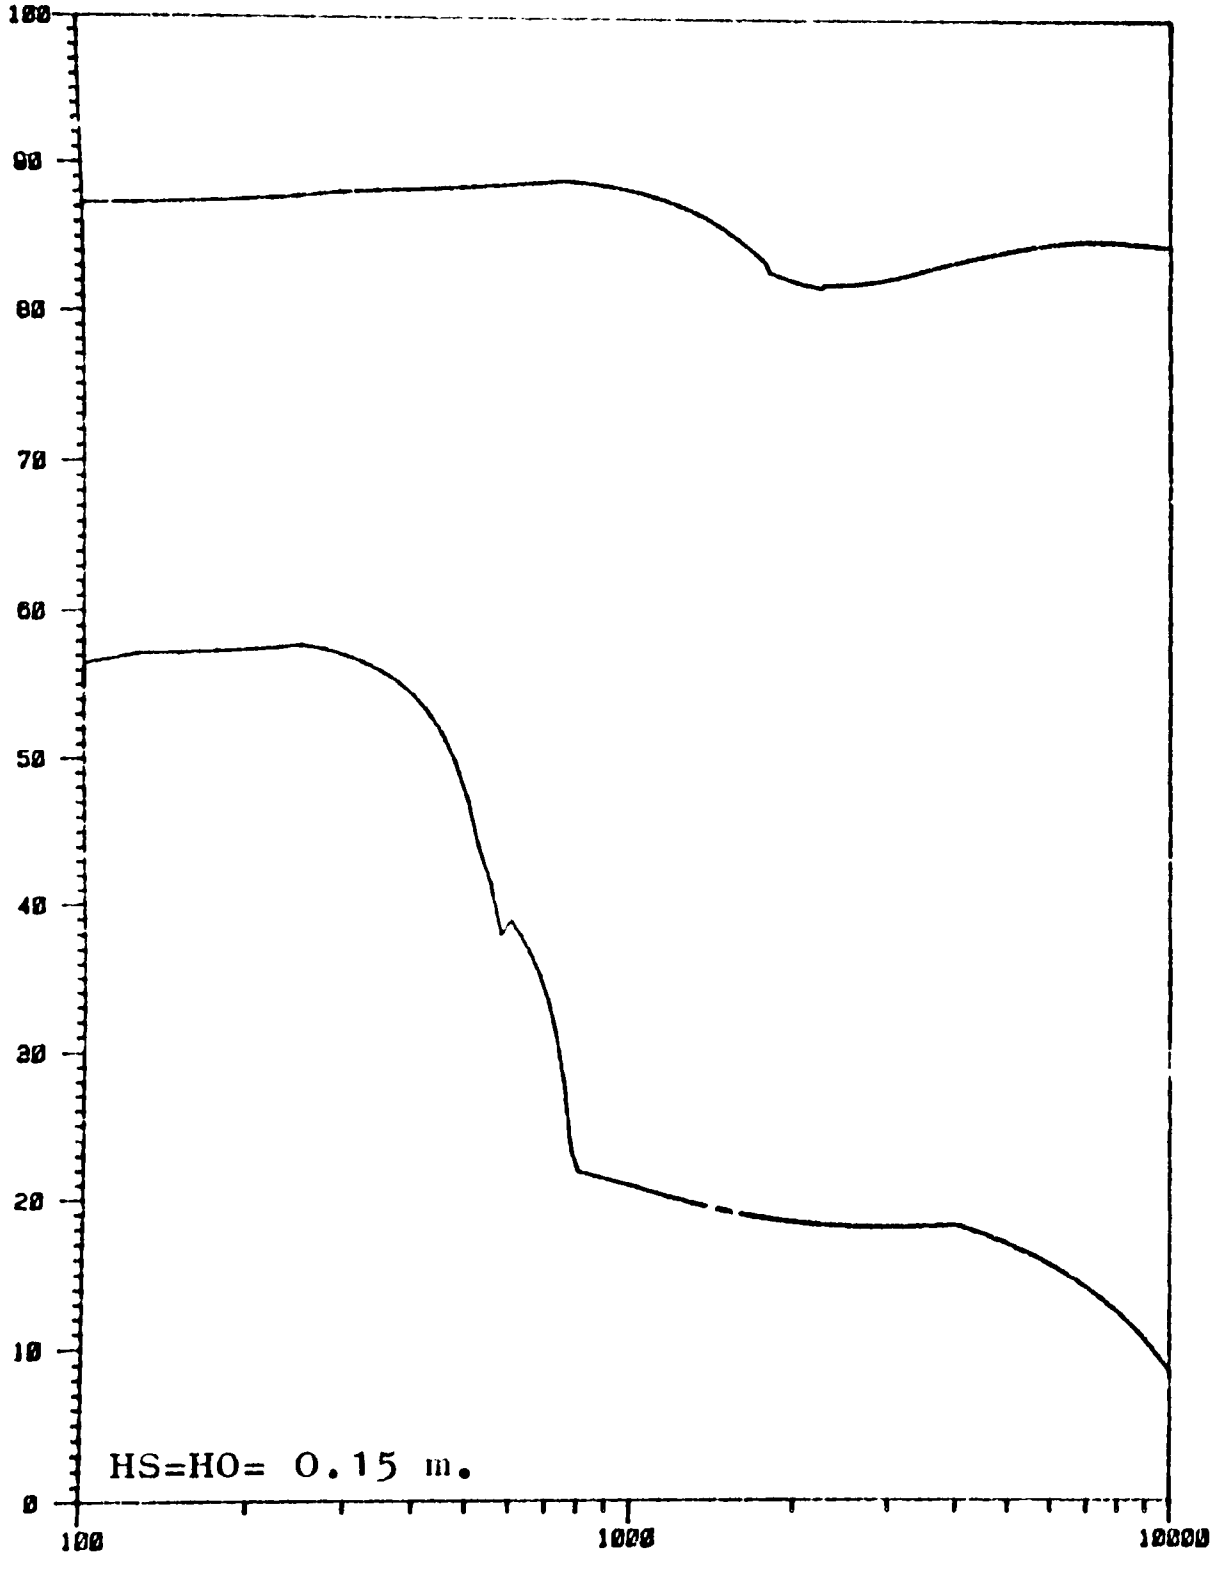

100
90
80
70
60
50
40
30
20
10
0
SPL - DB
HS=HO= 5.0 m.
50
500
FREQUENCY HZ

Fig. 2. Cont.


100
90
80
70
60
50
40
30
20
10
0
SPL - DB
HS=HO= 10.0 m.
50
500
5000
FREQUENCY HZ



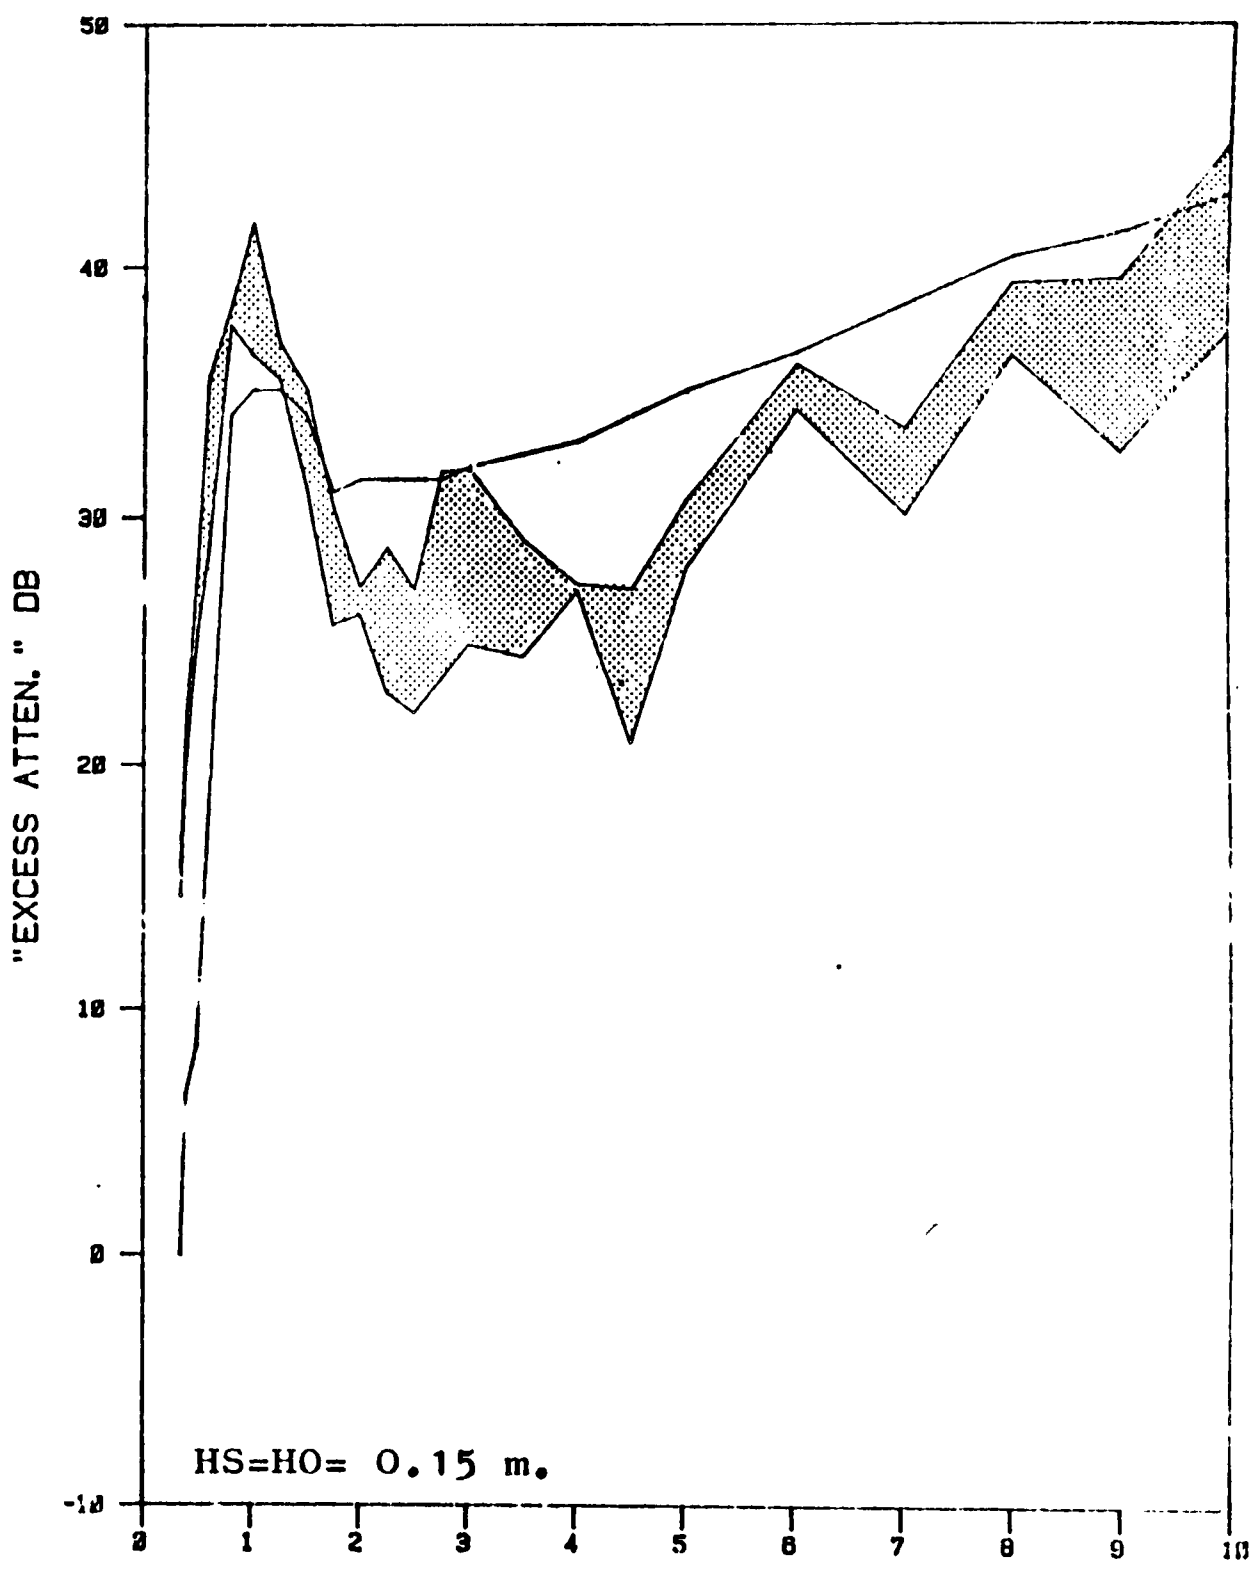
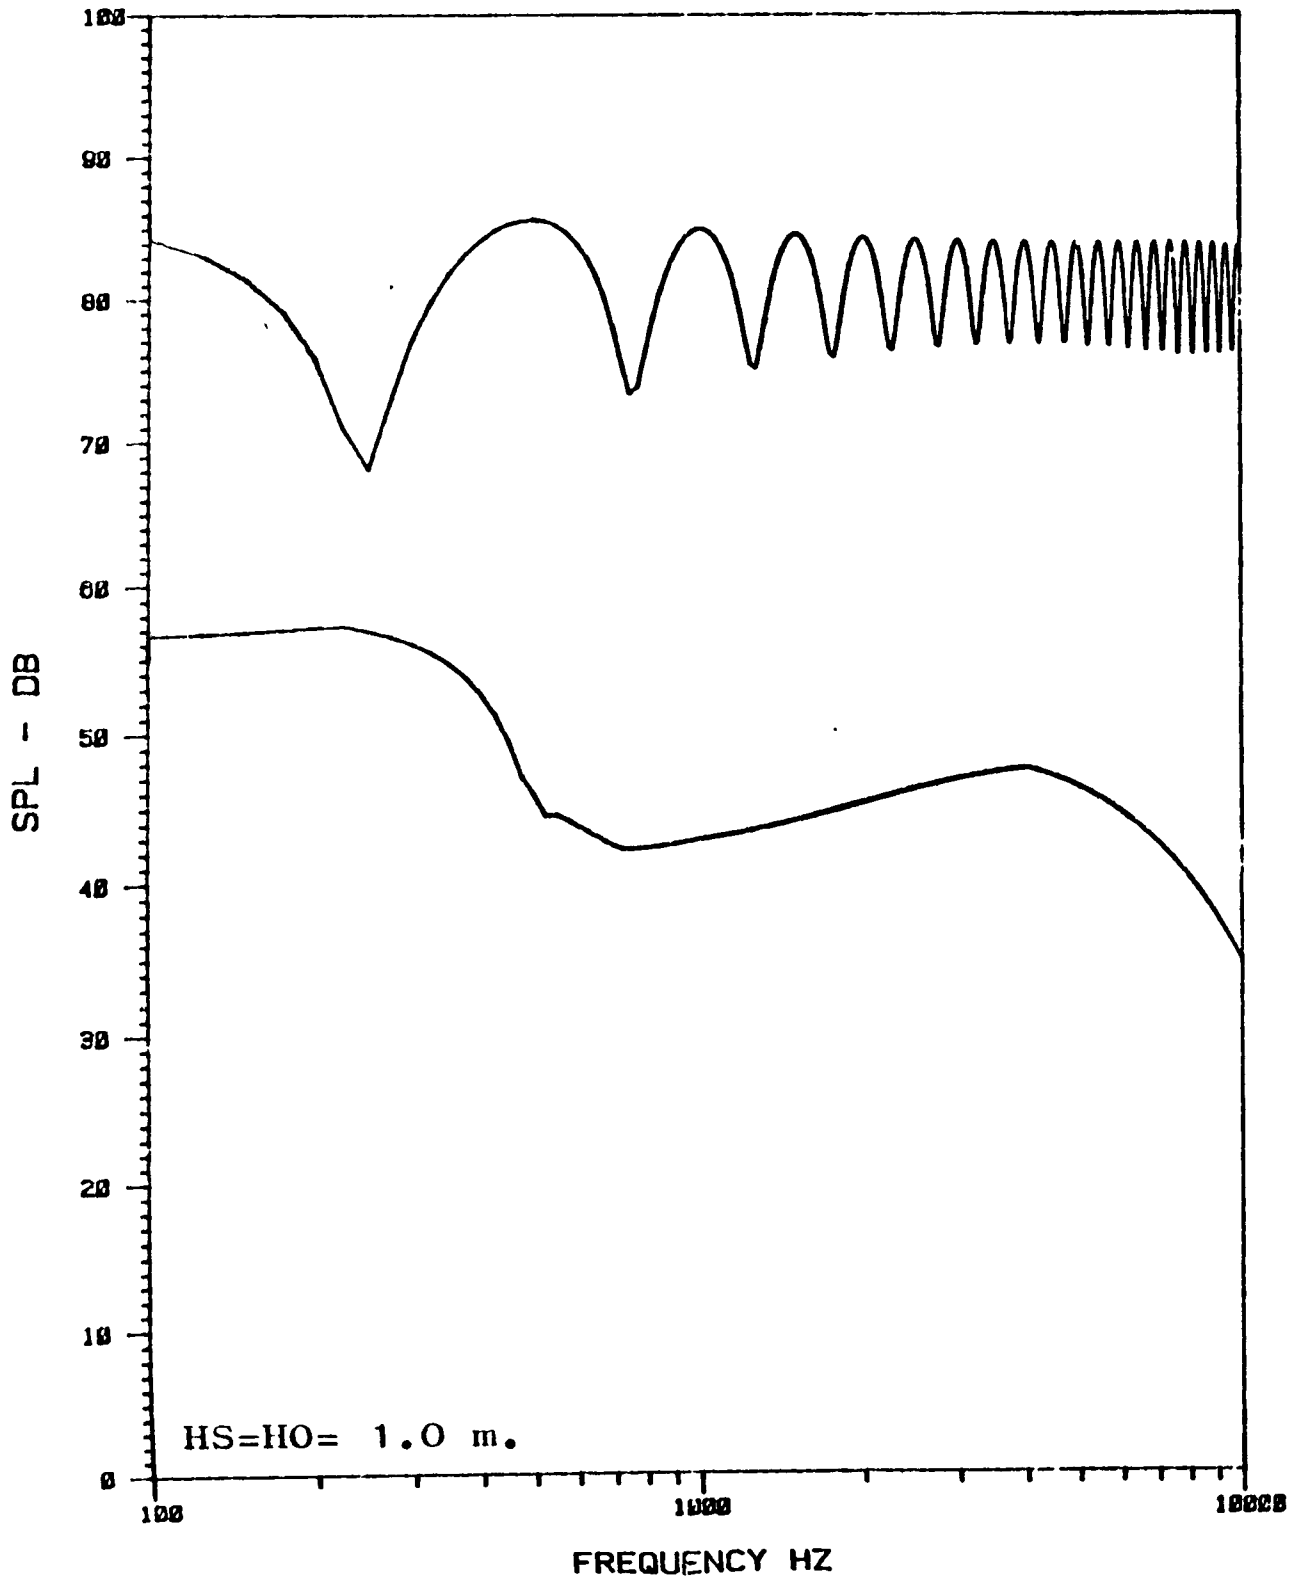


Fig.3. 'Excess attenuation' over grassland calculated from predicted SPLs (left hand frames) and measured by Marten and Marler. Each frame corresponds to one source-receiver height and receivers at 2.5 m. and 102.5 m.

SOUND PRESSURE LEVEL VS FREQUENCY



"EXCESS ATTENUATION" VS FREQUENCY

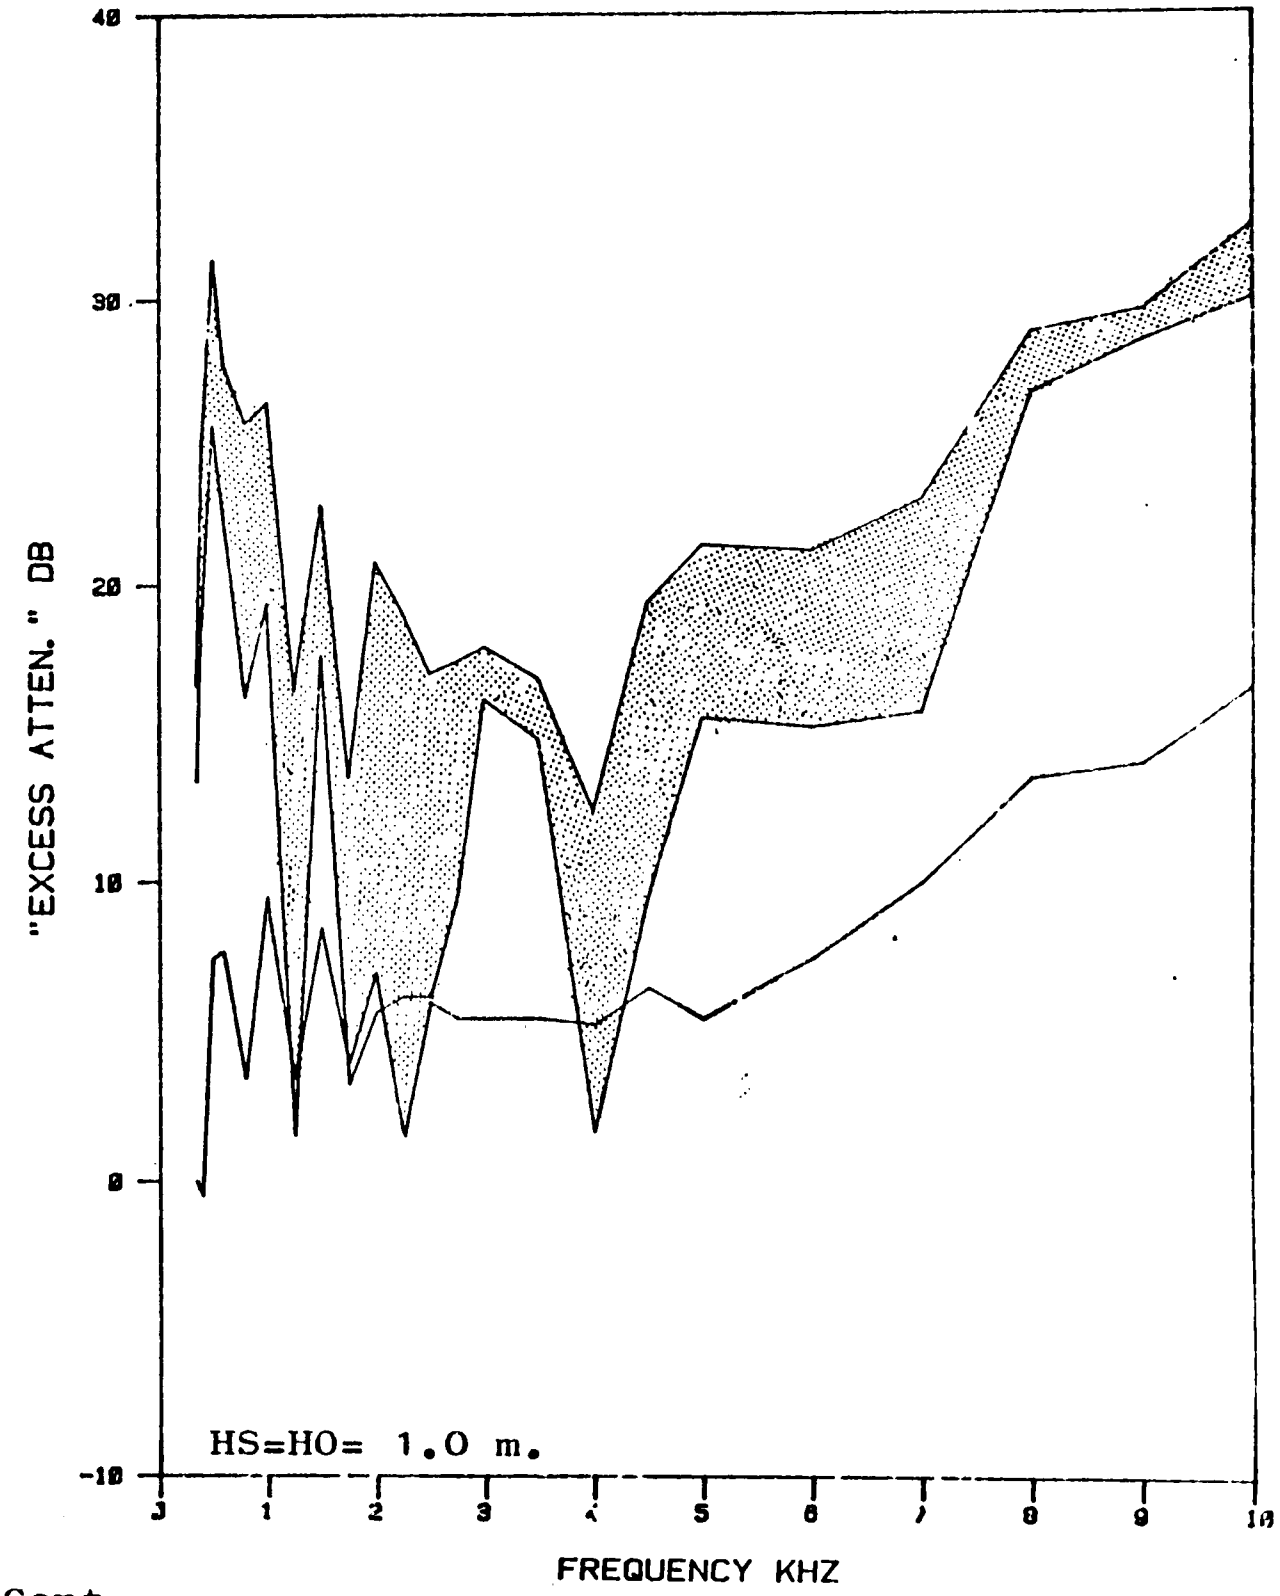
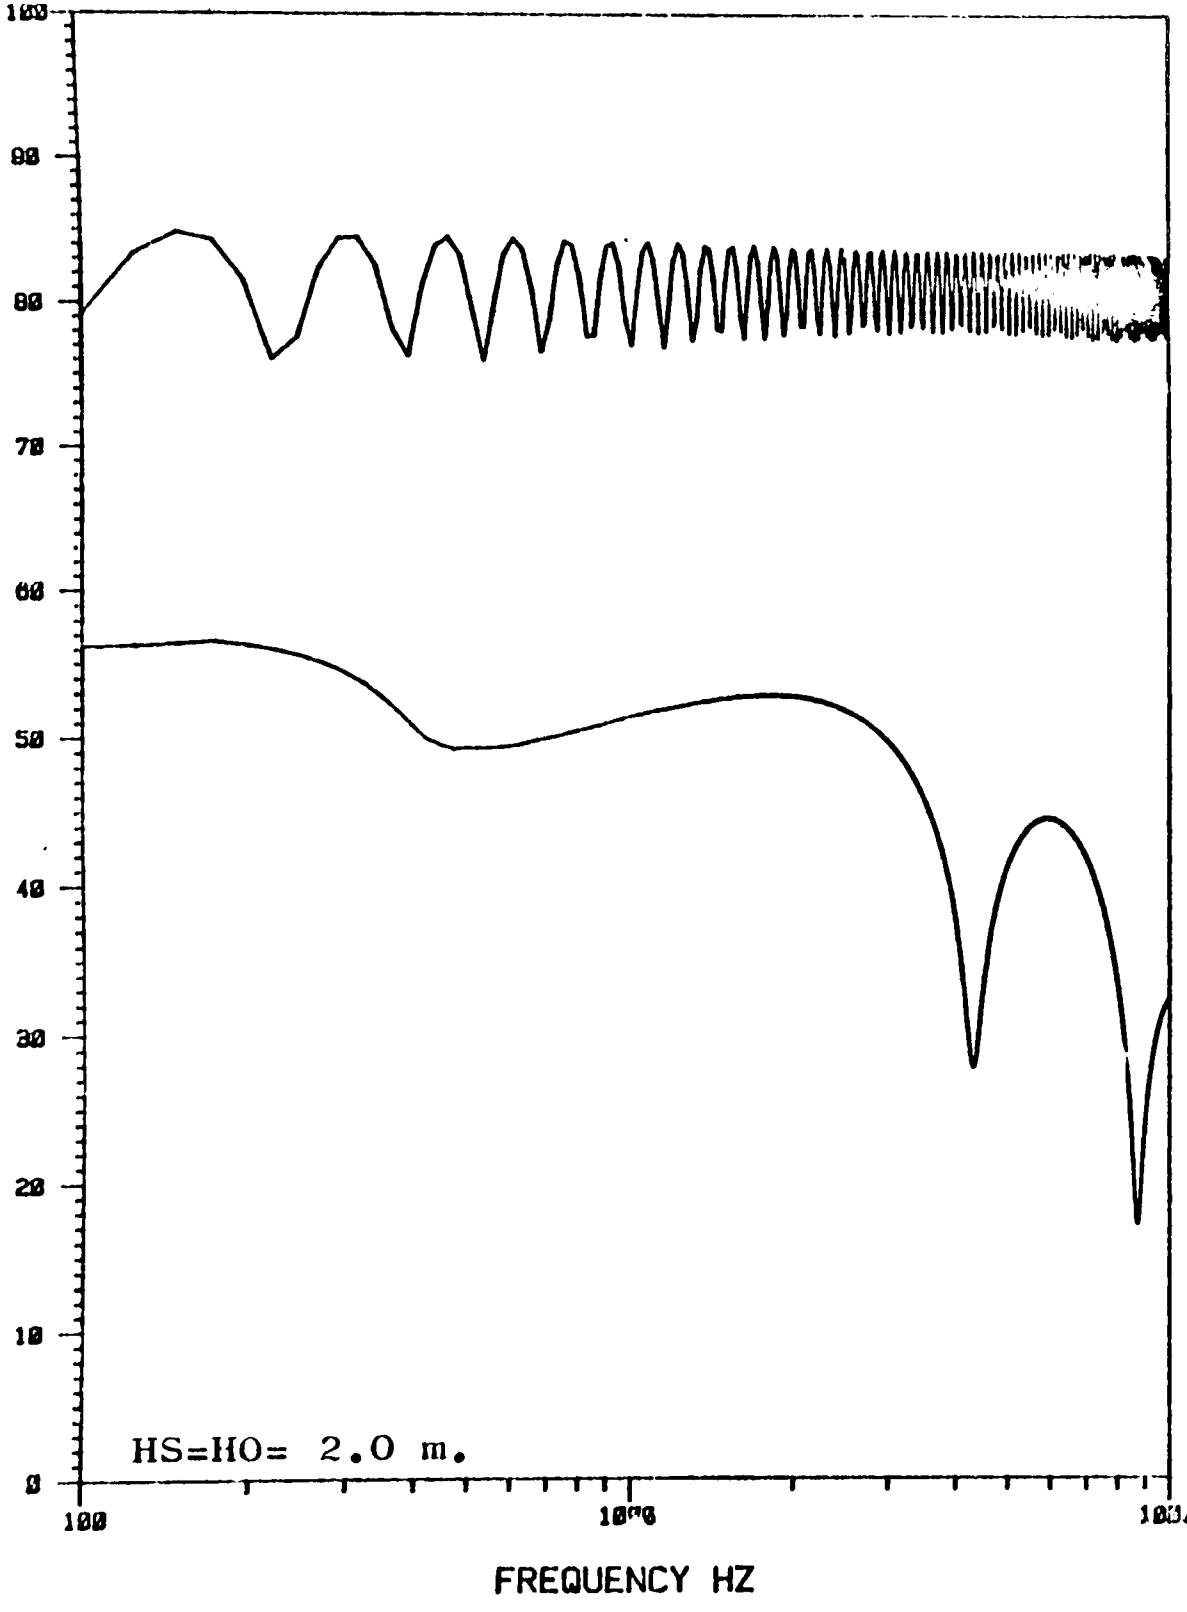


Fig.3. Cont.


100
90
80
70
60
50
40
30
20
10
0
HS=HO= 2.0 m.
100
FREQUENCY HZ

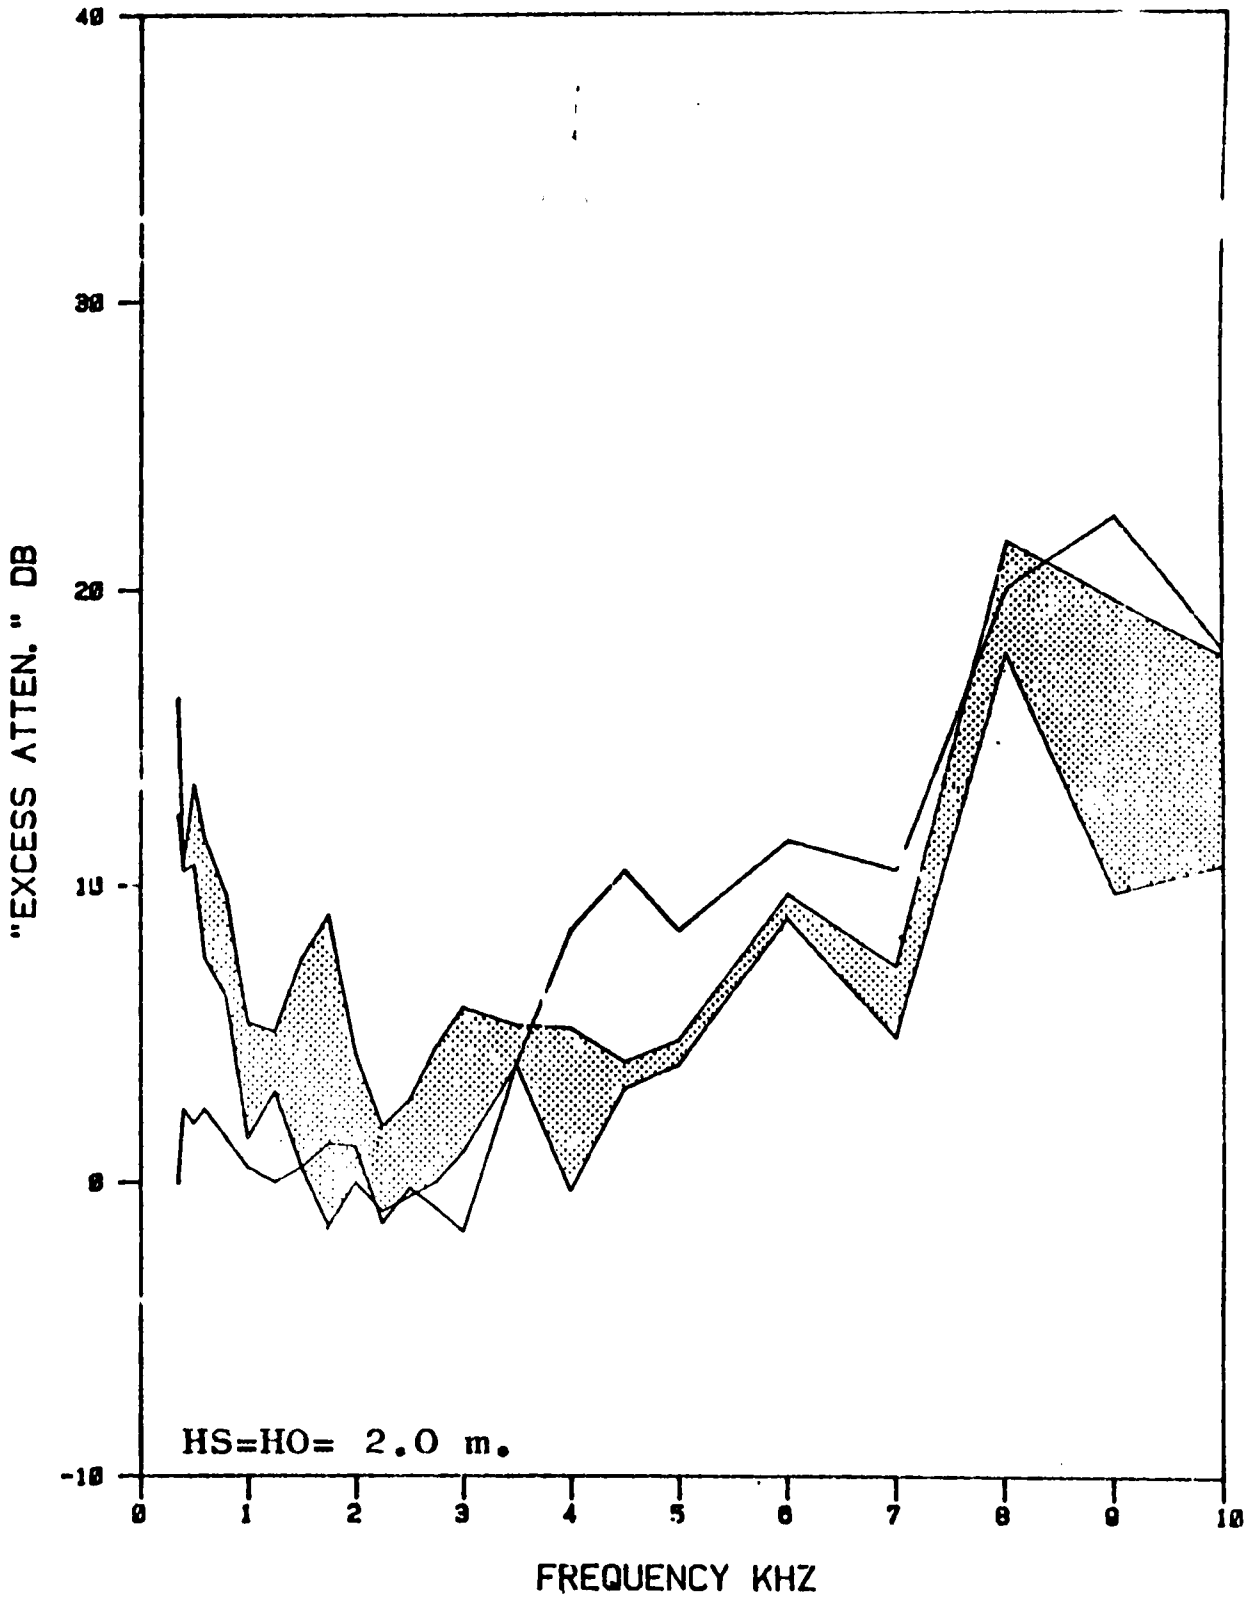
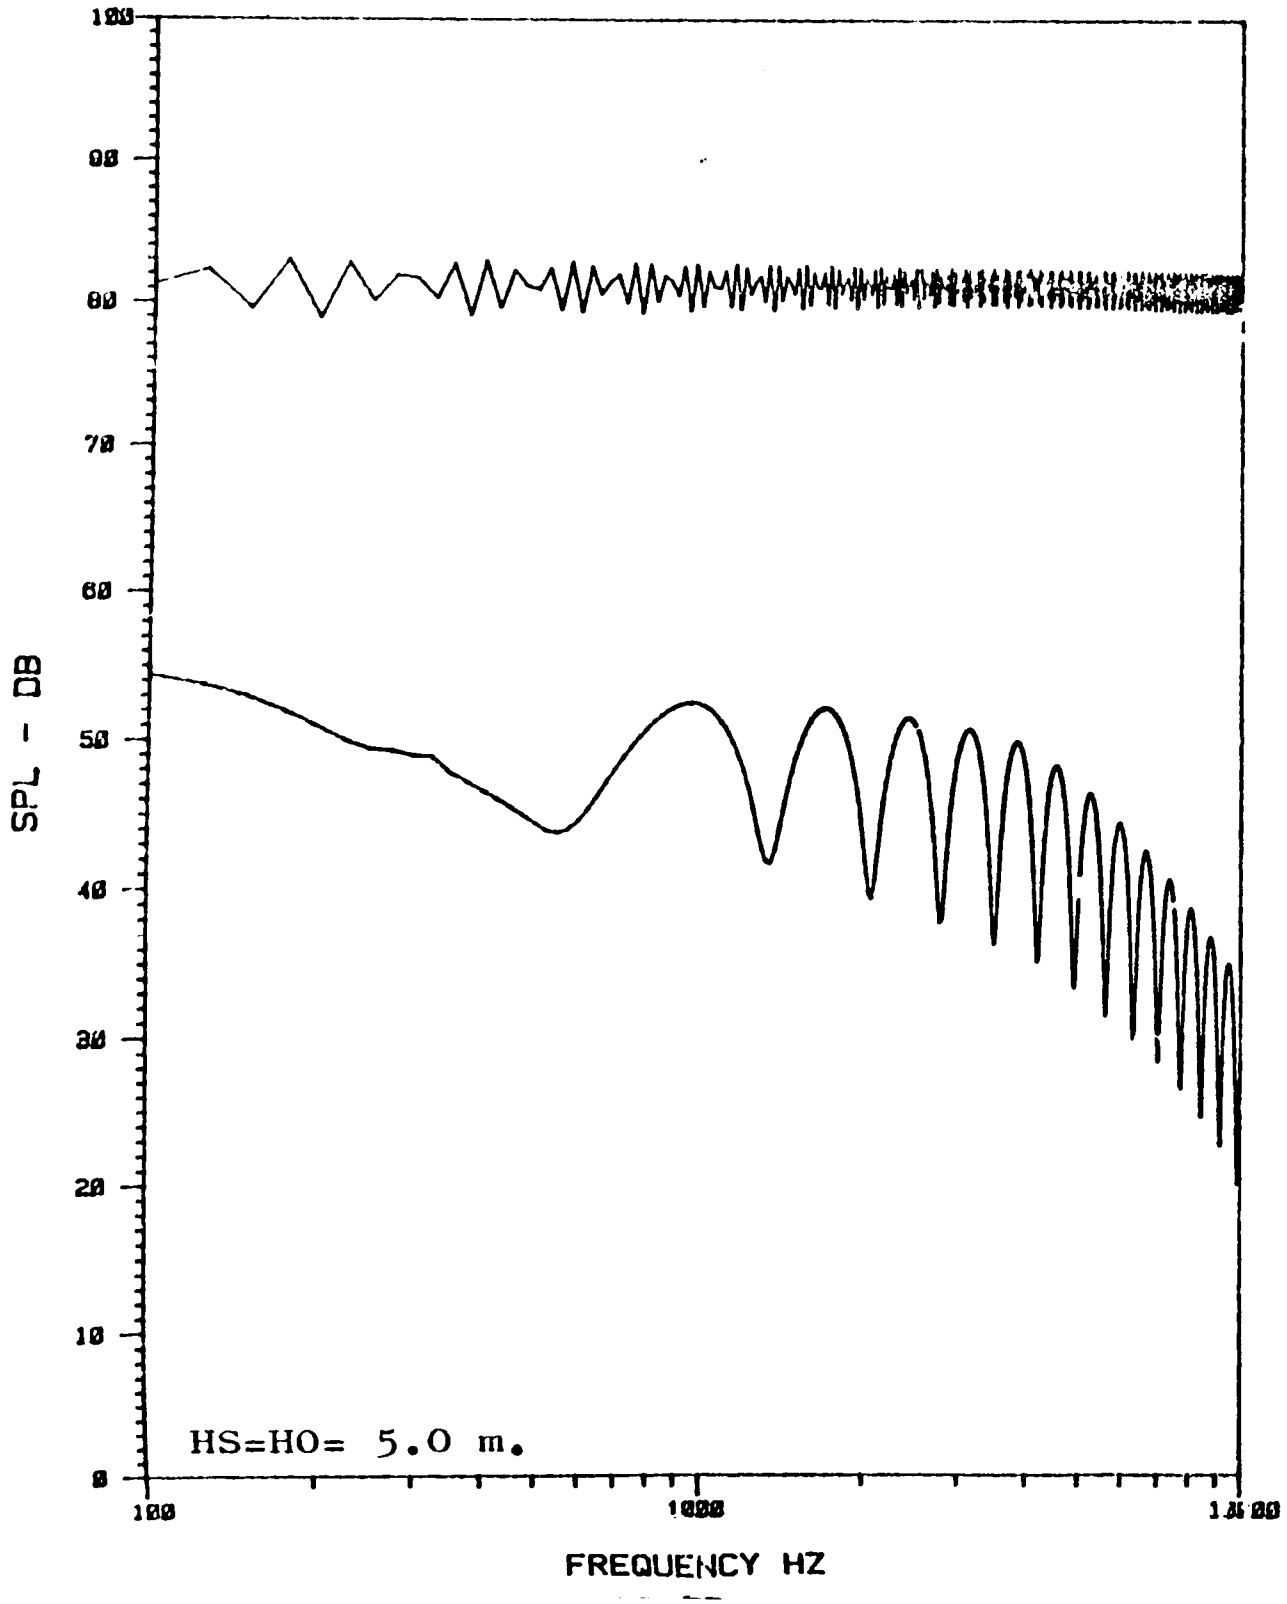
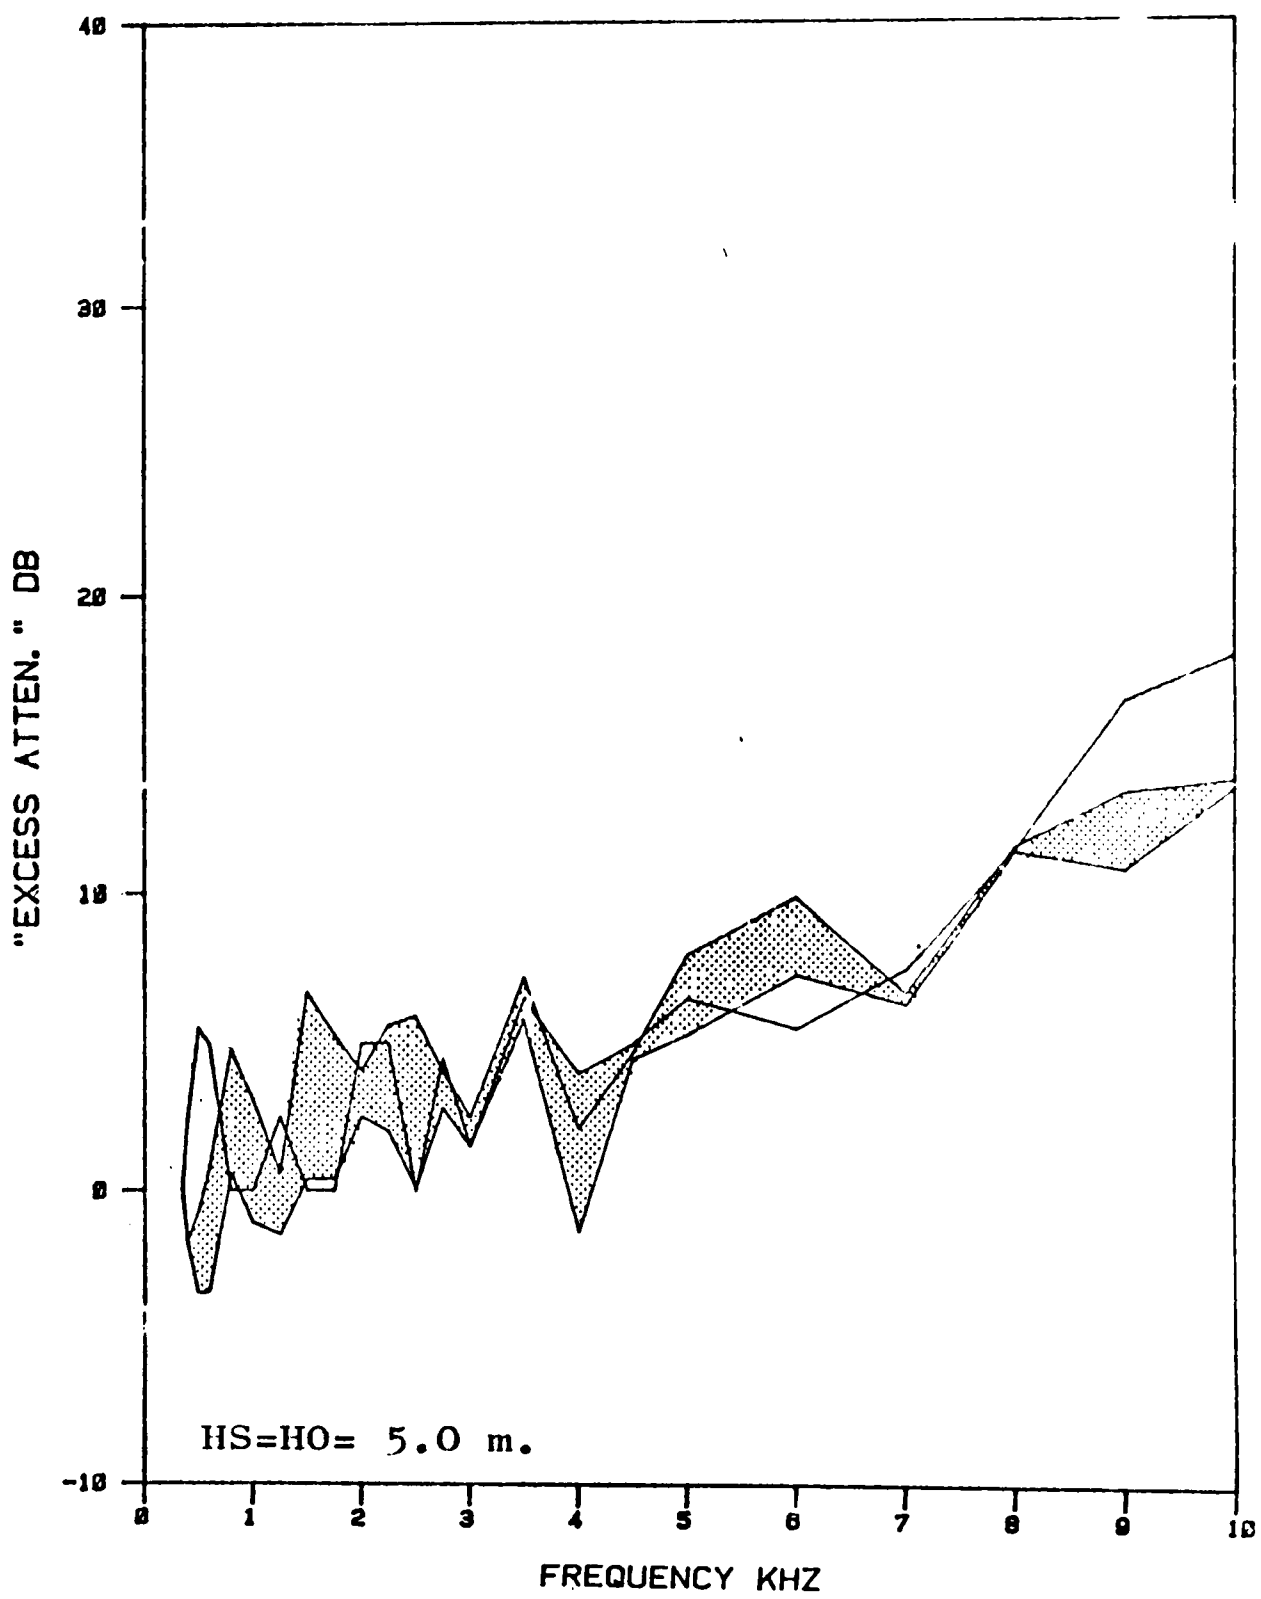

40
30
20
10
0
-10
"EXCESS ATTEN." DB
HS=HO= 2.0 m.
0
1
2
3
4
5
6
7
8
9
10
FREQUENCY KHZ

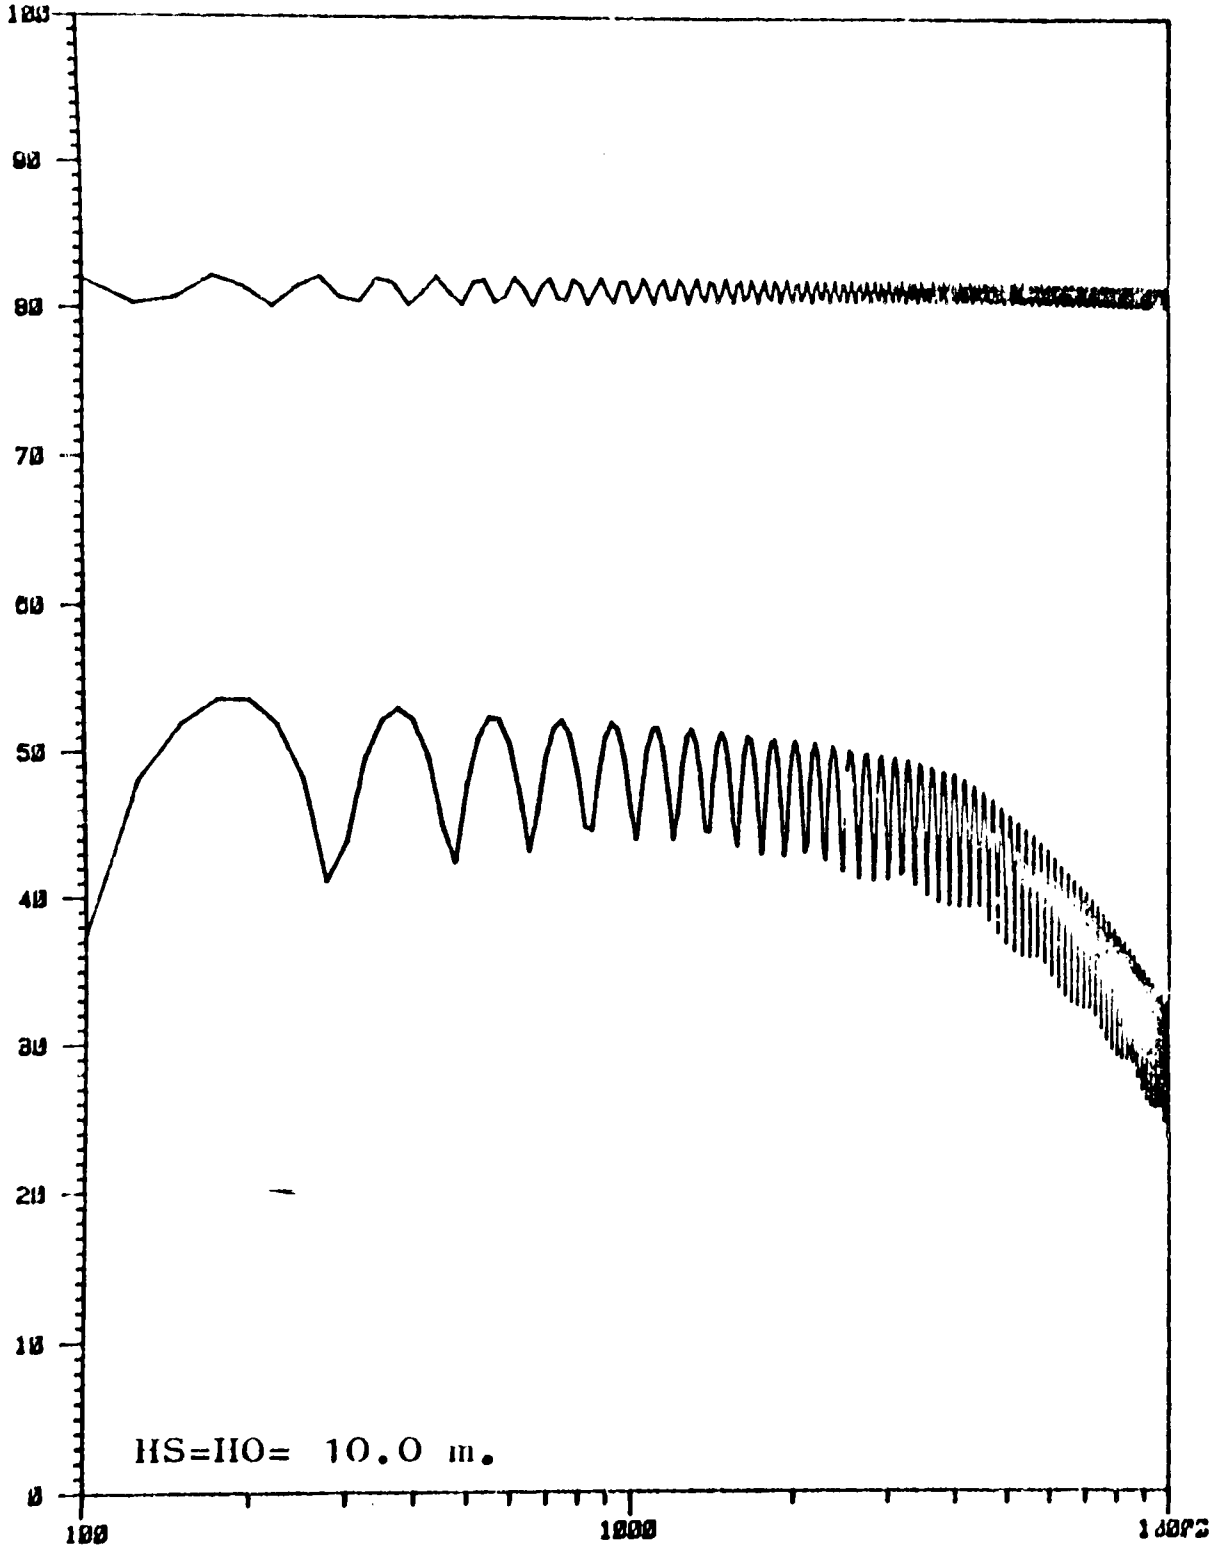
Fig.3. Cont.

SOUND PRESSURE LEVEL VS FREQUENCY



"EXCESS ATTENUATION" VS FREQUENCY




SPL - DB
100
90
80
70
60
50
40
30
20
10
0
HS=HO= 10.0 m.
100
1000
FREQUENCY HZ

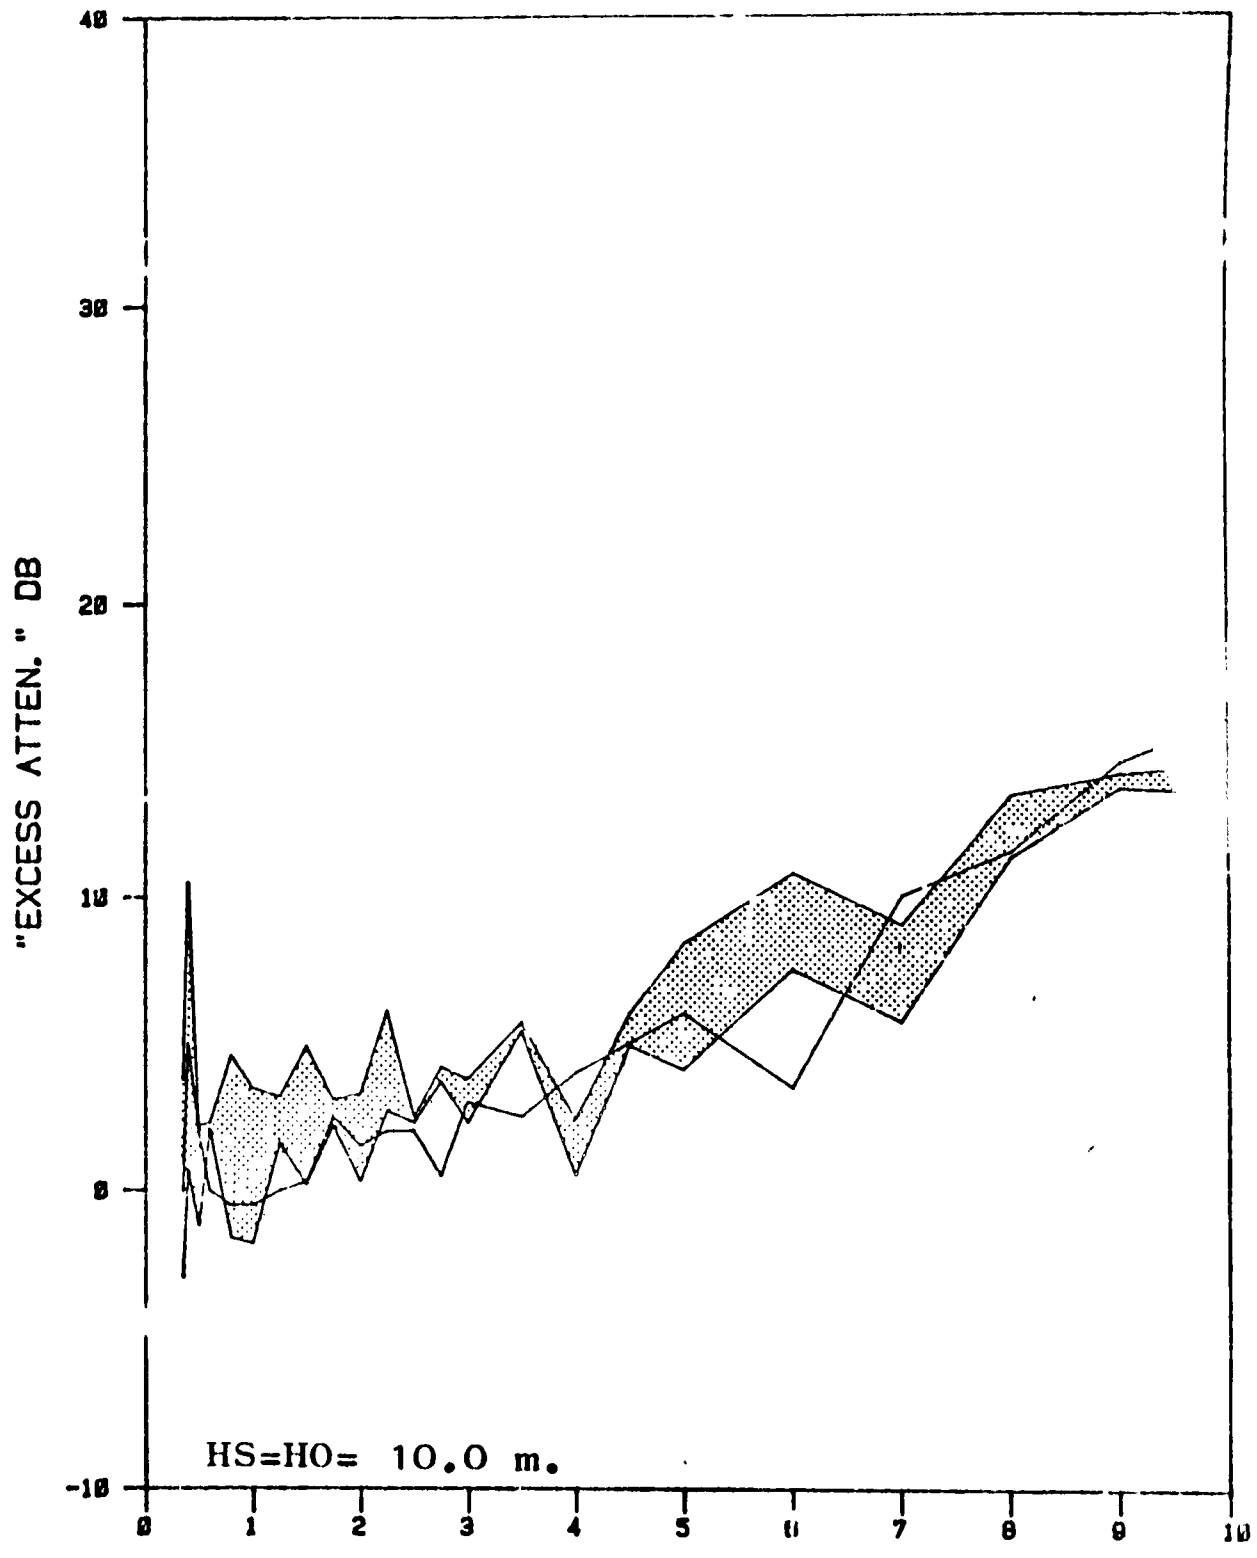
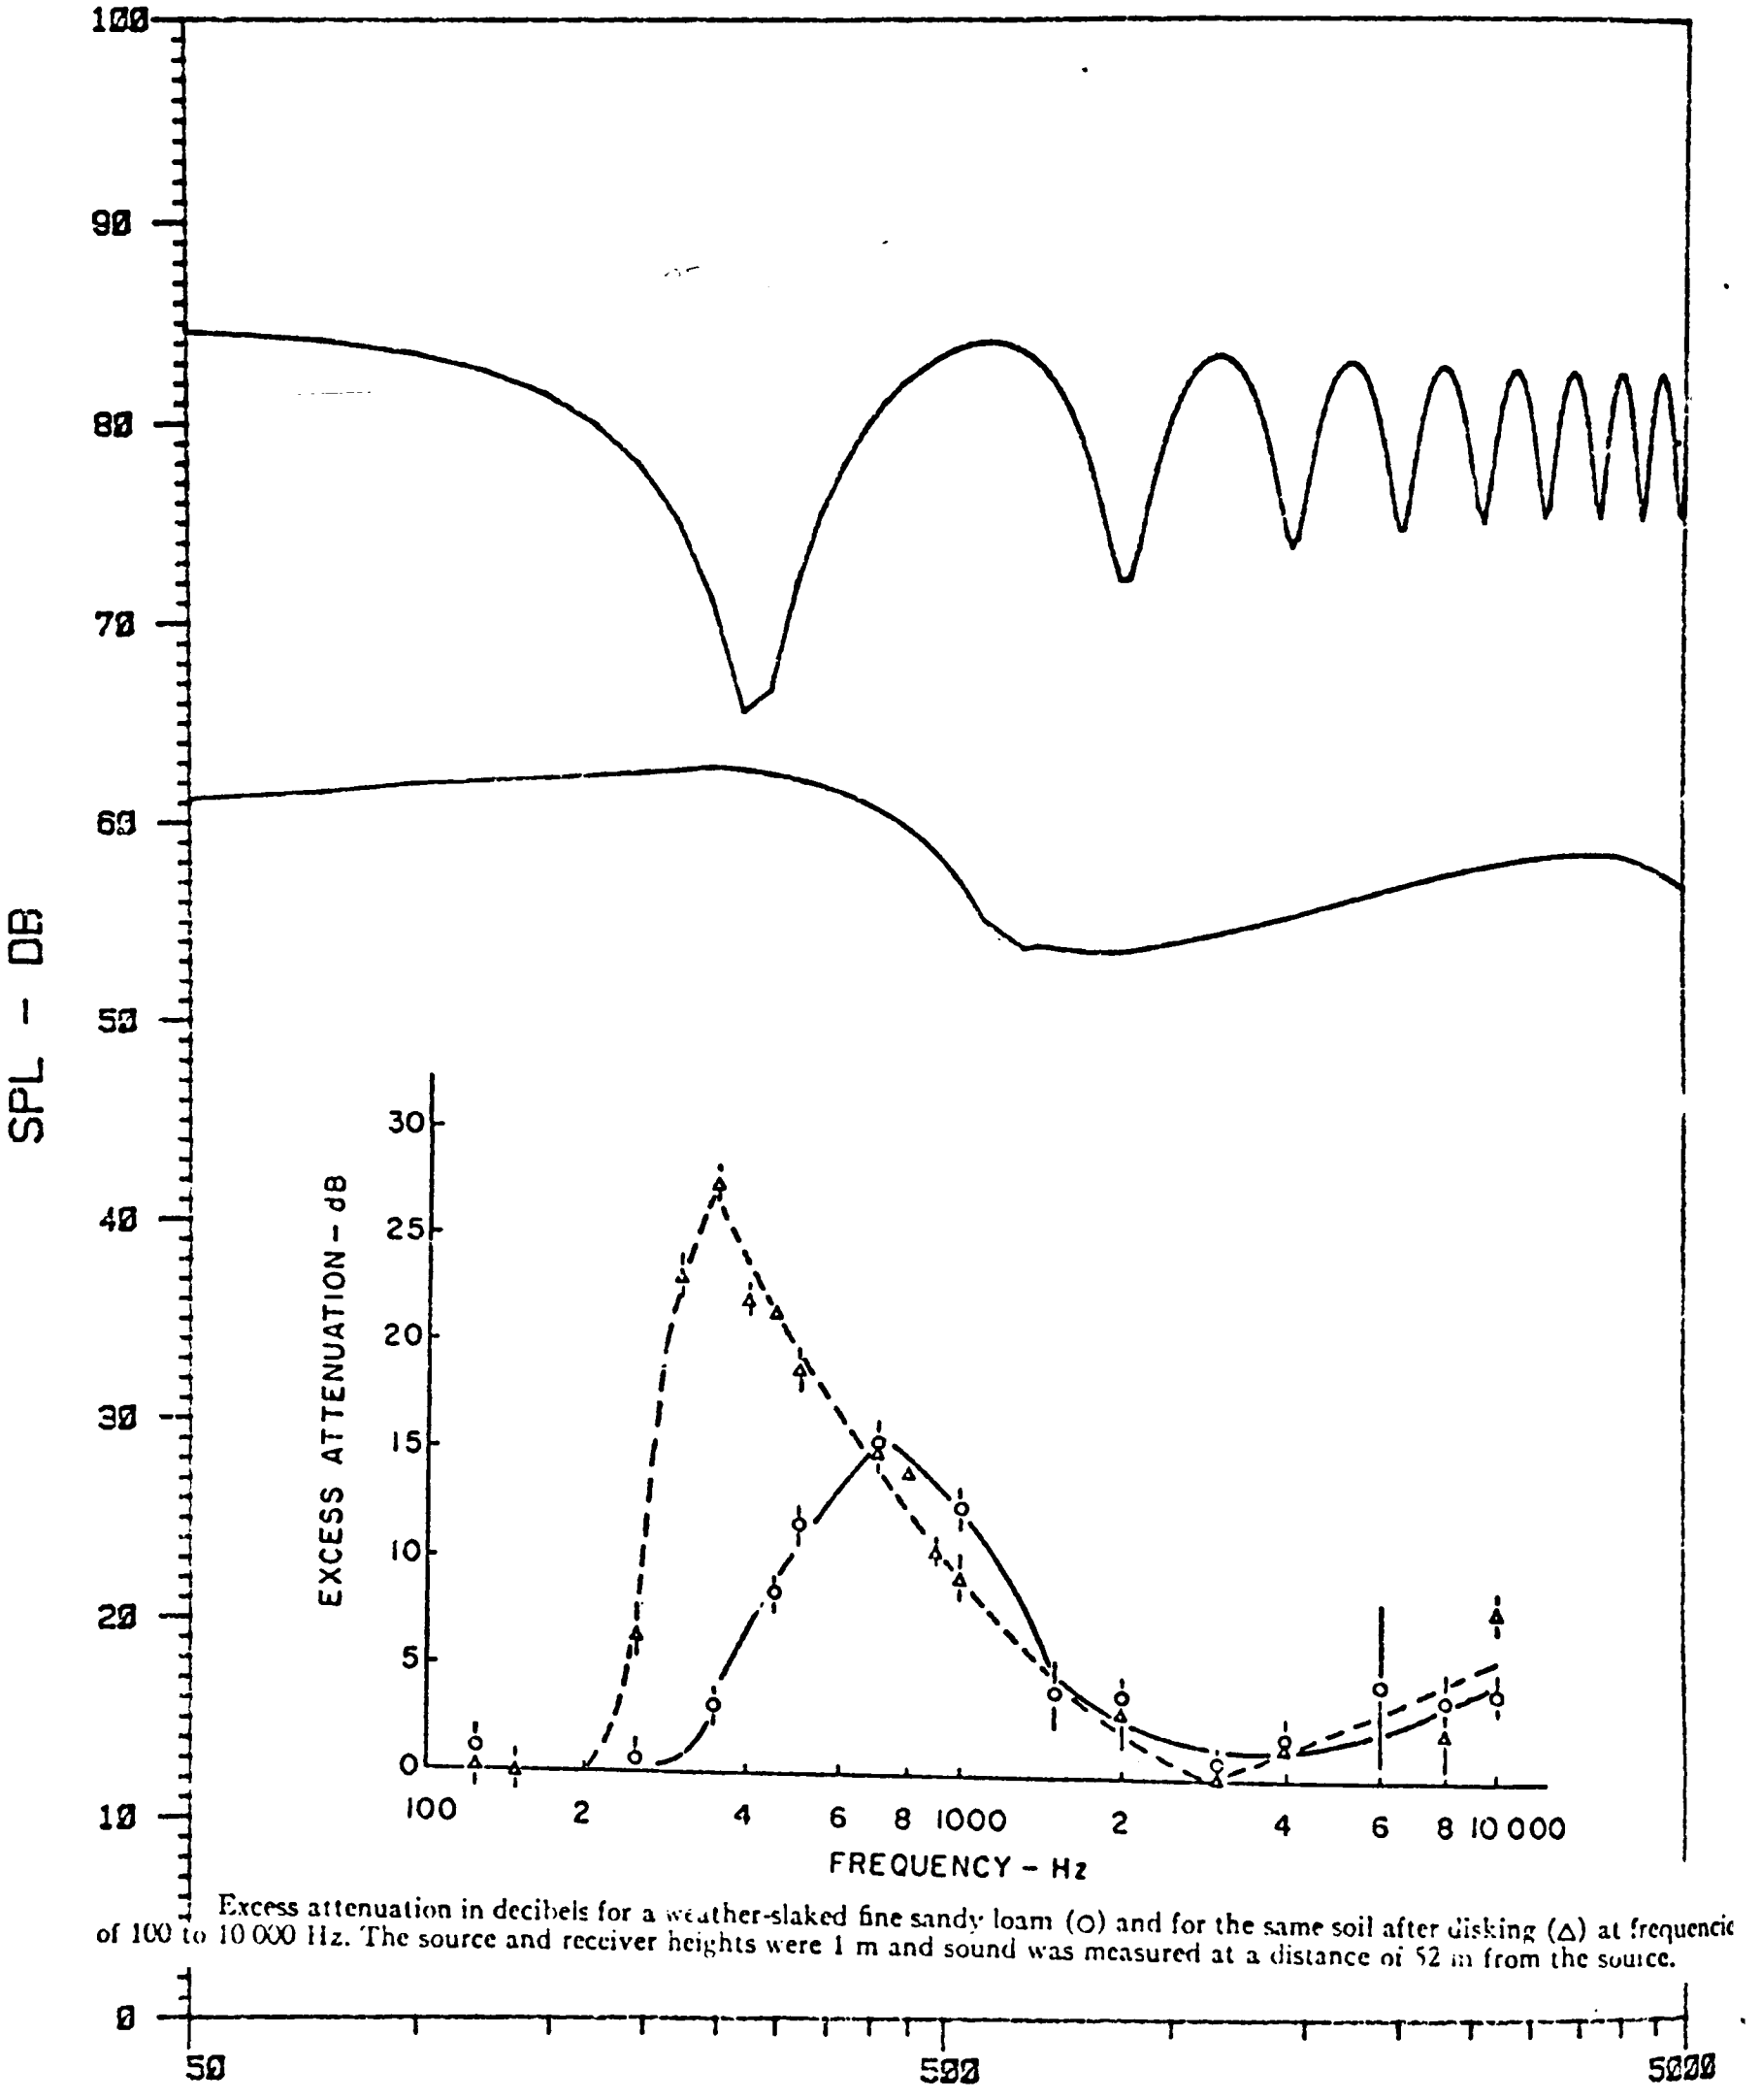

"EXCESS ATTEN." DB
40
30
20
10
0
-10
HS=HO= 10.0 m.
0
1
2
3
4
5
6
7
8
9
10
FREQUENCY KHZ

Fig.3. Cont.

SOUND PRESSURE LEVEL VS FREQUENCY



Excess attenuation in decibels for a weather-slaked fine sandy loam (O) and for the same soil after disking (Δ) at frequencie of 100 to 10 000 Hz. The source and receiver heights were 1 m and sound was measured at a distance of 52 m from the source.

Fig.4. SPLs predicted by equation (2) and 'excess attenuation' measured by Aylor over weather-slaked sandy loam for a source-receiver separtion of 3 and 52 m., and source-receiver height 1.0 m.

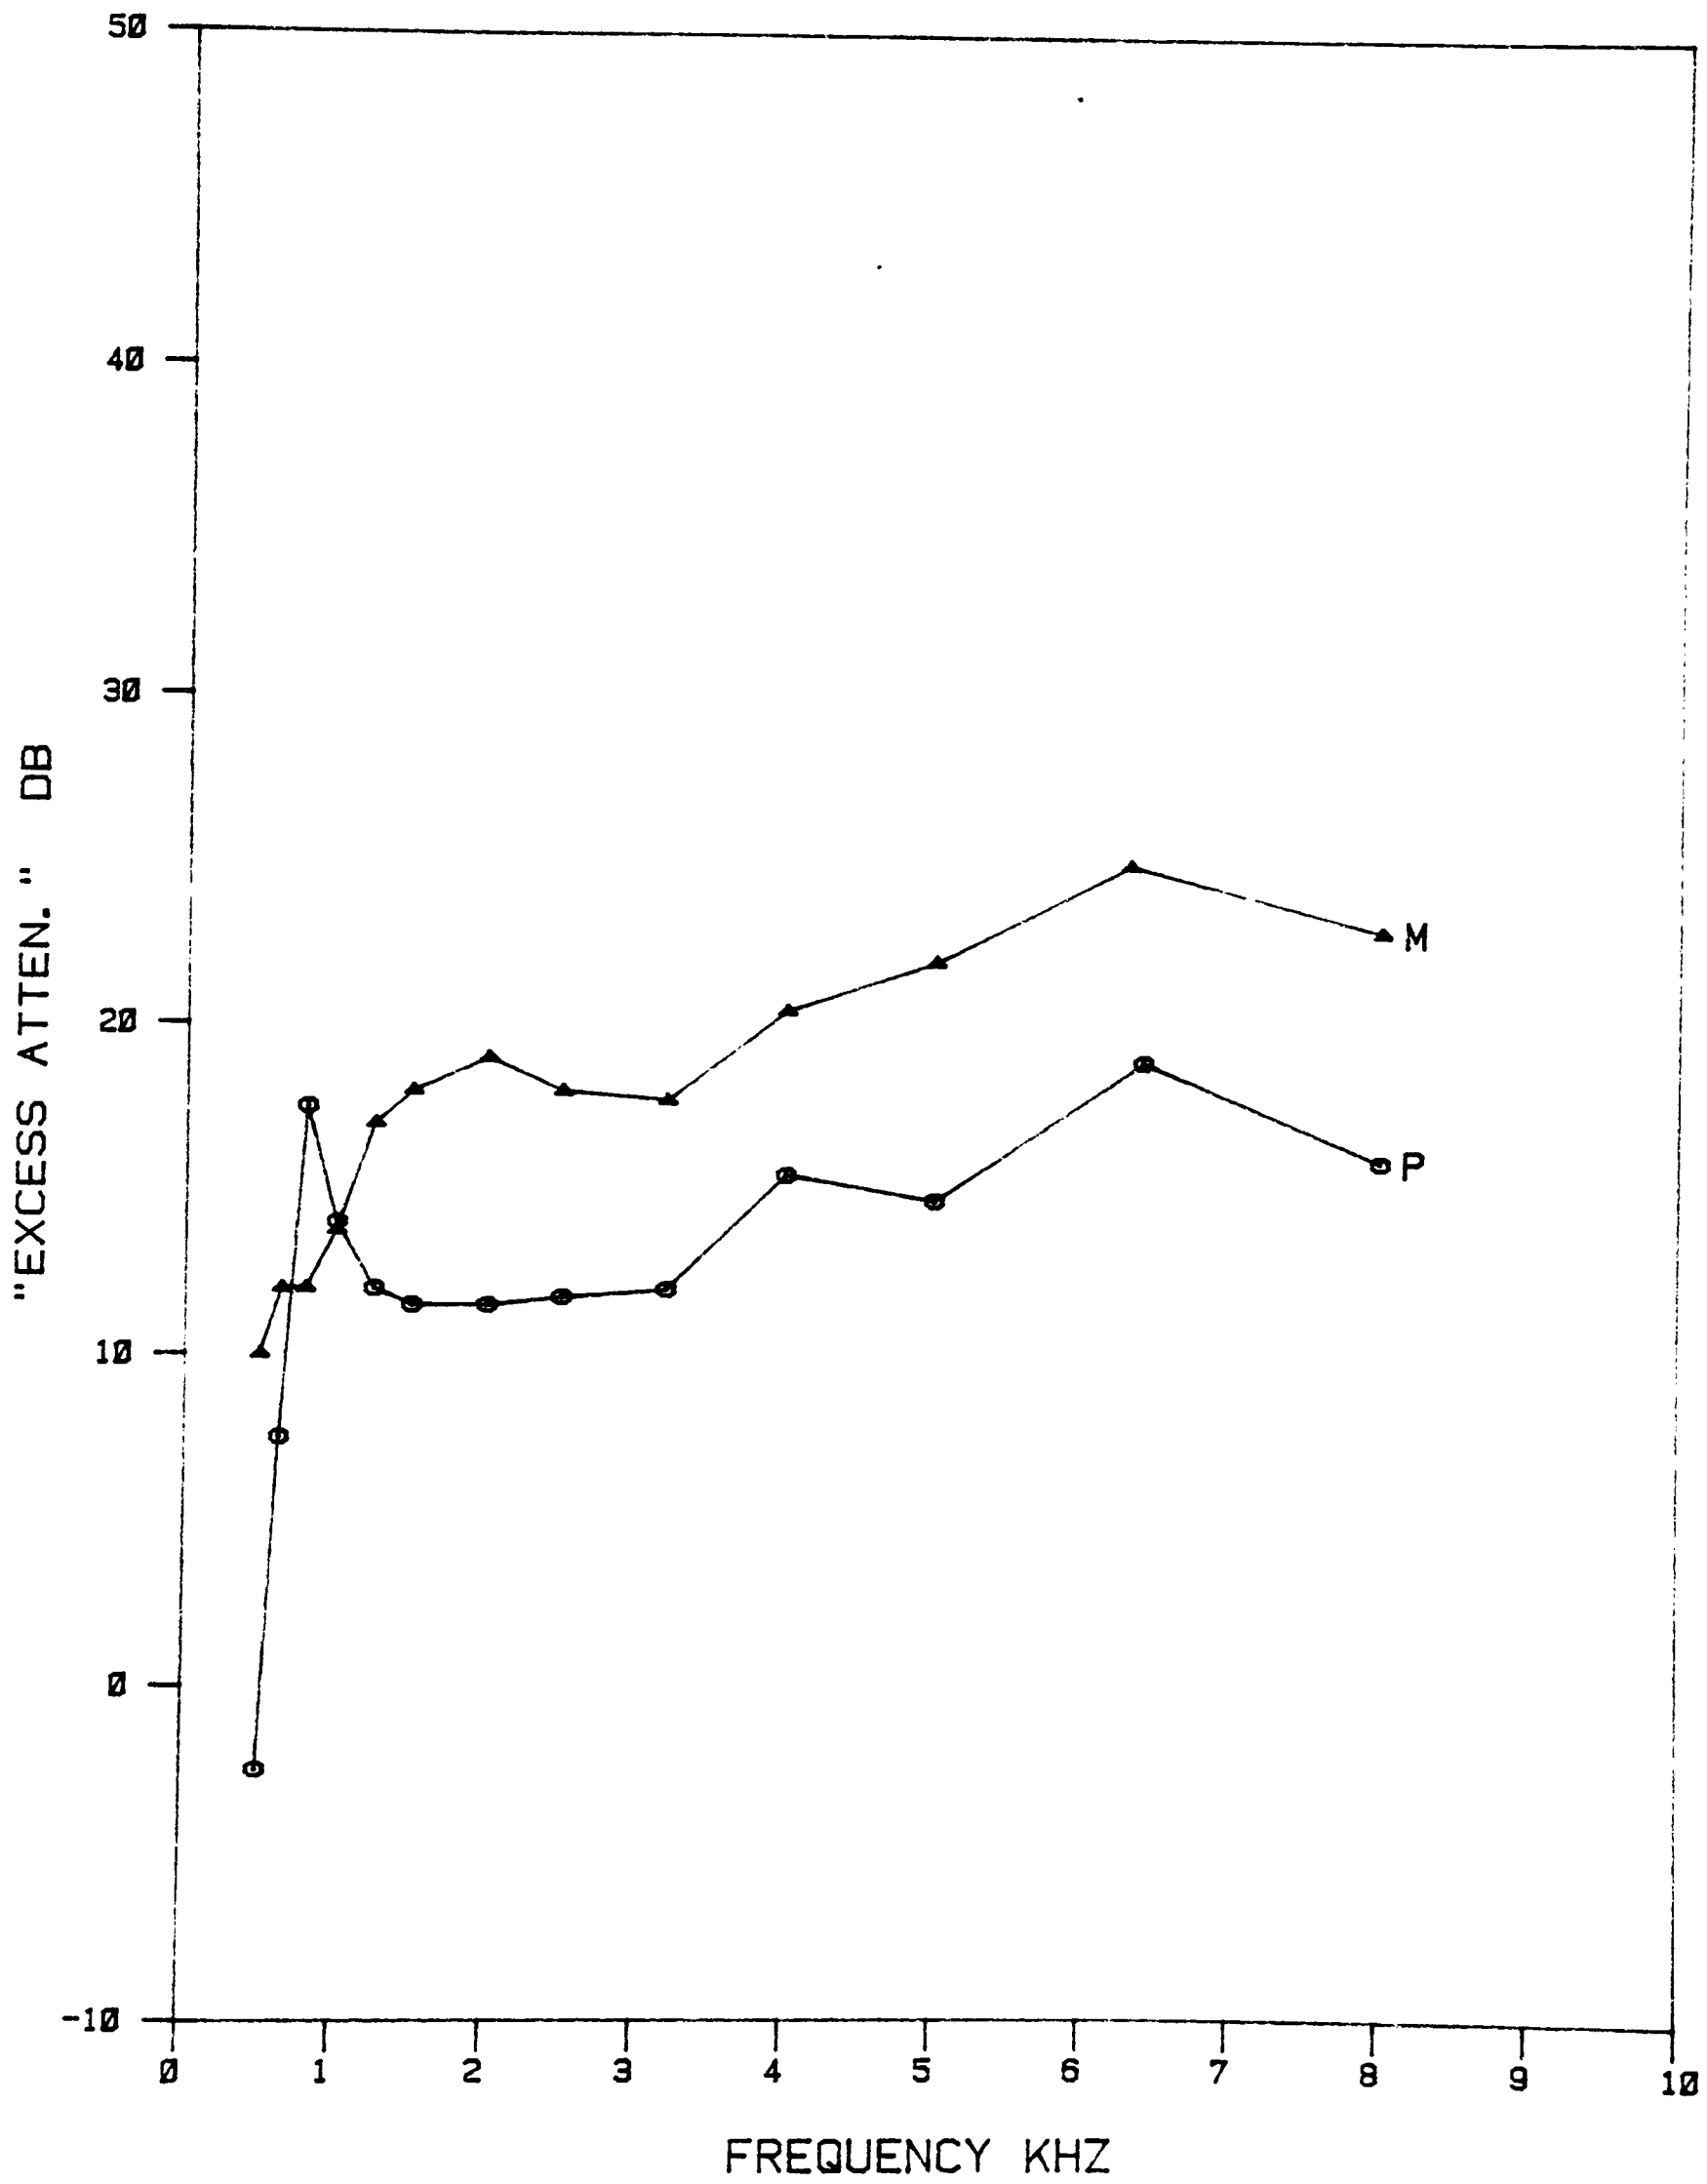


Fig.5. 'Excess attenuation' calculated from equation (4) modified to allow for absorption by vegitation and as measured by Morton for source-receiver height 0.3 m and receivers 7.62 and 91.44 m. from the source.

SOUND PRESSURE LEVEL VS FREQUENCY

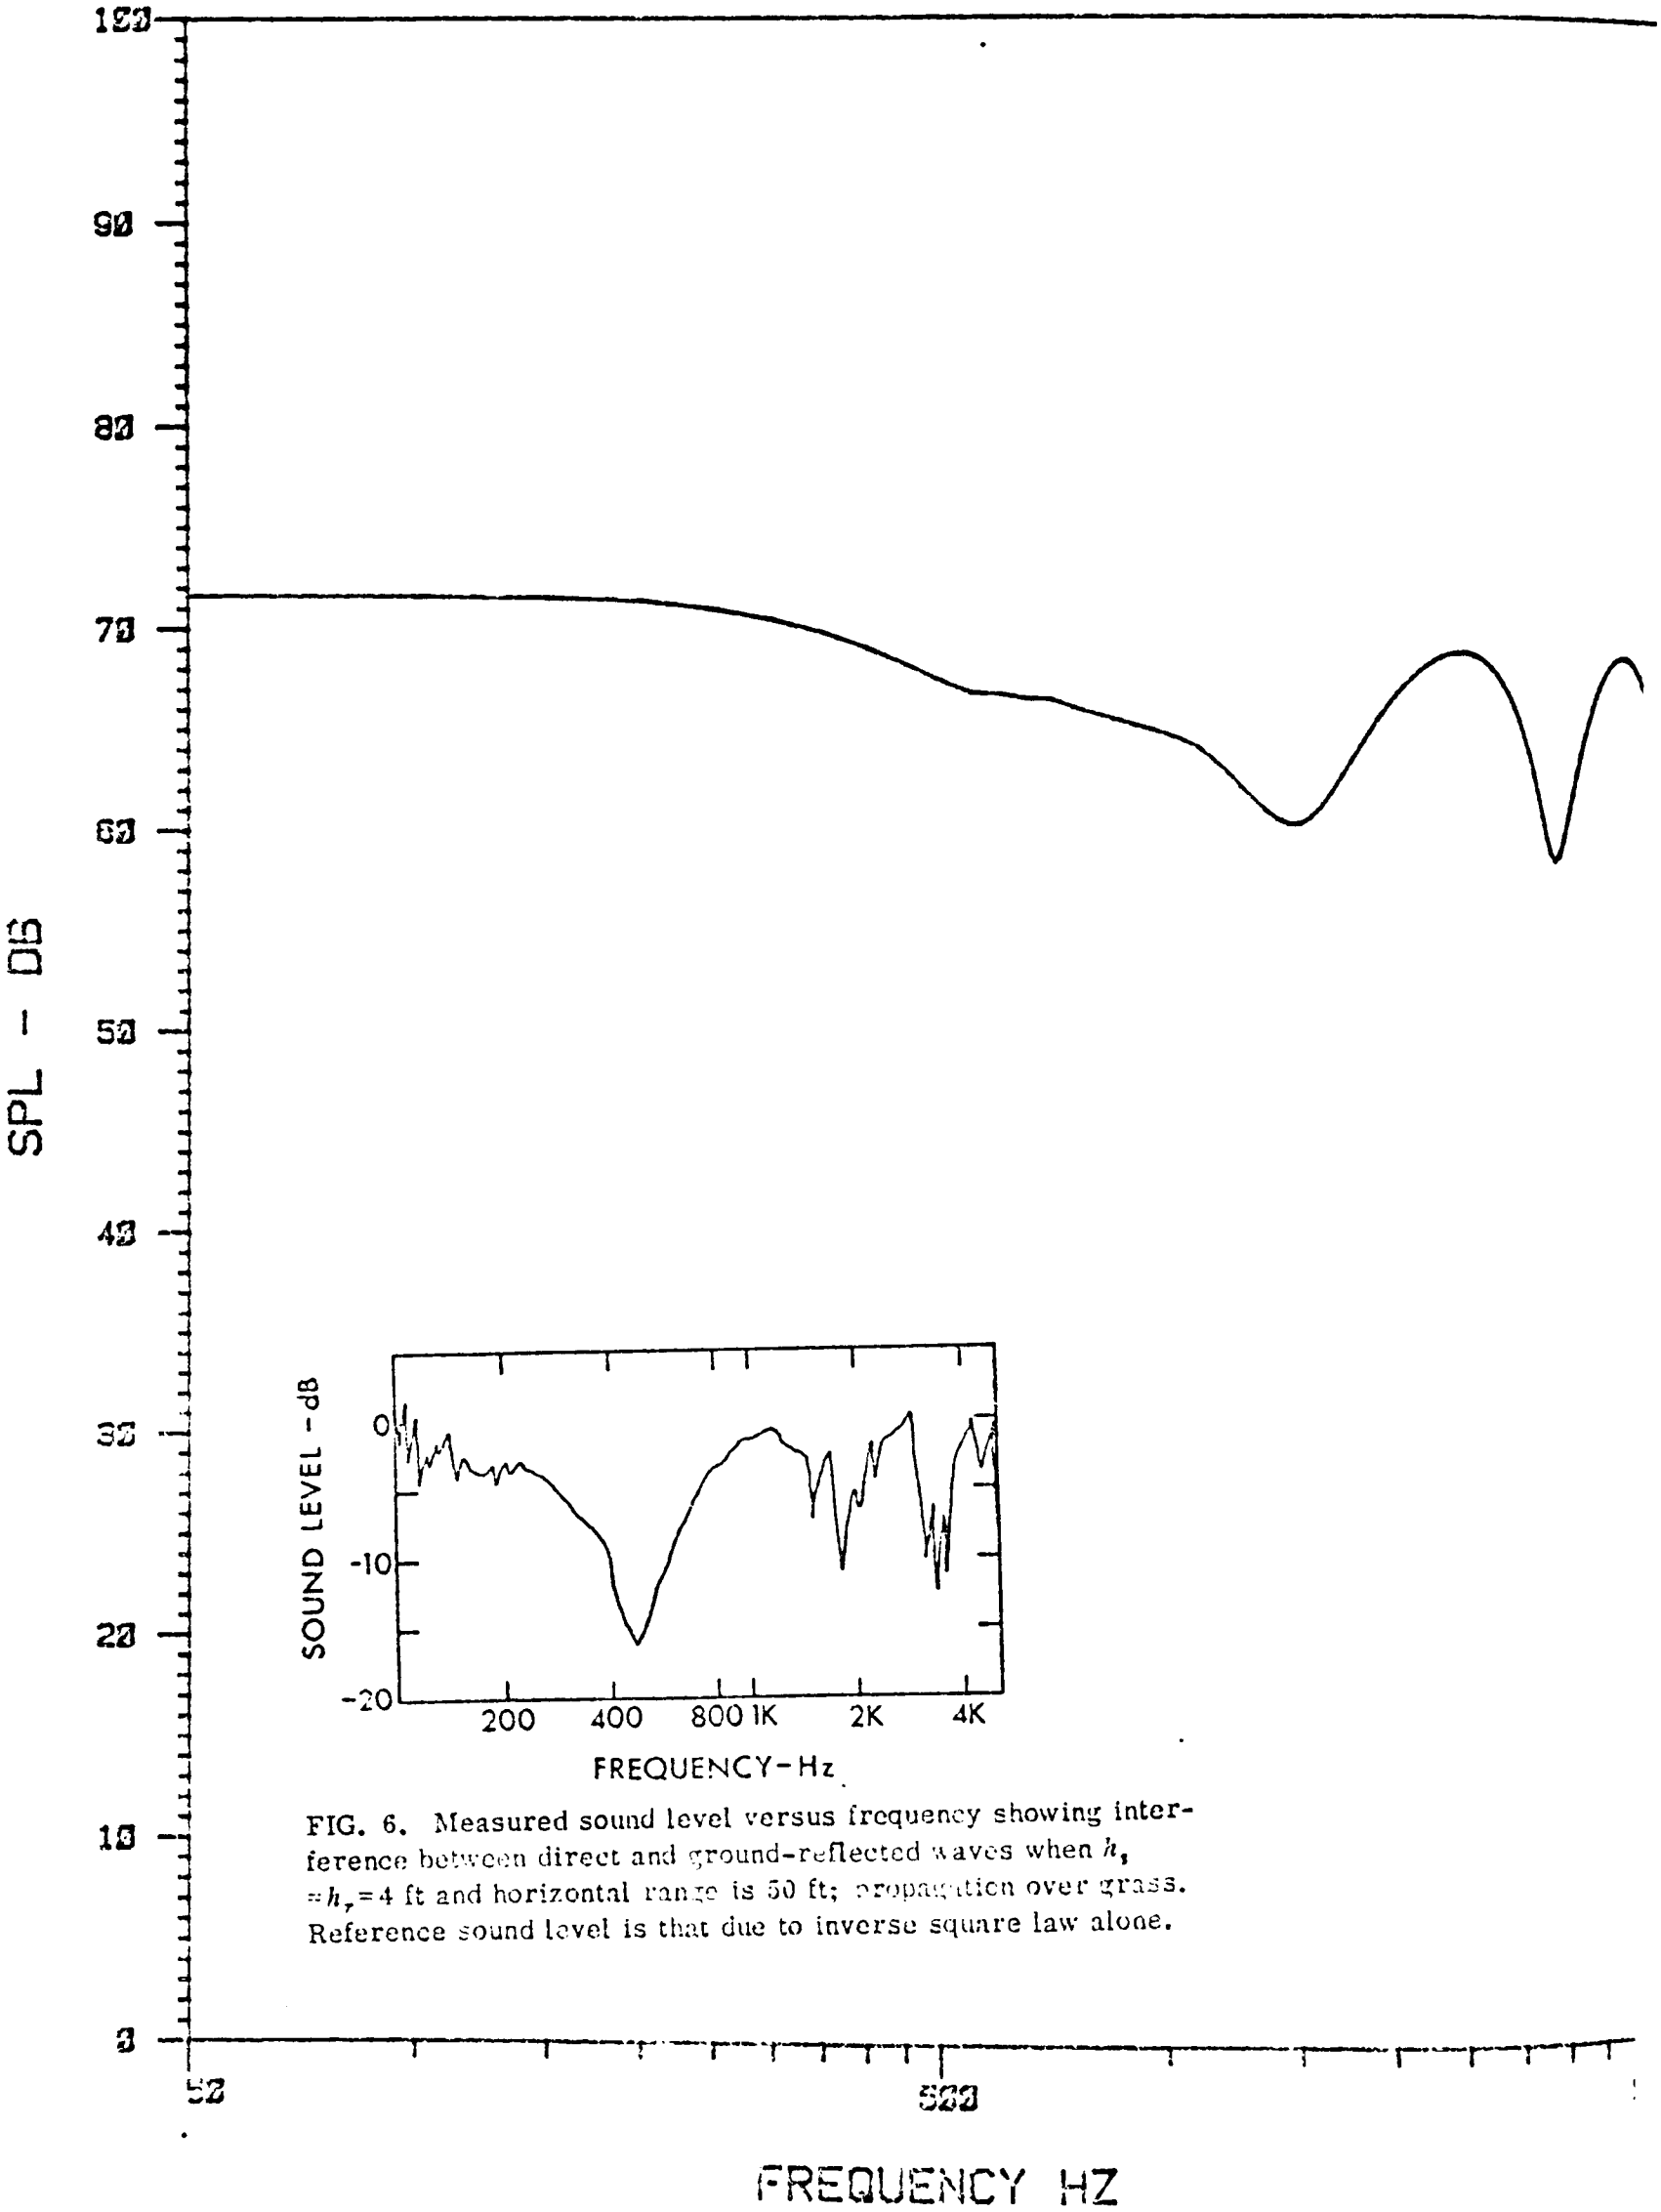


FIG. 6. Measured sound level versus frequency showing interference between direct and ground-reflected waves when $h_s = h_r = 4$ ft and horizontal range is 50 ft; propagation over grass. Reference sound level is that due to inverse square law alone.

Fig.6. SPLs over grassland predicted by equation (2) and measured by Embleton et al for source-receiver separation 15.24 m and height 1.2 m.

SOUND PRESSURE LEVEL VS FREQUENCY

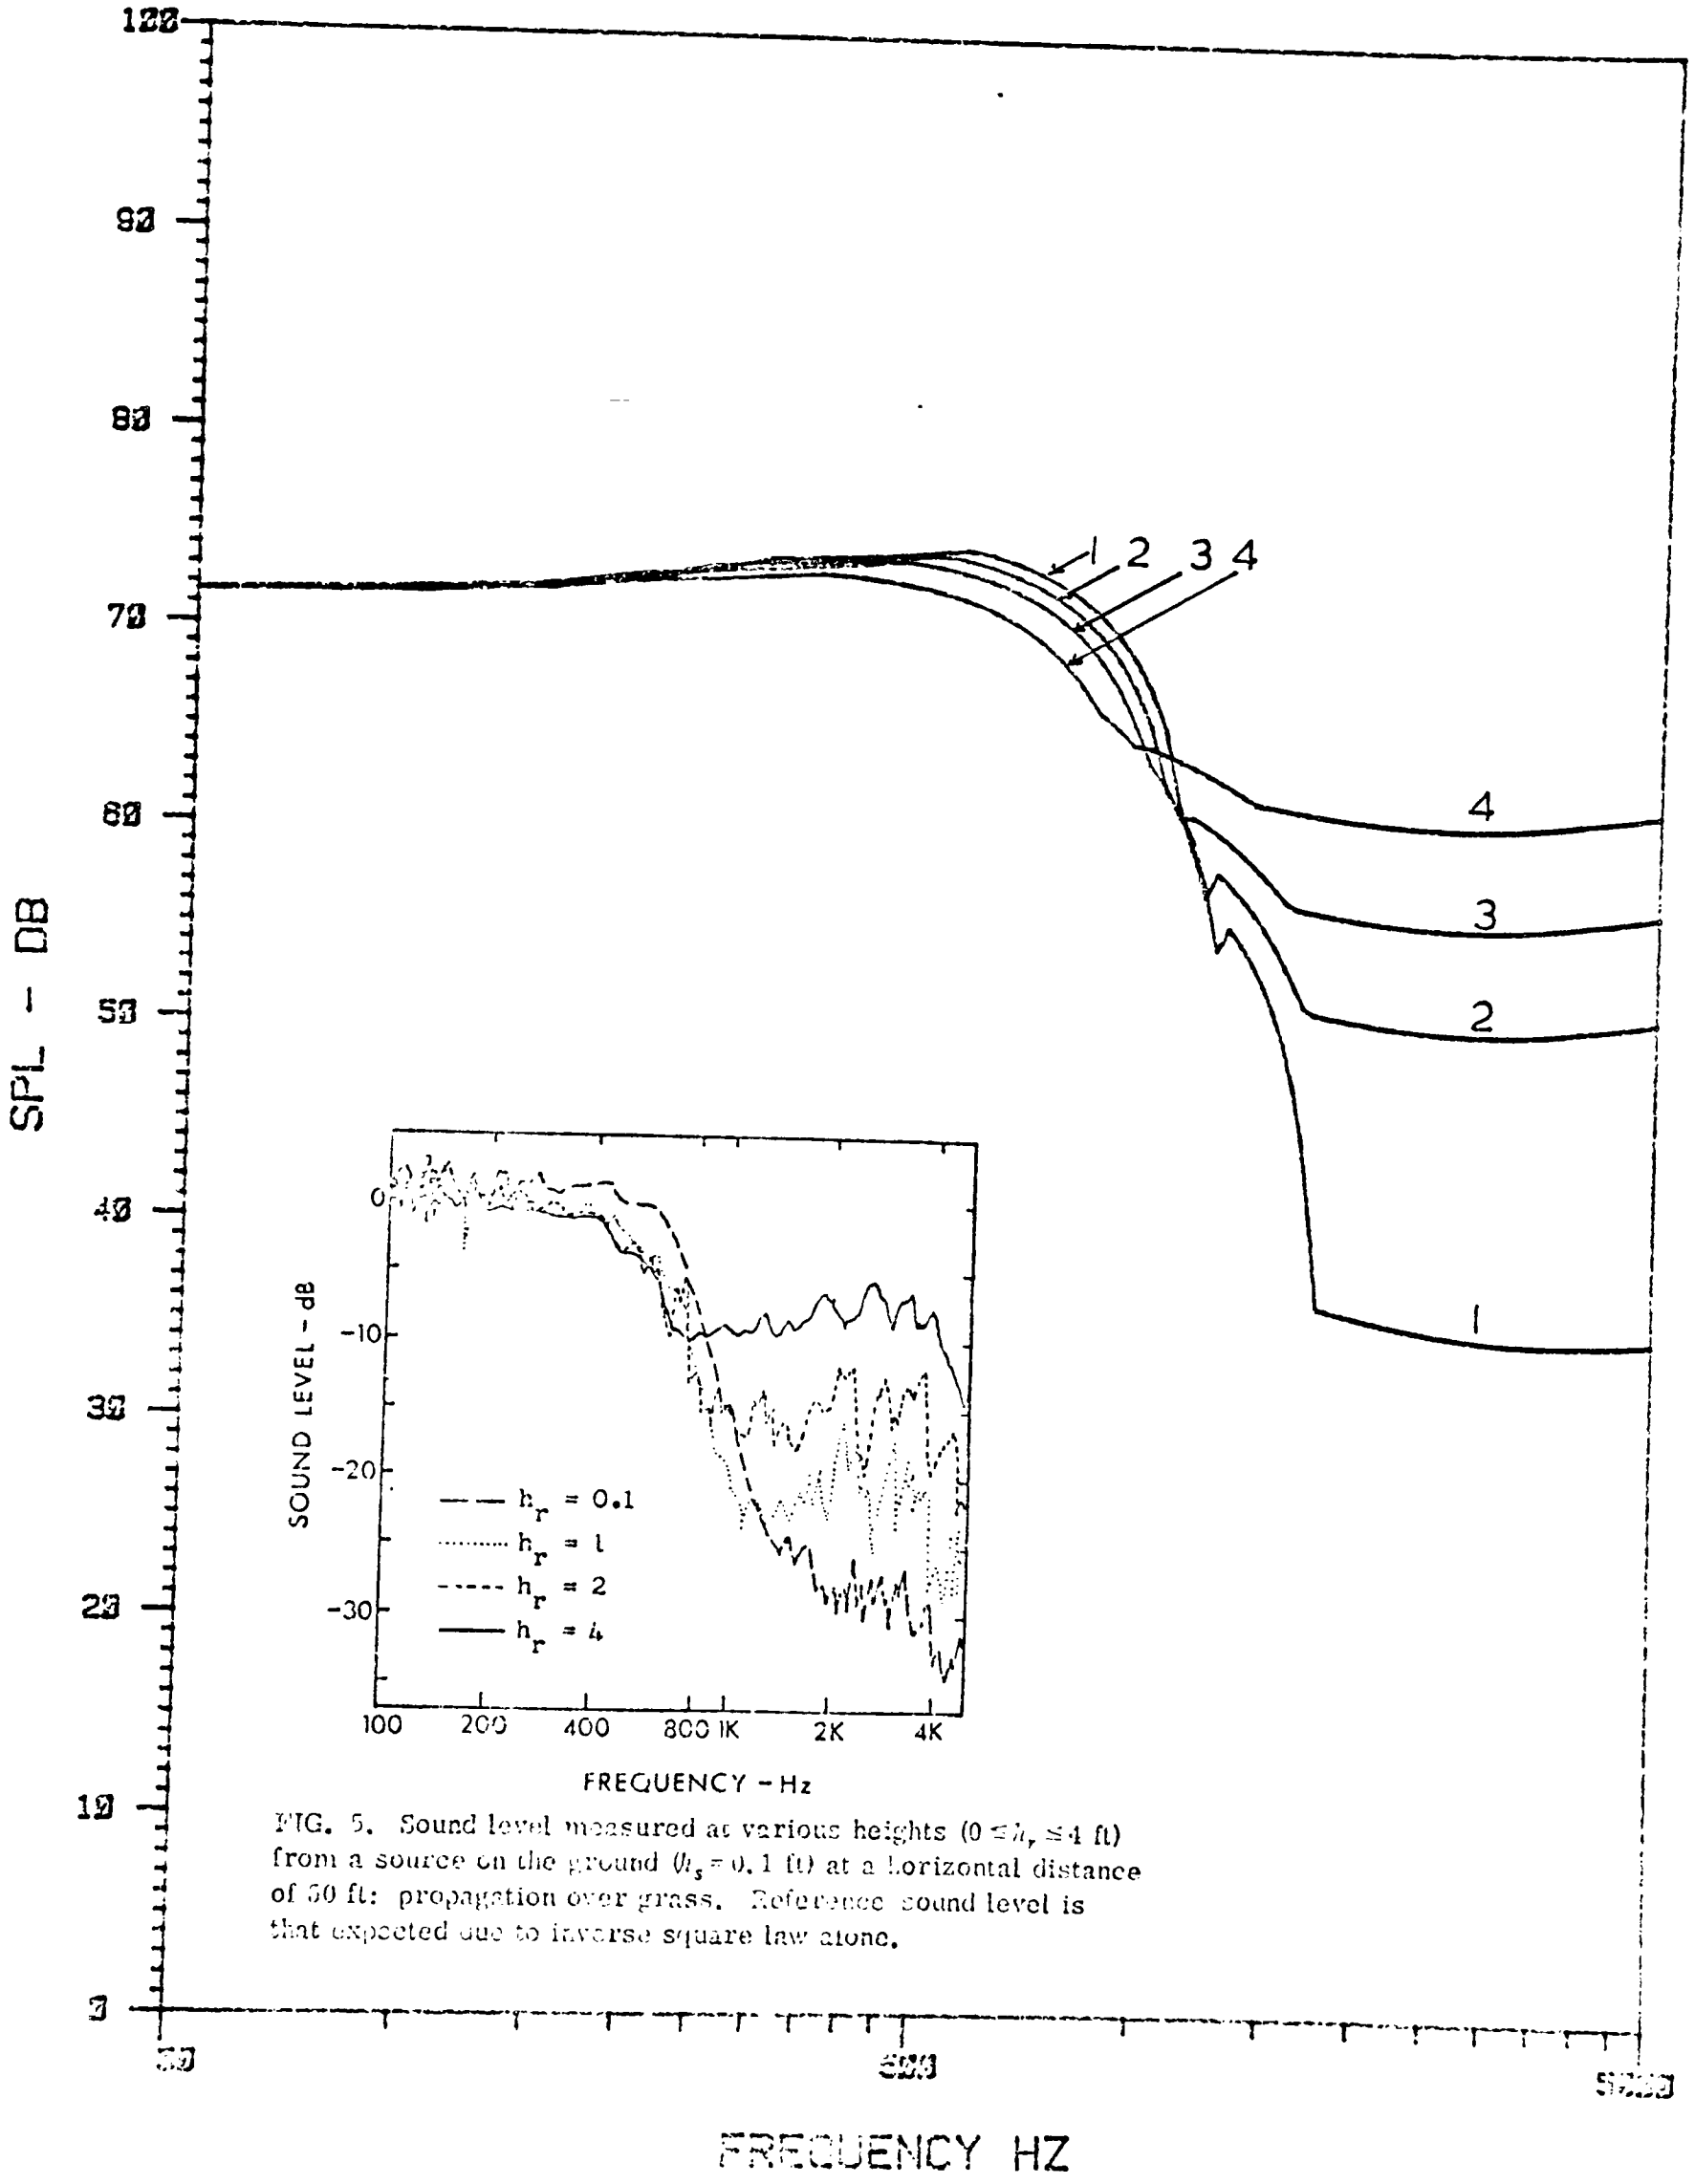


FIG. 5. Sound level measured at various heights ($0 \leq h_r \leq 4$ ft) from a source on the ground ($h_s = 0.1$ ft) at a horizontal distance of 50 ft: propagation over grass. Reference sound level is that expected due to inverse square law alone.

Fig.7. SPLs over grassland predicted by equation (2) and measured by Embleton et al. Source-receiver separation 15.24 m., source height 30.5 mm., and receiver heights (1) 30.5, (2) 305, (3) 610 and (4) 1220 mm.

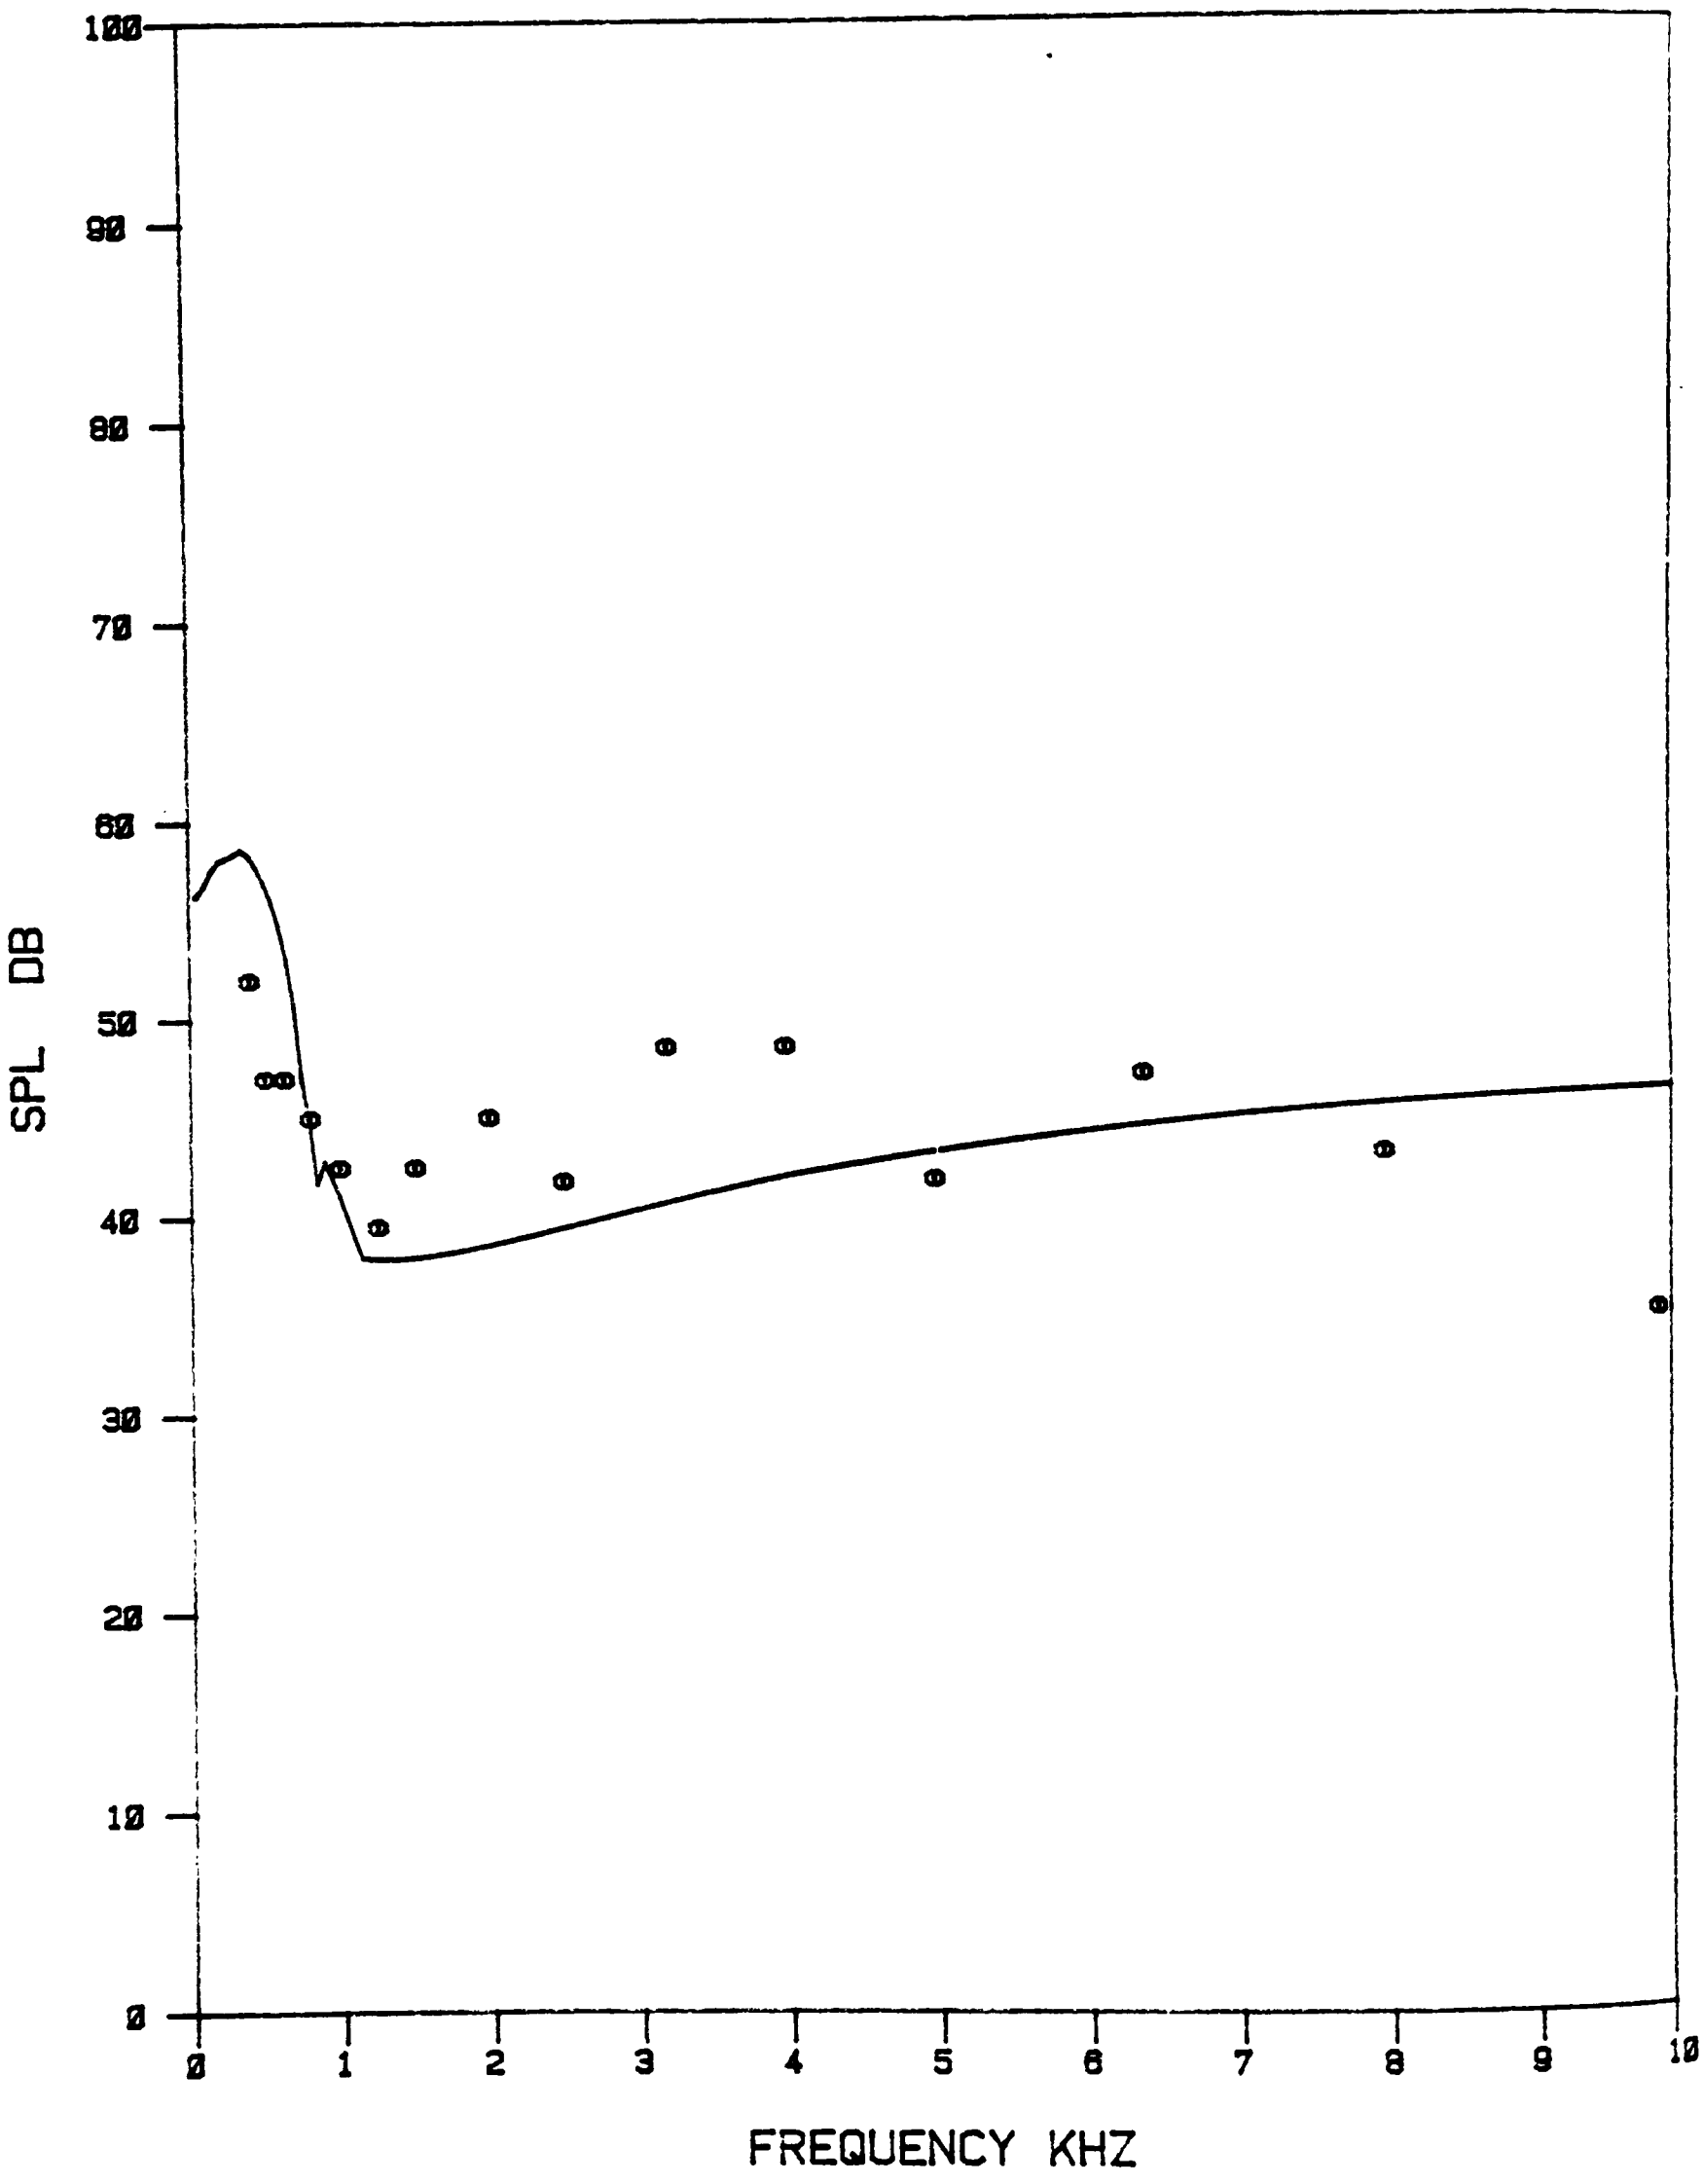


Fig.8. SPLs predicted by equation (2) and measured by Jilka and Leisler for source-receiver separation 16 m and height 0.2 m.

Appendix 4: ACCOUSTIC PROPERTIES OF THE ENVIRONMENT AS DETERMINANTS OF BIRDS' SONG PATTERNS

Several empirical studies have recently explored the possibility that animal vocalizations which are designed to maximize transmission distances are shaped by acoustic properties of the environment (e.g. Morton, Amer. Natur., 109, 17-34, 1975; Marten & Marler, Behav. Ecol. Sociobiol., 2, 271-290, 1977; Marten et al., Behav. Ecol. Sociobiol., 2, 291-302, 1977). Morton (1975) and Marten and Marler (1977) found that acoustic attenuation was particularly low for certain frequencies and Morton referred to these frequencies as the 'sound window'. This he considered an important factor in the evolution of animal acoustic signals. In this paper we attempt to explain its physical basis with a model that takes into account atmospheric attenuation and interference between the direct sound wave and the wave reflected from the ground.

In the model the signaller is assumed to be an omnidirectional source vocalizing in pure tones above a flat, unobstructed, locally reacting plane with no wind present. For the receiver there are two sound waves from any one signaller: the direct wave from the source and the reflected wave whose amplitude and phase depend upon the acoustic impedance of the ground.

If the signaller and receiver are both at height H, a distance D apart, with the sender vocalizing with sound power W, then the sound pressure heard by the receiver is (on plane wave theory with no air attenuation):-

$$\text{sound pressure} = -j\,\frac{\rho c W}{4\pi}\left[\frac{\exp(jkD)}{D} + \frac{\cos\phi-\beta}{\cos\phi+\beta}\cdot\frac{\exp(jk(D^2+4N^2)^{\frac{1}{2}})}{(D^2+4N^2)^{\frac{1}{2}}}\right] \qquad (1)$$

where $j = \sqrt{-1}$, ρc is the acoustic impedance of the air, k is the wave number of the acoustic signal, ϕ is the angle the indirect sound wave makes with the vertical and β is a function of the acoustic impedance of the ground which was measured on grassland using a Bruel and Kjaer type 4002 standing wave apparatus.

The value of β was ($\beta = A - jB$):

$A = 5.25 \times 10^{-3} + 0.134 \times f - 0.024 \times f^2 + 0.00129 \times f^3$

$B = 1.92 \times 10^{-2} + 0.125 \times f + 0.0224 \times f^2 - 0.00121 \times f^3$

f in kHz

With air attenuation included the Sound Pressure Level (SPL) due to the source may be calculated at any point in the sound field (Fig. 1).

Marten and Marler subtracted SPLs measured at 2.5 and 102.5 m to determine excess attenuation (EA) which they defined as the difference between the measured attenuation of the sound wave and the attenuation expected due to spherical expansion alone. Using equation (1) we calculated the expected SPLs at 2.5 m (SPL_1) and at 102.5 m (SPL_2) for a sound source of arbitrary power. Excess attenuation was determined from

$$EA = (SPL_1 - SPL_2) - 32.3 \quad dB \qquad (2)$$

where 32.3 dB results from spherical spread of the sound wave. This EA is the difference between the relevant SPL curves in Fig. 1.

Fig. 2 compares our predictions for EA with Marten and Marler's data. Both the trends and quantitative agreement are quite close; our predictions are generally lower because the model does not include wind attenuation or other such factors. The model and field measurements both indicate that attenuation is very strong close to the ground and decreases sharply with height up to about 1 m, after which further changes are negligible. They also agree that EA does not decrease monotonically with frequency; it decreases up to a certain frequency, which is dependent upon height and separation, then increases again. This minimum value for EA is Morton's sound window. However, our model indicates that sound windows should occur whenever interference between the direct and reflected waves at the farther receiver is constructive, provided air attenuation is not too great. Thus they depend upon both the relative microphone positions and the actual attentuation between source and receiver and in this sense are dependent on the method employed and may have little relevance for animal

vocalizations.

The model's predictions may have further implications for animal vocalizations. A specific pattern of peaks and troughs of SPLs exists for each source-receiver distance and height; at greater heights and shorter source-receiver distances the pattern is more pronounced (Fig. 1). Because animals are unlikely to be able to maximize the carrying power of their vocalizations by using solely frequencies which coincide with SPL maxima (they will know neither the distance to their receiver nor its height) the bandwidth of the signal should be sufficiently wide to avoid falling in a minimum in the interference pattern. Thus, all other things being equal, animals vocalizing at greater heights would find it advantageous to use signals of wider frequency span.

We tested this prediction by comparing the vocalization frequency span of the bird species listed by Morton (pers. comm.) whose singing perch heights were clearly defined. For 14 species singing from the ground the average span was 0.77 kHz (SD = 0.77 kHz), while for ten canopy species the span was 2.56 kHz (SD = 2.57 kHz). This difference is statistically significant ($p < 0.05$). This prediction, however, is not exclusive for our model. Heavy attenuation near the ground may also favour narrow band signals if concentrating the sound energy is important.

John Krebs and Henry Bennett-Clark commented on the manuscript. Gene Morton kindly let us analyze his unpublished data.

John Roberts *

Alejandro Kacelnik **

Malcolm L. Hunter, Jr. **

* I.E.S.T., Polytechnic of the South Bank,
London, U.K.

** Department of Zoology,
Oxford OX1 3PS, U.K.

Figure Captions

Fig. 1. Sound pressure levels (SPLs) predicted by the model for various source-receiver heights and separations. Temperature 0°C: relative humidity 90%.

Fig. 2. Excess attenuation from Marten and Marler (1977) and as predicted by the model. Fig. 2a shows deviation of mean EA at each height from grand mean; Fig. 2b shows the same for frequency. Temperature 24°C: relative humidity 60%.

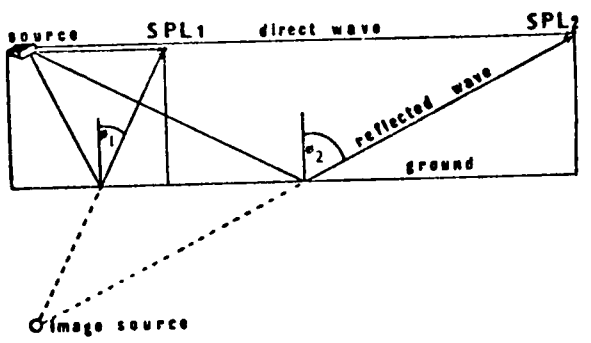
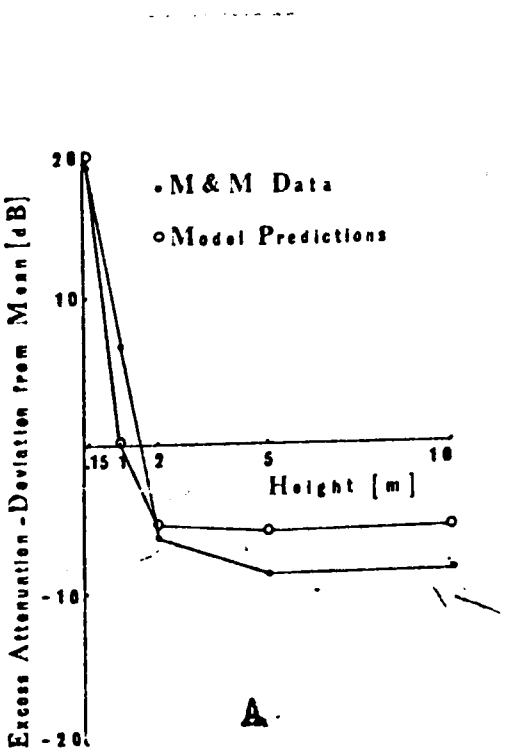


Fig. 1



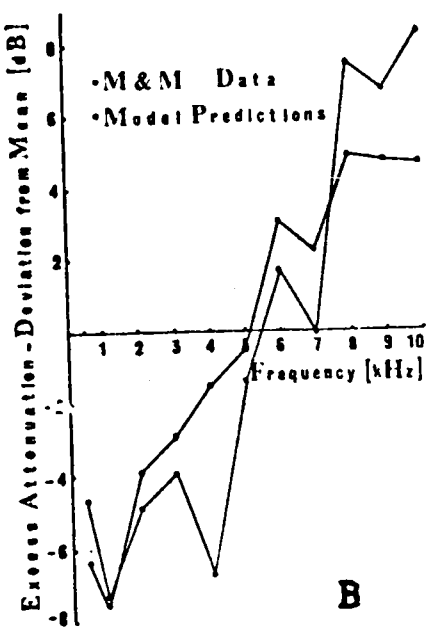


Fig 2

The effect of hunger on vigilance during prey handling in the great tit (Parus maj.:

When a foraging great tit (Parus major) finds a prey, it normally picks it up in its bill, moves to a suitable place to handle it, and starts pecking or tearing the prey while holding it with its feet. The pecking is frequently interrupted by the bird raising its head and moving it sideways. It would appear that this looking-around behaviour is performed in order to scan the surroundings for predators. Eventually, the bird either eats the last bit of food or drops the inedible portions such as the husks or wings in order to begin searching for a new food item. The interval between picking up the prey (or in some cases locating it) and starting to search for another one is termed handling time (H).

Handling time plays an important role in optimal foraging models. For example, diet selection has been predicted from the ratio of energy value to handling time (e.g. Krebs et al. 1977). Conversely, handling time itself has been predicted on the basis of the diminishing returns in rate of calories obtained from a single prey (Cook and Cockerell 1978). In all of these studies, it was implicitly assumed that handling time is independent of hunger since, according to the models, an optimally foraging animal would minimise H, and thus maximise energy intake per unit time, for any given level of food availability.

Milinsky and Heller (1978), however, recently showed that foraging strategy in sticklebacks is not only influenced by the rate of capture, but also by the presence of a predator. In general, they claimed that the optimal strategy for the stickleback lies somewhere between maximisation of food acquisition on the one hand and anti-predator behaviours on the other. The specific balance between these two behaviours will, of course, depend upon

the relative strength of the two tendencies (hunger and vigilance) at any given point in time (Sibly and McFarland 1976). Although Milinsky and Heller focussed on patch selection, their approach should also apply to other aspects of foraging such as handling time. If a certain amount of predator monitoring takes place while the prey is being handled, one would predict that H would remain constant only if the levels of risk from predators remain constant; when the risk of predators is not constant across conditions, H would compensate in some fashion.

We considered the same line of reasoning in terms of levels of hunger rather than predation risk. If an animal encountered two patches which ranked differently in terms of energy and predator risk, the optimal strategy would depend upon the combined risks (of starving and of being attacked). A very hungry animal, for instance, might increase its risk of attack from a predator in order to secure food and thereby reduce its risk of starvation; in contrast, a less hungry animal might choose a less risky feeding patch even though its net energy intake there was lower.

Handling time has indeed been shown to change during the course of feeding in birds (Smith and Dawkins 1971; Zachs and Falls 1978), but no functional explanation of this sort was considered. We hypothesised that as an animal gets less hungry, it is more likely to handle prey in a manner that allows for more anti-predator vigilance. Although the overall rate of food intake may decline, vigilance will increase since, as the animal nears its satiation point, it can afford to trade-off feeding efficiency for increased safety. Assuming, therefore, that the 'looking-around' pattern described above is, in fact, anti-predator behaviour (or vigilance), we devised a simple experiment in which we measured changes in handling time and vigilance during feeding.

We used two great tits as subjects. Each bird was tested four times under four deprivation levels: 45, 60, 75 and 90 minutes food deprived. The

sequence of sessions was such that all deprivation levels were tested in a random order before a replication was begun. In more formal terms, the experiment constituted a 2 x 4 x 4 random factorial design with subjects, deprivation levels, and replications as the factors.

Each test lasted 15 minutes and consisted of allowing the bird to obtain food (one fly pupa) by hopping onto a response perch (for a more detailed description of the apparatus and method, see Krebs et al. 1978). While the bird responded, the experimenter observed its behaviour from behind a one-way glass and recorded with an on-line computer device the occurrence and times for the following behaviours: hops onto the perch, picking up the prey, pecking at the prey, looking up, and dropping the prey.

The mean H time for fifths of the total number of prey eaten were first computed. We used this procedure because of the difference in the number of prey in the different sessions ( ). These values were then averaged over birds and replications (since those two main effects were not significant) to yield a five (of prey) by four (deprivation level) analysis matrix. As shown in Figure 1, H times did increase during feeding, thus confirming previous work (Smith and Dawkins 1971; Zachs and Falls 1978). Moreover, analysis of variance showed that fifths of total captures interacted with deprivation level indicating that the rate of increase of H times during feeding depended on the initial deprivation level - the greater the deprivation, the greater the rate of increase.

A similar analysis was done for the number of look-ups/prey and the average duration of looking-up (the means for the latter measure are shown in Figure 1). In both cases, the main effect of cumulative intake was significant thus supporting our hypothesis that looking-up occurs more frequently and lasts longer as a function of satiation. Similarly, for both measures, the interaction between satiation and initial deprivation was significant indicating that the rate of increase of vigilance during feeding depended on the initial eprivation - the greater the deprivation, the greater the rate of increase in vigilance during the feeding session.

Finally, we considered the looking-up times/prey as a proportion of the H time. The means, shown in Figure 1 (right ordinate), ranged from about 9 to 13%. An analysis of variance confirmed that the increase is significant. Vigilance not only increased during feeding, but it came to occupy a greater proportion of H time. In other words, the birds spent a greater proportion of their handling time performing anti-predator behaviours as they approached satiation.

We conclude from these results that (a) handling time does, indeed, increase with intake; (b) vigilance, as measured by the number and duration of look-ups, also increases during feeding; and (c) the latter, in part, accounts for the former. In other words, H time increases, in part, because the subjects are spending relatively more of the H time being vigilant.

The argument is clearly that animals have evolved to respond with respect to a number of different pressures. Energy per unit time and vigilance are merely two, albeit important, selective pressures that contribute to the observed behaviour pattern. Presumably animals seek to optimise not just one outcome such as energy, but rather a configuration of outcomes. The relative degree of optimisation for any given outcome then depends upon the configuration of the internal state and environmental conditions at that particular time. In other words, our study shows that the relative contribution of one activity (say vigilance) depends not simply upon the conditions that directly affect that activity (e.g. animals are more vigilant when the probability of predators is high, see Milinsky and Heller 1978), but also on conditions that affect other behaviours (i.e. animals are more vigilant when the risk of starvation permits).

Although these results support our hypothesis, we recognise that more detailed studies are required to specify the full relationship between hunger and vigilance. To this end, experiments in which both predation risk and hunger level are manipulated would be most important. In addition, however, we speculate

that more comprehensive and/or quantitative studies will provide not only greater accuracy but also will reveal further complexity. That is, foraging behaviour will be shown to be affected by many more parameters than just hunger and predation risk. Indeed, our own data suggest this point. As H increased with satiation, looking-up time increased, but 'eating time' per prey (defined as the duration of the uninterrupted pecking) and the time between capturing the prey and the beginning of eating also increased. Clearly the animals were becoming more vigilant as hunger decreased, but they were also serving other motives as indicated by this increase in these other components of H time.

Alex Kacelnik

University of Oxford

Roger M. Tarpy

Bucknell University

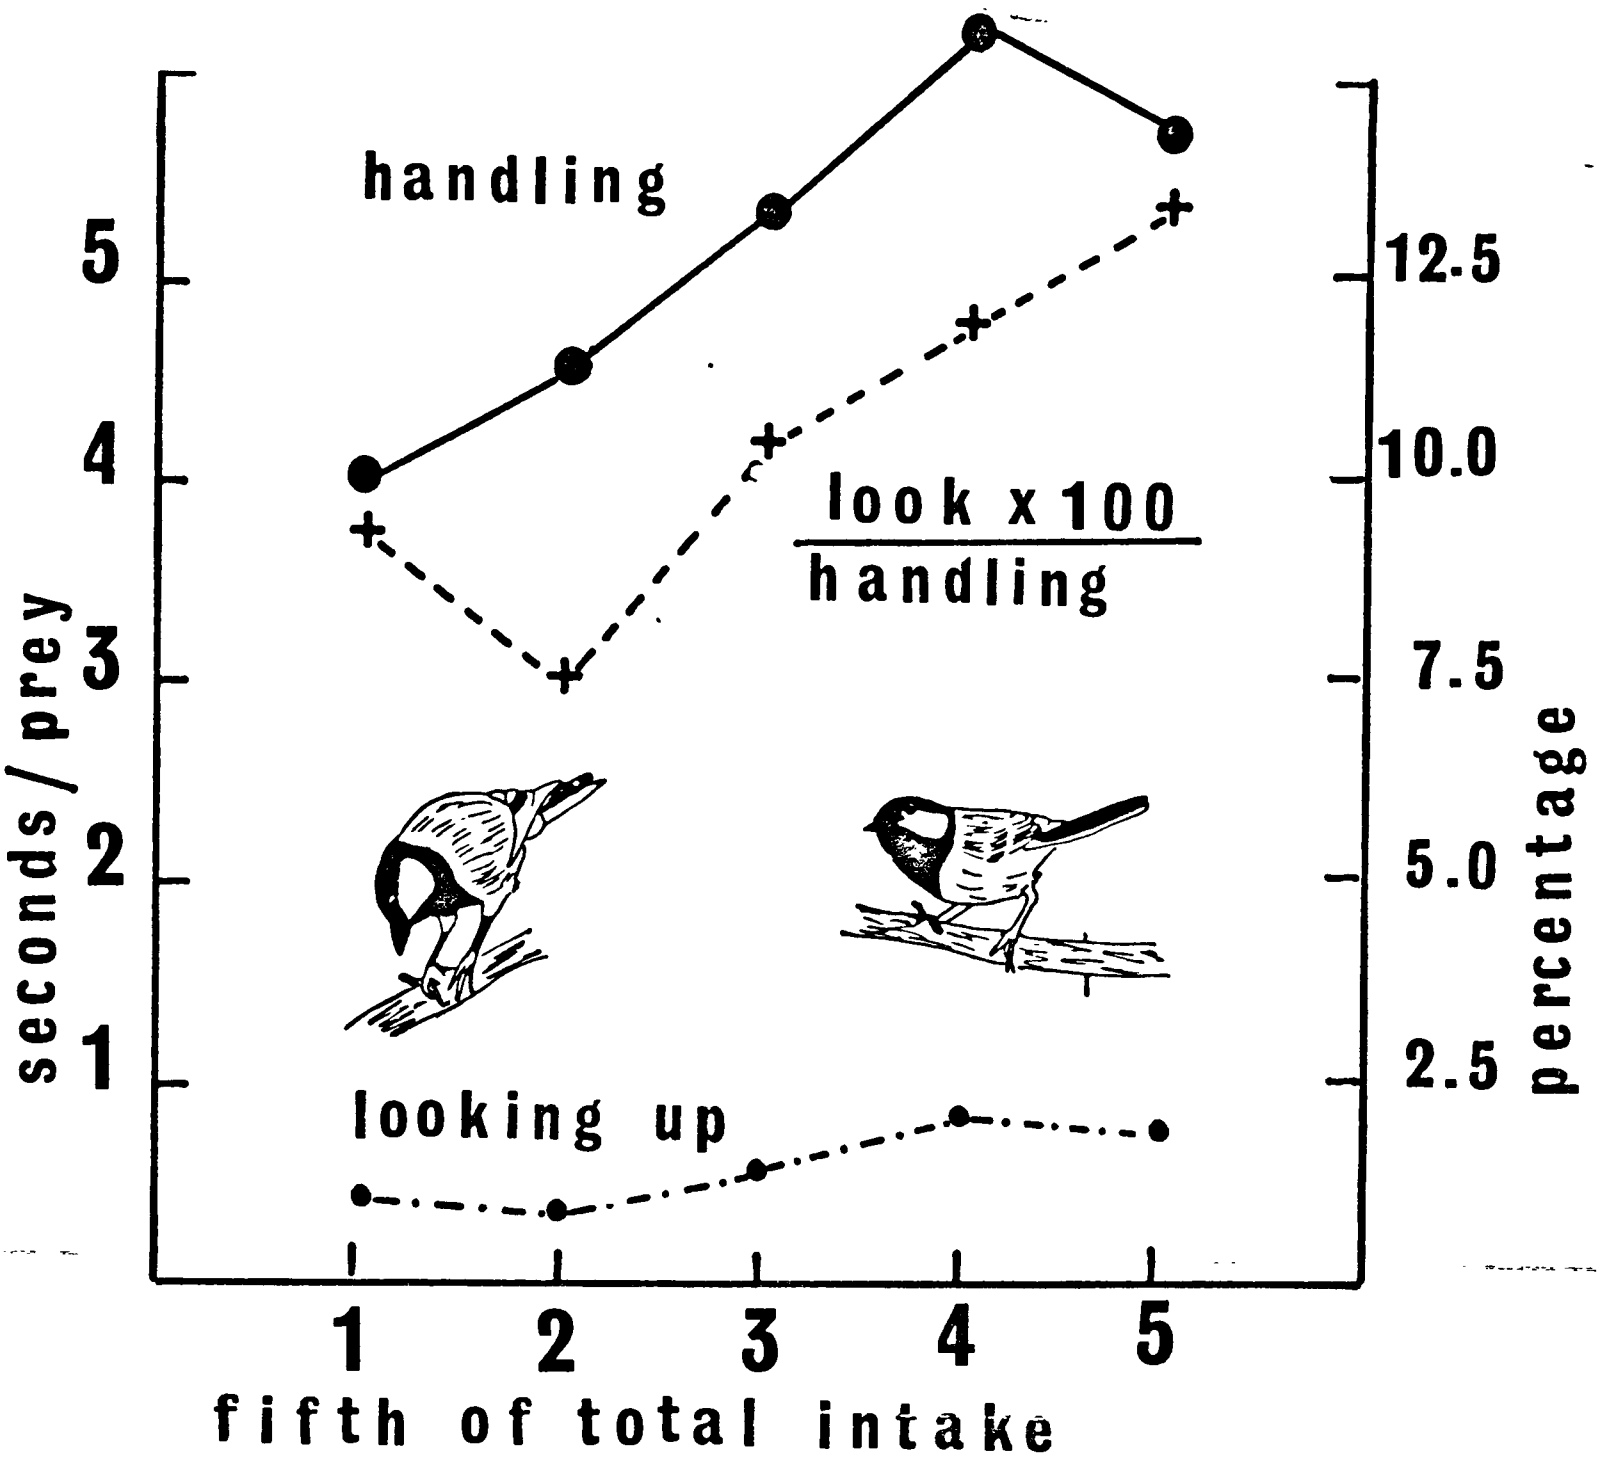


Figure Caption

Figure 1. Mean handling and look-up times (left ordinate) and mean look-up time as a percentage of handling time (right ordinate) for the cumulative fifth of total prey captures. Each point is a mean of 32 values.

Appendix 5

The effect of hunger on vigilance during prey handling in the great tit (Parus maj

When a foraging great tit (Parus major) finds a prey, it normally picks it up in its bill, moves to a suitable place to handle it, and starts pecking or tearing the prey while holding it with its feet. The pecking is frequently interrupted by the bird raising its head and moving it sideways. It would appear that this looking-around behaviour is performed in order to scan the surroundings for predators. Eventually, the bird either eats the last bit of food or drops the inedible portions such as the husks or wings in order to begin searching for a new food item. The interval between picking up the prey (or in some cases locating it) and starting to search for another one is termed handling time (H).

Handling time plays an important role in optimal foraging models. For example, diet selection has been predicted from the ratio of energy value to handling time (e.g. Krebs et al. 1977). Conversely, handling time itself has been predicted on the basis of the diminishing returns in rate of calories obtained from a single prey (Cook and Cockerell 1978). In all of these studies, it was implicitly assumed that handling time is independent of hunger since, according to the models, an optimally foraging animal would minimise H, and thus maximise energy intake per unit time, for any given level of food availability.

Milinsky and Heller (1978), however, recently showed that foraging strategy in sticklebacks is not only influenced by the rate of capture, but also by the presence of a predator. In general, they claimed that the optimal strategy for the stickleback lies somewhere between maximisation of food acquisition on the one hand and anti-predator behaviours on the other. The specific balance between these two behaviours will, of course, depend upon

the relative strength of the two tendencies (hunger and vigilance) at any given point in time (Sibly and McFarland 1976). Although Milinsky and Heller focussed on patch selection, their approach should also apply to other aspects of foraging such as handling time. If a certain amount of predator monitoring takes place while the prey is being handled, one would predict that H would remain constant only if the levels of risk from predators remain constant; when the risk of predators is not constant across conditions, H would compensate in some fashion.

We considered the same line of reasoning in terms of levels of hunger rather than predation risk. If an animal encountered two patches which ranked differently in terms of energy and predator risk, the optimal strategy would depend upon the combined risks (of starving and of being attacked). A very hungry animal, for instance, might increase its risk of attack from a predator in order to secure food and thereby reduce its risk of starvation; in contrast, a less hungry animal might choose a less risky feeding patch even though its net energy intake there was lower.

Handling time has indeed been shown to change during the course of feeding in birds (Smith and Dawkins 1971; Zachs and Falls 1978), but no functional explanation of this sort was considered. We hypothesised that as an animal gets less hungry, it is more likely to handle prey in a manner that allows for more anti-predator vigilance. Although the overall rate of food intake may decline, vigilance will increase since, as the animal nears its satiation point, it can afford to trade-off feeding efficiency for increased safety. Assuming, therefore, that the 'looking-around' pattern described above is, in fact, anti-predator behaviour (or vigilance), we devised a simple experiment in which we measured changes in handling time and vigilance during feeding.

We used two great tits as subjects. Each bird was tested four times under four deprivation levels: 45, 60, 75 and 90 minutes food deprived. The

sequence of sessions was such that all deprivation levels were tested in a random order before a replication was begun. In more formal terms, the experiment constituted a 2 x 4 x 4 random factorial design with subjects, deprivation levels, and replications as the factors.

Each test lasted 15 minutes and consisted of allowing the bird to obtain food (one fly pupa) by hopping onto a response perch (for a more detailed description of the apparatus and method, see Krebs et al. 1978). While the bird responded, the experimenter observed its behaviour from behind a one-way glass and recorded with an on-line computer device the occurrence and times for the following behaviours: hops onto the perch, picking up the prey, pecking at the prey, looking up, and dropping the prey.

The mean H time for fifths of the total number of prey eaten were first computed. We used this procedure because of the difference in the number of prey in the different sessions ( ). These values were then averaged over birds and replications (since those two main effects were not significant) to yield a five (of prey) by four (deprivation level) analysis matrix. As shown in Figure 1, H times did increase during feeding, thus confirming previous work (Smith and Dawkins 1971; Zachs and Falls 1978). Moreover, analysis of variance showed that fifths of total captures interacted with deprivation level indicating that the rate of increase of H times during feeding depended on the initial deprivation level - the greater the deprivation, the greater the rate of increase.

A similar analysis was done for the number of look-ups/prey and the average duration of looking-up (the means for the latter measure are shown in Figure 1). In both cases, the main effect of cumulative intake was significant, thus supporting our hypothesis that looking-up occurs more frequently and lasts longer as a function of satiation. Similarly, for both measures, the interaction between satiation and initial deprivation was significant indicating that the rate of increase of vigilance during feeding depended on the initial eprivation - the greater the deprivation, the greater the rate of increase in vigilance during the feeding session.

Finally, we considered the looking-up times/prey as a proportion of the H time. The means, shown in Figure 1 (right ordinate), ranged from about 9 to 13%. An analysis of variance confirmed that the increase is significant. Vigilance not only increased during feeding, but it came to occupy a greater proportion of H time. In other words, the birds spent a greater proportion of their handling time performing anti-predator behaviours as they approached satiation.

We conclude from these results that (a) handling time does, indeed, increase with intake; (b) vigilance, as measured by the number and duration of look-ups, also increases during feeding; and (c) the latter, in part, accounts for the former. In other words, H time increases, in part, because the subjects are spending relatively more of the H time being vigilant.

The argument is clearly that animals have evolved to respond with respect to a number of different pressures. Energy per unit time and vigilance are merely two, albeit important, selective pressures that contribute to the observed behaviour pattern. Presumably animals seek to optimise not just one outcome such as energy, but rather a configuration of outcomes. The relative degree of optimisation for any given outcome then depends upon the configuration of the internal state and environmental conditions at that particular time. In other words, our study shows that the relative contribution of one activity (say vigilance) depends not simply upon the conditions that directly affect that activity (e.g. animals are more vigilant when the probability of predators is high, see Milinsky and Heller 1978), but also on conditions that affect other behaviours (i.e. animals are more vigilant when the risk of starvation permits).

Although these results support our hypothesis, we recognise that more detailed studies are required to specify the full relationship between hunger and vigilance. To this end, experiments in which both predation risk and hunger level are manipulated would be most important. In addition, however, we speculate

that more comprehensive and/or quantitative studies will provide not only greater accuracy but also will reveal further complexity. That is, foraging behaviour will be shown to be affected by many more parameters than just hunger and predation risk. Indeed, our own data suggest this point. As H increased with satiation, looking-up time increased, but 'eating time' per prey (defined as the duration of the uninterrupted pecking) and the time between capturing the prey and the beginning of eating also increased. Clearly the animals were becoming more vigilant as hunger decreased, but they were also serving other motives as indicated by this increase in these other components of H time.

Alex Kacelnik

University of Oxford

Roger M. Tarpy

Bucknell University

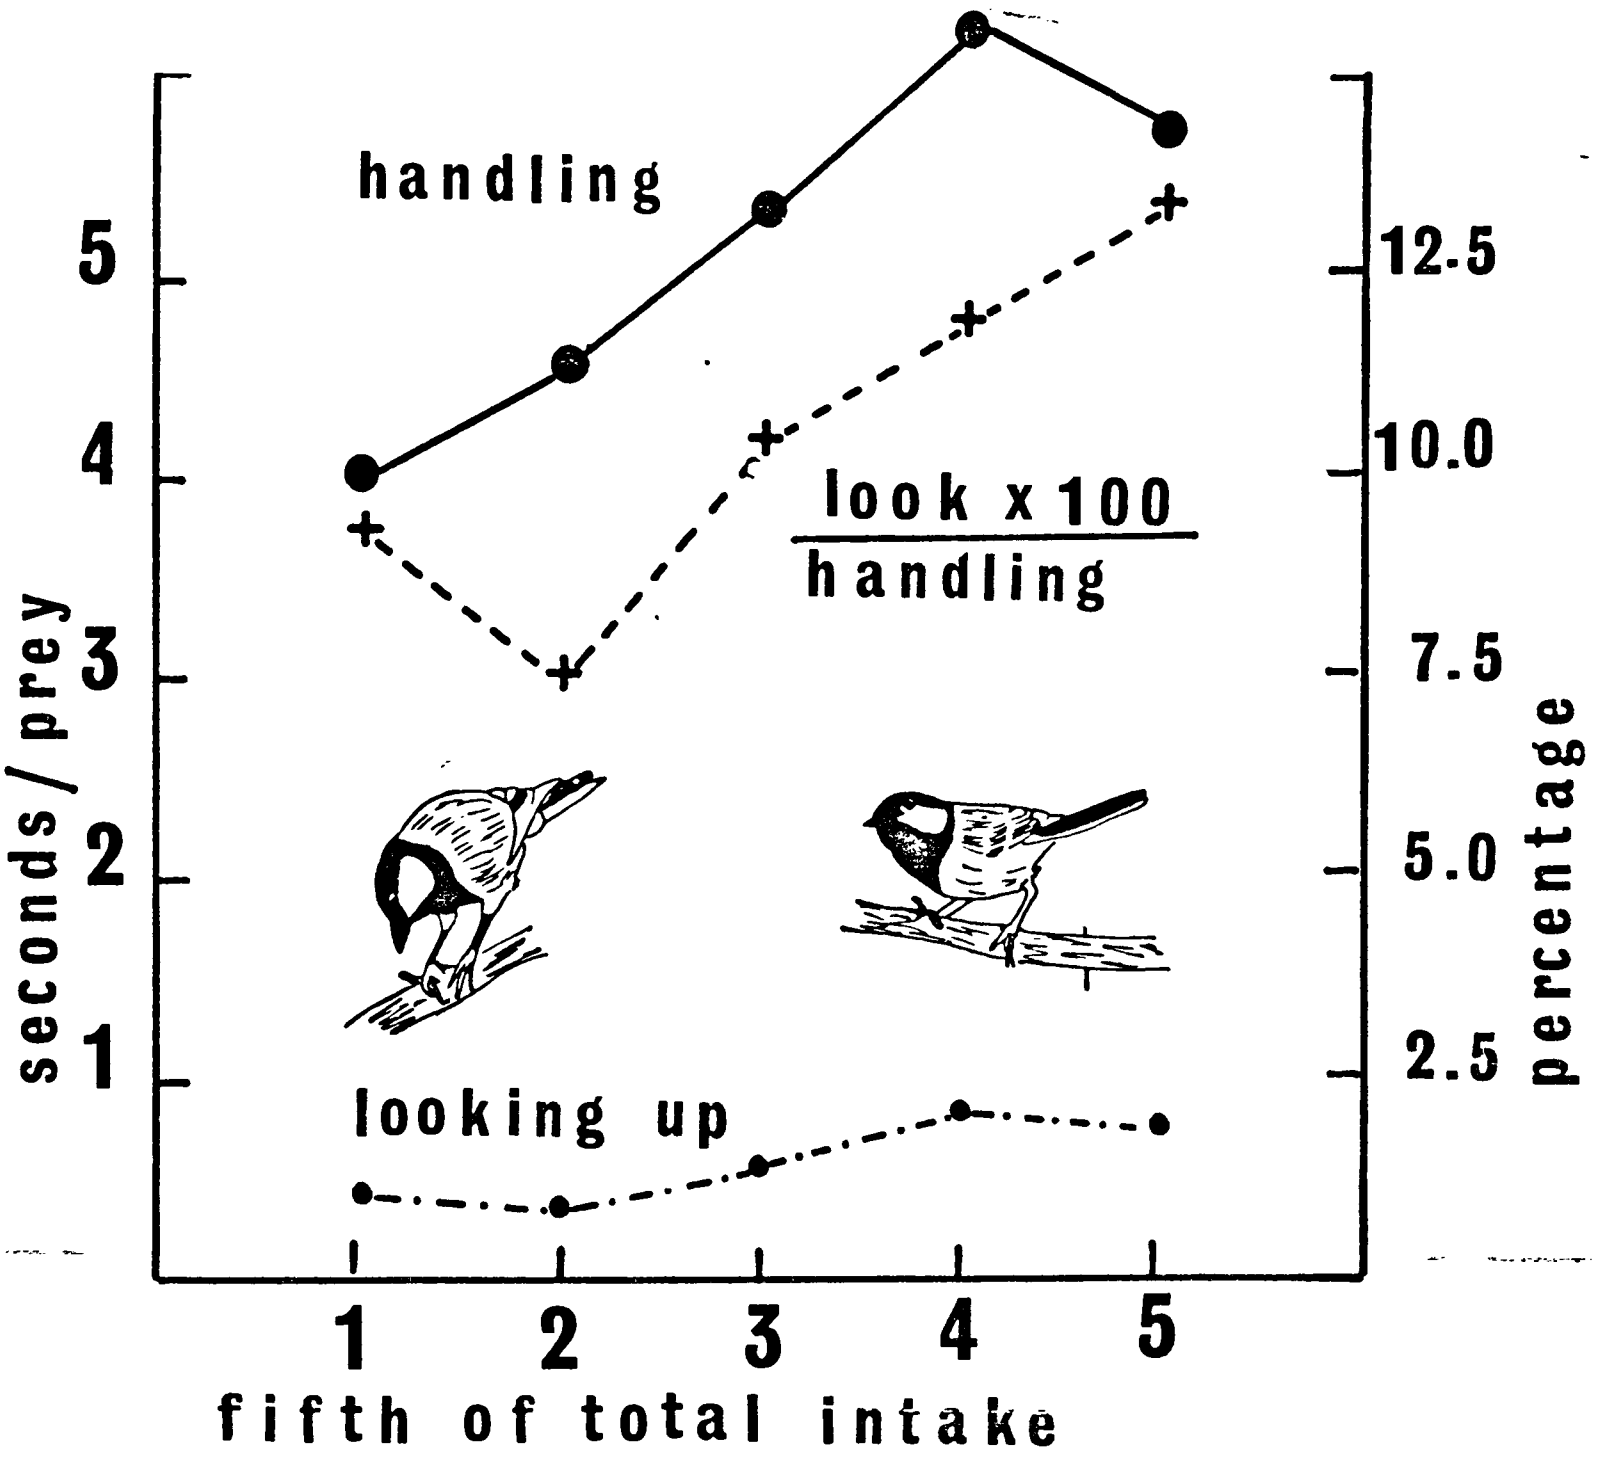


Figure Caption

Figure 1. Mean handling and look-up times (left ordinate) and mean look-up time as a percentage of handling time (right ordinate) for the cumulative fifth of total prey captures. Each point is a mean of 32 values.

(Reprinted from Nature, Vol. 275, No. 5675, pp. 27–31, September 7, 1978)
© Macmillan Journals Ltd., 1978

Test of optimal sampling by foraging great tits

John R. Krebs

Edward Grey Institute of Field Ornithology, Oxford, UK

Alejandro Kacelnik

Animal Behaviour Research Group, Department of Zoology, Oxford UK

Peter Taylor

Department of Mathematics, Queens University, Kingston, Ontario, Canada

When great tits forage in an unknown environment containing two feeding places of different profitability, they first sample the two places and then exploit the more profitable one. The balance between sampling and exploitation shown by the birds is close to an optimal solution for maximising the number of food-items obtained during a feeding period.

OPTIMAL foraging models, based on the premise that animals collect food in a way which maximises their net rate of food intake, have been quite successful in predicting the decision rules used by predators in laboratory experiments[1–9], but so far, little work has been done on the problem of how a predator samples the environment. Sampling is an implicit necessity of optimal foraging models, which assume that the predator behaves as if it knows the availability of different prey types or patches of food. We describe here an experiment designed to test whether or not the great tit, *Parus major*, uses a maximally efficient set of rules when sampling a patchy environment.

Two possible methods of obtaining maximal food intake with a choice of foraging area

Our experimental set up is extremely simple. The naïve bird is faced with a choice of two foraging places (patches) which differ in availability of food. The availabilities remain constant throughout a particular experiment, and the bird can only discover by exploration which of the two patches is better. Once the bird has distinguished between the quality of the two patches, it should concentrate on exploiting the better one; however, in our experiment we are concerned with how the bird makes the decision about which patch to exploit. We suggest that there are two simple types of maximising rules. On the one hand, the predator might attempt to maximise its rate of food intake at every instant in time by always foraging in the patch with the higher expected reward rate. We refer to this strategy as 'immediate maximising'. Alternatively, the predator might attempt to maximise its intake over the total foraging time and sacrifice short-term gain in order to acquire more information about the relative quality of the two patches. In this second strategy, an efficient predator has to choose the appropriate balance between exploration and exploitation. Our main aim was to test whether or not great tits follow this strategy and if so, whether they approximate an optimal balance between exploring and exploiting. The advantage of the sample-then-exploit strategy is that, in contrast to the strategy of immediate maximising, there is no risk of choosing to exploit the less profitable patch.

Simulation of the 'two-armed bandit' problem

The predator's problem of choosing the optimal balance between exploration and exploitation is similar to the classical 'two-armed bandit' problem, in which the player is faced by two machines, each with a characteristic, unknown reward rate on a random schedule, and the goal is to maximise the expected number of rewards in N trials[10,11]. This problem can be solved with a dynamic programming algorithm which works backwards from the Nth trial, at which point the machine with the higher percentage rewards is chosen. The algorithm involves analysing at each stage, n, the payoff for two possibilities (exploring or exploiting) in terms of the known optimal behaviour at stage $n+1$.

Our model is a modification of the optimal solution of the two-armed bandit problem[12]. We consider that the bird has a total of N trials (the meaning of which will become apparent), and at the beginning of an experimental session it samples each patch alternately. After M trials in each patch the bird commits itself to one patch or the other, so that all strategies have the form: M trials in each patch (exploration) followed by $N-2M$ trials on one patch (exploitation). In our model we introduced the constraint of equal sampling effort in the two patches during exploration, because this simplifies the model without significantly altering its predictions. The model predicts how the value of M, the point at which the predator decides to exploit, changes with p_1 and p_2, the reward probabilities in the two patches. Intuitively, one would expect M to decrease as $|p_1-p_2|$ becomes larger, and our model predicts the exact shape of the curve.

The details of our algorithm, which uses a Bayesian approach to describe the information possessed by the bird at each stage, are as follows. At the start of the experiment, the birds' prior estimate of p (the parameter of the Bernouilli process which is the reward rate in a patch) is considered to be β-distributed with parameters (α, β), and therefore, with mean $(\alpha+1)/(\alpha+\beta+2)$ and variance $(\alpha+1)(\beta+1)/(\alpha+\beta+2)^2(\alpha+\beta+3)$. The distribution has the property that after n trials with r successes, the posterior estimate of p is also β-distributed with new parameters $(\alpha+r, \beta+n-r)$. Thus, it is easy to follow the state of the bird's information about the two patches by updating α and β after each trial. In simulating the model, we used starting values of $\alpha=0$, $\beta=0$ for both patches. This means that the bird's initial estimate of the mean reward probability in the two patches is 0.50 with a uniform distribution between 0 and 1.0. Thus, after n trials with a_i successes, the bird's revised estimate of the reward rate p in patch i is $(a_i+1)/(n+2)$. We investigated the effect of varying the prior, by calculating our predictions for the following values of α and β: $\alpha=0, \beta=2$; $\alpha=0, \beta=0$; and $\alpha=2, \beta=0$. The model was rather insensitive to changes in the prior, most of the changes in prediction being between 10% and 20%.

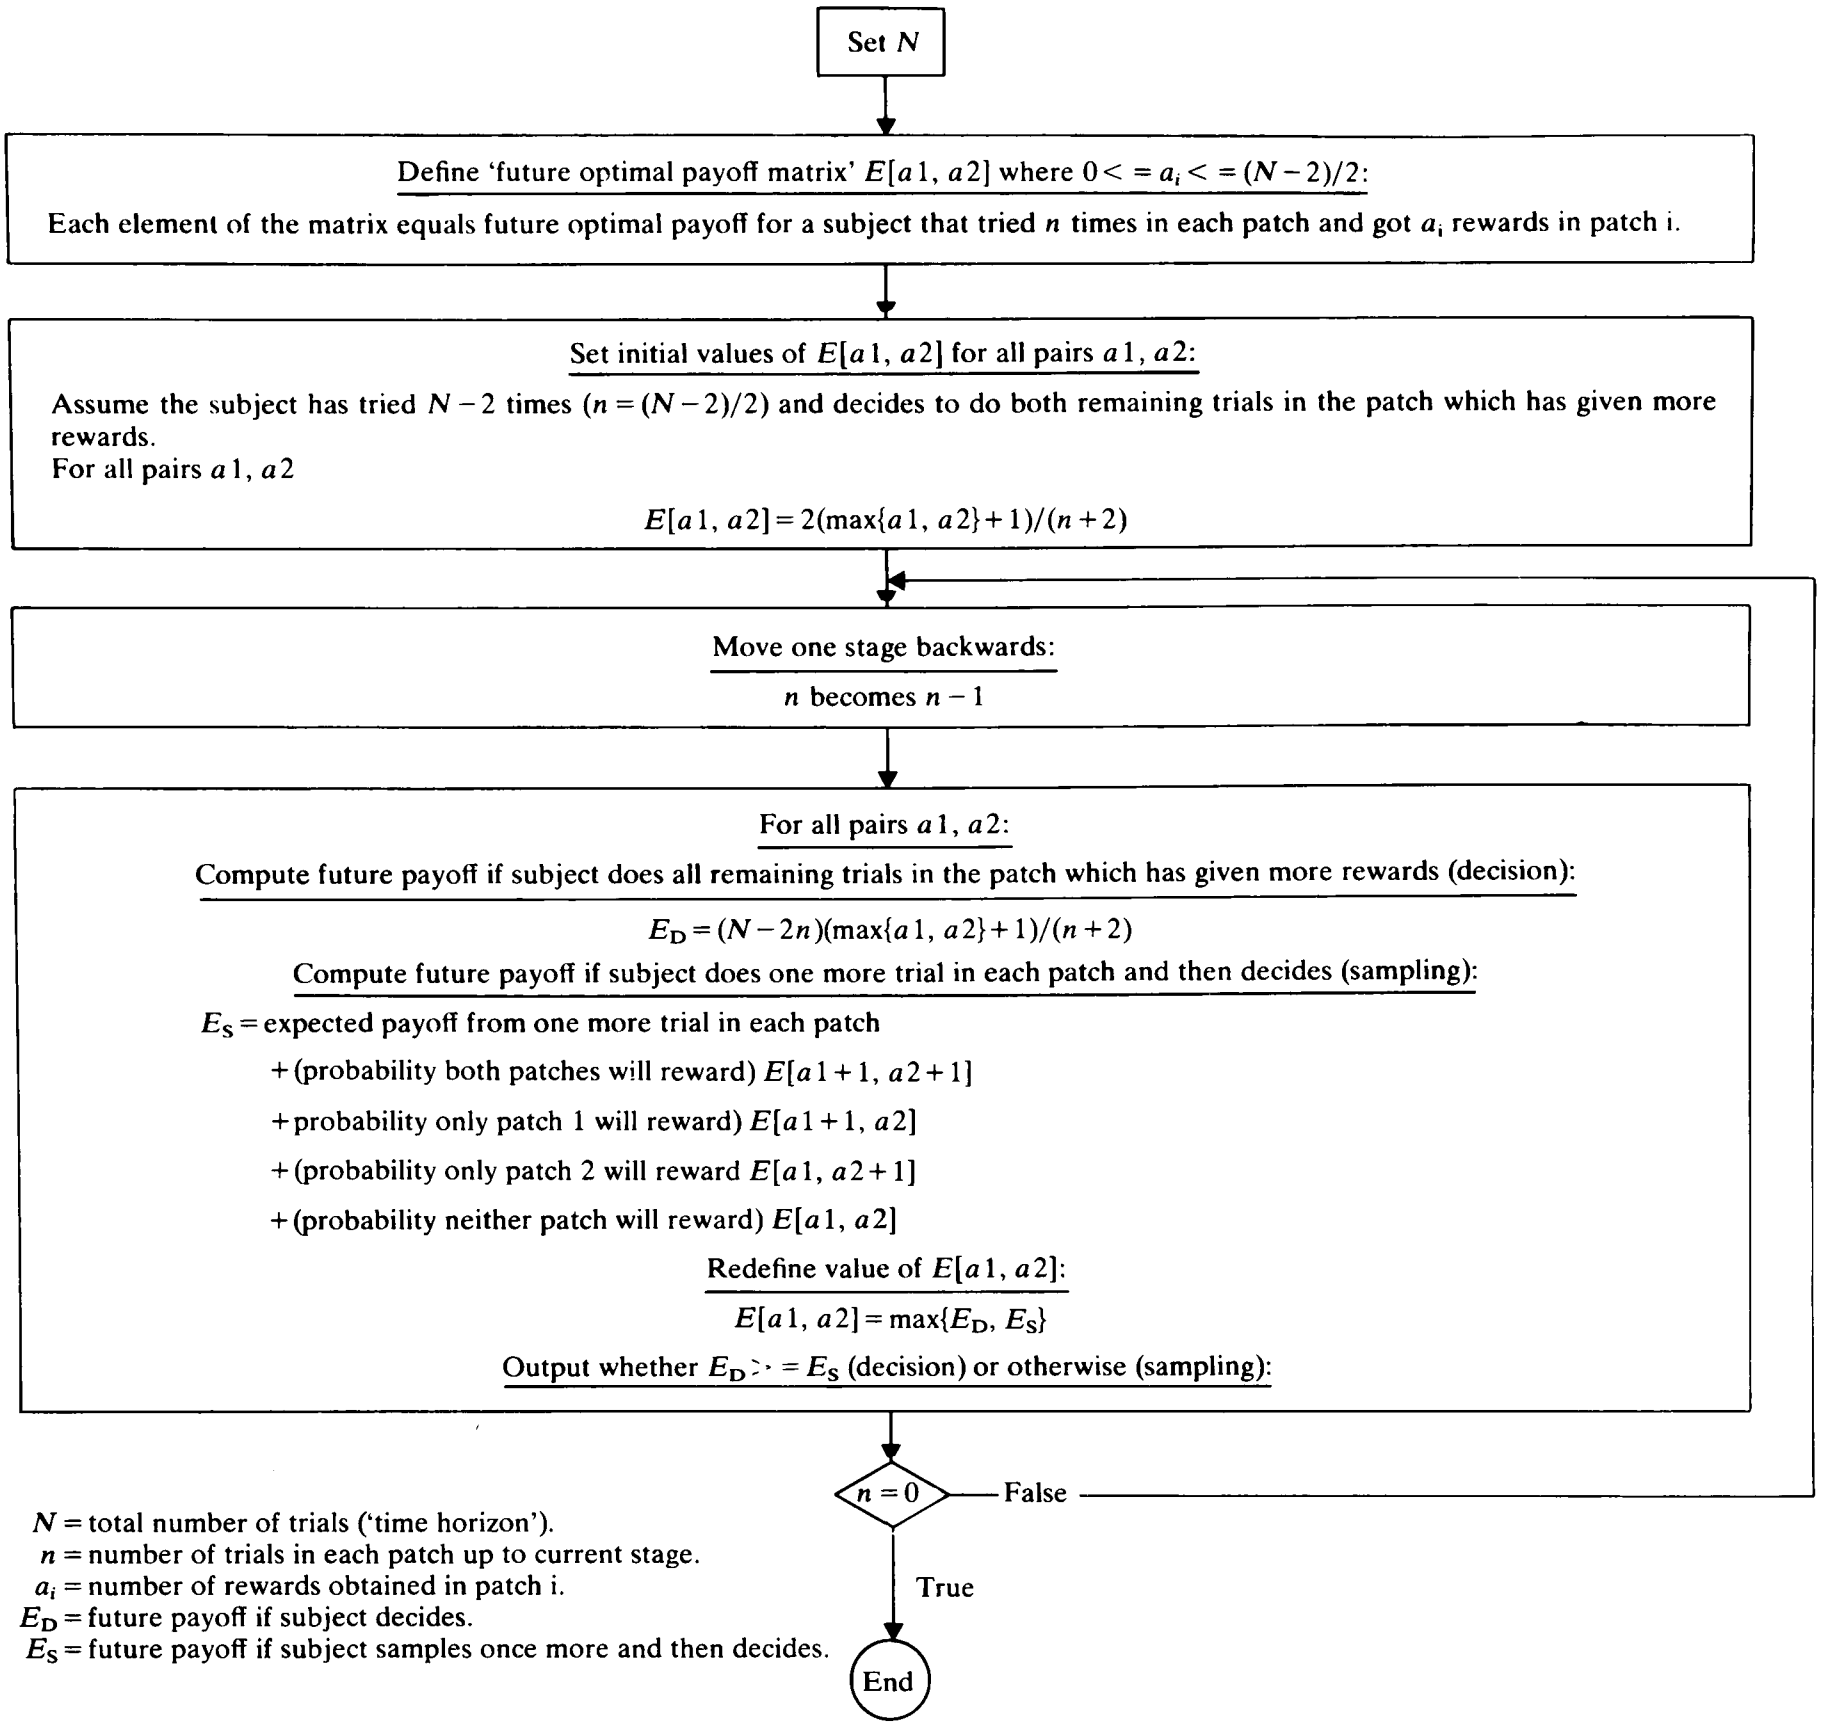


Fig. 1 'Two-armed bandit' simulation.

If, after n trials in each patch, the bird decides to sample once more, its expected payoff for the remainder of the foraging time is E_S, whereas if it decides to exploit, its payoff is E_D. The optimal strategy is to choose the larger (E) of these two values. E_S can be calculated in terms of the values of E at stage $n+1$ as follows:

Let

$$\pi_i = (a_i + 1)/(n+2)$$

Then

$$E_S(a_i, a_2, n) = \pi_1 + \pi_2 + \pi_2\pi_1 E(a_1+1, a_2+1, n+1) + \pi_1(1-\pi_2)E(a_1+1, a_2, n+1) + (1-\pi_1)\pi_2 E(a_1, a_2+1, n+1) + (1-\pi_1)(1-\pi_2)E(a_1, a_2, n+1)$$

and

$$E_D(a_1, a_2, n) = (N-2n)\max\{\pi_1, \pi_2\}$$

Figure 1 gives further details.

The model is simulated by starting with $2n = N$ and working backwards, generating E_S, E_D and E recursively. At each stage, the model produces an $n \times n$ matrix giving the value of $(E_D - E_S)$ for each combination of values of a_1 and a_2. When the bird is at the last pair of trials, E_D is clearly higher than E_S, assuming there is a difference in the two reward probabilities. As the simulation works backwards, it becomes more likely that, for any particular combination of a_1 and a_2, the payoff for sampling once more in each patch (and hence acquiring more information about the profitability of the two patches) is higher than the payoff for deciding to exploit. The exact point at which this change from sampling to decision occurs depends on the value of the total number of trials N, which in effect represents the bird's time horizon. For given values of a_1 and a_2, the larger the value of N, the longer the bird should sample before deciding. We ran the model with a range of values of N.

The experiment

In our experiments, the two patches consisted of two identical operant feeding places at opposite ends of an indoor aviary (measuring 4.3×3.7 m). The feeding places consisted of a perspex disk 35.5 cm in diameter enclosed in a shallow metal

box measuring 38.5 × 38.5 × 9 cm. The disk was drilled with 72 holes (diameter 0.7 cm, depth 0.4 cm) around the perimeter of its upper surface. Each of these holes contained a piece of mealworm (weight 0.07 g), and the bird had access to one hole at a time through a small horizontal window in the top of the box which enclosed the disk. To get at the next piece of food, the bird hopped on a perch next to the disk, which operated a solenoid-driven stepping cog to turn the disk through a small angle, bringing the next hole into line with the window. The reward rate was set on a pseudo-random variable ratio schedule using BRD modular logic units. This experimental set up is exactly analogous to that used in probability learning experiments[13].

The experiments were carried out between June 1976 and February 1977 and involved testing nine wild-caught, adult great tits. Each bird was first trained to perform the operant task and then given a series of tests of different treatments consisting of various percentage reward rates in the two disks. The values we used were 50:0; 40:10; 35:15; 30:20. It is important that although these ratios always summed to 50%, we think the birds did not learn this, as between each treatment they were subjected to a series of one to four 'neutralisation' tests in which both disks offered a 7.5% reward rate. These tests were always continued until the bird showed no marked preference for one disk over the other; we used a low reward rate during neutralisation sessions because we found this to be effective. If there was a slight preference, the less preferred disk was used as the more profitable place in the next treatment. A 'treatment' (for example, 20:30) consisted of a series of between 1 and 5 10-min tests. We had to run a series of short tests because the birds sometimes emptied one of the disks of rewards before a treatment was completed. When this happened, we interrupted the treatment for a short time while the disk was refilled. During these interruptions the bird was locked outside the test aviary without food, and in our analysis we always discounted the first 20 hops after an interruption. Any particular treatment was continued until the bird had 'decided' to exploit one disk or the other (nearly always the more profitable place).

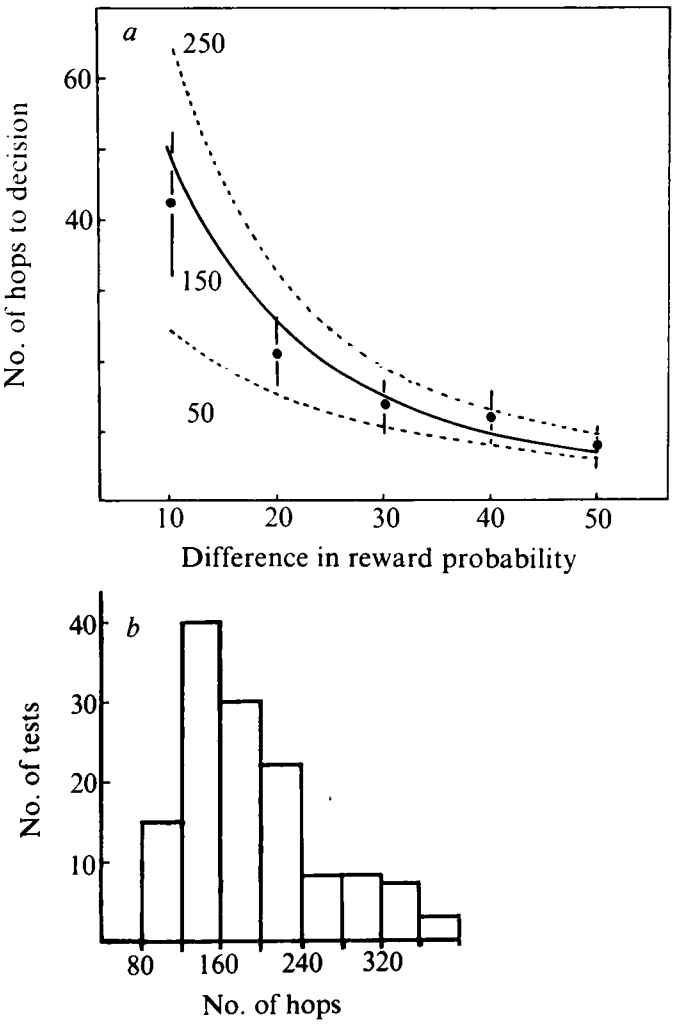


Fig. 2 *a*, The predicted (——, – – – –) and observed (●) number of hops before making a decision (nine birds). Three predicted curves based on the 'sample-then-exploit' model are shown, for $N = 50$, 150 and 250. The observed means are close to 150 (continuous line) curve. Also shown are the 95% confidence intervals of the observed geometric means (we used geometric means to normalise the distribution). The x axis is the difference in reward probability between the two perches for the five experiments (for example, 50:0, difference = 50; 35:15, difference = 20) *b*, A frequency distribution of lengths of tests. The modal value is close to $N = 150$, which gave the best fit in Fig. 3*a*. $\bar{x} = 199$.

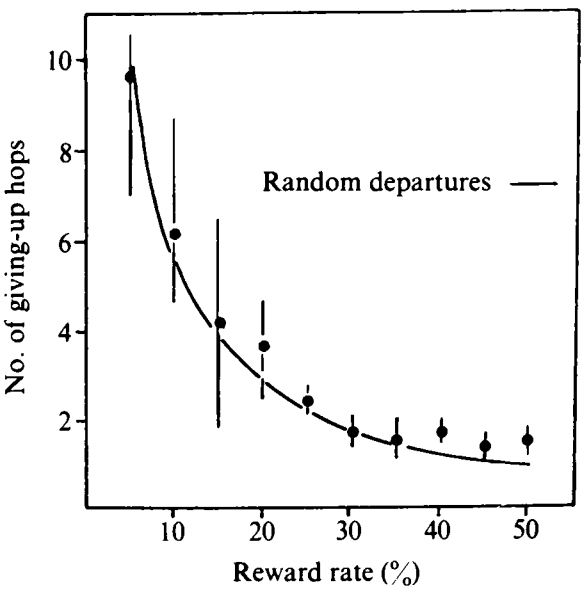


Fig. 3 The observed number of 'giving-up hops' (unsuccessful hops before a departure) plotted against reward probability. The random curve is calculated taking into account the fact that the birds tend to hop in short bouts. If p is the probability of getting a reward on a single hop and $q = (1-p)$, then the expected number of giving up hops for bouts of $5 = qp + 2q^2p + 3q^3p + 4q^4p + 5q^5p$. The observed points lie very slightly, but consistently above the random departure line, indicating that the birds tend to leave a perch after a run of bad luck. The observed points are ± standard error.

The decision criterion was that the bird should perform more than 90% of a sequence of 100 hops on the preferred perch. This decision criterion was chosen after calculating the probability of the bird going back to the less profitable perch for a bout of 20 or more hops after different lengths of hopping bouts on the profitable perch. With the 100-hop criterion the probability of this type of reversal is nearly asymptotic and is less than 5%. Each bird was tested twice with all five treatments, the sequence of testing being randomly chosen. At the beginning of each treatment, the test bird was deprived of its normal food and excluded from the indoor aviary while the disks were being loaded with prey. Preliminary trials showed that the birds worked at a constant rate (15 hops per min) throughout a series of tests after an initial 60-min deprivation, which was the minimum period we used. The data were recorded on a computer-compatible event recorder which stored the information on magnetic tape.

Treatment of results

The results of these experiments showed that the birds followed the general pattern of behaviour assumed by our version of the two-armed bandit model. There was an initial sampling period during which the birds performed on average 42% of hops on the more profitable and 58% on the less profitable perch. This proportion does not differ significantly from 50:50, and is consistent with the constraint of equal sampling in our model. The slight bias towards the less profitable perch resulted from our experimental procedure of always starting with the higher reward rate in whichever perch was used less in the preceding neutralisation test. After the decision point, the birds spent virtually all their time on the more profitable perch; on average, 95% of hops were on this perch, and the range for individual treatment means was 88–99%. Similar results, referred to as 'maximising' on concurrent variable ratio schedules, have been obtained with pigeons[13].

In terms of our optimal decision model, we can consider a hop as being equivalent to a single trial, so that we can predict, for each combination of reward rates, how many hops on each perch the bird should do before making the decision to exploit. Figure 2*a* shows the observed mean decision point in each experimen-

tal treatment and the predictions of the model for three values of N (total number of hops in the experiment), 250, 150 and 50. The fit to the model with $N = 150$ is good; not only is the observed curve the same shape as predicted, but the predicted values are all fairly close to the observed means. The other values of N give curves of the same general shape, but with predicted values too high or too low. The value of N which fits the data (150) is the modal value for the number of hops performed in a test (Fig. 1*b*). This suggests that birds are optimising over the 'time horizon' which they experienced most often during our experiments, but we do not know if the birds would optimise with a bigger value of N if they were accustomed to longer tests. We plan to test this possibility. There is a slight difference in the modal value of test lengths for different treatments, but as the treatments were presented in a random order there is no reason to suspect that the birds could modify their behaviour according to treatment.

One way in which the birds' behaviour differed from that implied in our model is that they did not sample by doing one hop on each perch at a time, but instead they hopped in bouts of about five. (The overall mean bout length before decision was 5.00, the range 3.00–16.5.) This makes no difference in predicting the average decision point, but if the birds had hopped in rigidly determined bouts, the predicted decision point for a particular experiment would have to be a multiple of 5, which would involve reading every fifth matrix of the model's output. This tendency to hop in bouts probably explains why the observed points are above those predicted at the right-hand end of the graph where the bird should make a decision after about 10 hops.

The alternative goal we considered above was the strategy of immediate maximising, in which the bird always switches to the perch with the higher current expected reward rate. We used the β-distribution and the same starting values of α and β as above to predict, for any given experiment, at which point the birds should no longer switch back and forth if they were following this strategy. This predicted point is reached when expected payoff in the current patch never drops below that expected in the other one. The observed point was based on the decision criterion discussed above. The expected number of switches depends on the actual sequence of obtaining rewards in the first few moments of the experiment. In calculating the expected values we made the conservative assumption that the birds would continue to switch until the difference between patches in expected reward rate was greater than 5%. The observed mean values for the five treatments (50:5, 45:5, 40:10, 35:15, 30:20) were 3.3, 3.6, 3.8, 6.4 and 9.6, whereas the corresponding expected values were 0.93, 1.3, 1.36, 1.3 and 3.25. If we relax the assumption that the birds cannot discriminate differences of less than 5% in our experiments, the difference between observed and predicted values becomes even greater. Clearly, the birds switch more than expected. In other words, they are not maximising in the very short term, but instead they behave as if they acquire information to achieve a longer-term optimum.

Although the birds do not seem to switch according to the rule of immediate maximising, there is some evidence that switches between perches are related to short-term changes in reward rate. In Fig. 3 we plot the number of unsuccessful hops before a switch during the sampling period as a function of reward rate. The points lie very close to, but slightly above, the random departure curve. This random curve takes into account the fact that the birds work in bouts rather than single hops. Eight out of 10 means lie just above the random curve ($P \simeq 0.05$). This shows that birds have a slight tendency to switch after a run of bad luck. In other words, they have a weak component of immediate maximising in their sampling strategy.

Discussion

The bird's behaviour approximates the prediction of a simple model based on a constrained optimal balance between exploration and exploitation, and does not fit the model of pure exploitation. The fit to the former model is close in spite of the fact that we have not yet directly incorporated a cost of switching between patches. In effect, we incorporated a small cost by using the rule that if the payoffs for sampling and exploiting were exactly equal, then the bird should exploit. In our experiment, this cost was small both in energetic terms and in time. However, in further developments of the model the cost of switching might be an important consideration.

Our comparison of the two maximising models has been in terms of predicting the details of the bird's behaviour, but an alternative approach is to compare the payoff, measured as number of rewards per test. On average, the two-armed bandit model achieves a slightly higher payoff than immediate maximising, except in very short experiments with less than five trials[14], but when the reward rates are very different in the two patches (for example, 50:0) and the risk of a wrong decision is negligible, immediate maximising is just as effective. The difference between the birds' behaviour and that predicted by the two-armed bandit model, although small, is sufficient for the birds to do slightly worse than either of the two models in terms of payoff. However, the difference is less than 5% of the predicted optimum of the two-armed bandit model (A.K. and J.K., in preparation). Thus, an analysis of payoffs enables us to say that the birds are close to the optimum but does not discriminate between the two models, unlike the analysis of the birds' behaviour.

We have presented averaged data for the nine birds we tested, but individual variation was very marked. We considered the possibility that some individuals are more cautious (do more exploring) but are less likely to make mistakes. Five of the nine individuals tested made one or more mistakes (that is, chose the poorer patch) in the 30:20 or 35:15 tests. These individuals did not make decisions after fewer hops than the other four birds. The average number of hops before a decision for all treatments was 23.9 ± 9.6 and 18.4 ± 9.8 for the fallible and infallible individuals, respectively. If, however, we examine individual treatments in which a bird made an incorrect decision, it is apparent that these tend to be cases in which the bird did little or no sampling. Examining the 18 '20:30' treatments (nine birds, two treatments each), only three out of the 13 in which a correct decision was made involved no sampling, whereas two out of the five in which the bird made a mistake were tests in which there was no sampling period at the beginning.

Although the birds in our experiments behaved on average as if they were calculating an optimal balance between exploration and exploitation, we suggest that they use some simple rule to approximate the optimum. We have simulated various potential rules including 'matching' future responses to rewards obtained so far, an empirically derived rule applying concurrent choice experiments[15]. The extent to which these proximate rules mimic the behaviour of the birds and match the total payoff achieved will be discussed elsewhere (A.K. and J.K., in preparation).

We thank the SRC for financial support, Mike Cullen and John Gittins for valuable advice, and Nick Styles, Janet Lomer and Ruth Ashcroft for help with the experiments. Ted Fleming and Nick Davies commented on the manuscript.

Received 20 March; accepted 18 July 1978.

1. Krebs, J. R. & Davis, N. B. (eds) *Behavioural Ecology: An Evolutionary Approach* (Blackwell, Oxford, 1978).
2. Krebs, J. R., Ryan, J. C. & Charnov, E. L. *Anim. Behav.* **22,** 953–964 (1974).
3. Krebs, J. R., Erichsen, J. T., Webber, M. I. & Charnov, E. L. *Anim. Behav.* **25,** 30–38 (1977).
4. Cowie, R. J. *Nature* **268,** 137–139 (1977).
5. Goss-Custard, J. D. *Anim. Behav.* **25,** 10–29 (1977).
6. Cook, R. M. & Hubbard, S. F. *J. Anim. Ecol.* **46,** 115–125 (1977).
7. Charnov, E. L. *Theor. Populat. Biol.* **9,** 129–136 (1976).
8. Charnov, E. L. *Am. Nat.* **110,** 141–151 (1976).
9. Werner, E. E. & Hall, D. J. *Ecology* **55,** 1042–1052 (1974).
10. Bellman, R. *Sankhyā* **16,** 221–229 (1956).
11. Barnett, V. *Comparative Statistical Inference* Wiley, New York, (1975).
12. Lindley, D. V. & Barnett, B. N. *Biometrika* **52,** 507–532 (1965).
13. Mackintosh, N. J. *The Psychology of Animal Learning* (Academic, New York, 1974).
14. Jones, P. W. *Biometrika* **62,** 523–524 (1975).
15. Herrnstein, R. J. & Loveland, D. H. *J. exp. Analysis Behav.* **24,** 107–116 (1975).

Printed in Great Britain by Henry Ling Ltd., at the Dorset Press, Dorchester, Dorset