

# **On the Relations between Behaviour, Mechanism, and Environment: Explorations in Artificial Evolution**

**Anil Kumar Seth**

Submitted for the degree of D. Phil.

University of Sussex

October, 2000

## **Declaration**

I hereby declare that this thesis has not been submitted, either in the same or different form, to this or any other university for a degree.

Signature:

## Acknowledgements

This thesis could not have existed without the advice, support, and criticism of many people. Phil Husbands and Hilary Buxton, my supervisors, not only provided bucketfuls of each, but gave me the opportunity to pursue this project in the first place. My first thanks must go to them.

Andy Philippides, Tom Smith, and Rebecca Ball have seen me through this and much else besides, they have contributed to every part of this thesis, and have - above all - made it a happy endeavour. I cannot thank them enough.

There are of course many others: John Anderson, Lionel Barnett, Seth Bullock, Ron Chrisley, Andy Clark, Steve Crossan, Ezequiel di Paolo, Joe Faith, Dario Floreano, Paul Graham, Inman Harvey, Nick Jakobi, Giles Mayley, Jason Noble, Matt Quinn, Julie Rutkowska, Linc Smith, Emmet Spier, Adrian Thompson, Peter Todd, and Friday Afternoon Football deserve particular mention; I thank all of them deeply. The contents of this thesis have also benefited from the consistent advice of friends from the humanities; Rafe Hallett, Phil Newman, and Ben Roberts have had the most influence; again, my thanks. Also, I must also acknowledge the various (and often anonymous) referees that have commented on the parts of this thesis that have been previously published; I have profited greatly from their constructive criticisms.

For keeping me in lucre for the past three-and-a-bit years I am grateful above all to the School of Cognitive and Computing Sciences for deciding to award me an EPSRC research scholarship. Also, Phil Husbands has been instrumental in arranging plenty of teaching, an activity perhaps more rewarding academically than financially, but which helped pay the bills nonetheless (thanks also to Danny Osorio for arranging teaching in the School of Biology). I must also acknowledge the contribution of the Centre for Computational Neuroscience and Robotics, not just for its bountiful facilities (here the Sussex High Performance Computing Initiative must also be mentioned) and peerless inhabitants, but for arranging, via Phil and Mick O'Shea, two very helpful bursaries when things were really tight. Thanks also to Phil's portfolio of grants for subsidising some of my conference excursions, all of which were, well, worthwhile.

Finally, my parents deserve, as always, my deepest gratitude. For their support, financial and otherwise, certainly. But also for coping with their only son seemingly intent on remaining a student for ever and ever. To them I say; I have finished.

## Preface

Many parts of this dissertation have previously appeared in single-author peer-reviewed publications. An early version of chapter 3 appeared in the *Proceedings of the Fourth European Conference on Artificial Life* (Seth, 1997, see also Seth, 1999a), and an early version of chapter 4 appeared in the *Proceedings of the First European Workshop on Evolutionary Robotics* (Seth, 1998c). Elements of both, together with the first outlines of chapter 2, appeared in the *Proceedings of the Sixth International Conference on Artificial Life* (Seth, 1998a). A version of chapter 5 recently appeared in *From animals to animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior* (Seth, 2000), and a version of chapter 6 was presented at the same conference two years beforehand (Seth, 1998b). Finally, parts of chapter 8 appeared in the *Proceedings of the Fifth European Conference on Artificial Life* (Seth, 1999b).

For ease of reference, I provide here a list of all abbreviations utilised in this dissertation that appear in more than one chapter (these abbreviations are also defined as and when they appear in the text):

- AL - Artificial Life
- ECT - Environmental Complexity Thesis
- GA - Genetic Algorithm
- IFD - Ideal Free Distribution
- IOS - Individual-based Optimal Situated
- IPD - Iterated Prisoner's Dilemma
- IPD/CR - Iterated Prisoner's Dilemma with Choice and Refusal
- LRV - Law of Requisite Variety
- SAB - Simulation of Adaptive Behaviour

# **On the Relations between Behaviour, Mechanism, and Environment: Explorations in Artificial Evolution**

**Anil Kumar Seth**

## **Abstract**

This thesis presents an externalist exploration of the relations between behaviour, mechanism, and environment, as they arise in a variety of agent-environment systems. It offers contributions at conceptual, methodological, and empirical levels of discourse.

Externalism describes the attempt to understand the internal in terms of the external, and the thesis begins by developing a conceptual framework justifying the use of artificial evolution models in the application of this perspective to agent-environment systems. In particular, it is argued that such models play a crucial role in elaborating the distinction between behavioural and mechanistic levels of description. There follows a series of models, of both game-theoretic and evolutionary-robotic character, which focus on explaining internal 'complexity' in terms of adaptation to (external) environmental variability. As part of this project, accounts of the 'evolution of complexity' in general are critiqued, and the practical importance of noise in artificial evolution is discussed.

The thesis continues with an integration of this externalist project with the well established theoretical biology methodology of 'optimal foraging theory'. A novel methodology - 'individual-based optimal situated modelling' - is described, which extends orthodox optimal foraging theory through (1) the use of artificial evolution as an optimisation procedure and, (2) modelling agent-environment interaction at the level of situated perception and action. The conceptual leverage afforded by this extension is illustrated in its application to the problem of behaviour coordination in a simple agent-environment system; for example, the need for a dedicated 'action selection mechanism' is questioned. The methodology is then addressed to a range of issues in contemporary theoretical biology and psychology: the 'interference function', the 'ideal free distribution', and the 'individual matching law', issues which are united by a concern with individual choice and its collective consequences. A series of models are presented which demonstrate, in these contexts, that (1) behaviours for which there is debate about the level of complexity required for their underlying mechanism, can be subserved by surprisingly simple mechanisms, and (2) behaviours which may be irrational when expressed by an isolated individual can be understood as rational in a group context.

Submitted for the degree of D. Phil.

University of Sussex

October, 2000

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Orientation . . . . .	5
1.2	Artificial evolution models . . . . .	6
1.3	Overview of the dissertation . . . . .	6
1.3.1	Background . . . . .	6
1.3.2	Complexity . . . . .	7
1.3.3	Towards a theoretical biology . . . . .	8
1.3.4	Interference and the matching law . . . . .	10
1.3.5	A return to orthodoxy . . . . .	12
1.4	Summary and contributions . . . . .	12
<b>2</b>	<b>Externalism, behaviour, mechanism, and environment</b>	<b>15</b>
2.1	Internalism, externalism, and enlightened interactionism . . . . .	16
2.1.1	Internalism and externalism . . . . .	16
2.1.2	Arbitrariness and enlightened interactionism . . . . .	17
2.1.3	Summary . . . . .	19
2.2	Cognitive science and cognitive mechanism . . . . .	20
2.3	Behaviour, mechanism, and environment . . . . .	21
2.3.1	A leitmotif . . . . .	21
2.3.2	Environmental structure . . . . .	22
2.3.3	Perception and action . . . . .	24
2.3.4	Summary . . . . .	24
2.4	From empiricism to behaviourism . . . . .	25
2.4.1	A brief history of association psychology . . . . .	25
2.4.2	Spencer’s ‘evolutionary associationism’ . . . . .	27
2.4.3	Behaviourist learning theory . . . . .	27
2.4.4	To the present day . . . . .	29
2.4.5	Summary . . . . .	30
2.5	Evolutionary psychology . . . . .	31
2.5.1	Contemporary evolutionary psychology . . . . .	31
2.5.2	Adaptive function and ecological ‘niche’ . . . . .	32
2.5.3	Optimal foraging theory . . . . .	32
2.5.4	Summary . . . . .	33
2.6	Complexity and the function of mind in nature . . . . .	33
2.6.1	The environmental complexity thesis . . . . .	34
2.6.2	Spencer’s version . . . . .	36

2.6.3	Dewey's version . . . . .	36
2.6.4	Accommodation versus response: a second look . . . . .	37
2.6.5	Summary . . . . .	38
2.7	A cybernetic ECT . . . . .	38
2.7.1	The law of requisite variety . . . . .	38
2.7.2	The LRV as a version of the ECT . . . . .	39
2.7.3	Homeostasis . . . . .	40
2.7.4	Summary . . . . .	41
2.8	Construction . . . . .	42
2.8.1	Causal and constitutive construction . . . . .	42
2.8.2	In defence of constitutive construction . . . . .	43
2.8.3	A potential confusion . . . . .	46
2.8.4	Summary . . . . .	46
2.9	The shape of things to come . . . . .	47
2.10	Overall summary . . . . .	48
<b>3</b>	<b>Complexity I: Definitions and game-theoretic models</b>	<b>49</b>
3.1	Definitions of complexity . . . . .	50
3.2	The evolution of complexity . . . . .	52
3.2.1	Three perspectives . . . . .	53
3.2.2	The trouble with teleology . . . . .	54
3.2.3	Evolution and progress . . . . .	55
3.2.4	The causal question . . . . .	56
3.3	Complexity in the iterated prisoner's dilemma . . . . .	57
3.3.1	Structure of the present model . . . . .	59
3.3.2	Results from the compulsory model . . . . .	60
3.3.3	The facilitation of genetic mobility . . . . .	64
3.3.4	IPD/CR models . . . . .	65
3.3.5	IPD/CR results . . . . .	66
3.3.6	Discussion . . . . .	70
3.4	'Win-stay, lose-shift' strategies . . . . .	71
3.4.1	WSLS strategies and the ECT . . . . .	73
3.5	Summary . . . . .	74
<b>4</b>	<b>Complexity II: Evolutionary robotics</b>	<b>75</b>
4.1	A study of homing navigation . . . . .	75
4.1.1	A description of the model . . . . .	76
4.1.2	Behavioural analysis . . . . .	78
4.1.3	Mechanistic analysis . . . . .	81
4.1.4	Discussion . . . . .	83
4.2	Noise and (artificial) evolution . . . . .	85
4.3	Summary . . . . .	87

<b>5</b>	<b>Unorthodox optimal foraging theory</b>	<b>88</b>
5.1	Optimal foraging theory . . . . .	89
5.1.1	What use is orthodox OFT? . . . . .	90
5.2	IOS models . . . . .	92
5.2.1	IOS as OFT: at first glance . . . . .	93
5.3	Difficulties with orthodox OFT . . . . .	93
5.3.1	Mathematical tractability . . . . .	93
5.3.2	Decision variables . . . . .	94
5.3.3	Historical constraints . . . . .	95
5.3.4	External environment and <i>Umwelt</i> . . . . .	96
5.3.5	The pragmatics of optimisation . . . . .	97
5.3.6	Proxies for reproductive fitness . . . . .	98
5.4	IOS as OFT: a second look . . . . .	98
5.5	Related Work . . . . .	100
5.5.1	Individual-based models in ecology . . . . .	100
5.5.2	Individual-based models in SAB/AL . . . . .	101
5.5.3	SAB/AL and theoretical biology . . . . .	101
5.6	Summary . . . . .	102
<b>6</b>	<b>Behaviour coordination</b>	<b>104</b>
6.1	The theoretical terrain . . . . .	104
6.1.1	Hierarchical control . . . . .	105
6.1.2	Non-hierarchical control . . . . .	106
6.1.3	Behavioural economics . . . . .	107
6.1.4	Behaviour and mechanism, perception and action . . . . .	109
6.1.5	Summary . . . . .	111
6.2	A simple IOS model of behaviour coordination . . . . .	112
6.2.1	A description of the model . . . . .	113
6.2.2	Genetic encoding scheme . . . . .	115
6.3	Results . . . . .	116
6.3.1	Behavioural analysis . . . . .	117
6.3.2	Mechanistic analysis . . . . .	119
6.3.3	Link properties . . . . .	121
6.3.4	Mechanism activity during behaviour . . . . .	122
6.3.5	Mechanism analysis summary . . . . .	128
6.4	Discussion . . . . .	128
6.5	Summary . . . . .	130
<b>7</b>	<b>An IOS model of the interference function</b>	<b>131</b>
7.1	Modelling the interference function . . . . .	132
7.1.1	A brief history of interference . . . . .	132
7.2	An IOS model of interference . . . . .	134
7.2.1	Agent and environment . . . . .	135



7.2.2	Evolving an interference function . . . . .	136
7.2.3	Experimental conditions . . . . .	137
7.2.4	Summary of model structure . . . . .	138
7.3	Results . . . . .	138
7.3.1	Interference functions . . . . .	138
7.3.2	‘Clumpiness’ . . . . .	139
7.3.3	Sensitivity: perceiving conspecifics . . . . .	143
7.3.4	Sensitivity: varying group sizes during evolution . . . . .	144
7.4	Discussion . . . . .	146
7.4.1	Modelling interference . . . . .	146
7.4.2	A broader view . . . . .	148
7.5	Summary . . . . .	149
<b>8</b>	<b>The matching law</b>	<b>151</b>
8.1	Individual choice and group choice . . . . .	152
8.2	An IOS model of within-patch matching . . . . .	154
8.2.1	Some caveats . . . . .	155
8.2.2	Extending the ‘interference’ model . . . . .	156
8.3	Results . . . . .	157
8.3.1	Simple matching in simple agents . . . . .	157
8.3.2	Adaptive sensors . . . . .	159
8.3.3	The role of interference . . . . .	160
8.3.4	The importance of history . . . . .	161
8.3.5	Historical resource abundances . . . . .	163
8.3.6	Sensitivity . . . . .	163
8.3.7	Summary . . . . .	166
8.4	Discussion . . . . .	166
8.4.1	Individual suboptimality . . . . .	166
8.4.2	The matching law . . . . .	166
8.4.3	Exploration versus exploitation . . . . .	168
8.4.4	A broader view . . . . .	168
8.5	Summary . . . . .	169
<b>9</b>	<b>Sampling strategies, matching, and the ideal free distribution</b>	<b>170</b>
9.1	The ideal free distribution in some detail . . . . .	171
9.1.1	Immediate consumption versus standing stock models . . . . .	171
9.1.2	Resource allocation . . . . .	173
9.1.3	Summary . . . . .	174
9.2	The matching law . . . . .	174
9.3	A Description of the model . . . . .	176
9.3.1	$\epsilon$ -sampling . . . . .	177
9.3.2	Model structure . . . . .	178
9.3.3	Results: $\epsilon$ -sampling . . . . .	179

9.3.4	Discussion: $\varepsilon$ -sampling . . . . .	179
9.3.5	$\omega$ -sampling . . . . .	181
9.3.6	Results: $\omega$ -sampling . . . . .	182
9.3.7	Discussion: $\omega$ -sampling . . . . .	183
9.4	The role of artificial evolution . . . . .	186
9.4.1	Near-optimal parameters . . . . .	186
9.4.2	Practical details . . . . .	187
9.4.3	Near-optimal versus pre-specified parameters . . . . .	188
9.4.4	Strategy memory and noise . . . . .	188
9.5	Summary . . . . .	189
<b>10</b>	<b>Conclusions</b>	<b>191</b>
10.1	Summary of contributions . . . . .	192
10.2	Perspective . . . . .	197
	<b>Bibliography</b>	<b>199</b>

## List of Figures

1.1	The artificial duck of Jacques de Vaucanson. . . . .	2
2.1	An agent and its environment (adapted from Beer, 1995). . . . .	18
3.1	IPD model encoding scheme. . . . .	59
3.2	Evolution of complexity in the compulsory model. . . . .	60
3.3	Fitness and memory in the compulsory model (1). . . . .	61
3.4	Fitness and memory in the compulsory model (2). . . . .	62
3.5	Compulsory model with 2% O-noise. . . . .	63
3.6	Cooperation in the compulsory model. . . . .	63
3.7	Compulsory model with functional memory limit of 1. . . . .	64
3.8	Compulsory model with genetic memory limit of 1. . . . .	65
3.9	Cooperation in the genetically limited model. . . . .	65
3.10	Evolution of complexity in the IPD/CR model. . . . .	67
3.11	Fitness and memory in the IPD/CR model (1). . . . .	68
3.12	Fitness and memory in the IPD/CR model (2). . . . .	69
3.13	2% O-noise and functional memory limit in the IPD/CR model. . . . .	69
3.14	Cooperation in the IPD/CR model. . . . .	70
3.15	Replication of Posch's (1999) WSLs model . . . . .	72
3.16	Influence of M-noise on average payoff and $\gamma$ , with fixed $p$ . . . . .	73
4.1	The Khepera mobile robot. . . . .	76
4.2	Sensor layout and network architecture for homing navigation study. . . . .	76
4.3	Khepera robot performing 'homing navigation'. . . . .	79
4.4	Trajectory plots for type <i>A</i> and type <i>B</i> robots. . . . .	80
4.5	Type <i>A</i> and type <i>B</i> robots in a 'no-wall' competition. . . . .	80
4.6	Neuron activation profile for a type <i>A</i> robot. . . . .	82
4.7	Neuron activation profile for a type <i>B</i> robot . . . . .	82
4.8	Hidden unit responses for type <i>A</i> and type <i>B</i> robots. . . . .	83
4.9	Hidden unit activity in no-wall test. . . . .	83
4.10	M-noise can accelerate evolution. . . . .	85
4.11	Population convergence in 'homing-navigation' model. . . . .	86
6.1	Baerends' 'hierarchy of motivations'. . . . .	106
6.2	Maes' 'bottom-up' action selection architecture. . . . .	108
6.3	Brunswik's lens model . . . . .	110
6.4	Braitenberg's 'decision making' vehicle. . . . .	112
6.5	Architecture of the agent. . . . .	114

6.6	Genetic encoding scheme . . . . .	116
6.7	A sample trajectory . . . . .	117
6.8	A balance between persistence and dithering. . . . .	118
6.9	‘Screen shots’ of behaviour coordination tests. . . . .	119
6.10	Opportunism. . . . .	120
6.11	Evolved transfer functions. . . . .	121
6.12	Motor output and battery levels. . . . .	123
6.13	Link input and output recordings. . . . .	124
6.14	Link input and output recordings: L-R difference. . . . .	125
6.15	Link input and output recordings: time difference. . . . .	126
7.1	Interference model agent and environment structure. . . . .	136
7.2	Interference model experimental design. . . . .	137
7.3	Interference functions for conditions $A, B$ , and $C$ . . . . .	139
7.4	Intake rates of isolated $A$ -, $B$ -, and $C$ -agents. . . . .	140
7.5	Intake rates of isolated $A$ -, $B$ -, and $C$ -agents (log-log plots). . . . .	140
7.6	Average $\bar{cl}$ for groups of 2-13 agents from $A, B$ , and $C$ conditions. . . . .	141
7.7	Average $cl_t$ of a group of 8 $A$ -agents and a group of 8 $B$ -agents. . . . .	142
7.8	Forward speed profiles of isolated $A$ - and $B$ -agents. . . . .	143
7.9	Speed ranges for conditions $A, B$ , and $C$ . . . . .	144
7.10	Forward speed profile of an isolated $C$ -agent. . . . .	144
7.11	Interference functions for conditions $D$ and $E$ . . . . .	145
7.12	Speed ranges for conditions $A - E$ . . . . .	145
8.1	Agent perception in the matching model. . . . .	157
8.2	Matching and ‘zero-one’ behaviour. . . . .	158
8.3	Evolved sensor properties. . . . .	159
8.4	Environmental instability in modified $S$ -type models. . . . .	162
8.5	Resource density robustness: same environment. . . . .	163
8.6	Resource density robustness: forced choice. . . . .	164
8.7	Agent density robustness of matching and zero-one behaviour. . . . .	165
8.8	Environmental complexity robustness of matching and zero-one behaviour. . . . .	165
9.1	Observed and predicted population distributions under $\epsilon$ -sampling. . . . .	179
9.2	Matching behaviour of isolated $\epsilon$ -sampling agents . . . . .	180
9.3	Matching behaviour of embedded $\epsilon$ -sampling agents. . . . .	180
9.4	Observed and predicted population distributions under $\omega$ -sampling. . . . .	182
9.5	Matching behaviour of isolated $\omega$ -sampling agents . . . . .	183
9.6	Matching behaviour of embedded $\omega$ -sampling agents. . . . .	184
9.7	Near-optimal values for $\epsilon$ and $\gamma$ . . . . .	188
9.8	Relation between noise and adaptation rate $\gamma$ . . . . .	189

# List of Tables

- 3.1 Standard IPD payoff table. . . . . 57
- 7.1 Summary of interference functions. . . . . 141
- 1 Parameters for the compulsory and IPD/CR models. . . . . 219
- 2 ‘Homing navigation’ simulation noise levels. . . . . 220
- 3 Error in observed population distributions with respect to IFD predictions. . . . . 225
- 4 Isolated agent matching to obtained resources. . . . . 226
- 5 Error in embedded matching to available resources with respect to strict matching. 226
- 6 Embedded agent matching to obtained resources. . . . . 227

*“The purpose of art is to lay bare the questions which have been hidden by the answers.”*

James Mark Baldwin

# Chapter 1

## Introduction

---

But there are a variety of motives other than the highly laudable one of giving pleasure that could provoke a man to devise artificial singing birds, dancing shepherds, and swinging flute players. One such motive might be the desire to understand how some part of the natural world actually works.

- Fryer and Marshall (1979, p.257)

I own freely, that I am surpriz'd myself to see and hear my *Automaton* play and perform so many and so differently varied Combinations: And I have been more than once ready to despair of succeeding; but Courage and Patience overcame every Thing.

- Jacques de Vaucanson (1742, p.24)

Throughout the year 1742, visitors to London's Opera House at the Haymarket had the opportunity to witness a phenomenon to "astonish all Europe".<sup>1</sup> The flute-player, tabor-pipe player, and (most renowned of all) artificial duck of the French automata-maker Jacques de Vaucanson were on display in the Long Room at "1, 2, 5, and 7 o'clock in the Afternoon" (de Vaucanson, 1742, p.1). These were objects of daunting appearance. The flute player, Vaucanson records, was "about six Foot and a half high, sitting upon a Piece of Rock, placed on a square Pedestal, four Foot and a half high, and three Foot and a half wide" (ibid., p.10). Both the pedestal and the figure were tightly packed with all kinds of machinery; bellows, pipes, reservoirs, pulleys, together controlling the movement of the lips and fingers of the figure, and governing the flow of air to the flute. The tabor-pipe player, a later creation, "stands upright on its Pedestal, dress'd like a dancing Shepherd [and] plays twenty Tunes, Minuets, Rigadoons, and Country-dances" (ibid., p.23). And the artificial duck, a picture of which appears in figure 1.1, "represent[ed] the Mechanism of the Intestines which are employed in the Operations of Eating, Drinking, and Digestion: Wherein the Working of all the Parts necessary for those Actions is exactly imitated" (ibid., p.21).

These descriptions come from translations, by J.T. Desaguliers (chaplain to the then Prince of Wales), of Vaucanson's deposition to the French Académie Royale du Sciences in 1738, describing the flute-player in detail, and of a subsequent letter to the Abbe de Fontaine, introducing the tabor-pipe player and duck. It is largely thanks to these documents that the efforts of Vaucanson and his contemporaries may be seen as more than gadgets developed purely for their entertainment value,

---

<sup>1</sup>Brewster (1834), p.201.



Figure 1.1: The artificial duck of Jacques de Vaucanson. From Chapuis and Droz (1958).

until quite recently a dominant view (Brewster, 1834; Raphael, 1976). Fryer and Marshall (1979), in particular, reassess Vaucanson's motivations, and portray him as an 'explanatory psychologist', intent on understanding behaviour, and the mechanisms underlying the generation of behaviour, through simulation. Of his description of the flute player, they have this to say: "There is a sense in which Vaucanson's seemingly 'objective description' is really a hypothesis concerning the mode of operation of man plus flute" (p.259). Vaucanson himself is quoted in support:

These, Gentlemen, have been my thoughts upon the sound of Wind-Instruments and the Manner of modifying it. Upon these Physical Causes I have endeavour'd to find my Enquiries; by imitating the same Mechanism in an *Automaton*, which I endeavour'd to enable to produce the same Effect in making it play on the German-Flute. (1742, p.12, quoted in Fryer & Marshall, 1979, p.259)

Perhaps most interestingly, however, the employment of automata for such purposes enabled Vaucanson to identify a variety of phenomena, the significance of which would not otherwise have been apparent - "Discoveries of Things which could never have been so much as guess'd at" - as he put it himself. He finds, for example, that playing the tabor-pipe requires a surprisingly great 'force of Wind': "Cou'd it have been thought, that this little Pipe shou'd, of all the Wind-Instruments, be one of the most fatiguing to the lungs?" (ibid., p.24). Another discovery is that the force required to produce a particular note is in part dependent on the *preceding* note:

[T]hat Wind, for example, which is able to produce a D following a C, will never produce it, if the same D is to be sounded next to the E just above it; and the same is to be understood for all other notes. (ibid., p.24, quoted in Fryer & Marshall, 1979, p.263)



Vaucanson is nevertheless extremely careful to delimit the extent of his investigations, and, in particular, is cautious of any confusion between the automaton itself and that which is being simulated by it. He advises the Abbe de Fontaine on the interpretation of his artificial duck:

I don't pretend to give this as a perfect *Digestion*, capable of producing Blood and nutritive Particles for the Support of the Animal. I hope no body will be so unkind as to upbraid me with pretending to any such Thing. I only pretend to imitate the Mechanism of that Action in three Things, *viz.* *First*, to swallow the Corn; *secondly*, to macerate or dissolve it; *thirdly*, to make it come out sensibly changed from what it was. [...] [T]he whole Mechanism of our artificial *Duck* is exposed to View; my Design being rather to demonstrate the Manner of the Actions, than to shew a Machine. (ibid., p.22)

Fryer and Marshall conclude, from their vantage, that “Vaucanson’s main achievement lies [...] in the clarity with which he perceived and articulated the explanatory mode he sought to attain” (p.264). How best to summarise this explanatory mode? For Fryer and Marshall it satisfies a series of desiderata; that theories must be explicit, that ‘occult entities’ may not interact with machines, that accounts of internal mechanism must accompany accounts of behaviour, and, finally, that the behavioural repertoires of organisms must be characterised. All this is so, yet underlying these specific contributions remains the methodological principle of using artefacts to understand, by simulation, biological phenomena; to articulate explicit hypotheses and - to recall Baldwin - to lay bare questions hidden by answers.<sup>2</sup>

Understood at this level, Vaucanson’s explanatory mode is strongly echoed in the present project. In the following chapters, artefacts - of both physical and computational constitution - are deployed in a series of interrogations of relations between behaviour, mechanism, and environment. Like Vaucanson, this dissertation is sensible of the need to understand behaviour in terms of dynamic agent-environment interaction (the flute-player as a system of ‘man plus flute’), and as distinct from its subserving (agent-side) mechanisms. Like Vaucanson, the artefacts are not taken to instantiate their targets of explanation, rather, they articulate specific hypotheses through carefully constrained simulation; the aim is never to “shew a Machine” for its own sake. And like Vaucanson, their use enables surprising and otherwise impenetrable phenomena - “Things which could have never been so much as guess’d at” - to be elucidated. Of course, many differences remain. The artefacts of this dissertation are unlikely to cause much public astonishment, nor indeed are they likely to find themselves on display in Opera Houses anywhere. Neither am I concerned with providing an improved theory of ‘Wind-Musick’ (or of the digestive system of ducks), and it is at this point that we must depart from Vaucanson and move on to matters of present concern.

This dissertation presents an exploration of relations between behaviour, mechanism, and environment, as they arise in agent-environment systems at a variety of different granularities. This is a project with conceptual, methodological, and empirical elements, and is characterised partly by the use of artefacts such as simulation models, and partly by the adoption, at a pragmatic level of commitment, of an ‘externalist’ perspective. Externalism describes the attempt to understand

---

<sup>2</sup>Why were nineteenth century (and many twentieth century) commentators so reluctant to award scientific merit to Vaucanson’s work? One possible reason is that the kind of science pursued by Vaucanson was not seen to contribute to the development of industry so important to Victorian society. However, this comment is itself rather unfair; Fryer and Marshall (1979) remind us of many industrial innovations directly accountable to Vaucanson, for example the design of an apparatus for the automatic weaving of brocades, and the use in machines of flexible tubes of India rubber.

internal properties in terms of external properties, and a recurring motif in this dissertation is the use of ‘artificial evolution’ (a strongly externalist paradigm, introduced in section 1.2) to partially automate artefact design.

The particular issues addressed cross-cut a number of disciplines, both new and old (some common themes will be isolated in the following section). The first part of the dissertation is concerned with matters of ‘complexity’, more specifically with understanding behavioural and mechanistic complexity in terms of adaptation to external (environmental) variability. Philosophical discourse is united with concrete models, of both game-theoretic and evolutionary robotic character, in an interrogation of the widely-held intuition that evolution by natural selection is responsible for the proliferation of biological complexity beheld in the history of life. In the second, and more extensive part, the focus is on the ‘coordination of behaviour’, its functional and mechanistic aspects, and on the relations between group behaviour and individual behaviour. Close contact is made with a variety of issues prominent in theoretical biology and experimental psychology, amongst them the ‘interference function’, the ‘ideal free distribution’, and the ‘matching law’ (details follow). The various models addressed to these issues deliver two kinds of insight in particular, firstly that behaviours for which there is debate about the level of complexity required for their underlying mechanism, can be subserved by surprisingly simple mechanisms, and secondly, that behaviours which may be irrational when expressed by an isolated individual can be understood as rational in a group context.

Elements of the subject matter and methodological practice of this dissertation recur throughout the currently fashionable disciplines of ‘artificial life’ (AL) and the ‘simulation of adaptive behaviour’ (SAB). For present purposes, the latter is best understood as a subfield of the former. SAB inherits from Vaucanson the use of artefacts to understand, by simulation, biological and behavioural phenomena (see, for example, Pfeifer, Blumberg, Meyer, & Wilson, 1998; Meyer, Berthoz, Floreano, Roitblat, & Wilson, 2000). AL is all this and more, famously summarised by Langton to extend beyond “the study of *life-as-we-know-it* into the realm of *life-as-it-could-be*” (Langton, 1989, p.1, emphasis in original). There is an important point to make here. Some practitioners of AL would, with Langton, consider their artefacts to ‘be alive’ themselves, to blur or traverse the distinction (carefully respected by Vaucanson) between simulation and instantiation. This stance, ‘strong AL’ in the terminology of Sober<sup>3</sup>, is *not* a part of the present project; ‘weak AL’ shares with SAB (and with Vaucanson, and with what follows) a clear emphasis on simulation. With regard to the subject matter of this dissertation, issues such as ‘complexity’ and ‘game-theory’ appear much more often in the discourse of AL than in that of SAB, whereas ‘artificial evolution’ and ‘behaviour coordination’ figure prominently in both.

How, then, to characterise the contents of this dissertation? The combined abbreviation SAB/AL offers itself as one possibility. This, however, seems to underestimate its expressly interdisciplinary nature, and its direct engagement with pre-existing intellectual currents. Other alternatives suggest themselves. Cybernetics, exemplified by W. Grey Walter’s construction of artificial ‘tortoises’, very directly mediates between Vaucanson and contemporary SAB/AL: Grey Walter’s creations were simple machines, wheels, wires, and little else, yet they engaged in complex, even social patterns of behaviours (Walter, 1950, 1953). There is also cognitive science, understood by

---

<sup>3</sup>See, for example, Sober (1996).

many as the study of the mechanisms of behaviour generation, undeniably relevant, yet for present purposes tarnished - as we will see in chapter 2 - by its association with computational theories of mind.

In the end, it does not much matter. Faith said of his thesis, following Quine, that it should face the tribunal of examination as a corporate body, not as a set of independent arguments; that it “may be easier to swallow whole than in pieces”.<sup>4</sup> This dissertation, by contrast, may be digested both ways. Each individual investigation will make its contributions and their context clear, and yet, if the whole appears greater than the sum of the parts (as is fervently hoped), it would be somehow churlish and contrary to the interdisciplinary nature emphasised throughout, to attempt to draw all together under a single banner.

## 1.1 Orientation

The contents of this dissertation are unashamedly diverse, exploring the relations between behaviour, mechanism, and environment in many different contexts and at many levels of abstraction. Nevertheless, as intimated above, some common themes prevail, and to provide some orientation for what follows it is worth isolating these themes, as far as is possible, at this early stage. To begin with, the terms behaviour, mechanism, and environment deserve some clear definition. Behaviour, in this dissertation, refers to observed patterns of agent-environment interactivity, mechanism to the agent-side internal structure subserving this interactivity, and environment can be interpreted intuitively as that which surrounds the agent. These definitions, although perfectly adequate for now, will be tightened up in the following chapter.

The importance of distinguishing between behavioural and mechanistic levels of description is emphasised throughout. Much of what follows is concerned with elucidating the consequences of misconstruing this distinction; in general these involve radical prejudice of how a given behaviour might be generated, significant overestimation of the degree of mechanistic complexity required, and denial of the close coupling of perception and action. Such consequences are particularly evident in the context of ‘behavioural choice’ (or, more generally, the ‘coordination of behaviour’), and many of the specific investigations that follow can be related to this context.

This dissertation offers an externalist conceptual framework, to be cashed out in the form of artificial evolution models (see below), as an effective means of interrogating the distinction between behaviour and mechanism. These models are also presented as a means of evaluating the hypothesis (deriving from the overarching conceptual framework) that behavioural and/or mechanistic complexity can be understood in terms of adaptation to environmental variability. The various models of this dissertation strongly support this hypothesis, however they also indicate - as suggested above - that complex behaviours need not be subserved by complex internal mechanisms; we will see that there need be no contradiction here. The utility of models which operate at the level of situated perception and action in affording these insights is emphasised, with particular reference to implications for orthodox modelling strategies in theoretical biology.

These general themes are representative but by no means exhaustive of the contents of this

---

<sup>4</sup>See Faith (2000), p.6. Quine’s original intention was to argue against extreme scientific reductionism: “our statements about the external world face the tribunal of sensory experience not individually but only as a corporate body” (1951).

dissertation. A more detailed preview follows, but before this, a few introductory remarks on the nature of artificial evolution models will be worthwhile.

## 1.2 Artificial evolution models

Artificial evolution models are the ‘artefacts’ of this dissertation. For present purposes, artificial evolution will refer to the use of genetic algorithms (GAs) as search processes operating over populations of randomly initialised ‘genotypes’ (usually strings of bits or real numbers). Each ‘generation’ of the GA, genotypes are decoded into ‘phenotypes’ (neural networks, for instance) which are evaluated according to some criteria of fitness. Then, in loose analogy with biological evolution, genotypes corresponding to ‘fit’ solutions progress to subsequent generations, either unchanged, or modified through processes analogous to mutation and recombination. Eventually, given a suitable genotype-phenotype mapping, the hope is that ‘good’ phenotypes will emerge in the population.<sup>5</sup>

In the present dissertation, these models are used to evolve internal structures for agents of various kinds, such that patterns of behaviour in need of explanation are reproduced at some level of abstraction. The evolved structures (and perhaps the search process dynamics by which they are arrived at) can then be used to assess pre-specified hypotheses, and, upon detailed analysis, can also often afford unanticipated insights. It should already be clear (at least superficially) that artificial evolution models are representative of externalism; properties of the internal are moulded, by the selective search process, to engage with properties of the external (this issue is discussed in greater detail in chapter 2).

The use of artificial evolution models outside the domain of SAB/AL remains controversial and is justified in detail in chapter 5, the point in this dissertation at which such issues become significant. Until then, the sketch given here will suffice.

## 1.3 Overview of the dissertation

### 1.3.1 Background

The first task of this dissertation is to describe in detail what is meant by an ‘externalist perspective’ on the relations between behaviour, mechanism, and environment. Chapter 2 undertakes this task by describing a conceptual integration of two closely related themes. The first is the relation of functional properties of behaviour to environmental structure. This theme, evidently externalist in character, is represented here most effectively by Peter Godfrey-Smith’s (1996) ‘environmental complexity thesis’ (ECT). This, the idea that the function of ‘cognition’ is to deal with ‘environmental complexity’, provides a philosophical foundation for this project, yet one that is not entirely unassailable.

Many of the problems with the ECT (at least of those identified here) have to do with the second theme, the need to carefully distinguish between behavioural and mechanistic levels of description, and what this entails for notions of ‘cognition’. Apprehended centuries ago by Vaucanson (and anticipated in section 1.1 above), this distinction is still far from universally appreciated. As Donald Hebb complains:

---

<sup>5</sup>See Mitchell (1997) for a detailed introduction to GAs.

It is inaccurate - worse, it is misleading - to call psychology the study of behavior: It is the study of the underlying processes, just as chemistry is the study of the atom rather than pH values, spectroscopy, and test tubes. (1980, p.1)

Although perhaps unfair on psychology - behaviour can certainly be of interest in and of itself - Hebb's point is clear; behaviour is what mechanism *does*, when observed in interaction with an environment.

The integration of these themes requires a number of associated issues to be discussed in some detail. The pattern of externalist explanation itself is addressed first, both from a contemporary philosophical perspective and in the form of an historical account; it is emphasised that the externalism of the present project is *pragmatic*. This is followed by a discussion of the concept of 'environment', in which the critical distinction is between the environment as perceived by a behaving agent, and as perceived by an external observer. Also important is the close coupling of perception and action in agent-environment interaction patterns (for example Gibson, 1966, 1979), and W. Ross Ashby's (1956) 'law of requisite variety' (LRV), an intuitive summary of which is that 'only variety can destroy variety'.

These themes are pulled together in an extended critique of Godfrey-Smith's formulation of the ECT, an exposition of which occupies the second half of chapter 2. Although this discussion cannot be fully anticipated here, the essential idea is as follows. The ECT as described by Godfrey-Smith faces the problem of accounting for the role organisms play in 'constructing' the features of the environment for which 'cognition' is supposed to be a response. The difficulty is that the slippery notion of 'cognition' covers both the functional properties of behaviour (to which the adaptationist claims of the ECT attach), and also the underlying behaviour generating mechanisms. In short, Godfrey-Smith's version of the ECT rides roughshod, in important places, over the distinction between behaviour and mechanism, with 'cognition', in his version, apparently standing for some notion of undifferentiated 'behavioural/mechanistic complexity'. The critique that I present tries to remedy this problem, a task which involves drawing a series of relations between the behaviour/mechanism distinction and the associated issues (introduced above) of 'construction', perception, action, and environment; another part of the argument involves an extended analogy between the ECT and Ashby's LRV.

With this integration, chapter 2 achieves three objectives. The first is to motivate the development of the artificial evolution models that follow; such models allow concepts such as 'construction' and 'mechanism' to be articulated with much greater facility than is possible at the level of philosophical discourse; they allow the relations between behaviour, mechanism, and environment to be empirically instantiated, and therefore to be explicitly traced. The second is to introduce a range of concepts and themes necessary for the interpretation of these models; each of the themes mentioned above reappears in a variety of contexts throughout the dissertation. In addition, the integration itself - described in terms of the critique of the ECT - constitutes a contribution in its own right, inasmuch as it provides an original way of thinking clearly about behaviour, mechanism, and environment and their interrelations. Importantly, the coherence of this conceptual integration (although I hope convincing) is *not* a necessary precondition for the coherence of the subsequent models, each of which can be judged on its own terms.

### 1.3.2 Complexity

Chapters 3 and 4 present the first empirical excursions of the dissertation. They explore a series of artificial evolution models addressed to the hypothesis, directly inherited from the ECT, that environmental complexity can promote the adaptive evolution of internal (mechanistic) complexity. The term ‘complexity’ itself is, of course, rarely free from controversy; what consensus there is seems to describe some kind of middle ground between randomness and order, something not entirely synonymous - nor entirely distinct - from ‘organisation’. After pondering various positions in the literature, chapter 3 settles on some necessarily model-specific interpretations; environmental complexity is identified with ‘variability’, and mechanistic complexity - in the first set of models at least - with ‘memory’.

These first models evolve strategies to play the ‘iterated prisoner’s dilemma’ (a central construct of game theory in which agents must decide whether to ‘defect’ or ‘cooperate’ with each other; mutual cooperation delivers high mutual payoff, but is open to exploitation by invading defectors). It is demonstrated that the evolution of long memories is indeed promoted by environmental variability (or ‘noise’). A detailed analysis is also given in terms of Ashby’s LRV, allowing novel hypotheses concerning the *locus* of variability to be framed and tested.

The second set of models, in chapter 4, involves the evolution of control structures for mobile robots engaged in ‘homing navigation’ behaviour (after Floreano & Mondada, 1996). These models engage closely with many issues prominent in chapter 2, exploring the influence of environmental variability on the complexity of the evolved behaviours *and* of their underlying mechanisms in a situated agent context. High levels of variability again encourage the evolution of complex (flexible, variable) behaviour patterns, and it appears that the underlying mechanisms in these cases undertake a broader integration of sensory data than their counterparts evolved in low-variability environments; these models therefore make use of a relatively broad conception of mechanistic complexity. Also, as part of this investigation it is observed that seductive behavioural level decompositions of the total activity of the robot do not map in any straightforward way onto the underlying mechanistic dynamics.

It is important to understand the contributions of each type of model in the context of the pervasive idea that (natural) evolution is a progressive force driving towards biological complexity. Chapter 3 critiques this idea, sharply separating the ‘causal’ question (how can evolution lead to complexity?) from the ‘teleological’ question (does evolution necessarily tend towards complexity?). I argue that only the former (causal) question can be of interest, since the latter can be both trivially true and trivially false at the same time. The ECT is taken to exemplify a conceptual-level response to the causal question, a response cashed out empirically in these early models.

Each of these enquiries requires the introduction of noise into artificially evolving systems. The consequences of such a practice are of both theoretical and practical significance, and three novel effects are identified here: (1) the evolution of adaptive complexity (as above), (2) the facilitation of the exploration of genotype space, and (3) the acceleration of evolutionary search. These influences are discussed in terms of the dynamics of artificial evolutionary search.

### 1.3.3 Towards a theoretical biology

The arguments of the dissertation up to this point take place largely within the discourse of SAB/AL, a domain in which the use of artificial evolution models is generally uncontroversial. The same cannot be said of their use in the context of theoretical biology, in which formal mathematical modelling is a well established methodological tradition. Artificial evolution models - and computer simulation models in general - offer both advantages and disadvantages to the theoretical biologist. They may be more flexible, better able to cope with complex situations, yet on the other hand they may display ‘explanatory opacity’; it can be hard to understand how simulation models do what they do, and hard to know what kinds of conclusions can be drawn from their behaviour.

From this point on, the dissertation begins to focus on issues current in theoretical biology and experimental psychology. Following Di Paolo (1999), a pragmatic approach is adopted for the use of computer simulation models in such contexts, in which “the best gain in knowledge comes from becoming aware of the model’s own limitations and not of the model’s closeness to some real pattern” (p.6). With this in mind, chapter 5 goes on to describe an original way in which artificial evolution models can explicate issues in theoretical biology. The particular concern is with models that operate at the level of situated perception and action in spatiotemporally structured environments; these models are called ‘individual-based optimal situated’ (IOS) models, and the idea is that these models constitute an unorthodox complement to the methodology in theoretical biology of ‘optimal foraging theory’ (OFT).

Orthodox OFT is an attempt to understand both functional and mechanistic aspects of animal (foraging) behaviour by treating observed behaviour as ‘optimally’ adapted to an environmental problem (Stephens & Krebs, 1986). It is argued that orthodox (classical, equational) OFT models carry with them several substantial ‘framework assumptions’ which limit their explanatory flexibility in a variety of ways; for instance, they require that internal mechanism be understood in terms of the operation of ‘decision variables’ arbitrating between distinct behavioural options. In each case the benefits offered by IOS models are discussed in detail, most of which flow from an increased ability to model subtle but significant agent-environment interactions mediated by situated perception and action. These benefits are set against the risk of explanatory opacity incurred by all simulation models, and one conclusion is that IOS models may find their best use in illuminating the functional potential of very *simple* mechanisms.

Despite the change in direction marked by chapter 5, the major themes of the dissertation persist. Both orthodox OFT and its unorthodox complement are strongly externalist, and both interrogate the relations between behaviour, mechanism, and environment (although, as I will argue, the latter may be expected to do so with greater success). Indeed, a useful way to think of the IOS modelling strategy is as a methodological expression of the (critiqued) ECT; many of the general explanatory targets of IOS models coincide with themes important in the revision described in chapter 2, notably the distinction between behaviour and mechanism, and the various ways in which agents may be considered to ‘construct’ their environments. Nevertheless, the independent contribution of chapter 5 remains its role in the methodological unification of artificial evolution models with theoretical biology, a unification that appears (justifiably) to be of increasing scientific significance (see, for example, Bullock, 1998; Noble, 1998; Di Paolo, 1999).

With this methodological foundation in place, chapter 6 goes on to describe the theoretical

context of ‘behaviour coordination’; a context which, in various guises, occupies the remainder of the dissertation. At its most general, behaviour coordination is the problem of ‘how to do the right thing’, and it fits into the overarching structure of the dissertation in that (a) it can be interpreted as a response to complex environments (*sensu* the ECT), and (b) much of the relevant discourse concerns the nature of the underlying mechanisms.

The chapter opens with a multi-disciplinary review of the behaviour coordination literature, with contributions from ethology (Tinbergen, 1950, 1951; Baerends, 1976), theoretical biology (McFarland & Sibly, 1972, 1975; Dawkins, 1976), artificial intelligence (Miller, Galanter, & Pribram, 1960), and SAB/AL (Brooks, 1986, 1994; Maes, 1990) amongst others. A problem with many of the ideas described in this review is their assumption of a need for an internal arbitration mechanism operating over a pre-existing repertoire of internalised behavioural correlates (the clear parallels with the ‘decision variable’ concept in OFT are discussed in detail). They also maintain unadventurous conceptions of the intimate relationship between perception and action, and - by and large - sit uneasily within externalist explanatory frameworks.

Chapter 6 then describes a novel alternative, originating in the work of Braitenberg (1984), in which behaviour coordination arises from continuous agent-environment interaction, mediated by tightly coupled perception and action. An IOS model is described and analysed in which GAs are used to evolve the shapes of a set of simple sensorimotor links operating continuously and in parallel, to control simple agents faced with a straightforward foraging task. The objective is to demonstrate that effective behaviour coordination can arise *without* explicit arbitration between internal behavioural correlates (or to put it another way, in the absence of decision variables). The significance of this model lies not in its biological fidelity (which is not asserted), but in the fact that it constitutes a proof of concept of an (externalist) account of behaviour coordination which avoids the theoretical obstacles that beset a large proportion of orthodox approaches. It also provides the empirical template for the relatively specific investigations that follow.

### 1.3.4 Interference and the matching law

The first of these investigations concerns ‘interference’, defined in the theoretical biology literature as the reversible decline in intake rate (of prey) due to the presence of competing predators (Goss-Custard, 1980; Sutherland, 1983). This is a concept deployed in many contexts throughout theoretical biology, but for the purposes of this dissertation its primary significance is with respect to the ‘ideal free distribution’. This, the optimal distribution of rate-maximising predators across a ‘patchy’ environment, is reached when all predators experience the same intake rate, such that some proportional relation exists between patch ‘quality’ and predator density (Fretwell & Lucas, 1970; Fretwell, 1972).<sup>6</sup> Any derivation of the ideal free distribution therefore requires knowledge of the per-predator intake rate in each patch, and this is the relation captured by the ‘interference function’.

Theoretical models of the interference function have a long history, yet the modelling strategies employed remain controversial (Van der Meer & Ens, 1997). Although the details of this debate will not be pre-empted here, in brief, chapter 7 argues that the IOS methodology is well placed to move matters forward in virtue of its distinguishing features of consistency with the principles of

<sup>6</sup>I am simplifying here. A relatively detailed discussion of the ideal free distribution will keep for the present.



OFT, and structuring at the level of perception and action.

The particular IOS model developed in this chapter is an extension of the behaviour coordination model of chapter 6, an extension in which relations between individual and collective behaviour play a significant role. Isolated agents are evolved to forage for resources (in a single patch environment), and, once effective foraging is established, an interference function is derived by recording intake rates in groups (of identical agents) of varying sizes.

The implications of this basic model of interference for theoretical biology are discussed in detail, but the most intriguing insights afforded in this context derive from an extension of this basic model in which foraging behaviour is not only *analysed* in group contexts, but also *evolved* in such circumstances. In addition to revealing influences on the nature of interference that would otherwise have remained hidden, this modification also provides an example of a behaviour, irrational for an isolated individual, yet rational for the individual as part of a group of conspecifics. This example is discussed in detail in terms of self-sustaining patterns of agent-environment interaction dynamics, and in terms of ideas of ‘construction’ as they appear in the (critiqued) ECT.

The explanation of irrational behaviour in terms of adaptation to group situations presents a different slant on the issue of behaviour coordination; instead of ‘how to do the right thing’, the onus is on accounting for doing the wrong thing. The same idea is taken further in chapter 8, in the context of the ‘matching law’. This law, first formulated in 1961 by the experimental psychologist R.J. Herrnstein, formalises the observation that many animals (including humans) often match the frequency of their response to different stimuli in proportion to the reward obtained from each stimulus type. Importantly, the matching law is descriptive rather than normative, and observed matching behaviour is not always optimal. Two questions therefore arise: Why match at all? What mechanisms might underlie matching behaviour?

Chapter 8 investigates the possibility that individual matching (and its potential suboptimality) might be a consequence of foraging behaviour adapted to a group context, an idea that derives indirectly from a series of analogies between the matching law and the ideal free distribution. The IOS model of this chapter - a direct extension of the preceding ‘interference model’ - provides strong support for this idea. It also offers an example of matching behaviour without there being any underlying ‘mechanism of matching’, a demonstration which again emphasises the importance of distinguishing between behaviour and mechanism; matching, in this model, arises from the same sensorimotor interactions constitutive of interference.

Both the interference and matching models offer insights difficult, if not impossible, to attain without the substrate of sensorimotor interaction provided by the IOS methodology. Not all of these insights can be anticipated at this stage, however they bring with them the consequence that the phenomena addressed by the models are not quite the same as those addressed by orthodox approaches. This is especially so in the case of matching; the matching behaviour of chapter 8, arising from simple situated foraging interactions, is distinct in many ways from the matching of psychology textbooks. The immediate disadvantage of this is that direct engagement with the orthodox literature is not easy; both chapters 7 and 8 are required to go into considerable detail to explain how their respective models relate to, and contribute to, the orthodox understanding of their subject matters. The benefits, however, derive from the suspicion that the designation of phe-

nomena as deserving of explanation is often inextricable from the methods addressed towards their elucidation. By addressing phenomena at unusual levels of abstraction, assumptions underlying their existence *as* phenomena may be exposed. The IOS model of matching, for instance, not only undermines the idea that matching behaviour need be subserved by some dedicated ‘mechanism of matching’, but also questions the preconception that the ‘matching law’ describes a relatively isolable pattern of behaviour in the first place.

Less controversially, both the interference and matching models also relate back to the ECT by illustrating relatively complex behaviours adapted to relatively complex (group) environments.

### 1.3.5 A return to orthodoxy

The final empirical installment of the dissertation is discontinuous with much - but by no means all - of what precedes it. The model described in chapter 9 does *not* operate at the level of situated perception and action (it is more akin in its agent-environment structure to the game-theoretic models of chapter 3 than it is to the preceding IOS models). Moreover, artificial evolution occupies only a subsidiary role (described below); the strategies are relatively well pre-specified, and thus the model itself more readily interpretable from an internalist perspective. The advantage of this approach is that, again in contrast to the preceding IOS models, it effortlessly engages with the relevant theoretical biology literature: the phenomena addressed in this model *are* the same phenomena addressed by orthodox approaches, for the approach itself is also (for the most part) orthodox.

Nonetheless, the objective of this chapter remains that of understanding individual suboptimal behaviour in terms of adaptation to group situations, and the focus is still on the matching law. Indeed, the idea at the heart of the chapter is similar to that explored in the IOS matching model; the suggestion is that (potentially suboptimal) individual matching behaviour may follow from foraging strategies that lead groups of foragers to the ideal free distribution. As before, this idea follows from a series of analogies between the matching law and the ideal free distribution, but unlike before, in this case the analogies are interpreted relatively directly.

In its orthodox guise, this idea is not original. In one example from the literature, Thuijsman, Peleg, Amitai, and Shmida (1995) claim that a particular foraging strategy -  $\epsilon$ -sampling - can underlie both phenomena. The first contribution of chapter 9 is an assessment of this claim, which finds it to be false. A related yet novel strategy,  $\omega$ -sampling, is then introduced and shown to succeed where  $\epsilon$ -sampling fails. Both strategies involve simple rules applied to patch-switching decisions, the significant difference is that the latter maintains a much more powerful representation of perceived environment value. These results are accompanied by a detailed discussion of how both strategies, and the kinds of matching they support, relate to the matching and interference of the preceding IOS models.

In marked contrast to the IOS methodology, artificial evolution is *not* used to engender strategy properties in this model, its role is rather to elucidate pre-existing properties, an application justified in detail in chapter 9 itself. There is, however, one final insight afforded by the model which *is* entirely dependent on artificial evolution, an insight which echoes the earliest findings of the dissertation. It is demonstrated that the introduction of environmental variability leads to the evolution of long strategy memories, just as observed in the context of the iterated prisoner’s

dilemma (chapter 3), and, as such, once again consonant with the externalist perspective of the ECT.

#### 1.4 Summary and contributions

This dissertation aims to elucidate, from an externalist perspective, the relations between behaviour, mechanism, and environment, as they arise in a variety agent-environment systems. This project is tackled at conceptual, methodological, and empirical levels, and original contributions are offered at each of these levels.<sup>7</sup> It should be warned that the empirical contributions of this dissertation in no sense exhaust the potential offered by its conceptual and methodological structure, however it is hoped that this is taken as testimony to the richness of this potential, and not as an indication of empirical indiligence. The main points of each individual chapter - whether pertaining to a concrete model or to an abstract discussion - can be understood in isolation, nevertheless their force is considerably enhanced in the context of the dissertation as a whole.

The conceptual framework offered consists primarily in an account and critique of Godfrey-Smith's (externalist) ECT, this is therefore the first contribution of the dissertation (chapter 2). The central features of this framework are the interpretation of behavioural and/or mechanistic complexity in terms of adaptation to environmental variability, and (of course) the distinction between behaviour and mechanism itself; these features constitute orienting themes of the project as a whole.

The first application of this framework arrives in a series of models that evolve strategies to play the iterated prisoner's dilemma (chapter 3). These models demonstrate that environmental variability can promote the evolution of strategy complexity, a contribution echoed in a subsequent model in which 'homing-navigation' behaviour is evolved for a mobile robot (chapter 4). In this case enhanced complexity is observed at both behavioural and mechanistic levels.

These early contributions are placed in the context of the semantics of complexity itself, and of the pervasive idea that (natural) evolution tends towards complexity. They are attended by the subsidiary discoveries that noise can abet artificial evolution as a search process by (a) facilitating the exploration of genotype space, and (b) accelerating the search process itself.

The conceptual foundations of chapter 2 are reformulated, in chapter 5, as a methodological framework portrayed as an unorthodox complement to optimal foraging theory. This reformulation, the IOS methodology, not only structures the remainder of the dissertation but constitutes a significant contribution to the ongoing - and urgent - task of integrating research in SAB/AL with the mature discipline of theoretical biology. IOS models are shown to be well-placed to address many of the shortcomings of orthodox OFT, in virtue of their ability to model subtle but significant agent-environment interactions mediated by situated perception and action.

Chapter 6 provides an example of an IOS model in the form of a minimal model of behaviour coordination. The primary contribution of this model is to illustrate that effective behaviour coordination can arise without the need for explicit arbitration over internal behavioural correlates.

Chapters 7, 8, and 9 are concerned with specific issues in theoretical biology and experimental psychology, issues united by their concern with individual choice and its collective consequences. Chapter 7 describes an IOS model of inter-predator interference which significantly extends those

---

<sup>7</sup>All contributions mentioned in this section are to be interpreted as original.

available in the literature, directly in virtue of its ‘situated optimal’ nature. Also, a novel situation is analysed in which behaviour adapted to a group context entails the irrational actions of isolated individuals.

Chapter 8 takes this idea further, supporting the hypothesis that adherence to Herrnstein’s matching law can be understood as a consequence of optimal foraging in a shared environment. The very idea of a ‘mechanism of matching’ is also challenged, the nature of this challenge a significant contribution in the light of the ongoing debate (in psychology) over the generation of matching behaviour. Additionally, both chapters 7 and 8 speak to the field of SAB/AL directly, potentially extending its grasp of both group behaviour and individual choice. These chapters also relate back to the ECT by illustrating examples of complex behaviours as adaptations to complex environments.

Chapter 9 is different, eschewing the IOS methodology by marginalising the role of artificial evolution and focussing on pre-specified foraging strategies that do *not* depend on situated perception and action. It nevertheless shares with previous chapters a desire to understand individual irrationality in terms of adaptation to group situations. Two strategies are compared on their ability both to lead groups to the ideal free distribution, and to lead individuals to match; the chapter’s central contribution is to repudiate previous claims of this kind for one of the strategies, and to describe a successful (and novel) alternative. It also contributes to the reconciliation of optimality-based and individual-based modelling in theoretical biology, albeit at a different level of abstraction than the IOS methodology. Finally, it recapitulates the findings of chapters 3 and 4 with respect to the influence of environmental variability on the evolution of strategy complexity.

Chapter 10 summarises these contributions with the benefit of hindsight, addresses the wider implications and general limitations of the project, and poses some challenges for future research.

## Chapter 2

# An externalist perspective on behaviour, mechanism, and environment

---

We don't see things as they are, we see them as we are.  
- Anais Nin

The purpose of this chapter is to explain what is meant by an externalist perspective on the relations between behaviour, mechanism, and environment, in order to provide a conceptual framework for the remainder of this dissertation. This will involve a conceptual integration of two closely related themes, the first of which is the relation of functional properties of behaviour to environmental structure. This theme is represented here most effectively by Peter Godfrey-Smith's (1996) 'environmental complexity thesis'; the idea that the function of 'cognition' is to deal with 'environmental complexity'. The second is the distinction between behavioural and mechanistic levels of description; recall from chapter 1 that behaviour is what mechanism *does*, when observed in interaction with an environment. The integration of these themes, which constitutes the primary original contribution of this chapter, will serve to motivate the empirical artificial evolution models of the chapters that follow.

A number of other themes, also prominent throughout this dissertation, are woven into this integration. These include the relationship between the external environment and the *Umwelt* of an agent, the essential inseparability of perception and action, W. Ross Ashby's (1956) 'law of requisite variety', and the application of externalist explanation to agent-environment systems in general. The elucidation of these themes is accompanied by an historical account of externalist explanations of behaviour, mechanism, and environment, from empiricist epistemology to behaviourist learning theory and beyond. This account is justified in part by the empirical diversity of this dissertation, but it is also through this history that a development of the conceptual integration itself may be traced, an integration which culminates in a detailed account - and critique - of the environmental complexity thesis itself.

The remainder of this chapter largely follows this division of labour. A discussion of key themes precedes the historical commentary, which is followed by a detailed exploration of the environmental complexity thesis. The chapter ends with a more specific discussion of its relation to the remainder of this dissertation.

## 2.1 Internalism, externalism, and enlightened interactionism

The perspectives of internalism and externalism can be used to classify large bodies of thought. Internalist theories give explanatory precedence to factors *internal* to a given system, and often entail appeals to rational, *a priori*, or perfect forms. Externalist theories avoid this kind of appeal, and attempt to explain properties of the internal in terms of their relation to the external. These two positions define a third - construction - in which properties of the external are explained in terms of their relation to the internal. My usage of these three terms - internalism, externalism, and construction - derives largely from Godfrey-Smith (1996). I will begin by discussing the first two perspectives, construction will receive further attention in section 2.8.

### 2.1.1 Internalism and externalism

Internalist and externalist positions may be discerned in almost all intellectual endeavours. This section indicates how these positions are defined in the context of biology, psychology, epistemology, and SAB/AL itself, these being disciplines in which the relations between internalism and externalism are cast in the context of agent-environment systems, and in which the terms behaviour, mechanism, and environment have intuitive application (this is less true for epistemology, which is included more for its historical significance).

In biology, externalism is exemplified by the explanation of complex organism properties in terms of their ‘fit’ to properties of the environment. Anti-externalist biological positions may take the form of critiques of unfettered externalism, as in Gould and Lewontin’s (1979) celebrated deconstruction of the ‘Panglossian’ adaptationist programme. More thoroughgoing internalism may be identified in the search for biological laws of form, a tradition presently associated with the structuralism of Goodwin (1994) and others, which has its historical roots in the ‘rational morphology’ of Georges Cuvier (1827) and D’Arcy Thompson’s *On Growth and Form* (1917).

In psychology, both nineteenth century associationism and twentieth century behaviourist learning theory are of externalist character, their essence is in explaining the structure of behaviour with reference to patterns of information derived from the environment (both will be discussed in more detail in section 2.4). More recently, the ecological psychology of James Gibson is also arguably externalist. For Gibson, ‘affordances’ of the environment provide sufficient information for the control of adaptive behaviour, however they are also argued to be perspective dependent, and therefore in some sense to transcend the dichotomy between agent and environment (Gibson, 1966, 1979).

Internalist psychology, on the other hand, stresses the importance of innate mental structure, a position epitomised by Chomsky’s argument from the ‘poverty of the stimulus’ (Chomsky, 1959). This is the argument that the organised complexity evident in the speech of a developing child is too great to be explained in terms of the structure available to the child in her linguistic environment. Chomsky believes that some kind of innate ‘language acquisition device’ must also contribute. Cognitive science in general is often considered internalist, but often this is more a matter of emphasis than of substance. Cognitive theories do indeed focus on internal psychological processes (usually construed as symbol manipulation of some kind), and to the extent that cognitive science is deployed against behaviourist learning theory it is certainly anti-externalist. However, it is rare that the environment is excluded from cognitive explanation, and indeed there

are prominent cognitive scientists whose emphasis is evidently externalist (for example Simon, 1956, 1988). Cognitive science will be discussed more thoroughly in section 2.2, and its relation to behaviourist learning theory in section 2.4.3.

The empiricist epistemology of Locke (1632-1704), Berkeley (1685-1753), and Hume (1711-76) is strongly externalist. Empiricism asserts that the origin of knowledge is to be located in patterns of sense data. An important qualification is that empiricist philosophers generally refrained from discussing the external environment itself, so for externalism to sensibly apply it is necessary to consider sensory experience as the ‘environment’ of the mind. Empiricism has been traditionally opposed by the rationalist school of epistemology, associated with Leibniz (1646-1716). Like Chomsky after him, Leibniz argued for the necessity of mental ‘pre-structure’ in the formation of beliefs and judgements.<sup>1</sup>

Closer to home for present purposes, internalism and externalism can be contrasted in the disciplines of SAB/AL. The search for ‘laws’ of complexity and/or self-organisation often flies an internalist flag. Indeed, cellular automata work of this flavour does not specify an environment of any kind (Langton, 1989). A less extreme position is taken by Kauffman (1993) who is concerned with locating the supposed ‘sources of order’ upon which (externalist) selective processes operate. Externalist work in these genres can be identified most obviously in the use of genetic algorithms to evolve internal structural properties (see, for example, Husbands & Meyer, 1998) and indeed this chapter will end by arguing that such models provide appropriate tools for the empirical application of the conceptual framework developed in what follows.

### 2.1.2 Arbitrariness and enlightened interactionism

Of course, the distinction between externalism and internalism presupposes a distinction between external and internal. As many have argued, this can be difficult to justify, and theories which take as their explanandum the *existence* (whether apparent or ‘real’) of a distinction between external and internal naturally resist classification as either externalist or internalist.<sup>2</sup> A relevant example is provided by Varela (1997), who argues that the attempt to understand the “cognitive self” of an organism is identical with the task of understanding how the boundaries between internal and external are drawn with regard to that organism:<sup>3</sup>

The cognitive self is the manner in which the organism, through its own self-produced activity, becomes a distinct entity in space, but always coupled to its corresponding environment from which it remains nevertheless distinct. A distinct coherent self which, by the very same process of constituting itself, configures an external world of perception and action. (1997, p.83)

Similarly, various readings of Heideggerian philosophy suggest that subject-object divisions between an organism and parts of its environment (for example a hammer) are only manifest when something about the system as a whole misfires (the head of the hammer falls off in the act of hammering, or the hammer is applied to the task of window cleaning). In such instances, the usual

<sup>1</sup>It is this very same Leibniz who was caricatured by Voltaire as ‘Doctor Pangloss’ - the inspiration for Gould and Lewontin’s (1979) savaging of *externalist* adaptationist biology!

<sup>2</sup>Explanandum: the target of explanation, as opposed to explanans: the means of explanation.

<sup>3</sup>This is just one aspect of a significant body of work developed over recent years by Francisco Varela and Humberto Maturana; see also Varela (1979), Maturana and Varela (1980), and Varela, Thompson, and Rosch (1991).




Figure 2.1: An agent and its environment (adapted from Beer, 1995).  $\mathcal{E}$  represents the environment,  $\mathcal{A}$  represents the agent, and  $S$  and  $M$  represent the processes of sensation and motor action by which agent-environment interaction is mediated.

mode of encounter with the world - what Heidegger calls “readiness-to-hand” - breaks down into a mode of encounter in which the properties of the system resolve into objects and subjects, into external and internal, into hammers and hammerers. Otherwise, when “readiness-to-hand” obtains, there is no external, no internal, and no division, just hammering (see, for example, Wheeler, 1996; Dreyfus, 1991). Nor is this kind of argument restricted to philosophy. In evolutionary biology, for example, the arbitrariness of an internal/external division is illustrated by the concept of the ‘extended phenotype’ (Dawkins, 1982). For Dawkins, the concept of phenotype is no longer constrained to apply only to aspects of an organism that fall within the boundary marked by the organism’s skin, rather, phenotypic effects are identified at *all* scales, ranging from intracellular biochemistry to the construction of beaver dams.

In the light of these arguments, this dissertation will adopt an externalist perspective to the extent that it remains conceptually consistent with the arbitrariness of a division between external and internal, and also whilst allowing for the possibility of accounting for interactions between external and internal factors. This kind of externalism can be called a *pragmatic* externalism.

How is this possible? To take the first qualification first, consider the conceptual framework illustrated in figure 2.1. In this diagram, adapted from Beer (1995), agent  $\mathcal{A}$  and environment  $\mathcal{E}$  are engaged in continuous mutual interaction, and, by externalism, one may attempt to explain properties of  $\mathcal{A}$  with reference to  $\mathcal{E}$ . Yet there is a certain arbitrariness associated with the way in which agent and environment are distinguished; why not select the area  $\mathcal{A}$  from the top-left corner of  $\mathcal{E}$ , or from the bottom-left corner? As Beer says, “an equally legitimate view is that [...]  $\mathcal{A}$  and  $\mathcal{E}$  are merely components of a single autonomous dynamical system  $\mathcal{U}$  whose state variables are the union of the state variables of  $\mathcal{A}$  and  $\mathcal{E}$  and whose dynamical laws are given by all of the internal relations (including  $S$  and  $M$ ) among this larger set of state variables and their derivatives. [...] Neither of these perspectives is intrinsically better than the other, we can switch between them as appropriate” (1995, p.131). Now, consider that some factor  $X$  may be identified as the key causal ingredient in the explanation of a certain pattern of activity. The point of this framework is to argue that the causal role identified for  $X$  in the explanation of the activity pattern should not



depend on how  $\mathcal{U}$  is carved up into  $\mathcal{A}$  and  $\mathcal{E}$ , if, indeed, it is carved up at all.

However, it need *not* be the case that the ease with which  $X$  can be identified in the first place is independent of this carving up. The role of  $X$  may be easier to apprehend if it is part of  $\mathcal{E}$  rather than  $\mathcal{A}$ ; that is, if it forms part of an externalist explanation. Of course, it remains possible that  $X$  could form part of an *internalist* perspective, were the boundaries between  $\mathcal{A}$  and  $\mathcal{E}$  to be redrawn, and, furthermore, it remains possible that other key ingredients (for example  $Y$ ) may be accessible only (or more easily) from an internalist perspective (ie. when analysed as part of  $\mathcal{A}$  rather than  $\mathcal{E}$ ). This is part of what is meant by the externalism of this dissertation being pragmatic. It is the attempt to locate key causal ingredients in the explanation of patterns of activity by resolving the system into components  $\mathcal{A}$  and  $\mathcal{E}$ , and focussing in large part on  $\mathcal{E}$ .

We are not quite there yet. It has been admitted that an externalist perspective may not - on its own - be able to elucidate all properties of the patterns of activity of a given system. Some may only be visible from an internalist vantage, and still others may only be understandable in terms of *interactions* between internal and external, as the doctrine of ‘enlightened interactionism’ has it (Oyama, 1985). But this is where the combination of a pragmatic externalist perspective with the construction of concrete models has particular purchase. The models of the following chapters separate intuitively into ‘agent’ and ‘environment’ components, and, by virtue of being concrete models, interactions between agent and environment can be captured in a way naturally amenable to empirical analysis, even though the explanatory emphasis of the models may remain externalist.

As a brief example of the significance of these ideas, consider Richard Lewontin’s critique of the concept of an ecological ‘niche’ (this example will be returned to in more detail in section 2.5.2). For Lewontin, the (strongly externalist) attempt to understand agent-side properties in terms of environmental structure (the niche) is fundamentally flawed because such structure can be continuously modified by agent activity. Agent and environment causally interpenetrate each other, so that niches cannot pre-exist (agent-side) adaptations.<sup>4</sup> However, although clearly anti-externalist, and convincingly so, this argument is in fact quite consistent with the pragmatic externalism proposed here. For present purposes it may simply be accepted that agents influence their environments by their own activity. Elucidation of this process is part of the explanatory duty of models that incorporate (enlightened) interaction and mutual specification between agent and environment, and may be cast as an empirical matter. Moreover, a denial - or indeed any straightforward refutation - of Lewontin’s argument, would necessitate the reification, and therefore defeat the conceptual arbitrariness, of the distinction between external and internal.

### 2.1.3 Summary

Varieties of externalism can be distinguished by their degree of consistency with the conceptual arbitrariness of distinguishing between external and internal (as we will see, some versions of externalism do *not* satisfy this condition, and should be regarded with suspicion). Internal/external divisions can nevertheless be of pragmatic value, but are necessarily accompanied by the suspicion of enlightened interactionism, a suspicion partially alleviated if an externalist perspective is to be cashed out in the form of concrete models. These ideas are relevant to the present dissertation both by providing an essential background for the conceptual integration of behaviour, mechanism, and

<sup>4</sup>This argument is not the same as his ‘Panglossian’ critique of adaptationism (Gould & Lewontin, 1979).

environment, and also by capturing the general mode of explanation adopted by the concrete models of the following chapters.

## 2.2 Cognitive science and cognitive mechanism

The development of the externalist project of this dissertation, and in this chapter in particular, will often be related to, and contrasted with, the relatively internalist perspective of orthodox cognitive science. It is worth expanding a little on what this perspective involves.

By ‘cognitive’ I refer to two characteristics in particular. First, the computational theory of mind associated with the ‘physical symbol system hypothesis’, which states, in essence, that a physically instantiated symbolic system (such as a computer) is both necessary and sufficient for general intelligent action (Newell & Simon, 1976). Second, the division of internal activity into ‘sense-think-act’ sequences. David Marr, for example, was concerned with how a three-dimensional world model could be constructed from two-dimensional raw images, such that ‘central reasoning devices’ could utilise the model in the generation of appropriate action (Marr, 1982). These characteristics come together in the classic definition of Lachman, Lachman, and Butterfield (1979), for whom cognitive science is about “how people take in information, how they recode and remember it, how they make decisions, how they transform their internal knowledge states, and how they translate these states into behavioral outputs” (p.99). Throughout this dissertation the term ‘cognitive’ will be used in this sense, unless explicitly qualified otherwise.

In the nineteen-sixties, these principles were embodied in the robot ‘Shakey’ (see Nilson, 1984).<sup>5</sup> Shakey was able to construct an internal world model of its environment of blocks and wedges and use this model to achieve pre-specified goals, for example pushing a block to a particular location. At first, Shakey was taken to vindicate the cognitive vision, but subsequent developments failed to live up to early expectations, and nowadays Shakey is remembered less as auguring a new dawn in intelligent robotics, and more often as an illustration of the severity of the problems encountered in constructing complete embodied cognitive agents (see, for example, Brooks, 1999).

The inglorious history of Shakey and its descendants partly accounts for the substantial criticism endured by the cognitive programme over recent years, both from within the field (for example Dreyfus, 1972; Pylyshyn, 1985; Dennett, 1987, 1991; Harnard, 1990), and from broader perspectives (Maturana & Varela, 1987; Varela et al., 1991; Port & van Gelder, 1995; Clark, 1997; van Gelder, 1998). It is not part of the present project to detail these criticisms; the objective is simply to contrast the cognitive perspective (which, despite these criticisms, remains influential) with the perspective of this dissertation. I share with cognitive science a desire to elucidate the mechanisms underlying the generation of behaviour, but when I speak of such mechanisms I am *not* implying that they are ‘cognitive’ in the sense above. In this dissertation, mechanism is to be interpreted quite straightforwardly as the agent-side internal structure that subserves the generation of behaviour. The following section fleshes this out with a more explicit discussion of the relationship between behaviour and mechanism.

<sup>5</sup>The origin of such attempts - and the nascence of the discipline of ‘artificial intelligence’ itself - can be traced at least as far back as the now legendary Dartmouth Conference of 1956, attended by (amongst others) such luminaries as McCarthy, Minsky, Newell, and Simon.

## 2.3 Behaviour, mechanism, and environment

This section introduces some of the key themes of this dissertation: the distinction between behaviour and mechanism, the relationship between the ‘external’ environment and the *Umwelt* of an agent, and the inseparability of perception and action.

### 2.3.1 A leitmotif

The essential distinction between behaviour and mechanism is a leitmotif of this dissertation. This distinction relies on two definitions (introduced in brief form in chapter 1). First, *behaviour* is defined as *observed ongoing agent-environment interactivity*. Second, *mechanism* is defined as the *agent-side internal structure* that subserves this interactivity. These definitions are offered as necessary commitments of this dissertation, not as self-evident truths, and the essential consequence of this leitmotif, in their light, is that the structure of mechanism should not be discussed in the language of behaviour. To do so would be to commit a category error.

To elaborate. All behaviours (eating, swimming, fleeing, building-a-house) depend on a continuous pattern of interaction between agent and environment; there can be no eating without food, no building-a-house without bricks, no swimming without water. In addition, it is ultimately up to the external observer to decide which segments of agent-environment interactivity warrant which behavioural labels. In principle, different observers may (i) privilege different junctures in observed activity, and (ii) label the same segments in different ways. Either would cause problems for any explanation of mechanism framed in the language of behaviour. To emphasise the general point: behaviour is a product of the joint activity of agent, environment, and observer, thus *the (agent-side) mechanisms underlying the generation of any behaviour should not be assumed to be identical to the behaviour itself*.<sup>6</sup>

One brief clarification must accompany these arguments. The distinction between behaviour and mechanism is clearly predicated on there existing a distinction between agent and environment, and to the extent that the former is arbitrary - as suggested in section 2.1.2 - there must exist the suspicion that the latter is also arbitrary. However, the present project is targeted at the understanding of relations between behaviour, mechanism, and environment, *as they arise in agent-environment systems*. For this project to make sense in the first place, as a project, the assumption must be in place that agent and environment can be pragmatically and reliably distinguished, with identifiable structure on both sides. In the context of this project, then, the leitmotif of distinguishing between behaviour and mechanism is *not* undermined by the threat of arbitrariness.

An instructive example of the importance of this leitmotif comes from Lorenz’s (1937) observations of the ‘parenting behaviour’ of ducks. Mother ducks engage in a number of different patterns of interaction with their offspring, and Lorenz subsumed all such observed patterns under the label of ‘parenting behaviour’. This is clearly valuable from the perspective of descriptive clas-

<sup>6</sup>For other discussions of the importance of this leitmotif, see Hendriks-Jansen (1996), Clancey (1991). The distinction between behaviour and mechanism is also related to the idea of ‘multiple realisability’, which is the argument that any given function could (in principle) be realised by a very large variety of different mechanisms (Putnam, 1967; Fodor, 1974; Sober, 1990). This notion is sharpened by Simon, who argues that to the extent that behaviours are effectively adapted, they will “reflect characteristics largely of the outer environment [...] and will record only a few limiting properties of the inner environment [ie. mechanism]” (1988). These arguments build on the basic distinction between behaviour and mechanism in interesting ways, however, it is beyond the scope of this chapter to develop them in the detail that they invite.

sification, however, it should *not* be taken as reason to believe in the existence of any ‘behavioural icon’ for parenting behaviour, internal to the mother duck, as a mechanistic explanatory locus. And indeed, as Lorenz subsequently discovered, from the perspective of the mother duck the various interaction patterns are all triggered by quite different stimuli. The only point at which they ‘intersect’ in any sense is on the duckling itself, as an object in the eyes of the external observer.

One important consequence of this leitmotif is that if a behaviour appears *complex* (to an external observer), this does *not* imply that the underlying internal mechanisms are also complex. The classic illustration of this is Herbert Simon’s description of an ant on a beach (Simon, 1988). The internal mechanism of this (hypothetical) ant consists of a simple obstacle-avoidance rule - if there is a clump of sand to the left, go right, and vice-versa. Thus the ant responds to every tiny clump of sand, veering first left then right as it negotiates its terrain. Simon’s point is that from the perspective of the external observer - who cannot perceive the small-scale heterogeneity of the beach surface, and who is not aware of the simplicity of the ant’s ‘algorithm’ - the trajectory traced by the ant is strikingly and perhaps irretrievably complex.<sup>7</sup>

The distinction between behaviour and mechanism is closely related to that made between *functional* and *operational* modes of discourse, which will also be appealed to in the course of this dissertation (Varela, 1979; Di Paolo, 1999). Briefly, operational descriptions consist of tracing lawlike (nomic) links between elements that comprise the domain in which the phenomena of interest occurs; as such they are mechanistic, but need not attach specifically to behaviour-generating mechanisms. Functional descriptions, by contrast, always *escape* the boundaries of this domain, imposing teleological criteria ultimately derived from the perspective of the external observer. Behavioural descriptions are therefore functional, in view of involving the external observer, but not all functional descriptions are behavioural. Consider, for example, the DNA molecule, for which a variety of functional roles can be assigned - that it encodes for particular amino-acid sequences, that it encodes for eye colour - none of which are behavioural in the sense outlined above. Di Paolo (1999) notes that functional and operational descriptions can constrain each other but can never be equated, in particular; “preferable functional explanations should be those which are constrained by a lack of contradiction with existing operational descriptions” (p.18).

### 2.3.2 Environmental structure

A series of questions immediately arise: how is the distinction between behaviour and mechanism to be explanatorily traversed? What is their relation? How can knowledge at one level lead to insight at the other? A first step may be taken by rephrasing Di Paolo’s assertion of the previous section, to say that behavioural (functional) descriptions should not contradict mechanistic (operational) descriptions. However, the context of agent and environment involved in any discussion of behaviour allows more to be said, through the employment of a refined concept of environment structure to tease apart the different ways in which behaviour and mechanism rely on environment and observer.

The essential distinction here is between the *external environment* and the *Umwelt*. The external environment is the environment as it appears to us, as external observers (following

<sup>7</sup>Agre and Chapman (1987) make a similar point, arguing that problems that appear intractable for a disembodied agent are often much simplified when the role of the environment is sufficiently appreciated.

Antonovics, Ellstrand, & Brandon, 1988; Brandon, 1990). The *Umwelt* comprises the space of sensorimotor cues relevant to an organism, it contains those environmental features which constitute stimuli for the organism, to which the organism can potentially muster a response. Jakob von Uexküll (1934), the originator of the concept, describes the *Umwelt* of a tick, which hangs from tree branches and drops onto passing mammals. This tick uses the stimuli of butyric acid, tactile contact, and heat, to guide the various stages of its behaviour (dropping, searching, burrowing):

Of all the influences that emanate from the mammal's body, only three become stimuli and those in definite sequence. Out of the vast world which surrounds the tick, three shine forth from the dark like beacons, and serve as guides to lead her unerringly to her goal. [...] The whole rich world around the tick shrinks and changes into a scanty framework consisting, in essence, of three receptor cues and three effector cues - her *Umwelt*. (1934, pp.11-12)<sup>8</sup>

In making this distinction, there is a potential confusion to be avoided regarding the interpretation of 'external'. Whereas the general perspective of externalism refers to the explanation of the internal (the agent) in terms of whatever surrounds it (see section 2.1.1), the term 'external environment' refers specifically to the environment as it appears to an observer of the entire system to which the perspectives of internalism and externalism may be applied by the observer; in other words, the external environment exists in the *Umwelt* of the observer.

How can this distinction help to clarify the relationship between behaviour and mechanism? Behaviours, being elements in the vocabulary of an observer, are located in the external environment. Consider Lorenz's duck: the descriptions of the various interaction patterns that constitute parenting behaviour are framed in terms of the external environment, and intersect only on the duckling *as an entity in the external environment*. For these behavioural descriptions to be related to mechanism, however, it can be helpful, or perhaps even necessary, that the relations between the external environment and the *Umwelt* be traced - Lorenz had to identify what constituted stimuli for the mother-duck in order to understand its parenting behaviour. To take another example, the catching behaviour of cricketers has been greatly elucidated by appreciating as a stimulus the acceleration of the tangent of elevation of gaze from player to ball; if this acceleration is kept at zero, the cricketer *will* meet the ball before the ball meets the ground (McLeod & Dienes, 1996).

There may nevertheless be a problem with this kind of argument. Should mechanism be understood in terms of underlying responses to stimuli present in the *Umwelt*, or in terms of the *generation* of the *Umwelt* from the external environment, or perhaps both? There is clearly a danger of circularity here, and this is an issue that will be revisited in greater detail in section 2.8.

It is worth noting that the distinction between external environment and *Umwelt* as stated here is equivalent to that sometimes made between *proximal* and *distal* situations (Brunswik, 1952; Heider, 1959; Nolfi, 1998). Distal situations pick out features of the external environment, whether objects, behaviours, or relational properties, and proximal situations refer to the sensorimotor

<sup>8</sup>Merleau-Ponty expresses a similar idea: "it is the organism itself - according to the proper nature of its receptors, the thresholds of its nerve centres and the movements of the organs - which chooses the stimuli in the physical world to which it will be sensitive" (1963). Bertrand Russell makes another similar point from a different angle: "[t]he observer, when he seems to himself to be observing a stone, is really, if physics is to be believed, observing the effects of the stone upon himself" (quoted in Arkin, 1998, p.244). Also, the poetic language of von Uexküll should not be taken to imply that the agent has to be *consciously aware* of a feature in order to respond to that feature. No distinctions are made in this dissertation between processes that may require consciousness, or awareness, and processes that do not.

space of the agent. Although the former terminology will be preferred wherever possible throughout this dissertation, in one or two cases fidelity to the literature will require appeal to the latter.

### 2.3.3 Perception and action

One more theme remains to be introduced; this is that *perception* and *action* should be viewed as ontologically inseparable, that they must be understood as different ways of looking at a single underlying process of sensorimotor interaction with an environment. There are various contemporary expressions of this idea (see, for example, Gibson, 1979; Powers, 1973; Ballard, 1991), the essence of which goes back a long way. Consider the following from John Dewey:

What is wanted is that sensory stimulus, central connections and motor responses shall be viewed, not as separate and complete entities in themselves, but as divisions of labour, functioning factors, within the single concrete whole, now designated as the reflex arc. (1896, p.137, quoted in Pfeifer & Scheier, 1994)

In saying this, Dewey could quite easily have been a contemporary critic of the sense-think-act cycle at the centre of cognitive science. Moreover, for Dewey, the close coupling of perception and action must always be seen in the context of an interaction, or a ‘coordination’, with an environment. As Dewey puts it, sensory stimulus and motor response “are always inside a coordination and have their significance purely from the part played in maintaining or reconstructing the co-ordination” (ibid., p.139). The connection to be made is this. To confuse behavioural and mechanistic levels of description to the extent that behaviours are somehow located ‘inside’ the behaving agent, is to attempt to internalise the entire process of observed ongoing agent-environment interactivity (by the definition of behaviour given in section 2.3.1). However, if this entire process is considered to be internal, there can be no meaningful interpretation of the close coupling of perception and action; this coupling is only significant, as Dewey notes, in virtue of its role in mediating agent-environment interaction. It follows that such significance can only be maintained for ‘internalised behaviours’ to the extent that the position is taken that perception and action are *also* fully internalised, and this is not an attractive position to be forced into. In other words, to internalise a behaviour is to obscure and illegitimate the entire process of (observed) agent-environment interaction upon which the significance of the intimacy of perception and action depends.

Of course, in complex organisms, there will be much to learn from the study of components of perception and action in relative isolation. For example, much is now known about the functioning of the primate visual attention system by which certain features of the visual field accrue much greater explanatory significance than others (Pashler, 1999; Parasuraman, 1998). However, it remains arguable that any story of visual attention must remain incomplete to the extent that it remains isolated from action, and from the ongoing processes of sensorimotor interaction constitutive of behaviour.

The inseparability of perception and action constitutes another running theme of this dissertation. For the remainder of this chapter, however, it is sufficient to keep in mind that this intimacy is itself inseparable from the distinction between behaviour and mechanism.

### 2.3.4 Summary

This section has identified important distinctions between behavioural and mechanistic levels of description, functional and operational explanations, and the external environment and the *Umwelt*. It has also asserted the inseparability of perception and action. These are themes that will recur both in the account of the environmental complexity thesis presented in this chapter, and throughout the dissertation as a whole.

## 2.4 From empiricism to behaviourism

With some of the key concepts of this dissertation now in place, I will turn to a brief history of externalist approaches to understanding the relations between behaviour, mechanism, and environment, from philosophical empiricism to behaviourist learning theory. This account is necessarily illustrative rather than complete, and I do not pretend that any of the claims made are unequivocal; more detailed, and more defensible histories can be found in Young (1970) and Boring (1950), volumes which have informed much of what follows. The account presented here is intended partly to provide context for the discussion of the environmental complexity thesis that follows in section 2.6, and partly to disentangle the themes of this dissertation from the abundance of theoretical positions that can be distinguished in discussions of behaviour, mechanism, and environment over recent decades.

### 2.4.1 A brief history of association psychology

The classic empiricist philosophers, Locke, Berkeley, and Hume, each formulated their own version of empiricism, the externalist doctrine that all knowledge originates in patterns of sense data. Bishop Berkeley, for example, came up with a handy ‘proof’ for the existence of God: Who else, after all, could ensure the “steadiness, order, and coherence” of our sensory input? (Berkeley, 1690, sec.30). A unifying characteristic, however, may be identified in their concern with epistemological issues, rather than with behaviour (or with the mechanisms underlying the generation of behaviour).

In the nineteenth century, empiricist philosophy gave rise to ‘association psychology’, the central claim of which was that all psychological properties could be explained in terms of the formation of associations in patterns of sense data, these associations thus constituting ‘mechanism’. Through the work of James Mill (1773-1836), his son John Stuart Mill (1806-1873), and Alexander Bain (1818-1903), a gradual transition may be discerned from an epistemological focus to a more modern concern with psychological properties: Thus the younger Mill had a theory of perception, and Bain’s discussion of motor phenomena provided associationism with a balanced sensorimotor perspective.<sup>9</sup>

As association psychology developed, so too did ideas about how psychological properties should be related to brain physiology. At the time, these ideas largely took the form of a debate over the localisation of function in the brain. Young (1970) traces this debate back to the phrenology of Franz Joseph Gall (1758-1828). Gall replaced the empiricist *tabula rasa* view of the mind

<sup>9</sup>Bain’s emphasis on action can be contrasted with Condillac’s “statue” (1754). Condillac tried to ‘prove’ a passive, sensory associationism by ‘adding senses’, one by one, to a marble statue (see Young, 1970, pp.13-15).

with “a theory postulating a set of innate, inherited instincts transmitted in the form of cerebral organs, whose activity varied with the size of the respective organs” (Young, 1970, p.15). Nowadays Gall is usually remembered for the complete failure of his cranioscopic method, but, as Young points out, it was Gall nevertheless who developed the first empirical approach both to the nature of ‘mental faculties’ and to their localisations within the brain. (*Ideas of cerebral localisation go back to the beginnings of anatomy and physiology in the fourth century A.D., but until Gall, were entirely without empirical attributes.*) Gall clearly wanted to understand the mechanisms underlying adaptive behaviour, and - crucially - wanted to locate these mechanisms explicitly in a physiological substrate, but his phrenology, apart from being wrong, was also of a very internalist character.<sup>10</sup> It therefore stood in conceptual opposition to the growing influence of associationism, and, as such, (externalist) associationism itself stood largely opposed to the idea of functional localisation - of tying behaviour to actual physiological mechanism - of any sort.

Empirical work in the early nineteenth century had also cast doubt on functional localisation. The ablation experiments of Flourens (1794-1867) and the experimental sensorimotor physiology of Magendie (1783-1855) were forthright in their rejection of functional localisation in the cortex, although they did determine that sensorimotor processes *were* closely tied to non-cortical parts of the nervous system. (The fact that these early experiments failed to discern any cortical functional localisation can be attributed partly to the pre-existing bias against finding such a result, and partly to the relatively primitive methodological practices of the day.) This position, established as orthodox after the publication of Johannes Müller’s *Handbuch der Physiologie des Menschen* (1833-40), was motivated largely by a vehement opposition to phrenology, and was closely allied - both in this opposition and regarding its dismissive treatment of cortical functional localisation - to sensorimotor associationism in psychology. For many years this brand of associationism was caught in a dualistic bind. Cortical structures were deemed to be entirely isolated from sensorimotor activity and involved only with ‘higher’ faculties such as ‘will’, ‘reason’ and ‘memory’, these being functional categories inherited largely unmodified from medieval times. ‘Non-cortical’ parts of the nervous system *were* experimentally tallied with sensorimotor processes, but there was no reasonable explanation for how this division was to be overcome. For Müller, for example, “the fibres of all the motor, cerebral and spinal nerves may be imagined as spread out in the medulla oblongata, and exposed to the influence of the will like the keys of a pianoforte” (1833, p.934).

All this changed with the growing influence of the theory of evolution in the latter half of the century. However it was not Darwin, but Herbert Spencer who was the first to realise that an evolutionary perspective sanctioned an extension of the principle of association *throughout* the nervous system, the cortex included. This paved the way for the empirical work of Fritsch and Hitzig (1870) - who finally established that electrical excitation of the cortex could elicit sensorimotor phenomena - and Ferrier (for example, 1876), who inaugurated the modern practice of experimental functional localisation (or cognitive neuropsychology as it is referred to today). Thus, at this point, it becomes possible to speak of the entire nervous system constituting a physiological substrate for the mechanisms of behaviour generation, with the operation of these mechanisms being

<sup>10</sup>I am simplifying here. Gall’s phrenology was also distinguished by a biological, and arguably externalist perspective, in relating the functions of the brain to the interactions of a behaving organism with its environment. However, Gall’s biology was pre-evolutionary. He identified the origin of these functions as following directly from the place of an organism in the ‘great chain of being’, hence the internalist focus on ‘innate, inherited instincts’.



understood in terms of sensorimotor association.<sup>11</sup>

#### 2.4.2 Spencer's 'evolutionary associationism'

Spencer's project, first apparent in his *Principles of Psychology* (1855), can be most simply described as the unification of associationism with the theory of evolution.<sup>12</sup> This had three important consequences, the first of which was the impetus to the empirical work of Fritsch and Hitzig described above. Another important consequence was that associationist theory, for the first time, was extended into the environment (thus constituting an essential development in any attempted externalist account of behaviour, mechanism, and environment). The classical empiricist philosophers had considered the 'environment' of the mind to consist of sense-data alone, but for Spencer, psychological properties were to be characterised in terms of adaptations to environmental structure. Finally, Spencer was able to reconcile externalist associationism with the existence of (innate) internal mechanistic structure, such structure being the result of the *inheritance* of associations:

The familiar doctrine of association here undergoes a great extension; for it is held that not only in the individual do ideas become connected when in experience the things producing them have repeatedly occurred together, but that the results of repeated occurrences accumulate in successions of individuals: the effects of associations are supposed to be transmitted as modifications of the nervous system. (Spencer, 1904, p.470, quoted in Young, 1970).

However, Spencer was notoriously noncommittal about the constitution of these supposed 'effects of associations', even by nineteenth century standards. His basic picture of both life and mind - "the continuous adjustment of internal relations to external relations" - was chastised as "vagueness incarnate" by William James (quoted in Young, 1970, p.194).

To summarise thus far, the externalist doctrine of associationism was set in opposition to internalist phrenology throughout most of the nineteenth century. This led to the ultimately untenable position of attributing explanatory significance only to non-cortical - and supposedly sensorimotor - parts of the nervous system. By the turn of the century, however, it was becoming clear that *all* parts of the nervous system could potentially constitute mechanisms of behaviour generation, with even the cortex subscribing to some sort of description in these terms. However, the extent to which a description of such mechanisms in terms of sensorimotor association would prove satisfactory was still open to question. Spencer's evolutionary associationism suggested both that the associationist principle *could* be extended *throughout* the nervous system, but also asserted the plausibility of the existence of innate, inherited psychological structure. However, it was the first of these suggestions, and the first alone, that was taken up with enthusiasm by the twentieth century behaviourists.

---

<sup>11</sup>Pierre Broca's localisation of language (Broca, 1869) may appear conspicuous by its absence from the above account. However Broca's view was certainly not original (Young, 1970), and indeed in conceptual terms it can be traced very directly to Gall - language, as a function, is much closer to medieval faculty psychology than it was to the prevailing sensorimotor associationism. Young notes that this conceptual (but not methodological) proximity to Gall certainly mitigated against less forceful exponents of language localisation than Broca, of whom he says that "his main contribution seems to have been a propaganda victory rather than an original discovery" (1970, p.135).

<sup>12</sup>Spencer's evolutionary perspective, especially in his early work (prior to the publication of Darwin's *Origin*, 1859), is usually associated with Lamarck's 'inheritance of acquired characteristics', a theory often located on the internalist side of the evolutionary fence. Spencer, however, was able to incorporate Lamarckian theory into a distinctly externalist framework (see Godfrey-Smith, 1996, pp.86-90).

### 2.4.3 Behaviourist learning theory

Following Wilhem Wundt's foundation of first formal laboratory of psychology in 1879, the rise of behaviourist learning theory in the early twentieth century lent empirical substance to associationism. Its aim was to establish that sensorimotor associations (or stimulus-response relations, as they came to be known) were sufficient to explain *all* behavioural structure. Experimenters would specify patterns of reinforcement and observe how stimulus-response relations were established. These investigations were united by the assumption that behaviour, and behaviour alone, was amenable to objective inquiry.

Behaviourism experienced a major impetus following Ivan Pavlov's discovery of the 'conditioned reflex' (see, for example, Pavlov, 1927), this being the capacity for learnt associations to form between motor responses and (relatively) arbitrary stimuli. Pavlov focussed on elucidating the *physiological basis* of the conditioned reflex, and in doing so provided the psychological entities of associationism with a firm mechanistic physiological foundation. (Until Pavlov the physiological foundation for sensorimotor associationism consisted largely of the so-called 'Bell-Magendie law', which asserted a distinction between afferent - sensory - and efferent - motor - nerve roots. However, nothing in this 'law' concerned the process of association itself; see Young, 1970, p.79.) The historical significance of Pavlov's achievement is recognised by Lorenz:

[T]hrough his [Pavlov's] investigation of the *physiological* side of a phenomenon that is *also psychological*, he radically dispensed with that ancient, deeply rooted prejudice that any process that is *analyzable* in causal physiological terms ipso facto cannot be a psychological process. (1948, p.204)

The empirical character of behaviourism, and its focus on the structure of *learnt* behaviour, also entailed an extension of the (associationist) concept of mechanism. Rather than mechanism consisting only of associations *per se*, behaviourism also considered mechanism insofar as it *constrained the formation* of associations in various ways. However, what was rarely at issue for behaviourists was the structure and underlying mechanistic basis of *unlearnt* behaviour, even though the grounds for there being mechanistic explanations of such behaviour had already been established by Spencer. Indeed, extreme behaviourism was often associated with a denial of the importance of pre-existing internal structure of any sort.<sup>13</sup> This stimulated both the (internalist) cognitive revolt spearheaded by Chomsky (section 2.2), and, earlier, the development of ethology associated with von Uexküll (1934), Lorenz (1948), and Tinbergen (1950).<sup>14</sup> These early ethologists were strongly motivated by the failure of behaviourism to explain (or, in their view, to even admit the existence of) innate species-specific behaviour patterns, or 'instincts'. It is useful here to distinguish the intellectual commitments of behaviourism, which - *pace* Spencer - did *not* exclude 'instinct', from its methodological commitments, which did; as noted above, behaviourism was committed to an explanatory monism in that all behavioural structure was to be accounted for in terms of stimulus-response relations. This conflict is well illustrated by the behaviourist

<sup>13</sup>I am simplifying again. The 'unconditioned responses' of Pavlov may well be construed as unlearnt behaviour. However, this type of response gained its explanatory significance in behaviourism not in virtue of this interpretation, but in virtue of its role in a theory of the acquisition of *conditioned* (ie. learnt) responses.

<sup>14</sup>A memorable perspective on the relationship between behaviourism and cognitive science is due to Lloyd (1989), for whom behaviourists treated mind as an opaque box in a transparent world, and cognitivists treated it as a transparent box in an opaque world.

B.F. Skinner, who, in response to a critique of behaviourism (Breland & Breland, 1961), wrote: “No reputable student of animal behavior has ever taken the position that ‘the animal comes to the laboratory as a virtual *tabula rasa*, that species differences are insignificant, and that all responses are about equally conditionable to all stimuli’ ” (Skinner, 1969, p.173). However, protestations such as these, although often made, had little effect on the *practice* of behaviourism.

The early ethologists were also concerned with a second, and related feature of behaviourism; the exclusion of agent-environment *interaction patterns* from explanations of behaviour. The environments of behaviourist learning experiments usually consisted of so-called ‘Skinner boxes’ which allowed only one or two responses (perhaps key-pecking or lever-pressing) and which presented only a limited range of stimuli (usually lights or sounds of various sorts). It was up to the experimenter to decide both the sequence of stimulus presentation and which stimulus-response relations to reinforce. The responses of the animal were significant only insofar as they satisfied, or failed to satisfy, these reinforcement conditions.<sup>15</sup> It is again worth quoting Lorenz:

The preconceived opinion that the reflex and the conditioned reflex are the only elements of behavior determined a quite special, scarcely varying kind of experimental setup in which the central nervous system under investigation had no *opportunity* to show that it was capable of anything other than responding to the influence of external stimuli. In this way, it was quite unavoidable that the opinion *necessarily* developed and became consolidated, that the functioning of the central nervous system is restricted to receiving and responding to external stimuli. (1948, p.211, emphasis in original)

This cavalier approach to environment prompted the early ethologists to study the behaviour of animals in their natural environments, and - in order to understand ‘instinct’ - to attempt to locate their observations in an evolutionary context. In ethology, reciprocal feedback relations between agent and environment are of great significance. The ways in which the actions of organisms influence their stimulus input are just as important as the ways in which such input is related to action. It is this criticism of behaviourism that is most germane to the present dissertation; both the philosophical arguments of this chapter and the concrete models that follow will lay considerable emphasis on agent-environment interaction patterns.

#### 2.4.4 To the present day

The history described in this section could of course have been told in many different ways, and to follow it through to the present day would require perhaps an even greater degree of judicious selectivity. For present purposes it will suffice to mention just a few of the major intellectual currents that flow from the beginnings of behaviourism and ethology towards the present (note that not all of these currents are externalist in character). The rise of classical cognitive science and ‘good-old-fashioned’ (symbolic) artificial intelligence (AI) in the mid-twentieth-century has already been discussed in section 2.2, and from the perspective of this dissertation these movements represent the significant misappropriation of intellectual resources. However, alongside the development of AI, there also rose to some prominence the discipline of cybernetics (Wiener, 1948; Walter, 1950,

<sup>15</sup>The usage of the past tense in the description of behaviourist practice refers to its early twentieth century heyday. Such practice, however, still persists to this day in many laboratories of animal psychology, even though few contemporary psychologists adhere to the tenets of behaviourism itself (see, for example, Dickinson, 1980).

1953; Ashby, 1952, 1956; Pask, 1961), the guiding principle of which was that adaptive behaviour should be construed as a control problem, to be understood through the application of concepts such as feedback, stability, and information. (Ashby's 'law of requisite variety' - a cornerstone of early cybernetic theory - will play a central role in the extension of the environmental complexity thesis in section 2.7.) The rise of 'new' (non-symbolic) AI, and of SAB/AL itself, can to some extent be traced to a positive reappraisal of these cybernetic principles, coupled with significant developments in both the understanding of the properties of neural networks (see, for example, Rumelhart, McClelland, & the PDP research group, 1986; Clark, 1989; Bechtel & Abrahamsen, 1991) and in the mathematical description of agent-environment systems as dynamical, rather than as computational systems (Port & van Gelder, 1995; Thelen & Smith, 1994). For detailed extensions of this (relatively recent) history the reader is referred to Pfeifer and Scheier (1999), Holland (1992), Langton (1995), and Brooks (1999).

A different thread can be followed through the diversification of ethology into the various disciplines of behavioural ecology (Krebs & Davies, 1991), optimal foraging theory (Stephens & Krebs, 1986), evolutionary psychology, (Barkow, Cosmides, & Tooby, 1992), cognitive ethology (Ristau, 1991), sociobiology (Wilson, 1975), and evolutionary anthropology (Boyd & Richerson, 1985), amongst others. These disciplines are each concerned with different aspects of drawing relations between environmental characteristics, evolutionary pressures, and human and animal behaviour, and I will not attempt to account for this diversity here. However, both evolutionary psychology and optimal foraging theory are of particular significance for present purposes, the former because it sets out very clearly the externalist explanatory projects at the heart of this dissertation, as we will see in the next section, and the latter because it provides a methodological bridge between the relatively abstract discussions of the first half of this dissertation and the relatively specific biological investigations that follow. This bridge will be built in chapter 5.

#### 2.4.5 Summary

To return to Spencer's evolutionary associationism is to return to a place in externalist intellectual history which, in some sense, represents the lowest conceptual common denominator of the various projects of this dissertation. For Spencer, the mechanisms underlying behaviour generation are necessarily adapted to environmental structure, are of sensorimotor constitution, and can comprise both of fixed structure and of acquired associations. After Spencer, this basic intellectual position becomes increasingly difficult to isolate in the *melée* precipitated by the development of behaviourist learning theory, the ethological and cognitivist responses to which have rebounded merrily throughout the twentieth century. It is of little surprise, and of some significance, that the environmental complexity thesis of Godfrey-Smith also finds early expression in the work of Spencer (as we will see in section 2.6.2).<sup>16</sup>

I turn now to a brief discussion of evolutionary psychology, a discipline which not only clearly articulates the two foundational themes of this dissertation (accounting for behavioural function

---

<sup>16</sup>I do not wish to overemphasise the intellectual importance of Spencer. Although he was unquestionably one of the most prominent intellectual figures of the Victorian era, his philosophical 'system' was often justifiably castigated, and his fall from grace remarkably rapid (Peel, 1971; Richards, 1987). As Godfrey-Smith puts it, "[a] very large proportion of Spencer's specific views on the relations between mind and environmental complexity, and on evolutionary and psychological mechanisms, are false. Most contemporary naturalistic philosophers would also find his overall world view a completely unacceptable one" (1996, p.7).

in terms of environmental structure, and relating behavioural and mechanistic levels of description) but which also provides the most direct link from the historical account given above to the environmental complexity thesis itself.

## 2.5 Evolutionary psychology

Evolutionary psychology is characterised by the (strongly externalist) claim that psychological properties can be understood as adaptations to an environment, in just the same way that evolutionary biology attempts to understand morphological and physiological characteristics (Barkow et al., 1992). Arguably, as intimated above, Herbert Spencer provided the earliest articulation of this claim with his unification of evolutionary theory with associationism, stressing the potential heritability of learned associations.

### 2.5.1 Contemporary evolutionary psychology

From the foundations laid down by Spencer and Darwin, evolutionary psychology has become a very specific project. Here is a recent definition: “The main claim of evolutionary psychology is that the human mind is a set of cognitive mechanisms that are adaptations to the environment of the Pleistocene” (Grantham & Nichols, 1999, p.47); typical examples of the explanatory targets of evolutionary psychology include male sexual jealousy and maternal preference between preterm twins (Barkow et al., 1992). In this section these specifics are relaxed; I am concerned not with any environment in particular, nor with any species in particular. What is retained is the principle of there being evolutionary accounts of behaviour and of the mechanisms underlying the generation of behaviour.<sup>17</sup>

The collected papers in Barkow et al. (1992) distinguish two explanatory projects for (human) evolutionary psychology, both of which remain controversial. The first is to assign functional roles to known psychological and/or behavioural properties. Much criticism has focussed, however, on the suspicion that evolutionary psychological accounts are often nothing more than Panglossian just-so-stories *sensu* Gould and Lewontin (1979), and, indeed, the frequent importance of *cultural* selective pressures in such accounts makes this suspicion particularly difficult to dislodge. Nevertheless, if it can be assured that a particular phylogenetic history of a trait is accurate, then this project is viable, just as it can be in the case of adaptationist evolutionary biology (although recall Lewontin’s critique of the concept of a ‘niche’, a critique revisited in the following section). In chapters 8 and 9 I will present a version of this project in which concrete models are utilised to explore potential functional roles for a particular pattern of choice behaviour reliably observed in humans and animals, the ‘matching law’ of Herrnstein (1961).

The second project is the use of a knowledge of the function of a behaviour to help to elucidate properties of the underlying behaviour-generating mechanisms. As Cosmides and Tooby (1987) advise: “It is nearly impossible to discover how a psychological mechanism processes information unless one knows what its function is, what it was ‘designed’ or selected to do” (p.285). The problem here is that although an appreciation of the function of behaviour may well be necessary

<sup>17</sup>In particular I wish to distinguish myself from the various ‘ethical’ misinterpretations of evolutionary psychology. That there may exist adaptationist accounts of particular behavioural phenomena does not remove individual culpability should those behaviours be morally unacceptable in the modern world.

in understanding the underlying mechanism, it is unlikely to be sufficient. Recall the leitmotif of section 2.3.1: agent-side mechanisms underlying the generation of behaviour should not be assumed to be identical to the behaviour itself (or to its function), and, as I argued in the same section, in order to effectively relate the two it can often be helpful to trace the relations between the external environment and the *Umwelt* of the agent. I shall again defer development of this argument until section 2.8.

These projects recapitulate the themes with which this chapter opened, the integration of which constitutes its primary duty. In the context of the history of externalism, what is important is that mechanism is no longer restricted to be considered solely in terms of associations and the constraints on association formation, but can be construed more broadly as the agent-side structure subserving the generation of behaviour, just as suggested by the definition of mechanism in section 2.3.1. In evolutionary psychology, externalism delivers explanations of the structure of internal mechanism (via behaviour), but, unlike the externalism of the associationists and the behaviourists, externalist concepts need *not* thereby monopolise descriptions of the *operation* of these mechanisms.

### 2.5.2 Adaptive function and ecological ‘niche’

This is an appropriate stage in the unfolding of this chapter to revisit Lewontin’s critique of the concept of the ecological ‘niche’, first discussed in section 2.1.2. The general idea is that agent and environment are in continual mutual specification, therefore any attempt to understand agent-side properties in terms of pre-existing environmental structure is fraught with difficulty. Although applicable to externalism in general, this argument can be seen to be particularly germane to the functional project of evolutionary psychology. A niche may be thought of as those features of the external/selective environments which encourage the ascription of a particular function, but, if functional descriptions necessarily include external environmental features, how can functional terms be used to *link* observed behaviours with aspects of environment structure? Lewontin (1983) argues that niches are always described in terms of the activity of an organism, in which case, “how can a niche exist *before* the species, so that the species can evolve into it? This contradiction is not resolved in the classical Darwinian theory of adaptation, which depends absolutely on the problem pre-existing the solution” (1983, p.98).<sup>18</sup>

There is of course no easy answer to this (although Lewontin’s decision not to distinguish ‘species’ from ‘organism’ may provide some hope). To my mind, his argument does indeed cause problems for the ascription of function in evolutionary biology, and also in human evolutionary psychology (for which a niche will have cultural attributes and therefore may be particularly hard to pin down). However, it does provide the opportunity to reassert the pragmatic nature of the externalism of this dissertation. As argued in section 2.1.2, for present purposes it may simply be accepted that agents influence their environments (and hence their niches) by their own activity, with the elucidation of this process constituting part of the explanatory duty of concrete models

<sup>18</sup>A more general definition of niche that is less tied to function - but to that extent less useful in the ascription of function - comes from Wilson (1975): “The range of each environmental variable such as temperature, humidity, and food items, within which a species can exist and reproduce” (p.317). Lewontin’s argument would attach most easily to the final attribute, since the pattern of food distribution is likely to depend to some extent on the foraging behaviour of the species (this issue will be explored in some detail in chapters 7 and 8).

that incorporate interaction and mutual specification between agent and environment.

### 2.5.3 Optimal foraging theory

It is also valuable at this stage to foreshadow the discussion of ‘optimal foraging theory’ in chapter 5. Optimal foraging theory is closely related to evolutionary psychology given its central premise that foraging behaviour is an optimal solution to an environmental problem (Stephens & Krebs, 1986). It may, however, be distinguished from evolutionary psychology in at least three ways: by the assumption that evolution is a process of optimisation, through a concern with animal (usually foraging) behaviour, and by a lack of emphasis on any distinction between present and past environments. A detailed discussion of optimal foraging theory will be left for chapter 5 itself; it will suffice to say for the present that it will figure prominently in this dissertation at the juncture between the relatively abstract investigations of behaviour and mechanism in the first part of this dissertation, and the more specific biological inquisitions of the later chapters.

### 2.5.4 Summary

Evolutionary psychology continues the externalist history of section 2.4 up to the present day, albeit to the exclusion of many intervening intellectual developments. A focus on evolutionary psychology allows the identification of two explanatory projects central to this chapter and to this dissertation as a whole: the assignment of functional roles to observed behaviours, and the elucidation of the relations between behaviour and mechanism. It also bears clear relation to the branch of theoretical biology known as optimal foraging theory, a relation which will reappear in chapter 5 as a methodological focus of the latter half of this dissertation.

## 2.6 Complexity and the function of mind in nature

Most of the conceptual groundwork has now been laid, and it is high time for these various strands to be pulled together - and for a final theme to be introduced - with a discussion of a very specific externalist project, Peter Godfrey-Smith’s *Complexity and the Function of Mind in Nature* (1996). The theme that has been waiting in the wings until now is that of *complexity*. Although a thorough discussion of the semantics of this term will be left for the following chapter, it may be said straightaway that an externalist perspective on complexity involves the attempt to explain the complexity of behavioural and/or mechanistic properties with reference to complexity properties of the environment. Godfrey-Smith calls these sorts of explanation ‘c-externalist’, the prefix standing, of course, for ‘complexity’. Such explanations stand in evident opposition to theories of ‘self-organisation’, which can be characterised as ‘c-internalist’; these are theories which try to explain complexity properties of a system with regard to the system itself.<sup>19</sup>

From this point on, this chapter largely pursues its own agenda; the integration of the externalist perspective on the relations between environment and behaviour with the relations between behaviour and mechanism. Although this integration is intended to motivate the empirical sections

<sup>19</sup>As Ashby (1962) has warned, the term ‘self-organisation’ should be treated with caution; used too loosely it is ambiguous, used too precisely and it becomes self-contradictory. Crutchfield (1994b) puts it in the following way: “it is the observer or analyst who lends the teleological ‘self’ to processes which otherwise simply ‘organize’ themselves according to the underlying dynamical constraints” (1994b, p.2).

of this dissertation, it is not necessary for their coherence. Much of what follows will be directed toward specific issues raised by Godfrey-Smith's formulation of his project, and will take the form of a selective critique of this formulation. The reader eager to encounter empirical substance may wish to briefly consult sections 2.6.1 and 2.7.1, for descriptions of the environmental complexity thesis and the law of requisite variety respectively, before proceeding to the final sections of this chapter (2.9 and 2.10).

### 2.6.1 The environmental complexity thesis

At the outset of his 1996 book *Complexity and the Function of Mind in Nature*, Godfrey-Smith states his 'environmental complexity thesis' (hereafter ECT) as follows:

The function of cognition is to enable the agent to deal with environmental complexity. (1996, p.3)

The ECT is evidently an externalist claim, and the intuition behind the claim is that 'cognition' is more complex than 'no cognition', thus 'cognition' is argued to be a complex response to a complex environment. The ECT is also a claim about necessity, not sufficiency. The idea is that environmental complexity is necessary for there to be cognition, but *not* that environmental complexity will always result in there being cognition. I will begin by briefly describing Godfrey-Smith's own analysis of the various components of the ECT: function, cognition, complexity, and the process of 'dealing with'. First, the vision of environmental *complexity* held by Godfrey-Smith is one of *heterogeneity*:

Complexity is changeability, variability; having a lot of different states or modes, or doing a lot of different things. Something is simple when it is all the same. (ibid., p.24)

This definition of environmental complexity is not without its difficulties (which will be discussed in some detail in the next chapter), but at least it has the virtue of being relatively clear. For now I will let it stand.

An important distinction for Godfrey-Smith is between 'instrumental' and 'teleonomic' versions of the ECT. These are distinguished by their treatment of *function*. The instrumental ECT is the claim that cognition is 'good for' dealing with environmental complexity. Instrumental functions are ascribed in virtue of their explanatory salience with respect to some capacity of a larger system (Cummins, 1975), in this case the behaviour of the organism. Godfrey-Smith admits that this version is too strong, since cognition can be 'good for' other things as well (for example, cognition is 'good for' aesthetic thought). The teleonomic ECT is the claim that cognition is 'for' dealing with environmental complexity. The teleonomic function of something is the effect it has which explains why it is there, usually in view of some selective process (Wright, 1973).<sup>20</sup> It is this version of the ECT that Godfrey-Smith defends; cognition is argued to be an *adaptation* to the problem of environmental complexity. But this does not make the ECT a part of evolutionary psychology, as Godfrey-Smith is careful to establish:

<sup>20</sup>Teleonomic functions are 'proper' functions, in the sense of Millikan (1989), if the justification of 'why it is there' is derived from *natural* selection. Note that all concepts of function are distinguished from mere effects in virtue of their explanatory salience.



[T]he environmental complexity thesis is construed as an adaptationist claim about the capacity for cognition itself; it is not itself an explanation for particular thoughts or patterns of belief. (ibid., p.24)

However it is not clear what Godfrey-Smith means by the ‘capacity for cognition’, or indeed by ‘cognition’ itself. Throughout his book, cognition is discussed at a variety of levels of abstraction. In one place it is interpreted “as a means to the production of behavioural complexity” (ibid., p.26). In another, as “a basic apparatus that makes possible perception, the formation of belief-like states, the interaction of these states with motivational states such as needs and desires, and the production of behaviour” (ibid., p.22). The possibility of cognition being an advanced kind of homeostatic device is advanced (ibid., pp.76-79), and there is also reference to cognition as a “basic mental tool-kit” (ibid., p.127). This variety makes it clear that although the ECT may provide a detailed discussion of how the functional properties of behaviour may relate to environmental structure, there will be less to say about how such behaviour is to be related to underlying mechanisms. For Godfrey-Smith, both levels of description are subsumed under the term ‘cognition’, a term which he loads with many of its customary connotations (see section 2.2).

This dalliance with cognitive terminology leads his version of the ECT into some unfortunate territory. Shapiro (1999), for example, gets things very much the wrong way around. He wants to know when the observation of a behaviour warrants the ascription of *psychological* - or cognitive *sensu* section 2.2 - mechanism, rather than just *physiological* - non-cognitive - mechanism. He suggests that the ECT can be used to answer this question; if there is sufficient environmental complexity, then cognitive mechanism must be involved, the implication being that non-cognitive mechanism just could not cope in such circumstances. However, it is clear from the outset that to ask this question is to revive the prejudice effectively dispelled by Pavlov; that any process analysable in physiological terms *ipso facto* cannot be a psychological process (see section 2.4.3).<sup>21</sup>

This, then, is the first element of my critique of Godfrey-Smith’s formulation of the ECT. I submit that, in the context of the ECT, the commitment that the mechanisms underlying behaviour generation are cognitive is damaging. This is so whether cognition is to be understood in Godfrey-Smith’s rather vague sense, or in the more specific sense of section 2.2. For me, the most appropriate definition of mechanism remains the ‘agent-side internal structure underlying the generation of behaviour’. It is the functional properties of behaviour which are the targets of the adaptationist claim of the ECT. Mechanisms are certainly required to subservise such behaviour, and so mechanisms will also be indirectly subject to the same forces of selection that operate directly on the functional entities of behaviour. But there is no necessary, direct link between complex behaviour and complex mechanism, cognitive or otherwise (recall Simon’s ant on the beach, section 2.3.1). Therefore, there is no way the ECT should be used to justify claims for the existence of ‘psychological mechanism’ as opposed to ‘physiological mechanism’. In other words, part of the appeal of the ECT comes from the apparent link between environmental structure and internal

<sup>21</sup>Dennett (1978, 1987) makes a similar point, arguing that behaviour of a sufficient apparent complexity drives observers to take the ‘intentional stance’ with respect to the behaving agent, such that in order to understand and/or predict the behaviour of the agent it is necessary to (or at least it *helps* to) ascribe cognitive states (*sensu* section 2.2) to the agent. However, as Dennett is careful to emphasise, the act of taking the intentional stance does *not* imply that the actual agent-side behaviour-generating mechanisms are cognitive in this sense; the intentional stance is an *observational* strategy.

mechanism, but part of the difficulty with Godfrey-Smith's formulation is that this link goes via the term 'cognition', and that his liberal use of this term obscures the essential distinction between behaviour and mechanism.

Moving on, for Godfrey-Smith, cognition (whatever this may be) 'deals with' environmental complexity by 'responding to' the environment rather than by 'accommodating' it. Godfrey-Smith organises much of his formulation of the ECT around this distinction, and so it is worth describing what he takes it to signify. In his book he illustrates the distinction by appealing to the historical roots of the ECT in the writings of Herbert Spencer and John Dewey. This is a good strategy, and I shall begin in the same way.

### 2.6.2 Spencer's version

According to Godfrey-Smith, it was Spencer who was responsible for articulating the earliest version of the ECT. Spencer believed in a general 'law of evolution' which prescribed a universal dynamic from states of "indefinite, incoherent homogeneity" to states of "definite, coherent heterogeneity" (or, in Godfrey-Smith's terms, complexity).<sup>22</sup> Increases in environmental heterogeneity were supposed to lead to corresponding increases in internal heterogeneity by "the continuous adjustment of internal relations to external relations", this process being the central plank of Spencer's evolutionary associationism (section 2.4.2). Godfrey-Smith interprets this form of externalist theorising as suggesting that "organic properties are not seen so much as solutions to environmental problems but rather as bearing the imprint of the environment's pattern" (ibid., p.89). There is therefore no sense, for Godfrey-Smith, in which a Spencerian ECT implies that cognition 'deals with' environmental complexity by way of response, rather, organism 'accommodates' environment. Interpreted this way, the Spencerian ECT is endlessly problematic. Not only does it deny enlightened interactionism, but the idea of organic properties accommodating environmental properties demands a reification of the distinction between external and internal: One cannot bring the external within the confines of the internal without presuming a pre-existing distinction between the two. Therefore, Spencer's version of the ECT, as understood by Godfrey-Smith, cannot maintain the conceptual arbitrariness of distinguishing between the external and the internal deemed important in section 2.1.

### 2.6.3 Dewey's version

It is in the later work of John Dewey - most particularly in his *Experience and Nature* (1929) - that Godfrey-Smith identifies a version of the ECT. For example:

The world must actually be such as to generate ignorance and inquiry; doubt and hypothesis, trial and temporal conclusions; [...] The ultimate evidence of genuine hazard, contingency, irregularity and indeterminateness in nature is thus found in the occurrence of thinking. (1929, p.69, quoted in Godfrey-Smith, 1996, p.100)

<sup>22</sup>Spencer struggled to reconcile this negentropic law with the (entropic) second law of thermodynamics, which was being formulated at around the same time. Spencer's insistence on the complete generality of his 'law', and the lack of a well developed theory of systems, led him into trouble with the postulation of various pseudo-physical 'laws' such as the 'conservation of forces'. William James - a persistent critic of Spencer - was particularly scathing of this aspect of his philosophical system (see James, 1897; Perry, 1935).

For Godfrey-Smith, Dewey's version of the ECT is, by and large, much more acceptable than Spencer's.<sup>23</sup> This is primarily because of Dewey's pragmatic outlook. As Godfrey-Smith explains, "Dewey was opposed to the idea that in solving problems organic systems merely accommodate environmental demands. Rather, they intervene in the environmental processes which generated the problem, and alter the environment's intrinsic course" (1996, p.139).<sup>24</sup>

Godfrey-Smith's interpretation of this version of the ECT, through sanctioning an interpretation of cognition as response, can be argued to improve on Spencer's version in several ways. First, enlightened interactionism can be entertained; responses "alter the environment's intrinsic course". Second, the division between external and internal can be retained as arbitrary, to the extent that the response is *not* framed in terms of the environment itself (more on this in the following section). A third advantage of Dewey's version is his idea that for cognition - *sensu* Godfrey-Smith - to be of adaptive value to an agent, the environment must be characterised by a mixture of predictability and unpredictability:

The incomplete and the uncertain give point and application to the ascertainment of regular relations and orders. (1929, p.160, quoted in Godfrey-Smith, 1996, p.130)

Godfrey-Smith restates this idea<sup>25</sup> in the form of two conditions to be satisfied if a role for cognition is to be mandated. First, that there exists *variability* "with respect to distal conditions that make a difference to the organisms's well-being" (1996, p.118). Second, that there be *stability* "with respect to the relations between these distal conditions and proximal and observable conditions" (*ibid.*, p.118, recall from section 2.3.2 that distal conditions concern features of the external environment, and that proximal conditions refer to the sensorimotor space - *Umwelt* - of the agent). This extension is valuable in two ways. First, any description of mechanism that is framed in terms of distal properties, in terms of Godfrey-Smith's accommodation rather than response, cannot sustain this distinction. If a mechanism accommodates environmental properties, then the distal (environmental) variability must be, or become, somehow *internal* to the organism. And if what is distal is already *ex hypothesi* internal, then the distinction between distal variability and stability in distal-proximal relations can have no sensible meaning. And of course if this distinction cannot be maintained, then - according to Godfrey-Smith - there is no mandate for cognition in the first place.<sup>26</sup>

Second, and more straightforwardly, the distinction between distal variability and stability in proximal-distal relations is an invitation to empirical exploration. It is not difficult to envisage the development of concrete models in which both distal variability and proximal-distal relations can be manipulated. Indeed, both the models discussed by Godfrey-Smith in the second part of his

---

<sup>23</sup>One problem for Godfrey-Smith is that, unlike Spencer, Dewey does not appear to subscribe to the view of environmental complexity as heterogeneity. Another is that Dewey conceives of 'mind' or 'cognition' as existing only within a sociolinguistic context.

<sup>24</sup>This view is challengingly reminiscent of Marx's eleventh thesis on Feuerbach (1848): "Philosophers have only *interpreted* the world in various ways, but the real task is to *alter* it".

<sup>25</sup>An idea which also finds expression in the philosophy of science of Levins and Lewontin (1985): "Things are similar: this makes science possible. Things are different: this makes science necessary" (p.141).

<sup>26</sup>Keijzer (1998) has recently argued from similar premises towards a different but connected purpose, pointing out that a confusion of behaviour with mechanism can lead to an unwarranted simplification of the sensorimotor capabilities underlying adaptive behaviour. He argues that *stability* in behaviour is often a desired condition (for a behaving agent), but that this may not necessarily imply, or require, sensorimotor stability. Keijzer translates these concerns into an intriguing criticism of the use of wheeled robots in SAB/AL research.

book and some of the models in the present dissertation do exactly this; they do so, however, in very different ways.

#### 2.6.4 Accommodation versus response: a second look

It is true that the distinction between accommodation and response, as different ways of ‘dealing with’ environmental complexity, lacks the epistemological bite of some of the other distinctions employed in this chapter (behaviour versus mechanism, external environment versus *Umwelt*). It will help the flow of argument in this chapter to try and alleviate this weakness a little, a weakness which arguably stems from the difficulty of distinguishing between the notions of ‘cognition’ and ‘dealing with’ as they appear in Godfrey-Smith’s formulation of the ECT. In particular there is some ambiguity surrounding the interpretation of accommodation as, on the one hand, *state*, and on the other, *process*. The above discussion tends to interpret accommodation as state; the external as ‘inside’ the internal. An interpretation in terms of process would weaken the arguments set out in this discussion insofar as the results of the process of accommodation are not assumed to be isomorphic to its objects (the relevant distal situations). However, to the extent that this is so, the motivation for distinguishing accommodation from response in the first place is less obvious. If, in accommodating (*sensu* process) distal situation *A*, the internal mechanism comes to display internal property *B*, which is not the same as *A*, then why not think of *B* as a response to *A*? For present purposes it is not necessary to commit to an interpretation of accommodation either as state or as process, however it *is* necessary that if an interpretation in terms of process is preferred, it should be taken to signify some situation in which *A* and *B* are isomorphic in some way. As such, any interpretation of accommodation (whether as state or as process) will remain distinct from response. This strategy is unabashedly pragmatic, the idea being simply to isolate a sense of accommodation which entails that accommodatory versions of the ECT are conceptually problematic.

#### 2.6.5 Summary

Godfrey-Smith’s ECT provides a very clear articulation of a contemporary externalist project that seeks to elaborate the relations between behaviour, mechanism, and environment. This project builds on the foundations of evolutionary psychology and is grounded both in the work of Spencer and Dewey, and in the more extensive history of externalist explanation described in section 2.4. However, as it has been described so far, the ECT blurs the essential distinction between behaviour and mechanism with its liberal use of the term ‘cognition’ (a term which Godfrey-Smith prefers to associate with notions of response rather than accommodation). The following sections continue with a critique of Godfrey-Smith’s formulation, in the next section with an exploration of W. Ross Ashby’s ‘law of requisite variety’.

### 2.7 A cybernetic ECT

The analysis of the ECT in terms of Ashby’s (1956) law of requisite variety (henceforth LRV) is worth pursuing largely because of its engagement with the concepts of variability and homeostasis. (Although mentioned by Godfrey-Smith, the LRV was never for him a focus of attention.) A product of the mid-twentieth century cybernetic school, the LRV may be considered intuitively

as asserting that ‘only [agent-side] variety can destroy [environmental] variety’. A more formal statement derives from the following account.

### 2.7.1 The law of requisite variety

For an agent to maintain relative stability in certain (internal) essential variables (for example heart rate, body core temperature), it must *prevent* the transmission of environmental variability through to these essential variables. In the same way that a good thermostat prevents the transmission of environmental variations in temperature through to a particular object (for example, the interior of a refrigerator should remain at a constant cool temperature despite the fluctuating temperature of a kitchen on a midsummer day), a well adapted agent prevents the transmission of certain environmental variables (for example the prevalence or scarcity of food, the proximity, or otherwise, of predators) through to such essential internal variables as blood sugar level or heart rate. With this situation in mind, the law of requisite variety (henceforth LRV) can be easily formulated: Consider a set of possible environmental disturbances  $D$ , a set of possible responses on the part of the agent,  $R$ , and a set of possible outcomes,  $O$ . Consider also that for each  $D_i$ , there is distinct outcome  $O_i$ , and a particular response  $R_i$ . Stability in the essential variables requires minimising the variation in  $O$ , and this then requires that the variety in  $D$  is matched by the variety in  $R$ . Ashby himself provides a more concise summary:

If  $R$ 's move is unvarying, then the variety in outcomes will be as large as the variety in  $D$ 's moves; only variety in  $R$ 's moves can force down variety in the outcomes. (1956, p.206)

A small qualification is necessary. It is not usually necessary for essential variables to be maintained at a precise value; variation within a certain range is usually permissible and often inevitable. In mammals, for example, changes in heart rate and blood sugar concentration are in fact essential contributions to adaptive behaviour, for example in ensuring physiological preparedness for ‘fight’ or ‘flight’ responses (Carlson, 1994). McFarland and Sibly (1972, 1975) speak of a physiological state-space, with essential variables defining axes in this space such that its dimensionality is determined by the number of essential variables. It would be more accurate to interpret the LRV in terms of maintaining the state of the organism *not* at a single point in physiological state-space, but rather in a region of this space circumscribed by a set of ‘lethal boundaries’. However, for the purposes of this chapter this distinction may be overlooked; I shall be concerned with a (broadly construed) stability in essential variables only.

Ashby freely admits that his LRV expresses the same basic idea as Shannon’s (1949) ‘tenth theorem’ in information theory, which states that the amount of noise that can be removed from a message by a ‘correction channel’ is limited to the amount of information that can be carried by that channel. Ashby notes that Shannon stated his theorem in the context of telecommunications, but that the formulation is just as applicable to biological regulation. And in this context, as I argue below, there are clear resonances with the ECT.

### 2.7.2 The LRV as a version of the ECT

It is worth restating Godfrey-Smith’s most general version of the ECT: “the function of cognition is to enable the agent to deal with environmental complexity” (1996, p.3). How can the LRV

express a version of this?

Most obviously, environmental complexity may be interpreted as the set of environmental disturbances  $D$ . Secondly, the ‘function’ can be stated as the maintenance of stability in essential variables (the minimisation of variation in  $O$ ). Just as with Godfrey-Smith’s ECT, this can be construed as a teleological function (*sensu* Wright, 1973); the maintenance of stability in essential variables is what ‘cognition’ does that explains why it is there. This leaves ‘cognition’, which instead of being associated with the variety of interpretations favoured by Godfrey-Smith (section 2.6.1), is here related directly to the regulation effected by the various responses (the set  $R$ ) deployed by the agent in ‘forcing down the variety’ in  $O$ .

In this way the LRV expresses a quite specific version of the ECT, and a version which makes several contributions. First, there is a clear emphasis on environmental (distal) variability. Ashby, however, views this emphasis in a slightly different light, arguing that environmental variety comes in two distinct forms:

There is that which threatens the survival of the gene-pattern. This part must be blocked at all costs. And there is that which, while it may threaten the gene-pattern, can be transformed (or re-coded) through the regulator  $R$  and used to block the effect of the remainder (1956, p.212).

This may be usefully compared with Godfrey-Smith’s dual condition for the adaptive significance of cognition. For Godfrey-Smith, without distal variability, cognition is not necessary, and without stability in distal-proximal relations, cognition is not possible. Ashby would have agreed with the first premise; without environmental variability that “threatens the survival of the gene-pattern”, regulation would not be necessary. However, instead of speaking of distal-proximal relations, Ashby identifies the importance of ‘potentially beneficial’ variability, the suggestion being that without this kind of variability, the agent would not be able to act - or respond - at all. One example of this kind of variability would be a change in odour intensity that correlates with proximity to a food source, another would be some kind of environmental stochasticity that enables protean behaviour in a predator-prey situation. The perspectives of Ashby and Godfrey-Smith are therefore complementary, and an interpretation of the ECT emerges in which all three characteristics combine: stability in distal-proximal relations, and distal variability that both threatens the agent and facilitates its activity.

A second advantage is the potential offered by the LRV for empirical exploration of the ECT. In framing the components of the ECT in terms of potentially *measurable* variety - in environment, response, and outcome, - the LRV points the way to the construction of models in which the empirical consequences of the ECT can be elaborated. For example, Fletcher, Zwick, and Bedau (1998, 1996) investigate how the manipulation of environmental texture in a toroidal grid-world environment relates to the variety of responses deployed by well-adapted agents. Fletcher et al. employ Shannon entropy to track three distinct measures of variety: the information content of the environment (from the perspective of the agent), the information content of look-up tables of sensorimotor rules that constitute agent mechanism (which represents the variety of response), and the *between-agent* variety in look-up table structure. They take two results to exemplify the LRV (they do not mention the ECT). First, that response variety approaches environment information content; thus variety in  $R$  is matching variety in  $D$ . Second, that between-agent variety falls to zero,

indicating that the variety in  $R$  is not random, such that for every  $D_i$  there is indeed a particular response  $R_i$ . This model (and similar work, see, for example, Menczer & Belew, 1996) is of the same spirit as the models of this dissertation, most notably the prisoner's dilemma models of the following chapter.

### 2.7.3 Homeostasis

As noted above, although Godfrey-Smith mentions the LRV in his 1996 book, it is given rather short shrift. One possible reason for this is that the LRV seems to imply that cognition is to be understood as *homeostasis*, and this is an implication that Godfrey-Smith tries hard to avoid (pp.76-79).<sup>27</sup> Ever since Cannon (1932), the term 'homeostasis' has been used to describe mechanisms in which stasis in certain properties of a mechanism is maintained by variation in others. Cannon originally discussed homeostasis in a physiological context, but it is clear that the notion generalises, and the LRV certainly has a homeostatic interpretation in terms of the maintenance of stability in essential variables (in whatever context) through the deployment of appropriate responses. Therefore let us briefly consider the question: can cognition - *sensu* Godfrey-Smith - be construed as homeostasis?

For Godfrey-Smith, the answer is 'sometimes'. In some cases, cognition will lead to actions that are genuinely homeostatic, for example in the intelligent use of fire to maintain bodily warmth. 'Genuine' homeostasis, for Godfrey-Smith, obtains when "there is some intermediate organic property [body temperature] such that complex activity contributes to the maintenance of stasis in this intermediate property, where this property makes a real contribution to survival" (1996, p.79). However, in other cases there will be no non-trivial homeostatic interpretation of the action of cognition, for example in the adept use of perception and coordination to evade a sudden rock slide. In these cases, cognition "is like hibernation - it is adaptive, but the explanation for why it is adaptive goes *directly* from organic variation to survival" (1996, p.79).

Godfrey-Smith is therefore right to be cautious of *equating* cognition with *non-trivial* homeostasis, even if some instances of cognition - or indeed many - can be understood in this way. However, it seems to me that reading the LRV as a version of the ECT does not necessitate this commitment. Admittedly, there is a temptation to equate Ashby's essential variables with Godfrey-Smith's 'intermediate organic properties' - body temperature, for instance, has been used in examples of both - and such an equivalence would indeed encourage an interpretation of the LRV as non-trivially homeostatic. But at the abstract level of the LRV, it does not appear to be possible to uphold any distinction between an 'intermediate organic property', and a property that is *constitutive* of survival. After all, for Ashby, the essential variables are *defined* as precisely those variables for which their stability is a condition of survival. And to the extent that essential variables are constitutive of survival, rather than mediatory of it, the LRV asserts only a 'trivial' homeostasis: as Godfrey-Smith says, the explanation for why a response is adaptive can go *directly* from variation to survival (survival being, by definition, stability in essential variables).

The mechanisms underlying the generation of behaviour may be genuinely homeostatic in some cases, not so in others. A version of the ECT drawn from the LRV is consistent with this

<sup>27</sup>Ashby himself would probably have preferred to associate cognition with his notion of 'ultrastability' (see, in particular, Ashby, 1952). This idea, although intriguing, will not be pursued further here.

position, and provides a useful conceptual tool for exploring those situations in which genuine homeostasis *does* apply.

#### 2.7.4 Summary

The second element of the present critique of the ECT is therefore that it fails to avail itself of the extended understanding offered by the LRV. This extended understanding consists of: (1) a particular emphasis on distal variability, together with the distinction between that which threatens the organism and that which facilitates its activity, (2) a potential for quantitative expression in terms of measurements of variety, and for empirical exploration in terms of instantiation of disturbances, responses, and essential variables, and lastly (3) an engagement with the idea of cognition as homeostasis, but *not* a blanket commitment to this proposition.<sup>28</sup> It is largely in virtue of these features that the LRV will be a central focus of the first concrete models of this dissertation (chapter 3).

## 2.8 Construction

Earlier in this chapter, in both section 2.3.2 and section 2.5.1, it was suggested that relating behavioural and mechanistic levels of description could require tracing the relations between external environment and *Umwelt*. However, the danger of circularity was also noted in these sections, in that an insight into such translation will require some prior knowledge of mechanism. In this section I resume the thread of this argument via a discussion of the role of ‘construction’ in the ECT, this being the third and final element of the present critique.

### 2.8.1 Causal and constitutive construction

Recall from the beginning of this chapter that alongside externalism and internalism there is also the perspective of *construction*, the explanation of properties of the external in terms of properties of the internal. The most obvious interpretation of construction is when the actions of an agent alter structures of the external environment, for example when a beaver builds a dam. Something about the external environment has changed, and, because this change also figures in the *Umwelt* of the beaver, it can elicit subsequent responses from the animal. For Godfrey-Smith, the key feature of this ‘narrow’, or ‘causal’ sense of construction is that “some change is made to an intrinsic property of something external to the organic system” (ibid., p.146). For reasons explored below, this is the only sense that Godfrey-Smith entertains in his formulation of the ECT.

However, construction can also be interpreted in a ‘constitutive’ or ‘ontological’ sense. For Godfrey-Smith this is the sense in which “[f]eatures of the environment which were not physically

<sup>28</sup>In concentrating on Godfrey-Smith and Ashby I am not suggesting that they are the only authors to engage with these issues. Jean Piaget, for example, is eminently consonant both with the notion of cognition as response, and also with importance of homeostasis. In his *Biology and Knowledge* he says that “life is essentially autoregulation” (1971, p.26), and with respect to response: “Knowing does not really imply making a copy of reality but, rather, reacting to it and transforming it (either apparently or effectively) in such a way as to include it functionally in the transformation of systems with which these acts are linked” (ibid., p.6). Piaget’s writings are notoriously difficult to interpret, and his biological discussions are perhaps the most refractory of all. They are, of course, considerably more complex than captured by the above quotations, but I shall not attempt an exegesis here. To satisfactorily integrate Piaget with the present argument would require many extra pages and, however desirable, is not necessary for present purposes.



put there by the organism are nonetheless dependent on the organisms's faculties for their existence, individual identity or structure" (ibid., p.145). The key features of this sense of construction, features also at the heart of Godfrey-Smith's definition, are that properties of an organism entail the existence of features in the *Umwelt* of that organism, that the relations between these features and features of the external environment need not be straightforward, and that no change to intrinsic properties of things external to the agent need be involved.

Constitutive construction can be manifest in several ways. Simply by moving around, an agent can influence what features of the external environment can influence its activity, without necessarily altering these features as they appear in the external environment (although relational properties of the external environment may have changed, intrinsic properties of things external to the agent will have remained the same). Indeed, even the fact that an agent is a particular *size* can influence the statistical structure of its *Umwelt*, for example by determining whether or not a field is homogeneous with respect to temperature, light intensity and so on. A third manifestation of constitutive construction attaches to the way in which organisms "transduce the physical signals that reach them from the outside world" (Lewontin, 1983, p.100). For example, an (external) environmental change - an approaching rattlesnake - that may entail a change in the rate of vibration of air molecules, which is then transduced by the organism into some feature of the agent's *Umwelt*, perhaps associated with - or identifiable with - changes in the concentration of particular chemicals, which are themselves "transformed by the neuro-secretory system into the chemical signals of fear" (ibid., p.100).<sup>29</sup>

Constitutive construction therefore describes the process by which *Umwelt* is generated from the external environment. As argued previously, an understanding of this process may be helpful - or indeed necessary - in effectively relating behaviour to mechanism. It therefore seems that any interpretation of the ECT that respects the distinction between behaviour and mechanism must entertain a role for constitutive construction.<sup>30</sup>

### 2.8.2 In defence of constitutive construction

Why, then, does Godfrey-Smith explicitly limit his formulation of the ECT to entertain only narrow, causal construction. One reason seems to be that he fears an admission of constitutive construction would necessitate an interpretation of the mechanisms underlying behaviour as accommodatory (recall section 2.6.4). This is evident from his choice of an example of constitutive construction:

Suppose an organism develops a way to detoxify some chemical in its environment which was formerly highly poisonous to it. This organism has made an internal change to its chemistry, and it has also made a change to the relational properties of the external chemical. There is now one less thing the chemical can poison. But if the organism has not, in doing this, made any change to an intrinsic property of

<sup>29</sup>Godfrey-Smith's definition of constitutive construction can be made consistent with this variety of manifestations by modifying the phrase "dependent on the organism's *faculties*" to "dependent on the organism's *properties*" (emphasis added). Also, the idea of the *Umwelt* depending in large part on the activity of an agent has much in common with the collection of ideas variously termed 'active vision', 'animate vision', or 'dynamic vision' (see, for example, Ballard, 1991).

<sup>30</sup>Both of these interpretations of construction are to be distinguished from the more general anti-realist use of the term 'constructivism', in which 'external reality' itself is supposed to be 'constructed' by thought, language, or scientific activity (see Feyerabend, 1988, with respect to the latter).

any external feature then this is a paradigm case of an internal accommodation of the environment. It is the type of thing to be *contrasted* with [narrow] constructive actions such as physically removing the chemical from the environment or spraying something on it to change its intrinsic nature. (1996, p.147)

There are two points to make here. The first is that this example is *not* a paradigm example of internal accommodation in the sense in which accommodation is understood in this chapter. The ‘internal change to its chemistry’ could perhaps be interpreted as a response, and if so, then the kind of constitutive construction alluded to in this example would *not* necessitate accommodation, and as such should not pose a problem for Godfrey-Smith.

The second point, and the more important, is that it is certainly not the case that this example is representative of constitutive construction in general. Consider another hypothetical example. Suppose an organism has a property such that some types of fluctuation in light intensity constitute relevant stimuli for it, whereas others do not. The *Umwelt* of the agent thereby contains these relevant stimuli (or transformations of them). The organism responds to these stimuli (or their transformations) in order to maintain some internal variable within a certain range (perhaps by moving away from certain kinds of bright lights in order to maintain body temperature). In this example, as in the examples described in the previous section, it is clear that constitutive construction plays a role - the generation of *Umwelt* - in determining how the organism should respond to (not accommodate) the environment. Godfrey-Smith is, however, reluctant to discuss this potentially useful interpretation of constitutive construction:

What do we say about the role which organisms play in determining which properties of the environment are relevant to them? [...] We should say that and nothing more: *relevance* is a good concept to capture these phenomena [...] The organism plays a role in making it the case that its environment contains *relevant complexity* or not. But what then is left of the ECT? (ibid., pp.148-154)

Godfrey-Smith addresses these concerns with a ‘concession’ and a ‘bet’. The concession - which he describes as a concession to internalism - is that “the organic system in question does play a role in determining whether or not a given environmental pattern is relevant to it” (ibid, p.155). The bet is that once this role has been played, “there will be other organic properties that can be explained in terms of this environmental pattern” (ibid., p.155). This ‘externalist bet’ is to be contrasted with its internalist counterpart, which would be the wager that once the first role has been played (of generating *Umwelt*), then there is little or nothing left to explain about the organic system. This position seems to suggest that constitutive construction be understood as a *precondition* of the ECT, not as something that can be explained *by* it. But this is a step Godfrey-Smith appears unwilling to take, for reasons that seem to stem from an attachment to the idea that construction must feature, if at all, as a consequence of the ECT, and that as such, it must involve accommodation. My argument is that neither of these suppositions are necessary, and that, first and foremost, constitutive construction describes the process of *Umwelt* generation.

But must constitutive construction always figure as a precondition? It seems just as plausible to consider the generation of an appropriate *Umwelt*, in some cases, as constituting a *response* to external environmental variability. However, this idea immediately encounters some difficulty concerning the place of mechanism in the ECT. As suggested above, to the extent that the generation of *Umwelt* is understood as a precondition of application of the ECT, those aspects of

mechanism involved in this generation can no longer be explained by the ECT; only those aspects of mechanism mediating responses, given the stimuli in the *Umwelt*, would fall within its explanatory domain. However, if the generation of *Umwelt* is construed as response, then those mechanism structures mediating the generation of *Umwelt* do fall within the explanatory domain of the ECT after all. The problem is that there may be no *a priori* way of disambiguating these two interpretations for any given agent-environment system. Furthermore, there may be no way of unambiguously identifying those aspects of mechanism involved in the generation of *Umwelt* and those involved in responding to stimuli in the *Umwelt*. This is the issue of circularity alluded at the beginning of this section, and previously in sections 2.3.2 and 2.5.1.

It is at this point that I believe that to advert to the construction of concrete models is wise and perhaps necessary. If there is no escaping the need for pre-existing internal structure in subserving the translation from external environment to *Umwelt* - and there isn't - then this should simply be accepted. Indeed, any concrete model of agents and environments in continuous interaction and mutual specification must *necessarily* endow agents with some initial structure, if the model is not to be infinitely trivial. In this way, once one starts building concrete models of this sort, issues of constitutive construction become attached to operational details of the model and cease to be philosophical obstacles. (A further consequence of this is that in grounding the ECT in particular patterns of behaviour, it naturally becomes part of evolutionary psychology, an eventuality Godfrey-Smith avoids by retaining only a vague characterisation of cognitive mechanism.)

To give some examples, hypothetical and otherwise. Models that endow agents with stable sensors (but other potentially modifiable aspects of mechanism) can treat the constitutive construction mediated by these sensors as a precondition of application of the ECT. This then clears the way for understanding - in terms of the ECT - the behaviour patterns that respond to features of the external environment, and for understanding - also in terms of the ECT - the mechanistic structures (apart from the sensors) that respond to features of the *Umwelt*. By extension, models that allow the sensors themselves to adapt can admit the constitutive construction associated with this adaptation into the explanatory domain of the ECT. Harvey, Husbands, and Cliff (1994), for example, explore the artificial evolution of sensor morphologies for robots faced with the task of discriminating between triangles and squares, as they appear in the external environment. This, on the face of it a difficult problem, becomes trivial for the robot once it has adapted its sensors to generate a very simple *Umwelt* (see also Husbands, Smith, Jakobi, & O'Shea, 1998).

In either case, constitutive construction associated with agent movement (see section 2.8.1 above) can fall within the explanatory domain of the ECT to the extent that this movement can itself be interpreted a response to the external environment. Todd and Yanco (1996), for example, explore several ways in which the adaptive significance of (externally apparent) resource 'clumps' depends upon the movement of simulated agents in a simple concrete model, and Nolfi and Parisi (1993) assess the performance of systems which have the ability to expose themselves only to sub-classes of stimuli to which they can effectively respond.<sup>31</sup>

<sup>31</sup>The model of Todd and Yanco comprises part of a research project, initiated by Todd and Wilson (1993) which is closely related to the present enterprise insofar as there is an emphasis on understanding how environmental structure influences the structure of well-adapted behaviour (see also Todd, Wilson, Somayaji, & Yanco, 1994). However their primary objective is to establish the properties of a minimal 'baseline' utilising stateless perception-less agents, and as such their project can be understood as complementary to the present research in which perception and/or internal state are important mediators of the relations between behaviour, mechanism, and environment.

In making these arguments, I am nowhere denying that ‘narrow’ or ‘causal’ construction may also have a role to play. I agree fully with Godfrey-Smith here; by intervening in formerly autonomous external environmental processes, an agent may well bring about modifications in that environment - which may entail changes in the *Umwelt* - which can then elicit distinctive responses. All I am saying is that causal construction on its own *is not always enough*.

A final point to make is that the various forms of constitutive construction - and indeed narrow construction - associated with agent activity all serve to recall Lewontin’s critique of the concept of the ‘ecological niche’. In the discussion of this concept in sections 2.1.2 and 2.5.2 it was argued that elaboration of the ways in which agent activity influences its environment falls within the explanatory duty of concrete agent-environment models. Exactly this point is echoed here, but with greater emphasis on the various kinds of influence that can be exercised.

### 2.8.3 A potential confusion

One small issue remains to be cleared up before bringing this discussion to a close. Back in section 2.6.3, the version of the ECT articulated by Dewey was contrasted with Spencer’s accommodatory version by highlighting the fact that agent response could intervene in environmental processes, and “alter the environment’s intrinsic course” (ibid., p.139). On this reading it seems that Dewey’s ECT is necessarily tied up with *causal* construction; a response is a response if and only if it alters the environment’s “intrinsic course”. With the benefit of hindsight, such a strong contrast between Spencer and Dewey is not necessary. As argued throughout the previous pages, responses can also consist in agent-side contributions to agent-environment interaction patterns that do *not* alter intrinsic properties of the external environment (this goes all the way back to the distinction between accommodation and response articulated in section 2.6.4). The original, excessively strong contrast between Spencer and Dewey may therefore be seen to originate from the suspicion that causal construction makes it more *obvious* - to an external observer - that a description in terms of accommodation is inappropriate.

### 2.8.4 Summary

The best way to summarise this section is to succinctly state the interpretation of the ECT that will be adopted throughout the remainder of this dissertation. At its most general, this is that *behavioural and/or mechanistic complexity can be understood in terms of mediating well-adapted responses to environmental variability*.

This interpretation differs from Godfrey-Smith’s formulation in three important ways, corresponding to the three elements of the critique offered above. The first is that there is no commitment to a cognitive interpretation of mechanism, an interpretation which obscures the essential distinction between behaviour and mechanism. This is important because the adaptationist claims of the ECT attach to behaviours, yet any behaviour can be subserved by a variety of mechanisms. This is why there is an obvious ambiguity in my interpretation of the ECT; the relations between behaviour and mechanism cannot be pre-specified in advance of consideration of any particular agent-environment system.

The second difference is that the interpretation of the ECT favoured here accepts - and exploits - the parallel with Ashby’s LRV. This extends the understanding of the role of homeostasis, en-

courages quantitative modelling, and highlights the distinction between environmental variability that threatens the agent and environmental variability that facilitates its activity.

The third difference, and the perhaps most difficult to articulate, is that the present interpretation of the ECT tries to explicitly incorporate constitutive construction. On the one hand as a precondition of application of the ECT (fixed sensor structure, for example), and on the other as a response explicable by the ECT (adaptive sensors, for example, or *Umwelt* generation through movement). This third difference is intimately related to the first insofar as understanding the relations between behaviour and mechanism can be facilitated - or can even require - tracing of the relations between the external environment and the *Umwelt*.

These differences indicate the importance of carrying out experimental work - of the design and analysis of concrete models - to see *just how* environmental variability influences internal mechanism via behaviour. Aside from these three differences, however, the thrust and detail of Godfrey-Smith's formulation is fully retained. This detail cannot be repeated in full here, but, for example, one important idea is that the ECT is a claim about necessity rather than sufficiency; in some situations environmental variability can be coped with perfectly well by simple behaviours subserved by simple mechanisms.

I have presented the ECT in some detail in this chapter, but my treatment has by no means been exhaustive. There are many other interesting and difficult issues surrounding the ECT which have not been dealt with here, and the inquisitive reader is directed primarily to Godfrey-Smith (1996), but also to the various commentaries in Todd (1996) and Hardcastle (1999).

## 2.9 The shape of things to come

With the work of this chapter done, it is worth stating explicitly the ways in which its ideas relate to the empirical material of the following chapters (although this material in no way exhausts the potential of these ideas).

Above all, this chapter motivates the use of artificial evolution models as explanatory tools (see chapter 1 for operational details of this kind of model). Artificial evolution models are strongly externalist, and can separate intuitively into agent and environment components (although this is not a necessary attribute). Most importantly, they are not required to prefigure the relations between behaviour and mechanism; the designer can specify a (functional) fitness function, and leave the evolutionary process to work out the details of the underlying mechanism. (It is still incumbent on the designer, of course, to specify in advance *some* aspects of mechanism structure.) Also, artificial evolution models are of course concrete in the sense that the relations between behaviour, mechanism, and environment are empirically instantiated. As such they allow an externalist perspective to be reconciled with enlightened interactionism, and they also permit the process of constitutive construction to be explicitly traced, both as a precondition to the application of the ECT and as an explanatory target of its application.

The remainder of this dissertation will explore artificial evolution models in a variety of related contexts (chapter 1 provides a thorough summary). The following two chapters are concerned with complexity. The first will employ the empirical context of the 'iterated prisoner's dilemma' to explore the central assertion of the ECT, that environmental complexity can promote 'complex' responses; the parallels between the ECT and the LRV (section 2.7) will be of particular impor-

tance in analysing the dynamics of the models of this chapter. The second chapter explores the same issues in the context of evolutionary robotics, a context in which the distinction between behaviour and mechanism is of immediate and inescapable significance.

The second part of this dissertation focuses on relatively specific issues in psychology and theoretical biology, and as a prelude to these investigations, chapter 5 provides a methodological bridge between the relatively abstract discussion of the present chapter and the firm biological foundation of optimal foraging theory (section 2.5.3). The chapters that follow this are concerned with various aspects of behavioural coordination, and these investigations are concerned firstly with the mechanisms underlying the generation of choice behaviour and how they are to be understood in view of the intimacy of perception and action, and secondly with the exploration of instances of apparently *irrational* behaviour in terms of variations in environmental structure.

## 2.10 Overall summary

The essential duty of this chapter has been to present a range of themes and concepts which structure the remainder of this dissertation. Foremost amongst these are the distinction between behaviour and mechanism, and the development of the (pragmatic) externalist perspective on the relation between environmental structure and the functional properties of behaviour, the latter being most effectively articulated by the ECT. The integration of these themes has been achieved by appeal to a three-pronged critique of Godfrey-Smith's formulation of the ECT in which constitutive construction is identified with the process of *Umwelt* generation, and in which 'cognitive' interpretations of mechanism are discarded, and in which the parallels between the ECT and Ashby's LRV are exploited. In its simplest form, the message is this: the ECT articulates the hypothesis that environmental complexity can incur behavioural complexity, but to understand how such behavioural complexity relates to underlying mechanism it is of enormous importance to understand how the agent perceives its environment, how the external environment is translated into *Umwelt*. This kind of understanding can figure both as a precondition, and as a consequence of application of the ECT.

This interpretation of the ECT - which constitutes the primary original contribution of this chapter - motivates the empirical exploration of concrete models in which the various dependencies between behaviour, mechanism, and environment can be empirically instantiated. However, the coherence of this version is *not* a necessary precondition for the coherence of the empirical excursions of the following chapters, each of which can be judged on its own merits.

There are two further themes which have accompanied this integration, these being the inseparability of perception and action, and complexity itself. Throughout this chapter I have tried to locate all these themes within a broad historical context, and although the genealogy of ideas that I have traced is necessarily incomplete, the hope is that it has been sufficient to point to an underlying unity in the diversity of issues that have been covered. This diversity will become even more apparent as this dissertation enters its empirical stages.

## Chapter 3

### **Evolving complexity I: Definitional issues and game-theoretic models**

---

This, and the following chapter of this dissertation, present the first empirical explorations of some of the ideas introduced in the previous chapter; most notably in this chapter the environmental complexity thesis (ECT, Godfrey-Smith, 1996), and the law of requisite variety (LRV, Ashby, 1956). The focus is on *complexity*, and the objective is to employ evolutionary game-theoretic models to explore the hypothesis that environmental variability can promote the adaptive evolution of internal complexity. A hypothesis such as this requires considerable clarification, and the first part of this chapter is dedicated to a discussion of definitional issues surrounding the term ‘complexity’ itself, and to a classification of theoretical positions concerning the ‘evolution of complexity’. It will become clear from this discussion that any empirical exploration of these issues must to some extent forfeit generality, a condition respected in the second part of the chapter.

The empirical context of this chapter is provided by the application of artificial evolution to the ‘iterated prisoner’s dilemma’ (IPD), and two distinct models are constructed and explored; one inspired by the variable memory strategies of Lindgren (1991), and one by a class of strategies known as ‘win-stay, lose-shift’ strategies (Colman, 1995; Posch, 1999). An advantage of the IPD is that it easily submits to quantitative analysis, and as such I will argue that the dynamics of IPD models can be usefully interpreted in terms of Ashby’s LRV; this is in contrast to most ‘applications’ of this law which remain at the level of loose analogy. A disadvantage is that the distinction between behaviour and mechanism, although certainly applicable at the abstract level of the IPD, is not as intuitively significant at this level than it can be at the level of spatially explicit agent-environment systems; we will see that this is so in the following chapter.

Results from both IPD models support the conclusion that complex strategies can evolve as an adaptation to environmental variability. However, not all the empirical data points directly to this conclusion; it is also argued in this chapter that variability in an evolving system can - in some circumstances - facilitate the exploration of genotype space by an artificial evolutionary process. This new theme is also taken further in the next chapter.

### 3.1 Definitions of complexity

Recall Godfrey-Smith's definition of complexity from chapter 2: "Complexity is changeability, variability; having a lot of different states or modes, or doing a lot of different things. Something is simple when it is all the same" (1996, p.24). It is fair to say that *definitions* of complexity themselves fall under this description, they are numerous and certainly diverse, and a comprehensive survey is beyond the ambit of this chapter. Only a small selection of definitions will be considered here, and these can be grouped into two broad camps; bio-philosophical and information-theoretic.

An important qualification, before we begin, is that definitions in both groups must account for the role of the external observer. What is simple for one observer may appear to be complex for another (recall Simon's 'ant on a beach' from the previous chapter). In general, therefore, one should not talk about the 'complexity' of  $X$  itself, but about the complexity of a *description* of  $X$ ;<sup>1</sup> I will address the potential arbitrariness of this qualification later in this section.

Bio-philosophical definitions cover both organism and environment. Emery (1967), for example, distinguishes four levels of environmental complexity which have to do with the structure and predictability of resource distributions; for example 'placid' environments present random unchanging distributions, but resources in 'turbulent' environments are influenced both by the behaviour of foraging animals and by intrinsic environmental dynamics. With respect to organism complexity, Bonner (1988) is very specific, defining his 'somatic complexity' as the number of different cell types present within an organism. A more abstract definition is provided by McShea (1993), for whom 'morphological complexity' is a function of "the number of different parts [...] and the irregularity of their arrangement" (p.731). Although McShea uses this definition predominantly to assess the complexity of extinct species - which he does through analysis of the arrangement of vertebrae in fossilised skeletons (McShea, 1996) - the definition is abstract enough to apply to both organisms and environments. Indeed, its level of abstraction is similar both to that of Godfrey-Smith (above) and to that of Simon (1988), for whom a complex system is "made up of a large number of parts that have many interactions" (pp.183-184); Simon also stresses the supposed hierarchic nature of complex systems. These various definitions may or may not hold intuitive appeal, but they are in any case either inappropriate outside the biological realm (Bonner), or, in virtue of their abstract nature, resist quantitative application altogether. A more widely applicable class of definition can be found under the auspices of information theory.

Information-theoretic definitions of complexity can only be applied to symbolic sequences, but this is not as restrictive as it may initially appear. A symbol sequence may well constitute a description (or observation) of a system, and it has already been conceded that any definition of complexity should always be taken to concern the complexity of a description of a system, rather than the complexity of the system itself. Most information-theoretic metrics of complexity derive in some way from Shannon entropy, a measure of the uncertainty present in any symbol sequence (Shannon & Weaver, 1949). The most widely used such metric is Kolmogorov-Chaitin (K-C) complexity, otherwise known as 'minimum description size' (Kolmogorov, 1965; Chaitin, 1966).

<sup>1</sup>See Pringle (1951) for an early exposition of this argument. More recently, Kolen and Pollack (1995) demonstrate, using an ingenious mathematical device, that "changes in the method of observation can radically alter both the number of apparent states and the apparent generative class [*sensu* Chomsky] of a system's behavioural description" (p.2), such that "[t]he holistic combination of the organism and symbolizing observer can create apparent computational systems independent of the actual internal behaviour producing processes" (ibid., p.17).



This defines the complexity of a given sequence as being the size of the minimal program which will reproduce the sequence when run on a universal Turing machine. This definition is certainly general; the K-C complexity of any system that can be described with a symbol sequence can be assessed, however, it carries the implication that purely random sequences are maximally complex.

Adami and Cerf (2000) argue that a metric of complexity should yield a vanishingly small measurement for both entirely regular and entirely random sequences, and suggest that the shortcoming of K-C complexity derives from neglecting the role of the environment. Their ‘physical complexity’ can be thought of as the size of the minimal program without the part necessary for producing the randomness, in the context of a given environment. More formally; given an environment  $U$  and a sequence  $S$ , the physical complexity of  $S$  is the length of  $S$  (this being the ‘unconditional complexity’), minus those bits in  $S$  which are random with respect to  $U$  (this being the ‘conditional complexity’).<sup>2</sup> This definition has the interesting property that the ‘observer’ of complexity is the environment itself, and indeed Adami and Cerf argue that the physical complexity is a measure of the information about the environment that is coded in  $S$ . However, a problem for the empirical application of this definition is that inspection of a given  $S$  will not tell you which bits, or even how many, correspond to information about  $U$ . In practice, the physical complexity must be estimated from the average number of fixed loci over an ensemble of sequences, assuming that they all occupy the same environment.

A similar approach has been developed over recent years by Crutchfield (1994a, 1994b) in the form of ‘statistical complexity’, measures of which “discount for randomness and so provide an indication of the regularities present in an object above and beyond pure randomness” (Feldman & Crutchfield, 1998, p.1). However, much debate attends how statistical complexity is to be measured, and this debate, in taking place largely within the domain of theoretical physics, cannot be described here. Suffice it to say that a general consensus has yet to be reached.

An important distinction between the approaches of Crutchfield and of Adami and Cerf concerns the evolutionary phenomena towards which their definitions are addressed. Adami and Cerf are interested in how an ensemble of sequences can come to encode information about an environment. This is a process of *adaptation*; the evolution of complex forms displaying an adaptive fit to properties of the environment. Crutchfield is more interested in how a given system can ‘inductively leap’ to a new ‘model class’. This is the problem of *innovation*, which in evolutionary biology refers to the emergence of strikingly new biological forms of qualitatively different levels of complexity. Maynard-Smith and Szathmáry (1995), for example, identify the genesis of multicellular life, the evolution of eukaryotes from prokaryotes, and the origin of sexual reproduction amongst what they call ‘the major transitions of evolution’.<sup>3</sup> I will return to this distinction between adaptation and innovation in section 3.2.

Many papers on complexity open by abstaining from the endorsement of any particular definition, and I too concur that no single definition can be satisfactory in all regards; it is too labile a concept, perhaps too prone to be associated with arbitrary properties. Information-theoretic definitions can be commendably mathematically precise, but perhaps the lasting impression from this

<sup>2</sup>If  $U$  represents the ‘usual rules of mathematics’, then the conditional complexity reduces to the K-C complexity. See also Tononi, Sporns, and Edelman (1994, 1996) for a similar complexity metric, designed to be applied directly to the complexity of brains (real and artificial).

<sup>3</sup>See also Eigen and Schuster (1982).

section is that complexity is best thought of - to paraphrase William James - as some kind of contextually-specific, observer-related, heterogeneity and interconnectedness.<sup>4</sup>

This does not mean, however, that it is necessary to resort to a stance of simple relativism by which any description of a system - to which measurements of heterogeneity and interconnectedness can be applied - is as good as any other. Definitions and measurements of complexity will always be associated with particular explanatory projects, and descriptions of systems are usually *not* arbitrary with respect to these projects. For example, if one's project is to understand the causal interactions underlying the generation of behaviour in an agent-environment system (this being the domain of the ECT), some descriptions of this system will elucidate these interactions, but others will render them opaque. The first class of description is to be preferred, and measurements of complexity attached to these descriptions are therefore not arbitrary with respect to the explanatory project at hand.

To relate this idea more directly to present circumstances, the central hypothesis of this chapter - derived from the ECT - is that environmental variability can promote the adaptive evolution of agent-side complexity. The ECT, discussed in detail in the previous chapter, is intended as a realistic (as opposed to relativistic) thesis, relating environmental complexity to agent-side complexity. What is needed in the remainder of this chapter is therefore a non-arbitrary description of agent-side complexity, with respect to the agent-environment interactions constitutive of the domain of the ECT. Armed with such a definition, observed increases in agent-side complexity according to this description can be taken as indications of increases in (unqualified) agent-side complexity itself.

It is indeed this strategy that is adopted in what follows. The metrics of complexity that are employed later in this chapter will be very specifically, but non-arbitrarily tied to the nature of the concrete models that are developed. Further details of these metrics will be deferred until the models themselves are described.

### 3.2 The evolution of complexity

A perennial topic in biological science is that of the 'evolution of complexity'. How can complexity evolve? Is there any reason to suppose that the process of evolution inherently engenders an increase in the complexity of the evolving organisms? The evolution of complexity is also the topic of this chapter, and I would like to continue this discussion by suggesting that the distinction between these two questions be made sharp:

- the *causal* question: how is it possible for evolution to produce structures of complexity (from structures of lesser complexity)?
- the *teleological* question: is evolution a necessarily progressive force that tends towards greater complexity?

I believe that the only interesting question is the causal question, but that much of the evidence that could otherwise help to answer this question is obscured and tarnished by association with the

<sup>4</sup>One of William James' numerous sideswipes at Herbert Spencer was his parody of the Spencerian view of evolution: "Evolution is a change from a nohowish untalkaboutable all-alikeness to a somehowish and in general talkaboutable not-all-alikeness by continual stickingtogetherations and somethingelseifications" (quoted in Perry, 1935, p.482).

teleological question. I will argue that the teleological question is uninteresting primarily because it can be both trivially affirmed and trivially negated at the same time. Before this, some general perspectives on the evolution of complexity will be set out.

### 3.2.1 Three perspectives

Crutchfield (1994b) provides a useful characterisation of three perspectives on the evolution of complexity, engaging with both adaptation and innovation. His ‘selectionists’ argue that all forms of biological complexity are adaptations produced by orthodox natural selection; this is a strongly externalist position. The ‘historicists’ also embrace Darwinian principles, but claim that much major structural change can be (and has been) non-adaptive and arbitrary. Finally, the ‘structuralists’ look to elucidate the ‘principles of organisation’, or ‘structural attractors’ that exist in the space of biological form. Evolution may or may not discover these attractors, but cannot change, create, or destroy them; this view is strongly internalist. Crutchfield himself does not commit to any of these views, but I shall discuss an example of each in the light of the distinction between the causal and teleological questions.

Adami and Cerf (2000, see also Adami, Ofria, and Collier, 2000), who can be located amongst the selectionists, believe that evolution acts as a ‘natural Maxwell demon’,<sup>5</sup> with adaptation the process by which a population of genotypes accumulates information about an environment. Critically, selection will ensure that, across an ensemble of genotypes, information (on average) is never lost; thus follows their conclusion that evolution naturally tends towards increasing (physical) complexity. In reaching this conclusion, Adami and his colleagues are addressing both the teleological and the causal questions, and answering the former in the affirmative. Their argument concerns adaptation, rather than innovation, but rests on the assumption of a static environment. If there is change in the environment, then information encoded in the ensemble of genotypes may no longer be information about the environment, and so will no longer necessarily be maintained by selection. Physical complexity can therefore decrease, rendering their positive answer to the teleological question insubstantial. However this is not to say that any causal understanding, provided by their ideas about information accumulation, is thereby devalued.

A more historical view is formulated by Schuster (1996). He is primarily interested in innovative episodes, and argues that two conditions must be fulfilled for evolutionary innovation to occur. The first has to do with evolutionary mechanics; he asserts as pre-requisites mechanisms for gene duplication and for symbiotic interaction. It is his second condition that is overtly historical; he argues that radical innovations can only occur during periods of environmental resource abundance. In epochs of scarcity, the argument goes, significant variations on already adaptive schema will be likely to suffer strong negative selection. But if competition is reduced (during periods of resource abundance) then significant variations may survive and provide pathways for the evolution of innovatively different organisms. The question then arises, how do such resource abundances come about? One scenario concerns those periods immediately following mass-extinction events. But

---

<sup>5</sup>A ‘Maxwell demon’ is a thermodynamic thought experiment in which random molecular kinetic energy is redistributed in such a way that energy can be extracted from a system that was previously in equilibrium. A common version involves the collection of slow moving molecules on one side of a demon-operated ‘trap-door’ which separates two otherwise sealed compartments. See Wiener (1948) for further details and discussion in the light of thermodynamics.

mass-extinctions are paradigmatically historical events. Consider a mass-extinction occasioned by an asteroid impact; its occurrence cannot be explained by evolutionary theory, nor does the event itself provide any explanation of the structure of the various evolutionary innovations that may follow. Schuster's discussion therefore cannot be addressed to the teleological question, and he himself never attempts to do so. His first condition, however, is an explicit response to the causal question, and his second suggests some environmental constraints on how the causal question might be answered.

The internalist theoretical position of the structuralists is, of course, quite deliberately beyond the remit of Darwinian evolutionary theory. Nevertheless, one suspects that a deep motivation for many structuralists is a belief that an evolutionary theory based on random mutation and selection *cannot* explain the complexity of life; in other words that the answer to the teleological question is *no*. It remains possible, however, that structuralist research can contribute to the causal question in the same way as the historicists, through identifying possible constraints on the evolutionary process. Kauffman (1993) is particularly vocal on this subject, arguing that his principles of 'self-organisation' underlie and influence the distribution of variety upon which natural selection operates.<sup>6</sup>

Crutchfield's three groups, then, can all potentially contribute to the causal question, but not at all to the teleological question. However, although the historicists and structuralists rather remove themselves from the debate by moving the theoretical focus beyond the range of evolutionary theory, what about the selectionists? The single case considered above should not be taken to represent all selectionist theories. But let us explore what it actually means to ask the teleological question in the first place.

### 3.2.2 The trouble with teleology

Over the course of biological evolution, the complexity of the most complex organisms *has* increased. This much may be conceded, but this does *not* mean that it is sensible to ask if evolution is responsible for this increase. Gould (1997) argues that the observable upward trend in the complexity of the most complex organism is an artefact of a *lower limit* on complexity. What some interpret as a trend towards increasing complexity is really just a fairly monotonic increase in *variation*. Take a species at random, says Gould. Maybe at the next speciation event there will be a more complex species, maybe a less complex one. But given a lower limit on species complexity, over time an increase in average complexity will *necessarily* be observed. Gould elucidates this point with a vivid analogy. A drunk stumbles out of a pub, the door of which is set into a wall. With each step he takes he is equally likely to stagger an equal distance, either forwards, or backwards, across a road. On the far side of the road is a ditch. How likely is it that drunk will end up in the ditch? The answer is; so long as the drunk cannot walk through walls (the lower limit), he will *always* end up in the ditch after some amount of time. His average position will necessarily 'trend' away from the pub towards the ditch.

Consider the size of horses. One of the most famous evolutionary 'trends', known as Cope's rule, is the steady increase in horse size over evolutionary time. However, recent palaeontological

<sup>6</sup>It is also possible that some (or all) biological structuralist principles 'merely' reflect very early adaptations (or 'historical accidents') rather than 'laws' of biological form (Burian & Richardson, 1996; Dennett, 1995). However, even if this is so, the interpretation of such principles as constraints on the evolution of complexity remains the same.

evidence suggests that horse speciation events were equally likely to deliver larger or smaller horse types (Arnold, Kelly, & Parker, 1995). Given some lower limit on horse size, this ‘trend’ is exactly analogous to the drunk in the ditch. We observe a ‘trend’, but there need not be a ‘mechanism’ within evolution that produces this trend.<sup>7</sup> On the other hand, consider also the kingdom of bacteria, which, whilst generally considered to be less complex than most other examples of life on the planet, constitute by far the largest proportion of distinct organisms and distinct species in existence, and which quite possibly are also responsible for the largest quantity of biomass (see Gold, 1992, for a speculative discussion). Furthermore, this situation has persisted largely unaltered ever since the evolution of bacteria in the first place, some billions of years ago.

Now it is quite clear just how awkward the teleological question is. *Of course* evolution will entail increasing complexity (the palaeontological evidence is incontrovertible, and just replace ‘size’ with ‘complexity’ in the discussion of Cope’s Rule). *Of course* evolution does *not* entail increasing complexity (the observation of a trend does not imply the existence of a mechanism dedicated to the entailment of that trend, and anyway consider the continuing predominance of bacteria). It is worth noting that these arguments are not arguments against there being teleological theories of evolution - after all, the observation of a trend does not imply the *non*-existence of a trend-generating mechanism - rather, they are arguments against the *relevance* of asking the teleological question, given some basic historical and dynamical facts about biological evolution.

### 3.2.3 Evolution and progress

These arguments against the relevance of the teleological question are neither novel nor difficult, yet the association of evolution with the emergence of complexity has proven remarkably tenacious, both in academic circles and for the general public. Levins and Lewontin (1985) trace this tenacity to the ideological background to evolutionary theory. They note that the development of biological evolutionary theory was part of a sea-change in nineteenth century thought from theories of stasis to theories of change, reflecting - as they argue - the alterations in European social structures that had seen the replacement of (unchanging) hereditary power, with power derived from entrepreneurial activity. In science, this sea-change was first evident in the uniformitarianist geology of James Hutton (1726-97) and Charles Lyell (1797-1895), and eventually found biological expression in Spencer’s ‘law’ of evolution: “From the earliest traceable cosmical changes down to the latest results of civilisation, we shall find that the transformation of the homogeneous into the heterogeneous is that in which progress essentially consists” ([1857] 1915, p.10).<sup>8</sup> And, as the quote from Spencer illustrates, attached to this notion of change - even equated with it - was the idea of progress; for Spencer, change of any kind was by its very nature progressive, “a beneficent necessity”. Evolution without progress was not seen as evolution at all, a view that prevailed well into the twentieth century. Witness A.N. Whitehead, writing in 1925: “Evolution,

<sup>7</sup>Compare the mistake of assuming the existence of a ‘trend-generating mechanism’ with the distinction between behaviour and mechanism elucidated in chapter 2. Also, there are a number of examples of evolution tending in the direction of *simplicity*; for example McShea (1991) observes a decrease in the ‘morphological complexity’ of several aquatic mammals, and there is the classic example of the cave fish lineage which became blind after many generations of darkness (Jeffery, Strickler, Guiney, Heyser, & Tomarev, 2000; Teyke & Schaerer, 1994), see also Diamond (1996).

<sup>8</sup>This is the same ‘law’ that was summarised in chapter 2. It is interesting to note that biology was one of the last areas of intellectual life to adopt an evolutionary perspective. And as Levins and Lewontin point out; “[t]o understand the development of the modern theory of organic evolution, it must first be realised that Darwin was the culmination and not the origin of nineteenth century evolutionism” (1985, p.27).

on the materialistic theory, is reduced to the rôle of being another word for the description of changes of the external relations between portions of matter. There is nothing to evolve, because one set of external relations is as good as any other set of external relations. There can merely be change, purposeless and unprogressive” ([1925] 1960, p.157, quoted in Levins & Lewontin, 1985, p.12). This equivalence of change and progress, according to Levins and Lewontin, can also be attributed to ideology. The breaking of the restrictive (and static) chains of hereditary power was bound to lead to the betterment of society, or so it seemed at the time. Finally, if evolution entailed progress, how else could this progress be manifest in biology if not in the emergence of the complex from the simple, of *Homo sapiens* from unicellular slime; this is the picture immortalised by the *Scala Naturae*, or ‘the great chain of being’. And Spencer once again obliges, with his progressive ‘law’ of evolution assuring an inevitable, beneficent, and upward change in the complexity of all systems.<sup>9</sup>

Levins and Lewontin provide one interpretation of the persistent tendency to affirm the teleological question. There are others.<sup>10</sup> However, what is important is that the redundancy of this question is understood, thus opening the way for a clear-headed consideration of the causal question. At the outset of this section I conceded that the complexity of the most complex organisms has increased over evolutionary time. There are three ways to treat this fact. First, as requiring no further explanation (thanks to the influence of a lower limit on complexity). Second, as the consequence of structuralist ‘laws of form’. Third, as an adaptive response to environmental complexity. The first way does not admit any further investigation, the second is thoroughly internalist and therefore contrary to the pragmatic externalism of this dissertation (although it may of course have much to contribute), and so it is the third way that occupies the remainder of this chapter.

### 3.2.4 The causal question

It is time to narrow this discussion down to the point at which empirical work can be usefully undertaken. This involves framing an externalist version of the causal question. Perhaps the most significant body of work concerning the causal question has to do with innovatory episodes very near the origin of life itself. For example, the hypercycle hypothesis of Eigen and Schuster (1977, 1978a, 1978b, see also Forst, 1997) explores how replication error thresholds may be overcome through the formation and persistence of ‘hypercycles’. However, although (externalist) selection does play a role in these models, they remain internalist in character, stressing the interplay of components internal to the evolving system. In what follows I will be concerned not with inno-

<sup>9</sup>In their 1985 book *The Dialectical Biologist*, from which these ideas are largely drawn, Levins and Lewontin are primarily concerned with applying the Marxist philosophy of dialectical materialism to evolutionary theory. However, the issue of progress is one point at which they demur. Literal Marxism, itself a product of the nineteenth century, is as committed to the idea of progress as was the incipient capitalism of the day.

<sup>10</sup>Consider, for example, Freud’s three strikes against the universal significance of the human condition: the heliocentric revolution of Nicolaus Copernicus, removing the earth from the centre of the universe, the Darwinian revolution, challenging the separation of *Homo sapiens* from the rest of the animal kingdom, and (modestly) his own theory of the unconscious, which will not be discussed here. With regard to the Darwinian revolution, the notion of the *Scala Naturae* and the associated ideas of evolutionary progress can be seen as a way to preserve human dignity; if we are not to be entirely separate from the animal kingdom, then let us at least be the highest form of animal, the ideal creature towards which all others are continually struggling via the tribulations of natural selection. Note also that similar suspicions could also potentially attend a structuralist position, in that it might be possible to construe the search for ‘laws of biological form’ as an attempt to justify human inevitability; does *Homo sapiens* represent a particularly ‘strong attractor’ in the space of biological form? An anthropological study performed on the scientists at the Santa Fe Institute would make for interesting reading (Helmreich, 1995).

	<i>B cooperates</i>	<i>B defects</i>
<i>A cooperates</i>	R(3)-R(3)	S(0)-T(5)
<i>A defects</i>	T(5)-S(0)	P(1)-P(1)

Table 3.1: Standard IPD payoff table for two players, *A* and *B*. Each entry gives the payoffs awarded to each player, the left-side figure for *A*, the right-side for *B*. The payoff values illustrated are those most commonly used, but this table will represent the IPD for all values such that  $T > R > P > S$  and  $2R > T + S$ .

vation, but instead with adaptation, and specifically with how variability in the environment can influence the evolution of complex adaptations. In other words, I will explore empirical versions of the externalist project epitomised by Godfrey-Smith’s ECT and Ashby’s LRV, as described in chapter 2. This will involve the use of artificial evolution, in this chapter to evolve strategies to play the IPD, and in the following chapter to evolve controllers for mobile robots. It is to be emphasised that although the relation between artificial evolution and biological evolution is necessarily loose, the ECT and the LRV concern adaptation in general, and do not rely on any particular instantiation of the adaptive process. (This loose relation will have to be tightened somewhat when we come to discuss more specific questions in theoretical biology, but this discussion may be deferred until chapter 5.)

### 3.3 Complexity in the iterated prisoner’s dilemma

Imagine that you and an alleged accomplice have both been arrested and accused of a terrible crime. You are held in separate cells, and upon interrogation you can either *cooperate* by denying all knowledge, or *defect* by implicating your accomplice, who is faced with the same decision. You have no means of communicating with each other, but if you both cooperate you will both be released (the reward  $R$ ), and if you both defect then you will both be jailed (the punishment  $P$ ). However if you defect and she cooperates, then you will receive a bounty (the temptation  $T$ ) and she will go to jail for longer (the sucker  $S$ ). But if she defects and you cooperate, then you yourself are the sucker. A dilemma is evident; in a single meeting you will always do best to defect, in doing so either receiving the bounty or avoiding an excessive jail term. But of course the logic is the same for your accomplice, and if you both defect then both of you will do worse than if you both had cooperated. This is the situation represented by the payoff table (3.1).

This dilemma - the prisoner’s dilemma - is well renowned for providing insight into how cooperation can emerge in a population of selfish individuals. In the one-shot dilemma cooperation is unlikely to arise, as the anecdote suggests, but if the same players meet time and time again, and if they can retain some memory of the outcomes of previous interactions, then it does become possible for cooperative strategies to establish themselves (Axelrod, 1984).<sup>11</sup> It is this ‘iterated’

<sup>11</sup>The most famous of these strategies is probably ‘tit-for-tat’: cooperate on the first move, and thereafter do whatever your opponent did on the previous iteration. (For a comprehensive overview of evolutionary game theory see Maynard Smith, 1982.) Note that the prisoner’s dilemma is employed to model cooperation only in terms of reciprocal altruism. Other evolutionary scenarios in which it has been argued that cooperation can arise include kin selection (Hamilton, 1964), group selection (for example Boerlijst & Hogeweg, 1991), and by-product mutualism, in which cooperation follows from immediate self interest (Dugatkin, Mesterton-Gibbs, & Houston, 1992). The prisoner’s dilemma is *not* concerned with these scenarios, and neither is this chapter.

prisoner's dilemma (the IPD) that constitutes the empirical locus of this chapter.

Many researchers have artificially evolved strategies to play the IPD, primarily for the purposes of understanding the conditions under which cooperative strategies may be expected to arise and persist in a population (see Langton, 1995, for a selection). In the present model, however, the evolution of cooperation *per se* is not of primary interest. Rather, the evolution of a cooperating population provides a useful platform upon which investigations of the evolution of strategy *complexity* can be pursued. Strategy complexity is interpreted as flexibility; a strategy that always cooperates or always defects is less complex than one which can take into account patterns of past interactions and produce a response based on these patterns. In terms of the distinction between behaviour and mechanism, this interpretation of complexity is therefore mechanistic. It is about the (mechanistic) potential of a strategy to produce a variable pattern of behaviour, not about behavioural variability itself. In the light of section 3.1 this metric can be argued to be appropriate for two reasons; it is consistent with the general notion of complexity as heterogeneity/interconnectedness, and it makes sense - as we shall see - in terms of the causal relations underlying the observable behaviour of the IPD system.

The general pattern of the models in the literature, also followed by the model developed in the present chapter, is that individuals in a population are represented by genotypes, each specifying a strategy for playing the IPD; whether to cooperate or defect in any given interaction. During each generation, individuals engage each other in a series of two player games each consisting of a number of IPD interactions, with the fitness of each individual simply being the total payoff accrued by that individual. The present model is distinguished from this general situation in two ways. First, by utilising variable length genotypes, with genotype length specifying the number of previous moves that can be taken into account when delivering each move. In this way, genotype length - strategy 'memory' - provides an easily measurable indication of strategy complexity. Second, by introducing environmental variability in the form of noise. During each interaction, and for each player, there is a certain probability that the move specified by the genotype is flipped (cooperation to defection or *vice-versa*), and only then are the payoff scores calculated with reference to the payoff table. This M-noise (for 'move'-noise) can be interpreted as environmental variability since the environment for any given individual consists simply of the moves made by the other members of the population with which that individual interacts.

Lindgren (1991) utilised variable length genotypes in an evolving IPD model, and demonstrated the emergence of strategies with increasingly long memories in noisy environments. However, he did not explore the causal role of this variability in the emergence of these strategies beyond suggesting that "noise may disturb the actions performed by the players, which makes the problem of the optimal strategy more complicated. This increases the potential for having long transients showing evolutionary behaviour" (1991, p.296). In particular, no control condition - without noise - is discussed. Variable length genotypes were instantiated by Lindgren through 'splitting' and 'doubling' mutations, which would respectively increment or decrement the potential memory of the mutated strategy by a single move. Under this scheme, strategy changes will only occur if the extra memory is subsequently recruited (through further point mutations or crossover in the new genotype segment) to discriminate between possible moves. This is arguably a good thing for present purposes because it means that extra memory cannot be selected for indi-



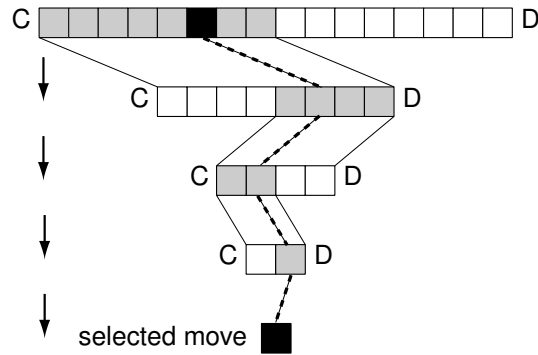


Figure 3.1: IPD model genotype encoding scheme. Each time a previous move in the game history (between two particular individuals) is considered, half of the genotype is (temporarily) discarded (the non-shaded areas in the example illustrated) - one half if the move had been cooperative, or the other if it had been a defection. In this way, the genotype illustrated here - of length 16 - can encode a strategy with a memory of 4 moves (after cutting a string of 16 characters in half 4 times, you are left with just a single character). The black square indicates which allele would be accessed for a [c,d,c,d] history. In practice the genotype must actually be longer than this in order to specify the initial moves up until this memory limit is reached. In this case the genotype would require an extra 15 alleles to code for the initial 3 moves before the final 16 alleles can be used.

rectly by virtue of automatically incurring some serendipitous strategy modification. Notice, however, that splitting (memory reductions) may well incur a strategy change if the discarded alleles were previously being used to dictate a strategy utilising the full memory potential. Nevertheless, Lindgren's model provides an appropriate starting point for the present investigation.

### 3.3.1 Structure of the present model

In the present model, each individual in a population of 30 consists of a genotype, comprising of a string of c's and d's, determining the strategy of that individual for playing the IPD. The longer the genotype, the more it can be influenced by the history of the game, thus the longer the 'memory' of the individual. Figure 3.1 illustrates how genotypes can encode strategies. The maximum genotype length employed was (arbitrarily) 127 alleles, allowing for a maximum memory of 6 moves.<sup>12</sup> The model uses a generational tournament selection genetic algorithm (GA), with point mutation, crossover, and 'splitting' and 'doubling' mutation operators. Doubling copies the latter half of the genotype twice again onto its own end, thereby preserving an appropriate section for initial move specification and also incrementing the memory by 1; splitting the genotype reverses this operation and its effect. (Appendix A provides further implementational details.)

Another feature of this model is that during each generation, each individual engages in the IPD with every other member of the population (for 60 iterations per contest) in what is often called a 'round-robin' tournament. As such it will be called the 'compulsory' model. The hypothesis

<sup>12</sup>This encoding scheme is also based on Lindgren (1991), with the difference that Lindgren did not employ additional genotype portions for initial move specification since he explored his model analytically, assuming an infinite number of interactions during each game. Also, note that a memory of 4 moves does *not* imply a memory of 4 entire IPD interactions; there are 2 moves per interaction; an individual remembers its own moves as well as its partner's.




Figure 3.2: Evolution of complexity in the compulsory model. Results from 12 evolutionary runs are superimposed.

under test, using this model, is that the presence of M-noise will promote the adaptive evolution of complex strategies. (See appendix A for full specification of the model parameters employed in this investigation.)

### 3.3.2 Results from the compulsory model

Twelve evolutionary runs, of 10,000 generations each, were performed either with or without M-noise. Figure 3.2 shows that without any M-noise, complex strategies never really evolve; the evolved memory tends to stay either at the initial level of 1, or drop to 0. With 2% M-noise, however, complex strategies *do* evolve; not always, and not always to the maximum, but it does happen, and applying t-tests to the sets of means from each condition demonstrates that there is significantly more evolved memory in the M-noise than in the zero noise condition ( $p < 0.05$ ).

These results support the hypothesis set out above, however, they do not by themselves elucidate the functional role of the evolved memory. The possibility cannot be excluded at this stage that the evolved memory is just a side-effect of evolution in the presence of M-noise, not an adaptive response to a variable environment. A first attempt to distinguish these possibilities can be made by looking at the fitness records of a single run of the compulsory model with 2% M-noise. Figure 3.3(a) scatter-plots the fitness and memory of every individual in the population every 200 generations during the run of 10,000 generations. At this gross level, no relation is discernible. Figure 3.3(b) represents the same data in a different way, showing only those generations for which there is a non-zero diversity in memory across the population, normalising the recorded memory range in such cases to  $[0.0,1.0]$ , and plotting this data against ranked and normalised fitness values. Once again, however, there is no discernible relation, and therefore no evidence for longer memories providing any reliable selective advantage over shorter memories.

An even finer analysis is provided in figure 3.4, which relates fitness and memory in four different ways. Figure 3.4(a) superimposes the fitness of the fittest individual over the mean fitness (across the whole population). Figure 3.4(b) superimposes the memory of the fittest individual over the mean memory. Figure 3.4(c) plots the difference between the highest fitness and the

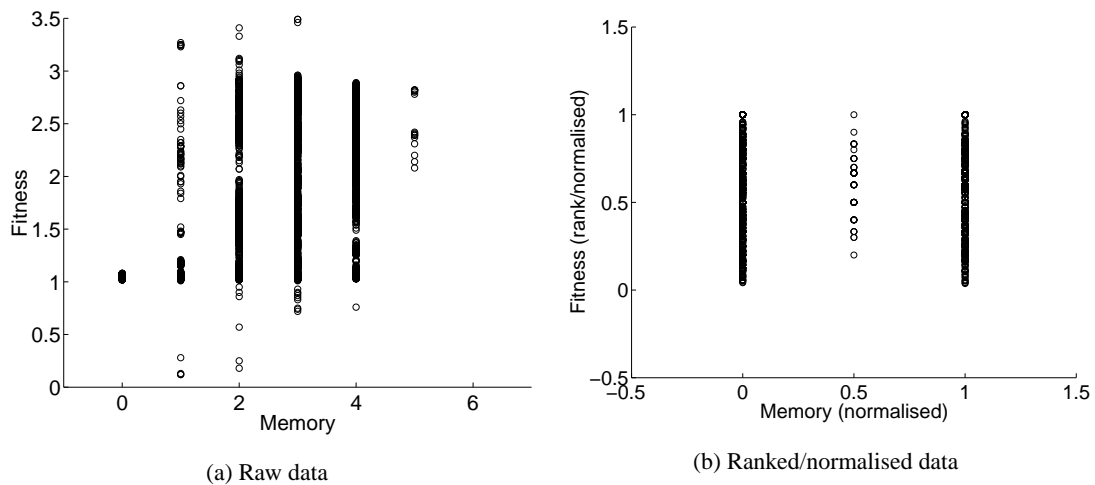


Figure 3.3: Fitness and memory in the compulsory model (1): 500 population snapshots are scatter-plotted using data recorded every 200 generations from a 10,000 generation evolutionary run with 2%  $M$ -noise. Plot (a) presents the raw data, plot (b) uses ranked fitnesses normalised to the range  $[0.0, 1.0]$  (tied fitnesses are given equal rankings), memory data is also normalised to the range  $[0.0, 1.0]$ . In (b) all ranking and normalisation is performed per-snapshot, and only those snapshots for which there is a non-zero diversity in memory are shown.

mean fitness, and, finally, figure 3.4(d) plots the difference between the memory of the fittest individual and the mean memory. The first thing to notice is that overall fitness is low and the fitness advantage of the fittest individual is only rarely apparent (a mean fitness of 1.0 represents mutual defection, and a mean fitness of 3.0 represents mutual cooperation). The second observation is that the relationship between fitness and memory is again tenuous at best. Although there is some degree of correlation between mean fitness and mean memory (there is a correlation coefficient of 0.52 across the entire run) and between the fitness and the memory of the fittest individual (0.48), fitness *differences* between fittest and mean do *not* correlate with differences between mean and fittest memory levels (-0.134). This is to say that during those generations where the fittest individual is noticeably better than average, the memory of this individual is not consistently above average. Indeed, across the entire evolutionary run, the fittest member is just as likely to have a below average memory as it is to have an above average memory.

For the compulsory model, therefore, there is no evidence from looking at fitnesses that evolved memory has evolved as an adaptation to environmental variability. However, fitness is not necessarily the most reliable or informative assay of the properties of an IPD system. This is because the fitness landscape involved can often be extremely rugged and rapidly changing, and mean fitness levels themselves are often considerably variable; in a population undergoing a rapid change from high mean fitness to low mean fitness it may not be appropriate to draw any conclusions from observing either the mean population memory, or the memory of the ‘fittest’ individual, at these times. It may be that the benefits of a longer memory in dealing with environmental variability show themselves only rarely and unpredictably. In other words, the large-scale population dynamics of an IPD ecology may well disguise any direct selective advantages of longer memories in dealing with environmental variability, without implying that these selective advan-




Figure 3.4: Fitness and memory in the compulsory model (2). Four representations of a single run of the compulsory model with 2% M-noise. Plot (a) shows mean and highest fitness, (b) shows mean memory and memory of the fittest individual, and (c) and (d) show the difference, at each generation, between best and mean fitness, and between the memory of the fittest individual and mean memory.

tages do not exist.

A different kind of assay is suggested by Ashby's law of requisite variety (LRV), which is briefly revised here. Consider a set of possible environmental disturbances,  $D$ , a set of possible responses on the part of the agent,  $R$ , and a set of possible outcomes,  $O$ . Consider also that for each  $D_i$ , there is a distinct outcome  $O_i$ , and a particular response  $R_i$ . There is also the 'system'  $S$ , which transforms a given  $(D_i, R_i)$  into the appropriate  $O_i$ . Stability in the essential variables requires minimising the variation in  $O$ , and this then requires that the variety in  $D$  is matched by the variety in  $R$ .

It is possible to think of an IPD model in terms of the components  $D$ ,  $R$ ,  $O$ , and  $S$ .  $D$  would be the set of all moves made over all iterations of the game, by all the opponents (of a given agent).  $R$  would be the strategy of the agent,  $S$  would be the IPD payoff table (table 3.1), and  $O$  would be the set of scores awarded to the agent. In a stable, cooperating population, the overall fitness of each agent (over many generations) will be maximised with minimum variation in  $O$  away from repeated cooperation payoffs (since mutual cooperation, in the IPD, maximises overall payoff in the population). Any environmental variation away from stable cooperation can then be expected to have a deleterious effect on fitness, unless it is countered by a strategy that can effectively cope with such variation. For example, an occasional 'accidental' defection will throw a population of 'tit-for-tat' players into continual mutual defection but a 'tit-for-two-tats' population will 'absorb' such a defection, permitting general cooperative behaviour to persist throughout the population.<sup>13</sup> In terms of the LRV, any variety in  $D$  can only be prevented from affecting  $O$  if it is countered through the action of  $R$  on  $S$ ; that is, if the variety in  $R$  (given  $S$ ) matches that of  $D$ . This con-

<sup>13</sup>See footnote 11 for a description of tit-for-tat. Tit-for-two-tats is a simple extension; the response of a player is only switched if the other player makes the opposite move twice in succession.




Figure 3.5: Compulsory model with 2% O-noise. Results from 12 evolutionary runs are superimposed.




Figure 3.6: Average and standard deviation cooperation in the compulsory model (statistics generated from 12 evolutionary runs).

strual of the IPD allows us to predict that noise (variety) on  $D$  may lead to the evolution of more complex strategies (supporting variety in  $R$ ) to counter this disturbance, but that noise on  $O$  will not. Noise on the outcome, after all, is precisely what the strategies would be expected to prevent, and if the noise is applied directly to  $O$ , then, by definition, *no* strategy can provide an effective response. More formally, a second type of noise can now be distinguished, this being O-noise (for ‘outcome’ noise): on each iteration, and for each player, there is a certain probability that the payoff awarded is altered. Each time this occurs, the actual payoff awarded is selected randomly from the four possible payoff values. The specific hypothesis afforded by this interpretation of the IPD is therefore that O-noise should *not* lead to the evolution of more complex strategies.

Unfortunately, figure 3.5 quite clearly illustrates that this hypothesis is not borne out. There is just as much evolved memory with 2% O-noise as there is with 2% M-noise. But perhaps this is not surprising. Earlier in this section it was noted that the average fitness levels (and hence average cooperation levels), for a single run of the compulsory model, were predominantly low. Indeed, this holds for all runs of the compulsory model in both zero noise and M-noise conditions (see figure 3.6; I will return to the noticeably higher cooperation levels under 2% O-noise in the following section). If, then, there is no stable cooperating population to begin with, there is no reason (with respect to the LRV) to expect complex strategies to evolve to cope with environmental variability. Again, then, it must be said that there is no evidence that the evolution of complex strategies in the compulsory model is an adaptation to environmental variability.

One could carry on indefinitely attempting to search for evidence that the evolved complexity in the compulsory model is adaptively significant. The final study in this section, by contrast, adopts the opposite approach by exploring a compulsory model in which, although the heritable complexity of the strategies varies as before, a *functional* memory limit of 1 is imposed during every game. In this model, then, there is no possible functional difference, with regard to the IPD, between a strategy of memory 1 and a strategy of memory 6. Nevertheless, as figure 3.7 illustrates, both 2% M-noise and 2% O-noise conditions result in the evolution of complexity as before. This is strong evidence for the position that the evolved memory in the compulsory model with 2% M-noise (or 2% O-noise) is *not* adaptively involved in responding to environmental variability.




Figure 3.7: Compulsory model with functional memory limit of 1. Results from 12 evolutionary runs are superimposed.

At this stage it therefore becomes worth asking whether another explanation can be found for the evolution of complex strategies in *either* O-noise or M-noise conditions?

### 3.3.3 The facilitation of genetic mobility

This section evaluates the hypothesis that the evolution of complexity observed with both M-noise and O-noise in the compulsory IPD model is an example of *enhanced genetic mobility* rather than an adaptationist response to environmental variability.

A compulsory IPD model is considered in which the memory of the strategies is *genetically* limited to being either 0 or 1 (the average population memory could then range continuously from 0.0 to 1.0), with the motivation of investigating differences between the three noise conditions that are relatively independent of the evolution of complex strategies.<sup>14</sup> In other words, to try to distinguish between the causes of the evolved complexity, and a mixture of the causes and consequences. If complex strategies are given free reign to evolve, the causes and consequences of their evolution cannot be rigorously distinguished. The following experiments allow such a distinction to be drawn - if not rigorously, then at least in outline. Twelve evolutionary runs were performed in each of the three conditions; zero noise, 2% O-noise, and 2% M-noise. Figure 3.8 illustrates that in both the O-noise and M-noise conditions, the average evolved memory is near to the expected average (0.5). But in the zero noise condition, in each case the average evolved memory (usually) either stays near the initial level of 1.0, or is very close to 0.0. This indicates that both O-noise and M-noise are permitting evolution to wander easily within the prescribed area of genotype space (determined by the maximum memory limit of 1), but with zero noise, the evolving population often becomes ‘stuck’ on either memory 1 or memory 0. Furthermore, the means of these average memories are not significantly different between any of the conditions. The average (of the average memory) in the M-noise condition is 0.26; in the O-noise, 0.34; and with zero noise, 0.41, and pairwise t-tests between all sets of means all give  $p$  values in excess of 0.05, indicating that neither type of noise is actually encouraging (or discouraging) the evolution

<sup>14</sup>Of course, this genetic limitation incurs a corresponding functional limitation, as in the previous experiment set.




Figure 3.8: Compulsory model; memory genetically limited to the range  $[0,1]$ . Both M-noise and O-noise permit average memory to approach the expected average; without noise, the average memory appears to be ‘stuck’ at one extreme or the other. A - zero noise, B - 2% M-noise, C - 2% O-noise.




Figure 3.9: Average and standard deviation cooperation in the genetically limited model (statistics generated from 12 evolutionary runs).

of complexity *per se* in this model.

Why might noise have this effect? At this point the investigation becomes speculative. One possibility is via an influence on the cooperation dynamics of the model: given a predominantly defecting population, both M-noise and O-noise might permit a cooperative move to occasionally score well, allowing the evolving system to be more tolerant of novel strategies. This intuition was tested in 12 runs of each condition of the model above (in which the maximum memory was genetically limited to 1 move) figure 3.9 presents the results. The pattern is clear, and strongly reminiscent of that obtained from the unmodified compulsory model (figure 3.6); only O-noise leads to any noticeable increase in average cooperation. Although clear, however, this observation does not permit any firm conclusions to be drawn. Perhaps it is the case that O-noise is more effective than M-noise at allowing occasional cooperative moves to score well. However, even if this can be shown to be the case, the enhanced evolution of complexity under M-noise would remain unaccounted for. Although resolution of these questions would be desirable, they remain tangential to thrust of this chapter, and shall therefore, for the present, be left open. The summary of this section that can be offered is therefore (i) that noise *does* facilitate genetic mobility, at least in the dimension reflected by strategy memory in the present model, and (ii) that this facilitation *may* be mediated by influences on cooperation dynamics - at least in the case of O-noise - but that the nature of any such influences remains to be established and clarified.

### 3.3.4 IPD/CR models

This chapter opened with the promise of empirical support for the hypothesis that complex strategies can evolve as an adaptation to environmental variability. By now, the reader may have begun to worry that this promise will remain unfulfilled. Here we attempt to assuage such worries with an exploration of a different kind of IPD model in which average cooperation - and hence fitness - is maintained at a much higher level, providing a potentially more suitable substrate for investigating

the evolution of *adaptive* complexity.

In the compulsory IPD model, every individual in the population was required to interact with every other member, every generation. IPD models can also be constructed in which individuals can exercise a degree of choice over which other individuals they interact with. Stanley, Ashlock, and Tesfatsion (1994) investigate the formation of ‘social networks’ in an IPD model in which individuals choose and refuse game partners with reference to continuously updated expected payoffs that each individual maintains for every other in the population.<sup>15</sup> They demonstrate that cooperation is evolved rapidly under these conditions, and they discuss the emergence of a variety of metastable networks displaying distinct patterns of cooperativity. The partner choice algorithm used in their model - which they call an IPD/CR model - can be described as follows. Each individual in the population maintains a set of ‘expectation values’  $E_i$  for every member in the population (indexed by  $i$ ), bar itself. These values reflect the expected outcome of a prisoner’s dilemma interaction, and each  $E_i$  is initially set to the mutual cooperation payoff. For each subsequent iteration, each individual signals its willingness to play to  $n$  of its most preferable partners (determined by the values of  $E_i$  for all  $i$ ;  $n$  is a constant); the recipients thus form lists of those individuals who have proposed to them. Every individual now has a list of potential partners, and engages in IPD interactions with those partners for whom  $E_i \geq \omega$ , where  $\omega$  is a fixed threshold value common to all members of the population; both participants are awarded payoffs according to table 3.1. If, however,  $E_i < \omega$ , the *refusee* is awarded a payoff of value 1.0 (the same that would have been awarded for mutual defection). After each interaction or refusal, values of  $E_i$  are updated according to:

$$E_i[t + 1] = \gamma E_i[t] + (1 - \gamma)p[t]$$

where  $\gamma$  is a (fixed) learning rate, and  $p[t]$  is the payoff awarded to that individual during the IPD interaction or refusal event at time  $t$  (note that only the refusee is updated after a refusal). After all accepted interactions have occurred, individuals who have not interacted at all are awarded a ‘wallflower’ payoff of value  $\omega$ . In this section, an IPD/CR model is developed as an extension of the compulsory model, using this algorithm, with the hope that such a model may provide a more stable background of cooperativity from which the evolution of complexity can be explored. I therefore begin as before, by testing the hypothesis that evolution in the presence of M-noise will lead to the emergence of complex strategies, but that evolution in the absence of noise will not. (See appendix A for full specification of the IPD/CR model parameters employed in this investigation.)

### 3.3.5 IPD/CR results

Just as with the compulsory model, 12 evolutionary runs were conducted, of 10,000 generations each, in both zero noise and 2% M-noise conditions. Informal initial observations suggested that the population rapidly reached maximum complexity in both conditions, and so, in order to differentiate the conditions, a small *cost* on complexity was applied - in all conditions - by levying a small fitness penalty on genotype length (0.0025 of total fitness per memory ‘unit’). The results from this model are clear. With zero noise complex strategies do not often arise, but with 2%

<sup>15</sup>See also Smucker, Stanley, and Ashlock (1996), Ashlock, Smucker, Stanley, and Tesfatsion (1995).






Figure 3.10: Evolution of complexity in the IPD/CR model. Results from 12 evolutionary runs are superimposed.

M-noise they do (figure 3.10). Again, t-tests reveal these observations to be highly significant ( $p < 0.01$ ).

Now comes the question of whether this evolved complexity can be interpreted as an adaptation to environmental variability. Again, the first task is to record and analyse fitnesses for a single run with 2% M-noise.

Figure 3.11(a) scatter-plots the fitness and memory of every individual in the population every 200 generations during the run of 10,000 generations. Figure 3.11(b) shows only those generations for which there is a non-zero diversity in memory across the population, normalising the recorded memory range in such cases to  $[0.0, 1.0]$ , and plotting this data against ranked and normalised fitness values. Just as with the compulsory model, no relation is discernible at these gross levels; there is no evidence here for longer memories providing any selective advantage.

Figure 3.12 gives the finer analysis, with figure 3.12(a) superimposing the fitness of the fittest individual on the average fitness (across the whole population), (b) superimposing the memory of the fittest individual on the average memory, and (c) and (d) plotting the difference between the highest fitness and the average fitness and the difference between the memory of the fittest individual and the average memory. There are some differences between the compulsory and IPD/CR models at this relatively fine-grained level of analysis. Most notably, fitness is now consistently high, and there is a reliable - if small - difference between best and mean fitness throughout the run. The correlation between mean fitness and mean memory is now negative (-21.6) and there is very little relation between the fitness and memory of the fittest individual (-0.06). The correlation between fitness difference (between mean and fittest) and memory difference (between mean and fittest) is also very small, however, it is at least positive (0.08).

What can be gleaned from these figures? As argued for the compulsory model, fitness *per se* may not be the best assay of an IPD system. Correlating mean values with mean values, and best values with best values (of fitness and memory) will not be particularly informative, because in a stable cooperative population (containing, it follows, well adapted strategies) there should be minimal variation in fitness. Potentially more informative is the final correlation measure between

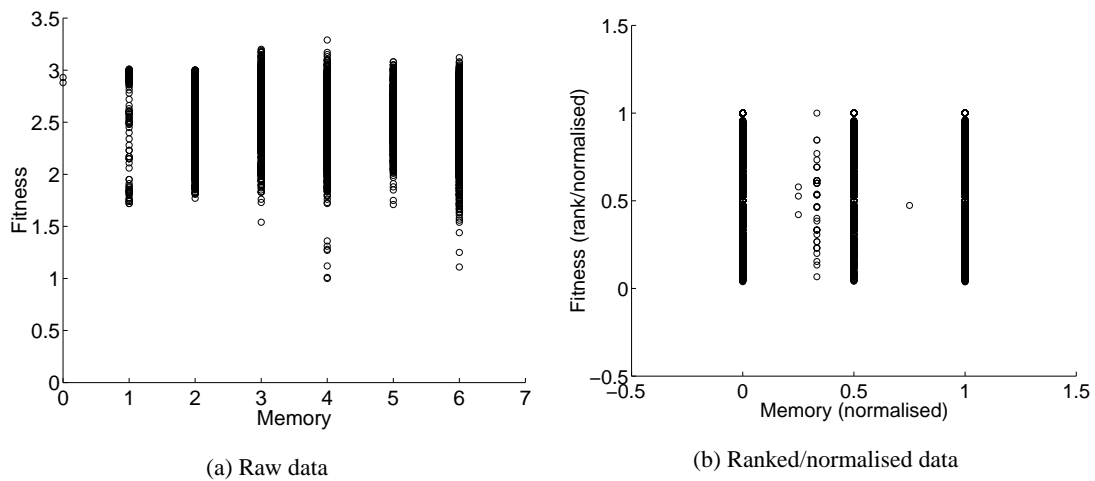


Figure 3.11: Fitness and memory in the IPD/CR model (1): 500 population snapshots are scatter-plotted using data recorded every 200 generations from a 10,000 generation evolutionary run with 2% M-noise. Plot (a) presents the raw data, plot (b) uses ranked fitnesses normalised to the range  $[0.0, 1.0]$  (tied fitnesses are given equal rankings), memory data is also normalised to the range  $[0.0, 1.0]$ . All ranking and normalisation is performed per-snapshot, and, also for plot (b), only those snapshots for which there is a non-zero diversity in memory are shown.

fitness differences and memory differences. A strong correlation here would support the idea that whenever the fittest individual is considerably better than average, it possesses a memory considerably longer than average. Although the correlation presented by this model is positive (unlike that observed for the compulsory model), it can hardly be described as strong.

Once again, then, fitness measurements fail to provide much evidence for the adaptive evolution of complexity. However, as argued above and in section 3.3.2, a lack of evidence of this kind does not demonstrate that the evolution of complexity is not adaptive. The volatility of IPD systems with respect to fitness may disguise the direct selective advantages afforded by memory in dealing with environmental variability, without implying that these selective advantages do not exist.

A different tack is therefore required, and the high levels of cooperation observed in the IPD/CR model (see figure 3.14) suggests that the analysis inspired by Ashby's LRV might be more appropriate than it was for the compulsory model. The hypothesis relating to this analysis is, of course, that if the evolution of complexity is adaptively significant, it can be expected to be promoted by M-noise but *not* by O-noise.

This hypothesis was tested using the IPD/CR model, and, in stark contrast to the results from the compulsory model, swapping 2% M-noise for 2% O-noise *does* completely abolish the evolution of complexity (figure 3.13a). This is strong evidence for the position that the evolved complexity in the M-noise condition is adaptive in some way. Further empirical support for this position comes from a repeat of the functional memory limit study first explored in section 3.3.2. There, it was found that for the compulsory model with 2% M-noise, the significant degree of evolved complexity was *not* diminished by rendering this complexity functionally impotent (through imposing an effective memory limit of 1 during each IPD interaction). However, for the IPD/CR model




Figure 3.12: Fitness and memory in the IPD/CR model (2). Four representations of a single run of the IPD/CR model. Plot (a) shows mean and highest fitness, (b) shows mean memory and memory of the fittest individual, and (c) and (d) show the difference, at each generation, between best and mean fitness, and between memory of the fittest individual and mean memory.




Figure 3.13: 2% O-noise (a) and functional memory limit (b) in the IPD/CR model (effective - but not heritable - memory limited to the range  $[0,1]$ ). 12 evolutionary runs superimposed in each case.




Figure 3.14: Average and standard deviation cooperation in the IPD/CR model (statistics generated from 12 evolutionary runs).

with 2% M-noise, such a functional limitation *does* entirely abolish the evolution of complexity (figure 3.13b). This strongly supports the position that the evolution of complexity observed in the IPD/CR model *is* functionally significant.

As a final observation, figure 3.14 illustrates quite clearly that, in contrast to the compulsory model, cooperation levels in the IPD/CR model are uniformly high. The slightly lower (but still high) average cooperation under M-noise may be understood to result from the evolved strategies failing to cope perfectly with the environmental variability. Nevertheless, the levels of cooperation are high enough such that an application of the LRV framework makes intuitive sense.

### 3.3.6 Discussion

The results described so far lend themselves to two observations in particular. The first is that both the compulsory model and the IPD/CR model provide clear examples of environmental variability engendering the (artificial) evolution of complexity. However, only in the IPD/CR model can this evolved complexity be appropriately interpreted as a *response* to environmental variability, an interpretation justified by (i) the abolition of the evolution of complexity if M-noise is replaced by O-noise, or under the application of a functional memory limit, and (ii) the evolution of complex strategies despite a fitness cost on complexity. However, it still cannot be said that the complexity of the strategies directly ‘matches’ the variability of the environment in the sense of the LRV, partly because the environmental variability cannot be quantified independently of the agents themselves, and partly because - as illustrated in figure 3.14 - the complex strategies do not succeed in fully overcoming the effects of noise on cooperation. At most, in the present case, the LRV provides a means of formalising the dynamics of the IPD, and of understanding the distinction between M-noise and O-noise and the nature of their distinct effects.

With respect to the distinction between behaviour and mechanism, the subject of much discussion in chapter 2, at this stage there is not much to be said. As noted in section 3.3, strategy complexity is mechanistic; it measures the potential for behavioural variability. Its relation to behavioural complexity has not been at issue here since adaptive behaviour in the IPD consists (usually) of mutual cooperation, and continuous mutual cooperation - like (maladaptive) continuous mutual defection - is not complex at all.

The second significant observation is that both M-noise and O-noise facilitate (non-functional)

genetic mobility in the compulsory model, at least in the dimension of genotype length. Although the evidence for the existence of this non-functional effect is strong, the evidence connecting it to the dynamics of population cooperation is only circumstantial. A more general discussion of the various roles of noise in artificial evolution will be picked up again in the following chapter. For the present, of greatest importance is the empirical support that has been lent to the externalist project outlined chapter 2 with regard to complexity as an adaptive response to environmental variability.

### 3.4 ‘Win-stay, lose-shift’ strategies

This section explores a re-implementation of a model by Martin Posch (1999), recently published in the *Journal of Theoretical Biology*. Using this model, Posch makes the claim - a claim consonant with the models of the previous section - that noisy environments can promote the evolution of long memories in certain types of strategy for playing the IPD. However, whilst one may wish to concur with his claim, the means by which he substantiates it are - it will be argued - unsound. In what follows his work is briefly critiqued and then a revised version is presented in which the original conclusion is reaffirmed.

Posch is interested in the performance of ‘win-stay, lose-shift’ (WSLS) strategies. These strategies repeatedly perform a ‘current’ action, but only whilst this action is considered ‘successful’. One such strategy for playing the IPD is given by the vector  $(i, a, p)$ , where  $i \in \{C, D\}$  gives the initial action,  $a \in [-1.0, 6.0]$  is the ‘aspiration level’, and  $p \in [0.0, 1.0]$  is the ‘update probability’. The strategy operates as follows. First, set the current action to  $i$ . Then, for every interaction, perform the current action, and with probability  $p$  update the current action by calculating the average payoff received since last update; if this average is below  $a$  then switch the current action. This strategy can be called ‘fixed’ since the aspiration rate  $a$  does not change. Posch also discusses a second strategy - an ‘adaptive’ strategy - which is given by the vector  $(i, a_0, p, \gamma)$ , in which  $a_0 \in [-1.0, 6.0]$  is the *initial* aspiration level, and  $\gamma \in [0.0, 1.0]$  determines the rate of adaptation of the aspiration level according to:

$$a[t + 1] = \gamma P[t] + (1 - \gamma)a[t],$$

which is calculated at every time-step  $t$ , where  $P[t]$  is the payoff received in round  $t$ .

After analytically exploring properties of both of these strategies in deterministic IPD models, Posch uses a GA to evolve the components of each strategy in non-deterministic, noisy models. In these models, noise is implemented as M-noise, and populations of 40 agents engage in round-robin IPD tournaments each generation. The number of IPD rounds in each contest is random; after each round the contest is terminated with a probability of 0.01. Genotypes encode all parameters of the strategies, with initial random values chosen from the ranges above (further details of the GA are given in appendix A). Posch investigated M-noise levels of 0%, 1%, and 4%, and figure 3.15 illustrates a set of results derived from a faithful replication of this model. These results are identical to those obtained by Posch himself.

Posch concludes from these results - especially figure 3.15(c) - that increasing levels of noise promote the evolution of longer memories for the *fixed* aspiration strategies, with long memories interpreted as low values of  $p$ . However, it is clear from figure 3.15(c) that although  $p$  significantly

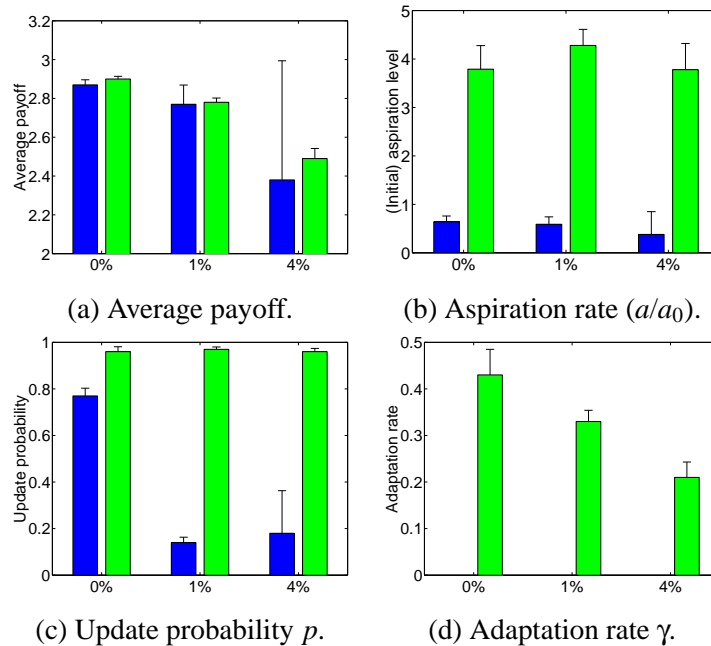


Figure 3.15: Replication of Posch's (1999) WSLS model. Each column represents average (with error) over 10 evolutionary runs, the dark columns represent the fixed strategies, and the light columns the adaptive strategies. For each set of 10 runs 4 statistics were measured; average payoff, initial aspiration rate  $a$  (fixed) or  $a_0$  (adaptive), update probability  $p$ , and adaptation rate  $\gamma$ . These statistics were measured at each of three M-noise levels; 0%, 1%, and 4%.

decreases when 1% M-noise is introduced, 4% M-noise leads, if anything, to a slight *increase* in  $p$ , although the relatively poor average payoffs awarded in this condition may lessen the significance of this observation. Furthermore, the opposite pattern is observed for the *adaptive* aspiration strategies, with (very) slight increases in  $p$  accompanying each increase in the noise level. Although Posch does not misrepresent this pattern, neither does he discuss it, an omission that is troubling given his emphasis on the relatively strong performance of adaptive strategies - both in his *Journal of Theoretical Biology* paper and in other work (Posch, Pichler, & Sigmund, 1999) - and given the generality of his overall conclusion: "I discuss the impact of noise and show that the memory length of the players increases with noise levels" (1999, p.1). This overall conclusion must be considered tenuous on the basis of his interpretation of the evidence as replicated in figure 3.15.

Notice, however, that in figure 3.15(d) there is a steady decrease in the rate of adaptation as noise levels increase. Notice also that there is just as much justification for interpreting low values of  $\gamma$  in terms of long memory as there is for understanding  $p$  in this way. A low  $\gamma$  means that changes in aspiration levels depend to a relatively greater extent on previous outcomes, in just the same way that a low  $p$  means that the decision to switch depends on the average payoff over many previous outcomes. Therefore, in Posch's study,  $p$  and  $\gamma$  are conflated as measures of memory.<sup>16</sup> Posch, however, does not discuss this and in particular does not pick up on the argument that the pattern of results in figure 3.15(d) may in fact illustrate the conclusion that he is seeking - the evolution of longer memories in noisy conditions - but with memory reflected by  $\gamma$  rather than by

<sup>16</sup>In a related study, for example, Thuisjman et al. (1995) model bee foraging decisions using a variable aspiration rate decision rule, and explicitly interpret memory in terms of the adaptation rate, not update probability. This study - a study actually cited in Posch (1999) - will be revisited in chapter 9.

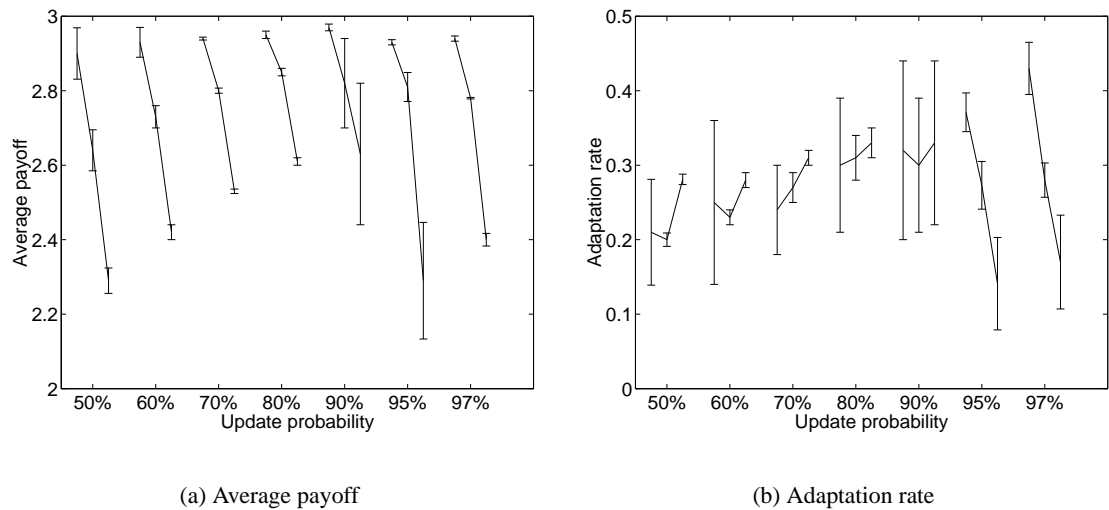


Figure 3.16: Influence of M-noise on (a) average payoff and (b)  $\gamma$ , with fixed values of update probability  $p$  (along the abscissa). Each group of 3 data points gives average and standard deviations for M-noise levels of 0%, 1%, and 4% (from left to right in each group), for a range of values of  $p$  from 50% to 97%, over 10 runs of 1000 generations each.

$p$ . An explicit test of this argument would require a modification of the model in which  $p$  is kept fixed whilst  $\gamma$  is allowed to vary over the course of evolution, the hypothesis being that increasing noise levels will lead to lower adaptation rates. Figure 3.16 illustrates a set of results from such a model with M-noise levels of 0%, 1%, and 4%, measuring both average payoff and  $\gamma$ , for a range of fixed  $p$  values. The most obvious pattern, in figure 3.16(a), is the degradation of strategy performance (in terms of average payoff) with increasing noise, for all values of  $p$ . The pattern of results in figure 3.16(b) is less clear. The hypothesis set out above, that increasing noise will lead to lower adaptation rates, is only supported for high  $p$  values (95% and above). For lower  $p$  values there is no discernible relationship between the level of noise and  $\gamma$ , the only reliable pattern being a general decline in  $\gamma$  as the fixed level of  $p$  also declines. However these results make sense if it is remembered that both  $p$  and  $\gamma$  are measures of memory. Low  $p$  values imply that the strategy already possesses significant memory, and so one may *not* expect to see the clear effects of noise on  $\gamma$  that are observed for high values of  $p$ . However, one might well expect to see a trade-off between  $\gamma$  and  $p$  as measures of memory, and indeed this is exactly what is observed - lower (fixed) values of  $p$  entail lower overall levels of  $\gamma$ . And recall that in the replication of the original model, in which both  $p$  and  $\gamma$  were allowed to vary over the course of evolution, the strong relationship between noise and  $\gamma$  (figure 3.15d) accompanied very high values of  $p$  (figure 3.15c).

To briefly summarise this section. Increasing noise levels *do* promote the evolution of WLSL strategies with longer memories, but this can only be established convincingly if it is appreciated that both the update probability  $p$  and the adaptation rate  $\gamma$  constitute measures of strategy memory.

### 3.4.1 WLSL strategies and the ECT

I said at the beginning of this section that Posch's claim was consonant with the models of the previous sections, but how far does this consonance go? Certainly, both kinds of model demon-

strate the evolution of longer memories in the presence of environmental variability. However, the fundamental hypothesis as set out at the beginning of this chapter - a hypothesis derived from Godfrey-Smith's ECT - had to do with internal complexity, not memory. And whereas in the first model set, the variable-length genotype mechanism permitted an intuitive association to be made between memory and mechanism complexity, for the WSLS models this association is not so clear; all that is changing is the value of a parameter ( $p$  and/or  $\gamma$ ), and there is no simply no scope for the kind of mechanistic change evident in the earlier models. Also, recall for a moment the concept of *Umwelt* from the previous chapter, a rough definition of which is the 'sensorimotor space of an agent'. Longer memories in the variable-length genotype models lead to potentially more variable *Umwelts*, but longer memories in the WSLS models lead, by averaging over time, to potentially *less* variable *Umwelts*, regardless of the actual variability in the moves expressed in the IPD system that may be recorded by an external observer (the 'external environment'). So, if there is any sensible association of longer memories with internal complexity in the WSLS models, it therefore has to do *not* with any heterogeneity in the description of the mechanism *per se* but rather to do with the construction of a relatively simple *Umwelt* out of a potentially variable external environment; the averaging process itself. As argued in the previous chapter, this kind of 'constitutive construction' can fall within the explanatory domain of the ECT, even if, in the present case, it does so at a relatively abstract level.

With this qualification in mind, the conclusions of this section do indeed accord with the primary conclusions of the first set of experiments in this chapter, thereby providing some evidence for their generality. Indeed, the WSLS strategy itself is an example of a very general class of learning rule that can be applied to many different situations. A version of this strategy will be explored in chapter 9, and it will be seen that the same relation between noise and strategy memory obtains in the novel context of a 'foraging game'.

### 3.5 Summary

This chapter has been about the hypothesis that environmental variability can lead to the evolution of adaptive complexity. The exploration of this hypothesis involved both its isolation from the larger debate surrounding complexity and its evolution, and also the construction of concrete game-theoretic models. The first task was achieved primarily by distinguishing the 'causal' question from the 'teleological' question, and the second by using IPD to formalise and quantify the hypothesis in the context of artificial evolution. Strong support for the hypothesis was provided by both IPD models, but it is important not to overplay the generality of the conclusions drawn; complexity and evolution are simply issues too large to be resolved into any single model or class of models. The results described in this chapter should be understood as an invitation to further investigation, not as the final word on the subject.

A novel feature of this investigation has been the use of Ashby's LRV to interpret the dynamics of the compulsory model and the IPD/CR model. This interpretation, motivated initially by the opacity of fitness statistics in these models, brought to light an additional influence of environmental variability in enhancing the mobility of the GA search process through genotype space. A detailed discussion of the role of noise in artificial evolutionary processes will be presented in the following chapter, a chapter which will also reintroduce the importance of the distinction between



behaviour and mechanism through the development of an evolutionary robotics model in which these terms have intuitive application.

## Chapter 4

### Evolving complexity II: Evolutionary robotics

---

This chapter continues the empirical exploration of the hypothesis of the previous chapter - that environmental variability promotes the evolution of internal complexity - but in the context of evolutionary robotics.<sup>1</sup> It is demonstrated in this context that environmental variability leads to the evolution of complex behaviour, in a ‘homing-navigation’ task, and that this complex behaviour is subserved by a mechanistic strategy of multimodal sensorimotor integration. The relations between behaviour and mechanism are explored in detail in this chapter, and provide a clear example of the methodological utility of artificial evolution models: the researcher need only specify functional fitness criteria, leaving the evolutionary process itself to come up with a suitable mechanism (given some initial pre-specification).

The chapter ends with a discussion of the various roles of noise in artificial evolution, drawing on both the iterated prisoner’s dilemma studies of the previous chapter and the present evolutionary robotic explorations. Three novel influences are identified - the evolution of behavioural/mechanistic complexity, the acceleration of evolutionary search, and the facilitation of genetic mobility - all of which are drawn together in a discussion of the dynamics of evolutionary search.

#### 4.1 A study of homing navigation

The empirical context for this chapter is drawn from Floreano and Mondada (1996). In this study, a ‘homing-navigation’ behaviour is evolved for a Khepera mobile robot (K-Team, 1993), controlled by a simple neural network, in which the robot explores a walled arena and periodically returns to a designated charging area marked by a light source. Floreano and Mondada stress the importance of performing artificial evolution in the ‘real world’ by evaluating candidate controllers on real robots, rather than in simulation.<sup>2</sup> They also emphasise the plausibility of evolving relatively complex behaviours using only simple fitness functions. The present focus is different: the ‘homing-navigation’ model of the present chapter is directed towards assessing the influence of

---

<sup>1</sup>For comprehensive introductions to the discipline of evolutionary robotics see Nolfi and Floreano (2000), Meyer, Husbands, and Harvey (1998), Husbands, Harvey, Cliff, and Miller (1997).

<sup>2</sup>The term ‘real world’ inhabits scare quotes because there are many differences between the ‘real world’ of a carefully constructed experimental arena and the ‘real world’ experienced by humans, other animals, and autonomous vacuum cleaners.




Figure 4.1: The Khepera mobile robot.

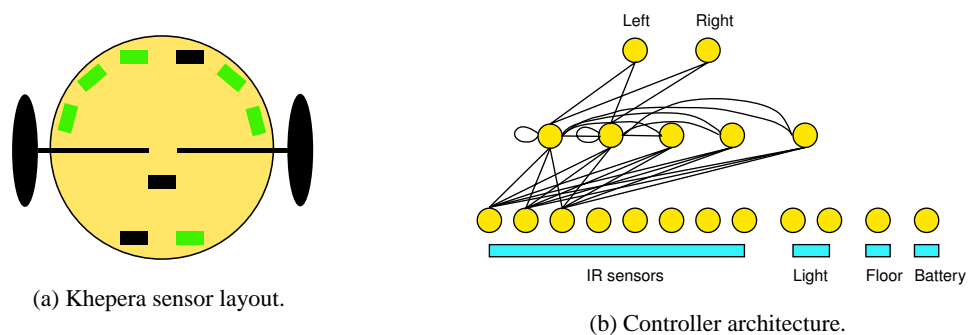


Figure 4.2: Sensor layout and network architecture. The lightly shaded rectangles in (a) represent IR sensors, and the darkly shaded rectangles represent sensors from which both IR and ambient light sense data is taken (apart from the ‘floor sensor’ in the middle, which transmits only ambient light sense data). The network architecture in (b) illustrates inputs from the 8 IR sensors, the 2 ambient light sensors, and the floor sensor and battery sensor. For clarity, not all network connections are shown; all input units connect to all hidden units, and all hidden units connect to all hidden units (including self-connection) and to both motor units.

environmental variability on the complexity of the evolved behaviours and their underlying mechanisms (in this chapter the term ‘environment’ is taken to cover everything external to the control structure of the robot). However, the use of a simple fitness function continues to be important in the present work, in permitting evolution a relatively free hand in designing mechanisms and behaviours.

#### 4.1.1 A description of the model

The model of the present chapter differs from that of Floreano and Mondada only in its use of simulation instead of the ‘real world’ for evaluating controllers (although, as we will see later, controllers evolved in simulation did transfer effectively to the ‘real world’). The significance of the use of simulation will be discussed later, following a description of the robot-environment system itself. The standard Khepera mobile research robot (figure 4.1) is equipped with 8 infra-red (IR) proximity sensors each of which is also capable of independently detecting ambient light levels, although only 2 are actually used for this purpose (figure 4.2a). Floreano and Mondada augment this capability with an additional ‘floor sensor’ attached to the undercarriage of the robot, thresholded so to be able to detect differences in floor reflectivity. The environment consists of a

40cm by 45cm walled arena, situated in a dark room, but with a small light tower placed in one corner. This corner (denoted the ‘charging area’) is painted with a poorly reflective black colour on the floor out to a radius of 8cm; the remainder of the floor is painted a highly reflective white. The robot possesses a simulated battery enabling it to ‘survive’ for a maximum of 50 time-steps, with each time-step corresponding to a single update of the controlling neural network. If, however, the robot passes over the charging area at any time, the battery is instantaneously recharged, enabling the robot to continue for a further 50 time-steps up to an arbitrary maximum of 150 time-steps. The arrangement of sensors and motors on the robot and the (fixed) architecture of the neural network controller are illustrated in figure 4.2. Sigmoid activation functions are employed at all layers except the input layer, which linearly scales the sensory inputs to the range [-0.5,0.5].

Floreano and Mondada use a GA to evolve the weights and thresholds of this network, under a very simple fitness function which is calculated incrementally at every time-step - except when the robot is directly over the charging area, during which time no score is awarded - according to the following formula:

$$\mathcal{F} = V(1.0 - i),$$

where  $V$  is absolute value of the average wheel speed linearly mapped from the range [-10.0,10.0] onto the range [0.0,1.0], and  $i$  represents the activity of the most active IR sensor, linearly mapped from the range [-0.5,0.5] also onto the range [0.0,1.0]. It is important to see that there is nothing in this fitness function that explicitly specifies periodic return to the charging area, it is maximised simply by high speed and low IR input. However, robots that come to adopt this strategy of ‘homing-navigation’ will tend to live longer and thus accrue higher fitness than those that do not. Such a simple fitness function allows evolution a relatively free hand in designing well adapted behaviours; in particular, the complexity of the potential solutions is not pre-specified. As Floreano and Mondada argue:

[C]onsider the fitness measure not as a detailed and complex function to be optimized in order to achieve a desired behavior, but rather as a general survival criterion that is automatically translated into a set of specific constraints by the characteristics of the interactions between the organism and the environment. (1996, p.9)

Floreano and Mondada performed artificial evolution in the ‘real world’, downloading candidate controllers onto real Khepera robots for each evaluation. In contrast, the present model develops a simulation model of the Khepera robot, controller, and environment. The motivation for this is threefold. First, to avoid the significant time costs incurred by ‘real world’ evolution (Floreano and Mondada required about ten days for evolutionary runs of about 200 generations, roughly equivalent to two hours of simulation time using a 143MHz Sun SparcUltra workstation). Second, the analysis of evolved behaviours and control systems is considerably easier in simulation than it is in the ‘real world’. Third, and most importantly, a simulation allows the degree of environmental variability to be carefully tuned.

The simulation model consists largely of a series of look-up tables that deliver appropriate values for the robot sensors and motors in any given situation (distance of the robot from a wall or corner, angle to wall or light source, orientation and speed of robot, and so on). Of central importance is the introduction of simulation *noise* (variability). Within each evaluation, noise

is applied to sensor input levels, robot position and orientation, wheel speeds, rate of change of orientation, and to the effects of collision with the edges of the arena. Following such collisions the robot is randomly repositioned within about 2-3cm of the wall, with a large and random orientation and speed change. Simulation noise is also applied *between* evaluations to the angle of acceptance of the ambient light sensors, the arena dimensions, the radius of the charging area, and to the background levels of IR and ambient light. In each case (apart from wall collisions) the application of noise involves modifying the value in question by a percentage drawn at random from a specific range (see appendix A). It is important to note that these loci of variability are all external to the control structure of the robot and therefore can be interpreted as aspects of *environmental* variability.

The liberal use of noise in simulation has been most forcefully advocated by Jakobi (see, for example, 1998a), who argues that the introduction of noise in the right places in a simulation model can facilitate the transference of behaviours evolved in simulation to the ‘real world’. Although a demonstration of transference is not an essential part of the argument of this chapter, the loci of variability listed above were nevertheless chosen in line with Jakobi’s methodology, partly to provide further evidence for the utility of his methodology, and partly because the ‘real world’ demonstration of evolved behaviour does help to alleviate any suspicion that the simulation model in question may be too abstract to be of broad interest.<sup>3</sup>

The experiment proceeded using a distributed GA, with a population of 100, to evolve the weights and thresholds for the network (this structure remains fixed for the duration of each individual), with evolution taking place in either high-noise or low-noise simulations. The weights and thresholds were specified as floating point numbers on a genotype of length 102 (see appendix A for details of noise levels and GA), with evolutionary runs of about 100 generations always producing very fit individuals. Successful transfer of evolved behaviours to the ‘real world’ was consistently observed when networks from the fittest robots (from high-noise simulations) were downloaded onto real Kheperas. Figure 4.3 presents a series of still images from a digitised video sequence of a real Khepera (powered externally, but with all processing on-board) engaged in evolved ‘homing-navigation’.<sup>4</sup>

#### 4.1.2 Behavioural analysis

For the purposes of evaluating the hypothesis that environmental complexity promotes the evolution of internal complexity, 12 evolved robots were analysed; 6 from high-noise simulations (A

<sup>3</sup>See also Jakobi, Husbands, and Harvey (1995), Jakobi (1998b). Jakobi distinguishes between ‘base set’ features - those aspects of an agent-environment system that are necessarily involved in the generation of the desired behaviour - and ‘implementation’ features - those aspects which are either simulation artefacts, or not relevant to the behaviour, or just real world properties that are simply too difficult to model in simulation. His methodology consists of two imperatives. First, that an ‘envelope of noise’ should be placed around simulation base-set features such that behaviours that evolve to rely on them will also retain their coherence in the ‘real world’. This follows from the intuition that simulation models can never be entirely accurate, but that an ‘envelope of noise’ around simulation base-set features can encapsulate the corresponding ‘real world’ feature. Evidence for the validity of this intuition can be found in Jakobi et al. (1995), Miglino, Lund, and Nolfi (1996). Jakobi’s second directive - and the distinguishing feature of his methodology - is that implementation features should be made *extremely* unreliable through the application of very high levels of noise, so that the evolutionary search process cannot come to incorporate them in *any* viable controller.

<sup>4</sup>This Khepera was fitted with a floor sensor as in Floreano and Mondada (1996). Many thanks to Adrian Thompson for help with this.




Figure 4.3: Twelve images taken at intervals of 3 seconds from a digitised video sequence of a (real) Khepera robot engaged in ‘homing-navigation’ behaviour (the sequence proceeds from left to right, top to bottom).

robots), and 6 from low-noise simulations (*B* robots).<sup>5</sup> Three environmental conditions were analysed for each robot; a ‘normal’ condition (*NO*), with both light source and charging area present, a ‘no charging area’ condition (*NC*), in which the black paint is removed and the robot cannot recharge, and a ‘no light source’ condition (*NL*), in which, although the charging area is present, the light source at the corner is removed.<sup>6</sup> The robots were all evolved in the *NO* condition, with the *NC* and *NL* conditions deployed only for test purposes. Low noise levels were employed during all testing.

Figure 4.4(a-c) illustrate typical overhead trajectory plots for the *A* robots in the three conditions, and (d-f) illustrate the same for the *B* robots. In the *NO* condition, both *A* and *B* robots can repeatedly find the charging area (situated in the lower left hand corner), and their trajectories are not obviously different. However in the *NC* and *NL* conditions, whereas the *B* robots maintain a pattern of behaviour that is qualitatively similar to that displayed in the *NO* condition, the *A* robots do nothing of the kind.

The *B* robots seem only to have evolved to move in straight lines and to turn upon encountering walls; a strategy which can indeed periodically return the robot to the charging area, but only reliably so in *NO* environments. By contrast, the *A* robots are clearly affected by the presence (or absence) of the black charging area and the light source. In the *NC* condition, these robots head towards the charging area and remain in its vicinity, as if ‘confused’ by the absence of charging; they appear to be *searching* for the missing charging area. In the *NL* condition, the robots begin, as in normal conditions, with a semilinear trajectory, but after a while they begin to *circle*. This circling behaviour makes good sense if the robot is considered to be attempting to orient to a light

<sup>5</sup>Robot controllers evolved in low-noise conditions did *not* transfer effectively to the real world, as is to be expected from Jakobi (1998a), but - as noted in section 4.1.1 above - this is of no relevance to the central argument of this chapter.

<sup>6</sup>These tests were also performed by Floreano and Mondada (1996), whose observations were similar to those made of the type *A* robots in the present study.

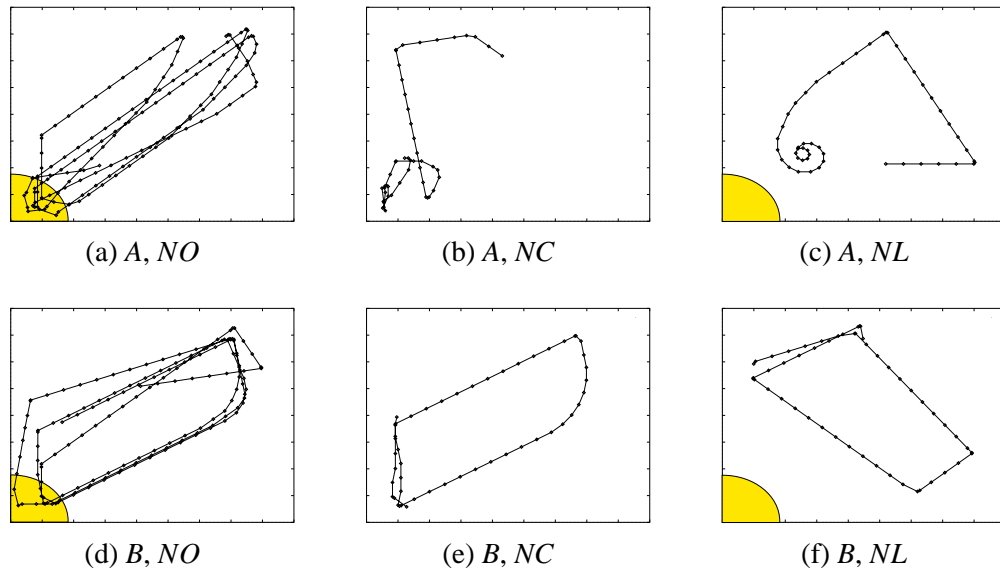


Figure 4.4: Trajectory plots of robots evolved under high-noise (a-c) or low-noise (d-f) conditions in either *NO* (a,d), *NC* (b,e), or *NL* (c,f) conditions. The charging area (if present) is indicated by the shaded area of each plot. The *B* robots maintain a simple trajectory regardless of the environmental manipulations, but the *A* robots deploy more complex ‘searching’ and ‘circling’ behavioural strategies.




Figure 4.5: Type *A* and type *B* robots in a ‘no-wall’ competition. The charging area is indicated by the shaded area of each plot.

source using front and rear light detectors. All 6 *A* robots presented qualitatively similar searching and circling behaviours, and all 6 *B* robots displayed the same simple behaviour (as in figure 4.4d-f). Also, these searching and circling behaviours are observed in real world Khepera behaviour, when controllers from *A* robots are downloaded to real Kheperas, with the ‘real’ environment manipulated in the appropriate way.

Are the *A* robots behaving in a more complex fashion than the *B* robots? Arguably they are. Their behaviour is certainly more heterogeneous, flexible, and variable. In the present context these are good enough indications of behavioural complexity.

Similar observations can be drawn from a competition between an *A* robot and a *B* robot in a condition (in a low-noise simulation) in which the arena walls are removed, and in which the

charging area extends in a complete circle around the light source. Figure 4.5(b) illustrates that the *B* robot is impotent in such circumstances, hinting at a reliance on IR stimulation (one typical run is shown, and out of 40 test runs the robot only managed a single visit to the charging area). By contrast, the *A* robot found the charging area on 10 attempts out of 40, and in 4 cases returned more than once. Figure 4.5(a) illustrates a particularly impressive *A* robot trajectory, and although in general the robot is undeniably adversely affected by the lack of walls, the considerably greater success rate enjoyed by the *A* robot strongly suggests that a greater range of environmental stimuli is being assimilated in the determination of its behaviour.

### 4.1.3 Mechanistic analysis

The purpose of this section is to illustrate that the mechanistic dynamics of the *A* robots are more complex than is the case for the *B* robots (this again has to be a qualitative judgement, as in the behavioural case), and, further, that the enhanced mechanistic complexity of the *A* robots makes sense in terms of the more complex behaviour displayed by these robots. Initially, activation plots for all 19 neurons in all 3 conditions (*NO*, *NC*, and *NL*) for all of the 12 robots (6 type *A*, and 6 type *B*) were compiled. These plots illustrate that whereas the hidden units (HUs) of the *B* robot neural networks react almost solely to IR stimulation, those of the *A* robots react much more strongly to light, battery, and floor sense data. Figures 4.6 and 4.7 provide examples from one robot of each type in the *NO* condition. Notice, at time-steps 35, 60, 95, and 125 in figure 4.6 (an *A* robot), that HU and motor output can be seen even when there is *no* IR input. This is never the case in figure 4.7 (a *B* robot). Also, in the ‘no-wall’ condition discussed in section 4.1.2 - and therefore in the absence of any IR input - only the *A* robots display any significant HU activity (figure 4.9).

To pursue this analysis in a non-behavioural context, short periods of activation were injected into 6 combinations of input units, with the activations of the HU and motor units being recorded during each such period. The first two combinations consisted of IR inputs only, with either all 8 inputs active, or all except the rear two. The next two tested combinations of ambient light inputs in the absence of IR input. Finally, a negative floor sensor input was injected (as if the robot were over the charging area), and (separately) a negative battery input (signifying an empty battery), both in the absence of IR. These 6 combinations were tested on each of the 12 robots, and figure 4.8 presents summary data for all 12 robots over all the 6 combinations, in terms of the HU activity elicited by the various inputs. For example, for the input combination involving (negative) battery unit activation, 60% of the type *A* robot HUs responded strongly, compared to 20% of the type *B* robot HUs. Thus, figure 4.8 makes it clear that the *A* robots take greater account than the *B* robots of the light and battery sense data. These results are statistically significant according to Mann-Whitney U tests ( $U = 57.0; df = 6, 6; p < 0.01$ ,  $U = 56.5; df = 6, 6; p < 0.01$  respectively). Statistical significance cannot, however, be asserted on this sample for the apparent tendency of *B* robots HUs to take greater account of IR sense data.

The reliance of *A* robots on ambient light information is particularly clear in the ‘no-wall’ condition discussed beforehand, in which only the *A* robots display any significantly varying HU activity (figure 4.9). The mechanistic dynamics of the *A* robots are therefore taking into account a wider variety of environmental stimuli than is the case for *B* robots, and forging them into a coherent and complex behaviour. Furthermore, these observations are consonant with the behavioural






Figure 4.6: Neuron profile for type *A* robot (evolved in high-noise conditions), tested in the *NO* condition. HUs (H1-H5) and motor units display activity *not* correlated with IR input (notably H1,H3,H5 - see time-steps 35, 60, 95, and 125, marked by the shaded bars).




Figure 4.7: Neuron profile for type *B* robot (evolved in low-noise conditions), tested in the *NO* condition. HUs (H1-H5) seem to be predominantly reacting to IR input.

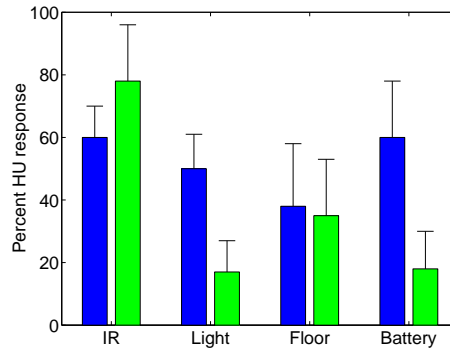


Figure 4.8: Hidden unit response patterns for type *A* (dark columns) and type *B* (light columns) robots. Error bars represent standard deviations.

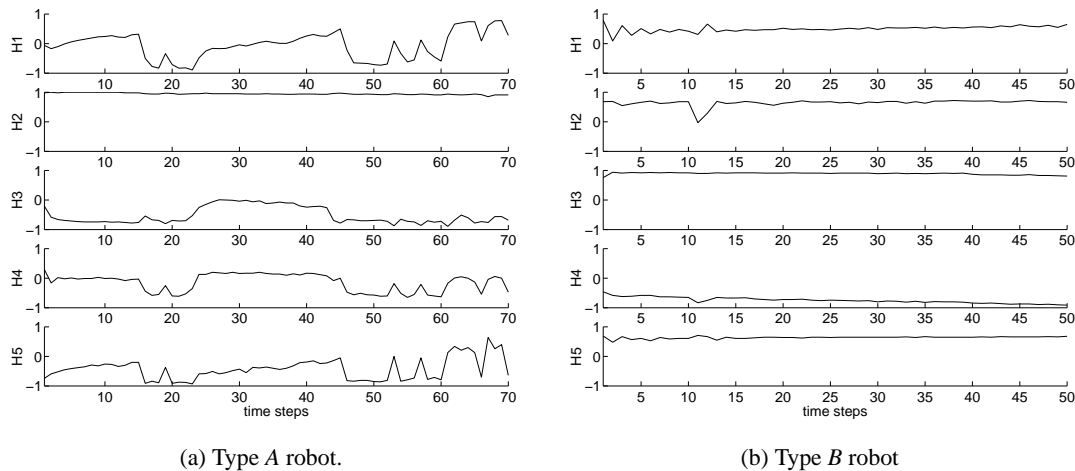


Figure 4.9: Type *A* and type *B* robot hidden unit activation in no-wall test (only the first 70 time-steps are shown).

analysis of section 4.1.2, in that only the type *A* robots are strongly affected by manipulations of the light source and charging area.

#### 4.1.4 Discussion

This section has extended the empirical explorations of the previous chapter into the domain of evolutionary robotics. Further support has been provided for the hypothesis - a version of Godfrey-Smith's ECT - that environmental variability can promote the (artificial) evolution of complexity, but in contrast to the quantitative evidence provided by the previous chapter, the evidence of this chapter is qualitative.

The context of evolutionary robotics does however allow a return to some of the major themes of chapter 2. Recall Dewey's contention that 'cognition' is only useful in environments characterised by a mixture of variability and stability. This condition is illustrated in the present model in that the variable environment necessary for the evolution of the more complex type *A* behaviour also maintains stability in certain properties (for example the presence of a charging area and the modal reliability of the robot sensors). Contact is also made with the relation between be-

haviour and mechanism: A fixed neural network architecture generates ‘homing-navigation’ in two very different ways, distinguished largely by the extent to which environmental stimuli of different modalities influence the dynamics of the network. It is this kind of analysis that targets the circularity of the version of the ECT outlined in chapter 2; a certain (and requisite) degree of mechanistic pre-structure makes it the case that certain features of the external environment can comprise features of the *Umwelt* of the robot, and the mechanistic dynamics through which the robot responds to these features can then be understood in terms of the variability presented by the ensemble of such features. It is apparent here that the ECT cannot be applied to the complexity of behaviour alone, nor only to the complexity of the underlying mechanism. In the present model, the relative behavioural complexity of the *A* robots is only revealed following radical restructuring of the environment, and the complexity of the mechanistic dynamics only makes sense in terms of subserving the observable ‘homing-navigation’ behaviour.

More directly, it may be noted that although type *A* robots display distinct ‘searching’ and ‘circling’ behaviours, there is no sensible way in which these behaviours map onto internal behavioural ‘icons’ of any sort. Nor can descriptions of *A* robots as ‘attending’ to the light (or battery, or floor sensor) at some times but not at others, or as ‘selecting’ between distinct behavioural options (circling, searching, exploring, homing), be mechanistically justified. The light source, for example, appears to influence mechanistic dynamics (and hence the behaviour) *not* just during phototaxis, but at all times, and indeed the same may be said for all the other sensory modalities (figures 4.6 and 4.7). Also, the apparent junctures in behavioural dynamics are *not* imposed from within by any putative ‘behavioural arbitration device’, but simply arise from the continuously unfolding agent-environment interaction patterns. These issues - of ‘attention’ and in particular ‘behavioural coordination’ - will be the focus of much of what follows in the remainder of this dissertation.

From the perspective of evolutionary robotics itself, this chapter has demonstrated a new way of promoting the evolution of complex behaviours and mechanisms, through the judicious application of simulation noise. The nature of the mechanistic complexity of the *A* robots is also of considerable interest; the integration of multimodal sensorimotor data is an issue of importance in its own right, and in this example has been achieved without having been explicitly specified in the fitness function (or hand-coded in any way).<sup>7</sup>

On a methodological note, the construction of the simulation required that collisions with walls be modelled, and modelled noisily, despite the fact that interactions with walls played no part in the eventual evolved behaviour. The suggestion here is that effective artificial evolution in simulation may depend on the simulation of agent-environment interaction patterns that serve to *scaffold* the development of the desired behaviour. There is an extensive literature on scaffolding in developmental psychology (see, for example, Vygotsky, 1986), and it may be speculated that the development of a theoretical relation between evolutionary robotics and developmental psychology might be of considerable utility.<sup>8</sup>

To conclude, the ECT provides a theoretical response to the *causal* question - of how is it

<sup>7</sup>See, for example, Hager (1990), Stein and Meredith (1994), Brooks and Richard (1998), and Arkin (1998).

<sup>8</sup>Some background to such an endeavour may be obtained from Rutkowska (1994a, 1994b, 1997). Although this body of work does not directly address evolutionary robotics, it does discuss relations between developmental psychology and SAB/AL in general. Also relevant, but again distinct, is the ‘neural Darwinism’ of Edelman (1987), which is a selective theory of neural development that has been implemented on robotic platforms.




Figure 4.10: Each figure plots the lifetime of the fittest robot of each generation against the generation number, overlaying results from 5 simulations in each condition.

possible for evolution to produce structures of complexity from structures of lesser complexity - and both this chapter and the preceding chapter have provided empirical versions of this response, in the context of artificial evolution. It is *not* a conclusion of this dissertation that evolution will always lead to complexity, or even that environmental variability will always lead to the evolution of complexity; I have merely presented examples of situations when this is the case.

## 4.2 Noise and (artificial) evolution

This final section entertains a more general discussion of the role(s) of noise in artificial evolution, with a pragmatic - and theoretically speculative - emphasis. This discussion will require a generalisation of the distinction between M-noise and O-noise (introduced in the previous chapter) with M-noise now referring to environmental noise, and O-noise to random fitness adjustments.

As noted in section 4.1.1, perhaps the most prominent role for M-noise in artificial evolution is in the facilitation of transference from simulation to reality, as developed in detail by Jakobi (1998b); this will not be discussed further. The present chapters have focussed on the role of M-noise in engendering the evolution of complexity, but have also touched upon the possibility of both M-noise and O-noise facilitating the exploration of genotype space by an evolutionary search process. There remains a final influence to reveal; that M-noise can *accelerate* an evolutionary search process. This influence is illustrated in figure 4.10 in the present context of the evolution of ‘homing-navigation’ behaviour. Multiple fitness plots for evolutionary simulations both with and without M-noise clearly indicate that the presence of M-noise entails a more rapid attainment of an optimal behaviour.

How might the application of noise entail these diverse effects? One possibility is that M-noise, selectively applied to only parts of a simulated agent-environment system, may alter the large-scale structure of the corresponding fitness landscape. For example, in the ‘homing-navigation’ model, by making the turning rate of the (simulated) robot sufficiently unreliable, the simple (type *B*) behaviour along the lines of ‘turn through  $\theta$  degrees every time you approach a wall, and you’ll get back to base’ is rendered infeasible, and points on the fitness landscape corresponding to implementations of this strategy will no longer be maxima. If this is correct, it also provides a




Figure 4.11: Convergence statistics for 4 runs of a ‘homing-navigation’ simulation with M-noise. Population converges rapidly to an average Hamming distance of about 20%.

hypothesis for why M-noise might accelerate evolution in this model: It may be that the application of M-noise can lead to a reduction in the overall density of maxima in the fitness landscape - selective large scale alterations of the landscape structure - allowing the evolutionary search to proceed with greater rapidity towards its ultimate destination. These ideas are aired as informal speculations, and for any plausibility require - at the very least - that the search process operate over a relatively converged population. Figure 4.11 illustrates that this is arguably the case for the ‘homing navigation’ model of this chapter, although it must be conceded that higher degrees of convergence are often observed in artificial evolution experiments.

To understand how noise may facilitate genetic mobility, an analogy with the Baldwin effect may be helpful. The Baldwin effect describes the way in which lifetime learning can influence the course of (non-Lamarckian) evolution by ‘smoothing’ sharp peaks in a fitness landscape (Baldwin, 1896, see also Belew and Mitchell, 1996). The idea is that if individuals vary genetically in their ability to learn some useful trait, then those most able to do so will leave the most descendants, and the genes responsible for the relatively rapid learning will spread through the population. Over many such generations this process can lead to the genetic determination - or ‘assimilation’ - of a trait which originally had to be learned from scratch each generation. Hinton and Nowlan (1987) provide a simple computer simulation demonstration of this effect, and the learning algorithm they use is simply random search, which can be informally construed as O-noise since it has the effect that any given genotype can be awarded a greater diversity of fitness values than would otherwise have been the case. By this analogy, O-noise may ‘smooth’ certain dimensions of the fitness landscape, allowing the GA to search the corresponding dimensions of genotype space more extensively. And to the extent that M-noise entails variation in the fitnesses awarded to a given genotype, it will have a similar effect. This putative ‘smoothing effect’ may also be implicated both in how noise can facilitate the evolution of qualitatively different behaviours, by ‘smoothing out’ any steep valleys that may lie between maxima representing the different behaviours, and also in the acceleration effect, by making any given maxima easier to find and easier for a population to climb.

Despite the admittedly speculative nature of this discussion, the relation between M-noise and O-noise should now be a little clearer. O-noise incurs M-noise to the extent that ‘smoothing’ is thought of as a change in fitness landscape structure. M-noise incurs O-noise in that, with M-

noise, individuals may (but need not) accrue a wider variance of scores than they would without M-noise. But O-noise will only ever incur that M-noise associated with smoothing, and will never incur M-noise associated with selective large-scale structural alterations of the fitness landscape (for example the selective removal of certain local maxima). And it is this engineering of the fitness landscape that may solicit adaptive *response* which is reflected in the evolutionary search reaching maxima corresponding to qualitatively different phenotypic expressions (the robot now has to integrate sensorimotor information rather than relying on a single modality).

Finally, although the distinction between M-noise and O-noise derives from Ashby's LRV, this discussion hints at roles for noise that go beyond Ashby's original conception. Recall his position on variety: "There is that which threatens the survival of the gene-pattern. This part must be blocked at all costs. And there is that which, while it may threaten the gene-pattern, can be transformed (or re-coded) through the regulator *R* and used to block the effect of the remainder" (1956, p.212). Thus for Ashby, if variability is significant, it is so in virtue of its (positive or negative) functional salience. The present discussion extends this conception with variability implicated in the facilitation of genetic mobility and in the acceleration of evolutionary search.

### 4.3 Summary

Environmental variability can promote the (artificial) evolution of complexity. This is the primary message of this chapter - expressed in the medium of evolutionary robotics, and also of the previous chapter - expressed in the context of evolutionary game theory. The flavour of these chapters has been exploratory, with models developing empirical aspects of the conceptual enterprise described in chapter 2. This interplay has been most prominent in the use of Ashby's LRV to formalise the dynamics of the IPD, and in the distinction between behaviour and mechanism in the analysis of the mechanisms underlying 'homing-navigation' behaviour in this chapter.

Although these models may be considered to have reached a satisfactory stage of completion, it is in their nature that they generate more questions than they resolve. How exactly does noise influence the dynamics of artificial evolution? Is the multimodal sensorimotor integration observed in the type *A* robots an instance of a more general strategy for coping with environmental complexity? In any case, three novel roles for noise in artificial evolution have been identified, and may be summarised thus: the evolution of complex behavioural/mechanistic structure, the acceleration of the evolutionary process, and the facilitation of the genetic mobility.

From now on the flavour of the dissertation is somewhat different; less exploratory, and directed more specifically towards the elucidation of particular biological and psychological phenomena. The primary themes that have occupied us so far will continue to predominate, but in a methodological context that is more rigorous, and in empirical contexts of greater precision; in particular, the issue of behavioural coordination, alluded to in section 4.1.4 above, will become a central concern. The following chapter lays the necessary methodological groundwork.

## Chapter 5

### Unorthodox optimal foraging theory

---

[T]he capacity to analyze organic entities is directly proportional to the willingness to shift to a different explanatory principle.  
- Konrad Lorenz (1948, p.207)

In the homing-navigation model of the previous chapter, a GA was employed to optimise a control architecture for a situated agent. This chapter describes a generalisation of this methodology in the form of an extended analogy with the (theoretical) biological methodology of ‘optimal foraging theory’ (OFT). As such, this chapter marks a change of direction in this dissertation. Until now the emphasis has been on developing concrete models to explore issues that have significance largely in view of their appearance in the discourse of SAB/AL, notably complexity and its emergence. From now on the empirical content of this dissertation will engage with problems that are relatively well-described in the literature of contemporary biology and psychology. This chapter lays the necessary methodological groundwork, developing an unorthodox complement to OFT through the description of a class of SAB/AL models - of which the homing-navigation model is one example - that shall be referred to as ‘individual-based optimal situated’ (IOS) models.<sup>1</sup> This groundwork not only justifies the structure of the various empirical models that follow, it also allows several fundamental problems that arise directly out of the framework of orthodox<sup>2</sup> OFT to be addressed, although, as with any trade-off, not without disadvantages of its own. As such, the IOS modelling strategy offered here is of significance to any researcher interested in exploiting the relationship between SAB/AL and theoretical biology. For this reason, the arguments of this chapter will be illustrated with examples drawn from a diverse literature, although links to empirical work in the later chapters of this dissertation will of course be pointed out where appropriate.

Despite the change of direction initiated by this chapter, the major themes of this dissertation persist. Both orthodox OFT and its unorthodox complement are strongly externalist in character, and both interrogate the relations between behaviour and mechanism, although, as I shall argue, IOS models may be expected to do so with greater success. Indeed, the IOS modelling strategy

---

<sup>1</sup>I employ a novel acronym with reluctance and only for want of an appropriate existing abbreviation.

<sup>2</sup>As G.K. Chesterton quipped, “[t]he word ‘orthodoxy’ not only no longer means being right, it practically means being wrong”; the present use of the word is, however, *not* intended to carry any such pejorative connotations.

may be construed as a methodological expression of the revised version of the environmental complexity thesis (ECT) described in chapter 2.

The essence of the argument is as follows. Orthodox OFT (introduced in section 5.1) is an attempt to understand both functional and mechanistic aspects of animal (foraging) behaviour by treating observed behaviour as ‘optimally’ adapted to an environmental ‘problem’. I argue that orthodox (classical, equational) OFT models carry with them several substantial ‘framework assumptions’ which both limit their explanatory flexibility, and undermine the rationale behind the ‘incremental revision’ process that characterises the orthodox OFT research cycle. I concentrate on identifying the problematic consequences of these framework assumptions and indicating, in each case, the benefits offered by IOS models (introduced in section 5.2); most of these flow from an increased ability to model subtle but significant agent-environment interactions mediated by situated perception and action. Some balance is introduced by pointing out that this extended explanatory flexibility will often be accompanied by reduced explanatory transparency; it can be difficult to understand how IOS models do what they do. I also suggest that IOS models do not necessarily operate in the mode of incremental revision typical of orthodox OFT, rather that they find value in focussing attention on components of foraging behaviour that are either concealed, or exaggerated, through the framework assumptions of orthodox models. One consequence of this is that IOS models facilitate the generation and evaluation of hypotheses concerning the functional potential of very *simple* mechanisms.

## 5.1 Optimal foraging theory

OFT starts from the assumption that foraging behaviour can be treated as a Darwinian adaptation, and, as such, can be considered to be an ‘optimal’ solution to a ‘problem’ posed by (present and/or past) environments. A strongly externalist project, and similar in many ways to evolutionary psychology (see chapter 2), it seeks to understand the foraging behaviour of animals through the identification of what would constitute optimal foraging behaviour given the constraints on the animal and of its (present and/or past) environment. As an example to follow through this discussion, consider the redshank *Tringa totanus*, a shorebird, foraging for both large and small worms (Goss-Custard, 1977).<sup>3</sup>

Orthodox OFT models comprise three primary components: a decision variable, a currency, and a set of constraints (Stephens & Krebs, 1986). Their specification does not require that individual foraging agents be explicitly instantiated; orthodox models draw behavioural inferences simply from mathematical (equational) relationships between the components. All orthodox OFT models consider the ‘best’ way to make a foraging decision, and the *decision variable* captures the type of choice that the animal is assumed to make (or that natural selection has made for it). In the redshank example, the decision is whether or not to eat small worms upon encounter (it is assumed that large worms will always be eaten), and the decision variable has two possible states; ‘eat’ and ‘reject’. The *currency* specifies the quantity that the animal is supposed to be maximising. The ultimate mandate for OFT in Darwinian natural selection requires that this currency be identifiable with reproductive fitness, but since this quantity is more or less impossible to assess in any given study, practitioners of OFT invariably employ a proxy. In the redshank example, the

<sup>3</sup>A more detailed exegesis of this example can be found in Krebs and Kacelnik (1991).



rate of energy intake is the currency, with (it is assumed) large worms providing more energy than small worms. *Constraint* assumptions limit and define the relationship between the currency and the decision variable. Intrinsic constraints derive from properties of the animal; the speed with which the redshank can search for prey and the handling time required for each captured worm both qualify. Extrinsic constraints derive from the environment, for example the densities of the two types of worm. Orthodox OFT models also include a *choice principle* which defines how optimisation of a currency is to be interpreted; this is almost always ‘maximisation’ but in principle could be some other function (for example a step function would lead to a ‘satisficing’ formulation of OFT). A *currency function* is also required to transform the decision variable into a value commensurable with the currency. And finally there must be a method of optimisation, ranging from analytic solution to the numerical methods of ‘stochastic dynamic programming’ which can cope with state-dependent decisions and uncertain environments (Mangel & Clark, 1988).

These components together constitute what may be called the framework assumptions of orthodox OFT, and are usually defined after a period of observing the behaviour of a focal animal. An equational model is then built, which is used to predict how the animal should behave in different environments in order to be ‘optimal’, and the extent to which observed and predicted behaviours differ is taken to indicate the degree of revision (in terms of the incorporation of additional components - usually constraints) that the model need undergo; these conclusions are framed as testable hypotheses of the form ‘(additional) component *P* plays such-and-such a role in observed foraging behaviour *Q*’. This mode of operation may be termed ‘incremental complexity’ (to use the terminology of Maley, 1999); that is, an initially simple model is augmented with additional constraints until (if ever) it satisfactorily predicts real world behaviour. For example, Goss-Custard (1977) predicted that the foraging redshank should *either* always accept *or* always reject small worms. He observed, however, that under some conditions redshanks would accept *some* small worms. Revised models, proposed to account for such ‘partial preferences’, incorporated additional constraints of imperfect discrimination (Rechten, Avery, & Stevens, 1983) and minimisation of starvation (Houston & McNamara, 1985).

It must be stressed that orthodox OFT models *cannot* tell us whether or not an animal ‘optimises’. Some authors have criticised OFT for being ‘unscientifically irrefutable’ in the sense that ‘alternatives to optimality’ are never considered, rather, practitioners of OFT insist on “shoring up the cracked facade of optimality with a scaffolding of *ad hoc* modifications” as the imaginary critic of Stephens and Krebs (1986, p.207) has it (for a less lyrical but substantively identical dissension see Pierce & Ollason, 1987). But of course optimality is not the hypothesis under test, it is a way of generating testable hypotheses. In the long run OFT will stand or fall on the knowledge gained from the progressive confirmation or refutation of these hypotheses, and *not* because of any ‘demonstration’ that animals do, or do not optimise.

### 5.1.1 What use is orthodox OFT?

OFT models can ask what animals are ‘designed’ to do, and they can also attempt to unravel the behaviour-generating mechanisms involved. Notice that these explanatory projects are equivalent to those identified for evolutionary psychology in Barkow et al. (1992).

The first kind of question invokes a functional perspective, and involves finding a set of con-

straints and a currency (or currencies) such that an observed behavioural pattern may be considered optimal. A focus on currency assesses the validity of a chosen proxy for reproductive fitness. A behaviour that appears suboptimal in the context of maximising intake rate may be seen as optimal in the context of a different currency, as in the case of the currency of minimisation of starvation proposed for the redshank by Houston and McNamara (1985). A focus on constraint addresses aspects of environment structure that have functional significance, these being *extrinsic* constraints. That is, we accept that  $X$  is maximising intake but we want to understand how this general goal translates into specific actions in environment  $Y$ . An observed behaviour may maximise intake rate only in the presence of certain extrinsic constraints presented by  $Y$ ; we hypothesise that  $X$  is ‘designed’ to interact with these features of  $Y$  so as to maximise intake. From this perspective, instances of suboptimal behaviour may be attributed either to environmental features that may have formed part of past, but not of present environments (see, for example, Kacelnik, 1997), or to the impact of previously unidentified extrinsic constraints.

There is a problem here. Given a discrepancy between model prediction and observed behaviour, there is no general rule for deciding which of these foci to attend to, and, worse, it is not even possible to maintain them as distinct options in principle. Attention to extrinsic constraints can be considered to be an elaboration of currency; for example a goal of intake rate may multiply into a set of sub-goals (explore, find food, hoard food, etc.), each associated with (possibly overlapping) sets of extrinsic constraints. On the other hand, recourse to multiple currencies (quantity of food located, quantity of food hoarded, etc.) can represent the interpolation of extrinsic constraints between reproductive fitness and behaviour. The utility of these currencies as proxies for reproductive fitness can then be seen to depend upon the presence of particular extrinsic constraints (food hoarding may only make sense in particular kinds of environment). This conceptual overlap between ‘currency’ and ‘constraint’ is revisited in section 5.3.6.

Hypotheses about internal mechanism follow from an analysis of decision variables and *intrinsic* constraints, both of which may be considered to represent aspects of the internal mechanism of an animal. Such a strategy requires that both extrinsic constraints and currency be relatively well specified; a condition that is also important in evolutionary psychology, recall from chapter 2 that specification of function is an essential part of the formulation of accounts of mechanism. An example of this strategy can be found in Cheverton, Kacelnik, and Krebs (1985), who assume that bumblebees are maximising intake rate and then use ‘errors’ in performance to infer properties of the mechanism controlling movement between flowers; these inferences constitute mechanistic hypotheses. Application of this explanatory strategy accounts for seemingly suboptimal behaviour in terms of the impact of previously unsuspected *intrinsic* constraints.

The IOS methodology, introduced below, retains both of these explanatory projects. Examples of the functional perspective appear in chapter 7 of this dissertation, which deals with the so-called ‘interference function’, and chapter 8, which explores the ‘matching law’ of experimental psychology. In both cases, situations are described in which the apparently suboptimal foraging behaviour of isolated individuals can be understood in terms of adaptation to a group context. The mechanistic perspective is most thoroughly explored in both the following chapter and in chapter 8; these chapters suggest that the mechanisms underlying effective choice behaviour, in a variety of situations, need not be as complex as is often thought.

## 5.2 IOS models

What, then, are IOS models? The homing navigation model of the previous chapter is one example. More generally, a biological perspective would view IOS models as derivatives of the larger class of ‘individual-based models’ which “[t]reat individuals as unique and discrete entities with at least one property [...] that changes during the life cycle” (Grimm, 1999, p.130), and as distinguished from this larger class in two ways. First, they are endowed with a dimension of optimality through the use of GAs to design agent architectures. Second, they deal exclusively with *situated* agents. Note that the adjective ‘situated’ has a particular meaning throughout this dissertation; it connotes not just that agents are contextualised in some way (as will be the case for many individual-based models), but that they interact with a spatiotemporally structured environment by means of perception and action (this does *not* require, however, that agents be *physically* instantiated or embodied). A point of agreement between individual-based models and IOS models is that they both offer alternatives to ‘classical’ models that consist primarily of state variables and systems of equations that average over properties of individuals. The essential components of IOS models are therefore (i) a computer simulation model of agent and environment framed in terms of situated perception and action, (ii) a supervening GA - or alternative optimisation procedure - which can potentially modify aspects of agent structure, and (iii) a fitness function (a metric by which the performance of agents can be evaluated and compared).

In both the homing-navigation model and the prisoner’s dilemma models of previous chapters, the relation between GAs and biological evolution was not explored in any detail, given that the theoretical contexts of the ECT and the LRV required only a loose analogy. However, the present use of GAs in IOS models must be justified more carefully, given the relatively strict analogy proposed between the IOS modelling strategy and orthodox OFT. As stated in section 5.1, orthodox OFT assumes that foraging behaviour is an ‘optimal’ solution to an environmental ‘problem’, and as such GAs are utilised in IOS models above all for their optimisation properties. It is patently *not* intended for any significance be drawn from the ‘evolutionary’ connotations of GAs; IOS models can in principle accommodate the use of any optimisation algorithm capable of exploring a space of agent structure and delivering ‘near-optimal’ solutions. Note that ‘near-optimality’, in optimisation theory, describes the results of processes which, although designed to find optimal solutions, cannot be guaranteed to do so; ‘near-optimal’ solutions are (almost) always very good, even if they may not always be the best (Beightler, Philips, & Wilde, 1979).

Perhaps the most obvious application of IOS models - to rehearse a dominant theme in this dissertation - is in the interrogation of the relations between behaviour and mechanism. As exemplified by the homing-navigation model of the previous chapter, the use of a sufficiently simple fitness function gives the evolutionary search process a relatively ‘free hand’ in designing mechanisms to fulfil observer-defined functional criteria. In other words, IOS models allow the relations between behaviour and mechanism to be the object of study rather than assumption. As Nolfi (1998) says (with respect to the closely related discipline of evolutionary robotics):

Evolutionary robotics, by relying on the evaluation of the system as a whole and of its global behavior, releases the designer from the burden of deciding how to break the desired behavior down into simple basic behaviors. (1998, p.6)

Although this general strategy will continue to be of importance throughout this dissertation, the

remainder of this chapter focuses on developing a more specific relationship between the IOS methodology and orthodox OFT.

### 5.2.1 IOS as OFT: at first glance

IOS models can be construed as unorthodox OFT models in the following way. The method of deriving an optimal solution (and the choice principle) in orthodox OFT clearly correspond to the GA in IOS models. Currency corresponds to the fitness function. The constraints are determined by a combination of environment (extrinsic) and phenotype/mechanism (intrinsic) - both those aspects open to variation during optimisation and those that remain fixed. The currency function corresponds to the process of evaluating an agent within its environment and giving it a fitness score. All these equivalences should be obvious. Decision variables, by contrast, do *not* have necessary counterparts in IOS models. It is entirely up to the modeller whether or not the phenotypic architecture utilised mandates interpretation in terms of the operation of a decision variable. Some may do, for example architectures involving simple rules that hinge on particular internal state variables, and some may not, for example some classes of simple neural network.

It must be emphasised that IOS models are nowhere proposed to be ‘better’ than orthodox OFT models. Their relationship is complementary, and a trade-off is already apparent: whereas orthodox OFT runs the risk of overly restricting the space of possible hypotheses through the biases inherent in the orthodox framework assumptions, IOS models risk confounding *post-hoc* analysis through the complexity engendered by the relatively large number of components and parameters (and interactions between them) constitutive of even simple instances. In the following section we explore reasons for favouring the IOS side of this trade-off, but it should not be forgotten that compelling arguments for engaging in orthodox modelling will always exist, some of which will be outlined in section 5.4.

## 5.3 Difficulties with orthodox OFT and ways around them with IOS models

Orthodox OFT, in spite of (or perhaps because of) its indisputable popularity amongst many biologists, has remained controversial. Some criticisms are frequently rehearsed and futile (most notably the debate surrounding ‘alternatives to optimality’), of these we will hear no more. Others are well formed and significant, some of these are described and expanded upon below, and still others are novel to the present discussion. No comprehensive coverage is attempted here; there is neither the space nor the need. Many criticisms focus on aspects of OFT that are best addressed from within the orthodoxy; see Pierce and Ollason (1987), and for a defence, Stearns and Schmid-Hempel (1987).

### 5.3.1 Mathematical tractability

Of those criticisms of OFT that have found expression in the literature, perhaps the most common address the implausibility of specific constraints often employed in orthodox models for reasons of mathematical tractability. For example, many conventional foraging models (including the original redshank model of Goss-Custard, 1977) assume (i) the exclusivity of search (for prey) and exploitation (of prey), (ii) sequential pseudo-random encounters with prey, and (iii) that the animal has complete information about the state of the environment. Although the relaxation of these

constraints can be, and has been explored from within the orthodox framework (Stephens & Krebs, 1986), IOS models offer avenues for the progressive relaxation of other constraints that would be very difficult to achieve in any other way. For example, a constraint of homogeneous within-patch food distribution may resist manipulation in a mathematical model, but is immediately open to relaxation in an IOS model in which items of food may possess explicit spatial locations. There is a continuum of flexibility to be identified here. At one extreme lie the relatively rigid orthodox equational models, in the middle the general class of individual-based models allow assumptions of homogeneous group properties to be modified, and finally IOS models, being situated, permit the relaxation of constraints that embody assumptions about perception and action (remember that IOS models are *defined* as individual-based models that incorporate optimality and situated perception and action).

### 5.3.2 Decision variables

The remaining criticisms concern the framework assumptions of orthodox models in general, and are therefore difficult to resolve from within the orthodoxy itself. The first of these concerns the requirement that internal mechanism be understood in terms of the operation of some kind of decision variable. Although the mathematical instantiation of a decision variable does not necessarily imply that it is implemented as such in the animal, in practice, orthodox OFT provides no way of interrogating this implication. It is then all too easy to read orthodox OFT models as making significant and possibly erroneous claims about mechanism, a suspicion well articulated by Gallistel, Brown, Carey, Gelman, and Keil (1991):

What is striking about these optimality models is that in spelling out the decision processes underlying the optimisation of foraging behaviour they credit the animal with complex representational and computational abilities. (1991, p.18)

Also, the distinction between a decision variable and an *intrinsic* constraint does not appear to have any *a priori* justification. Both can be components of internal mechanism, and it may be that optimal behaviour can arise solely from the (externally observed) interplay of constraints (intrinsic and extrinsic) and currencies. There are also the attendant dangers of associating the operation of a decision variable with junctures in behavioural dynamics that may seem significant from the perspective of the external observer but which need not indicate the existence of mechanistic structures dedicated to the entailment of these junctures. This is once again to highlight the leitmotif of distinguishing between behavioural and mechanistic levels of description; as argued in chapter 2, internal mechanistic correlates for behaviours (or ‘choices’ between behaviours) should not be presumed to exist (of course this does not *prove* that decision variables cannot exist; see, for example, Platt & Glimcher, 1999, who claim to have isolated neural correlates for decision variables in the parietal cortex of rhesus monkey brains).

IOS models avoid these difficulties by permitting the construction of models *without* obvious decision variables (for example using simple neural networks). Such models broaden the space of possible internal mechanism, and, in particular, facilitate the generation and evaluation of hypotheses concerning the behavioural potential of very *simple* underlying internal mechanisms. In particular, the risk is no longer run of decision variables being confounded with junctures in behaviour that appear significant from the perspective of an external observer, and by the same token,

the awkward (and arbitrary) distinction between decision variable and intrinsic constraint can be abruptly dispensed with.

For the present, consider an example from evolutionary robotics. Mondada, Franzi, and Ienne (1993) found that it was much easier to evolve a *single* network controller for a complex task (to find and stay close to a target object) than it was to evolve four separate networks corresponding to the separable behavioural elements into which the task could be divided from the point of view of an external observer (these were; explore, avoid walls, approach target, discriminate target from walls). Indeed, they found the latter method impossible, concluding that at least in this case mechanistic and (at least some) behavioural decompositions were incommensurable. Any attempts to associate (mechanistic) decision variables with the observable junctures in behaviour would therefore have been wholly inappropriate. A similar argument may be stated for the homing navigation model of chapter 4, insofar as mechanistic analysis did not reveal any internal behavioural icons for ‘circling’ or ‘searching’ behaviour patterns. However, this analysis did not focus on switching between behaviours, and so its relation to the postulation of a decision variable is less clear. This is not so for the models that follow in chapters 6 and 8, which *are* explicitly concerned with behavioural switching. The contexts of ‘action selection’ and the ‘matching law’ explored in these chapters provide very clear articulations of the explanatory superfluosity of decision variables.

### 5.3.3 Historical constraints

Another problematic distinction in orthodox OFT is the separation of constraints into mutually exclusive ‘intrinsic’ and ‘extrinsic’ categories. Such a scheme of classification misses out any constraints that may arise through the history of interaction between agent and environment; these can be called *historical* constraints. One can talk about such constraints at a variety of timescales, notably behavioural, developmental, and evolutionary. The present focus on optimal *behaviour* encourages concentration on behavioural historical constraints, but evolutionary historical constraints will receive some attention in section 5.3.5.<sup>4</sup>

Consider, as an example, the construction of termite mounds. Some of the essential constraints on the behaviour of an individual termite that lead to it contributing towards the construction of a termite mound can only be framed in terms of the past history of interaction of that termite with its conspecifics and with the surrounding environment (in this instance the phenomenon is also known as ‘stigmergy’, Grassé, 1959).<sup>5</sup> As an example from the SAB/AL literature, consider the simple model of clustering behaviour in groups of artificial ‘ants’ in Deneubourg, Goss, Franks, Sendova-Franks, Detrain, and Chretien (1990). Individual ants wander at random and pick up or put down objects with probabilities determined by the density of objects in their neighbourhood. Over time more and more locations in the shared environment become ‘blocked’ to the formation of clusters since (if the above probabilities are appropriately selected) no objects are deposited in empty locations; these locations then constrain both the possible locations of clusters, and the fact that clusters will form at all (as they must do as the number of ‘unblocked’ locations diminishes).<sup>6</sup>

<sup>4</sup>The present use of the term ‘historical’ does not need to be yoked to notions of ‘contingency’ or of ‘frozen accidents’, as it is to an extent in some discussions of evolutionary history (Gould, 1989) and in the explanatory strategies Crutchfield’s ‘historicists’ (Crutchfield, 1994a), as described in chapter 3.

<sup>5</sup>I am thinking in particular of the process of ‘arch formation’ described in detail by Clark (1997, p.75).

<sup>6</sup>The fact that this model does not utilise an optimisation procedure, and is therefore not strictly an IOS model, does not detract from its role in illustrating a behavioural historical constraint.

Although this behaviour is somewhat simpler than termite mound building, the principle of historicity is the same and indeed the cluster patterns generated by Deneubourg et al.'s ants are very similar to the clustered 'graveyards' of ant corpses built by real ants *Messor sancta* (and others). Chapters 7 and 8 of the present dissertation develop further examples in which behavioural historical constraints play an important explanatory role: patterns of 'resource instability' created by groups of foraging agents are seen to impose constraints on the individual foraging strategies that generate these patterns in the first place.

In general, IOS models can elucidate behavioural historical constraints by providing a sufficiently rich medium in which agent-environment interactions can create dynamical invariants which constrain the future dynamics of the system, and which are able to alter their own conditions of realisation. These are constraints that need not be, and in some cases *cannot* be imposed from the start in either 'extrinsic' or 'intrinsic' guises. An essential qualification is that such constraints can arise in individual-based models at a variety of levels of description. Again there is a continuum; the richer the medium of interaction, the greater the potential for historicity. Individual-based models that do not incorporate perception and action may nevertheless permit the emergence of historical processes contingent on the (broadly construed) interaction of distinct agents, but only *situated* models (in the present sense of IOS models) can possibly engender historical processes from a substrate of *sensorimotor* interaction.

#### 5.3.4 External environment and *Umwelt*

It is difficult, from within the framework of orthodox OFT, to appreciate the structure of the environment from the perspective of the agent itself. As Pierce and Ollason (1987) surmise:

It will always be possible to identify a set of environmental characteristics with respect to which observed behaviour is consistent with a particular functional hypothesis, but this process is entirely circular. By asserting that animals perceive the environment in a particular way it would be possible to show that observed foraging behaviour was consistent with any functional hypothesis. (1987, p.114)

The problem is that orthodox OFT models, in not dealing with explicitly instantiated and situated agents, are framed in terms of the external environment (the environment as it might appear to us as observers), whereas many components of such models (notably the decision variable - if there is to be one - and the currency) should properly be framed in terms of the *Umwelt* of the agent. Recall from chapter 2 the argument that the use of the ECT to interrogate the relations between behaviour and mechanism requires that the process of constitutive construction - the tracing of relations between external environment and *Umwelt* - be first understood. The same argument applies here. Without an understanding of the process of constitutive construction, the opportunity remains, as Pierce and Ollason conclude, for researchers to propose that "animals perceive the environment in a particular way" to suit their functional convenience.

Orthodox OFT models, again because they do not deal with explicitly instantiated and situated agents, do not provide an effective means of tracing the relations between external environment and *Umwelt*. By contrast, IOS models allow the researcher to measure all those features of the external environment that come to constitute the *Umwelts* of explicitly instantiated and situated

agents; the levels of activation of the sensors in various conditions, and so on. Thus, the IOS modeller is *not* free to assert that animals perceive the environment in an arbitrary fashion (although, of course, the structure of the environment itself *is* under the control of the modeller). Instead, the *Umwelt* of the agent can be reconstructed by the researcher from the interplay of extrinsic, intrinsic, and historical constraints. Particular functional hypotheses can be entertained only to the extent that they do not conflict with these constraints, and, as mentioned above, such reconstruction may play an essential part in understanding the relations between behavioural and mechanistic levels of description.

### 5.3.5 The pragmatics of optimisation

By construing behaviour as an ‘optimal solution’ to an ‘environmental problem’ (section 5.1), orthodox OFT makes two implicit commitments: (i) that process and outcome (the optimal behaviour) are conceptually distinct, and (ii) that evolution is a process of optimisation. These commitments are separate; each will be dealt with in turn.

First, and in contrast to orthodox OFT, IOS models can be constructed in which the process and outcome can continuously interact and influence each other. Indeed, as soon as optimisation is *instantiated* as a process there is the potential for this to happen. Such interaction may be particularly significant if it entails the formation of behavioural historical constraints; recall from section 5.3.3 that such constraints often resist pre-specification as part of a problem for which a solution may be found. In these cases the optimal behaviour at any time (the outcome) cannot be understood without an appreciation of the preceding optimisation dynamics (as is the case, for example, in chapters 7 and 8 of this dissertation).<sup>7</sup> However, in order to maintain the methodological parallel between IOS and OFT, it is important to ensure that the GA itself is doing nothing more than optimisation. This means ensuring, as far as possible, that there are no interactions *between* genotypes that could affect the fitness values awarded to the genotypes. In this sense, IOS models retain an independence from the particularities of the method of optimisation.

Second, there are many GA-based models in which the evolutionary (process) mechanics of the GA itself *can and do* directly influence genotype fitness (consider, for example, the co-evolution of pursuit and evasion strategies in Miller & Cliff, 1994, also the prisoner’s dilemma experiments of chapter 3). In such cases, neither the outcome nor any process-outcome interactions can be held independent of the process itself; it is no longer a safe assumption that the GA is doing only optimisation. Indeed, historical constraints can arise directly from the (process) mechanics of a GA and can influence its dynamics in ways not necessarily associated with the dynamics of any actual behaviours, even if the specification of behaviour is still the outcome (see, for example, Di Paolo, 1999, ch.5).<sup>8</sup> Such models exceed the minimum specification of IOS models (laid down in section 5.2) in that the ‘evolutionary connotations’ of the GA cannot now be dis-

<sup>7</sup>Similar issues can arise when dynamic optimisation techniques are applied within an orthodox framework (Mangel & Clark, 1988). However, the rich potential for behavioural historical constraints in IOS models (section 5.3.3) justifies the present emphasis. For other perspectives on the interaction of behaviour with the dynamics of evolution see Waddington (1942), Bateson (1963), Piaget (1971) and the collected papers in Belew and Mitchell (1996).

<sup>8</sup>It is this historical potential - whether behavioural as in Miller and Cliff (1994), or attached to GA mechanics as in Di Paolo (1999) - that partly distinguishes such models from equational game-theoretic models, in which optimisation is frequency dependent but in which optimal solutions (now described as ‘evolutionary stable strategies’) are still considered independent of the mechanics of the (optimisation) process (Maynard Smith, 1982). Note that the prisoner’s dilemma models of chapter 3 did *not* require analysis in terms of evolutionary historical constraints.



counted, and indeed often become prominent (in the guise of process mechanics) in explanations of model behaviour. As such, they may be thought of as ‘extended’ IOS models. At the extreme, GA models can be used to explore pure evolutionary phenomena (see, for example, Hinton & Nowlan, 1987). These models usually place little if any emphasis on the situated behaviour of individual agents and therefore have little to do with foraging theory, orthodox or otherwise. They are not IOS models.

It is important to note that *even if* the method of optimisation bears explanatory weight, this still does not mean that such models can now test whether or not animals optimise (see section 5.1). Rather, such models should be interpreted as extensions of IOS models in which aspects of the method of optimisation can constitute constraints on optimisation; such constraints can only modify the *space of generable hypotheses*. This contrasts with both orthodox OFT and ‘unextended’ IOS models (section 5.2), for which constraints attach only to outcomes or to (optimisation) process-outcome interactions; thereby *generating hypotheses*.

In all cases, the parallel with OFT requires an emphasis on outcome (optimal behaviour) over process. However, whereas orthodox models encourage a conceptual separation of process from outcome, IOS models can assess the impact of this separation, in particular through their ability to elucidate behavioural historical constraints, which may interact with the dynamics of optimisation. And ‘extended’ IOS models, although less related to orthodox OFT models through their involvement with ‘evolutionary’ mechanics, permit the consideration of a suite of phenomena concerning relationships between optimisation, evolution, and adaptive behaviour. However these are issues beyond the remit of the present discussion, and for what remains of this chapter IOS models should be understood as ‘unextended’.

### 5.3.6 Proxies for reproductive fitness

The practice of adopting a proxy for reproductive fitness is widely appreciated as troublesome (Pierce & Ollason, 1987; Dawkins, 1995, see also section 5.1.1). IOS models are at least explicit in directly associating (in the GA) the fitness function (short-term optimality) with reproductive fitness (long-term optimality). There is also the potential for IOS models to employ progressively general fitness functions, ultimately to the extent that no explicit fitness function need be stated at all, with differential reproductive success arising solely out of the interplay of agents and their environments (see, for example, Todd & Yanco, 1996). Thus, at least in principle, IOS models afford a means of evaluating the consequences of making assumptions about how long-term and short-term optimality relate.<sup>9</sup>

## 5.4 IOS as OFT: a second look

It is therefore possible to consider IOS models as OFT models in which the constituent assumptions are dramatically restructured. Decision variables become optional, and constraints can be articulated in terms of sensorimotor processes and spatiotemporally defined and structured simulated environments. Agent *Umwelts* can be constructed, and the pragmatic value of currencies,

<sup>9</sup>Bullock (1999) has advanced a similar claim with regard to the benefits of ‘evolutionary simulation modelling’, a related methodology that may be distinguished from IOS modelling by (i) an equal emphasis on process and outcome, and (ii) no insistence on modelling situated perception and action.

and indeed of optimisation itself, can be interrogated.

As with orthodox OFT models, and as with evolutionary psychology, IOS models can be directed towards both functional and mechanistic modes of inquiry, engaging with both well adapted and apparently suboptimal patterns of behaviour. Functional hypotheses for observed patterns of behaviour can be derived from the fitness function, given the nature of the constraints. Insight into mechanistic structures potentially able to underlie observed behaviour patterns can be gleaned from a knowledge of the intrinsic constraints and their (possibly historical) interplay with extrinsic constraints, abetted by an understanding of the process of constitutive construction. In particular, IOS models allow for the functional potential of *simple* mechanisms to be appreciated through (i) abandoning the commitment to decision variables and (ii) distributing the task of explanation across agent, environment, and time. More generally, IOS models allow the relations between behaviour and mechanism to be the object of study rather than assumption.

The IOS models described in the remainder of this dissertation exploit - between them - all of these features. They are also representative of three further profound differences between the IOS methodology and orthodox OFT, described below.

First, orthodox models are usually of ‘incremental complexity’ (see section 5.1). Predictive failures are normally attributed to an inadequate capture of pertinent constraints (or, more radically, currency), prompting a revision of the model. The process ends (if ever) when sufficient constraints have been incorporated such that there is no significant difference between model and real-world behaviour. There is, however, an inconsistency in this approach. The incremental addition of constraints, in orthodox OFT, is always in the context of the significant set of framework assumptions concerning decision variables, extrinsic versus intrinsic constraints, and so on. These assumptions are already so strong that it is not clear that *incremental* revision of models in such a context is always going to be appropriate. Of course in many cases progress can be made in just this way, as the continuing prosperity of orthodox methods attests. But it may be, in other cases, that the aspects of model behaviour that prompt revision are in fact consequences of the framework assumptions themselves, and not of inadequate constraint capture. In IOS modelling, these framework assumptions can be relaxed; and so, rather than generating predictive hypotheses of incremental verisimilitude, IOS models can bring into focus aspects of orthodox models that are either explicit but potentially unnecessary (e.g. decision variables) or implicit and usually ignored (e.g. situated perception and action). Unlike orthodox models, IOS models need not operate in a mode of incremental complexity.

Second, individual-based models in general (of which IOS models are a subset) often present ‘explanatory opacity’ (Di Paolo, Noble, & Bullock, 2000; Grimm, 1999). Even relatively simple IOS models will usually be of sufficiently complex constitution that their operational description is not obvious. In general, work has to be done to understand the behaviour of the model itself. This requires the formulation and assessment of hypotheses - not about the real world - but about how the various structures instantiated by the model entail its behaviour. This may be particularly so, and the process of investigation particularly valuable, when historical constraints are prominent, or when aspects of process mechanics figure in operational descriptions of ‘extended’ IOS models (see section 5.2). This contrasts with orthodox OFT models which, given sufficient mathematical fluency on the part of the researcher or student, do not hide the means by which model structure

entails behaviour.

Third, individual-based models (and therefore also IOS models) have occasionally been touted as offering greater explicitness and clarity with regard to their assumption structure than mathematical formalisations (Miller, 1995). However, it may equally be argued that individual-based models run a *greater* risk of artefactual results than mathematical models because of the large array of components and potentially manipulable parameter settings of which they normally comprise. Assumptions in mathematical models, although they may sometimes be more theoretically pernicious than those of their individual-based counterparts, at least have the virtue of being explicitly stated.

A further general criticism of the IOS modelling approach, when held in comparison with orthodox OFT, may be levelled at its apparent exploratory character; it may seem that IOS models are best suited to situations in which coherent hypotheses, framed in terms of a well-defined solution space, have *not* been formulated. However this would not be entirely fair. In principle, IOS models are equally amenable to the formulation of explicit hypotheses, it is just that hypotheses particular to IOS models will engage with issues of historical constraints, of sensorimotor interactions (and so on), and as yet there is little precedent for concocting such hypotheses. Therefore IOS modelling may well appear exploratory (not a bad thing in itself), but this may be a consequence of novelty rather than inherent disposition and ought not to be assumed to be a permanent state of affairs.

## 5.5 Related Work

This section discusses how IOS models relate to three distinct bodies of work; individual-based models in ecology, ostensibly similar models in SAB/AL, and general notions of how SAB/AL models can contribute to biology.

### 5.5.1 Individual-based models in ecology

The previous section began to consider some properties of IOS models (explanatory opacity, a questionable explicitness of assumptions) that inherit from the larger class of individual-based models, and some of the benefits offered by IOS models identified earlier also attach, to some extent, to this larger class (historical constraints, and the relaxation of constraints in general). The fundamental distinguishing features nevertheless remain. Recent review papers indicate that individual-based models in ecology do not generally incorporate optimality (that is to say individual structure is not the result of any optimisation process), nor do they operate at the level of situated perception and action (Grimm, 1999; Durrett & Levin, 1994; Giske, Huse, & Fiksen, 1998). Thus they differ in both regards from IOS models. Consider, as a representative case study, the recent history of modelling the optimal distribution of foragers in patchy environments. Equational models have been applied to this problem ever since the seminal efforts of Fretwell (1972); such models have been reasonably successful in predicting what optimal distributions should look like, but have offered little insight into the nature of the mechanisms by which such distributions might be arrived at. A subsequent wave of individual-based models (see, for example, Moody & Houston, 1995) explored the relationships between individual behaviour and population distribution, but such models have been criticised for not incorporating optimality and thus being inconsistent

with the foundations of Fretwell's original intuition (Van der Meer & Ens, 1997). Most recently, Stillman, Goss-Custard, and Caldwor (1997) attempt to model individuals that make optimal decisions, but their mechanisms of decision-making are complex and there is still no engagement with situated perception and action. In a related example, Bernstein, Kacelnik, and Krebs (1991) analyse an individual-based model in which agents inhabit a spatially structured environment, but in which their behaviour is not mediated by any explicit perception of this structure. A fair summary would be that, although biologists are well aware of the need to reconcile individual-based models with optimality, and also to account for the effects of situated perception and action, principled strategies for achieving these aims are still lacking. I believe that IOS models are well placed to step into the breach, and examples of IOS models which tackle some of these issues directly appear in chapters 7 and 8 of this dissertation.

### 5.5.2 Individual-based models in SAB/AL

Although a good deal of SAB/AL work falls under the individual-based model rubric, only a small proportion engages with the same issues as OFT. One positive example is provided by Spier and McFarland (1998) who use situated individual-based models to compare the merits of several foraging strategies. However, although their strategies "hav[e] claim to perform optimally when analysed mathematically", their point is exactly that such analysis takes place *outside* the context of a situated individual-based model, and they explore how such supposedly 'optimal' strategies fare in situated environments. This work may be taken to illustrate the need for IOS models which, by contrast, allow the exploration of strategies that *are* optimal in the context of situated environments.

Many of the individual-based models that populate the SAB/AL literature and that also use GAs are concerned with issues of evolutionary dynamics (Hinton & Nowlan, 1987), and to the extent that they are, their aims should be distinguished from those of IOS models, which are directed squarely at the mechanistic and functional aspects of behaviour. Recall from section 5.3.5 that IOS models are characterised by their emphasis on outcome over process and on the independence of outcome from the mechanics of the optimisation process itself. Of course there do exist many examples of models in the SAB/AL canon that *can* be interpreted, in both substance and emphasis, as IOS models (see, for example, Koza, Rice, & Roughgarden, 1992, and, of course, the models of the following chapters), however a detailed survey is beyond the present remit.

### 5.5.3 SAB/AL and theoretical biology

Finally, it is important to locate the relationship between IOS and orthodox OFT models in the context of general debate over the nature of *possible* methodological relationships between SAB/AL models and theoretical biology (both broadly construed).

The incremental complexity of orthodox OFT resonates with the 'virtual biology' approach of Kitano, Hamahashi, Kitazawa, Takao, and Imai (1997) who, although not concerned with optimal strategies of any sort, attempt to introduce as much detail as possible into their models such that there is no (significant) difference between real-world and model behaviour. They are interested in specific biological processes, for example the genetic and biochemical processes underlying the development of the fruit fly *Drosophila melanogaster*. Their unreserved commitment to verisimil-

itude encourages them to assert that divergences between model behaviour and actual *Drosophila* development can point to extra factors at play which have yet to be identified by the biologists themselves. Of course the validity of this assertion depends entirely on the theoretical and informational fidelity of their model to the biological state of the art. This approach exemplifies one extreme of the possible relationships between SAB/AL and biology. A weaker version of this manifesto is outlined by Maley (1999), who proposes a cycle of activity again similar to OFT in the construction, evaluation, and incremental revision of models. However, Maley does not insist upon the same degree of fidelity as Kitano et al., a trade-off which affords greater modelling scope but detracts from the confidence with which specific predictions about empirical biology can be formulated.

The other extreme, by tradition associated with Christopher Langton (see, for example, Langton, 1989), considers that SAB/AL creations provide *instantiations* and not just *models* of biological phenomena, such that observations of their behaviour should be accorded the same status as observations of similar phenomena in the natural world. This ‘strong AL’ position has little contact with the methodological stance of the present enterprise. Perhaps the only direction it provides is in emphasising the utility of analysing behaviours in IOS models that demonstrably do not correspond to observations of the world. However such analysis should be used only to further an operational understanding of the model itself, and should have nothing to do with the postulation of additional empirical data. (Recall from chapter 1 that the distinction between instantiation and simulation was appreciated very early in the history of the use of artefacts to understand biological phenomena; I refer of course to the eighteenth century automata maker Jacques de Vaucanson, who came down - as does the present dissertation - very much on the side of simulation.)

Most consonant with the approach as outlined in these pages is the notion of the ‘opaque thought experiment’. Di Paolo et al. (2000) develop the argument that simulation models in SAB/AL provide a way of “re-organising and probing the internal consistency of a theoretical position” such that “theoretical terms may be shown to stand in different relationships than previously thought.” This position does not locate the benefits of SAB/AL in furthering the verisimilitude of models, nor does it interpret these models as ‘instantiations’ of phenomena. Instead, “the researcher may be forced to focus on facts or processes that were at the periphery of her conceptual structure and place them in novel relationships with other theoretical terms” - but only after some work has been done to overcome the explanatory opacity of simulation models. IOS models fit in well with this conception, harnessing the versatility of modelling at the level of perception and action to restructure the conceptual apparatus of the researcher with particular regard to the framework assumptions of orthodox OFT.<sup>10</sup>

## 5.6 Summary

The methodological groundwork laid in this chapter derives from an extended analogy between orthodox OFT and the application of GAs in situated, individual-based SAB/AL models. This

---

<sup>10</sup>Bedau (1999) has explicated a related idea in which his ‘emergent, computational thought experiments’ are considered to instantiate target phenomena. However, Bedau restricts his interest to ‘deep’ phenomena, for example ‘multi-level emergent activity’, ‘open-ended adaptive evolution’, and ‘unbounded complexity/diversity growth’. Such ‘deep’ phenomena are well insulated from observation in the real world, allowing the distinction between instantiation and model to be somewhat glossed over.

approach highlights the difficulties associated with reconciling the incremental complexity ethos of orthodox OFT with the strong framework assumptions that partially constitute all orthodox models. IOS models allow these framework assumptions to be radically restructured, with decision variables, historical constraints, ecological environments, and the pragmatics of optimisation all receiving attention. These models can be characterised as varieties of thought experiment, such that any knowledge gained concerns the researcher's own conceptual space, the reorganisation of which can encourage the formulation of novel empirical hypotheses. Emphasis is placed on the conceptual leverage obtained by modelling subtle agent-environment interactions mediated by perception and action, and in particular, attention is drawn to how IOS models can illuminate the functional potential of relatively simple internal behaviour-generating mechanisms.

IOS modelling is closely related to the version of the ECT developed in chapter 2, and indeed the former may be considered to be a methodological expression of the latter. They are both strongly externalist, they both emphasise the importance of understanding the process of constitutive construction, and each can be directed towards both functional and mechanistic inquiry, and above all to the relations between these two levels of description. The most salient difference is that the ECT is formulated at a relatively abstract level and is largely concerned with complexity, whereas IOS models are concrete projects designed to engage with the subject matter of orthodox OFT. Also, the IOS methodology encourages a strict interpretation of the role of GAs in terms of optimisation, in contrast to the loose analogy between artificial and biological evolution sanctioned by the relatively abstract ECT.

This chapter has served primarily to lay the necessary methodological groundwork for the remainder of this dissertation, and the following chapter complements this groundwork with a discussion of the theoretical context within which the IOS methodology will be applied. This context is best described by the term 'behaviour coordination', and subsequent chapters explore its manifestation in the particular situations of interference amongst groups of foragers, the 'matching law' of experimental psychology (which pertains to individual choice behaviour), and the optimal distribution of groups of foragers over 'patchy' environments.

This chapter has also elaborated a specific and practical programme of research capable of contributing to the general maturation of the field of SAB/AL, a programme that may - and ought to - extend beyond the confines of the present dissertation. (Although the investigations mentioned above are certainly representative of the IOS methodology, in no sense do they constitute an exhaustive exploration of its potential.) From the perspective of theoretical biology, IOS models can be seen as a new generation of individual-based model in which both optimality and situated perception and action adopt fundamental roles, and from the vantage of SAB/AL, IOS models offer a principled strategy for engaging with issues of real currency in the natural sciences.

## Chapter 6

### Behaviour coordination

---

With the methodological framework of the previous chapter in place, the purpose of this chapter is to develop the theoretical context for the remainder of this dissertation. The focus will be on *behaviour coordination*, a topic integral to much of biology and psychology, and also to SAB/AL itself. At its most general, the problem of behaviour coordination is the problem of ‘how to do the right thing’, and it fits into the general structure of this dissertation in that (a) the coordination of behaviour may be considered to be a response to a challenging environment (*sensu* Godfrey-Smith’s ECT), and (b) much of the relevant discourse concerns the nature of the agent-side mechanisms underlying the observed (or desired, in the case of autonomous agent construction) coordination. The latter, as it will be seen, is closely related to the concept in optimal foraging theory of a ‘decision variable’.

This chapter is structured into two parts; first, a brief tour of the relevant theoretical terrain together with a conceptual critique of this terrain in light of the themes of this dissertation introduced in chapter 2, and second, a simple model - constructed in line with the methodology of the previous chapter - exploring the properties of a minimal mechanism for behaviour coordination, and constituting an empirical response to the foregoing conceptual critique. In particular, it will be demonstrated that effective behaviour coordination can be achieved *without* explicit arbitration between internalised behaviour correlates. This model introduces a novel combination of evolutionary algorithms and Braitenberg-style architectures (Braitenberg, 1984) which will be extended in various ways in the remainder of the dissertation; these later models will connect closely with the subject matter of optimal foraging theory as it relates to behaviour coordination.

#### 6.1 The theoretical terrain

In the field of SAB/AL the problem of behaviour coordination is usually referred to by the term ‘action selection’, a standard interpretation of which is given by Pattie Maes:

Given an agent that has multiple time-varying goals, a repertoire of actions that can be performed [...] and specific sensor data, what actions should this agent take next so as to optimise the achievement of its goals? (1994, pp.15-16)

Action selection - behaviour coordination - can therefore be treated both as an engineering problem and as a theoretical challenge. The following discussion adopts the second perspective, and in doing so covers a diverse theoretical terrain before returning to question Maes' conception of the problem.

### 6.1.1 Hierarchical control

In a seminal paper, Rosenbleuth, Wiener, and Bigelow (1943) identified action selection with the problem of how to plan sequences of actions in order to fulfil pre-specified and functionally defined goals. Their solution involved the cybernetic principle of negative feedback, such that the difference between the state of an agent at any given time and the 'goal-state' determines the appropriate behaviour. Miller et al. (1960) gave substance to this proposal in the form of the TOTE (test-operate-test-exit) unit, which tracks proximity to a goal-state, and performs a behaviour ('operates') until the goal-state is reached ('exit'). Complex behaviour coordination is achieved through the operation of *hierarchies* of TOTE units, and they argued that all behaviour could be understood in these terms: "The reflex is not the unit we should use as the element of behavior: the unit should be the feedback loop itself" (p.27) and "the underlying structure that organizes and coordinates behavior is [...] hierarchical" (p.34).<sup>1</sup> This conception of behaviour coordination has underpinned much theoretical and empirical work in cognitive psychology, and, as McFarland and Bösser (1993) note, has "persisted in artificial intelligence and in robotics up to the present day" (p.viii). As such, it invites criticism from all sorts of angles, perhaps most obviously in view of the internal mechanistic significance attributed to (functionally specified) 'goals'. However, even if this confusion of functional and mechanistic discourse is overlooked, Hinde and Stevenson (1970) note that there are many instances of behaviour which *cannot* be understood in terms of goal-directed negative feedback at all, one example being the final strike behaviour in the catching of prey in many animals; an extremely rapid behaviour not subject to error-correction of any kind. Moreover, in an earlier paper, Hinde (1957) questioned the validity of framing a useful definition of a 'goal-state' in the first place, given the variety of factors that can influence the cessation of any given activity (see also Toates & Jensen, 1990). The ethologist Niko Tinbergen was also suspicious of the mechanistic reification of goals:

There has been, and still is, a certain tendency to answer the causal question by merely pointing to the goal, end, or purpose of behaviour, or of any life process. This tendency is, in my opinion, seriously hampering the progress of ethology. (1951, p.4)

Tinbergen's own conception of hierarchical behaviour coordination rejected the idea that animals have explicit knowledge of their 'goals', and appealed to natural selection to explain the apparent purpose of observed behaviour, with environmental 'sign stimuli' releasing activity in 'behaviour centres' at a succession of hierarchical levels (Tinbergen, 1950, 1951). For Tinbergen, natural selection ensured that activity at any level of the hierarchy would (eventually) lead the animal to encounter the 'sign stimuli' necessary to release behaviours at the next level down.

---

<sup>1</sup>Miller et al. were motivated by their dissatisfaction with behaviourist learning theory, of which the 'reflex' - or stimulus-response connection - is a central concept (see chapter 2). A similar dissatisfaction motivated Lashley (1951) who argued that some rapid action sequences, such as those performed by accomplished pianists, are simply too fast to be explained in terms of the behaviouristic 'chaining' of reflexes, in which each reflex is triggered by its predecessor; Lashley's own models of the generation of action sequences were also hierarchical.



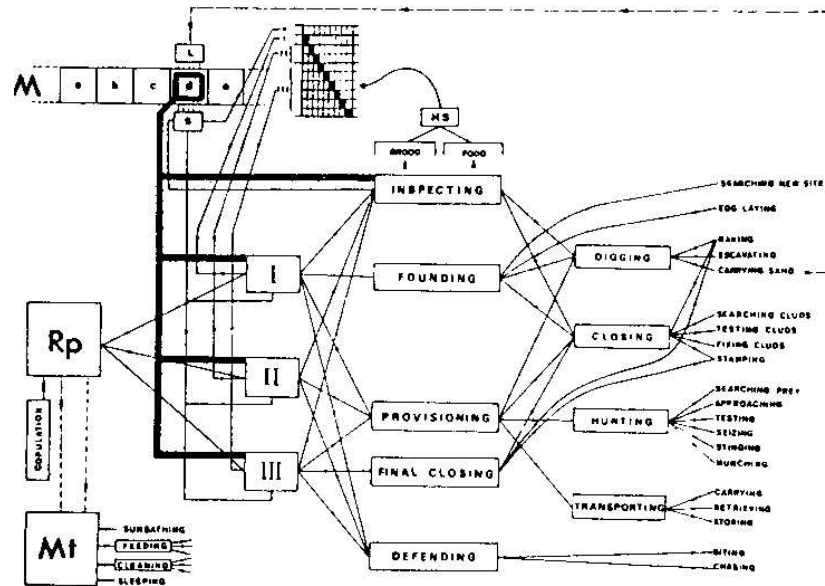


Figure 6.1: Baerends' 'hierarchy of motivations' (from Baerends, 1976).

Although this model has no need of 'goals', Tinbergen could only explain how his hierarchy actually generated behaviour by postulating a form of 'action-specific energy' which would flow downwards through the hierarchy when released by the appropriate 'sign stimuli', activating the various 'behaviour centres' on its way. Many of the subsequent criticisms of Tinbergen's model focussed on the scientific untenability of this 'energy' (Hinde, 1956); the notion of hierarchical control, however, persisted.<sup>2</sup> In ethology, Baerends (1976) proposed a 'hierarchy of motivations' in which 'action-specific energy' is clothed in a relatively respectable style as 'motivational state' (figure 6.1), and Dawkins (1976) developed a 'software model' in which higher hierarchical levels 'call' lower levels as if they were program subroutines; hierarchical models of action selection have also become common in the SAB/AL canon (Rosenblatt & Payton, 1989; Tyrrell, 1993; Blumberg, 1994; Blumberg, Todd, & Maes, 1996).

The theoretical status of hierarchical control remains in dispute. Substantial criticism continues to be levelled at the requirement - entailed by the orthodox cognitive structure of these architectures - that all situations and responses be entirely pre-specified, leaving no space for the emergence of behavioural coordination from agent-environment interaction patterns (Steels, 1990; Hendriks-Jansen, 1996). However, it is neither easy nor necessarily advisable to discard the arguments of the early ethologists, that the complexity of natural behaviour is simply too great to be explained in any other way. I will not attempt to resolve, or even fully describe this debate here, and although the simple model described in the second half of this chapter is non-hierarchical, this should not be taken to be a criticism of hierarchical control *per se*.

### 6.1.2 Non-hierarchical control

There are alternatives to hierarchical control. Writing before Tinbergen, Lorenz (1935, 1937, 1939) conceived of behaviour coordination in terms of 'fixed action patterns' - temporally ex-

<sup>2</sup>See Hendriks-Jansen (1996), chapters 11-13, for further discussion of Hinde's position with respect to both 'action-specific' energy and goal-directed behaviour.

tended patterns of activity such as the egg retrieval behaviour of mother geese (Lorenz & Tinbergen, 1938) - released by a combination of environmental sign stimuli and internal 'action-specific energy'.<sup>3</sup> However, this conception does little to explain the *coordination* of behaviour - there is an absence of hierarchy rather than an alternative - and the notion of 'energy' (a concept Tinbergen in fact inherited from Lorenz) again constitutes a theoretical weakness.

More recently, the 'subsumption architecture' of Rodney Brooks has achieved prominence as a non-hierarchical alternative to the problem of behaviour coordination, expressed not in ethology, but in the medium of 'behaviour-based robotics' (Brooks, 1986, 1994; Arkin, 1998). Subsumption involves the decomposition of an agent's control structure into a set of task-achieving 'competences' which are organised into layers, with competences at lower layers (for example 'avoid obstacles') able to override - or 'subsume' - those at higher layers (for example 'explore'). Although Brooks developed the subsumption architecture primarily from the perspective of engineering, an evolutionary rationale can be suggested insofar as lower layers may be considered to represent earlier phylogenetic stages of evolution; Prescott, Redgrave, and Gurney (1999) take this idea to an extreme, interpreting the layered neural architecture of the rat basal ganglia as a subsumption-esque action selection mechanism (see also Redgrave, Prescott, & Gurney, 1999). In practical terms, Brooks and his colleagues certainly enjoyed some early successes (Brooks, 1989; Ferrell, 1994), however, their ambitious project to construct the torso of a humanoid robot with a repertoire of extremely sophisticated competences (including, for example, 'mental rehearsal', 'body mimicry', and 'proto-language', according to Brooks & Stein, 1993) remains unfulfilled (although see Marjanovic, Scassellati, & Williamson, 1996, for a description of the development of simple sensorimotor control of the robot).

In an influential paper, Maes (1990) considers a distributed 'bottom-up' architecture for action selection, in which "a creature is viewed as consisting of a set of behaviors" and in which "behavior selection emerges in a distributed fashion by parallel local interactions among the behaviors and between behaviors and the environment" (pp.238-239).<sup>4</sup> The 'behaviours' to which Maes refers include 'avoid-obstacle', 'explore', and 'flee-from-creature' among others, and are connected through a carefully pre-specified network of 'predecessor', 'successor', and 'conflicter' links, each of which specifies how the activity of a focal behavioural node will influence the activity of those to which it is connected (figure 6.2). At any moment in time, only a single behaviour will have control over the motor outputs of the agent. Maes observes some interesting ethological phenomena in simulated agents controlled by such networks, for example instances of 'displacement behaviour' in which the most relevant behaviours eliminate each other. However, Tyrrell (1993, 1994) argues - on the basis of a series of comparisons carried out in simulated environments - that Maes' architecture suffers from many shortcomings, some of which are alleviated by the use of Tyrrell's own preferred 'free-flow' hierarchy in which higher levels of the hierarchy express 'preferences' rather than - as is the case for the hierarchies of Tinbergen *et seq.* - imposing strict 'winner-take-all' constraints on the activity of lower levels.

<sup>3</sup>See Halperin (1990) for a SAB/AL model faithful to the fixed action pattern framework.

<sup>4</sup>Maes cites Minsky's (1986) 'society of mind' theory as a significant influence on her model. See also Sahota (1994) for a similarly-inspired action selection mechanisms based on 'inter-behaviour bidding'.




Figure 6.2: Maes' 'bottom-up' action selection architecture (from Maes, 1990).

### 6.1.3 Behavioural economics

A different description of the problem of behaviour coordination can be found in the 'behavioural economics' of McFarland (see, for example, 1996). This is the idea that agents will, in virtue of the action of natural selection, allocate their behaviour in such a way to maximise 'utility', where utility is an economic concept usually - and perhaps unhelpfully - defined *a posteriori* as 'that which is maximised'. Insofar as utility is equated with fitness, and this is indeed McFarland's stance, behavioural economics can be identified with optimal foraging theory, and is therefore closely related to the methodological perspective of this dissertation. However, as argued in some detail in chapter 2, criteria of optimality apply only to functional entities, and the relations between functional and mechanistic levels of description is not simple. McFarland himself has, on occasion, taken what might be an overly strong position, for example: "Such [optimality] considerations are purely functional. They specify what animals ought to do to make the best decisions under particular circumstances, but they do not say anything about the mechanisms that animals might employ to attain these objectives" (1985, p.456).

Nevertheless, one mechanistic possibility explored by McFarland and Sibly (1972, 1975) is that agents maximise their utility by performing simple 'hill-climbing' over a landscape of internal signals (motivations) and external signals (environmental cues).<sup>5</sup> A simple example is given by an agent that must divide its time between eating and drinking; an instance of the so-called 'two resource' problem (Toates, 1982, 1986). Sibly and McFarland (1976) derive a decision rule for such an agent, that operates over the levels of hunger and thirst of the agent (the motivations), and the relative availability of food and water (the environmental cues), and that minimises a cost function that is also framed in these terms.<sup>6</sup> Recent extensions suggest that instances of 'planning' can

<sup>5</sup>See also Townsend and Busemeyer (1995) for a dynamical model directed towards the extension of utility-based models of decision making, with particular concern for the temporal dynamics of decision making (vacillation, indecisiveness, inconsistency of preferences, etc.).

<sup>6</sup>A cost function is the economic equivalent of a fitness function, but whereas fitness is usually maximised, cost is usually minimised.

be obtained in this way from agents controlled by mechanisms that react only to stimuli immediately available in the environment (Spier, 1997; Spier & McFarland, 1998). From the perspective of this dissertation, a problem with this kind of model is that it requires that the (mechanistic) elements of the decision rule are directly associated with behavioural entities (eating, drinking). As Hendriks-Jansen (1996) puts it in his general discussion of behavioural economics: “The fact that cost-benefit bird (an actuarial abstraction like ‘economic man’) necessarily optimizes or maximizes its utility [...] does not imply that individual birds contain choice mechanisms that work by optimization on entities like foraging and incubation” (p.128). McFarland would presumably not disagree with these sentiments, but - as was argued to be the case with respect to the decision variable in optimal foraging theory - the methodological and conceptual tools of behavioural economics do not provide any way of challenging such assumptions.

#### 6.1.4 Behaviour and mechanism, perception and action

The above description of behavioural economics highlights what is a common factor to all of the approaches discussed so far, that they each assume a need for an internal arbitration mechanism operating over a pre-existing repertoire of behavioural options. Such an assumption requires that there be internal correlates to behaviours in exactly the sense disputed in chapter 2; recall the leitmotif that behavioural and mechanistic levels of description should not be conflated. Just as behaviours themselves are the product of the joint activity of agent, environment, and observer, so the apparent interstices between behaviours should also be considered to be products of this joint activity.<sup>7</sup>

A related difficulty with many of the approaches discussed so far is their unadventurous conception of the intimate relation between perception and action. Recall from chapter 2 the argument that any attempt to internalise entire behaviours necessarily obscures the process of agent-environment interaction upon which the close coupling of perception and action depends. Given the focus on observable actions or behaviours in models of behaviour coordination in SAB/AL, it is not surprising to discover that the role of perception is often significantly underplayed; in the model of Maes (1990), perceptual variables are constrained to hold only predicate values, and Tyrrell (1993), in a description of the model in which the comparisons alluded to in section 6.1.2 took place, writes that “[i]t was decided that the small amount of added realism [from considering the role of perception] was not worth the large amount of time it would have taken to include”. The conceptual consequences of such a position are well illustrated by the venerable ‘lens model’ of Brunswik (1952, see also Keijser, 1998) in which the problem of behaviour coordination is described as the maintenance of stable relations in what is called the ‘distal’ environment, achieved through the operation of ‘proximal’ sensorimotor couplings (figure 6.3). Recalling chapter 2, ‘distal’ features are non-local to the agent, and may be considered to inhabit the external environment (the environment as it appears to an external observer), and the ‘proximal’ stimulation may be considered to reflect features of the agent’s *Umwelt*. Brunswik interprets distal ‘stability’ homeostatically, for example in terms of the maintenance of a minimum distance from predators or of a regular proximity to food; organisms use proximal stimulation to guide proximal actions - the

<sup>7</sup>See Kelso (1995) for an example of clean switching patterns in human behaviour without the need for identifiable switching elements. At the opposite end of the scale, it is interesting to note that Tinbergen fully expected to find neural correlates for his hypothetical ‘behaviour centres’.

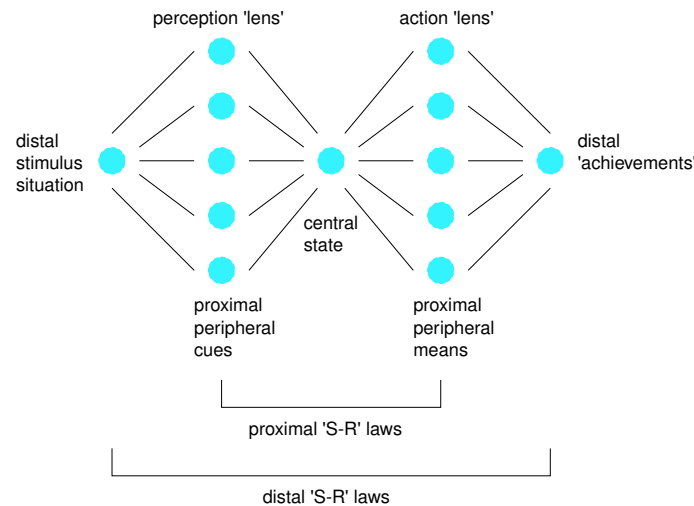


Figure 6.3: Brunswik's lens model (adapted from Brunswik, 1952, and Keijcer, 1998). 'S-R' laws refer to 'stimulus-response' laws.

rotation of wheels or the movement of limbs - in the mediation of this stability. It is abundantly clear from this model that the role of perception cannot be summarily excluded from any account of behaviour coordination. (The postulation of stimulus-response laws, an indication of the vintage of the model, is not essential to its coherence. It is quite possible to replace such laws with the idea of continuous agent-environment interaction, with responses 'feeding back' to stimuli; its relevance is unaffected.)

But what about the constitution of the 'lenses' themselves, or of the 'central state'? Brunswik intended each lens to represent the potential many-to-one arrangement between proximal relations and distal situations, suggesting that distal situational stability is mediated by a constantly changing and highly context dependent - yet somehow 'focussed' - set of proximal relations. However, in the present theoretical context this picture affords an alternative interpretation. The 'action' lens may be considered to represent a repertoire of executable 'actions', and the 'perception' lens may likewise represent the features of the *Umwelt* extracted from the distal (external) environment by the processes of constitutive construction (see chapter 2). Then it is easy to imagine a process of 'action selection' choosing the most appropriate action, given the distal situation, and a corresponding process of 'selective attention' by which the relevant features of the *Umwelt* are picked out to influence the generation of behaviour. However, not only does this interpretation leave the constitution of the 'central state' unspecified, but it also introduces an unwarranted ontological division between perception and action. From the perspective of this dissertation, a more congenial solution would be to dispense with the lenses - and the 'central state' - altogether, and to understand both action selection and selective attention as different descriptions of an underlying process of sensorimotor interaction with the environment.<sup>8</sup>

<sup>8</sup>Pfeifer (1996) provides an alternative expression of these arguments in the form of two of his principles for building autonomous agents; the principle of 'parallel, loosely coupled processes', and the principle of 'sensory-motor coordination' (see also Pfeifer & Scheier, 1999). A powerful demonstration of sensory-motor coordination - of the importance of the unity of perception and action - can be found in the categorisation of objects by mobile robots in Scheier and Pfeifer (1995). See also the model of housefly navigation by exploiting motion parallax in Franceschini, Pichon, and Blanes (1992).

The use of the term ‘selective attention’ is not incidental. It is normally employed in connection with a large body of work in cognitive psychology, cognitive neuropsychology, and more recently in artificial intelligence, that deals with exactly the process of establishing how certain features of the perceptual world of organisms come to have explanatory significance in the generation of behaviour (Pashler, 1999; Parasuraman, 1998; Tsotsos, Culhane, Wai, Lai, Davis, & Nuflo, 1995). Of particular interest is that most accounts of selective attention leave the role of *action* relatively unspecified (in many experiments with human subjects, for example, the subject is required only to press either button *A* or button *B*; see, for example, Posner, 1978, 1980), and also that many proposed mechanisms for selective attention are hierarchical (Treisman & Gelade, 1980; Moran & Desimone, 1985; Tsotsos et al., 1995). A disturbing symmetry is therefore evident between orthodox conceptions of action selection and selective attention; they both underplay the significance of their opposite number, and both have utilised hierarchies to trace the path from an array of potentialities to an isolated, and causally privileged state. At worst, a picture arises in which attention involves (hierarchically) isolating some ‘central percept’ from a varied *Umwelt*, which is then somehow transformed into an internal behavioural correlate, which is then expressed, in action, via an action selection hierarchy. This is a picture with strong cognitivist overtones; there is a sensory hierarchy, some central ‘magic’, followed by a motor hierarchy, and in such a picture the central wizardry must necessarily disturb the intimate coupling of perception and action. To some readers this may appear to be a wilful misconstrual of the well established problems of action selection and selective attention, but my purpose is not to deny that these problems exist, nor to denigrate the invaluable experimental work that has been performed under their auspices (especially in the area of selective attention), but merely to suggest these problems are conceptually inconsistent with each other to the extent that they are considered in isolation from an ongoing process of perception and action, and that the various notions of ‘central states’ and internal behavioural correlates may to some extent be symptoms of this inconsistency.<sup>9</sup>

### 6.1.5 Summary

Many conceptions of the problem of behaviour coordination, and many of the proposed solutions, are inconsistent with the themes of this dissertation. Perhaps the most obvious conflict is with the idea that (functional) goals can have causal, mechanistic influence. However, both the hierarchical responses to this problem of Tinbergen *et seq.*, and the non-hierarchical models of Brooks (1986) and Maes (1990) continue to confound behavioural and mechanistic levels of description, proposing the existence of internal behavioural correlates and supervenient arbitration mechanisms. Moreover, insofar as these accounts confer explanatory precedence on internal properties of the agent, they are inconsistent with the externalist perspective of the version of the ECT outlined in chapter 2. Internalism may appear to attach more readily to hierarchical accounts than to non-hierarchical accounts, in that hierarchies explicitly locate the sources of organisation inside the agent, but it has not been my intention in this discussion to criticise hierarchical theories of control in view of their hierarchical nature. It is true that, in the examples discussed above, hierarchical models do tend to exacerbate the conceptual problems associated with the various relations

<sup>9</sup>There is a rapidly increasing quantity of work in SAB/AL directed towards selective attention, see, for example, Scheier and Lambrinos (1994), Foner and Maes (1994), Grossberg (1995), and Tani, Yamamoto, and Nishi (1997).

between behaviour, mechanism, perception, and action, but I have *not* demonstrated that this will necessarily be the case for hierarchical control structures in general.

An alternative, externalist perspective, well aligned with both the ECT and with OFT, is provided by McFarland's 'behavioural economics'. However, as with OFT, behavioural economics as it stands provides no means of tracing the relations between behavioural and mechanistic levels of description. What is required is a version of McFarland's framework - of OFT - that *does* permit such relations to be traced, and that also respects the intimacy of perception and action. My contention is that the IOS methodology, introduced in the previous chapter, is well placed to serve this function. In what follows I describe a simple IOS model in which effective behaviour coordination is achieved by a simple situated agent, the internal mechanism of which resists decomposition in terms of internal behavioural correlates or decision variables, and which enshrines the inseparability of perception and action. This model serves as the empirical foundation for the less abstract investigations of the subsequent chapters.

To return at last to the definition of action selection offered by Maes (1994), the idea of explicit goal-directed arbitration between internal behavioural correlates may be replaced by the notion of behavioural coordination emerging from a process of continuous agent-environment interaction, mediated by tightly coupled perception and action. In this view, both action selection and selective attention are skewed descriptions of this underlying activity.


## 6.2 A simple IOS model of behaviour coordination

Not all SAB/AL models of action selection rely on internal arbitration between behavioural correlates, and not all deny the intimacy of perception and action. In chapter 4 of his 1984 volume *Vehicles*, Valentino Braitenberg describes an imaginary artefact in which a number of simple connections, linking sensors to motors, operate continuously and in parallel. He argues that such a vehicle may exhibit what to an external observer would be describable as behaviour coordination insofar as, at different times, different combinations of sensorimotor activity lead to the expression of qualitatively distinct behaviours (figure 6.4).

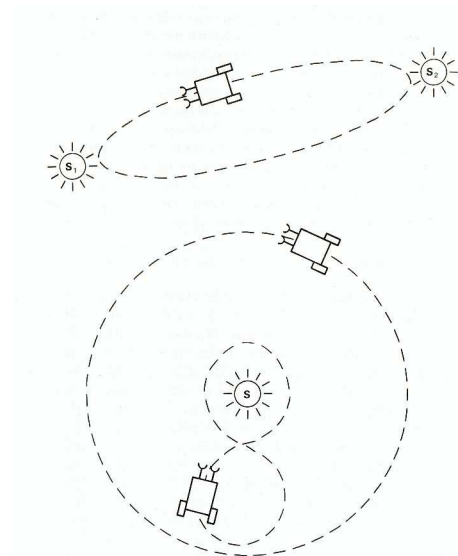
Under the banner of what they call 'extended Braitenberg architectures', Lambrinos and Scheier (1995) describe an experiment in which a Khepera mobile robot is able to achieve a successful balance between collecting pegs, returning the pegs to a 'home base', and periodically visiting a separate 'charging area' (see also Pfeifer & Scheier, 1999, and for a related investigation, Steels, 1992). This balance is achieved through coordination between several continuously active sensorimotor processes; each of which continuously influences the states of the motors. Pfeifer and Scheier (1999) draw out the significance of this:

[I]t will not be trivial to infer from observing the agent's behavior what the underlying mechanisms are, since the behavior is the result of many processes running in parallel [...] the agent's behavior is emergent from the joint activation of the processes and not determined by some sort of selection mechanism [...] [p]rocesses are not behaviors that can be chosen for execution. (p.352)

Interestingly, this kind of analysis is not restricted to the discipline of SAB/AL. Kien and Altman (1992), for example, consider insect motor systems to comprise of several sensorimotor loops acting in parallel, such that it is "the consensus of the activity in these loops that regulates behavior



(a) Architecture.



(b) Behaviour.

Figure 6.4: Braitenberg’s ‘decision making’ vehicle (from Braitenberg, 1984). The image (a) is his idea of the kinds of connections that might exist between a simple sensor and a simple motor, and (b) represents his idea of the kinds of ‘decision making’ behaviour that might result.

and provides the basis for decision making [...] decisions are an emergent property of the whole system; they are the outcome of the total activity in all the loops at any time, where each loop regulates different aspects of motor outputs and hence of behavior” (1992, pp.164-165, quoted in Pfeifer & Scheier, 1999, p.352).<sup>10</sup>

### 6.2.1 A description of the model

The IOS model described in this chapter is related to the model of Lambrinos and Scheier (1995), and comprises a novel combination of evolutionary algorithms with Braitenberg-style architectures. It involves agents that are controlled by sets of direct, autonomous, sensorimotor links, with neither interconnections between links nor intervening artificial neurons of any kind. Artificial evolution is used to specify the transfer functions supported by the links, all of which are continuously active. The purpose of this model is to illustrate how such a minimal mechanism can underlie effective behaviour coordination in a simple environment.

One difference between the model of Lambrinos and Scheier (1995) and the present model is that the former utilises sensorimotor ‘processes’, and the latter employs sensorimotor ‘links’. For Lambrinos and Scheier, the term ‘process’ is a general term for a connection between sensor and motor, which may include intermediate stages. By contrast, the term ‘link’, as employed here, signifies that the sensorimotor connections consist of nothing more than simple transfer functions, each of which underlies an instantaneous transformation of sensory input into motor output. An-

<sup>10</sup>A discussion of parallel and continuously active pathways in *perceptual* systems can be found in Milner and Goodale (1995).



other important difference between the two models is that, in the present case, the connections between sensors and motors are evolved rather than designed by hand. An advantage of this approach is that it ensures that the distinction between behaviour and mechanism is made completely explicit. For Lambrinos and Scheier, the reliance on hand-design requires the artificial separation of the problems of process design (what each individual process should do) and process fusion (how to combine these processes to deliver coherent behaviour). Once these problems are separated, there is the potential for each sensorimotor process to be considered in terms of an eventual behaviour to be subsequently ‘fused’, and indeed their processes carry the (behavioural) labels ‘go-to-station’, ‘home’, and so on. Although there is no explicit arbitration between these processes (they are all autonomous and concurrently active) there remains a certain tension between their affirmation that behavioural and mechanistic levels of description should not be confused, and this potential influence of behavioural descriptions on the design of the internal mechanism. The IOS methodology avoids this problem by allowing process (or link) design and process (or link) fusion to proceed concurrently, providing no scope for behavioural categorisation to influence the structure of internal mechanism. This is a consequence of the general point, belaboured in the previous chapter, that IOS models allow the relation between behaviour and mechanism to be the object of study rather than assumption.

There are other differences. The present model eschews the real world, locating all evolution and analysis in simulation. Also, the task environment is different and is more closely related to the two resource problem (section 6.1.3) than it is to the relatively arbitrary task chosen by Lambrinos and Scheier. The present task environment requires the agent to balance its intake of ‘food’ and ‘water’ whilst avoiding ‘traps’, allowing the performance of the agent to be evaluated on a list of desiderata for an effective action selection mechanism. The list presented below is extracted from Werner (1994), but only those items relevant to the present model are included; extended versions can be found in either Werner (1994) or Tyrrell (1993). According to this list, an effective action selection mechanism should:

- Prioritise behaviour according to current internal requirements.
- Allow contiguous behavioural sequences to be strung together.
- Exhibit opportunism; for example by diverting to a nearby food item even if there a greater immediate need for water.
- Balance dithering and persistence; for example, by drinking until full and then eating until full instead of oscillating between eating and drinking.
- Interrupt current behaviour; for example, by changing course to avoid the sudden appearance of a dangerous object.
- Privilege consummatory actions (those that are of immediate benefit to the welfare of the agent, for example eating) over appetitive actions (those that *set up* the conditions in which consummatory actions become more likely).
- Use all available information.
- Support real-valued sensors and produce directly usable outputs.




Figure 6.5: Architecture of the agent: there are 3 sensor types, with each sensor connected directly to the wheel on the same side. Each illustrated connection actually comprises 3 genetically specified links, so the agent consists of 18 concurrent and direct transformations of sensory input into motor output. The two internal ‘batteries’ are also indicated.

The simulated environment is a continuous (in space, but not in time) unbounded area within which 3 types of objects can exist - food, water, and ‘traps’ - in addition to the agent itself. Each object is a circle of radius 16.0 units (except for the agent, which is a circle of radius 5.0 units), and all objects appear within a 200.0 by 200.0 unit area of the environment (the agent is not constrained to remain within this area). There are 3 items of food, 3 of water, and 9 traps.

The architecture of the agent is inspired by Braitenberg (1984) and is illustrated in figure 6.5. The agent possesses 2 wheels, and 3 sensor pairs, with each sensor pair responding to a different object type. It also possesses 2 internal ‘batteries’ - one for food ( $\mathcal{B}_f$ ) and one for water ( $\mathcal{B}_w$ ) - which diminish at a steady rate during the ‘lifetime’ of the agent. Encounters with food or water replenish the corresponding battery, but if both reach zero, or if the agent encounters a trap, then it will ‘die’. The 3 sensor pairs respond to the distance from the agent to the nearest instance of each type of object, with each sensor ranging linearly from 100.0 (at the object) to 0.0 (200.0 or more units distant). If an object is to the left of the agent, the corresponding sensor on the left of the agent will respond with 20% greater activation (subject to the maximum output value of 100.0) than the sensor on the right, and *vice-versa* if the object is to the right of the agent.

The links between the sensors and the motors simply transform the sensor input signal (range [0.0,100.0]) into an output signal (range [-1.0,1.0]) in a manner specified by a transfer function (figure 6.6). It is the shape of this transfer function that is evolved, and it is also possible for this shape to be modified *during the lifetime of the agent* by the values of either  $\mathcal{B}_f$  or  $\mathcal{B}_w$ . The manner in which this may happen, and the details of the genetic encoding scheme, is described in section 6.2.2 below. Each connection illustrated in figure 6.6 actually stands for 3 independent (and concurrently active) such links, and a left/right symmetry is imposed such that both sets of 9 links (both sets of 3 illustrated connections) are identical.<sup>11</sup> The link outputs are combined at the wheels; for each wheel the relevant link outputs are summed, passed through a sigmoid function, and then scaled to the range [-10.0,10.0] to set the wheel speed.

The model is initialised by placing the objects and the agent randomly within the environment (within a 200.0 by 200.0 unit range). The movement of the agent is then calculated on the basis of the wheel speeds (if both wheel speeds are set to 10.0, then the agent moves forward at a maximum speed of 2.8 units per time-step). Each battery has a maximum (and initial) level of 200, which decreases by a single unit each time-step. If the agent encounters a food or water

<sup>11</sup>This figure was arrived at by trial and error; effective agents simply could not be evolved with fewer.

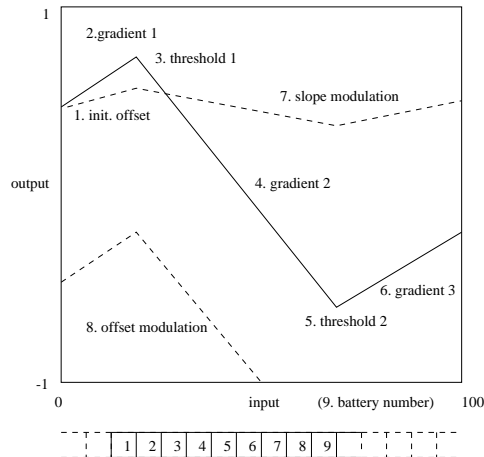


Figure 6.6: Genetic encoding scheme: each sensorimotor link requires 9 integers to specify the various link parameters. The first 6 integers specify the basic shape of the transfer function transforming the sensor input into an output signal, and the final 3 integers specify how this shape can be modified by battery levels. The shape of each link can be modified by either  $\mathcal{B}_f$  (if integer 9 is even) or  $\mathcal{B}_w$  (if integer 9 is odd), and the level of this battery can then influence both the overall gradient of the function (to a degree specified by integer 7), and the offset (to a degree specified by integer 8).

object, the appropriate battery level is restored to 200, and the object is replaced at a different random location. The agent has a maximum ‘lifetime’ of 800 time-steps.

### 6.2.2 Genetic encoding scheme

For each agent a total of 9 links need to be specified, 3 for each connection. Each genotype consists of 83 integers in the range [0,99]; thus 9 integers for each of the 9 links, and one integer each for the left wheel and right wheel sigmoid threshold values (recall that only 9 links need to be genetically specified since left/right symmetry is enforced). Figure 6.6 illustrates how the 9 integers for each link specify the shape of the transfer function. The offset and thresholds are set by scaling the first 6 integers onto the range [-100.0,100.0], with the restriction that the second threshold must follow the first. Note that all scaling is linear and maps onto continuous (non-integer) ranges. The gradients are set by scaling the relevant integers to the range  $[-\pi/2, \pi/2]$ , and then taking the tangent. The sigmoid thresholds are set by scaling to the range [-3.0,3.0]. Specifying the potential influence of the battery levels on the shape of the transfer function is slightly more complicated. If the 9th integer is even, then the shape of the function can be influenced by  $\mathcal{B}_f$ , and if odd, then by  $\mathcal{B}_w$ . The relevant battery level modifies the shape of the transfer function in two ways. First, through ‘offset modulation’  $O$ , where  $O$  is obtained by scaling the 8th integer to the range [-1.0,1.0]. The equation below describes how the output ( $\theta$ ) of the transfer function is influenced (at every time-step) by offset modulation:

$$\theta = \theta + \frac{O\mathcal{B}_{f/w}}{2}.$$

The second potential influence is ‘slope modulation’  $\mathcal{S}$ , where  $\mathcal{S}$  is obtained by scaling the 7th integer to the range [0.0,1.0]. The equation below describes how  $\theta$  is influenced (at every time




Figure 6.7: A sample trajectory; the agent passes through a series of food and water items, displaying opportunistic behaviour at point A, and backtracking to avoid traps at points B and C. Food items are represented by pale circles, water by darker circles, traps by black circles; consumed items are indicated by circles of small radius. An arbitrary north is indicated.

step) by slope modulation:

$$\theta = \theta + \theta S \left( \frac{\mathcal{B}_{f/w} - 100}{100} \right).$$

Note that there is no requirement that any given transfer function, connecting a particular sensor type to a wheel, should be influenced by the battery corresponding to the object type that its sensor responds to.

### 6.3 Results

A distributed GA was employed (see appendix A), with a population size of 100, to evolve the shapes of the transfer functions with a fitness function that simply rewards a high average battery level, calculated incrementally for each time-step that the agent is ‘alive’:

$$\mathcal{F} = \frac{\mathcal{B}_f + \mathcal{B}_w}{400}$$

This function rewards agents that live long (by keeping at least one battery level above zero and by avoiding traps), and that visit food and water items as often as possible. It does *not* specify how an agent should act in any particular situation. Fit individuals consistently evolved in about 200 generations and displayed effective behaviour coordination. The results described in this section are drawn from the fittest individual of the 430th generation of a successful evolutionary run.

### 6.3.1 Behavioural analysis

The effective behaviour coordination of this evolved agent is illustrated in the sample trajectory of figure 6.7. The agent begins by collecting a series of food and water items. At point A, it displays opportunistic behaviour by consuming a nearby food item even though it has only just ‘eaten’. At points B and C, the threat of nearby traps is dealt with by backtracking and turning towards some other food and water items. The agent then continues towards some other food and water items before reaching its maximum lifetime of 800 timesteps. Recalling the list of desiderata for an effective action selection mechanism (section 6.2.1), this example demonstrates the ‘effective prioritisation of behaviour according to current needs’, by generally heading towards the type of object for which there is the greatest current need, whilst, at the same time, avoiding encounters with any traps. The performance of ‘contiguous action sequences’ is illustrated by the way the agent moves smoothly from visiting one item to visiting another, and ‘opportunistic behaviour’ is illustrated at point B. These observations are admittedly anecdotal; the following tests illustrate, more formally, that the evolved agent performs very well with respect to the full list of requirements:

*A balance between dithering and persistence.* The agent would always briefly slow down if equidistant between items for which there was an equal immediate need, before unequivocally plunging one way or the other. Figure 6.8 presents a screen shot from a contrived situation in which the agent was placed an equal distance from a cluster of 3 food items (on the left), and a cluster of 3 water items (on the right). Both  $\mathcal{B}_f$  and  $\mathcal{B}_w$  were set to 100. The agent moves towards the food first, and collects all 3 food items before returning for the 3 items of water; the agent is *persisting* at each task for an appropriate duration, rather than *dithering* between eating and drinking.




Figure 6.8: A balance between persistence and dithering (screen-shot); the agent collects all 3 food items (on the left) before collecting all 3 water items (on the right). The agent does not ‘dither’ between the food and water items.

*Interrupts current behaviour if necessary.* The agent was placed at the (arbitrary) position (0.0,0.0), with a food item to the right at the position (100.0,0.0). Both  $\mathcal{B}_f$  and  $\mathcal{B}_w$  were set to 150. As the agent approaches the food item, and passes the position (30.0,0.0), a trap is suddenly introduced at the position (65.0,0.0). Figure 6.9(a) illustrates the trajectory of the agent as it changes course to successfully avoid the trap whilst still reaching the food, (b) illustrates the simple straight line

trajectory traced by the agent if no trap is introduced, and (c) and (d) illustrate the trajectory immediately before the trap appears, and immediately after its appearance. These images clearly indicate that the agent is interrupting its direct navigation towards the food in order to avoid the trap.

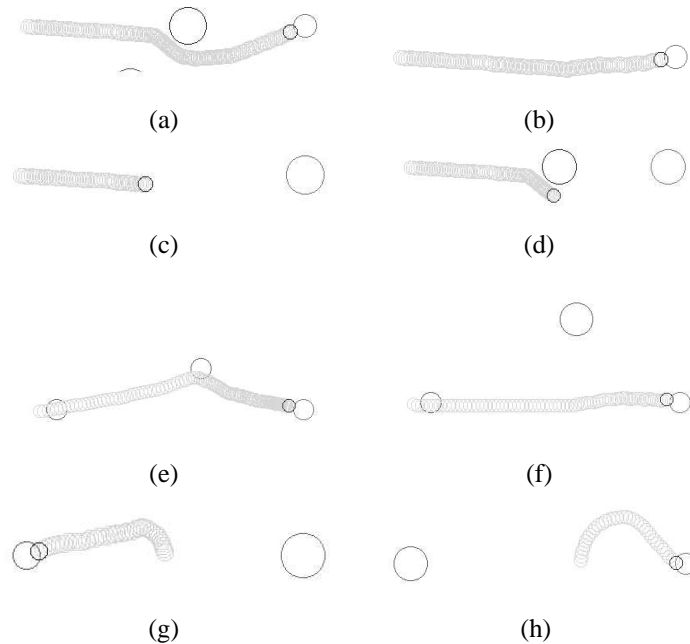


Figure 6.9: Screen shots of behaviour coordination tests. In each picture, consumed items appear as smaller circles. (a) a trap suddenly appears, and the agent interrupts its progress towards the food in order to avoid it. (b) with no trap, the agent progresses directly to the food item. (c) the trajectory of the agent at the point just before the trap appears ... (d) ... and just after the trap appears. (e) the agent collects a first food item, and will divert to a second food item (*en route* to water) if the diversion is not too great. (f) if the second food item is too far away, the agent will proceed directly from the first food item to the water. (g) the agent starts between two equidistant items, and will head towards the item for which there is the greatest immediate need. (h) if the ‘most needed item’ is too far away, then the agent will collect the other item instead.

*Opportunism.* The agent was positioned at  $(-100.0,0.0)$ , with a food item at the position  $(-90.0,0.0)$ , and a water item at the position  $(30.0,0.0)$ . Both  $\mathcal{B}_f$  and  $\mathcal{B}_w$  were set to 150. An additional food item was placed with an  $x$  position of  $-20.0$ , and a  $y$  position that was varied from trial to trial in the range  $[0.0,30.0]$ . The observed pattern of behaviour was as follows. After collecting the first food item, the agent is in greater need of water than of food. However, if another food item is *en route* to the water, the agent may ‘opportunistically’ visit it (figure 6.9e) - unless this second item is too far out of the way (f). Twenty trials were performed at each of 8 different  $y$  positions of the second food item, and figure 6.10 illustrates that as the second food item becomes increasingly out of the way, the agent diverts to it less frequently. It is worth noting that this demonstration of opportunism is also illustrating a preference for consummatory actions over appetitive actions, another *desideratum* of section 6.2.1, and an issue further explored in the mechanistic analysis of the following section.




Figure 6.10: Opportunism. Each column indicates the number of trials (out of 20) in which the agent ‘opportunisticly’ diverted to a second food item; as the diversion distance increased this occurred less frequently.

*Prioritisation.* The agent was placed between an item of water (to the left) and an item of food (to the right), facing away from both. The battery  $\mathcal{B}_f$  was set to 150, and  $\mathcal{B}_w$  to 100. If the items are equidistant (figure 6.9g), the agent will approach the water, but if that item is moved too far away (h), the agent will approach the food instead. Thus the agent is striking a balance between prioritisation and opportunism.

*Other.* The remaining criteria for effective behaviour coordination listed in section 6.2.1 are evidently satisfied in this model. The agent does indeed use all available information, since all sensorimotor links are always influencing the motor output to some extent. And it is also the case that the agent is supplied with real valued sensors and produces directly usable output signals, although of course this condition is rather easy to satisfy in such a simple simulated environment.

### 6.3.2 Mechanistic analysis

So far, this model has illustrated that nothing more is required (under the conditions of the model) for effective behaviour coordination than a set of independent sensorimotor links, and the influence of some internal state. But what else can be said about the mechanism? An easy answer is that its most important features have to do with what it is *not*. It is not a mechanism which performs explicit arbitration between pre-specified internal behavioural correlates (unlike that of Maes, 1990, for example). It, in fact, does not pre-specify such correlates at all. Nor does it decompose the problem of behaviour coordination into ‘sense’, ‘select’, and ‘act’ sub-problems. In fact the opposite is the case. The direct connections from sensor to motor emphasise the importance of intimacy between perception and action, and the parallel operation of these connections precludes any interpretation of the mechanism in terms of explicit arbitration.

To ensure that these claims rest on more than an intuitive appeal, however, it is worth pursuing a more positive analysis of how the mechanism does what it does. This can begin with an exploration of the properties of the links themselves.

Two kinds of link can be distinguished: those that slope downward (from left to right), and those that slope upward. Given a symmetry between left- and right-sided links, and between left and right sensor input, the former will tend to generate ‘approach’ behaviour, by preferentially decreasing the motor output on the side of the agent with the highest sensor input (this is just what

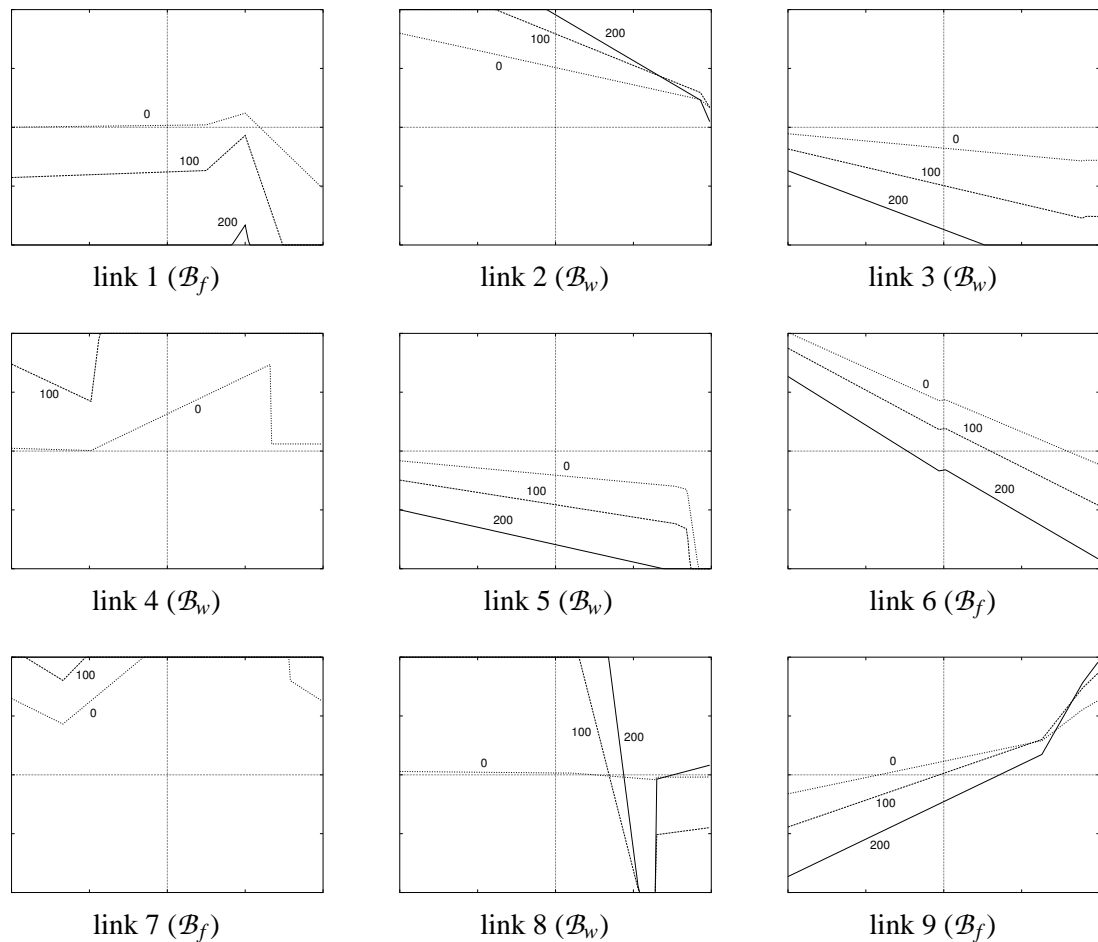


Figure 6.11: Evolved transfer functions; links 1-3 connect the ‘food’ sensors to the wheel (on one side), links 4-6 the ‘water’ sensors, and links 7-9 the ‘trap’ sensors. All axes are scaled to the range  $[-100,100]$ , with input along the abscissa, and output along the ordinate. The 3 lines on each graph represent the transfer function at battery levels of 200,100 and 0, and the battery that influences each function is shown in parentheses. During the operation of the agent, all links are active all of the time. Note that each set of 3 links is influenced by *both*  $\mathcal{B}_f$  and  $\mathcal{B}_w$ .

happens in the very simplest Braitenberg vehicles). Similarly, upward sloping links will tend to generate ‘avoidance’ behaviour. Figure 6.11 illustrates the evolved transfer functions for the links of the evolved agent analysed in the previous section, and already some sense can be made out of their shapes. With some exceptions, those responding to food and water tend to be downward sloping, and those responding to traps tend to be upward sloping. Also, there is some tendency for the gradient to increase with higher sensor input values, indicating that the closer the agent is to a particular item, the more vigorously the agent will attempt to approach or avoid it. To the extent that this is so, it would tend to encourage a discrimination on the part of the agent between consummatory (final approach/avoidance) and appetitive (general taxis) stages of behaviour.

It is also evident that every link is influenced by either  $\mathcal{B}_f$  or  $\mathcal{B}_w$ , even the links connecting the ‘trap’ sensors to the wheels, and also that each set of 3 links is influenced by *both* batteries. That is, the food-sensitive links are not only influenced by the current food deficit, but also by the current water deficit. And the same applies also to the water-sensitive links. This suggests the



absence of any simple mapping from behavioural descriptions such as ‘get food’, ‘get water’, and ‘avoid trap’ onto mechanistic descriptions in terms of the operation of the links.

The remainder of this analysis will provide further support for this suggestion with a discussion of the properties of each link in turn, followed by an exploration of how their joint activity contributes to the behaviour of the agent, with particular regard to behavioural episodes which may - from the perspective of an external observer - suggest the making of decisions.

### 6.3.3 Link properties

To begin with links 1-3, those which are sensitive to the proximity of food items. The output of link 1 is mostly flat until a food item is nearby, at which point which a steep negative gradient encourages approach behaviour. This link is affected by  $\mathcal{B}_f$ , and as  $\mathcal{B}_f$  diminishes, there is an upward shift in the function which would tend to make the agent move faster, perhaps to enhance its search for food. Link 2 also encourages approach to food, with this tendency being very strong if food is very near, and weaker for low values of  $\mathcal{B}_w$ . This property makes good sense because if  $\mathcal{B}_w$  is low, the agent should increase its responsiveness to water. Link 3 is similar to link 2, encouraging approach to food to a degree affected by the level of  $\mathcal{B}_w$ . Again the gradient is shallower for lower levels of  $\mathcal{B}_w$ , and there is also an upward shift in the function as  $\mathcal{B}_w$  diminishes, just as with link 1. Each of these links therefore contributes in a sensible way towards a balanced search for food, encouraging approach behaviour when appropriate and even to some extent distinguishing between consummatory and appetitive behaviours by virtue of there being a stronger tendency to approach when in close proximity to food.

Links 4-6, those sensitive to water, also afford - for the most part - a sensible interpretation. The shape of link 4 is the most peculiar, however it can only produce anything less than maximum output when  $\mathcal{B}_w$  is dangerously low (or when water is very far away with high  $\mathcal{B}_w$ ; an unlikely situation). It may seem that when this is so the link encourages *avoidance* behaviour, however it is also true that the discontinuity evident for high sensor values may encourage very sharp approach (consummatory) behaviour when the agent is very close to water. Further analysis of this link will await the next section. Links 5 and 6 are easier to understand, encouraging general approach behaviour (with a strong consummatory element to link 5). Both of these links present an upward shift as their relevant battery diminishes, encouraging - as before - faster movement.

Links 7-9 are sensitive to traps, and should therefore be expected to encourage avoidance behaviour. Link 7, however, does not seem to do this, being only rarely able to deliver anything other than maximum output. Link 8, by contrast, is very interesting. As with link 4 there is a strong discontinuity when the agent is very close to a trap, which may well encourage sharp avoidance on such occasions. Moreover, there is a very strong dependence on  $\mathcal{B}_w$ . When  $\mathcal{B}_w$  is low, this function is essentially flat, therefore any avoidance behaviour generated by the discontinuity will only happen when  $\mathcal{B}_w$  is high. This makes sense: if  $\mathcal{B}_w$  is very low then it may pay the agent to focus on foraging rather than on avoiding traps. The last link, link 9, is clearly dedicated to generating avoidance behaviour, with this avoidance being sharper for both high  $\mathcal{B}_f$  and if the trap is very close. These links therefore represent a striking example of implicit prioritisation. Imagine for a moment that both  $\mathcal{B}_w$  and  $\mathcal{B}_f$  are very low. Should this be the case, their only contribution would be a slight taxis towards traps if they were very far away (link pair 7/16), and

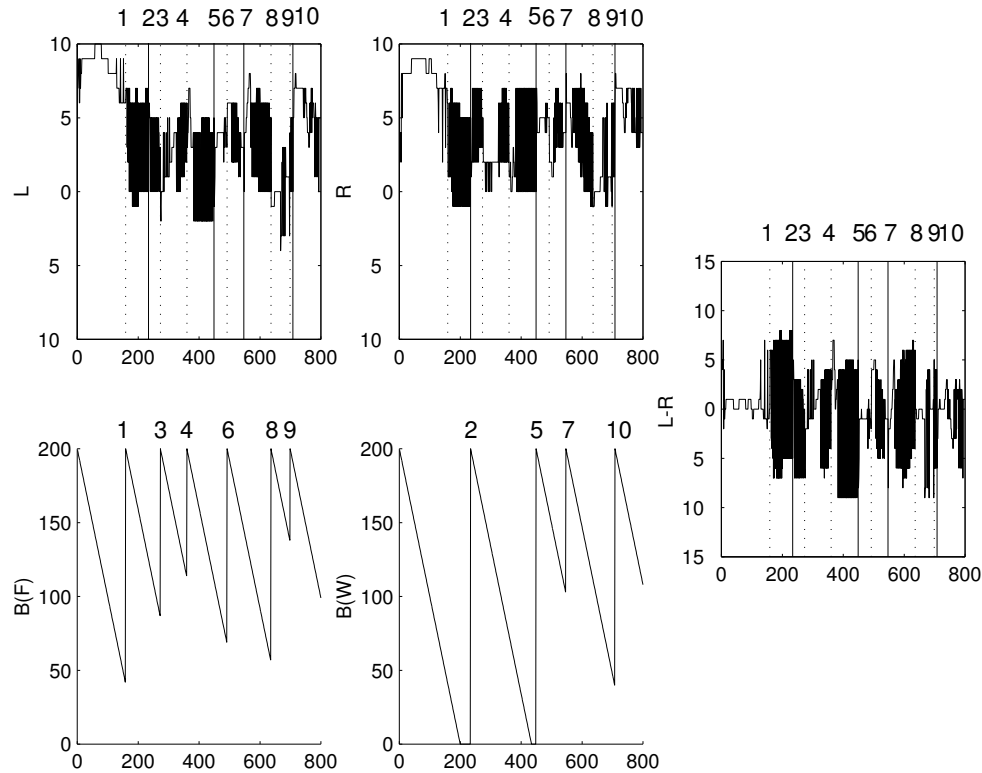


Figure 6.12: Motor output and battery levels. The top two plots show the left and right motor speeds over time, the rightmost plot shows the difference between left and right motor speeds, and the lower two plots show the levels of  $\mathcal{B}_f$  and  $\mathcal{B}_w$ . All plots correspond to the trajectory illustrated in figure 6.7. The motor output plots are marked with dotted lines representing encounters with food items, and solid lines representing encounters with water. All encounters are numerically labelled, with each label corresponding to those of the trajectory plot (figure 6.7). The battery plots are labelled only for encounters with the appropriate item type.

slight avoidance if they were very close (link pair 9/18). The behaviour of the agent would depend mostly on the food-sensitive and water-sensitive links. In the eye of an external observer, trap avoidance would be almost entirely ‘subsumed’ by food and water related behaviour.

### 6.3.4 Mechanism activity during behaviour

Even though it is possible to understand to some extent what each link is doing, it is important to remember that during the ‘natural’ behaviour of the agent, all links are active all of the time. Further insight into how the mechanism works can therefore be gained by tracing the activity of the links during this ‘natural’ behaviour. This section offers a series of observations correlating the sample trajectory (figure 6.7), the link activity patterns, and the motor outputs and battery levels (figure 6.12) during this trajectory. Link activities are illustrated in three ways: absolute activity over time (figure 6.13), activity difference between left-right symmetrical pairs (figure 6.14), and activity difference between successive time-steps (figure 6.15). The observations themselves can be divided into a number of stages corresponding to the various encounters with the food and water items (numerically labelled on each figure).

It is worth opening with a few general remarks. First, figure 6.13 makes it clear that the links

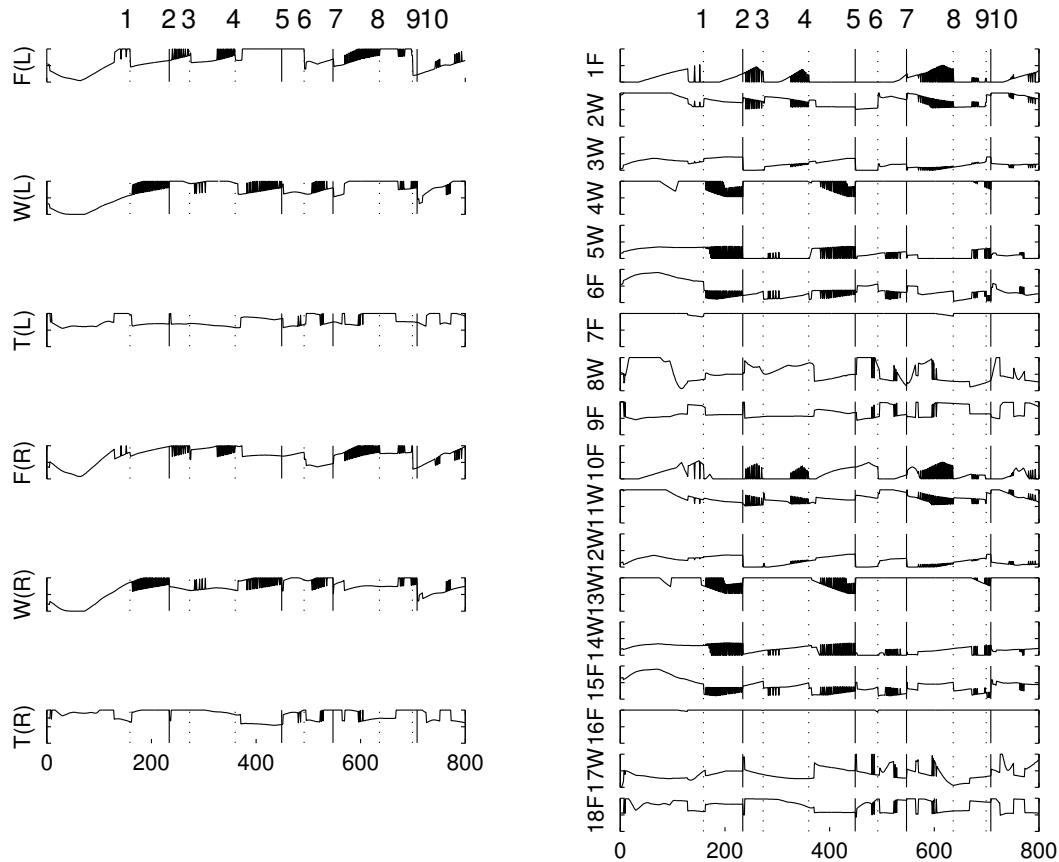


Figure 6.13: Link input and output recordings. The left hand side illustrates the 6 link input signals over time; the top 3 show the food, water, and trap sensor values of the left sensors - F(L), W(L), and T(L) - and the bottom 3 show the same information for the right sensors - F(R), W(R), and T(R). The right hand side shows the link outputs; the set 1-9 control the left wheel, and the symmetrical set 10-18 control the right wheel (each link is labelled with the battery level that influences its shape, F for  $\mathcal{B}_f$  and W for  $\mathcal{B}_w$ ). The links can be divided into groups of 3, with each group taking the same input. For example, links 1-3 take their input from F(L), links 4-6 from W(L), and so on. Both sides are marked with dotted and solid lines and numerically labelled as in figure 6.12, and all plots correspond to the trajectory illustrated in figure 6.7.

characterised by relatively uncomplicated slopes do exactly what would be expected of them; namely to map input to output with varying degrees of fidelity (links 2,6, and 9, for example, exhibit relatively high fidelity). Figure 6.13 also confirms the property of link 8, discussed in the previous section, of relative impotence when  $\mathcal{B}_w$  is low (in particular before encounters 2 and 5, see figure 6.12 for levels of  $\mathcal{B}_w$ ).

Figure 6.12 illustrates that encounters with food and water are often preceded by periods of deceleration and oscillation between left and right wheel dominance, a pattern of activity which makes sense in terms of the agent guiding its final approach to an object (in this model, an object directly ahead will appear, to the agent, to oscillate rapidly between being to the left and to the right). Strong deceleration is also notable in the close proximity of traps; see, for example, the wheel speeds either side of encounter 8.

Final approach to food and water is also easily recognisable in figure 6.15 and 6.14, because of

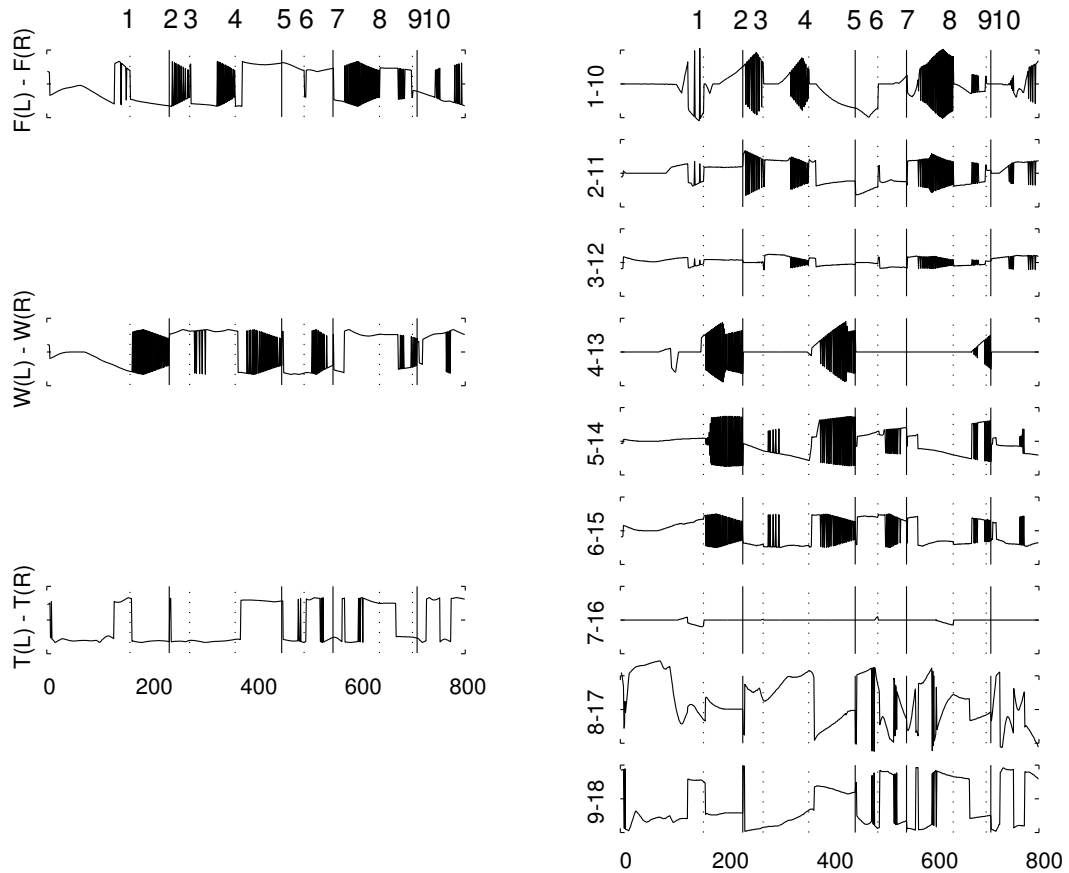


Figure 6.14: Link input and output recordings: left-right difference. These plots are to be interpreted as in figure 6.12, and again correspond to the trajectory illustrated in figure 6.7. These plots show the difference in activity between the left side and right side link outputs for each symmetrical pair (and between the left side and right side sensor input for each of the 3 sensor pairs).

the large per-time-step and left-right differences in link activity that occur during these occasions. Figure 6.15 is particularly useful for isolating these instances (together with sharp avoidance reactions to traps), since such large per-time-step changes do not occur at other times. Figure 6.14, by contrast, shows relatively long-lasting biases of left over right (or *vice-versa*) and so can reflect smoother changes in the agent's trajectory.

With these general remarks in mind some specific observations can now be discussed, beginning with encounter 1: the consumption of a food item. Leading up to this event, link 1 records a general disinhibition of forward speed (figure 6.13; high speeds are evident in 6.12), culminating in a series of sharp spikes as the final approach to the food is generated. This is entirely consistent with the discussion above, and similar activity is observed in its (right-hand side) symmetrical counterpart, link 10. Links 2 and 3 also contribute to the final approach behaviour, but their counterparts 11 and 12 also record slightly stronger activity *before* the final approach (see figure 6.14); this may contribute to the slight leftward bend in the trajectory as the agent approaches the food (left and right are used in this section relative to the heading of the agent). On this evidence it might be argued that link pairs 2/11 and 3/12 are preferentially concerned with generating appetitive be-

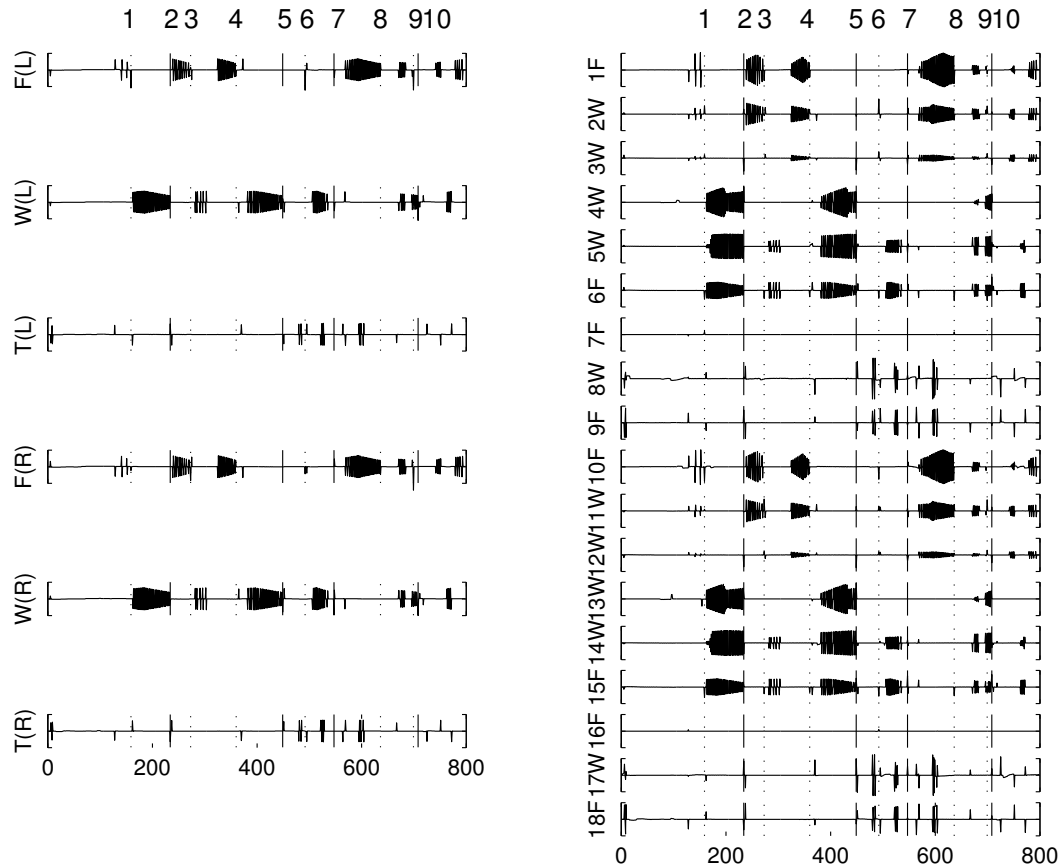


Figure 6.15: Link input and output recordings: time difference. These plots are to be interpreted as in figure 6.12, and again correspond to the trajectory illustrated in figure 6.7. However, instead of showing absolute input and output values (as in figure 6.13), these plots show the difference in activity from one time-step to the next.

haviour (over consummatory behaviour), and *vice-versa* for pair 1/10. The link shapes themselves support this idea, with a sharper gradient change for 1/10 than for the other pairs (figure 6.11), as do the time-difference plots (figure 6.15), with 1/10 displaying by far the most changing activity during the consummatory stage.

Link pairs 4/13, 5/14, and 6/15 are mostly quiet during this first phase, as might be expected for water-sensitive links. Pairs 8/17 and 9/18, on the other hand, are very busy coping with the proximity of traps (7/16 provides an essentially constant positive output throughout). There is plenty of changing sensor information which the activity of 9/18 reflects fairly closely (figure 6.13, see also 6.14; this observation holds throughout the time course of the trajectory). Pair 8/17 responds in a much more complex fashion thanks to its discontinuous structure. Link 8 is initially much more active than 17 (figure 6.14) due to the presence of traps to the right, which helps to bring the agent round from its initially inappropriate northerly heading; perhaps in this case taxis *towards* traps is being used to keep the agent in the general area in which things happen. As the agent approaches the food, this surplus of activity falls away, and, together with a jump in activity from pair 9/18, the agent is guided safely between the first two traps that lie more or less directly in its path.

After consuming the food item, the agent proceeds to encounter 2, a water item. Water-sensitive pairs 4/13, 5/14, and 6/15 are dominant here, delivering sharp taxis towards the water; their changing activity is clear in all representations of link activity (especially figure 6.15). However, although these links are dominant, it is not true that they are the *only* links influencing behaviour. Even though pairs 1/10, 2/11, and 12/3 are generally quiet (figure 6.13), a rising difference can be detected between links 1 and 10 (figure 6.14) tending to push the agent round to the right, where a food item - 3 - is to be found. And even though pairs 8/17 and 9/18 are also quiet (there is not much change in sensor data, and  $\mathcal{B}_w$  is very low meaning that 8/17 is relatively impotent) they compensate for this effect with a slight bias pushing the agent to the left, due to the proximity of a trap to the right (figure 6.14). (It is also worth noting here that link 4, despite its peculiar shape, in practice rarely departs from maximum output, and when it does it is to contribute to final approach to water.)

Instead of continuing this analysis throughout the entire trajectory (a lengthy task), this discussion will now concentrate on just a few moments of behavioural significance in what remains. The first of these moments occurs between encounters 3 and 4. Here, the agent has just consumed a food item, and appears to be heading for a water item (5) but then alters course and collects another food item (4). This ‘opportunistic’ display is significant for its representation of a period of activity to which descriptions of ‘decision making’ behaviour might potentially attach. However, an inspection of the mechanistic dynamics during this period reveals no corresponding arbitration between internal behavioural correlates for, for example, ‘going towards food’ and ‘going towards water’. Link pairs 1/10 - 3/12 certainly encourage taxis towards food as encounter 4 becomes imminent (figure 6.13, 6.15), but there is nothing special to distinguish this activity from that observed prior to other encounters with food. Also, links 2 and 3 are more active than their counterparts (11 and 12) during the earlier phase, helping (appetively) push the agent around to the right, away from the water and towards the food. At the same time, link pairs 5/14 and 6/15 are trying to initiate a final approach to the water item (see in particular the spikes in figure 6.14). However, this does not happen, and as the agent turns to the food, activity in these links settles down (they do, however, continue to maintain a significant bias towards the left, the direction of water item 5, with this bias increasing over time for 6/15). But it is the trap-sensitive links that perhaps, and perhaps surprisingly, hold the key to understanding this period of behaviour. Despite an almost entirely static sensor input (greater on the left, see figure 6.14), there is some changing activity in link pairs 8/17 and 9/18 as a consequence of the varying battery levels. Pair 8/17 displays a particularly strong bias pushing the agent towards the right (and the food). This is offset somewhat by an opposite bias from 9/18, but as time goes by the bias from 8/17 increases whilst that of 9/18 decreases. The ‘decision’ between going towards water and going towards food may, on this occasion, therefore depend to some extent on the activity of these trap-sensitive links. Not that these links were necessarily evolved for this purpose, nor indeed solely responsible for the final behaviour. It is equally plausible, as suggested above, for their contribution here to be understood in terms of keeping the agent in the ‘general area in which things happen’, with decision making - on this occasion - a happy consequence of this when taken in conjunction with the activity of the other 6 pairs of links. The lesson, in this instance, is that any decision making that an external observer may discern in this episode is certainly *not* underpinned by explicit arbitration

between distinguishable internal behavioural correlates of any kind.

The period of behaviour following encounter 4 is also instructive. The agent collects a water item (5) before bearing left towards another food item (6). Link pairs 4/13 - 6/15 are heavily involved in the approach to the water (figure 6.15), but even during this time there is some bias towards turning left (towards the food), both from the food-sensitive links (1/10 and 2/11), and also from the trap-sensitive pair 8/17 (again offset to some degree by 9/18), as evidenced in figure 6.14. In the midst of unambiguous water foraging, then, the mechanism is beginning to generate the conditions required for the next behavioural episode. Indeed this pattern of activity has already been described twice before, between encounters 1 and 2 and again between 3 and 4, and can lay some claim to being a motif throughout the trajectory: after encounter 5, again, the activity of links 4-6 (13-15) imposes a slight *rightward* bias, in preparation for the next *water* encounter (7).

After encounter 8, the agent turns sharply leftward in order to avoid a trap (point C on figure 6.7). However, the majority of the changing activity in the trap-sensitive links occurs well *before* the encounter (see in particular figure 6.15), as if the agent is cautious of getting too close to the trap in the first place. This activity is reminiscent of the preparatory activity of the food-sensitive and water-sensitive links described above. Also, the rapid deceleration upon consumption, and sharp turning afterwards, is as much due to link pairs 1/10 - 6/15 (sensitive to food and water) as to pairs 8/17 and 9/18; indeed, the activity of these trap-sensitive links is strikingly continuous during the pre-encounter and post-encounter period. What this is suggesting is that trap avoidance behaviour, on this occasion, is not something that can be uniquely localised to the trap-sensitive links.

A final episode worth exploring follows the final encounter (10). The agent, heading northerly, appears to veer first towards a pair of food items to its right before altering course towards some water on its left. This episode is similar to that described above between encounters 3 and 4, and may again invite interpretation in terms of explicit decision making. As before, however, the actual mechanism dynamics belie such an interpretation, but this time no particular links hold the key to understanding the behaviour. All are equally important. Pairs 1/10 - 3/12 initiate taxis twice; the first time when the agent is veering rightwards, and the second time on the final approach to the water. (This second taxis may be explained by the presence of another food item slightly beyond the water.) The water-sensitive links are quiet for the first phase of this episode, but then after a brief period of high activity they impose a steady leftward bias on the agent, turning it towards the water. And the trap-sensitive links are also involved, helping the agent turn away from two traps in succession, and in doing so tending to favour the water over the food. Once again the agent appears to have made a decision without this decision being represented within the agent in terms of explicit arbitration.

### 6.3.5 Mechanism analysis summary

To briefly summarise this rather detailed mechanistic analysis. The behaviour of the agent during any externally isolable episode can be understood to have its mechanistic basis in shifting coalitions of link activity, with apparent ‘decisions’ not identifiable with any explicit arbitration between distinguishable internal behavioural correlates. Some properties of the mechanism which help to explain how it can underlie effective behaviour coordination are (i) a general balance be-

tween approach and avoidance behaviour across the set of all links, (ii) the fact that each set of links is influenced by *both* battery levels, (iii) the way in which the links can express the distinction between appetitive and consummatory activity, (iv) the preparatory activity exhibited by the links as they begin to generate the conditions required for successive behavioural episodes, and (v) the implicit prioritisation effected by the links, for example the weakening of trap avoidance and general disinhibition of forward speed with low battery levels. These properties, taken together, explain how the agent - in the eye of an external observer - may appear to make decisions, and to make them effectively.

## 6.4 Discussion

As I have already argued, this model has illustrated that nothing more is required (under the conditions of the model) for effective behaviour coordination than a set of independent sensorimotor links, and the influence of some internal state. In this model, perception and action are inseparable, and there is no arbitration between internally represented behaviours. This claim, and in particular its second part, is supported both by appeal to the assumptions (or lack of them) that constitute the *a priori* structure of the mechanism, and also by the detailed analysis of the properties of an actual evolved mechanism of this type. By these lines of argument, it is clear that an understanding of how link activity underlies effective behaviour coordination does not require, and indeed actively undermines the identification of any straightforward mapping from behavioural descriptions onto unique mechanistic entities.

Of course, there would be no reason to expect such a mapping to have emerged, given that the IOS methodology allows the design of each sensorimotor link to proceed concurrently with the fusion of all links into a coherent global behaviour. As noted before (section 6.2.1), this avoids the potential influences of behavioural descriptions on internal mechanism that may arise in hand-designed architectures, where process design and process fusion are necessarily carried out separately.

An obvious criticism of this model may be levelled at its apparent simplicity. In such a simple environment, not only are the processes of perception and action somewhat trivial, but the problem of behaviour coordination itself is constrained by the limited behavioural flexibility of the system. Rather than balancing and temporally sequencing a wide variety of potential actions, the agent is required only to move over a two-dimensional surface, and in particular is *not* required to string together complicated contiguous appetitive action sequences prior to performing consummatory actions. This simplicity is in marked contrast to many of the (mostly hierarchical) action selection mechanisms in the literature that purport to be either biologically faithful (Baerends, 1976) or practically useful (Tyrrell, 1993). The conclusions to be drawn from the present model must therefore be stated with care. I have shown that it is possible for a *simple* form of effective behaviour coordination to emerge from a process agent-environment interaction mediated by tightly coupled perception and action, *without* the existence of internal behavioural correlates, and *without* any explicit process of internal arbitration (the simplicity of the present model is therefore a virtue rather than a handicap insofar as it enhances the clarity of this demonstration). I have *not* shown that such a minimal mechanism can ‘outperform’ orthodox mechanisms, whether hierarchical or not (although it does do rather well), and it cannot be concluded that the architecture of the present



model is in any sense more ‘biologically plausible’ than orthodox architectures.<sup>12</sup>

However, neither should the significance of this model be understated. Whilst admittedly simple in many ways, it nevertheless affords a set of behaviour-mechanism relations of intriguing complexity. The evolved behaviour coordination is extremely effective, and the evolved mechanism accomplishes this coordination in a non-trivial fashion. Even though the ‘biological plausibility’ of the mechanism is not the issue, it may well be that some of the specific mechanistic insights delivered in the previous section will have implications for how biologists come to analyse organic systems. More generally, and perhaps more importantly, the present model speaks to the theoretical critique of traditional frameworks of behaviour coordination developed in section 6.1. Here I have shown that an alternative is available, not only in principle, but also in practice.<sup>13</sup>

## 6.5 Summary

What is the difference between externalist and internalist accounts of behaviour coordination? There is no hard and fast distinction, but internalism may be divined in any account that relies, to some extent, on any of the following: internal behavioural correlates, internal arbitration mechanisms, and internal hierarchical structure. Examples of such accounts were discussed in the first part of this chapter, and significant theoretical problems with these accounts were identified with respect to the necessary distinction between behaviour and mechanism, and the essential unity of perception and action. By contrast, the present model has developed a relatively *externalist* account of behaviour coordination in the form of an IOS model in which coordination emerges from agent-environment interaction patterns mediated by closely coupled perception and action. To the extent that the internal mechanism of the agent in the model is ‘minimal’ (and it would appear to be difficult to propose any mechanism with equivalent functionality that could be any more ‘minimal’) it must resist any attribution of internalism. The significance of this model lies not in its practical utility (which, whilst promising, has not been demonstrated) nor in its biological fidelity (which is not asserted), but in the fact that it constitutes existence proof of the coherence of a pragmatically externalist account of behaviour coordination which avoids the theoretical obstacles that beset a large proportion of orthodox approaches.

In the previous chapter it was emphasised that IOS modelling enabled accounts of decision-making to be developed in the absence of explicit decision variables. The model of the present chapter can certainly be interpreted in this way, but so far we have engaged with the problem of decision-making - behaviour coordination - only at a relatively abstract level. In the following

<sup>12</sup>For recent evidence in support of an orthodox position, see Biró and Ziemke (1998) and Platt and Glimcher (1999). Biró and Ziemke use a GA to evolve the parameters of a recurrent neural network for an agent to perform a simple visual task. They analyse the dynamics of effective networks and argue that these dynamics decompose into behavioural categories; in other words they appear to have evolved internal behavioural correlates. Along similar lines, but working with real monkeys rather than with simulated agents, Platt and Glimcher claim to have discovered neural correlates for decision variables in the parietal cortex of rhesus monkey brains (as noted in chapter 5).

<sup>13</sup>An interesting precursor to the present model is described by Herbert Simon, who asks “[h]ow simple a set of choice mechanisms can we postulate and still obtain the gross features of observed adaptive choice behavior?” (1956, p.129). However, instead of constructing computer-based simulation models - a difficult task in 1956 - Simon attempted to mathematically formalise a simple foraging situation. His model, whilst persuading him to conclude that “we should be skeptical in postulating [...] elaborate mechanisms for choosing among diverse needs” (ibid., p.137), is therefore unable to represent mechanism at the same level of situated perception and action as the present model, or to interrogate assumptions concerning internal behavioural correlates in the same way. It is, nevertheless, a wonderfully prescient and perspicuous document.

chapters the present model is modified in various ways to allow the exploration of less abstract instantiations of the problem that engage relatively directly with the theoretical biology literature.

## Chapter 7

### An IOS model of the interference function

---

When multiple predators forage for prey in an environment characterised by limited prey availability, the intake rate per predator can be considered to be a function of both prey density and predator density. In cases in which the influence of predator density is non-negligible, this function describes an *interference* relationship, where interference is defined as the more-or-less immediately reversible decline in intake rate due to the presence of conspecifics (Goss-Custard, 1980; Sutherland, 1983). Although such situations have been extensively modelled in contemporary theoretical biology, the modelling strategies generally employed remain the focus of debate (Van der Meer & Ens, 1997). This chapter will outline and critique this contemporary work, and it will be argued that the IOS modelling strategy is well placed to move the debate forward, in virtue of (a) modelling at the level of situated perception and action, and (b) consistency with the principles of optimal foraging theory.

The IOS model of interference developed in this chapter is related to that described in the previous chapter. However, instead of dealing with the question of ‘how to do the right thing’, and in addition to modelling interference itself, the present model explores (within the context of interference) an instance of *suboptimal* - or ‘irrational’ - behaviour.<sup>1</sup> Recall from chapter 5 that both orthodox optimal foraging models and IOS models may be directed towards functional as well as mechanistic explanation, and that one way of understanding apparently suboptimal behaviour is in terms of adaptation to environmental features that formed part of past, but not present environments. Kacelnik (1997), for example, draws our attention to the way in which people discount the value of future rewards, arguing that what seems in the laboratory to be a suboptimal underestimation of future value would, in fact, have been optimal had future rewards been unreliable, a feature perhaps characteristic of our ancestral habitats. This chapter, and those that follow, begin to investigate a similar idea; the hypothesis that behaviours adapted to a *group* situation may be suboptimal - ‘irrational’ - when expressed by an isolated individual.

---

<sup>1</sup>The reader may prefer to think of behaviour as ‘rational’ or ‘irrational’ rather than as ‘optimal’ or ‘suboptimal’; in this context the terms are interchangeable, but in this dissertation as a whole the language of optimality is preferred.

## 7.1 Modelling the interference function

It may seem obvious that any model of interference should be consistent with the principles of optimal foraging; that modelling the reduction in intake due to predator density should take place in the context of predators attempting to maximise their intake. This perspective is mandated not only by orthodox optimal foraging theory itself, the IOS methodology, and the general externalist emphasis of this dissertation, it is also more-or-less demanded in view of the importance of interference in optimality-based ‘ideal-free’ theory.

The central premise of ideal-free theory is that an equilibrium distribution of rate-maximising predators across a ‘patchy’ environment of varying prey availability will be reached when no predator can improve its intake by moving to a different patch. At this equilibrium, referred to as the ‘ideal free distribution’, all predators necessarily experience the same intake rate, and there will be some proportional relation between prey availability in a patch and the number of predators inhabiting that patch (Fretwell & Lucas, 1970; Fretwell, 1972). The derivation of the ideal free distribution therefore requires knowledge of the relation between predator density and per predator intake in each patch, and this, of course, is the relation captured by the interference function. (Strictly speaking this is only an accurate description of so-called ‘standing-stock’ ideal free distribution models, however, for present purposes this qualification may be overlooked; chapter 9 will furnish further details.)

It should also be stressed that in spite of the biological tenor of this chapter, the relevance of interference and its application in ideal free theory is by no means limited to the biological domain. A theory of the spatial distribution of predators (and, by implication, a model of interference) may apply as much to an oystercatcher-mussel ecosystem as to a collection of robots foraging for rock samples, as to a group of information agents scouring the internet.<sup>2</sup>

Despite these strong incentives, the development of models of interference consistent with principles of optimality still “stands out as a clear theoretical challenge”, as Van der Meer and Ens argue in a recent review (1997, p.857). The following section briefly outlines the range of interference models that can be found in the literature, and the remainder of this chapter is devoted to the development of an IOS model that, amongst other things, addresses itself to Van der Meer and Ens’ challenge.

### 7.1.1 A brief history of interference

In their recent survey of the biological literature on the subject, Van der Meer and Ens (1997) distinguish two approaches to modelling the interference function. One involves the construction of simple individual-based models in which the actions of individual agents follow pre-specified rules. The way in which the application of these rules interacts with agent density in the determination of intake rate can then be derived (Ruxton, Gurney, & Roos, 1992; Moody & Houston, 1995; Tregenza, Parker, & Thompson, 1996; Holmgren, 1995). The second method concentrates on identifying some empirical relationship between intake rate and agent density (Hassell & Varley, 1969; Sutherland & Koene, 1982). These two approaches are labelled ‘mechanistic’ and ‘phenomenological’ respectively by Van der Meer and Ens. For the sake of consistency with the rest

<sup>2</sup>For an example of autonomous agents research that attempts to employ the concept of interference, but without reference to its biological foundation, see Goldberg and Matarić (1997).

of this dissertation, however, ‘mechanistic’ models will here be referred to (equally appropriately) as individual-based models.

Most individual-based models make reference to the ‘chemical reaction kinetics’ model of Ruxton et al. (1992), in which a population is divided into a number of mutually-exclusive states, (for example, prey-handling, searching, fighting, time-wasting). Sets of ‘reactions’ are defined between these states, for example, upon encountering a food item, an agent would move from a searching state to a prey-handling or feeding state, and differential equations are constructed expressing transition rates between these states. The equilibrium solutions of these equations (obtained either analytically or numerically) yield the interference function. The intuition underlying this approach is that individual agents behave like ‘aimless billiard balls’, or molecules in a chemical reaction, moving randomly and interacting upon collision in ways determined by the reaction matrix. Interference is explicitly associated with particular states, for example time-wasting, or state transition sequences, such as kleptoparasitism.<sup>3</sup> It is important to note that these individual-based models do *not* generally involve simulations of agents actually moving around in spatially defined environments (in other words they are not IOS models); rather, agents are abstract entities characterised only by being in particular states (drawn from a pre-determined repertoire), with the ‘environment’ consisting only of the set of probabilities for the various state transitions. The various descriptions in the literature of this class of model as ‘mechanistic’, or, as is sometimes the case, as ‘behaviour-based’ (Ruxton et al., 1992), may therefore be potentially misleading for those for whom such terms connote modelling at the level of agent perception and action.

Phenomenological approaches originate from the identification, from empirical data, of a linear relationship between the logarithms of search efficiency (closely related to intake rate) and agent density (Hassell & Varley, 1969), the slope of which was identified with the level of interference and given the label  $m$ . This model and its derivatives, although widely used, have been criticised both for lacking any interpretation of interference at the level of the underlying behaviour-generating mechanisms (Van der Meer & Ens, 1997), and failing to adequately describe subsequent empirical data, some of which, for example, describes a non-linear log-log relationship and hence a level of interference that varies with agent density (Ruxton et al., 1992; Moody & Houston, 1995). These criticisms are naturally linked; if one is to account for patterns of empirical data that differ from those previously observed, one must possess a model at a level capable of generating data, not just of description of observed data. As a consequence, generative individual-based models along the lines of Ruxton et al. have, over recent years, come to predominate.

However, as Van der Meer and Ens point out, these individual-based models are not fully consistent with the optimal foraging foundations of ideal free theory, since the inflexible nature of the state transitions experienced by the agents provides no guarantee that individual agents will follow sequences of states that maximise their intake rate. Stillman et al. (1997) have recently addressed this challenge with an individual-based model in which agents follow ‘optimal’ decision-rules. Here, state transitions are no longer entirely pre-determined, rather, they can be made according to the relative costs and benefits they offer. For example, an agent that encounters a feeding conspecific will decide between continuing to search independently or fighting for the half-consumed prey (again, these events concern only disembodied states and transition probabilities). Many

<sup>3</sup>A kleptoparasitic event is when a predator captures a prey item only for it to be stolen by another predator.

factors can be involved in such a decision, and Stillman et al. consider, amongst other things, the relative competitive abilities of the two agents and the probability of encountering other prey. They use data from an oystercatcher-mussel ecosystem to show that their approach provides enhanced predictive power over a control model with fixed state-transition rules.

This ‘optimal’ model is a valuable step forward, but it can itself be criticised on at least two fronts. First, the continual assessment of relative costs and benefits imputes considerable computational power and statistical nous to the agents. This criticism is a version of Gallistel et al.’s (1991) argument that the concept of a ‘decision variable’ in optimal foraging theory, in spelling out the processes underlying the optimisation of behaviour, “credit[s] the animal with complex representational and computational abilities” (p.18, see also chapter 5). Of course, one response to this kind of criticism would be that the point at which agent complexity becomes theoretically awkward is merely a matter of opinion, but such a response is not likely to be universally satisfactory.

A second, and perhaps more telling criticism, and one that applies also to the predecessors of Stillman et al., is that agents are not modelled at the level of situated perception and action, nor do they operate in a spatially explicit environment. The isolation of a set of behavioural states and a set of potential state transitions, and the explicit identification of a subset of these states or state transition sequences with interference, is a phenomenological process in just the same sense as is the identification of interference directly from empirical data. In other words, existing individual-based models can be criticised on exactly the same grounds that proponents of these models criticised earlier phenomenological models: the components of the models require explanation in terms of behaviour-generating mechanisms. To return to the distinction between behaviour and mechanism first expressed in chapter 2; models that claim to invoke behaviour-generating mechanisms should operate at a level other than that of the behaviours themselves. In the present case, this would seem to require agents to be modelled at the level of individual perception and action in a spatially explicit environment. Interference would then be *observable* as a consequence of the interaction of such agents with each other and with the environment, rather than being introduced as an *a priori* component of the model itself.

In summary, phenomenological models which aim to *describe* the data have given way to individual-based models which aim to *generate* the data from underlying behavioural rules. Recent emphasis has been placed on locating these behavioural rules within a framework of optimality, and the work of Stillman et al. represents a significant development in this direction, albeit at the expense of requiring considerable agent complexity. However, no individual-based models to date operate at the level of perception and action in a spatially explicit environment, a level that is arguably a requirement for any interpretation of an agent-based model in terms of ‘behaviour-generating mechanisms’.

## 7.2 An IOS model of interference

The model of the present chapter is presented as a response to the above critique. It is best introduced as a multi-agent version of the behaviour coordination model developed in the previous chapter. To rehearse the essential common features: (a) operation at the level of perception and action in a spatially explicit environment, and (b) the use of GAs to partially specify agent mech-

anism, together with a very simple and general fitness function which requires only that agents obtain high intake rates, without prefiguring the structure of the behaviour through which this is to be achieved. These features, representative of the IOS modelling strategy in general, are clearly significant in the context of modelling interference. The use of GAs entails consistency with the principles of optimal foraging, without necessarily requiring that agent-side mechanisms be complex. Modelling at the level of situated perception and action means that interference no longer need be explicitly identified with specific isolable agent components, as is the case, for example, with the states and state transition sequences of Ruxton et al. (1992). Instead, interference can be observed as a consequence of the interaction of agents with each other in the context of a shared environment; in this way - as with IOS models in general - the relationship between behaviour and mechanism becomes the object of study rather than constituting an *a priori* component of the model.

The structure of the model is designed to capture the simplest possible scenario in which interference can be explored in an IOS context. Simple neural network controllers are embedded in equally simple agents that move around a simulated single-patch spatial environment, collecting food items, with artificial evolution specifying network parameters such that agents forage near-optimally. The observed relationship between agent density and intake rate then yields the interference function. This model is manipulated in a number of ways to explore the relations between individual behavioural strategies and the patterns of interference they entail. Before discussing the details of these conditions, however, I will describe the basic model itself.

### 7.2.1 Agent and environment

In the basic model, each agent possesses 3 sensors, 2 responding to the distance to the nearest food item, and one reflecting an internal battery level (agents in this basic model therefore cannot directly sense the presence of other conspecifics). The distance sensors range linearly from 100.0 (at the item) to 0.0 ( $\geq 200.0$  arbitrary distance units away), and the battery sensor ranges linearly from 0 to 200 (the maximum, and initial, battery level). If the nearest food item is to the left of the agent, the relevant left sensor responds with 20% greater activation (and *vice versa* if the object is to the right). The internal architecture of each agent comprises a simple feedforward neural network, fully interconnected between layers, but with no recurrency; agents, therefore, can only engage in *reactive* behaviour. The 3 inputs (corresponding to the 3 sensors) feed through to a 5 unit hidden layer, and then to a 2 unit ‘motor’ layer. The input units linearly scale the sensor values to the range [0.0,1.0], and all interconnecting weights are in the range [-1.0,1.0]. Each neuron in the hidden and motor layers applies the standard sigmoid transfer function (see Rumelhart et al., 1986) to the sum of its inputs (plus a threshold value), with the outputs scaled to the range [0.0,1.0]. The outputs of the motor layer are scaled to the range [-10.0,10.0] to set the wheel speeds; a value of 10.0 on both wheels translates to a speed of 2.8 distance units per time step.

Real valued genotypes of length 32, with each locus constrained to the range [0.0,1.0], determine network structure as follows: the first 15 loci map onto the range [-1.0,1.0] to specify the weights from the input layer to the hidden layer, the following 10 loci similarly specify the weights from the hidden layer to the motor layer, and the final 7 loci retain the range [0.0,1.0] and specify the thresholds of the hidden and motor units.




Figure 7.1: Depiction of a typical environment populated by 3 agents (filled circles) and 8 food items (clear circles). Both agents and food items are always initialised within the 200.0 by 200.0 patch (dotted square), but agents are *not* constrained to remain within this patch. The structure of a single agent is also depicted, with 3 sensor input units, 5 ‘hidden’ units, and 2 motor output units arranged in a feedforward network architecture.

The environment comprises a (simulated) infinite plane containing stationary food items (radius 8.0). Both food items and agents (radius 2.5) are always initialised within a 200.0 by 200.0 ‘patch’ centred at the origin. Each time an agent encounters a resource item, the item disappears to be immediately replaced in another random location in the patch, thus ensuring a constant density of available resources (this is important in view of the distinction between interference and the straightforward depletion of resources, see Free, Beddington, & Lawton, 1977). Each encounter with a resource item fully replenishes the agent’s battery level, which otherwise depletes at a rate of one unit per time-step; if the battery level reaches zero the agent ‘dies’. Encounters with conspecifics (if any) have no effect. Figure 7.1 illustrates a typical environment populated by 3 agents, together with a representation of the structure of a single agent.

### 7.2.2 Evolving an interference function

In each of the experimental conditions described below, near-optimal foraging behaviour is evolved over 800 generations of a distributed GA (see appendix A), operating on a population of 100 randomly initialised genotypes. Each generation of the GA, each genotype is evaluated 4 times, with a final fitness score for the genotype derived from the average fitness across all 4 evaluations. Each evaluation begins by decoding the genotype into either a single agent, or into a number of identical ‘clonal’ agents. Consider, for the moment, the single agent case only. After decoding, the agent is placed at random in the patch along with 8 randomly scattered food items (with a minimum spacing of 25 distance units between objects). The fitness of the agent (and therefore of the genotype) is then assessed by means of the incremental fitness function:

$$\mathcal{F} = \sum_{t=1}^{800} \frac{\mathcal{B}}{200},$$

where  $t$  indexes time-steps and  $\mathcal{B}$  represents the battery level (at time  $t$ ); each evaluation lasts for a maximum of 800 time-steps. This function rewards agents that live long and forage efficiently, without in any way specifying the structure of the behaviour through which this is to be achieved.

After a sufficient number of generations have elapsed to ensure near-optimal foraging (800 is quite enough, fitness reliably asymptotes after 400), an interference function can be derived by taking the fittest genotype from the final generation, decoding it into an agent, and assessing the



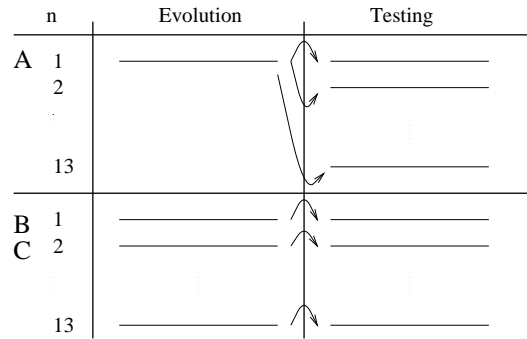


Figure 7.2: Interference model experimental design;  $n$  represents group size. In condition *A* all groups in the testing phase are derived from a single agent, this agent being the product of a single GA which assessed genotypes in isolated agents. In conditions *B* and *C* each group in the testing phase derives from a distinct agent, each of which is the product of a separate GA, each of which assessed genotypes in groups of size  $n$ .

intake (averaged over 500 evaluations) of the agent both in isolation (as it was assessed in the GA itself), and also in the presence of  $n$  ‘clonal’ conspecifics for all  $n \in \{1, 12\}$ . In this way the relationship between agent density and intake rate - the interference function - can be recovered.

### 7.2.3 Experimental conditions

The basic model described above involves a period of evolution in which agents are evolved in isolation, and a period of testing in which the intake rate of a focal agent is assessed in (clonal) groups of various sizes. It is important to stress that all groups in this basic model - which may be called condition *A* - comprise of agents that were *evolved in isolation*. Two further conditions are described here as extensions of this basic model.

Condition *B*, in contrast to condition *A*, requires 13 distinct GAs in the evolutionary phase (indexed by  $n$ ), one for each group size from  $n = 1$  to  $n = 13$ . Consider the case, for example,  $n = 8$ . In the corresponding GA each genotype in the population is assessed by being decoded into 8 clones which coexist in the environment. The fitness of the genotype is determined by that of a randomly selected agent clone from within this group. The testing phase proceeds as before, but of course each group is now derived from the fittest genotype from the final generation of the corresponding GA. Thus, all groups in condition *B* comprise of agents that were *evolved in the presence of the same number of conspecifics also present during testing*. Figure 7.2 graphically illustrates the difference between conditions *A* and *B*.

Condition *C*, an extension of condition *B*, was introduced to assess the sensitivity of the model to a significant variation in its structure; the agents in *C*, unlike those in *A* or *B*, are able to perceive one another. The structure of the model in this condition differs only in that each agent possesses an extra pair of sensors (and an additional pair of neural network input units) that respond to the bearing and distance of the nearest conspecific (if any).

Why bother with these various conditions? After all, the desired features of the present model, as set out at the beginning of this section - consistency with the principles of optimal foraging and modelling at the level of situated perception and action - are well catered for by condition *A* alone. However, it is also an objective of this chapter to explore the hypothesis that behaviours adapted to

a group context may be suboptimal when expressed by an isolated individual, and it is here that the distinction between conditions *A* and *B* becomes significant. In what follows, it is demonstrated that agents evolved in the presence of conspecifics deploy a behavioural strategy which lessens the toll of interference for all agents, even though this strategy entails that isolated individual agents behave suboptimally. Condition *C* is designed to assess the robustness of this observation to a lessening of the extreme sensorimotor poverty experienced by agents in conditions *A* and *B*.

#### 7.2.4 Summary of model structure

Three major departures from the behaviour coordination model of the previous chapter can now be identified: (1) the use of multiple agents (as required by any model of interference), (2) the presence of only a single resource type and the absence of ‘traps’ (for the sake of simplicity), and (3) the use of simple neural networks as opposed to sensorimotor ‘link’ controllers. The last of these changes is motivated by pragmatic rather than theoretical reasons. The link architecture is, in practice, rather clumsy to work with and was intended in the previous chapter to emphasise the intimacy of perception and action, and the arbitrariness of assumptions of ‘internal behavioural correlates’ underlying behavioural choice. The present model reverts to the use of standard neural networks primarily because they are simple to use and their mechanistic properties are well understood (Rumelhart et al., 1986). The present networks, being feedforward and non-recurrent, nevertheless retain the essential simplicity of the links of the previous chapter.

The present model is indeed simple in all respects; all that agents can do is *move*, with food being consumed instantaneously upon encounter. Many activities that are usually incorporated into individual-based models of interference have been omitted. Agents do not spend time handling prey, they do not fight each other, and they do not distinguish between ‘searching’, ‘avoiding’, or ‘time-wasting’ (although an external observer might wish to describe their activity using such terms). They are able only to perceive the direction and distance of the *nearest* food item (although this poverty of perception is relaxed in condition *C*) and are able to use this information only to control their heading and speed. The benefits of such simplicity have been discussed at length in chapter 5 - in essence they attend the preservation of explanatory transparency - however, the inevitable concession is that the omission of behaviours which patently *are* evident in real animal populations (for example fighting and prey-handling) means that evaluating the present model on its predictive accuracy with regard to biological field data may not be particularly informative.

### 7.3 Results

Agents evolved in conditions *A*, *B*, and *C* shall from now on be referred to as *A*-agents, *B*-agents and *C*-agents respectively. Given that the primary aim of this model was to compare conditions *A* and *B*, a discussion of the sensitivity condition *C* will be deferred until section 7.3.3. The results presented below derive from 8 complete repetitions of both evolution and testing phases of each condition, from which means and standard deviations of intake rate at each group size (in each condition) were calculated.

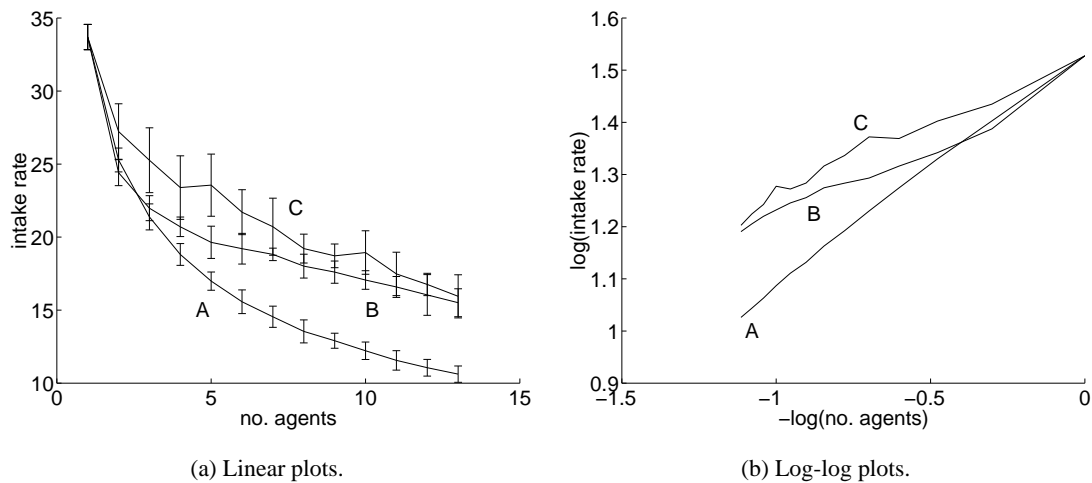


Figure 7.3: Interference functions for conditions *A*, *B*, and *C*. Each data point is the average of 8 runs of each condition. In (a) error bars represent standard deviations, and in (b) logarithms are taken to base 10.

### 7.3.1 Interference functions

Figure 7.3 illustrates the interference functions obtained in conditions *A*, *B*, and *C*, with the linear scale of 7.3(a) permitting error bars to be superimposed, and with 7.3(b) charting the log-log relationships (lines of best fit for this figure are given in the first column set of table 7.1). In figure 7.3(b) there is a strikingly clear linear relationship for *A*-agents between the logarithms of agent density and intake rate (and a corresponding smooth curve in figure 7.3a, just as first identified in the field by Hassell & Varley, 1969). The slope of this line indicates a constant level of interference,  $m = 0.45$ , and the antilogarithm of the *y*-intercept recovers the interference free intake rate, 34.67. (This terminology derives from Hassell and Varley's original model as described in section 7.1.1. The interference free intake rate is simply the predicted intake of a forager in the absence of competition.)

For *B*-agents, this linear relationship is no longer apparent. Not only is the observed level of interference generally lower than for *A*-agents, but this difference becomes more apparent at higher agent densities. Two observations follow immediately from these results. The first is that interference - both constant and density dependent - can indeed be effectively modelled using only the sensorimotor interactions of groups of agents, without recourse to explicit behavioural states and associated transition rules. Secondly, that somehow, *B*-agents - evolved in the presence of conspecifics - are deploying behavioural strategies which allow them to forage more effectively in the presence of conspecifics than those (*A*-agents) evolved in isolation.

Figures 7.4 and 7.5 illustrate the foraging performance of *isolated* agents from each condition, with the abscissa representing the group sizes within which the agents were evolved (for conditions *B* and *C*). Condition *A* will of course provide only a single point on these graphs, but this is extended into a line in order to aid comparison. The second column set of table 7.1 summarises the lines of best fit for figure 7.5. It is immediately evident that the larger the group that a given *B*-agent was evolved with, the worse it performs when on its own. The same behavioural strategy that delivers an advantage over *A*-agents in group situations is a handicap for isolated agents.

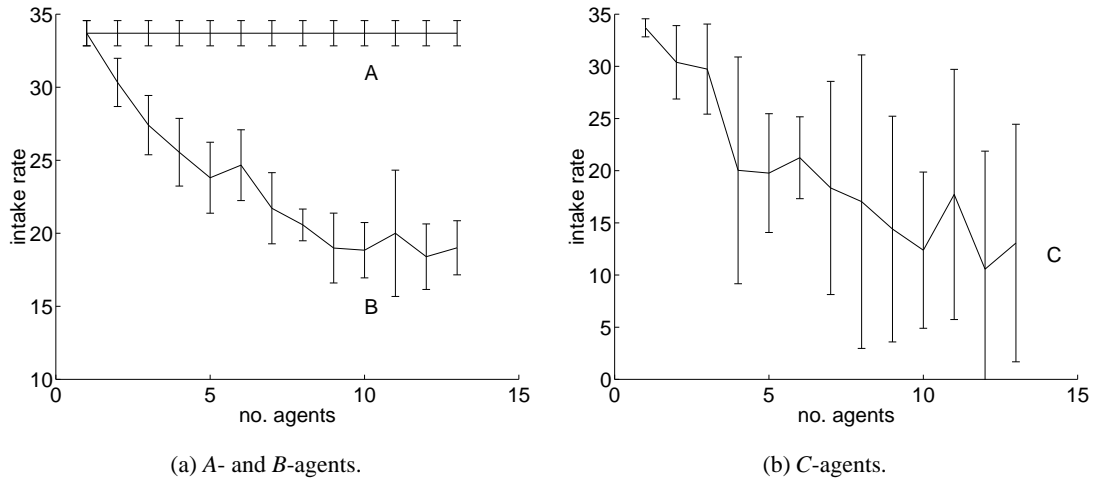


Figure 7.4: Intake rates of isolated A-, B-, and C-agents; A- and B-agents are illustrated in (a), and C-agents, for clarity, in (b). The abscissa of each plot represents the groups sizes within which the B- and C-agents were evolved (condition A is actually a single point, but is extended into a line to aid comparison). Each data point is the average of 8 runs of each condition, and error bars represent standard deviations.




Figure 7.5: Intake rates of isolated A-, B-, and C-agents. The abscissa represents the groups sizes within which the B- and C-agents were evolved (condition A is treated as in figure 7.4). Each data point is the average of 8 runs of each condition. Logarithms are taken to base 10.

condition	interference		isolation	
	best-fit line	error	best-fit line	error
A	$y = 0.45x + 1.54$	0.007	$y = 1.53$	0.000
B	$y = 0.27x + 1.49$	0.018	$y = 0.25x + 1.55$	0.021
C	$y = 0.27x + 1.54$	0.017	$y = 0.43x + 1.59$	0.070

Table 7.1: This table summarises the log-log plots of figure 7.3(b) (interference) and figure 7.5 (isolated foraging), giving lines of best fit and errors. For the first column set the slope of each line represents the average interference level and the antilogarithm of the y-intercept recovers the interference free intake rate, and for the second column set the slope represents the extent to which isolated intake declines as a function of group size during evolution. All lines are fitted to the averages of 8 evaluations of each condition, and each error measure specifies the range around any point which contains at least 50% of the data points.




Figure 7.6: Average  $\bar{cl}$  for groups of 2-13 agents from A, B, and C conditions. Each line represents the average of 6 runs of each condition, with each run having provided the  $\bar{cl}$  measure as described in the text.

### 7.3.2 ‘Clumpiness’

What is it about the behaviour of *B*-agents that enhances their performance (relative to *A*-agents) when in a group, but degrades their performance when in isolation? A first observation of *A*-agents suggests that they all follow the very simple and entirely intuitive strategy (given their sensorimotor poverty) of making for the nearest food item as rapidly as possible. However, further observation suggests that groups of agents following this strategy tend to become ‘clumped’ together. By contrast, groups of *B*-agents, despite suffering the same sensorimotor poverty as *A*-agents, appear to distribute themselves more evenly throughout the patch. This section examines the possibility that it is this avoidance of clumping that both lessens the toll of interference in group situations and leads to inefficient foraging by isolated agents.

It is possible to be more formal with the notion of ‘clumpiness’. Let us define a measure of the ‘instantaneous clumpiness’ of a group of agents as  $cl_t \in [0.0, 1.0]$  such that  $cl_t = 0.0$  indicates a maximally dispersed group, and  $cl_t = 1.0$  indicates a maximally clumped group. Given  $d_i$  the distance from each (alive) agent  $i$  to the nearest (alive) conspecific,  $r_a$  the number of alive agents, and  $f(\cdot)$  a Gaussian function of height 1.0, mean 0.0, and radius 15.0, then:




Figure 7.7: Average  $cl_t$  of a group of 8 *A*-agents and a group of 8 *B*-agents. Each line represents the average  $cl_t$  of the group at each of 800 time-steps over 500 evaluations, standard deviations are shown. Note that the average  $cl_t$  depicted here is not to be confused with  $\overline{cl}$ ; the former is the average *instantaneous* ‘clumpiness’ over many evaluations, the latter is the average ‘clumpiness’ over many time-steps of a *single* evaluation.

$$cl_t = \frac{1}{r_a} \sum_{i=1}^{r_a} f(d_i).$$

The ‘overall clumpiness’  $cl$  of the behaviour of a group of agents is then the average  $cl_t$  over all time-steps for which at least one agent was alive (note that the ‘instantaneous clumpiness’  $cl_t$  of a single agent ‘group’ is zero). This metric was used to compare the clumpiness of groups (of evolved agents) of sizes ranging from  $n = 1$  to  $n = 13$  from each condition. Each group was evaluated 500 times, with a final clumpiness  $\overline{cl}$  derived from the average  $cl$ . Figure 7.6 indicates, as hypothesised, that groups of *A*-agents present much higher  $\overline{cl}$  values than groups of *B*-agents. An example is illustrated in figure 7.7, which contrasts the clumpiness profile (the average of 500 evaluations of 800 time-steps each) of a group of 8 *A*-agents with a group of 8 *B*-agents. Both groups, being initially randomly distributed, display an initially low average  $cl_t$ , but for the majority of the evaluation the *B*-agents present much the lower average  $cl_t$ .

We must now ask how, given the sensorimotor poverty of the agents, does this difference in propensity to clump come about, and also how this could lead to a difference in the level of interference. A little more observation reveals that after an encounter with a food item, a *B*-agent will decelerate dramatically for a few time-steps, before accelerating away towards another target (an example is illustrated in figure 7.8b). By contrast, *A*-agents rarely deviate from maximum velocity (figure 7.8a). This in itself is enough to ensure that the *B*-agents become more dispersed. Consider two agents  $p$  and  $q$  heading towards a single food item from a single starting point. Both agents will remain in close proximity as they approach the item, but only one of them, let us say  $p$ , can benefit from it. The item now disappears and both  $p$  and  $q$  perceive, and begin to move towards, a new nearest item. If they both travel always at maximum speed (*A*-agents), they will again remain in close proximity as they approach this new item and, again, only one will benefit. However, if  $p$  slows down following consumption of the first item, then the situation is different




Figure 7.8: Forward speed profile of isolated *A*- and *B*-agents over one evaluation of 800 time-steps. The circles represent encounters with food items. The *A*-agent does not decelerate rapidly following encounters, the *B*-agent does.

and can be seen to work to the benefit of both agents, especially if they coexist with several others. From the perspective of any individual agent the environment will now be less stable, with the nearest food item suddenly changing location as other agents move around and consume. Now, if  $p$  slows down after consumption, it becomes possible that after a few time-steps  $p$  and  $q$  will perceive *different* nearest food items, and so will no longer be in direct competition. And the more dispersed a group is, the more unpredictable the pattern of food depletion will be, and hence the greater individual benefit there is to be gained from being part of a dispersed group. Of course the strategy of slowing down after consumption is clearly *not* the optimal strategy for an *isolated* agent for whom any deviations from maximum speed can only serve to reduce intake rate.

It is important to stress that this strategy difference between *A*- and *B*-agents is reliable, and is not dependent on whether agents are analysed in isolation or in groups. Figure 7.9 shows average maximum deviations from maximum speed across for evolved agents across 8 runs of each condition, with agents tested both in isolation and in groups of 10 conspecifics. In all cases, *A*-agents maintain relatively stable (and high) speeds, with *B*-agents consistently exhibiting large deviations from maximum speed.

### 7.3.3 Sensitivity: perceiving conspecifics

There are many systems in which agents will not, or need not, suffer the same degree of sensory poverty as do *A*- and *B*-agents. How might the above pattern of results be affected with agents that are able to perceive each other? Recall figure 7.3, which illustrates that the interference levels experienced by such agents (*C*-agents) are the lowest of any condition, although the difference is not as dramatic as that between *A*- and *B*-agents. And the performance of isolated *C*-agents is similar to that of *B*-agents but erratic (notice the large errors in figure 7.4b), a simple reason for this being that some isolated *C*-agents are unable to forage at all, dependent as they often are on receiving some sensory input from conspecifics. *C*-agents also display the least proclivity for forming clumps (figure 7.6).




Figure 7.9: Speed ranges for conditions *A*, *B*, and *C*. Each column shows average and standard deviation of maximum deviation from maximum speed. Agents from each condition are tested both in isolation and in groups of 10. Maximum possible speed is 10.




Figure 7.10: Forward speed profile of an isolated *C*-agent over one evaluation of 800 time-steps. The circles represent encounters with food items. The agent decelerates rapidly following each encounter.

These results are perhaps not surprising since it is reasonable to suppose that agents that can perceive each other stand a better chance of avoiding each other than agents that cannot perceive each other. However, figure 7.10 illustrates that the speed profile of a single *C*-agent is similar to that of the *B*-agent in figure 7.8(b), if less extreme, and figure 7.9 suggests that this similarity is reliable. It *cannot* therefore be said, on this evidence, that *C*-agents are reducing interference by means of a qualitatively different behavioural strategy to the *B*-agents. Nevertheless the general result remains that *C*-agents, with their enriched *Umwelt*, demonstrate the same phenomenon of (near) optimal group behaviour and suboptimal isolated behaviour as do *B*-agents.

### 7.3.4 Sensitivity: varying group sizes during evolution

One last question to ask is what happens if agents are evolved in environments in which they are sometimes on their own, and other times surrounded by varying numbers of conspecifics, this being a situation with perhaps a greater natural plausibility than the previous conditions analysed. In particular, would such agents be able to switch between maintaining a continuous speed and decelerating after consumption depending on the prevailing density of conspecifics?

Two final conditions were analysed: condition *D*, identical to condition *A* with the exception that each genotype is evaluated with each of 0, 2, 4, and 6 clonal conspecifics each generation (in a random order), and condition *E*, identical to *D* with the exception that agents - like those in condition *C* - are able to perceive the presence of conspecifics. Note that the distinction remains between these new conditions and conditions *B* and *C* that, like condition *A*, they do *not* require distinct GAs for each (tested) group size.

Figure 7.11 illustrates the interference functions derived for conditions *D* and *E*, superimposed over those originally derived for conditions *A*, *B*, and *C*. It is clear that there are considerable similarities between conditions *B* and *D* (and also between *C* and *E*), suggesting that the differences in interference originally observed between conditions *A* and *B* have less to do with the precise densities



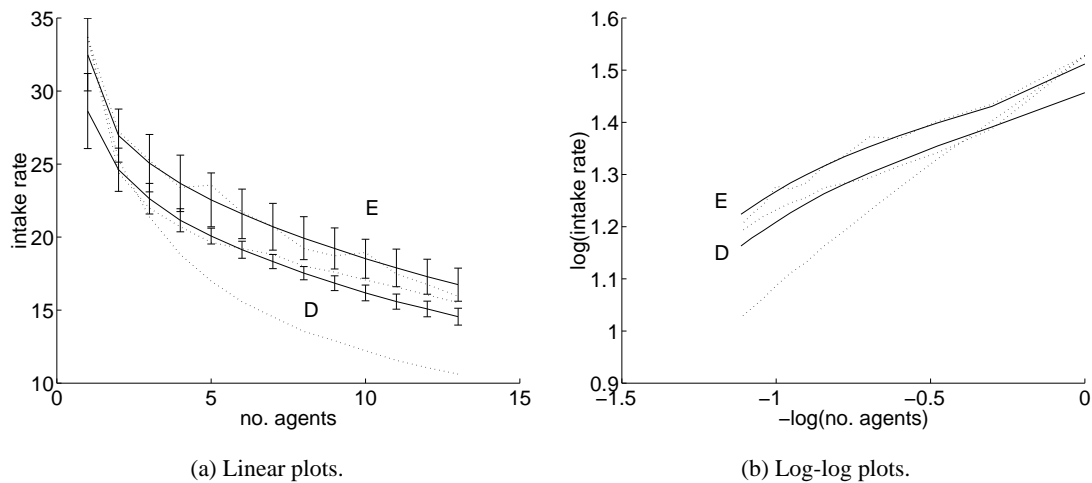


Figure 7.11: Interference functions for conditions *D* and *E*. Each data point is the average of 8 runs of each condition. The dotted lines represent the interference functions derived previously for conditions *A*, *B*, and *C*, illustrated in fig 7.3. In (a) error bars represent standard deviations, and in (b) logarithms are taken to base 10.




Figure 7.12: Speed ranges for conditions *A* – *E*, to be interpreted as figure 7.9.

of conspecifics present during evolution, and more to do with the very fact that they *are* present. This supports, and attests to the robustness of the conclusion that agents evolved in the presence of conspecifics experience lower levels of interference than those evolved in isolation.

More pertinently, figure 7.12 extends the analysis of deviation from maximum speed (in figure 7.9) to encompass conditions *D* and *E*. The first thing to notice is that *D*-agents follow the same pattern as *B*- and *C*-agents; that is, they exhibit large deviations from maximum speed whether in isolation or in the presence of conspecifics. *E*-agents, by contrast, are able to modify their behaviour according to prevailing conditions. When on their own, they generally move at steady speeds (like *A*-agents), when in groups, they frequently slow down (like *B*-agents). Not only does this exemplify the flexibility of what is a strikingly simple internal mechanism, it also buttresses the argument that consistent deviation from maximum speed is a good strategy when foraging amongst conspecifics (in the context of this model).

In contrast to the previous analyses of this chapter, the ability to perceive conspecifics is in this

case highly significant. *D*-agents - blind to their presence - have no way of adjusting to prevailing conditions, and do not; *E*-agents can, and do. So, although direct sensitivity to conspecific density can enhance foraging performance (the difference between conditions *D* and *E*), such sensitivity is only selected for if varying conspecific densities are experienced during evolution (the difference between conditions *C* and *E*). Note that this point is entirely separate from the *indirect* sensitivity to conspecific density (via intake rate) that, through interaction with the evolutionary process, leads to the original differences observed between conditions *A* and *B*.

## 7.4 Discussion

### 7.4.1 Modelling interference

The contributions of this chapter may be assessed both with respect to the interference function itself, and also in terms of the themes of this dissertation as a whole. With respect to the former, several observations may be made. First, we have seen how interference, both constant and density dependent, can be modelled using an agent-based (IOS) model operating at the level of perception and action in a spatially defined environment, without being explicitly associated with aspects of the agent mechanism. Interference in the present model is a consequence of agents behaving (near) optimally in a spatially explicit limited resource environment, interacting with each other via their shared environment.

Second, optimality has been incorporated into an agent-based approach without making great demands on the computational agility of the agents themselves, in significant contrast to Stillman et al. (1997). This has been achieved by means of the IOS methodology applied in conjunction with a simple internal mechanistic substrate, as opposed to the (orthodox) methodology of ‘building in’ mechanisms of optimality which supervene on pre-determined repertoires of internal behavioural correlates and transitions rules. (It may also be the case that the location of the ‘behaviour-generating mechanisms’ of agents in a sensorimotor space that is continuous but of few dimensions, is, in some sense, easier to optimise than the discrete, many-dimensional behavioural state-spaces of the individual-based model of Stillman et al. and its predecessors. If so, it can then be asked whether this may account for the relative simplicity of the ‘sensorimotor mechanism’. These questions are, for the present, left open.)

These first observations capture the most immediate implications of this model for theoretical biology, and it is worth specifying exactly where they lie. The model does *not* challenge any particular assumption or prediction of orthodox models, rather it suggests a different way of building models of interference that avoids some of the conceptual pitfalls associated with orthodox approaches. Most obviously, the call of Van der Meer and Ens (1997) for a model of interference consistent with the principles of optimality (in terms of individual rate-maximisation) is answered. Beyond this, as noted above, interference is properly treated as a behavioural level phenomenon rather than, improperly, as a (mechanistic) pre-determined component of the model, and, as part of this, individual rate-maximisation can be modelled without resorting to overly complex and architecturally dubious internal mechanisms.

Moving on, perhaps the most striking observation of all is the clear linear form of the log-log interference function of *A*-agents (figure 7.3b), both for its congruence with the original observations of Hassell and Varley (1969), and for its contrast with conditions *B-E*. To understand

this contrast, one may consider condition *A* as a reductionist ideal; the (optimal) behaviour of a single agent is first determined and the behaviour of the group is then derived from the independent (additive, linear) combination of the behaviours of all constituent agents. Conditions *B-E*, by contrast, do not permit such a conceptual separation of individual and group behaviours. In these conditions, patterns of agent-environment interaction contingent on group behaviour (for example resource instability patterns) can be present *during* optimisation, and hence can influence the trajectory and outcome of the optimisation process, in terms of both individual and group behaviours. Such patterns may be considered to represent ‘behavioural historical constraints’ insofar as they constitute dynamical invariants which constrain the future dynamics of the system as a whole, and which are able to alter their own conditions of realisation (see chapter 5, and also Di Paolo, 1999, p.52). Therefore, whilst condition *A* treats the process of optimisation and the outcome as entirely separable, conditions *B-E* allow for historical constraints to arise and persist from their continual interaction. Perhaps, then, the idealised linear interference function observed in condition *A* is indicative of a rigidly reductionistic, ahistorical approach to group foraging, an approach from which dissent would not have been possible under orthodox modelling regimes that necessarily deny the potential for historicity inherent in conditions *B-E* of the present (IOS) model. This carries clear implications for theoretical biology; it may well be that historical processes of this kind are essential in understanding the dynamics of interference in real situations. Of course it remains an open question whether the departures from ‘ideal’ interference that have been observed since Hassell and Varley (see, for example, Ruxton et al., 1992; Moody & Houston, 1995) can be attributed to the kind of historical processes elucidated here. It should also be mentioned that, in line with the IOS methodology outlined in chapter 5, no extra explanatory burden need be placed on the GA other than as a process of (dynamic) optimisation, since in neither condition can the mechanics of the GA itself influence the fitness values awarded to the constituent genotypes.

Lastly, it has been observed that agents evolved with conspecifics experience lower levels of interference (and hence higher intake rates) than those evolved in isolation and subsequently tested with conspecifics, and conversely that agents evolved with conspecifics experience *lower* intake rates when tested in isolation than those evolved in isolation. A strategy was identified (deceleration after consumption), suboptimal for the isolated individual, yet optimal for the individual in a group of conspecifics. Agents evolved under conditions of varying conspecific density, and able to perceive the presence of these conspecifics, were able to switch between this strategy and a continuous speed strategy depending on prevailing conditions. In the domain of artificial agents these observations are of clear importance, whether for the purposes of design or of prediction of spatial distributions, broadly construed. Whether or not they will be relevant to understanding the behaviour of real organisms must remain an open question, although the fact that most encounters with food items require a certain ‘handling time’ suggests that, in many real situations, slowing down upon consumption is not something that can be avoided.

It is important to note that the hypothesis that adaptation to group situations can lead to isolated individual suboptimality is easily supported and entirely uncontroversial in groups in which there is a necessarily heterogeneous division of tasks. For example in colonies of meerkats, *Suricatta suricatta*, duties of vigilance, foraging, nursing, and reproduction are divided amongst individuals (Clutton-Brook & Harvey, 1977), and an isolated meerkat engaged in just one of these tasks - for

example vigilance - is not likely to be well adapted to the challenges of its environment.<sup>4</sup> However, the present context is different. The models of this chapter, and also of the following two chapters, are directed towards situations in which groups are faced with only a single task (collecting food), and, moreover, in which individuals are *homogeneous*; that is, they must all subscribe to identical functional and operational descriptions. These are situations in which it cannot be trivially asserted that adaptation to group situations will lead to individual suboptimality.

#### 7.4.2 A broader view

The most obvious way in which this chapter engages with the general themes of this dissertation is in illustrating a concrete application of the IOS methodology to a problem of contemporary concern in theoretical biology, with all the various consequences described above. However, the engagement runs deeper than a demonstration of methodological efficacy. In particular, the brief history of interference modelling described in section 7.1.1 makes it clear that the present model is motivated by a confusion between behavioural and mechanistic levels of description instilled at the heart of most orthodox models of interference, whether individual-based or ‘phenomenological’.

The issue of behaviour coordination, set out in some detail in the preceding chapter, has been treated with a different slant here. Rather than looking at ‘how to do the right thing’, the present chapter has explored an instance of apparently suboptimal behaviour. This exemplifies an important part of the functionalist projects of optimal foraging theory, evolutionary psychology, and IOS modelling: the interpretation of suboptimal behaviour in terms of adaptation to impermanent environmental features. However, behavioural choice itself and the mechanisms underlying the generation of choice behaviour have received less attention, with the exception of *E*-agents and their ability to ‘switch’ between continuous speed and variable speed strategies depending on perceived conspecific density. The focus of this dissertation will return to behavioural choice and its mechanistic basis more directly in the following chapter.

The discussion of the environmental complexity thesis of chapter 2 is relevant to the present chapter in several ways, not only through its methodological connection with the IOS modelling strategy, but also by virtue of issues of ‘construction’. Recall that ‘narrow’, or ‘causal’ construction involves agents making changes to the ‘external’ environment *sensu* Brandon (1990). The patterns of resource instability evoked by multiple foraging agents may be considered to be instances of this kind of construction, and in conditions *B-E* of the present model such construction plays an essential role in determining the form of the near-optimal behaviour. Recall also from chapter 2 Lewontin’s (1983) critique of the concept of ‘ecological niche’, in which - in view of the various forms of construction - he questions how any set of environmental features, taken to provide a functional mandate for an observed behaviour, can pre-exist the activity of the behaving agent itself. The response to this critique, in chapter 2, was to allow that agents certainly do influence their environments by their activity, but to suggest that this should form part of the explanatory duty of models that incorporate mutual specification between agent and environment. The present model has provided a clear example of this strategy in terms of the ‘behavioural historical constraints’ constituted by the resource instability patterns.

<sup>4</sup>At the extreme, consider eusocial groups such as the Hymenopteran order of insects (Wilson, 1971), or the naked mole rats *Heterocephalus glaber* (Jarvis, 1981), in which a majority of sterile individuals altruistically support a reproductive élite.

The other side of construction, when discussed in the context of the environmental complexity thesis, is of course *constitutive* construction; the translation of external environment into *Umwelt*. One can interpret the distinction between conditions  $B(D)$  and  $C(E)$  in the present model as investigating the influence of altering the *Umwelt* of the agents; recall that  $C$ - and  $E$ -agents were able to perceive conspecifics, but that  $B$ - and  $D$ -agents were not. The most salient observation here was that this extra ability, when combined with an evolutionary history involving variable conspecific densities (condition  $E$ ), led to the ability of agents to modify their behaviour in response to prevailing conditions. Constitutive construction is here involved in *lifetime* adaptation, in the switching between established behaviour patterns, but not in the constitution of the patterns themselves, which are largely the same as observed in conditions  $A-D$ . It may also be argued that the behaviour of  $B-E$ -agents is in some sense more ‘complex’ than that of  $A$ -agents, in virtue of deploying ‘adaptive decelerations’; the environmental complexity thesis may suggest that this be attributed to the relatively complex environment presented by the evolutionary phases of these conditions. Indeed, arguably the most complex behaviour of all - that of  $E$ -agents - derives from arguably the most complex environment of all (varying densities of observable conspecifics).<sup>5</sup>

Nevertheless, despite the intriguing nature of  $B-E$ -agents, it is perhaps the ‘ideal’ interference function of condition  $A$  that remains the most striking observation of this chapter. Such ‘ideal’ functions, with exponentially arranged curves and linear logarithm plots, are a common currency in contemporary science. And so to end this discussion on a speculative note: perhaps the significance often attributed to such functions derives in part from the tendency of simple modelling strategies (whether analytical or otherwise) to furnish the researcher with these functions in the first place. New strategies of modelling - such as the IOS methodology - may be needed to elucidate the complex dynamical and historical processes which constitute the less than ideal real situations we wish to understand.

## 7.5 Summary

This chapter has presented an IOS model of interference; the relation, well described in theoretical biology, between the intake rate of a focal foraging agent in a single ‘patch’ and the density of agents in the patch. The present model extends orthodox modelling strategies in several significant ways. It answers the call for interference models interpretable in terms of optimality, but *without* requiring that internal structure be framed in the conceptually dubious terms of internal behavioural correlates and complex supervening arbitration mechanisms. Along the same lines, it respects the status of interference as a behaviour, arising from agent-environment dynamics, rather than as an *a priori* component of agent mechanism. Finally, it encompasses behavioural historical constraints, and by virtue of this has revealed an example of a behaviour, suboptimal for the isolated individual, yet optimal for the individual in a group of conspecifics. It is hard to see how these kinds of insights could have been arrived at from within the orthodox modelling community described at the outset of this chapter. This may explain why field biologists have not (to my knowledge) explicitly looked for the kind of suboptimal behaviour that has been described

<sup>5</sup>In the terminology of Emery (1967) the environment of  $B-E$ -agents may be considered to approach ‘turbulence’, in contrast to the ‘placid’ environment of  $A$ -agents; see chapter 3. See also Floreano (1993) for a relatively early example from the SAB/AL canon concerning the differences in behaviour observed when agents in a shared environment could, or could not, see each other.

here: hopefully there is now some motivation for them to do so.

The present model nevertheless deals only with very simple interference situations. It should, however, be extensible to support the exploration of some of the more complex situations of interest to biologists, for example groups in which there are individual competitive differences (Milinski & Parker, 1991; Holmgren, 1995). On the other hand, the implications of the model are by no means restricted to biology; an understanding of the dynamics of interference is of interest any situation in which groups of agents - biological or artificial - pursue a shared goal - cooperatively or competitively - in a limited resource environment. Indeed, the structure and dynamics of the model as it stands (and in particular specific interaction patterns such as those associated with ‘clumping’ and its avoidance) may well have closer analogues in artificial situations than in biological contexts.

The next chapter returns the focus to behavioural choice, extending the IOS model of interference described above to incorporate multiple resource types in a single patch (as in the original behaviour coordination model of chapter 6). Of central importance in this investigation is the ‘matching law’; the observation that animals and humans often distribute their behaviour in proportion to the reward they obtain from each choice alternative. The primary concerns are with (a) demonstrating that the same sensorimotor interactions constitutive of interference can also underlie matching (thereby questioning the assumption of a dedicated ‘mechanism’ of matching), and (b) using the relations between matching and the ‘ideal free distribution’ (introduced in section 7.1) to illustrate another context in which isolated individual suboptimality can be understood in terms of optimal group behaviour. The role of interference in determining these relations in the more usual context of the distribution of groups of agents across *multiple* patches will be a focus of the final empirical installment of this dissertation, chapter 9.

## Chapter 8

### The matching law

---

In 1961, the psychologist R.J. Herrnstein famously observed that pigeons *match* the frequency of their response to different stimuli in proportion to the reinforcement obtained from each stimulus type (Herrnstein, 1961). This observation has been found, over subsequent years, to be very general; pigeons pecking at a pair of differently coloured discs, rats pressing one or another lever, even humans faced with a variety of buttons to choose between, will all distribute their responses in proportion to the reward obtained from each option (Davison & McCarthy, 1988; Herrnstein, 1997). This chapter, and the following chapter, concentrate on two questions that arise from this ‘matching law’. First, why ‘match’ at all? The matching law is descriptive rather than normative; matching is not necessarily optimal, and indeed there is considerable evidence to indicate that when faced with a situation in which matching and maximisation (of reward) are incompatible, isolated individual animals (including humans) will tend to match and not maximise (Mazur, 1981; Herrnstein, 1990, 1997). The experimental psychology literature has approached this question with various attempts to redefine that which is being optimised, for example replacing maximisation of intake rate with minimisation of inter-reward intervals (Staddon, 1992). However, no such alternative ‘currency’ has been justified independently, and for every candidate so far proposed, empirical exceptions have always been found.

The second question concerns the behaviour-generating mechanism(s) that might underlie matching. In the psychology literature there is not even consensus over the level at which such a mechanism should be interpreted. Some authors argue that matching is the product of ‘underlying learning rules’, others that matching itself is the rule followed (Williams, 1994, provides a review of these debates).

A focus on matching marks a return to the issue of behaviour coordination, described in detail in chapter 6, and most often associated in SAB/AL with the problem of ‘action selection’ (Tyrrell, 1993; Maes, 1994). However, whilst action selection analyses choices between alternatives that satisfy *distinct* requirements (for example feeding and sleeping), matching concerns situations in which choice operates between different ways of satisfying the *same* requirement. In what follows, an extension of the IOS ‘interference’ model of the previous chapter is used to explore the hypothesis that matching (and its potential suboptimality) is a consequence of optimal foraging in a *shared* environment, and also to suggest that some instances of matching behaviour can arise

without there being any dedicated ‘mechanism of matching’. Motivating both of these objectives (but in particular the former) is the idea - first explored in the previous chapter - that *isolated individual* suboptimal behaviour can, in some cases, be understood in terms of the operation of mechanisms adapted to *group* situations.

### 8.1 Individual choice and group choice

The idea that individual matching might follow from patterns of choice behaviour adapted to a group context derives from a series of analogies between the matching law and the ideal free distribution (hereafter the IFD). Recall from chapter 7 that the IFD describes the equilibrium distribution of rate-maximising foragers across a patchy distribution of resources, such that no forager can profit by moving elsewhere, regardless of the local resource quality. Even though the IFD and the matching law derive from different disciplines - the IFD from behavioural ecology and the matching law from experimental psychology - the analogies are striking. First, the matching law is to do with individual choice and the IFD is to do with the collective consequences of individual choice. Second, the laboratory environments employed by experimental psychologists are often interpreted as abstractions of natural foraging environments (Dallery & Baum, 1991; Shettleworth, 1988). Third, there is an evident congruence in their mathematical forms (Gray, 1994; Baum & Kraft, 1998). Consider the ‘generalised matching law’ (Baum, 1974):

$$\log \frac{B_A}{B_B} = s \cdot \log \frac{R_A}{R_B} + \log(b).$$

in which  $B_A$  and  $B_B$  represent the rates of response to options  $A$  and  $B$ , and  $R_A$  and  $R_B$  represent the rewards obtained from  $A$  and  $B$  ( $s$  and  $b$  are ‘sensitivity’ and ‘bias’ constants). Under this law, each choice alternative receives a proportion of the total response according to the proportion of total reward it provides. Compare this with the ‘generalised habitat matching law’ expressed in the context of the IFD (Fagen, 1987):

$$\log \frac{N_A}{N_B} = \frac{1}{m} \log \frac{F_A}{F_B},$$

in which  $N_A$  and  $N_B$  represent the number of predators on patches  $A$  and  $B$ , and  $F_A$  and  $F_B$  represent the prey availabilities on patches  $A$  and  $B$  ( $m$  is the interference constant). Under this law, each patch in a patchy environment will receive a proportion of the total number of predators according to the proportion of prey it provides.

As striking as these analogies may be, several profound differences remain. As stated before, whereas IFD predictions are normative, the matching law is an observed relation, noted for not always being optimal. Also, whereas the habitat matching law of the IFD is expressed in terms of *available* resources  $F_i$ , the individual matching law is expressed in terms of *obtained* resources  $R_i$  (I will return to this distinction later). Even the mathematical congruence illustrated above may be less surprising than it might appear at first glance; many situations in which ratios are to be related to each other will be describable using equations of this form.

Nevertheless, when considered together these analogies remain seductive, and they have been taken, on several occasions, to suggest that there might exist a single behavioural dynamic that can both lead populations to the (optimal) IFD, and also cause individuals to adhere to the (potentially



suboptimal) matching law (Houston, 1986; Gallistel et al., 1991; Thuisjman et al., 1995). This is the idea taken up in both this and the following chapter, however, whereas chapter 9 interprets the idea in the direct manner suggested above, the remainder of the present chapter explores a rather unorthodox interpretation, described below.

Before we continue, it must be said that although there are extensive literatures treating both matching and the IFD separately, only a surprisingly small number of studies directly address their relationship. Those of Houston, Gallistel et al., and Thuisjman et al. represent some of the few that do, most of which consider the appropriate level of analysis to be that of ‘switching rules’ that individuals may use to decide whether or not to leave a patch; such rules can also be applied to switching between two choice options in a ‘Skinner box’. This remains a useful perspective, and indeed is the level of analysis that will be adopted in the following chapter. For now, it will suffice to say that although there are indeed switching rules which can support both matching and the IFD, they remain poorly characterised in the modelling literature (see in particular Thuisjman et al., 1995), and field observations have failed to reach consensus over their use in real situations (Gray, 1994; Baum & Kraft, 1998). Chapter 9 will contribute to the modelling enterprise, if not to the accumulation of data from the field.

The present chapter approaches the relationship between the matching law and the IFD in a very different manner, focusing on patterns of behaviour coordination in a *single* patch. This novel approach retreats from direct engagement with the IFD itself, but, at the same time, returns the issue of *interference* to centre stage.

If, as is orthodox, the analogy between the IFD and the matching law is interpreted at the level of patch-switching rules, then interference - the decline in intake rate due to the presence of competing predators - is relevant only insofar as it influences the relative values of the different patches. The primary responsibility for individual matching would remain, in such cases, with the switching rules themselves, in their use of these values in determining the distribution of behaviour. In the context of a single patch, the situation is markedly different. Consider a single patch containing both rich and poor resources. In such a (heterogeneous) patch, there will be the opportunity for agents foraging in the presence of conspecifics to experience different levels of interference with respect to each resource type, and in order to maximise overall intake, to adopt a distribution of responses across resource types rather than exclusively concentrating on the rich resource. This suggests an informal ‘single patch’ interpretation of Fagen’s habitat matching law: each resource type (rather than each patch) may sustain a level of attention from agents as a function of its relative value. Although this interpretation lacks the formal background of Fagen’s law, it nevertheless suggests that the analogies between the IFD and the matching law may continue to hold in the context of a single patch. The questions asked in this chapter are whether groups of agents in such heterogeneous single patches do in fact distribute their responses in this way, and whether, as individuals, they adhere to the matching law in doing so. It turns out that both questions can be answered with a qualified ‘yes’, and the importance of this turns on the suggestion that some instances of matching behaviour can arise from the same sensorimotor interactions that give rise to interference; no supervening ‘mechanism of matching’ is necessarily required.

These ideas clearly require a model of interference consistent with the principles of optimal

foraging, and therefore build on the arguments, and the model, of the previous chapter. It is important to note, however, that this chapter does *not* directly address the suggestion that a single mechanism underlies both the IFD and the matching law, simply because no single patch model can claim to address the IFD itself. The role of the IFD here has instead been to motivate and provide context for postulating a link between individual matching and optimal group behaviour within a single, heterogeneous patch. As noted above, a relatively straightforward investigation of the relations between matching and the IFD itself is left for the following chapter.

## 8.2 An IOS model of within-patch matching

The model in this chapter utilises a heterogeneous (single patch) environment containing *two* resource types  $X$  and  $Y$ , in which items of type  $X$  are (usually) in some sense more ‘valuable’ than those of type  $Y$  (‘value’ will be defined more carefully in section 8.2.2). The response rate (or, synonymously, encounter rate) of a foraging agent in such environments can therefore be described as a distribution over  $X$  and  $Y$  items, and one can imagine a spectrum of such distributions, from indifference between  $X$  and  $Y$  through to exclusive choice of either  $X$  or  $Y$ .

The extremes of this spectrum fit what is known as the ‘zero-one’ foraging rule, which has long been established as an optimal foraging rule for a single agent in the context of the ‘prey’ model (Charnov & Orians, 1973; Maynard Smith, 1974). Agents in this model are presented with a succession of prey items of different types, and upon each presentation must decide whether to consume (respond to) the prey or continue searching. The optimal policy in a simple formulation of this model, the zero-one rule, is that a given prey type is either always consumed or always ignored upon presentation.<sup>1</sup> Intermediate points on the spectrum, on the other hand, would reflect a distribution of response that in some way mirrored the relative values of  $X$  and  $Y$  and so would be describable in terms of (generalised) matching. With this background in mind the primary hypothesis of this chapter can be accurately stated. Isolated agents evolved in a heterogeneous environment should behave in accordance with the zero-one rule, but agents evolved with conspecifics should adopt some intermediate response distribution describable in terms of matching. This hypothesis derives from the intuition that as agent density increases there is more to be gained from paying attention to the less valuable resources, due to the relatively high interference levels likely to be associated with the more valuable resources. Validation of this hypothesis would in turn support the following claims. First, that some instances of matching behaviour can be understood as a consequence of choice behaviour adapted to a group context. Second, that the same behaviour-generating mechanisms that support zero-one behaviour can also support matching behaviour. And third, a related claim, that matching behaviour can arise from the same sensorimotor interactions that give rise to interference, without necessarily requiring a dedicated supervening ‘mechanism of matching’.

Some terminological clarification may be useful at this point. In this chapter the terms ‘encounter’ and ‘response’ are used interchangeably, signifying the physical overlap of an agent with an item (or another agent). The ‘intake rate’ of an agent is, by contrast, the actual energy, re-

<sup>1</sup>The standard ‘prey’ model assumes that the agent is presented with (by the environment, it is assumed) prey types (indexed by  $i$ ) at rates  $\lambda_i$ , that each prey type has an expected handling time  $h_i$ , and an expected net energy gain  $g_i$ . The output of the model is the probability that items of type  $i$  will be consumed upon presentation.

ward, or resources accrued by an agent over time. In the previous chapter, intake rates and encounter/response rates (with/to resource items) were identical, since every encounter/response yielded a fixed reward. However, this need not always be the case. If, as in this chapter, encounters/responses do *not* always yield the same reward, the intake rate with respect to a resource type may differ from the corresponding encounter/response rate.

### 8.2.1 Some caveats

The above claims must be accompanied by a series of qualifications which serve to refocus, rather than to restrict their significance. The most obvious is that none attach specifically to the IFD itself, as is to be expected in view of the discussion above; the IFD returns to the limelight in the following chapter. Perhaps more important, however, is that the matching law, as understood throughout the psychological literature, describes a relationship between response frequency and reward *obtained* ( $R_A, R_B$ ) whereas the present model focuses on the relationship between response and reward *available* ( $F_A, F_B$ ), for reasons that will be explained in a moment.<sup>2</sup> A third caveat is that none of the claims are meant to be exclusive; there may well be other rationales for matching which complement the present argument, and there are certainly other mechanisms capable of supporting matching (see, for example, the ‘switching rule’ mechanisms in the following chapter). Furthermore, the present model, because it does not involve agent plasticity, cannot address the problem of maintaining adaptive fit by the tracking of changing reinforcement contingencies, a problem which arguably constitutes a significant part of the task of any mechanism of matching.

Nevertheless the claims still stand, and indeed these caveats help to mark out the novel territory into which they extend. The idea that matching can arise from the most fundamental of situated foraging interactions is very different from the orthodoxy of recognising matching only in situations in which repeated choices are presented between discrete and mutually exclusive alternatives, be they keys in a Skinner box or patches in a meadow. The present model also underlines the potential value of a distinction between an account of the mechanisms underlying the expression of a behaviour, and of those responsible for maintaining its adaptive fit. Although this distinction may not often be enforceable in real systems, it can be made in the abstract and, as will be argued, in this case it extends the understanding of the processes involved.

It is also this novel territory that sanctions an interpretation of matching in terms of available rather than obtained resources (in terms of  $F_i$  rather than  $R_i$ , for choice options indexed by  $i$ ). In the present model a reliable contingency is always arranged between response and reward; that is, more encounters with resource items mean more reward (although exactly *how much* reward may depend on the resource type). This may seem sensible enough, but in fact most paradigms in experimental psychology go to great lengths to avoid such a contingency (see, for example, the various reinforcement schedules described in Herrnstein, 1970), since otherwise matching to obtained resources could only ever occur if either (a) item types were equally valuable, or (b) response to one or other resource types was negligible (see appendix B for the maths behind this claim).<sup>3</sup> Although many interesting results proceed from experiments within such paradigms,

<sup>2</sup>The reader should be warned that this distinction, despite its indubitable importance, has not always been respected in the matching literature.

<sup>3</sup>Somewhat paradoxically, in the present model it is therefore only ‘zero-one’ behaviour which can be described, if trivially, in terms of matching to obtained resources.

some of which will be discussed in the following chapter, the present model is concerned with the simplest formulation of foraging behaviour in a single patch in which a strong link between response and reward is essential, and must therefore consider matching to available resources. In what follows, as in the claims set out above, the term ‘matching’ should therefore be interpreted carefully, *not* as matching to obtained resources, but as matching in the more general sense of a distribution of response which mirrors in some way the relative values of the response options, as opposed to the extremes of either indifference or exclusive choice.

### 8.2.2 Extending the ‘interference’ model

The model developed in this chapter is a direct extension of the IOS interference model described in the previous chapter. The 200.0 by 200.0 patch now contains 4  $X$  items and 4  $Y$  items, with each item type associated with a probability that encounter will lead to full replenishment of the agent’s internal battery. These probabilities ( $F_X, F_Y$ ) therefore represent the resources available from each item type; this is what is meant by item ‘value’. In the present model,  $X$  items *always* replenish the battery ( $F_X = 1.0$ ), and  $F_Y$  is experimentally manipulated. As before, resource items (radius 8.0) and agents (radius 2.5) are always randomly initialised within the (single) patch, and consumed resources are immediately replaced. Agents ‘die’ if their battery expires, and encounters with conspecifics (if any) have no effect.

Each agent now possesses 5 sensors, 4 of which are sensitive to resource items (in 2 left/right pairs) and one of which reflects the battery level. The internal architecture of each agent again comprises a reactive feedforward network, the only difference from the previous chapter being the addition of 2 extra input units corresponding to the extra sensor pair. As before, input units scale sensor values to the range [0.0,1.0], with all interconnecting weights in the range [-1.0,1.0]. Each unit applies a sigmoid function to the sum of its inputs (plus a threshold value), with each output scaled to the range [0.0,1.0]. Motor outputs are scaled to the range [-10.0,10.0] to set wheel speeds, and again as before, a value of 10.0 on both wheels translates to a speed of 2.8 distance units per time-step. A genotype of length 46 is required in this model, 42 loci to specify the weights and thresholds of the network (5 input units, 5 hidden units, and 2 output units), and 4 for something less conventional; to specify how well the agent is able to discriminate between the food item types.

In the general case each agent has  $i$  sensor pairs, each of which is associated with a set  $\{d_X, d_Y, \dots, d_g\}$ , with  $g$  indexing resource types. Each  $d$  lies in the range [0.0,1.0], and specifies the probability with which the associated sensor will perceive an item of type  $g$ ; all  $d$  values are genetically specified. In the present case, with 2 sensor pairs and 2  $d$  values per pair, 4 additional loci are required. This scheme functions in the following way. Every time an item is initialised it is tagged with the identity of each sensor pair that can perceive it. For example, if an agent has  $[d_X = 1.0, d_Y = 0.0]$  for its first sensor pair ( $s_1$ ) and  $[d_X = 0.5, d_Y = 1.0]$  for its second ( $s_2$ ), then an  $X$  item initialised within the range of the agent will be tagged as perceivable by  $s_1$ , and also by  $s_2$  if (and only if)  $\mathcal{R} < 0.5$  ( $\mathcal{R}$  being a random number in the range [0.0,1.0]). During each sensorimotor cycle, each sensor pair of each agent responds to its nearest perceivable item (if any), ranging, as before, from 100.0 (at an item) to 0.0 ( $\geq 200$  distance units away). The rationale for introducing this admittedly unusual process (further illustrated in figure 8.3), and its




Figure 8.1: An agent is depicted in relation to an  $X$  item and a  $Y$  item, both within range. Consider a sensor pair  $s_1$  with  $[d_X = 1.0, d_Y = 0.0]$ ;  $s_1$  will respond to  $X$ . Now consider  $[d_Y = 0.3]$ ; there is now a 30% chance that  $s_1$  will respond to  $Y$  instead of  $X$ . This probability will have been evaluated when  $Y$  was initialised so that  $s_1$  (and all sensor pairs) respond to items in a consistent manner. If the positions of  $X$  and  $Y$  were reversed then  $s_1$  would respond to  $X$  no matter what the status of  $Y$  with respect to  $s_1$  was.

consequences, are discussed in detail in section 8.3.2.

The experiments themselves proceeded as follows. A distributed GA (appendix A) was used to evolve populations of genotypes in each of 8 conditions; 4 involving a single isolated agent (the  $S$  - single agent - condition set) and 4 involving 3 ‘clonal’ agents derived from the same genotype (the  $M$  - multiple agent - condition set). The fitness function used - in all conditions - was the same as in the previous chapter:

$$\mathcal{F} = \sum_{t=1}^{800} \frac{\mathcal{B}}{200},$$

where  $t$  indexes time-steps and  $\mathcal{B}$  represents the battery level (at time  $t$ ), with each evaluation lasting, as before, for a maximum of 800 time-steps. This function rewards agents that live long and forage efficiently, without in any way specifying the structure of the behaviour through which this is to be achieved. Each genotype - also as before - was evaluated 4 times each generation.

Evolved agents from the various conditions (described below) were used to test the hypotheses of section 8.2: that agents evolved in isolation would obey the zero-one rule and that those evolved with conspecifics would behave in accordance with the matching law.

## 8.3 Results

### 8.3.1 Simple matching in simple agents

Both the  $S$  and  $M$  condition sets involved evolving and testing agents in environments distinguished by the value of  $F_Y$ . Four values of  $F_Y$  were employed in each condition set; 1.0, 0.66, 0.33, and 0.0, with  $F_X = 1.0$  in all conditions. Two conditions were expected to produce ‘zero-one’ behaviour in both sets:  $F_Y = 1.0$  and  $F_Y = 0.0$ . If  $F_Y = 1.0$ ,  $X$  and  $Y$  are functionally equivalent, and so both  $S$ - and  $M$ -agents should respond to each equally. If  $F_Y = 0.0$ ,  $Y$  is always worth nothing, and so both  $S$ - and  $M$ -agents should respond exclusively to  $X$  items. In the remaining conditions,  $F_Y = 0.66$  and  $F_Y = 0.33$ , it can be predicted that the single  $S$ -agent will *either* continue to respond to  $X$  and  $Y$  equally, *or* will switch to exclusively responding to  $X$  (thus continuing to follow the zero-one rule), and, by contrast, that  $M$ -agents will respond to  $Y$  in some proportion to the difference in value between  $X$  and  $Y$ .

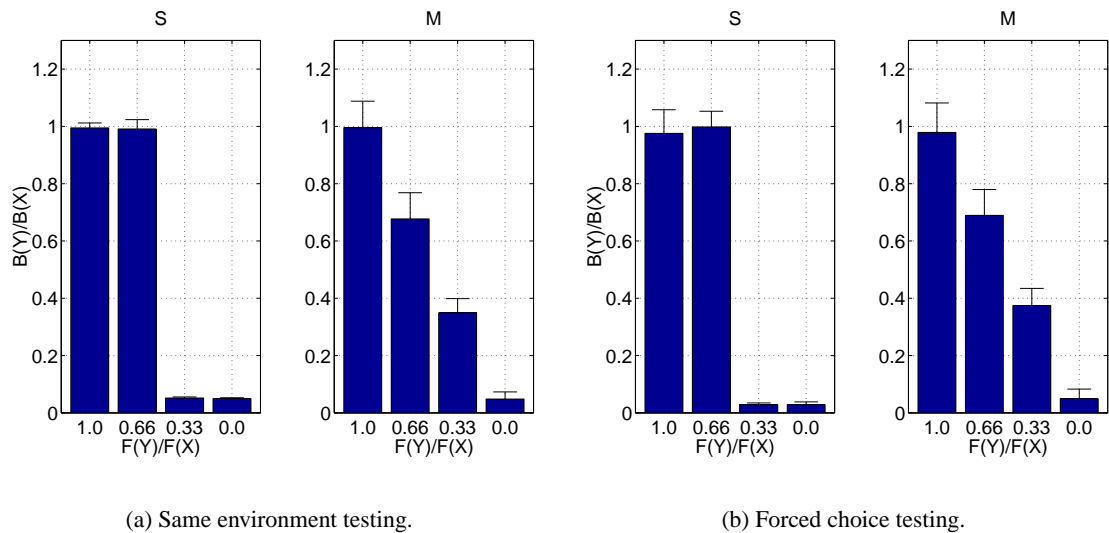


Figure 8.2: These graphs show the average (and standard deviation) rate of response to  $Y$  ( $B_Y$ ) as a fraction of rate of response to  $X$  ( $B_X$ ) over 12 evolutionary runs in each condition. From each run, a single value was obtained by testing the fittest agent 1000 times in either (a) the same environment as evolution, or (b) in a forced-choice environment. Matching behaviour is observed in  $M$ -agents, and zero-one behaviour in  $S$ -agents, under both testing regimes.

Genotypes of high fitness, in both  $S$  and  $M$  condition sets, reliably evolved after about 400 generations in each condition, but in each case the population was left until 1000 generations had been completed. The fittest genotypes from each condition were then evaluated, in the same conditions as experienced during evolution in each case, with the average number of responses to  $X$  and  $Y$  items (over 1000 evaluations) being recorded. The entire set of evolutions (and analyses) was repeated 12 times to obtain overall averages.

It can clearly be seen from figure 8.2(a) that, when tested in the same environment as evolution,  $S$ -agents do indeed follow the zero-one rule, and  $M$ -agents do match to available resources. It may be surmised from this that, given the constraints of the model, zero-one behaviour constitutes near-optimal foraging for an isolated agent, and matching to available resources constitutes near-optimal foraging in a shared environment.

$S$ -agents and  $M$ -agents were also assessed in what may be called a ‘forced-choice discrimination’ task. In this analysis, the fittest genotype from each condition was decoded into a *single* agent, which was then assessed *in isolation* by being placed equidistant from a single  $X$  item and a single  $Y$  (no other items were present). Each trial was stopped as soon as one or other of the items had been visited, and again each agent was tested 1000 times. It is important to emphasise that these tests always involved a *single* agent, even if evolution had occurred in a multi-agent environment. Figure 8.2(b) illustrates that the pattern of results is unchanged. In these tests, however, the optimal policy for all conditions for which  $F_X > F_Y$  is exclusive choice of  $X$ , therefore the matching behaviour of isolated agents from the  $M$  condition is suboptimal.<sup>4</sup>

<sup>4</sup>The zealous reader may have noticed that the behaviour of isolated  $S$ -agents in the condition  $F_Y = 0.66$  is, by this argument, also suboptimal. Recall, however, that the zero-one rule - followed by the  $S$ -agents in all conditions - specifies only that a given resource type should be always consumed or always ignored. Which of these is optimal in

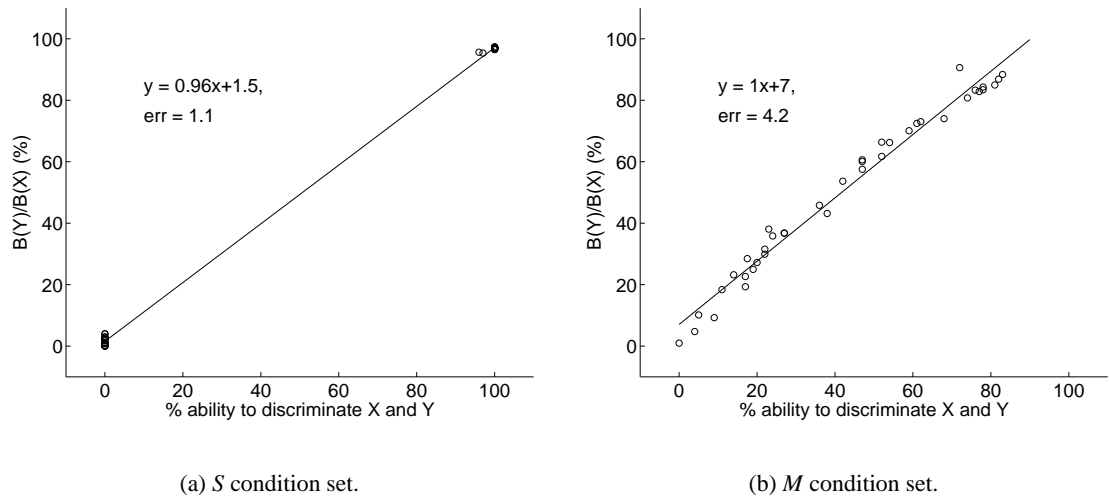


Figure 8.3: Evolved sensor properties. The correlation between the ability of the sensors to discriminate between  $X$  and  $Y$  and the response ratio is shown; each data point represents the fittest agent from a single run of either the  $S$  condition or the  $M$  condition. Best fit lines are shown, with the error value specifying the range around any point on the line that contains at least 50% of the data points.

### 8.3.2 Adaptive sensors

A question may immediately be asked of the non-deterministic behaviour of the evolved agents. How does this come about? The answer has to do with the only source of stochasticity in the model, the sensors. Figure 8.3 illustrates that there is a strong correlation, in both  $S$  and  $M$  condition sets, between the ability of agents to discriminate between item types  $X$  and  $Y$  and the difference in response to these types (the correlation value is less meaningful for the  $S$  condition set because of the lack of intermediate values). It is this adaptive potential of the sensors which enables the agents, in both condition sets, to adapt their response distributions to the reward structure of the environment.

This explains the perhaps mysterious methodological step of introducing this kind of sensor. Without such a source of stochasticity, there would be no way for the agents to generate the (matching) behaviour that the model was designed to investigate. To the extent that we are concerned with functional properties of matching, the particular way in which stochasticity is introduced does not much matter. It only matters that the GA can search a sufficient volume of behaviour space to allow functional hypotheses to be adequately assessed, in the present case both matching and zero-one behaviour must be feasible. However, the source of stochasticity *is* important for any discussion of the *mechanisms* underlying matching. Here, the adaptive sensors form an essential component of the mechanism, and their significance must be acknowledged.

Of course, other strategies for introducing stochasticity can be conceived, some of which may

---

any particular instance depends on the overall resource density (Charnov & Orians, 1973), an aspect of the environment that is not conserved in the forced choice analysis. In other words, the behaviour of  $S$ -agents in figure 8.2(a) represents a near-optimal expression of the zero-one rule, given the resource density levels in the ‘natural’ environment of the agents. This particular expression of the zero-one rule is no longer optimal in the case where  $F_Y = 0.66$  in figure 8.2(b) because resource density levels have changed. Importantly, the (suboptimal) matching of isolated  $M$ -agents cannot be explained in this way because the matching law is *not* framed in terms of resource density dependence.

perhaps be more intuitively satisfying than the present scheme (one idea might involve an additional hidden unit with randomly varying output), however adaptive sensors at least have the advantages of emphasising (i) the importance of perceptual mechanisms in choice behaviour, and (ii) the fact that less than perfect sensors can sometimes be a good thing (for more on this issue see Smithers, 1994). The investigation of alternative sources of stochasticity is therefore left for future work.

### 8.3.3 The role of interference

The results of section 8.3.1, which constitute the primary empirical contributions of this chapter, certainly support the hypothesis that individual matching behaviour (and its potential suboptimality) can arise from foraging behaviour adapted to a group context. But is it the case that these results can be attributed to high levels of interference attaching to high quality resources? To address this question directly, a series of *S*-type models (in which a single agent is involved in both evolution and testing) is considered, in which the environment is modified in various ways to mimic the influence of interference in *M*-type models. A secondary hypothesis is evaluated using these models; that matching can be attributed to the patterns of *resource instability* entailed by interference. Resource instability, in the present context, is taken to describe the unpredictable (from the perspective of the agents) fluctuations in the patterns of resource distribution caused by the activity of conspecifics. It is therefore taken to be a direct correlate of interference, an assumption not directly put to the test here, but one which is strongly supported both by informal observation of the *M* condition set of the present model, and by the interference model of chapter 7 (recall that in condition *B* of this model, agents were adapted to exactly this kind of instability).

In what follows, various levels of resource instability are artificially induced in *S*-type models, the hypothesis being that relatively high levels of resource instability (such as is observed in unmodified *M*-type models) may lead to individual matching, but that relatively low levels (such as is observed in unmodified *S*-type models) may not. The level of instability in each case is determined by the parameters  $I_X$  and  $I_Y$ , such that, for example,  $I_X = 0.02$  means that for each *X* item in the patch, there is a 2% chance, on each time-step, that the item will disappear and reappear at another random location in the patch, exactly as if it had been consumed. The total instability introduced to the patch,  $I_X + I_Y$ , is designated by the parameter  $I_{tot}$ .

Five levels of overall resource instability  $I_{tot}$  were investigated in a series of modified *S*-type models, ranging from  $I_{tot} = 0.02$  to  $I_{tot} = 0.06$ , and for each value of  $I_{tot}$  three experimental conditions were distinguished, for reasons described in the next section. In all cases,  $F_X = 1.0$ ,  $F_Y = 0.33$ , and all other aspects of each model are as before. (The choice of  $F_Y = 0.33$  over  $F_Y = 0.66$  was arbitrary; for present purposes it is only necessary to explore a single value indicative of matching in the *M*-type models.)

In condition *A*,  $I_{tot}$  is assigned exclusively to the (high quality) resource type *X*, thus  $I_X = I_{tot}$ ,  $I_Y = 0.0$ . This, a benchmark condition, is the simplest way in which resource instability can be introduced. In condition *B*, by contrast,  $I_{tot}$  is *initially* assigned entirely to *X*, but the values of  $I_X$  and  $I_Y$  may vary over generations, and different individuals (within a single generation) may experience different levels of resource instability. The instability parameters are rewritten in this condition as  $I_{X_j}(g)$  and  $I_{Y_j}(g)$ , representing the instabilities of *X* and *Y* items during generation



$g$  for individual  $j$ . In condition  $B$ , then,  $I_{Xj}(0) = I_{tot}$  and  $I_{Yj}(0) = 0.0$ , for all  $j$ , but for every subsequent generation, each individual  $j$  inhabits an environment in which

$$\frac{I_{Xj}(g)}{I_{Yj}(g)} = \frac{B_{Xj}(g-1)}{B_{Yj}(g-1)},$$

where  $B_{Xj}(g-1)$  and  $B_{Yj}(g-1)$  represent the rates of encounter with  $X$  and  $Y$  items of (either one of) individual  $j$ 's immediate ancestors. Finally, in condition  $C$ , *all* individuals of each generation experience identical resource instability profiles determined by:

$$\frac{I_X(g)}{I_Y(g)} = \frac{\overline{B_X}(g-1)}{\overline{B_Y}(g-1)},$$

where  $\overline{B_X}(g-1)$  and  $\overline{B_Y}(g-1)$  are the population average encounter rates of the previous generation, and in which the  $j$  suffix may be dropped from all terms. Twelve evolutionary runs were performed in each of these conditions for each value of  $I_{tot}$ , and the results reported below concern only testing in a forced-choice discrimination environment (with no resource instability); results from testing in the same environment as evolution are not reported.

### 8.3.4 The importance of history

The crucial distinctions between conditions  $A$ ,  $B$ , and  $C$  lie in the distinct ways in which they engage with behavioural historical constraints. In condition  $A$ , a control condition, resource instability is imposed from the outset and remains a constant feature of the environment. There is nothing historical about this; the structure of the optimal behaviour (the outcome) remains conceptually separate from the optimisation process, and figure 8.4(a-e) indicates that this condition always leads to zero-one behaviour. Agents switch from exclusive preference for  $X$  when  $I_{tot}$  is low, to indifference when  $I_{tot}$  is high. It is not enough to consider matching as a 'solution' to an environmental 'problem' of resource instability.

However, *pace* Lewontin's (1983) critique of the concept of ecological niche, resource instability should *not* be considered to pre-exist a foraging strategy. It is clear that in  $M$ -type models any resource instability will be intimately tied to the foraging behaviour of the agents at any time. Also, as was discussed in the case of the interference model in the previous chapter, such patterns of agent-environment interaction invite interpretation as behavioural historical constraints to the extent that they interact with the process of optimisation in influencing their own structure and persistence. Condition  $B$ , the most important of this series, evaluates the significance of this interaction in the generation of matching by establishing a direct connection between the response profiles of single agents and the patterns of resource instability imposed on them, albeit with an (unavoidable) generational lag. Figure 8.4(a-e) clearly illustrate that this condition *does* lead to matching of a sort, and although it is less well defined than that of the  $M$ -agents of section 8.3.1, the departure from zero-one behaviour is indisputable. Also, a correlation is evident between  $I_{tot}$  and the centre of the distribution of results, that is, the higher the value of  $I_{tot}$ , the lower the mean preference for  $X$ . Furthermore, as with both the interference model and the  $S$ -type and  $M$ -type models above, the mechanics of the GA still cannot influence the fitness values awarded to the constituent genotypes, thus the interpretation of the GA may again be confined to optimisation. (This state of affairs is not altered by the inheritance of resource instabilities since such inheritance

is restricted entirely to individual lineages.) The evidence from this condition therefore suggests that patterns of resource instability *can* engender matching, but *only when construed in terms of behavioural historical constraints*.

Condition *C*, another control condition, preserves the notion that resource instability should not pre-exist foraging behaviour, but differs significantly from condition *B* in its employment of population averages. The direct link between the response profile of an agent and the resource instability encountered by that agent (or its immediate descendant) is now broken, abolishing the potential for historicity described above. Figure 8.4(a-e) illustrates that this condition does *not* lead to matching, with agents always retaining exclusive preference for *X*. Note also that, in this condition, the mechanics of the GA *can* influence fitness values; the GA is still optimising, but the way in which this optimisation occurs must now bear explanatory weight because attributes of the GA that concern population distribution (convergence, size, and so on) can influence population average response rates and therefore affect the resource instability and hence the fitness values of individual agents.

### 8.3.5 Historical resource abundances

Along slightly different lines, and again as something of a control condition, a final variation was analysed: a modified *S*-type model in which conditions *A*, *B*, and *C* were attached to resource abundances, rather than to resource instabilities ( $I_{tot} = 0.0$ ,  $F_X = 1.0$ , and  $F_Y = 0.33$  in all cases). The idea was to see whether individual matching could arise simply from a relative scarcity of high quality resources. Condition *A* comprised an *S*-type model with 2 *X* and 4 *Y* resource items. In condition *B*, the relative abundances of *X* and *Y* were determined, for each individual, by the response profile of (either one of) its parents. Condition *C*, as before, employed population averages. In conditions *B* and *C* abundances were limited to range from 0 to 4, providing a total of 9 possible resource distributions (an initial default of 4 *X* and 4 *Y* items, and 8 further distributions for diminishing quantities of either *X* or *Y*); response profiles were linearly mapped onto this range. Figure 8.4(f) illustrates that none of these conditions entailed departure from zero-one behaviour when agents were tested under forced-choice discrimination conditions.

### 8.3.6 Sensitivity

Now that the generation of matching has been explored in some detail, this section returns to assess the robustness of the results from the original *S*- and *M*-type models.

#### *Resource density*

The first manipulation involved small variations in overall resource density (this time maintaining equal quantities of *X* and *Y* items). All 8 conditions were re-evolved with either 3 or 5 items of each type in the environment (12 complete re-evolutions in each case), with analysis performed as before. Figures 8.5 and 8.6 illustrate quite clearly that these variations in overall resource density do not affect the overall pattern of results at all (compare with figure 8.2). Evolution in the presence of conspecifics still leads to matching, and evolution in isolation still leads to zero-one behaviour.

#### *Agent density*

Here the distinction between the *S* and *M* condition sets is extended with a series of similar sets in which the number of coexisting agents varies from 3 to 8. The intuition is that as agent density

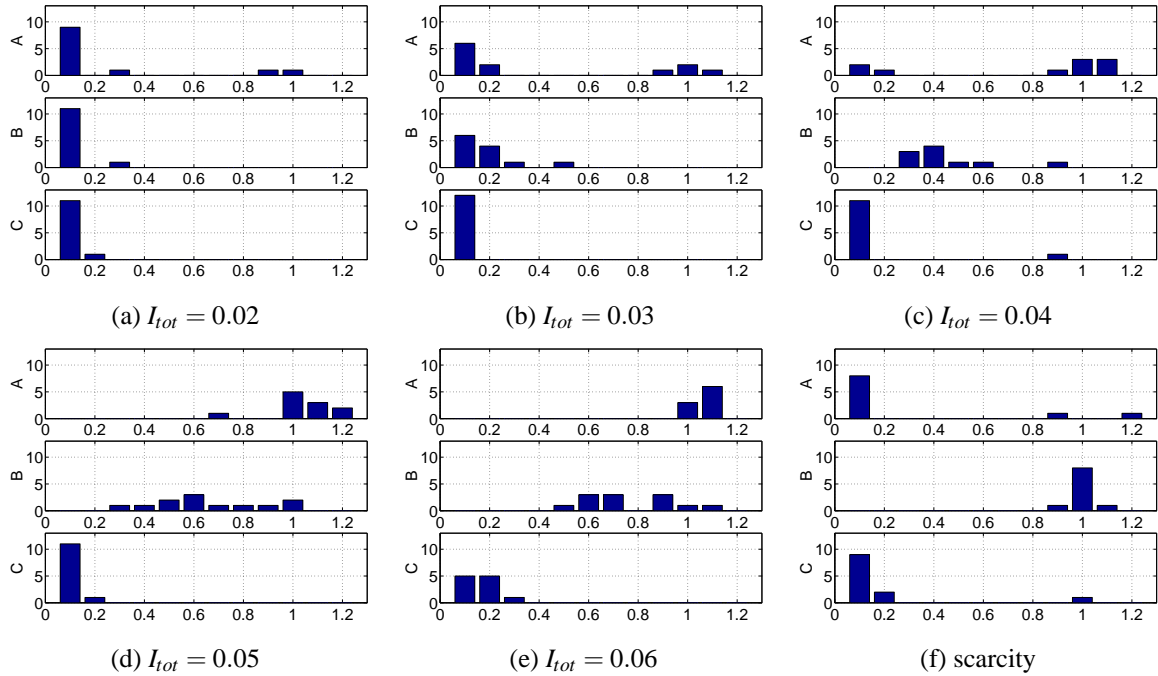


Figure 8.4: Environmental instability in modified *S*-type models. In each plot the abscissa represents  $B_Y/B_X$  (the ratio of responses to *Y* items over responses to *X* items), and each column indicates the number of GA runs (out of a total of 12) that led to each value of  $B_Y/B_X$  (in ranges of breadth 0.1) when tested in forced-choice discrimination conditions (with  $I_{tot} = 0.0$ ). Each plot contains 3 sub-plots representing conditions *A*, *B*, and *C*. Zero-one behaviour is indicated by the bars clustering around the extremes (0.0, 1.0) of the abscissa, and matching behaviour is indicated by a distribution of results across intermediate values of  $B_Y/B_X$ . Plots (a) to (e) illustrate a range of values of resource instability ( $I_{tot}$ ), and clearly only condition *B* ever results in matching. Plot (f) illustrates a similar experiment in which resource abundance (rather than instability) is modified in three analogous conditions. Matching never results. The various conditions are explained in the text.

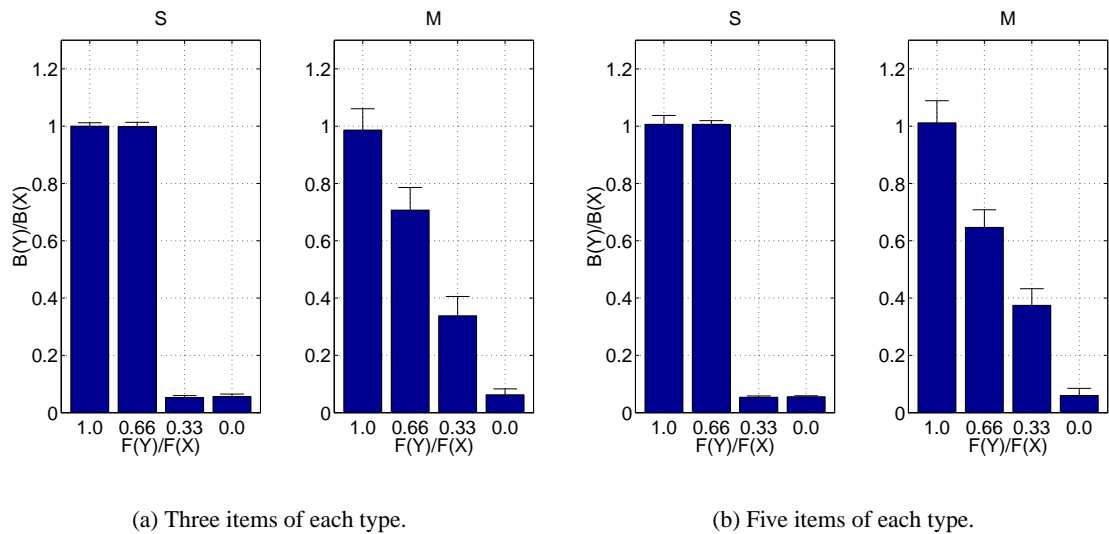


Figure 8.5: The distinction between matching and zero-one behaviour remains clear with altered food item densities. These graphs concern testing in the same environment as evolution, and are to be interpreted as in figure 8.2(a); 12 evolutionary runs were performed in each condition.




Figure 8.6: The distinction between matching and zero-one behaviour remains clear with altered food item densities. These graphs concern testing in a forced-choice discrimination environment, and are to be interpreted as in figure 8.2(b); 12 evolutionary runs were performed in each condition.

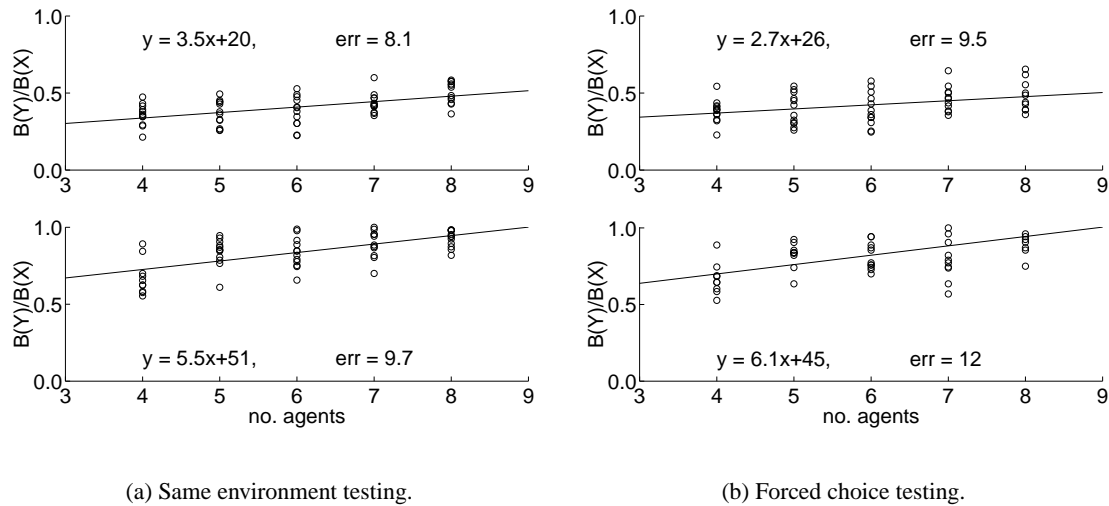


Figure 8.7: Relative response rates (of a randomly selected agent from each group) for  $F_Y = 0.33$  (top graphs) and  $F_Y = 0.66$  (bottom graphs). As the number of agents increases, foraging behaviour shifts towards indifference between  $X$  and  $Y$  (most obviously for conditions in which  $F_Y = 0.66$ ). Each data point represents a single evolutionary run (12 for each agent density) and the average of 500 evaluations of the fittest genotype. Lines of best fit through these data points are drawn, and each error measure specifies the range around any point which contains at least 50% of the data points. Results are shown both for testing in the same environment as evolution (a) and in forced-choice environments (b).

(and hence competition for *both* resource types) increases, the optimal response distribution (in all conditions except those in which  $F_Y = 0.0$ ) should shift towards indifference between  $X$  and  $Y$ . (It is worth remembering that indifference is *not* the same thing as the zero-one rule, which describes density dependent switching between indifference and exclusive choice.)

Ten condition sets were investigated, each a variation of the  $M$  condition set, 5 of which consider  $n$  conspecifics foraging in an environment with  $F_Y = 0.33$ , and 5 with  $F_Y = 0.66$ , for all  $n \in \{4, 5, 6, 7, 8\}$ . Twelve entire sessions of evolution and testing were carried out in each condition as before, with testing carried out both in the same environment as evolution and also in forced-choice contexts. Figure 8.7 suggests that as agent density increases,  $X$  and  $Y$  items are indeed collected with increasing lack of discrimination, underlining the point that matching is *not* an inevitable outcome of multi-agent foraging in a heterogeneous environment (admittedly the trend is less evident for conditions in which  $F_Y = 0.33$  than it is for conditions in which  $F_Y = 0.66$ ).

#### Complex environments

In this final extension, the  $M$  condition set is re-implemented with varying densities of a third item type  $T$  (for ‘trap’) in the patch; any encounter with a trap immediately terminates an agent, just as in the original behaviour coordination model of chapter 6. The hypothesis here is that the increased complexity of this environment will again weaken the adherence of agents to the matching law since a significant component of agent behaviour should now have to concern avoiding traps, and this should constrain the freedom with which the collection of  $X$  and  $Y$  items can be balanced to maximise overall intake.




Figure 8.8: Relative response rates (of a randomly selected agent from a group of 3) for  $F_Y = 0.33$  (top graph) and  $F_Y = 0.66$  (bottom graph). As the number of traps increases, foraging behaviour shifts towards indifference between  $X$  and  $Y$  (at least for conditions in which  $F_Y = 0.33$ ). Each data point represents a single evolutionary run (12 for each agent density) and the average of 500 evaluations of the fittest genotype. Lines of best fit through these data points are drawn, and each error measure specifies the range around any point which contains at least 50% of the data points. Results are shown both for testing in the same environment as evolution (a) and in forced-choice environments (b).

Consider the 3 agent  $M$  condition set of section 8.2. Let  $n_T$  represent the number of (randomly distributed) traps that share the patch with the 4  $X$  and 4  $Y$  items. A total of 18 conditions were analysed, 9 with  $F_Y = 0.33$  and 9 with  $F_Y = 0.66$ , with each set of 9 representing  $n_T = \{1, 2, \dots, 9\}$ . In these conditions, agents are endowed with a third pair of sensors (and a corresponding pair of additional input units) which, unlike the original two sensor pairs, are dedicated to trap detection; this third sensor pair will *always* respond to the nearest trap, and never to resource items. Thus, the genotype needs only to be extended by 6 alleles in order to code for the extra weights between the input and hidden layers (no  $d$  values need be encoded). Twelve sessions of evolution and testing were carried out in each condition as before, and figure 8.8 illustrates that as the number of traps in the environment increases, individuals do indeed become increasingly indifferent between  $X$  and  $Y$ . (In contrast to figure 8.7, however, this trend is only clearly evident for conditions in which  $F_Y = 0.33$ ; for the  $F_Y = 0.66$  conditions agents are close to being indifferent between  $X$  and  $Y$  items for all values of  $n_T$  above 3.)

### 8.3.7 Summary

The results of this chapter may be summarised as follows. Agents evolved in isolation in heterogeneous environments exhibit either exclusive choice or indifference, this being the zero-one rule. Evolution in the presence of conspecifics (of a certain density) leads individuals, both when in groups and when in isolation, to deviate from the zero-one rule and match (strictly or otherwise) to available resources. This can be traced to patterns of resource instability (interference) when

construed as historical processes. These results are robust to variations in resource abundance, but as the density of conspecifics increases, agents tend to become increasingly indifferent and matching becomes less evident, an effect that can also be elicited by the introduction of additional environmental survival contingencies.

## 8.4 Discussion

### 8.4.1 Individual suboptimality

Consider first the theme of individual suboptimality. This is most clearly illustrated in figure 8.2(b) in which isolated agents face a repeated forced-choice discrimination test and in which those agents evolved in the presence of conspecifics (*M*-agents) behave suboptimally (similar results are evident in figures 8.6, 8.7b, and 8.8b). The optimal policy in this situation is to always choose the high value item type, but *M*-agents instead distribute their responses across item types in proportion to their relative values. Of course, this pattern of behaviour is - indeed it must be - the near-optimal solution to foraging in the multi-agent environment in which the *M*-agents evolved. *S*-agents, on the other hand, display zero-one behaviour both in the environment of their evolution and in the forced-choice test. Together, these observations clearly describe a case in which suboptimal individual behaviour can be understood as a consequence of adaptation to group environments. But how far do they contribute to an understanding of the matching law itself?

### 8.4.2 The matching law

Earlier in this chapter (section 8.2.1) it was carefully established that matching in this model was to be assessed in terms of available resources, rather than (as is orthodox) obtained resources. Therefore, it cannot be claimed that the internal mechanisms of *M*-agents are ‘mechanisms of matching’ in the specific sense of delivering matching to obtained resources in environments characterised by distinct patches or choice options. Nor can it be asserted that the potential suboptimality of the (orthodox) matching law has been explained by the individual suboptimality observed in figure 8.2(b). Instead, the present model shows how individual agents, with very simple internal mechanisms (admittedly incorporating adaptive sensors), can match to available resources as a consequence of foraging near-optimally in a shared, heterogeneous resource, single patch environment; an environment in which there is also a (necessarily) strong connection between resources consumed and reward received. This matching behaviour incurs its own potential for individual suboptimality, and is sensitive to both the level of competition from conspecifics and to the severity of additional survival contingencies (traps). Moreover, this behaviour, and the concomitant patterns of resource instability, are best understood as historical in the sense of influencing the conditions of their own realisation. In particular, the (unorthodox) matching identified here *cannot* be seen as a ‘solution’ to the ‘problem’ of interference. This kind of insight would have been difficult to explore, if not impossible, in the absence of the substrate of sensorimotor interaction provided by the IOS methodology. This point will become particularly clear in the following chapter, in which an orthodox interpretation of matching naturally involves a model in which interference is *a priori* part of the environmental problem to be adapted to.

It is possible to locate the matching of the present model at the simplest extreme of a proposed continuum of matching phenomena, an extreme which has no need of distinct patches, nor of any

artificial separation of response from reward, nor of adaptive internal state. The other extreme represents the matching of psychology textbooks, replete with complex and ecologically implausible reinforcement schedules (that often explicitly deny response-reward contingencies) and hypotheses of powerful and ultimately rational (for the isolated individual) ‘mechanisms of matching’; an extreme which nevertheless engages with, and continues to generate, a wealth of empirical data. In the middle there reside, amongst other things, models in which the operation of simple ‘switching rules’ build directly on the analogy between matching and the IFD (and which may tend towards the encapsulation of interference as part of the environment; the following chapter explores a region of this middle ground in some detail). I believe that the identification of this continuum casts doubt on the plausibility of a single dedicated ‘mechanism of matching’ of any sort. It may well be more profitable to consider a multiplicity of mechanisms, of differing adaptive stature, all potentially contributing to a determination of a response distribution, but in principle separable and present to varying degrees in different varieties of agent.

The claims with which this exploration of matching commenced can now be revisited (see section 8.2). Matching (of one kind at least) *can* be understood in terms of adaptation to group choice, although it is not an inevitable consequence of such situations. The same sensorimotor interactions that give rise to interference can also give rise to both matching and zero-one behaviour, a claim dependent upon a model of interference consistent with optimal foraging and permissive of certain historical processes, and qualified in the present case by the importance of the adaptive sensors in generating the matching behaviour. Moreover, at the same time as establishing a mechanistic homology between the structures capable of underlying both zero-one and matching behaviour, their distinct functional roles have been highlighted; the former in isolated foraging, the latter in competitive foraging. Finally, it is only through the separation of the machinery of adaptation or optimisation (the GA) from the mechanisms of behavioural expression (the neural network and the adaptive sensors) that these claims have been substantiated and a continuum of matching described.

### 8.4.3 Exploration versus exploitation

Behaviour in environments presenting multiple choice options is often associated with the trade-off between ‘exploitation’ and ‘exploration’, a central idea in the reinforcement learning literature which recognises that agents must balance exploitation - selecting the best known action - against exploration - sampling alternative actions to find better options (see, for example, Kaelbling, Littman, & Moore, 1996). Although the present model involves a multiple choice option environment, it is important to stress that the agents themselves do *not* explicitly strike any such balance. Responses of *M*-agents to apparently poor resource items (for example figure 8.2a) are not to be interpreted as exploratory moves, rather as part of a near-optimal foraging strategy given the presence of conspecifics. And the suboptimal persistence of this behaviour in forced-choice discrimination tests (figure 8.2b) again reflects *not* exploration, but the inability of these agents to modify their response patterns on the basis of experience.

These issues will be of greater significance in the following chapter, in which reinforcement learning is instrumental in the adaptation of agents to their environment.



#### 8.4.4 A broader view

It is apparent - with regard to the general themes of this dissertation - that many of the issues raised by the present chapter are similar to those discussed in the context of interference. The distinction between behaviour and mechanism is most clearly reflected in the challenge to the preconception that matching behaviour must be subserved by a dedicated 'mechanism of matching'. As with the interference model, the theme of behaviour coordination has emerged in the guise of suboptimal behaviour. In the present chapter, however, the focus on matching provides a close connection with action selection in SAB/AL, and in doing so serves as a clear example of how IOS models may contribute to a tightening of relations between SAB/AL, psychology, and biology; IOS models of this kind provide both theoretical structure and a unifying discourse for the modelling of phenomena (such as the matching law) that straddle these disciplines.

First and foremost then, this chapter represents another example of the application of the IOS methodology, and, through the analysis of the relations between resource instability and matching, the present model has exemplified the utility of this methodology in elucidating the properties of behavioural historical constraints. As with the interference model, these historical constraints can be related to Lewontin's critique of the 'ecological niche'; behavioural historical constraints - like ecological 'niches' - cannot be presumed to pre-exist the activity of the 'constrained' agent itself.

Godfrey-Smith's environmental complexity thesis also relates to the present model more-or-less as it did to the interference model. There is the part played by causal construction (alterations in the external environment) in the generation of behavioural historical constraints. This, of course, remains closely connected to Lewontin's critique. Constitutive construction - the translation of external environment into *Umwelt* - associates closely with the use of adaptive sensors. There is also the abstract explanation of the relative 'complexity' of matching over zero-one behaviour, insofar as an environment replete with conspecifics may be considered to be more 'complex' than one populated only by an isolated individual.

### 8.5 Summary

This chapter has continued the development of the IOS model of interference into the realm of the matching law. This context has provided a clear example of a behaviour - a simple yet unorthodox kind of matching - optimal in a group context, yet suboptimal for an isolated individual. This example of apparent irrationality is arguably of greater clarity and potential biological and psychological significance than that provided in the context of interference, and it has also helped to extend the understanding of choice behaviour in SAB/AL, which until now has been largely confined within the framework of action selection. Furthermore, this chapter has demonstrated that this simple version of matching (and indeed also zero-one behaviour) can arise from the same sensorimotor activity responsible for interference, without distinct 'patches', internal adaptive state, or artificial separation of response from reward (but *with* the addition of adaptive sensor capabilities). However it is *not* the case that this matching can be considered to be a 'solution' to the 'problem' of interference. (Recall that part of the explanatory power of the interference model itself followed from its potential for historicity; it is clear the same may also be said in the present case.)

These observations when drawn together support the identification of a continuum of matching

behaviour, not all of which is well described by the original matching law of Herrnstein. This is both a contribution to the understanding of behavioural choice in general, and an illustration of the limitations of the present model, which remains unable to engage directly with the matching literature. The next chapter faces up to these limitations, departing from the IOS methodology that has characterised much of the dissertation, but in doing so locating itself squarely in the ‘middle ground’ between the present model and the matching of psychology textbooks, a region in which the relations between the IFD and matching can be directly addressed.

## Chapter 9

# Sampling strategies, matching, and the ideal free distribution

---

In the previous chapter, the many parallels between the ideal free distribution (IFD) and the matching law were explored indirectly, in the context of a single patch IOS model containing both rich and poor resources. Matching to available resources was observed as a consequence of the levels of interference associated with each resource type, but the relations between matching and the IFD itself were not directly addressed. In this chapter, the focus returns to these relations by assessing a claim in the literature (Thuisjman et al., 1995) that a particular foraging strategy, ‘ $\epsilon$ -sampling’, is capable of (a) leading groups to the IFD (in a dual-patch environment), and (b) leading isolated individuals to obey the matching law. As in the previous chapter, this represents an attempt to understand potentially suboptimal behaviour in terms of adaptation to a group context.

The primary contribution of this chapter is to show that Thuisjman et al.’s claim is false, and to describe an alternative (and novel) strategy,  $\omega$ -sampling, which succeeds where  $\epsilon$ -sampling fails. Both strategies involve simple rules applied to patch-switching decisions, the significant difference is that whereas the former only maintains a single estimate of environment ‘value’, the latter maintains value estimates for each distinct patch.

This chapter departs from the IOS modelling strategy that has underpinned much of this dissertation. The operational descriptions of both  $\epsilon$ -sampling and  $\omega$ -sampling reside at a higher level of abstraction than the sensorimotor mediated agent-environment interaction patterns of the previous chapters. Although this approach necessarily forfeits some of the explanatory flexibility offered by the IOS methodology, it has the advantage of bearing close relation to existing work in theoretical biology. Another departure has to do with the role of artificial evolution, which, for the most part, is utilised only to elucidate performance differences (by determining values of parameters required for the strategy operation), and *not* to engender the strategy properties themselves. Importantly, the primary contributions of this chapter do *not* depend on this methodological innovation.

The remainder of the chapter is organised as follows. First, the IFD and the matching law are rehearsed in some detail. The model itself is then described, and  $\epsilon$ -sampling fully defined. The first set of results demonstrate the falsity of the claims of Thuisjman et al. (1995). The  $\omega$ -sampling strategy is then defined and analysed. A discussion of the role of artificial evolution follows, and

the chapter ends with a brief analysis of the influence of environmental ‘noise’ on the evolution of strategy memory, an analysis which evokes some of the themes prominent at the outset of this dissertation.

## 9.1 The ideal free distribution in some detail

In a recent edition of *Trends in Ecology and Evolution*, Thomas Weber (1998) writes that “[i]n the seemingly limitless supply of quantitative models in behavioural ecology, efforts sold under the label of the ‘ideal free distribution’ [...] enjoy an almost unrivalled popularity” (p.89). This popularity, however, is attended by a high level of controversy, with authors repeatedly disagreeing about the extent to which theory matches observation (Parker & Sutherland, 1986; Kennedy & Gray, 1993; Milinski, 1994), and attempting to sort out the relations between the various theoretical terms implicated in ideal free theory (Tregenza, 1994; Lessells, 1995; Van der Meer & Ens, 1997). In what follows I will attempt to provide a detailed account of the IFD (some of which will be familiar from previous chapters) whilst maintaining some distance from the intricacies of the controversies themselves.

### 9.1.1 Immediate consumption versus standing stock models

Fretwell and Lucas (1970, see also Fretwell, 1972) described the IFD only in very general terms, assuming (a) that animals are ‘ideal’, in that they are able to choose the habitat (or ‘patch’) that maximises their fitness rewards, and (b) that animals are ‘free’, in that there are no costs associated with movement between habitats. Additionally, habitats are assumed to be ‘density dependent’, in that the fitness value of a habitat is taken to be inversely related to the number of animals in that habitat. Given this general situation, the IFD predicts that a stable distribution of animals will be reached when no animal can improve its fitness rewards by moving to a different habitat, and that in this state all animals will obtain equal rewards.

Recent interpretations of the IFD have focussed on the distribution of predators across patches with different availabilities of prey (Milinski & Parker, 1991; Kennedy & Gray, 1993). In order to make specific predictions about the distributions of predator groups, these models require a way of relating the per predator intake rate  $W_i$  ( $s^{-1}$ ) to both the number of predators  $N_i$  and the number of prey  $F_i$  on each patch  $i$ . Two ways of doing this have emerged from the literature, labelled by Van der Meer and Ens (1997) as ‘immediate consumption’ and ‘standing stock’. Immediate consumption models describe situations in which there is a constant rate of prey input into each patch in the environment, with each prey item being consumed as soon as it arrives (Parker & Stuart, 1976; Parker, 1978). A typical way to write down  $W_i$  in an immediate consumption model is given in Milinski and Parker (1991):

$$W_i = \frac{V_i}{N_i},$$

in which  $V_i$  ( $s^{-1}$ ) represents that rate of prey input, and  $N_i$  represents the number of predators in patch  $i$ . Across  $i$  patches, an equilibrium IFD will be reached when  $W_i$  is equal for all  $i$ . In an environment with two patches  $A$  and  $B$ , the immediate consumption IFD will therefore describe a situation in which the number of predators in each patch is directly proportional to the rate of prey

input; this being the ‘input matching rule’ (Parker, 1978), also known as the ‘habitat matching rule’ (Pulliam & Caraco, 1984):

$$\frac{N_A}{N_B} = \frac{V_A}{V_B}. \quad (9.1)$$

Immediate consumption situations are notoriously rare in the field. Two examples that have been claimed to exist are dung flies, *Scatophaga stercoraria*, competing to mate with arriving females (Parker, 1970), and fish feeding on items of food drifting downstream (Milinski, 1979).

The second class of IFD model, known either as ‘standing stock’ models or (more usually) as ‘interference’ models, deal with situations in which there is a relatively constant prey density in each patch (Sutherland, 1983; Parker & Sutherland, 1986). One way to write down the per predator intake rate in these models is given below, adapted from Sutherland (1983) and Milinski and Parker (1991):

$$W_i = \frac{QF_iF^*}{N_i^m}, \quad (9.2)$$

in which  $Q$  ( $ms^{-1}$ ) is a measure of patch-independent predator search efficiency,  $F_i \in [0.0, 1.0]$  (dimensionless) represents the fraction of prey (the resource level) on patch  $i$ ,  $F^*$  represents the total prey available, and  $m$  (dimensionless) is the interference constant, which is usually taken to vary between 0.0 (no interference) and 1.0 (high interference); recall from chapter 7 the definition of interference as the more-or-less immediately reversible decline in intake due to the presence of conspecifics (Goss-Custard, 1980; Sutherland, 1983).<sup>1</sup> This equation can be easily rearranged to predict the (ideal free) distribution of predators across two patches  $A$  and  $B$ , assuming  $W_A = W_B$ :

$$\log \frac{N_A}{N_B} = \frac{1}{m} \log \frac{F_A}{F_B}, \quad (9.3)$$

this being the ‘generalised habitat matching rule’ of Fagen (1987). Also, taking the total number of predators to be  $N_T (= N_A + N_B)$ , it is possible to predict the value of both  $N_A$  and  $N_B$  directly (see appendix B for details; the following equation also appears in Tregenza et al., 1996):

$$N_A = \frac{N_T}{(10^{-c} + 1)}, \quad c = \frac{\log \frac{F_A}{F_B}}{m}. \quad (9.4)$$

We will call this the ‘simple’ form of the standing stock model. Notice that if  $m = 0.0$ , this model describes an interference-free situation in which all predators would be expected to congregate in the richest patch. For this reason the label ‘standing stock’ is preferred here to ‘interference’. Notice also, if  $m = 1.0$ , the generalised habitat matching rule collapses to a version of Pulliam and Caraco’s original habitat matching rule (equation 9.1), in which prey input rate is replaced by (standing stock) resource levels:

$$\frac{N_A}{N_B} = \frac{F_A}{F_B}.$$

<sup>1</sup> Values of  $m$  in excess of 1.0 are possible, and can be expected in cases in which prey items can be *lost* (for example, by fleeing) as a result of interference.

A more general version of the standing stock model is also described in Sutherland (1983). It derives from a combination of the ‘type II functional response’ (Holling, 1959), and the interference relationship identified by Hassell and Varley (1969). Holling’s functional response expresses how  $W_i$  is related to prey availability:

$$W_i = \frac{aF_iF^*}{1 + ahF_iF^*},$$

where  $a$  ( $ms^{-1}$ ) is a measure of predator search efficiency in the absence of interference, and  $h$  is the handling time (expressed as a fraction of unit time) required for each captured prey item. This can be combined with Hassell and Varley’s model of interference whereby:

$$a = QN_i^{-m},$$

with  $Q$  defined as above, and this leads to the following statement of intake rate:

$$W_i = \frac{QF_iF^*N_i^{-m}}{1 + QhF_iF^*N_i^{-m}}. \quad (9.5)$$

This expression is employed in many IFD models, including those of Bernstein et al. (1988, 1991) discussed later. It also can be rearranged to reveal the generalised habitat matching law (equation 9.3); the necessary algebra is laid out in appendix B. Notice that if  $h = 0.0$ , this equation is equivalent to the simple form (equation 9.2), and indeed it is this ‘simple’ standing stock model that is used in this chapter to generate IFD predictions, for reasons which will be discussed in a moment.

Lessells (1995) has pointed out that the distinction between immediate consumption and standing stock models may not be entirely clear in real situations, arguing that ‘immediate consumption’ is a mathematical idealisation, and that in any real situation some persistence of prey must be anticipated. By the same token, most (if not all) standing stock situations will involve some rate of prey input or regeneration. Consider, for example, the ‘single patch’ models of the previous chapters. In these models, a constant (and non-zero) density of prey is ensured by immediately replacing consumed prey. To the extent that prey input rate and standing stock levels are treated as different sides of the same coin, then, it is possible to collapse the distinction between the the immediate consumption model and the simple standing stock model with  $m = 1.0$ . It is left for the reader to decide whether or not this is a useful interpretation, the arguments of this chapter are unaffected either way.

### 9.1.2 Resource allocation

The foraging strategies described in this chapter operate over discrete time intervals, and as such it is possible to interpret the resource level  $F_i$  (in the simple standing stock model) in at least two ways. The first is simply to take it as specifying a resource level that contributes to the intake of every agent (or predator) in the patch at every time-step. This process of ‘continuous allocation’ (C-allocation) is the usual interpretation in the literature (see, for example, Bernstein et al., 1988, 1991; Tyler & Hargrove, 1997). In this case, equation 9.2 can be used at every time-step, exactly as it is written. The second approach is to understand  $F_i$  as specifying a probability that patch

$i$  will yield the fixed resource quantity  $F^*$  to each agent in the patch at each time-step. Under this process of ‘probabilistic allocation’ (P-allocation),  $W_i$  becomes a random variable:

$$W_i = \begin{cases} \frac{QF^*}{N_i^m}, & p(F_i) \\ 0, & p(1 - F_i) \end{cases} \quad (9.6)$$

The IFD condition of equal intake rates across all patches in this case must apply to *expected* intake rates over many time-steps. We can write:

$$E(W_i) = \left( \frac{QF^*}{N_i^m} \right) F_i.$$

from which the condition  $E(W_A) = E(W_B)$  leads to the same generalised habitat matching law described above (equation 9.3). Notice, however, that this is *not* the case for the general form of the standing stock model with non-zero handling time. The expected intake rate under this model would be:

$$E(W_i) = \left( \frac{QF^* N_i^{-m}}{1 + QhF^* N_i^{-m}} \right) F_i.$$

Appendix B shows that the condition  $E(W_A) = E(W_B)$  has no simple analytical solution for  $\frac{N_A}{N_B}$  unless  $QhF^*$  is very small; the use of this general form with non-zero handling time would therefore be likely to require numerical techniques for the derivation of predicted IFD distributions. Although this is certainly possible in principle, in this chapter the simple form is preferred (equation 9.2), since with this form IFD predictions can be derived analytically from the generalised habitat matching law (equations 9.3 and 9.4) under either allocation method.

No claims are made for the biological relevance of the distinction between C-allocation and P-allocation; the distinction is motivated by analogous resource allocation methods often employed in ‘matching law’ choice experiments, described below. One possible intuition, however, is that it may reflect the difference between relatively accessible and widespread types of resource (grass, for example), and relatively inaccessible yet potent types of resource (truffles, for example).

### 9.1.3 Summary

The IFD literature describes two kinds of model; immediate consumption and standing stock. This chapter will develop an individual-based standing stock IFD model which will be analysed under two interference levels (1.0 and 0.3), and under both C-allocation and P-allocation.

## 9.2 The matching law

Even though the individual matching law was the focus of the previous chapter, its relatively unorthodox interpretation therein requires that it be rehearsed in some detail here. Consider the concise definition of Krebs and Kacelnik (1991): “the matching law states that the animal allocates its behaviour to two alternatives in proportion to the rewards it has obtained from them” (p.131). If the proportionality is direct, this is known as ‘strict’ matching (Davison & McCarthy, 1988):

$$\frac{B_A}{B_B} = \frac{R_A}{R_B}, \quad (9.7)$$

where  $B_A$  and  $B_B$  represent the rate of response to options  $A$  and  $B$ , and  $R_A$  and  $R_B$  represent the resources obtained in each case. The ‘generalised’ matching law (Baum, 1974) includes parameters for bias ( $b$ ) and sensitivity ( $s$ ) to account for the departures from strict matching often observed in empirical data:

$$\log \frac{B_A}{B_B} = s \cdot \log \frac{R_A}{R_B} + \log(b). \quad (9.8)$$

Recall the parallels that may be drawn between the matching law and the IFD: (1) the former is to do with individual choice, the latter with the collective consequences of individual choice, (2) the controlled environments of psychological experiments are often treated as abstractions of natural foraging environments (Shettleworth, 1988; Dallery & Baum, 1991), and (3) their mathematical congruence; compare Pulliam and Caraco’s habitat matching law (equation 9.1) with the strict matching law above (equation 9.7), and also compare the generalised habitat matching law (Fagen, 1987, equation 9.3) with the generalised individual matching law above (equation 9.8). Recall also the differences that remain: whereas habitat matching predictions are normative, the individual matching law is an observed relation, and whereas habitat matching is expressed in terms of *available* resources ( $F_i$ ), the individual matching law is expressed in terms of *obtained* resources ( $R_i$ ). Of course,  $F_i$  can still be used in the context of individual matching even if it is not represented in the matching equations themselves, and indeed it is necessary to do so in order to describe the various ‘schedules of reinforcement’ by which resources are allocated in matching experiments.<sup>2</sup>

Experimental psychologists have investigated the matching law under many different schedules of reinforcement. In this chapter, four such schedules are considered for rewarding responses to two options  $A$  and  $B$  (with associated resource availabilities  $F_A$  and  $F_B$ ):

- *Basic*: Each response is rewarded with an amount determined by the relative values of  $F_A$  and  $F_B$ . Responses are rewarded at every time-step. This is analogous to the C-allocation method of section 9.1.2.
- *Concurrent (conc) VR VR*: A variable ratio (VR) schedule indicates that an option must receive a certain number of responses before a reward is given. This number can vary around a mean value, and can therefore be implemented by associating a probability of reward with each option.  $F_A$  and  $F_B$  are here interpreted as the mean values, so that conc VR VR is analogous to the P-allocation method in section 9.1.2.
- *Concurrent (conc) VI VI*: A variable interval (VI) schedule requires that a certain delay elapse after a reward on a given option until that option can be rewarded again. This delay time can vary around a mean, and these means can differ between response option ( $F_A$  and  $F_B$  are interpreted as the delays).
- *Concurrent (conc) VI VR*: This is a ‘mixed’ schedule in which one choice option is rewarded under a VI schedule, and the other under a VR schedule.

There is a substantial literature surveying human and animal behaviour under these schedules (Davison & McCarthy, 1988). Under both basic and conc VR VR schedules, the general consensus is that exclusive choice for the most profitable option is observed (although see Myers, 1976, and

<sup>2</sup>The terms ‘reinforcement’ and ‘reward’ are used interchangeably; ‘reinforcement’ is employed only when it helps to maintain consistency with the psychological literature.



parts of Sutherland and Mackintosh, 1971, for data to challenge this consensus). There is nothing counterintuitive about this; if repeatedly offered a choice between 80p and 40p, any sensible subject would presumably choose the former 100% of the time, and the same would apply to repeated choices between odds of 3:1 and odds of 5:1. Observations of exclusive choice, although consistent with the matching law, are only trivial instances of its applicability, as such these schedules present relatively undemanding assessments of matching behaviour.

The conc VI VI schedule provides a much more interesting assay. Indeed, it was with this schedule that the individual matching law was first formulated (Herrnstein, 1961). Recall from the previous chapter that most paradigms in experimental psychology are designed to eliminate direct contingencies between response and reward, since otherwise matching to obtained resources can only occur with either resource types of equal value and/or exclusive choice (appendix B). The basic and conc VR VR schedules described above maintain this kind of direct contingency, hence the trivial predictions and unsurprising observations of exclusive choice. Under conc VI VI, however, the reward rate can be largely independent of the response rate, such that matching to obtained resources can be achieved with a variety of response distributions, including - but by no means limited to - exclusive choice. Furthermore, under conc VI VI, exclusive choice is no longer the optimal response pattern (Herrnstein, 1970). Matching to obtained resources under conc VI VI has been observed for both non-humans animals (Davison & McCarthy, 1988) and human subjects (Conger & Killeen, 1974), in all cases *without* exclusive choice.

The final schedule, conc VI VR, also leads to observations of matching to obtained resources, in some cases in the trivial form of exclusive choice, and in other cases non-trivially, depending on the relative productivities of the two component schedules (Herrnstein & Heyman, 1979; Herrnstein & Vaughan, Jr., 1980). The most important feature of this schedule is its relation to the maximisation of reward. Under basic, conc VR VR, and conc VI VI schedules, matching to obtained resources also leads to maximisation. Under conc VI VR, matching (whether trivial or not) is *not* the optimal behaviour pattern. It is therefore only with this schedule that the ambit of this chapter extends to the potential individual suboptimality associated with matching. The consensus in the literature, on the basis of experiments with many species of animal subject, is that matching to obtained resources - not maximisation - is observed under conc VI VR (Herrnstein & Heyman, 1979; Herrnstein, 1997).

### 9.3 A Description of the model

The present model is essentially an individual-based IFD model operating at the level of patch-switching rules. Models of this kind, which are steadily increasing in popularity, are mostly motivated by the insight they can afford into the situations in which groups will actually settle into the IFD. Bernstein et al. write:

[T]he IFD describes the distribution of a population, but it is not immediately obvious that without assuming integrated behaviour in which a whole group of consumers 'agree' on the equilibrium distribution, the population will end up in the IFD. (1991, p.207)

Perhaps the most extensive body of individual-based IFD modelling work remains that of Bernstein et al. (1988, 1991), who consider how individual foragers, using decision rules based on past

experience, might distribute themselves between patches of varying resource density. Foragers are assumed to continually assess the average availability of resource in the environment, according to the ‘learning rule’:

$$E(t+1) = \gamma E(t) + (1-\gamma)r(t), \quad 0 \leq \gamma \leq 1$$

where  $E(t)$  is the estimate of overall environmental quality at time  $t$ ,  $r(t)$  is the resource acquired during time-step  $t$ , and  $\gamma$  is the ‘adaptation rate’. They are also assumed to move from their current patch, to a randomly selected alternative, according to the following ‘decision rule’: “migrate whenever gain rate in the current patch is lower than the expected mean gain rate for the environment as a whole” (1991, p.207), i.e. whenever  $r(t) \leq E(t)$ . This ‘decision rule’ is an expression of the ‘marginal value theorem’ (Charnov, 1976), and its combination with the above learning rule has features in common with both the strategies explored in this chapter, and also with the ‘adaptive aspiration rate’ strategy for playing the iterated prisoner’s dilemma discussed in chapter 3 (see also Posch, 1999).

Bernstein et al. demonstrate that, in many situations, groups of foragers following these ‘learning’ and ‘decision’ rules do indeed settle into the IFD. In particular, they show that the IFD can be attained without the assumption of omniscience on the part of the foragers, and they extend their basic model to explore (amongst other things) the influences of travel costs, prey depletion, and environmental structure. They also make a point of claiming that their work connects with optimal foraging theory (see chapter 5) inasmuch as their choice of agent strategy “has been shown to be the optimal learning rule under specific forms of environmental stochasticity” (1988, p.1019). However, be this as it may, they can provide no guarantee that their strategy is in fact optimal given the particular constraints of their model. This debate has already been treated at length in the context of the IOS methodology (chapter 5), and as a consequence no similar claims are made here for either  $\varepsilon$ -sampling or  $\omega$ -sampling.

### 9.3.1 $\varepsilon$ -sampling

Despite their thorough treatment of the IFD, Bernstein et al. do not consider the individual matching law at all. This is why the present chapter instead focuses on the work of Thuisjman et al. (1995), who claim that a related foraging strategy, which they call  $\varepsilon$ -sampling, leads groups of agents to the IFD *and* leads individual agents to obey the individual matching law. A primary objective of this chapter is to show that these claims are false.

The idea behind  $\varepsilon$ -sampling is that agents stay on a ‘current’ patch, and occasionally ‘sample’ other patches. If the ‘sampled’ patch is better than the ‘current’ patch, the agent will switch and the ‘sampled’ patch will become the new ‘current’ patch. This is similar to the approach of Bernstein et al. insofar as there is a ‘learning rule’ by which the agent maintains a continuously updated estimate of environmental quality. However, rather than using this estimate to directly drive patch switching via the marginal value theorem,  $\varepsilon$ -sampling uses it only to decide whether or not to accept a stochastically instigated ‘sampling’ switch.

More formally, given two alternatives  $A$  and  $B$ , the  $\varepsilon$ -sampling strategy is implemented as follows. An agent initially selects  $A$  or  $B$  at random. At each subsequent time interval, the agent abides by its choice with probability  $(1 - \varepsilon)$ , and samples the other option with probability  $\varepsilon$ .

The agent will remain with the new option (with probability  $1 - \varepsilon$ ) if the reward from this option exceeds the agent's 'critical level' ( $E$ ), which is a dynamic estimate of the 'value' of the environment, in which more recent rewards are more strongly represented to a degree specified by the *adaptation rate*  $\gamma$  of the agent. The operational definition of  $\varepsilon$ -sampling given below is from Thisjman et al. (1995):

**Definition 1** Let  $\gamma, \varepsilon \in (0, 1)$ , let  $M(t) \in A, B$  represent the option selected and let  $r(t)$  be the resources obtained at time  $t \in \{1, 2, 3, \dots\}$ . Define  $E(1) = 0$  and

$$E(t+1) = \gamma E(t) + (1 - \gamma)r(t)$$

for ( $t \geq 1$ ). Then  $E(t)$  is called the critical level at time  $t$ . Let  $A_\varepsilon$  denote the behaviour of choosing  $A$  with probability  $(1 - \varepsilon)$  and  $B$  otherwise. Let  $B_\varepsilon$  be defined similarly. The  $\varepsilon$ -sampling strategy is then defined by playing:

at ( $t = 1$ ) use  $A_{0.5}$ ,

at ( $t = 2$ ) use  $M(1)_\varepsilon$ ,

at ( $t > 2$ ) use  $M(t-1)_\varepsilon$  in case  $M(t-1) \neq M(t-2)$  and  $r(t-1) > E(t-1)$ , otherwise use  $M(t-2)_\varepsilon$ .

In terms of the balance between exploitation (of the best known option at any time) and exploration (of perhaps better alternatives), the sampling frequency  $\varepsilon$  can be understood as specifying explicitly the level of exploration.

### 9.3.2 Model structure

To investigate  $\varepsilon$ -sampling in the context of an individual-based model it is necessary to specify values for  $\varepsilon$  and  $\gamma$  for each agent. Usually this would involve choosing 'plausible' values which would then be held constant across all experimental conditions. Here we develop a novel approach, using a GA to find near-optimal values for  $\varepsilon$  and  $\gamma$  given the constraints of each strategy and each experimental condition. Whilst a full discussion of the practical details and merits of this approach awaits section 9.4, it can be said immediately that the results described below do *not* depend on this innovation. Analysis of the model was in fact carried out twice: once using near-optimal parameters, and once using a pre-specified and fixed parameter set (plus variations on this set to assess sensitivity) derived from the average near-optimal values across all experimental conditions:  $\bar{\varepsilon} = 0.052$ , and  $\bar{\gamma} = 0.427$ . No significant differences were observed.

The first stage of analysis in each case involved recording the equilibrium distribution (after 1000 time-steps) of populations of 100  $\varepsilon$ -sampling agents, for each of 9 different resource distributions across two patches  $A$  and  $B$ . Four separate populations were analysed, one for each combination of interference level (1.0 or 0.3) and allocation method (C-allocation or P-allocation). In each case, agents were initially randomly allocated to either  $A$  or  $B$ . Then, each time-step, the resource obtained by each agent was calculated (equation 9.2 for C-allocation and equation 9.6 for P-allocation), the  $\varepsilon$ -sampling strategy applied, and the new distribution of agents determined. The final equilibrium distributions were compared with the predictions of the IFD (equation 9.4).

The second stage involved the analysis of isolated individual behaviour under the various reinforcement schedules described in section 9.2 in terms of the predictions of the individual

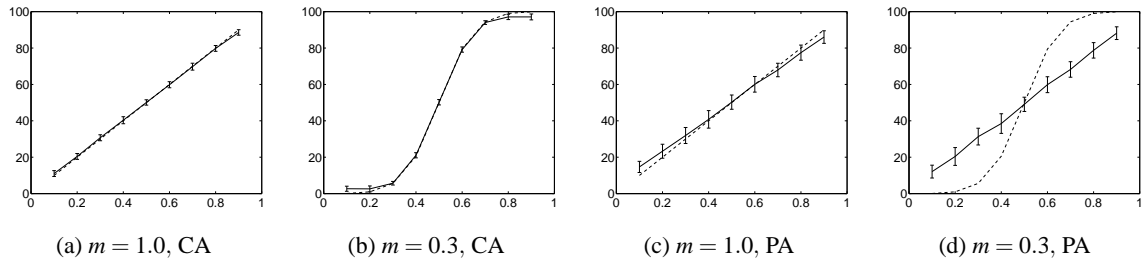


Figure 9.1: Observed (solid) and predicted (dashed)  $\epsilon$ -sampling population distributions under 9 different resource distributions (IFD predictions obtained using equation 9.4). Each observation derives from the mean of 30 distributions, standard deviations are shown. Each abscissa represents  $F_A$  and each ordinate represents the percentage of agents on patch A. Four conditions are shown, defined by all combinations of interference level (1.0 or 0.3), and C-allocation or P-allocation (CA or PA).

matching law. For the basic and conc VR VR schedules, single  $\epsilon$ -sampling agents were allowed to forage in isolation, under C-allocation or P-allocation respectively, for 1000 time-steps under each of 9 different resource distributions. Conc VI VI was implemented by using  $F_i$  to set delay intervals ( $D_i$ ) such that  $D_i = 20(1.0 - F_i) + r$ , with  $r \in [-2, 2]$  an integer random number. The first response to option  $i$  on each evaluation procured the full reward  $F^*$  and initialised  $D_i$ . Subsequent responses to  $i$  went unrewarded until  $D_i$  time-steps had elapsed, after which a response would again procure  $F^*$  and re-initialise  $D_i$ , with the incorporation of  $r$  ensuring that the schedule was indeed ‘variable interval’. The implementation of conc VI VR consisted, naturally, of applying the VI schedule to one option (A), and P-allocation to the other (B). In both cases, as before, isolated agents were allowed to forage for 1000 time-steps under each of 9 different resource distributions. Note that evolved near-optimal parameter values could not be used for assessing performance under conc VI VI or conc VI VR, since no populations were evolved under these conditions. The average values  $\bar{\epsilon}$  and  $\bar{\gamma}$  were used instead (see section 9.4).

Although the psychological matching literature concentrates almost exclusively on isolated agents, in many biological studies the emphasis is on the matching behaviour of individuals embedded in groups (Gray, 1994; Baum & Kraft, 1998). The final stage of analysis therefore consisted of recording the behaviour of embedded individuals from each of the 4 original populations over the full 1000 time-steps, under each of the 9 resource distributions, comparing their behaviour with the predictions of the individual matching law.

All three stages of analysis were repeated 30 times enabling means and standard deviations to be calculated.

### 9.3.3 Results: $\epsilon$ -sampling

Figure 9.1 compares observed distributions of  $\epsilon$ -sampling agents to the predictions of the simple standing stock IFD (equation 9.4). Although in most cases there is a good match,  $\epsilon$ -sampling agents are unable to find the IFD under P-allocation with  $m = 0.3$ .

With regard to individual matching (figure 9.2),  $\epsilon$ -sampling agents exhibit exclusive choice (trivial matching) under basic reinforcement, as predicted by the psychological data (9.2a,e). Under conc VR VR, however, although they continue to match to available resources, they no longer

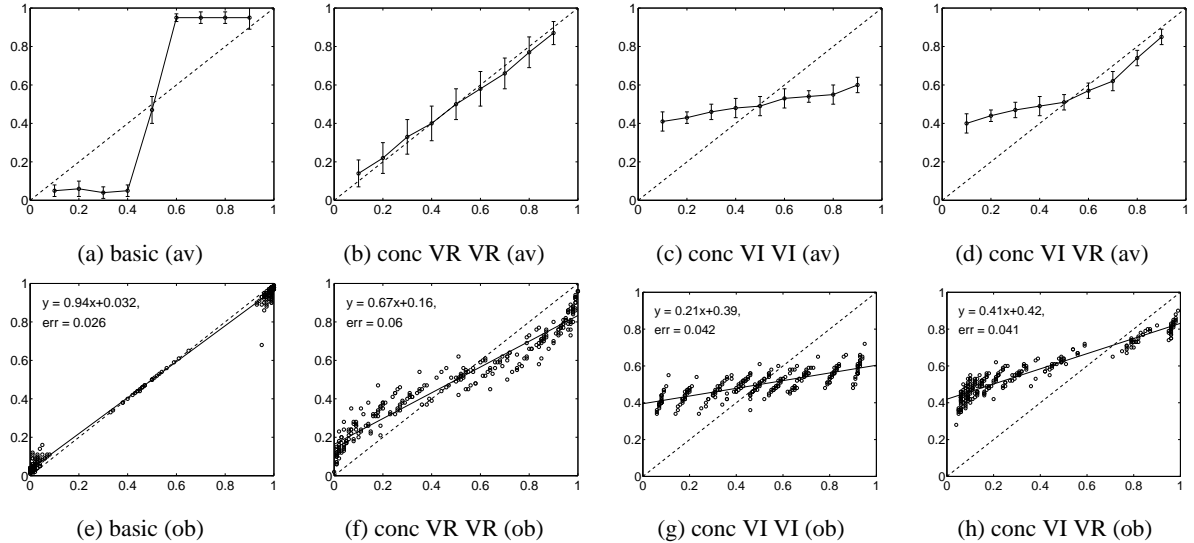


Figure 9.2: Matching behaviour of isolated  $\epsilon$ -sampling agents. Data for each plot is collected from 30 analyses at each of 9 values of  $F_A$ , with dashed lines indicating strict matching. Plots labelled (av) concern matching to *available* resources; solid lines show mean proportion of time spent on A (ordinate) as a function of  $F_A$  (abscissa), standard deviations are shown. Plots marked (ob) concern matching to *obtained* resources; mean proportion of time spent on A (ordinate) is scatter-plotted as a function of proportion of resources obtained from A (abscissa), with best-fit lines superimposed. The equation of each best-fit line is given together with a measure of goodness-of-fit (this ‘error’ measure specifies the range around any point on the line that contains at least 50% of the predictions).

match to obtained resources, and certainly do not exhibit exclusive choice (9.2b,f). This conflicts with individual matching predictions. Performance is no better under conc VI VI or conc VI VR; in both cases there are clear departures from strict matching to obtained resources (9.2g,h).

Embedded  $\epsilon$ -sampling agents in most cases match closely to obtained resources (figure 9.3), although there is some divergence from strict matching when  $m = 0.3$  under P-allocation (9.3h).

### 9.3.4 Discussion: $\epsilon$ -sampling

The above results demonstrate that  $\epsilon$ -sampling can neither reliably lead populations of agents to distribute according to the IFD, nor reliably lead individual agents to match to obtained resources. These findings directly repudiate the claims of Thuisjman et al. (1995) that  $\epsilon$ -sampling is capable of underlying both matching and the IFD. Why, one must ask, did they make these claims in the first place? One likely reason is that they considered only a small set of analytically tractable special cases. With respect to the IFD they explored only C-allocation with an interference level  $m = 1.0$ . Here we have seen that  $\epsilon$ -sampling does indeed lead populations to the IFD in this condition, but that there is at least one other condition (representative of many others) in which it does not. With respect to the matching law, they analysed isolated  $\epsilon$ -sampling agents only under the equivalent of the conc VR VR schedule of the present model. This is not a useful way to explore individual matching, since under this schedule one should expect only trivial adherence to the matching law in the form of exclusive choice. Moreover,  $\epsilon$ -sampling does not even conform to

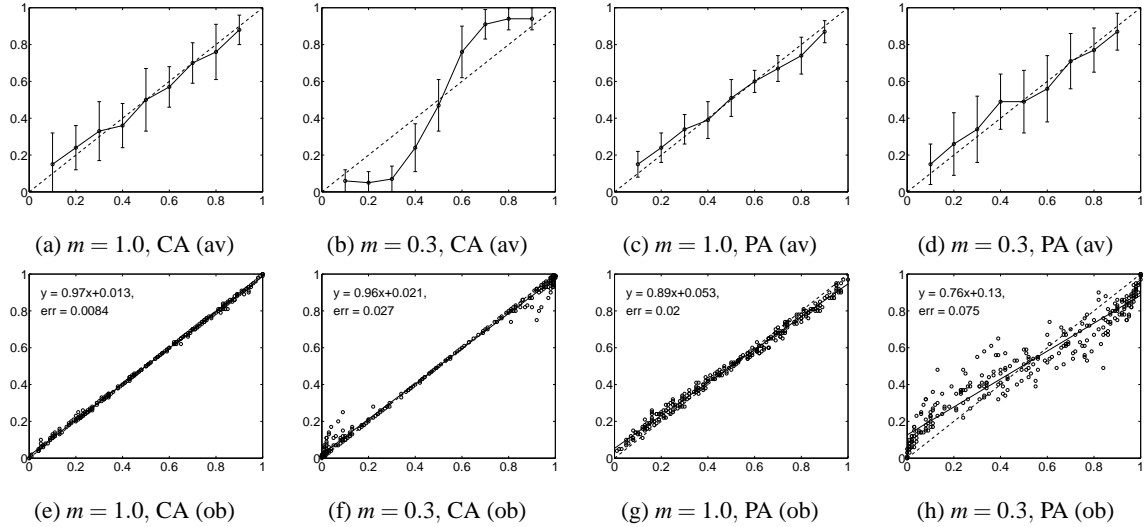


Figure 9.3: Matching behaviour of embedded  $\varepsilon$ -sampling agents. Dashed lines represent strict matching. Each plot shows mean proportion of time spent on  $A$  (ordinate) as a function of resources obtained from  $A$  (abscissa), with data collected from 30 analyses at each of 9 values of  $F_A$ . Best fit lines are superimposed and equations (with goodness-of-fit) are given as in figure 9.2.

this prediction, instead leading agents to match to available resources. Unfortunately, this result, which was also obtained by Thuisjman et al. (1995), was wrongly asserted by them to be consistent with the individual matching law. This is testament to their misinterpretation of the matching law as describing matching to available resources, a misinterpretation evident throughout their paper, and their claim that  $\varepsilon$ -sampling can underlie matching rests entirely on this confusion. The present model has illustrated, to the contrary, that  $\varepsilon$ -sampling matches to obtained resources only under basic reinforcement. It should be stressed that the results obtained from assessing  $\varepsilon$ -sampling in the present model are in agreement with those derived by Thuisjman et al. (1995) *in those special cases considered by them*. The problem is that these special cases are not appropriate for an assessment of their claims.

### 9.3.5 $\omega$ -sampling

One possible explanation for the poor performance of  $\varepsilon$ -sampling is its extremely conservative use of environmental information, involving only a single measure of overall environmental quality. In this section we extend our investigation to analyse a related, and novel strategy,  $\omega$ -sampling, which is much more aggressive in its accrual and usage of information. In what follows it is demonstrated that  $\omega$ -sampling, unlike  $\varepsilon$ -sampling, is capable of reliably leading groups of agents to the IFD, and also of reliably leading individuals to obey the matching law.

The  $\omega$ -sampling strategy combines features from  $\varepsilon$ -sampling and from the strategy described in Bernstein et al. (1988, 1991). The use of sampling to drive patch switching derives from  $\varepsilon$ -sampling, but switching can also be driven by a version of the marginal value rule, as in Bernstein et al.. However, the most significant feature of  $\omega$ -sampling is that agents maintain concurrent estimates of *each* (visited) patch in the environment, rather than (as is the case both for  $\varepsilon$ -sampling and for Bernstein et al.) a single estimate of environmental quality as a whole. The two patch environment of the present model means that this is not a particularly demanding assumption, the

implications of its relaxation in relatively complex environments are discussed in section 9.3.7.

For a two patch (or two choice) environment this strategy is implemented as follows. A forager initially selects  $A$  or  $B$  at random. At each subsequent time interval, the forager samples the other option with probability  $\varepsilon$ , otherwise (with probability  $1 - \varepsilon$ ) it compares the estimate of its current selection with that of the unselected option, and switches if the former is the lower of the two. An operational definition of the strategy is given below:

**Definition 2** Let  $\gamma, \varepsilon, M(t), r(t)$  be as in Definition 1, let  $E_A(t)$  and  $E_B(t)$  represent the estimated values of options  $A, B$ , and let  $N(t)$  represent the unselected option at time  $t \in \{1, 2, 3, \dots\}$ . Define  $E_A(1) = E_B(1) = 0$ . For  $(t \geq 1)$  then if  $M(t) = A$ :

$$E_A(t+1) = \gamma E_A(t) + (1 - \gamma)r(t), \quad E_B(t+1) = E_B(t),$$

otherwise (if  $M(t) = B$ ):

$$E_A(t+1) = E_A(t), \quad E_B(t+1) = \gamma E_B(t) + (1 - \gamma)r(t).$$

Let  $\mathcal{R} \in (0, 1)$  be a random number. Let  $A_\varepsilon$  and  $B_\varepsilon$  be as in Definition 1. The  $\omega$ -sampling strategy is then defined by playing:

at  $(t = 1)$  use  $A_{0.5}$ ,

at  $(t = 2)$  use  $M(1)_\varepsilon$ ,

at  $(t > 2)$  if  $(\mathcal{R} < \varepsilon)$  use  $N(t - 1)$ , else if  $(E_{M(t-1)} < E_{N(t-1)})$  use  $N(t - 1)$ , otherwise use  $M(t - 1)$ .

There is not much difference between  $\omega$ -sampling and  $\varepsilon$ -sampling with regard to the balance between exploitation and exploration. In both cases, the sampling frequency  $\varepsilon$  explicitly specifies the level of exploration. The hope for  $\omega$ -sampling is that a more effective use of information can lead to greater proficiency in both domains. Rather than simply staying on a ‘current’ patch,  $\omega$ -sampling agents should be able to more effectively determine the ‘best’ option at any time, through maintaining value estimates of each. This can only help exploitation. By the same token,  $\omega$ -sampling agents should be better equipped for exploration by being able to build up a picture of the environment over time, rather than relying, as  $\varepsilon$ -sampling does, on one-shot samples.

I do not wish to overstate the novelty of  $\omega$ -sampling. It remains very close in its operational detail to many strategies described in the theoretical biology literature, and is certainly a comparatively trivial strategy in relation to the many reinforcement learning algorithms that populate the computer science literature (see, for example, Kaelbling et al., 1996). Both of these properties - proximity to theoretical biology and relative simplicity - are in fact good reasons for selecting  $\omega$ -sampling from the myriad of possible alternatives. And what is important for present purposes is not its novelty *per se*, but its novelty with respect to the context of the IFD and the matching law and the ways in which it differs from  $\varepsilon$ -sampling.

### 9.3.6 Results: $\omega$ -sampling

The analysis of  $\omega$ -sampling exactly recapitulated that of  $\varepsilon$ -sampling. It is immediately clear that  $\omega$ -sampling strategy outperforms  $\varepsilon$ -sampling in leading groups to the IFD. Figure 9.4 shows that




Figure 9.4: Observed (solid) and predicted (dashed)  $\omega$ -sampling population distributions under  $\omega$ -sampling, to be interpreted as figure 9.1.

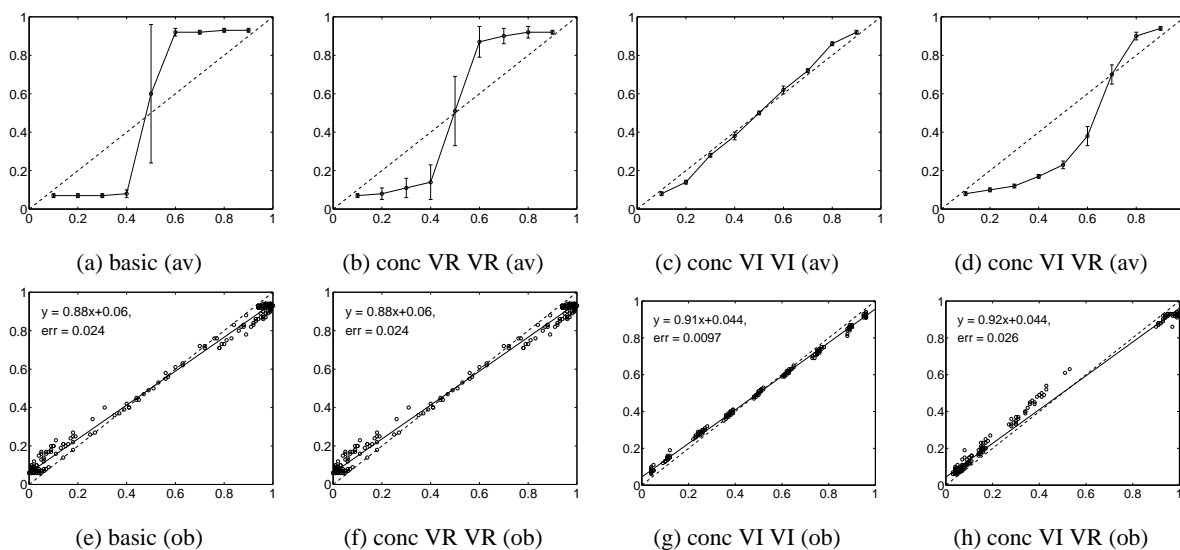


Figure 9.5: Matching behaviour of isolated  $\omega$ -sampling agents, to be interpreted as figure 9.2.

in all 4 conditions ( $m = 1.0$  or  $0.3$ , C-allocation or P-allocation) populations of  $\omega$ -sampling agents closely fit the predictions of equation 9.4.

Adherence to the individual matching law is also improved by  $\omega$ -sampling. Figure 9.5 shows that  $\omega$ -sampling agents exhibit exclusive choice (trivial matching to obtained resources) under basic and conc VR VR reinforcement, in agreement with the psychological literature.<sup>3</sup> Furthermore, matching to obtained resources is also observed under conc VI VI and conc VI VR schedules (figure 9.5g,h). (The slight deviations from strict matching entailed by  $\omega$ -sampling under conc VI VR are in the *opposite* direction to that expected if agents were maximising reward; see Herrnstein & Heyman, 1979.)

Embedded  $\omega$ -sampling agents also reliably match to obtained resources under all 4 test conditions (figure 9.6). Notice, however, that these observations are not reflected in the relatively accessible (in the field) statistic of matching to *available resources*. The significance of this is that if embedded agents are *not* observed to match to available resources, it *cannot* be concluded that isolated agents using the same strategy would fail to match to obtained resources. In other words, matching behaviour in the field may not be a reliable indicator of the performance of isolated individuals with regard to the individual matching law.

<sup>3</sup>Careful inspection of figure 9.5(a,b) reveals that the exclusive choice of  $\omega$ -sampling under conc VR VR is not quite as exclusive as it is under basic reinforcement. Although this deviation is slight, it is interesting to note that similar deviations have also been observed in real situations (see, for example, Sutherland & Mackintosh, 1971).



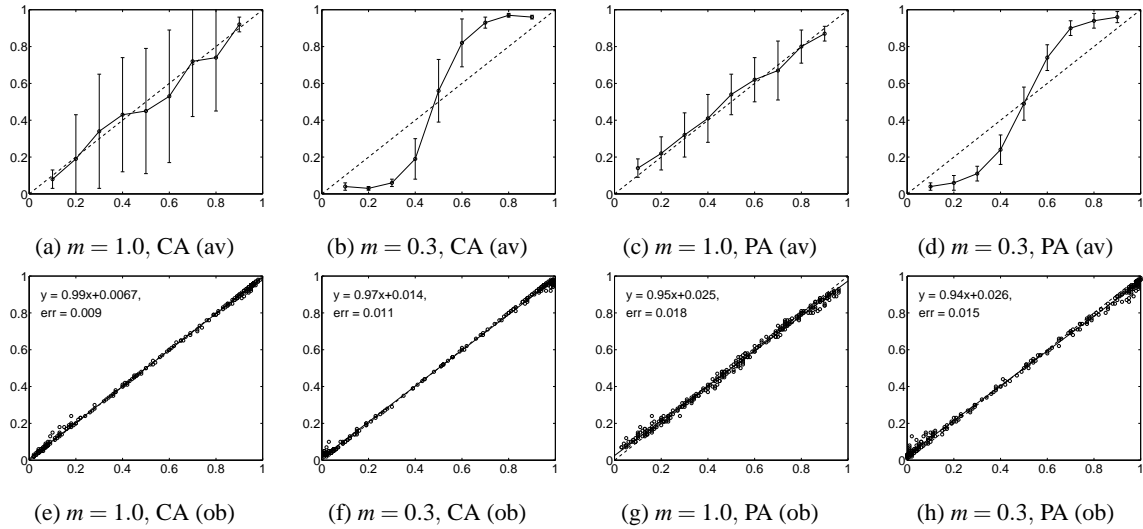


Figure 9.6: Matching behaviour of embedded  $\omega$ -sampling agents, to be interpreted as figure 9.3.

### 9.3.7 Discussion: $\omega$ -sampling

As a candidate mechanism underlying both the IFD and individual matching behaviour,  $\omega$ -sampling is clearly more successful than  $\varepsilon$ -sampling. This is perhaps unsurprising: in virtue of maintaining multiple estimates and employing a marginal value switching rule in addition to stochastic sampling,  $\omega$ -sampling is much less likely than  $\varepsilon$ -sampling to be adversely affected by the indeterminacy of P-allocation (with respect to the IFD) or of the conc VR VR, conc VI VI, and conc VI VR schedules (with respect to the matching law). And it is of course in these very conditions that the inadequacies of  $\varepsilon$ -sampling are revealed. The conclusion may therefore be drawn that there do indeed exist patch-switching rule-based strategies capable of underlying both the IFD, and the individual matching law, in several non-trivial situations;  $\omega$ -sampling is such a strategy,  $\varepsilon$ -sampling is not. This conclusion is broadly representative of a recent movement in evolutionary psychology towards understanding choice behaviour in terms of collections of simple heuristics adapted to ‘biologically plausible’ situations; ‘bounded’ rationality as opposed to ‘general-purpose’ rationality (see, for example, Gigerenzer, Todd, & the ABC research group, 1999).

Of particular importance is that  $\omega$ -sampling entails matching to obtained resources in at least one situation in which such behaviour is suboptimal (the conc VI VR schedule). This returns to a unifying theme of the latter half of this dissertation, that individual suboptimal behaviour can (in some cases) be understood in terms of the operation of mechanisms adapted to a group context. It is, however, to be conceded that the matching behaviour of  $\omega$ -sampling under conc VI VR is not perfect. It may be that slight modifications of this strategy would lead to improved fidelity. Another possibility is that, in the same way that the present model has exposed the inadequacies of  $\varepsilon$ -sampling through the exploration of a wider set of conditions than considered by Thuisjman et al. (1995), it could be that further extensions to the present model would reveal conditions in which  $\omega$ -sampling also fails to lead populations to the IFD, or fails to lead individuals to adhere to the matching law. This would not necessarily be a bad thing; the literature abounds with descriptions of experimental situations in which deviations from strict matching and from the IFD are routinely observed, and it would be a telling observation if  $\omega$ -sampling also deviated in similar ways.

A significant computational cost of  $\omega$ -sampling is of course that it requires agents to concur-

rently maintain more than one value estimate, a cost which would seem to become exacerbated the more patches (or choice options) there are. Certainly, many ‘mechanisms of matching’ proposed in the psychology literature avail themselves of multiple estimates (for example, ‘melioration’, Herrnstein, 1982, or ‘momentary maximisation’, Hinson & Staddon, 1983), but is it necessary to assume that agents are able to maintain concurrent estimates for *every* patch (or option) in their environment? The answer is: only to the extent that agents can specify which patch (or option) to choose. This may seem trite, but, in fact, many patch-switching strategies specify only when to *depart* from a current patch (or option), without specifying where to go afterwards (Bernstein et al., 1988, 1991, being a case in point). Let us concentrate on the IFD for a moment. If only random movement is possible, and all areas of the environment can be accessed with equal ease, then it would only be necessary to maintain a single estimate of environmental quality (to prompt departure). If directed (non-random) movement is possible to any part of the environment (again with equal ease) then concurrent estimates of every patch would be valuable. However, if it is assumed that movement is somewhat restricted, but non-random, then some intermediate solution is likely to be best, at which the agent only maintains a few functionally relevant estimates.

At this point, it is worth recalling the continuum of matching phenomena identified in the previous chapter, ranging from the IOS model of that chapter to the complex reinforcement schedules and dedicated matching mechanisms of the psychology literature. The patch-switching strategies of the present chapter represent a ‘middle ground’ in this continuum, but the discussion so far suggests that as the agent-environment system becomes less abstract, the patch-switching rules themselves may become less separable from the perception (of patch value) and action (movement between patches) of the agents. In other words, the apparent disadvantage of  $\omega$ -sampling, in requiring multiple value estimates, can be thought of as a call for the development of models that explore the space between the models of the present and previous chapters.

To illustrate this further, consider the performances of  $\epsilon$ -sampling and  $\omega$ -sampling under conc VR VR:  $\epsilon$ -sampling matches to available resources,  $\omega$ -sampling matches (trivially) to obtained resources. Now, the conc VR VR schedule in the present model is very similar to the ‘forced-choice’ analysis of the previous chapter, in which situated agents were faced with two resource items and each trial was stopped as soon as an agent reached either one of them. This analysis found that agents evolved with conspecifics (so-called *M*-agents) matched to available resources, thus behaving in a similar way to the present  $\epsilon$ -sampling agents. The difference, of course, is that in the previous chapter, matching to available resources was considered an appropriate interpretation of the matching law, and in the present chapter, it is not. This only goes to emphasise the importance of exploring the terrain that separates the two models. In the previous model, agents could not alter the distribution of their behaviour on the basis of experience, those of the present model do so continually. The situated environment of the previous model demanded reliable contingencies between response and reward, the abstract environment of the present model replaces these contingencies with relatively complex reinforcement schedules, enabling direct engagement with the matching literature. The previous model located the ‘mechanisms of matching’ in the sensorimotor interactions constitutive of interference; in the present model, interference is only indirectly linked to the generation of behaviour. In short, the features of the agent-environment system of the previous model which mandated an interpretation of matching in terms of available

resources, are simply not evident in the present model; the overt similarity between the previous forced choice analysis and the present conc VR VR schedule is easily overshadowed by the many differences that remain.

On the matter of interference, in the previous chapter it was argued that the observed matching behaviour could *not* be understood as a ‘solution’ to the ‘problem’ of interference, given the complex historical interplay between the two. In the present model, by contrast, interference is confined within an equational (standing stock) model, and can interact with matching only indirectly, by specifying the intake rate available from a given patch. Thus, in the present model, interference *a priori* forms part of the (environmental) problem for which matching is a solution. This, once again, illustrates that the benefits offered by the present model (effective engagement with existing literature, for example) come at the cost of a curtailed domain of inquiry.

## 9.4 The role of artificial evolution

### 9.4.1 Near-optimal parameters

The final section of this chapter is largely separate from what has gone before. We discuss here the use of artificial evolution to specify parameter values, this being the main point of divergence of the present model from the ‘orthodox’ individual-based approach represented by Bernstein et al. (1988, 1991) and Thuisjman et al. (1995). Orthodox modelling would advocate the selection of ‘plausible’ values which would then be kept constant across all experimental conditions. However in this study (and arguably in many other individual-based models) it is not obvious what these ‘plausible’ values should be, any selection will be somewhat arbitrary through necessarily relying on information from outside the model itself. An alternative approach is to use an optimisation algorithm to set these values for each experimental condition. Given a fitness function requiring maximisation of reward, near-optimal values can be found for  $\epsilon$  and  $\gamma$  given the constraints of each strategy and each condition.

To consider the relative merits of these approaches in the present case, it must be kept in mind that the objective was to compare the performance of two strategies across a range of experimental conditions. Such a comparison requires some equivalence criteria to be drawn in terms of the strategy parameters  $\epsilon$  and  $\gamma$ . If a fixed (but somewhat arbitrary) parameter set is chosen, it could be argued that because the parameters themselves are identical in all conditions, any performance differences must be due to inherent strategy properties. The problem with this is that the arbitrary parameter set may be more appropriate for some conditions than others, and so any performance differences may, to some extent, reflect imbalances in parameter suitability rather than inherent strategy properties. An alternative equivalence criteria is that of optimality. Optimal parameter values may well vary across conditions, but on the other hand it can now be asserted that, in each condition, each strategy is performing as well as it possibly can, therefore any performance differences really must reflect inherent strategy properties, and cannot be explained away in terms of parameter (un)suitability.

How does this relate to the IOS methodology? Essentially, both modelling strategies employ GAs to attempt a reconciliation of individual-based modelling with optimality modelling at some level. The present approach is conservative in this respect, taking well-specified strategies and ensuring only that they perform as well as they possibly can in any given condition. The IOS

methodology, by contrast, is relatively radical, allowing the strategies themselves to be partially specified by the optimisation process. A consequence of this is that it will not normally be appropriate, in an IOS context, to assert a distinction between ‘strategy parameters’ and the strategies themselves. For the present model, by contrast, the issue of optimal strategy parameters (ensuring that each strategy does as well as it can) can indeed be distinguished from that of whether or not a strategy produces optimal behaviour (how well a strategy does in the space of possible strategies). In other words, in an IOS context, strategies may have some claim to be optimal with respect to a larger space of possible strategies, *in virtue of the use of GAs*. This is not so for the present model, in which GAs may elucidate the relative performances of different strategies, with respect to optimality criteria, but in which it does *not* follow from this that the strategies themselves are ‘optimal’ with respect to larger strategy spaces. The  $\omega$ -sampling strategy may indeed be more ‘optimal’ than  $\varepsilon$ -sampling in particular situations, and although GAs can help reveal to this fact, they do not thereby make it the case.

An additional difference is that a strict interpretation of the IOS methodology would require that genotypes be evaluated in isolation from each other, so that the evolutionary mechanics of the GA may be excluded from accounts of the outcome. The present model, by contrast, evaluated all genotypes in a population *simultaneously*, with interactions between genotypes potentially influencing fitness. Each agent in a group is represented by a separate genotype, so that the (aclonal) population of genotypes evolving under the GA, and the population of agents foraging for resources, were one and the same (and similar in this way to the prisoner’s dilemma models of chapter 3). Nevertheless, during analysis of the present model, clonal groups were reintroduced, for the reason that they enabled *average* near-optimal parameters (over many evolutionary runs) to be utilised. The results did not, however, depend on populations being clonal; identical results (not presented here) proceed from analyses of aclonal populations, each such population being randomly selected from those evolved in each experimental condition.

#### 9.4.2 Practical details

Here is how artificial evolution was actually deployed in the model. For both  $\varepsilon$ -sampling and  $\omega$ -sampling, populations of agents were evolved in each of 4 cases ( $m = 1.0$  or  $0.3$ , C-allocation or P-allocation): a total of 8 conditions. Each agent (in each condition) possessed a genome of 2 real numbers in the range  $[0.0, 1.0]$  specifying  $\varepsilon$  and  $\gamma$ . Initial populations (size 100) were randomly generated for each condition, and in each case a tournament-style GA applied for 100 generations (see appendix A). During each generation the population was evaluated 10 times, with the final fitness of each agent calculated as the average over all 10 evaluations.

Each evaluation began by randomly assigning values for  $F_A$  and  $F_B$  ( $F_A + F_B = 1.0$ , total resource  $F^* = 200.0$  in all conditions), and randomly allocating each agent to either  $A$  or  $B$ . The behaviour of each agent was then followed over 1000 time-steps, just as described in section 9.3.2. The fitness of each agent was determined on the basis of total accumulated resources at the end of the 1000 cycles.

The evolved near-optimal values for  $\varepsilon$  and  $\gamma$  are shown in figure 9.7, which illustrates average values over the 10 evolutionary runs in each condition. These are the values that were used to conduct the analyses in sections 9.3.3 and 9.3.6. That there is considerable variation between

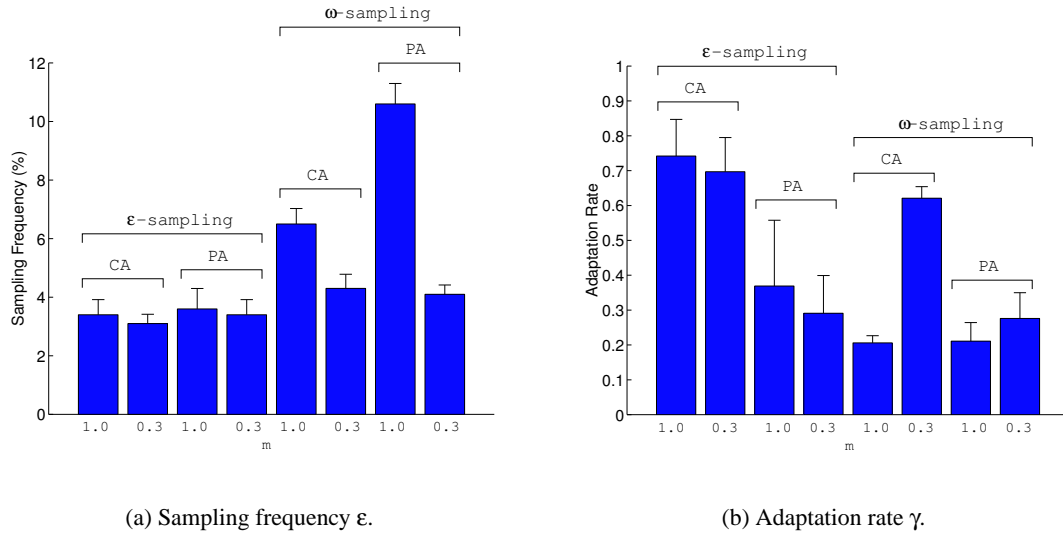


Figure 9.7: Near-optimal values for  $\varepsilon$  and  $\gamma$ . The 8 columns of each graph represent the 8 conditions, and each indicates the average near-optimal value (with standard deviation) over 10 evolutionary runs (CA represents C-allocation, PA, P-allocation).

conditions, and consistency within conditions, is clear. This is important because if all conditions demonstrated very similar, or entirely random near-optimal values, there would have been little justification by the above arguments for their use in the first place. Note that the specific evolved values of  $\varepsilon$  and  $\gamma$  are not of interest here: the idea is to use artificial evolution to *elucidate* strategy properties, not to *engender* them.

### 9.4.3 Near-optimal versus pre-specified parameters

As has been mentioned at least twice already, the primary contributions of this chapter do not depend on the use of artificial evolution. This can be confidently asserted because both  $\varepsilon$ -sampling and  $\omega$ -sampling were analysed using both near-optimal parameters and a pre-specified parameter set derived from the average near-optimal values across all conditions:  $\bar{\varepsilon} = 0.052$  and  $\bar{\gamma} = 0.427$ . Sensitivity to variation of these values was tested by repeating the analyses for 4 variations in  $\bar{\gamma}$ , keeping  $\varepsilon$  constant, (in step sizes of 0.1 above and below  $\bar{\gamma}$ ), and for 4 (small) variations in  $\varepsilon$ , keeping  $\gamma$  constant, (in step sizes of 0.01 above and below  $\bar{\varepsilon}$ ).

The detailed results from these analyses are presented in a series of tables in appendix C; the comments here will be restricted to a short summary. Most importantly, the use of the pre-specified set  $(\bar{\varepsilon}, \bar{\gamma})$  does indeed deliver the same overall pattern of results as obtained with near-optimal parameters. Furthermore, this observation is robust; it holds for all the variations described above. Interestingly, however, the near-optimal parameters do lead to a closer fit to the IFD than pre-specified parameters, even though the qualitative pattern is the same. Similarly, the near-optimal set leads both embedded and isolated  $\omega$ -sampling individuals to most closely approximate strict matching; again the difference is quantitative rather than qualitative, but a difference it remains. These observations strongly suggest that the near-optimal results are of greater clarity than those obtained with any fixed parameter combination, a clarity which serves to emphasise the value of using near-optimal parameters as equivalence criteria across experimental conditions.




Figure 9.8: Relation between noise and adaptation rate  $\gamma$ . Higher noise levels entail longer memories (as measured by  $\gamma$ ; low  $\gamma$  implies long memory). Each column represents the average (and standard deviation) of 10 evolutionary runs.

#### 9.4.4 Strategy memory and noise

One final analysis remains, in which - in contrast to the majority of this chapter - artificial evolution is absolutely essential. Notice, first of all, that there is considerable similarity between  $\omega$ -sampling and the ‘adaptive aspiration rate’ strategy used by Posch (1999) for playing the iterated prisoner’s dilemma. In particular, both share a ‘memory’ parameter (for  $\omega$ -sampling this is the ‘learning rate’,  $\gamma$ ). Chapter 3, in the early part of this dissertation, presented results demonstrating, in the context of the iterated prisoner’s dilemma, that increased environmental noise led to the evolution of longer ‘memories’ for an adaptive aspiration strategy. It therefore provides a satisfying coda to the empirical work of this dissertation to recapture these results in the context of the present model, through an analysis of the performance of  $\omega$ -sampling in the presence of environmental noise.

Noise in the present model - the equivalent of M-noise in chapters 3 and 4 - is implemented by a probability (on each time-step) that the resources obtained by each agent is calculated as if the agent were in the opposite patch to that specified by the strategy. Note that this is *not* equivalent to sampling; sampling visits inform the strategy about the value of other patches, M-noise misleads the strategy about the value of the current patch. Populations (size 100) were evolved under C-allocation with  $m = 1.0$ , exactly as described in section 9.4.2, in 7 conditions distinguished by the level of noise, ranging from 0% to 10%. In each case the adaptation rate  $\gamma$  of the fittest agent (after 100 generations) was recorded, and 10 entire evolutionary runs were performed for each condition enabling means and standard deviations to be calculated. Figure 9.8 illustrates that as the level of noise increases, so too does the average memory of the agents.

## 9.5 Summary

This chapter has investigated two patch-switching foraging strategies,  $\epsilon$ -sampling and  $\omega$ -sampling, on their ability to (a) lead groups of agents to distribute according to the predictions of the IFD, and (b) lead individual agents to adhere to the matching law, in each case under a number of non-trivial conditions. Contrary to the claims of Thuisjman et al. (1995),  $\epsilon$ -sampling proved inadequate at

both;  $\omega$ -sampling, by contrast, was very successful, even to the extent of entailing matching under conditions in which such behaviour is suboptimal. These results therefore support  $\omega$ -sampling as a candidate ‘mechanism of matching’, and also support the general hypothesis, first explored in the previous chapter, that matching behaviour can be understood as a consequence of optimal foraging in a shared environment. Also, with respect to the assessment of matching in the field, it was suggested that the matching behaviour of embedded agents to available resources does *not* provide a reliable guide to their matching behaviour when in isolation. These observations represent the primary contributions of this chapter.

The present model has entered the middle ground identified in the previous chapter, lying somewhere between IOS models of matching to available resources, and the complex and arcane matching experiments that populate the psychology literature. Although this is a rich territory, the many differences between the present model and the previous IOS model - along with the potentially problematic requirement of  $\omega$ -sampling for multiple value estimates - encourage the development of models at intermediate levels of description.

One obvious similarity that remains is that both models employ GAs to specify aspects of agent architecture. In the present model, the use of near-optimal parameters has been shown to deliver results of impressive clarity, and, more generally, to continue the reconciliation of optimality modelling and individual-based modelling exemplified by the IOS methodology. Also of interest is that the evolution of near-optimal parameters under noisy conditions recovers a result prominent in the early part of this dissertation; that environmental variability can promote the evolution of longer memories. Nevertheless, the role of GAs in the present model should not be overstated. Apart from these specific observations, its insights do *not* depend on their use.

This chapter has, in many ways, digressed from the central themes of this dissertation. Concessions to internalism are evident in the adoption of pre-specified internal mechanisms (the strategies), and matters of situated perception and action have been set to one side. The compensations are that the present model has provided a particularly close point of contact with contemporary theoretical biology, and an extremely clear demonstration of a situation in which behaviours which may be irrational when expressed by an isolated individual can be understood as rational in a group context. It remains, in the following and final chapter, to reassert the central themes of this project, and, with the benefit of hindsight, to arrange its various empirical contributions (including those of the present chapter) in their light.

# Chapter 10

## Conclusions

---

The art of being wise is knowing what to overlook.  
- William James

This dissertation has been about the use of artefacts to understand the relations between behaviour, mechanism, and environment, as they arise in a variety of agent-environment systems. Such a broad aim as this has required the project to span conceptual, methodological, and empirical levels of discourse. The time has now arrived to assess the contributions made in each of these domains, and the extent to which their combination offers a cohesive picture of adaptive behaviour. This final chapter will offer some general remarks before recapitulating the individual contributions of the dissertation.

The importance of distinguishing between behavioural and mechanistic levels of description has been emphasised from the outset, given a definition of behaviour as ‘observed ongoing agent-environment interactivity’, and of mechanism as ‘agent-side structure subserving such interactivity’. I argued that to ignore this distinction was to radically prejudge how a given behaviour might be generated, and often to significantly overestimate the degree of mechanistic complexity required. Many instances were exposed in which disrespect of this distinction had exactly this consequence.

An externalist conceptual framework, to be cashed out empirically in the form of artificial evolution models, was proposed as a way of interrogating the consequences of this distinction. These models were also presented as a means of evaluating the hypothesis (deriving from the overarching conceptual framework) that behavioural and/or mechanistic complexity could be understood in terms of adaptation to environmental variability. The many examples of such models throughout this dissertation provided strong support for this hypothesis, but also demonstrated - as suggested above - that complex behaviours need not be subserved by complex internal mechanisms. However, it was only ever possible to clearly elucidate these various relations in models that operated at the level of situated perception and action (not all the models of this dissertation operated at this level, but not all were addressed specifically to this issue). From chapter 5 onwards, considerable effort was dedicated to locating the implications of these insights in the context of current theoretical biology and experimental psychology.



These general conclusions are representative of the various contributions of this dissertation, but by no means capture them entirely. On the other hand, neither do the empirical contributions detailed in the next section come close to fulfilling the potential of the conceptual and methodological frameworks in which they reside. This, I hope, will be taken to attest to the richness of this potential and not to be a shortcoming of this dissertation.

### 10.1 Summary of contributions

The first task of the dissertation was to make clear the conceptual foundations of an externalist perspective on the relations between behaviour, mechanism, and environment. This took the form of an integration of two closely related themes; the relation of functional properties of behaviour to environmental structure, and the essential distinction between behavioural and mechanistic levels of description.

The terms ‘behaviour’, ‘mechanism’, and ‘environment’ are bartered with a disturbing recklessness in the current conceptual economy, and so much of the first part of chapter 2 was devoted to some historical clearing of the decks. Godfrey-Smith’s (1996) environmental complexity thesis (ECT) was then adopted as a suitable focus in virtue of its formalisation of the (pragmatic) externalist position that ‘cognition’ can be understood in terms of providing responses to environmental variability, heterogeneity, and/or complexity. In its standard form, however, this idea can be related only to the functional properties of behaviour and makes liberal use of the term ‘cognition’ to gloss over the distinction between behaviour and mechanism; it also denies the importance of constitutive construction (the translation of features of the external environment into features of the *Umwelt*). The second part of chapter 2 carefully exposed the origin and nature of these difficulties in the form of an extended critique of the ECT, the three elements of which focused on ‘cognitive’ interpretations of mechanism, the importance of constitutive construction, and the clear parallels between the ECT and Ashby’s law of requisite variety (LRV). This critique, bringing behaviour, mechanism, and environment together in a novel conceptual union, constituted the primary independent contribution of the chapter.

There is not the opportunity here to rehearse the critique of the ECT in detail, however one of its main concerns must be mentioned. This has to do with the interpretation of internal mechanism, on the one hand as underlying the behaviour which constitutes the adaptation to environmental complexity, and on the other as contributing to the mediation of the process of constitutive construction. The problem is that the first idea of mechanism *derives from* the ECT, in virtue of providing the responses to the agent’s *Umwelt* which underlie the behaviour in question, but the second idea - pertaining as it does to the generation of the *Umwelt* itself - can be interpreted both as following from the ECT, but also to some extent as being a *pre-condition* of its application. There is clearly some danger of circular argument here, a danger addressed by chapter 2 in the form of a call for the development of concrete models in which the relations between behaviour, mechanism, and environment can be empirically instantiated, and in which the issue of constitutive construction ceases to be a philosophical obstacle and can be cashed out in terms of operational details. In particular, artificial evolution models were recommended in view of their inherent externalism, and because they allow the relations between behaviour and mechanism to be the object of study rather than assumption.

The first such models of the dissertation appeared in chapters 3 and 4, and were concerned with exploring the hypothesis, directly inherited from the ECT, that environmental variability can promote the adaptive evolution of internal complexity. This investigation required its own episode of deck clearing at the outset of chapter 3; complexity and its evolution are issues well acquainted with controversy. It was decided, after some discussion, to hold back from endorsing any universal definition of complexity, and to rely instead on a series of model-specific interpretations consistent only with the general notion of contextually-specific, observer-related, heterogeneity and interconnectedness.<sup>1</sup> Chapter 3 also sharply distinguished between the ‘teleological’ and ‘causal’ questions with respect to the biological evolution of complexity. The teleological question, which asks whether or not (biological) evolution is a naturally progressive force tending towards greater complexity, was argued to be ill posed because it can be both trivially true and trivially false at the same time. The causal question, on the other hand, was deemed worthy of study: it asks *how* it is possible for evolution to produce structures of complexity (from structures of lesser complexity). One instance of this question of course recapitulates the adaptationist hypothesis with which the chapter opened, the question of whether environmental variability can promote the adaptive evolution of complexity.

Empirical support for this hypothesis was provided by a series of experiments introducing environmental noise into an evolutionary iterated prisoner’s dilemma (IPD) model, where complexity was measured by the memory length of the constituent strategies. Two versions of the model were compared, a ‘compulsory’ model and an ‘IPD/CR’ model, differing most significantly in the stability of cooperation normally observed in their behaviour. Although noise promoted the evolution of complexity in both models, it was found to be adaptively significant only in the (cooperatively stable) IPD/CR model; a distinction made by appeal to an interpretation of the dynamics of each model in terms of Ashby’s LRV. The observed complexity in the compulsory model was identified, after further analysis, with the facilitation of the exploration of the search space by the evolutionary algorithm.

Chapter 3 closed with a replication of a very different IPD model, due to Posch (1999), which utilised pre-specified strategies defined in part by an explicit memory parameter. Posch’s results were accurately reproduced, but were shown to be insufficient to support his (congenial) claim that noise promotes the evolution of long memories. After exposing the flaw in Posch’s reasoning, his model was extended in various ways such that his original conclusions were reclaimed on a sound empirical basis.

The evolution of complexity remained the focus of attention in chapter 4, however the empirical context of the IPD was replaced with that of the evolution of homing-navigation behaviour for a mobile robot (Floreano & Mondada, 1996). It was found that evolution in high-noise environments augmented both the complexity of the behaviour and the complexity of the underlying mechanistic dynamics. Behavioural complexity was revealed by adaptive flexibility in the face of environmental perturbation, mechanistic complexity by the integration of sensory data across a number of modalities in the determination of motor output (neural network controllers evolved in low-noise conditions relied almost exclusively on data from the infra-red proximity sensors). These results were discussed in terms of the (critiqued) ECT, with the distinction between be-

---

<sup>1</sup>The reader may recognise, once again, the influence of William James here.

haviour and mechanism and the importance of understanding the structure of the *Umwelt* receiving particular attention. One important observation was that junctures between distinct behaviours - from the perspective of an external observer - were *not* subserved at a mechanistic level by 'internal behavioural correlates' or 'supervening arbitration devices'.

Chapter 4 concluded with a pragmatic (and relatively speculative) discussion of three novel roles of noise in artificial evolution: (1) the facilitation of search space exploration (as identified in chapter 3), (2) the *acceleration* of the search process itself (identified in the homing navigation model), and (3) the promotion of the adaptive evolution of complexity (as hypothesised at the outset). The origins and interrelation of all three phenomena were discussed in terms of the dynamics of evolutionary search.<sup>2</sup>

Chapter 5 marked a change in the direction of the dissertation, laying the methodological foundations required to address issues current in theoretical biology and experimental psychology. This methodological framework, referred to as the 'individual-based optimal situated' (IOS) modelling strategy, was billed as an unorthodox complement to optimal foraging theory (OFT), and comprises of a generalisation of the modelling strategy expressed in the preceding chapter; the artificial evolution of control structures for agents engaged in agent-environment interaction mediated by situated perception and action.

Orthodox OFT is an attempt to understand both functional and mechanistic aspects of behaviour by treating observed behaviour as optimally adapted to an environment (it is *not* an attempt to discover whether organisms 'optimise'). IOS models maintain these broad aims, but are well placed to interrogate a number of 'framework assumptions' intrinsic to orthodox OFT models. In particular, IOS models escape the need for orthodox 'decision variables' which frequently trample over the behaviour/mechanism distinction, reifying junctures in behavioural dynamics that may only be significant from the perspective of an external observer. There are a number of other advantages: IOS models naturally encompass behavioural historical constraints (dynamical invariants in patterns of agent-environment interaction), the dynamics of the optimisation process itself, and they inherently capture the distinction between the external environment and *Umwelt*. Accompanying these advantages is the inevitable cost of situated individual-based modelling; it can be hard to understand why IOS models do what they do, and difficult to know what conclusions can be drawn from their behaviour.

These issues, and others, were discussed in chapter 5 as part of an extended analogy between the IOS methodology and orthodox OFT. The IOS methodology was also set in the context of related work in theoretical biology, SAB/AL itself, and the general debate over the nature of possible methodological relationships between SAB/AL and theoretical biology. The conclusion offered was that IOS models are best construed as a form of computational thought experiment (following Di Paolo et al., 2000) with the capacity to reveal the impact of implicit and explicit assumptions unavoidable in orthodox OFT. This enhancement of the methodological foundations of SAB/AL was therefore identified as the main independent contribution of the chapter.

Despite this change in direction, in many ways chapter 5 remained faithful to the conceptual infrastructure laid down in chapter 2. Both IOS models and orthodox OFT models are strongly

---

<sup>2</sup>The *general* idea that noise - or variability - can promote the evolution of complexity is of course *not* novel; Godfrey Smith's ECT, for example, expresses a version of it. However I am referring here to the specific incarnations of this idea in the context of artificial evolution.

externalist, and many of the benefits offered by IOS models have to do with making explicit, and elucidating the consequences of, the distinction between behaviour and mechanism. A useful interpretation of this chapter is as a reformulation of the call for the development of artificial evolution models voiced at the end of chapter 2; a methodological expression of the (critiqued) ECT.

Chapter 6 followed this methodological discourse with a discussion of the theoretical context of behaviour coordination; the problem, at its most general, of ‘how to do the right thing’. A survey of the relevant literature revealed some deep - yet extremely widespread - conceptual difficulties, most notably the preponderance of internal behavioural correlates and the accompanying ontological divorce of perception and action. To illustrate how these difficulties might be circumvented, an IOS model was described (inspired by Braitenberg, 1984, and Lambrinos & Scheier, 1995) in which artificial evolution is used to specify the shapes of the direct, independent, and continuously active sensorimotor links controlling simple agents engaged in simple foraging behaviour, in environments containing two types of resource and dangerous ‘traps’. The point of this extremely minimal (and novel) architecture was to make it abundantly clear that the behaviour of the agents could *not* rely on internal behavioural correlates or ‘decision variables’ of any kind, and, to the contrary, that the underlying mechanism enshrined the intimacy of perception and action. Evolved agents nevertheless engaged in effective behaviour coordination - or ‘action selection’ - when assessed on a battery of tests drawn from the SAB/AL literature, an assessment complemented by a detailed analysis of the evolved mechanism itself.

Although criticisms could be levelled at the simplicity of the structure of the model, its purpose was neither to be of practical utility, nor to pretend to biological fidelity. Rather, the objective was to provide a proof of concept that effective behaviour coordination need not rely on those relations between behaviour, mechanism, and environment commonly supposed. This constituted the primary contribution of chapter 6, a contribution consistent with the overarching externalism of the dissertation insofar as the sources of organisation in the model can be located in the agent-environment interactions themselves, and not - as is sometimes assumed in models of behaviour coordination - in hierarchical structures and/or arbitration mechanisms internal to the agent.

Chapters 7, 8, and 9 presented the investigations of the dissertation most closely related to current theoretical biology and experimental psychology. Chapter 7 explored the ‘interference function’, the relation in theoretical biology between predator density and per-predator intake of prey. A survey of the literature revealed a strong and entirely appropriate desire for interference models consistent with the principles of OFT, however the most significant extant developments in this direction suffer from a now familiar conceptual shortcoming: the conflation of behaviour and mechanism. The example of Stillman et al. (1997) was used, in which strikingly complex internal arbitration mechanisms are assumed to supervene on predefined behavioural entities, and indeed in which *interference itself* is explicitly associated with *a priori* components of the model, in this way largely prefiguring the conclusions that it can support.

As a response to this unsatisfactory state of affairs, an IOS model was described - a multi-agent extension of the preceding behaviour coordination model - in which interference is *not* preregistered into the model structure, but instead arises from agent-environment interaction patterns. In this way the call for interference models interpretable in terms of optimality was answered *without*

requiring that internal structure be framed in terms of internal behavioural correlates and complex arbitration mechanisms; indeed, from an orthodox point of view, the agent-side mechanisms of this model are surprisingly simple. Furthermore, by virtue of its operation at the level of situated perception and action, the model revealed ways in which behavioural historical constraints can influence interference situations, influences which present some challenges to the significance awarded in theoretical biology to ‘ideal’ interference functions (linear log-log relationships between agent density and intake rate).

From a biological perspective, the significance of this model is not to do with overturning any particular predictions or assumptions of orthodox models, rather that an entirely new way of modelling interference has been shown to be feasible, conceptually congenial, and capable of elucidating previously opaque phenomena of importance in understanding the dynamics of interference situations. Moreover, interference is not restricted to the biological sphere, but is of significance in any situation in which multiple agents compete for limited resources; a further contribution of chapter 7 can therefore be identified in its introduction of a useful biological concept into the empirical discourse of SAB/AL.

As part of all this, and of most immediate relevance to interference in an artificial agent context, chapter 7 also analysed in detail a behaviour, suboptimal for isolated individuals, yet optimal for individuals in groups. This behaviour, slowing down following consumption, promoted the avoidance of ‘clumping together’ of agents which, in a group context, can reduce foraging efficiency. Agents were evolved that were able to switch between this strategy and a simple strategy of continuous high speed, according to prevailing environmental conditions.

Chapter 8 continued the investigation of suboptimal behaviour, using an extension of the interference model to explore the hypothesis that (potentially suboptimal) individual adherence to Herrnstein’s ‘matching law’ can be a consequence of optimal foraging in a shared environment, a hypothesis derived from a series of analogies between the matching law and the ideal free distribution (IFD). The IOS model described in this chapter supported this hypothesis, and in doing so challenged the idea that matching behaviour must be subserved by a dedicated ‘mechanism of matching’ (an idea which, of course, has its origins in a confusion of behaviour and mechanism). Moreover, the idea of matching itself, as an isolable pattern of behaviour, was brought into question. The model described a kind of matching originating in the same simple foraging interactions constitutive of interference, without distinct environmental patches, internal adaptive state, or artificial separation of response from reward; a kind of matching for these reasons very different from that found in psychology textbooks, and suggestive of a *continuum* of matching phenomena, not all of which is well described by the original matching law of Herrnstein.

Perhaps more than any other, the model of chapter 8 exemplified the benefits of IOS modelling. It provided support for a functional hypothesis, accounting for the potential irrationality of matching in a way consistent with the premises of OFT. It bestowed mechanistic insight, and beyond all this it engendered a reconceptualisation of the relations between behaviour, mechanism, and environment, associated with matching; the understanding of matching itself, as a phenomenon deserving of explanation, was significantly extended. Furthermore, the operational details of the model required interpretation in terms of behavioural historical constraints; in particular, matching was not to be understood as a ‘solution’ to the ‘problem’ of interference. The model also exem-

plified the IOS methodology in its location at the intersection of theoretical biology, experimental psychology, and SAB/AL. Not only did it provide a bridge from biology to psychology, but in doing so it gave new depth to the understanding of choice behaviour in SAB/AL, an understanding too often confined within the framework of action selection.

The benefits of IOS modelling are also its limitations. In reshaping the identity of matching, the empirical results afforded by the model of chapter 8 were unable to smoothly engage with the orthodox literature. Chapter 9, by contrast, returned to a largely orthodox terrain and *directly* explored the hypothesis that adherence to the matching law (as it is understood in psychology) can be understood in terms of leading groups of foragers to the IFD. The claim of Thuisjman et al. (1995), that the ‘ $\epsilon$ -sampling’ foraging strategy could underlie both the IFD and the matching law, was assessed and found to be false. A related (and novel) strategy,  $\omega$ -sampling, was introduced and found to succeed where  $\epsilon$ -sampling failed. Both strategies specified simple patch switching rules and differed primarily in their capacity to maintain representations of the environment;  $\epsilon$ -sampling only maintains a single value estimate,  $\omega$ -sampling maintains multiple estimates. These detailed analyses constituted the primary independent contributions of chapter 9, contributions which - in contrast to those of earlier chapters - integrate effortlessly with the context of theoretical biology.

Even though the model of chapter 9 represented a departure from the conceptual and methodological framework of the bulk of the dissertation, many themes persisted. The model provided additional support for the idea that the potential suboptimality of matching can be understood in terms of adaptation to group situations. It continued the use of GAs in the reconciliation of optimality and individual-based approaches to modelling, although in this case the GAs were employed only to facilitate comparison between predefined strategies, rather than to specify the strategies themselves. Indeed, the only result of the model that *did* depend entirely on the use of GAs was the recapitulation of the conclusions of chapter 3, that environmental variability can promote the evolution of strategy complexity, when measured in terms of ‘memory’.

Perhaps most importantly, the model - being orthodox in constitution - made clear the relationship between orthodox and IOS approaches to phenomena such as matching and the IFD. The model of chapter 9 did *not* reconstrue the relations between behaviour, mechanism, and environment, and, as a consequence of this, matching *was* treated as a potential solution to the problem of interference, with respect to the goal of the IFD. This afforded the opportunity, taken towards the end of the chapter, to discuss in greater detail the constitution of the continuum of matching identified in the preceding chapter, and to make some suggestions for other regions of this continuum deserving of exploration.

## 10.2 Perspective

The contributions described above represent a small part of what promises to be a very long story, the outlines of which may be discerned in the conceptual and methodological foundations set down in chapters 2 and 5. However, it is not my intention to rehearse here the many ways in which these beginnings may unfold in the future. Each of the empirical sections of this dissertation has its own momentum, a momentum in each case best appreciated from within its proper empirical context. These final remarks instead concern the general balance of the dissertation.

If there is one area of the conceptual framework of chapter 2 that could benefit most from a

tighter empirical focus in the future, it is the relationship between the external environment and the *Umwelt*. As an example of one way to proceed, it would be interesting to explore models in which the process of constitutive construction could be modified in various ways, with the corresponding relations between behaviour and mechanism subject to analysis. Biró and Ziemke (1998), for instance, make a start in this direction by trying to reconcile cluster analyses of mechanistic activity with behavioural level decompositions, in the context of simple visual approach behaviour for a situated agent. They claim to achieve this, but the suspicion remains that their apparent success may rest on their use of ‘magic’ sensors by which behaviourally relevant features of the external environment are transported wholesale into the input space of the controlling neural network.

More generally, many of the investigations in this dissertation have focussed on mechanisms that do not possess the capacity to *learn* (the exceptions being game-theoretic models in chapters 3 and 9). Although motivated largely for reasons of simplicity, the frequent exclusion of learning may be considered to be a general limitation of this research, since many interesting phenomena concerning the relations between behaviour, mechanism, and environment are framed in terms of lifetime adaptation (see, for example, Connell & Mahadevan, 1993; Barker, 1994; Sutton & Barto, 1998).<sup>3</sup> This is of course particularly so for psychological phenomena, and although the rule-based matching model of chapter 9 did consider learning agents, the IOS matching model of chapter 8 did not. It is likely that further research intermediate between these models would need to consider learning agents in an IOS context.

Lifetime adaptation is not only associated with learning, it is of course also associated with ontogeny, or development. And it is here that a deep challenge may be posed to the present project. Whereas theories of learning are often explicitly externalist (recall the associationism of nineteenth century psychology from chapter 2), theories of ontogeny generally concern the unfolding of internal dynamics. They are perhaps guided to some extent by environmental constraints (see Waddington, 1966; Piaget, 1971, and the collected papers in Johnson, 1993), but nonetheless remain relatively internalist in character. It is no accident that, out of Tinbergen’s celebrated four questions that may be asked of features of animal behaviour or morphology, it is only the question of ontogeny that has been entirely shirked in this dissertation (the other three concern function, mechanism, and phylogeny, see Tinbergen, 1963). The integration of ontogeny with the picture of adaptive behaviour offered by this dissertation is therefore a challenge that must be met at many levels: conceptual, methodological, and empirical.

---

<sup>3</sup>Nor can learning be considered the exclusive preserve of relatively ‘advanced’ organisms. Sensitisation and habituation can be observed in the simplest of creatures (for example *Aplysia californica*, Kandel & Schwartz, 1982), and it has even been announced - as I am framing these final paragraphs - that the slime mold *Physarum polycephalum* “has the ability to find the minimum-length solution between two points in a labyrinth” (Nakagaki, Yamada, & Toth, 2000, p.470).

## Bibliography

- Adami, C., & Cerf, N. (2000). Physical complexity of symbolic sequences. *Physica D*, *137*, 62–69.
- Adami, C., Ofria, C., & Collier, T. (2000). Evolution of biological complexity. *Proceedings of the National Academy of Science (USA)*, *97*, 4463.
- Agre, P., & Chapman, D. (1987). Pengi; an implementation of a theory of situated action. In *Proceedings of the Sixth National Conference on Artificial Intelligence*, pp. 268–272 Los Altos, CA. AAAI/Morgan Kaufmann.
- Antonovics, J., Ellstrand, N., & Brandon, R. (1988). Genetic variation and environmental variation: Expectations and experiments. In Gottlieb, L., & Jain, S. (Eds.), *Plant evolutionary biology*, pp. 275–303. Chapman Hall, London.
- Arkin, R. (1998). *Behavior-based robotics*. MIT Press, Cambridge, MA.
- Arnold, A., Kelly, D., & Parker, W. (1995). Causality and Cope's rule: Evidence from the planktonic *Foraminifera*. *Journal of Paleontology*, *69*, 203–10.
- Ashby, W. (1952). *Design for a brain: The origin of adaptive behaviour*. Chapman Hall, London.
- Ashby, W. (1956). *An introduction to cybernetics*. Chapman Hall, London.
- Ashby, W. (1962). Principles of the self-organizing system. In von Foerster, H., & Zopf, G. (Eds.), *Principles of self-organization*, pp. 255–278. Pergamon Press, New York.
- Ashlock, D., Smucker, M., Stanley, E., & Tesfatsion, L. (1995). Preferential partner selection in an evolutionary study of prisoner's dilemma. *Biosystems*, *1-2(37)*, 99–125.
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books, New York.
- Baerends, G. (1976). The functional organisation of behaviour. *Animal Behaviour*, *24*, 726–738.
- Baldwin, J. (1896). A new factor in evolution. *American Naturalist*, *30*, 441–451.
- Ballard, D. (1991). Animate vision. *Artificial Intelligence*, *48*, 57–86.
- Barker, L. (1994). *Learning and behavior: A psychological perspective*. Macmillan College Publishing, New York.
- Barkow, J., Cosmides, L., & Tooby, J. (Eds.). (1992). *The adapted mind: Evolutionary psychology and the generation of culture*. Oxford University Press, Oxford.
- Bateson, G. (1963). The role of somatic change in evolution. *Evolution*, *17*, 529–539.
- Baum, W. (1974). On two types of deviation from the matching law: Bias and undermatching. *Journal of the Experimental Analysis of Behavior*, *22*, 231–242.
- Baum, W., & Kraft, J. (1998). Group choice: Competition, travel, and the ideal free distribution. *Journal of the Experimental Analysis of Behavior*, *69(3)*, 227–245.
- Bechtel, W., & Abrahamsen, A. (1991). *Connectionism and the mind*. Blackwell, Oxford.



- Bedau, M. (1999). Can unrealistic computer models illuminate theoretical biology?. In Wu, A. (Ed.), *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program*, pp. 20–23 San Francisco. Morgan Kaufmann.
- Beer, R. (1995). Computational and dynamical languages for autonomous agents. In Port, R., & van Gelder, T. (Eds.), *Mind as motion: Explorations in the dynamics of cognition*, pp. 121–149. MIT Press, Cambridge, MA.
- Beightler, C., Philips, D., & Wilde, D. (1979). *Foundations of optimisation*. Prentice-Hall, London.
- Belew, R., & Mitchell, M. (Eds.). (1996). *Adaptive individuals in evolving populations: Models and algorithms*, Vol. XXVI of *Santa Fe Institute studies in the sciences of complexity*, Reading, MA. Addison-Wesley.
- Berkeley, G. (1690). Principles of human knowledge. (D.M. Armstrong, ed.) Macmillan, 1965.
- Bernstein, C., Kacelnik, A., & Krebs, J. (1988). Individual decisions and the distribution of predators in a patchy environment. *Journal of Animal Ecology*, 57, 1007–1026.
- Bernstein, C., Kacelnik, A., & Krebs, J. (1991). Individual decisions and the distribution of predators in a patchy environment II: The influence of travel costs and structure of the environment. *Journal of Animal Ecology*, 60, 205–225.
- Biró, Z., & Ziemke, T. (1998). Evolution of visually-guided approach behaviour in recurrent artificial neural network robot controllers. In Pfeifer, R., Blumberg, B., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior*, pp. 73–77 Cambridge, MA. MIT Press.
- Blumberg, B. (1994). Action selection in hamsterdam: Lessons from ethology. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 107–116 Cambridge, MA. MIT Press.
- Blumberg, B., Todd, P., & Maes, P. (1996). No bad dogs: Ethological lessons for learning in hamsterdam. In Maes, P., Mataric, M., Meyer, J., Pollack, J., & Wilson, S. (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pp. 295–304 Cambridge, MA. MIT Press.
- Boerlijst, M., & Hogeweg, P. (1991). Spiral wave structure in prebiotic evolution. *Physica D*, 48, 17–28.
- Bonner, J. (1988). *The Evolution of complexity by means of natural selection*. Princeton University Press, Princeton.
- Boring, E. (1950). *A history of experimental psychology*. Appleton, New York. 2nd edition.
- Boyd, R., & Richerson, P. (1985). *Culture and the evolutionary process*. The University of Chicago Press, Chicago.
- Braitenberg, V. (1984). *Vehicles: Experiments in synthetic psychology*. MIT Press, Cambridge, MA.
- Brandon, R. (1990). *Adaptation and environment*. Princeton University Press, Princeton.
- Breland, K., & Breland, M. (1961). The misbehavior of organisms. *American Psychologist*, 16, 681.

- Brewster, D. (1834). *Letters on natural magic*. John Murray, London.
- Broca, P. (1869). Sur le siège de la faculté du langage articulé. *La Tribune Médicale*, 74, 245–256.
- Brooks, R. (1999). *Cambrian intelligence: The early history of new AI*. MIT Press, Cambridge, MA.
- Brooks, R. (1986). A robust layered control system for a mobile robot. *IEEE Journal of Robotics and Automation*, RA-2, 14–23.
- Brooks, R. (1989). A robot that walks: Emergent behaviors from a carefully evolved network. *Neural Computation*, 1, 153–162.
- Brooks, R. (1994). Coherent behavior from many adaptive processes. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 421–430 Cambridge, MA.
- Brooks, R., & Stein, L. (1993). Building brains for bodies. Tech. rep. A.I. Memo no. 1439, MIT AI Lab.
- Brooks, R., & Richard, R. (1998). *Multi-sensor fusion: Fundamentals and applications with software*. Prentice Hall, London.
- Brunswik, E. (1952). *The conceptual framework of psychology*. University of Chicago Press, Chicago.
- Bullock, S. (1999). Jumping to bold conclusions: A review of Amotz and Avishag Zahavi's *The handicap principle: A missing piece of Darwin's puzzle*. *Adaptive Behavior*, 7(1), 129–134.
- Bullock, S. (1998). *Evolutionary simulation models: On their character, and application to problems concerning the evolution of natural signalling systems*. Ph.D. thesis, University of Sussex.
- Burian, R., & Richardson, R. (1996). Form and order in evolutionary biology. In Boden, M. (Ed.), *The philosophy of artificial life*, pp. 146–173. Oxford University Press, Oxford.
- Cannon, W. (1932). *The wisdom of the body*. Norton, New York.
- Carlson, N. (1994). *Physiology of behavior*. Allyn and Bacon, Boston. 5th edition.
- Chaitin, G. (1966). A theory of program size formally equivalent to information theory. *Journal of the Association of Computing Machinery*, 13, 547.
- Chapuis, A., & Droz, E. (1958). *Automata, a historical and technological study*. B.T. Batsford Ltd., London. Translated by A. Reid.
- Charnov, E. (1976). Optimal foraging: The marginal value theorem. *Theoretical Population Biology*, 9, 129–136.
- Charnov, E., & Orians, G. (1973). Optimal foraging: Some theoretical explorations. Unpublished manuscript.
- Cheverton, J., Kacelnik, A., & Krebs, J. (1985). Optimal foraging: Constraints and currencies. In Holldobler, B., & Lindauer, M. (Eds.), *Experimental behavioral ecology*, pp. 109–126. Fischer-Verlag, Berlin.
- Chomsky, N. (1959). A review of Skinner's *Verbal behavior*. *Language*, 35, 26–58.

- Clancey, W. (1991). The frame of reference problem in the design of intelligent machines. In van Lehn, K. (Ed.), *Architectures for intelligence: Proceedings of the 22nd Carnegie-Mellon Symposium on Cognition*, pp. 357–423 Hillsdale, NJ. Erlbaum.
- Clark, A. (1989). *Microcognition: Philosophy, cognitive science, and parallel distributed processing*. MIT Press, Cambridge, MA.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. MIT Press, Cambridge, MA.
- Clutton-Brook, T., & Harvey, P. (1977). Primate ecology and social organization. *Journal of Zoology*, 183, 1–39.
- Colman, A. (1995). *Game theory and its applications in the social and biological sciences*. Butterworth-Heinemann, Oxford. 2nd edition.
- Condillac, E. (1754). *Condillac's treatise on the sensations*. Faval. trans. G. Carr, London, 1930.
- Conger, R., & Killeen, P. (1974). Use of concurrent operants in small group research. *Pacific Sociological Review*, 17, 399–416.
- Connell, J., & Mahadevan, S. (Eds.). (1993). *Robot learning*. Kluwer Academic Publishers, Boston.
- Cosmides, L., & Tooby, J. (1987). From evolution to behavior: Evolutionary psychology as the missing link. In Dupré, J. (Ed.), *The latest on the best: Essays on evolution and optimality*, pp. 277–306. MIT Press, Cambridge, MA.
- Crutchfield, J. (1994a). The calculi of emergence: Computation, dynamics, and induction. *Physica D*, 75, 11–54.
- Crutchfield, J. (1994b). Is anything ever new? Considering emergence. In Cowan, G., Pines, D., & Melzner, D. (Eds.), *Integrative themes*, Vol. XIX of *Santa Fe Institute studies in the sciences of complexity*, pp. 479–497. Addison-Wesley, Redwood City.
- Cummins, R. (1975). Functional analysis. *Journal of Philosophy*, 72, 741–764.
- Cuvier, G. (1827). *Essay on the theory of Earth*. William Blackwood, Edinburgh. 3rd edition, translated from French. Additional material by R. Jameson.
- Dallery, J., & Baum, W. (1991). The functional equivalence of operant behavior and foraging. *Animal Learning and Behavior*, 19(2), 146–152.
- Darwin, C. (1859). *On the origin of species by means of natural selection: Or the preservation of favoured races in the struggle for life*. Reprinted with introduction by E. Mayr. Cambridge, MA. Harvard University Press.
- Davison, M., & McCarthy, D. (1988). *The matching law*. Erlbaum, Hillsdale, NJ.
- Dawkins, M. (1995). *Unravelling animal behaviour*. Longman, London. 2nd edition.
- Dawkins, R. (1976). Hierarchical organisation: A candidate principle for ethology. In Bateson, P., & Hinde, R. (Eds.), *Growing points in ethology*, pp. 7–54. Cambridge University Press, Cambridge.
- Dawkins, R. (1982). *The extended phenotype*. Oxford University Press, Oxford.
- de Vaucanson, J. (1742). *An account of the mechanism of an automaton or image playing on the German-flute*. London. trans. J.T. Desaguliers.

- Deneubourg, J.-L., Goss, S., Franks, N., Sendova-Franks, A., Detrain, C., & Chretien, L. (1990). The dynamics of collective sorting: Robot-like ant and ant-like robots. In Meyer, J., & Wilson, S. (Eds.), *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pp. 356–363 Cambridge, MA. MIT Press.
- Dennett, D. (1978). *Brainstorms*. MIT Press, Cambridge, MA.
- Dennett, D. (1987). *The intentional stance*. MIT Press, Cambridge, MA.
- Dennett, D. (1991). *Consciousness explained*. Little, Brown, Boston.
- Dennett, D. (1995). *Darwin's dangerous idea: Evolution and the meanings of life*. Penguin Books, London.
- Dennett, D. (1987). Cognitive wheels: The frame problem of AI. In Hookaway, C. (Ed.), *Minds, machines, and evolution: Philosophical Studies*, pp. 129–151. Cambridge University Press, Cambridge.
- Dewey, J. (1896). The reflex arc concept in psychology. *Psychological Review*, 3, 357–370. Reprinted in: J.J. McDermott (1981), ed., *The philosophy of John Dewey*. University of Chicago Press.
- Dewey, J. (1929). *Experience and nature*. Dover, New York. revised edition.
- Di Paolo, E. (1999). *On the evolutionary and behavioural dynamics of social coordination: Models and theoretical aspects*. Ph.D. thesis, University of Sussex.
- Di Paolo, E., Noble, J., & Bullock, S. (2000). Simulation models as opaque thought experiments. In Bedau, M., McCaskill, J., Packard, N., & Rasmussen, S. (Eds.), *Artificial Life VII: Proceedings of the Seventh International Conference on the Simulation and Synthesis of Living Systems*, pp. 497–506 Cambridge, MA. MIT Press.
- Diamond, J. (1996). Regulatory biology - competition for brain space. *Nature*, 382, 756–757.
- Dickinson, A. (1980). *Contemporary animal learning theory*. Cambridge University Press, Cambridge.
- Dreyfus, H. (1972). *What computers can't do: A critique of artificial reason*. Harper and Row, New York.
- Dreyfus, H. (1991). *Being-in-the-world: A commentary on Heidegger's Being and time*, division 1. MIT Press, Cambridge, MA.
- Dugatkin, L., Mesterton-Gibbs, M., & Houston, A. (1992). Beyond the prisoner's dilemma: Toward models to discriminate amongst mechanisms of cooperation in nature. *Trends in Ecology and Evolution*, 7, 202–205.
- Durrett, R., & Levin, S. (1994). The importance of being discrete (and spatial). *Theoretical Population Biology*, 46(3), 363–394.
- Edelman, G. (1987). *Neural Darwinism*. Basic Books, New York.
- Eigen, M., & Schuster, P. (1977). The hypercycle. A principle of natural self-organization. Part A: Emergence of the hypercycle. *Naturwissenschaften*, 64, 541–565.
- Eigen, M., & Schuster, P. (1978a). The hypercycle. A principle of natural self-organization. Part B: The abstract hypercycle. *Naturwissenschaften*, 65, 7–41.

- Eigen, M., & Schuster, P. (1978b). The hypercycle. A principle of natural self-organization. Part C: The realistic hypercycle. *Naturwissenschaften*, 65, 341–369.
- Eigen, M., & Schuster, P. (1982). Stages of emerging life - five principles of early organization. *Journal of Molecular Evolution*, 19, 47–61.
- Emery, F. (1967). The next thirty years: Concepts, methods, and anticipations. *Human Relations*, 20, 199–237.
- Fagen, R. (1987). A generalized habitat matching law. *Evolutionary Ecology*, 1, 5–10.
- Faith, J. (2000). *Emergent representations*. Ph.D. thesis, University of Sussex.
- Feldman, D., & Crutchfield, J. (1998). Measures of statistical complexity: Why?. *Physics Letters A*, 238, 244–255.
- Ferrell, C. (1994). Robust adaptive locomotion of an autonomous hexapod. In Gaussier, P., & Nicoud, J. (Eds.), *Proceedings: From Perception to Action Conference*, pp. 66–77 Los Alamitos, CA. IEEE Computer Society Press.
- Ferrier, D. (1876). *Functions of the brain*. Smith, Elder, London.
- Feyerabend, P. (1988). *Against method*. Verso, London. Third edition.
- Fletcher, J., Zwick, M., & Bedau, M. (1996). Dependence of adaptability on environmental structure. *Adaptive Behaviour*, 4(3-4), 275–307.
- Fletcher, J., Zwick, M., & Bedau, M. (1998). Effect of environmental texture on evolutionary adaptation. In Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.), *Proceedings of the Sixth International Conference on Artificial Life*, pp. 189–199 Cambridge, MA. MIT Press.
- Floreano, D. (1993). Patterns of interaction in shared environments. In *Pre-Proceedings of the Second European Conference on Artificial Life*, Vol. II Brussels.
- Floreano, D., & Mondada, F. (1996). Evolution of homing navigation in a real mobile robot. *IEEE transactions on systems, man, and cybernetics: Part B; cybernetics*, 26(3), 396–407.
- Fodor, J. (1974). Special sciences, or the disunity of science as a working hypothesis. *Synthese*, 28, 97–115.
- Foner, L., & Maes, P. (1994). Paying attention to what's important: Using focus of attention to improve unsupervised learning. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior* Cambridge, MA. MIT Press.
- Forst, C. (1997). Molecular evolution of catalysis. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 83–92 Cambridge, MA. MIT Press.
- Franceschini, N., Pichon, J., & Blanes, C. (1992). From insect vision to robot vision. *Philosophical Transactions of the Royal Society of London: Series B*, 337, 283–294.
- Free, C., Beddington, J., & Lawton, J. (1977). On the inadequacy of simple models of natural interference for parasitism and predation. *Journal of Animal Ecology*, 46, 543–554.
- Fretwell, S. (1972). *Populations in seasonal environments*. Princeton University Press, Princeton.
- Fretwell, S., & Lucas, H. (1970). On the territorial behaviour and other factors influencing habitat distribution in birds. *Acta Biotheoretica*, 19, 16–36.

- Fritsch, G., & Hitzig, E. (1870). On the electrical excitability of the cerebrum. trans. G. von Bonin, in: *Some papers on the cerebral cortex*, pp. 73-96, Thomas-Springfield, (1960).
- Fryer, D., & Marshall, J. (1979). The motives of Jacques de Vaucanson. *Technology and Culture*, 20, 257–269.
- Gallistel, C., Brown, A., Carey, S., Gelman, R., & Keil, F. (1991). Lessons from animal learning for the study of cognitive development. In Carey, S., & Gelman, R. (Eds.), *The epigenesis of mind*, pp. 3–37. Lawrence Erlbaum, Hillsdale, NJ.
- Gibson, J. (1966). *The senses considered as perceptual systems*. Houghton-Mifflin, Boston.
- Gibson, J. (1979). *The ecological approach to visual perception*. Houghton-Mifflin, Boston.
- Gigerenzer, G., Todd, P., & the ABC research group (1999). *Simple heuristics that make us smart*. Oxford University Press, New York.
- Giske, J., Huse, G., & Fiksen, O. (1998). Modelling spatial dynamics of fish. *Reviews in Fish Biology and Fisheries*, 8, 57–91.
- Godfrey-Smith, P. (1996). *Complexity and the function of mind in nature*. Cambridge University Press, Cambridge.
- Gold, T. (1992). The deep, hot biosphere. *Proceedings of the National Academy of Sciences (USA)*, 89, 6045–49.
- Goldberg, D., & Matarić, M. (1997). Interference as a tool for designing and evaluating multi-robot controllers. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pp. 637–642 Providence, Rhode Island. AAAI Press.
- Goodwin, B. (1994). *How the leopard changed its spots: The evolution of complexity*. Phoenix, London.
- Goss-Custard, J. (1977). Optimal foraging and size selection of worms by redshank *Tringa totanus* in the field. *Animal Behaviour*, 25, 10–29.
- Goss-Custard, J. (1980). Competition for food and interference amongst waders. *Ardea*, 68, 31–52.
- Gould, S. (1989). *Wonderful life*. Penguin, London.
- Gould, S. (1997). *Life's grandeur*. Vintage, London.
- Gould, S., & Lewontin, R. (1979). The spandrels of San Marco and the Panglossian paradigm. *Proceedings of the Royal Society of London: Series B.*, 205, 581–598.
- Grantham, T., & Nichols, S. (1999). Evolutionary psychology: Ultimate explanations and Panglossian predictions. In Hardcastle, V. (Ed.), *Where biology meets psychology: Philosophical essays*, pp. 47–67. MIT Press, Cambridge, MA.
- Grassé, P. (1959). La reconstruction du nid et les coordinations interindividuelles chez *Bellitermes natalensis* et *Cubitermis* sp. La théorie de la stigmergie: Essai d'interprétation du comportement des termites constructeurs. *Insectes Sociaux*, 6, 41–81.
- Gray, R. (1994). Sparrows, matching, and the ideal free distribution: Can biological and psychological approaches be synthesised?. *Animal Behaviour*, 48, 411–423.
- Grimm, V. (1999). Ten years of individual-based modelling in ecology: What have we learned and what could we learn in the future?. *Ecological Modelling*, 115, 129–148.

- Grossberg, S. (1995). Neural dynamics of motion perception, recognition learning, and spatial attention. In Port, R., & van Gelder, T. (Eds.), *Mind as motion: Explorations in the dynamics of cognition*, pp. 449–491. MIT Press, Cambridge, MA.
- Hager, G. (1990). *Task directed sensor fusion and planning: A computational approach*. Kluwer Academic, Boston, MA.
- Halperin, J. (1990). Machine motivation. In Meyer, J., & Wilson, S. (Eds.), *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pp. 213–221 Cambridge, MA. MIT Press.
- Hamilton, W. (1964). The genetical evolution of social behaviour. *Journal of Theoretical Biology*, 7, 1–52.
- Hardcastle, V. (Ed.). (1999). *Where biology meets psychology: Philosophical essays*. MIT Press, Cambridge, MA.
- Harnard, S. (1990). The symbol grounding problem. *Physica D*, 42, 335–346.
- Harvey, I., Husbands, P., & Cliff, D. (1994). Seeing the light: Artificial evolution, real vision. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 392–401 Cambridge, MA. MIT Press.
- Hassell, M., & Varley, G. (1969). New inductive population model for insect parasites and its bearing on biological control. *Nature*, 223, 1133–1136.
- Hebb, D. (1980). *Essay on mind*. Lawrence Erlbaum, Hillsdale, NJ.
- Heider, F. (1959). On perception and event structure and the psychological environment. *Psychological Issues*, 1(4), 1–123.
- Helmreich, S. (1995). *Anthropology inside and outside the looking-glass worlds of artificial life*. Ph.D. thesis, Stanford University.
- Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence, and human thought*. MIT Press, Cambridge, MA.
- Herrnstein, R. (1961). Relative and absolute strength of responses as a function of frequency of reinforcement. *Journal of the Experimental Analysis of Behavior*, 4, 267–272.
- Herrnstein, R. (1970). On the law of effect. *Journal of the Experimental Analysis of Behavior*, 13(2), 243–266.
- Herrnstein, R. (1982). Melioration as behavioural dynamism. In Commons, M., Herrnstein, R., & Rachlin, H. (Eds.), *Quantitative analyses of behavior, vol II: Matching and maximizing accounts*, pp. 433–458. Ballinger Publishing Co., Cambridge, MA.
- Herrnstein, R. (1990). Rational choice theory: Necessary but not sufficient. *American Psychologist*, 45, 356–367.
- Herrnstein, R. (1997). *The matching law*. Harvard University Press, Cambridge, MA. A posthumous collection of the papers of R.J. Herrnstein, edited by H. Rachlin and D.I. Laibson.
- Herrnstein, R., & Heyman, G. (1979). Is matching compatible with reinforcement maximization on concurrent variable interval, variable ratio?. *Journal of the Experimental Analysis of Behavior*, 31, 209–223.

- Herrnstein, R., & Vaughan, Jr., W. (1980). Melioration and behavioral allocation. In Staddon, J. (Ed.), *Limits to action: The allocation of individual behavior*. Academic Press, New York.
- Hinde, R. (1956). Ethological models and the concept of drive. *Science*, *6*, 322–331.
- Hinde, R. (1957). Consequences and goals: Some issues raised by Dr. Kortland's paper on aspects and prospects of the concept of instinct. *British Journal of Animal Behaviour*, *5*, 116–118.
- Hinde, R., & Stevenson, J. (1970). Goals and response control. In Aronson, L. (Ed.), *Development and evolution of behavior*, pp. 216–237. W.H. Freeman, San Francisco.
- Hinson, J., & Staddon, J. (1983). Hill-climbing by pigeons. *Journal of the Experimental Analysis of Behavior*, *39*, 25–47.
- Hinton, G., & Nowlan, S. (1987). How learning can guide evolution. *Complex Systems*, *1*, 495–502.
- Holland, J. (1992). *Adaptation in natural and artificial systems*. MIT Press, Cambridge, MA.
- Holling, C. (1959). Some characteristics of simple types of predation and parasitism. *Canadian Entomologist*, *91*, 395–398.
- Holmgren, N. (1995). The ideal free distribution of unequal competitors: Predictions from a behaviour-based functional response. *Journal of Animal Ecology*, *64*, 197–212.
- Houston, A. (1986). The matching law applies to wagtails' foraging in the wild. *Journal of the Experimental Analysis of Behavior*, *45*, 15–18.
- Houston, A., & McNamara, J. (1985). The choice of two prey types that minimises the probability of starvation. *Behavioral Ecology and Sociobiology*, *17*, 135–141.
- Husbands, P., Harvey, I., Cliff, D., & Miller, G. (1997). Artificial evolution: A new path for artificial intelligence?. *Brain and Cognition*, *34*, 130–159.
- Husbands, P., & Meyer, J. (Eds.). (1998). *Proceedings of the First European Workshop on Evolutionary Robotics*, Vol. 1468 of *Lecture notes in computer science*, Berlin. Springer Verlag.
- Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, *10*(3/4), 185–211.
- Jakobi, N. (1998a). Evolutionary robotics and the radical envelope of noise hypothesis. *Adaptive Behaviour*, *6*(2), 325–368.
- Jakobi, N. (1998b). *Minimal simulations for evolutionary robotics*. Ph.D. thesis, Sussex University.
- Jakobi, N., Husbands, P., & Harvey, I. (1995). Noise and the reality gap: The use of simulation in evolutionary robotics. In Moran, F., Moreno, A., Merelo, J., & Chacon, P. (Eds.), *Proceedings of the Third European Conference on Artificial Life*, pp. 704–720 Berlin. Springer-Verlag.
- James, W. (1897). *The will to believe, and other essays in popular philosophy*. Longmans, New York.
- Jarvis, J. (1981). Eusociality in a mammal: Cooperative breeding in naked mole-rat colonies. *Science*, *212*, 571–573.



- Jeffery, W., Strickler, A., Guiney, S., Heyser, D., & Tomarev, S. (2000). Prox 1 in eye degeneration and sensory organ compensation during development and evolution of the cavefish *Astyanax*. *Development, Genes, and Evolution*, 210(5), 223–230.
- Johnson, M. (Ed.). (1993). *Brain development and cognition*. Blackwell, Oxford.
- K-Team (1993). Khepera: the user's manual. Tech. rep., LAMI-EPFL.
- Kacelnik, A. (1997). Normative and descriptive models of decision making: Time discounting and risk sensitivity. In *Characterising human psychological adaptations*, pp. 51–70. Wiley, Chichester.
- Kaelbling, L., Littman, M., & Moore, A. (1996). Reinforcement learning: A survey. *Journal of Artificial Intelligence*, 4, 237–285.
- Kandel, E., & Schwartz, J. (1982). Molecular biology of learning: Modulation of transmitter release. *Science*, 218, 433–443.
- Kauffman, S. (1993). *The origins of order. Self-organisation and selection in evolution*. Oxford University Press, Oxford.
- Keijser, F. (1998). Some armchair worries about wheeled behaviour. In Pfeifer, R., Blumberg, B., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 5: Proceedings of the Fifth International Conference on Simulation of Adaptive Behavior*, pp. 13–22 Cambridge, MA. MIT Press.
- Kelso, J. (1995). *Dynamic patterns: The self-organisation of brain and behavior*. MIT Press, Cambridge, MA.
- Kennedy, M., & Gray, R. (1993). Can ecological theory predict the distribution of foraging animals? A critical analysis of experiments on the ideal free distribution. *Oikos*, 68, 158–166.
- Kien, J., & Altman, J. (1992). *Neuroethology of action selection*. MIT Press, Cambridge, MA.
- Kitano, H., Hamahashi, S., Kitazawa, J., Takao, K., & Imai, S. (1997). The virtual biology laboratories: A new approach of computational biology. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 274–283 Cambridge, MA. MIT Press.
- Kolen, J., & Pollack, J. (1995). The observer's paradox: Apparent computational complexity in physical systems. *Journal of Experimental and Theoretical Artificial Intelligence*, 7, 253–277.
- Kolmogorov, A. (1965). Three approaches to the definition of the concept 'quantity of information'. *Problems Inform. Transmission*, 1, 1–7.
- Koza, J., Rice, J., & Roughgarden, J. (1992). Evolution of food foraging strategies for the caribbean *Anolis* lizard using genetic programming. *Adaptive Behavior*, 1(2), 171–199.
- Krebs, J., & Davies, N. (Eds.). (1991). *Behavioural ecology: an evolutionary approach*. Blackwell Scientific Publishers, Oxford. 3rd edition.
- Krebs, J., & Kacelnik, A. (1991). Decision making. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology*, pp. 105–137. Blackwell Scientific Publishers, Oxford. 3rd edition.
- Lachman, R., Lachman, J., & Butterfield, E. (1979). *Cognitive psychology and information processing*. Erlbaum, Hillsdale, NJ.

- Lambrinos, D., & Scheier, C. (1995). Extended Braitenberg architectures. Tech. rep. AI Lab no. 95.10, Computer Science Dept., University of Zurich.
- Langton, C. (Ed.). (1995). *Artificial life: An overview*. MIT Press, Cambridge, MA.
- Langton, C. (1989). Artificial life. In Langton, C. (Ed.), *Proceedings of the Interdisciplinary Workshop on the Synthesis and Simulation of Living Systems*, Vol. VI of *Santa Fe Institute studies in the sciences of complexity*, pp. 1–48. Addison-Wesley.
- Lashley, K. (1951). The problem of serial order in behavior. In Jeffress, L. (Ed.), *Cerebral mechanisms in behavior*. John Wiley and sons, New York.
- Lessells, C. (1995). Putting resource dynamics into continuous input ideal free distribution models. *Animal Behaviour*, 49, 487–494.
- Levins, R., & Lewontin, R. (1985). *The dialectical biologist*. Harvard University Press, Cambridge, MA.
- Lewontin, R. (1983). The organism as the subject and object of evolution. Reprinted in Levins and Lewontin (1985), pp. 85-106.
- Lindgren, K. (1991). Evolutionary phenomena in simple dynamics. In Langton, C., Farmer, J., Rasmussen, S., & Taylor, C. (Eds.), *Artificial life II*, Vol. XVI of *Santa Fe Institute studies in the sciences of complexity*, pp. 295–312. Addison-Wesley, Reading, MA.
- Lloyd, D. (1989). *Simple minds*. MIT Press, Cambridge, MA.
- Lorenz, K. (1935). Companionship in bird life. In Schiller and Lashley (1957), pp. 83-128.
- Lorenz, K. (1937). The nature of instinct: The conception of instinctive behaviour. In Schiller and Lashley (1957), pp. 129-175.
- Lorenz, K. (1939). Comparative study of behaviour. In Schiller and Lashley (1957), pp. 239-263.
- Lorenz, K. (1948). *The natural science of the human species*. MIT Press, Cambridge, MA. edited from author's posthumous works (1944-48) by A. von Cranach, trans. R.D. Martin, 1996.
- Lorenz, K., & Tinbergen, N. (1938). Taxis and instinct: Taxis and instinctive action in the egg-retrieving behaviour of the Greylag goose. In Schiller and Lashley (1957), pp. 176-208.
- Maes, P. (1990). A bottom-up mechanism for behaviour-selection in an artificial creature. In Meyer, J., & Wilson, S. (Eds.), *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pp. 238–246 Cambridge, MA. MIT Press.
- Maes, P. (1994). Modelling adaptive autonomous agents. *Artificial Life*, 1, 135–162.
- Maley, C. (1999). Methodologies for the use of computational models in theoretical biology. In Wu, A. (Ed.), *Proceedings of the 1999 Genetic and Evolutionary Computation Conference Workshop Program*, pp. 16–19 San Francisco. Morgan Kaufmann.
- Mangel, M., & Clark, C. (1988). *Dynamic modeling in behavioral ecology*. Princeton University Press, Princeton.
- Marjanovic, M., Scassellati, B., & Williamson, M. (1996). Self-taught visually-guided pointing for a humanoid robot. In *From animals to animats 4: Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pp. 35–44 Cambridge, MA. MIT Press.

- Marr, D. (1982). *Vision*. W.H. Freeman, San Francisco.
- Marx, K. (1848). Theses on Feuerbach. In *Karl Marx and Frederic Engels: Selected works, [1968]*. International Publishers, New York.
- Maturana, H., & Varela, F. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel Publishing, Dordrecht, Holland.
- Maturana, H., & Varela, F. (1987). *The tree of knowledge: The biological roots of understanding*. Shambhala, Boston, MA.
- Maynard Smith, J. (1974). *Models in ecology*. Cambridge University Press, Cambridge.
- Maynard Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press, Cambridge.
- Maynard-Smith, J., & Szathmáry, E. (1995). *The major transitions in evolution*. W.H. Freeman, San Francisco.
- Mazur, J. (1981). Optimization theory fails to predict performance of pigeons in a two-response situation. *Science*, 214, 823–825.
- McFarland, D. (1985). *Animal behaviour: Psychology, ethology, and evolution*. Pittman Publishing, London.
- McFarland, D. (1996). Animals as cost based robots. In Boden, M. (Ed.), *The philosophy of artificial life*, pp. 179–209. Oxford University Press, Oxford.
- McFarland, D., & Bösner, M. (1993). *Intelligent behavior in animals and robots*. MIT Press, Cambridge, MA.
- McFarland, D., & Sibly, R. (1972). ‘Unitary drives’ revisited. *Animal Behaviour*, 20, 548–563.
- McFarland, D., & Sibly, R. (1975). The behavioural final common path. *Philosophical Transactions of the Royal Society of London: Series B*, 270(907), 265–293.
- McLeod, P., & Dienes, Z. (1996). Do fielders know where to go to catch the ball, or only how to get there?. *Journal of Experimental Psychology: Human Perception and Performance*, 22, 531–543.
- McShea, D. (1991). Complexity and evolution: What everybody knows. *Biology and Philosophy*, 6(3), 303–324.
- McShea, D. (1993). Evolutionary change in the morphological complexity of the mammalian vertebral column. *Evolution*, 47, 730–740.
- McShea, D. (1996). Metazoan complexity and evolution: Is there a trend?. *Evolution*, 50, 477–492.
- Müeller, J. (1833-40). *Elements of physiology*. Taylor and Walton, London. trans. Wm. Baly.
- Menczer, F., & Belew, R. (1996). From complex environments to complex behaviours. *Adaptive Behaviour*, 4(3-4), 309–355.
- Merleau-Ponty, M. (1963). *The structure of behavior*. Beacon Press, Boston, MA.
- Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.). (2000). *From animals to animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.

- Meyer, J., Husbands, P., & Harvey, I. (1998). Evolutionary robotics: A survey of applications and problems. In Husbands, P., & Meyer, J. (Eds.), *Proceedings of the First European Workshop on Evolutionary Robotics*, pp. 1–22. Springer Verlag.
- Miglino, O., Lund, H., & Nolfi, S. (1996). Evolving mobile robots in simulated and real environments. *Artificial Life*, 2(4), 417–434.
- Milinski, M. (1979). An evolutionary stable feeding strategy in sticklebacks. *Z.Tierpsychol.*, 51, 36–40.
- Milinski, M. (1994). Ideal free theory predicts more than only input matching - a critique of Kennedy and Gray's review. *Oikos*, 71(1), 163–165.
- Milinski, M., & Parker, G. (1991). Competition for resources. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology*, pp. 137–168. Blackwell Scientific Publishers, Oxford. 3rd edition.
- Miller, G., Galanter, E., & Pribram, K. (1960). *Plans and the structure of behaviour*. Holt, Rinehart and Winston, New York.
- Miller, G. (1995). Artificial life as theoretical biology: How to do real science with computer simulation. CSRP 378. Tech. rep., School of Cognitive and Computing Sciences, Sussex University.
- Miller, G., & Cliff, D. (1994). Protean behavior in dynamic games: Arguments for the co-evolution of pursuit-evasion tactics. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior.*, pp. 411–420 Cambridge, MA. MIT Press.
- Millikan, R. (1989). In defence of proper functions. *Philosophy of science*, 56, 288–302.
- Milner, A., & Goodale, M. (1995). *The visual brain in action*. Oxford University Press, Oxford.
- Minsky, M. (1986). *The society of mind*. Simon and Schuster, New York.
- Mitchell, M. (1997). *An introduction to genetic algorithms*. MIT Press, Cambridge, MA.
- Mondada, F., Franzi, E., & Ienne, P. (1993). Mobile robot miniaturization: A tool for investigation in control algorithms. In Yoshikawa, T., & Miyazaki, F. (Eds.), *Proceedings of the Third International Symposium on Experimental Robots*, pp. 501–513 Berlin. Springer-Verlag.
- Moody, A., & Houston, A. (1995). Interference and the ideal free distribution. *Animal Behaviour*, 49, 1065–1072.
- Moran, J., & Desimone, R. (1985). Selective attention gates visual processing. *Science*, 229, 782–784.
- Myers, J. (1976). Probability learning and sequence learning. In Estes, W. (Ed.), *Handbook of learning and cognitive processes*, Vol. 1, pp. 171–205. Erlbaum, Hillsdale, NJ.
- Nakagaki, T., Yamada, H., & Toth, A. (2000). Maze solving by an amoeboid organism. *Nature*, 407, 470–471.
- Newell, A., & Simon, H. (1976). Computer science as empirical enquiry: Symbols and search. *Communications of the ACM [Association for Computing Machinery]*, 19, 113–126.
- Nilson, N. (1984). Shakey the robot. SRI A.I. Center Technical Note 323.
- Noble, J. (1998). *The evolution of animal communication systems: Questions of function examined through simulation*. Ph.D. thesis, University of Sussex.

- Nolfi, S. (1998). Evolutionary robotics: Exploiting the full power of self-organisation. *Connection Science*, 10(3/4), 167–185.
- Nolfi, S., & Floreano, D. (2000). *Evolutionary robotics: The biology, intelligence, and technology of self-organizing machines*. MIT Press, Cambridge, MA.
- Nolfi, S., & Parisi, D. (1993). Self-selection of input stimuli for improving performance. In Bekey, G. (Ed.), *Neural networks and robotics*, pp. 403–418. Kluwer Academic, Boston, MA.
- Oyama, S. (1985). *The ontogeny of information*. Cambridge University Press, Cambridge.
- Parasuraman, R. (Ed.). (1998). *The attentive brain*. MIT Press, Cambridge, MA.
- Parker, G. (1978). Searching for mates. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology: An evolutionary approach*, pp. 214–244. Blackwell Scientific Publishers, Oxford. 1st edition.
- Parker, G. (1970). The reproductive behaviour and the nature of sexual selection in *Scatophaga stercoraria* L. ii. the fertilization rate and the spatial and temporal relationships of each sex around the site of mating and oviposition. *Journal of Animal Ecology*, 39, 205–228.
- Parker, G., & Stuart, R. (1976). Animal behavior as a strategy optimizer: Evolution of resource assessment strategies and optimal emigration thresholds. *American Naturalist*, 110, 1055–1076.
- Parker, G., & Sutherland, W. (1986). Ideal free distributions when individual differ in competitive ability: Phenotype limited ideal free models. *Animal Behaviour*, 34, 1222–1242.
- Pashler, H. (1999). *The psychology of attention*. MIT Press, Cambridge, MA.
- Pask, G. (1961). *An approach to cybernetics*. Hutchinson, London.
- Pavlov, I. (1927). *Conditioned reflexes*. trans. G.V. Anrep (Dover, NY, 1960).
- Peel, J. (1971). *Herbert Spencer: The evolution of a sociologist*. Basic Books, New York.
- Perry, R. (1935). *The thought and character of William James*. Little, Brown, and co., Boston.
- Pfeifer, R. (1996). Building “fungus eaters”: Design principles of autonomous agents. In Maes, P., Mataric, M., Meyer, J., Pollack, J., & Wilson, S. (Eds.), *From animals to animats 4: Proceedings of the Fourth Conference on the Simulation of Adaptive Behavior*, pp. 4–12 Cambridge, MA. MIT Press.
- Pfeifer, R., Blumberg, B., Meyer, J., & Wilson, S. (Eds.). (1998). *From animals to animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior*, Cambridge, MA. MIT Press.
- Pfeifer, R., & Scheier, C. (1994). From perception to action: The right direction?. In Gaussier, P., & Nicoud, J. (Eds.), *From perception to action*, pp. 1–11 Los Alamitos, CA. IEEE Computer Society Press.
- Pfeifer, R., & Scheier, C. (1999). *Understanding intelligence*. MIT Press, Cambridge, MA.
- Piaget, J. (1971). *Biology and knowledge: An essay on the relations between organic regulations and cognitive processes*. Edinburgh University Press, Edinburgh.
- Pierce, G., & Ollason, J. (1987). Eight reasons why optimal foraging theory is a waste of time. *Oikos*, 49(1), 111–117.

- Platt, M., & Glimcher, P. (1999). Neural correlates of decision variables in parietal cortex. *Nature*, 400, 233–238.
- Port, R., & van Gelder, T. (1995). *Mind as motion: Explorations in the dynamics of cognition*. MIT Press, Cambridge, MA.
- Posch, M. (1999). Win-stay, lose-shift strategies for repeated games - memory length, aspiration levels and noise. *Journal of Theoretical Biology*, 198, 183–195.
- Posch, M., Pichler, A., & Sigmund, K. (1999). The efficiency of adapting aspiration levels. *Proceedings of the Royal Society of London: Series B*, 266, 1427–1436.
- Posner, M. (1978). *Chronometric explorations of the mind*. Lawrence Erlbaum Associates Inc., Hillsdale, NJ.
- Posner, M. (1980). Orienting of attention. *Quarterly Journal of Experimental Psychology*, 32, 3–25.
- Powers, W. (1973). *Behaviour: The control of perception*. Aldline/de Gruyter, Greenwood, NY.
- Prescott, T., Redgrave, P., & Gurney, K. (1999). Layered control architectures in robots and vertebrates. *Adaptive Behavior*, 7, 99–127.
- Pringle, J. (1951). On the parallel between learning and evolution. *Behaviour*, 3, 90–110.
- Pulliam, H., & Caraco, T. (1984). Living in groups: Is there an optimal group size?. In Krebs, J., & Davies, N. (Eds.), *Behavioural ecology: An evolutionary approach*, pp. 122–147. Blackwell Scientific Publishers, Oxford. 2nd edition.
- Putnam, H. (1967). The mental life of some machines. In Castañeda, H. (Ed.), *Intentionality, minds, and perception*. Wayne State University Press, Detroit, MI.
- Pylyshyn, Z. (Ed.). (1985). *The robot's dilemma: The frame problem in artificial intelligence*. Ablex, Norwood, NJ.
- Quine, W. (1951). Two dogmas of empiricism. *The Philosophical Review*, 60, 20–43.
- Raphael, B. (1976). *The thinking computer: Mind inside matter*. Freeman, San Francisco.
- Rechten, C., Avery, M., & Stevens, J. (1983). Optimal prey selection: Why do great tits show partial preferences?. *Animal Behaviour*, 31, 576–584.
- Redgrave, P., Prescott, T., & Gurney, K. (1999). The basal ganglia: A vertebrate solution to the action selection problem?. *Neuroscience*, 89, 1009–1023.
- Richards, R. (1987). *Darwin and the emergence of evolutionary theories of mind and behavior*. University of Chicago Press, Chicago.
- Ristau, C. (Ed.). (1991). *Cognitive ethology: The minds of other animals - essays in honor of Donald R. Griffin*. Lawrence Erlbaum, Hillsdale, NJ.
- Rosenblatt, K., & Payton, D. (1989). A fine-grained alternative to the subsumption architecture for mobile robot control. In *Proceedings of the IEEE/INNS International Joint Conference on Neural Networks*, pp. 317–323 Washington. IEEE Press.
- Rosenbleuth, A., Wiener, N., & Bigelow, J. (1943). Behavior, purpose, and teleology. *Philosophy of science*, 10, 18–24.

- Rumelhart, D., McClelland, J., & the PDP research group (1986). *Parallel distributed processing: Explorations in the microstructure of cognition*. MIT Press, Cambridge, MA.
- Rutkowska, J. (1994a). Emergent functionality in human infants. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 179–187 Cambridge, MA. MIT Press.
- Rutkowska, J. (1994b). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior*, 2, 349–373.
- Rutkowska, J. (1997). What's value worth? Constraints on unsupervised behaviour acquisition. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 290–298 Cambridge, MA. MIT Press.
- Ruxton, G., Gurney, W., & Roos, A. (1992). Interference and generation cycles. *Theoretical Population Biology*, 42, 235–253.
- Sahota, M. (1994). Action selection for robots in dynamic environments through inter-behaviour bidding. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 138–142 Cambridge, MA. MIT Press.
- Scheier, C., & Lambrinos, D. (1994). Categorization in a real-world agent using haptic exploration and active perception. In Maes, P., Mataric, M., Meyer, J., Pollack, J., & Wilson, S. (Eds.), *From animals to animats 4: Proceedings of the Fourth International Conference on the Simulation of Adaptive Behavior*, pp. 65–74 Cambridge, MA. MIT Press.
- Scheier, C., & Pfeifer, R. (1995). Classification as sensory-motor coordination. In Moran, F., Moreno, A., Merelo, J., & Chacon, P. (Eds.), *Proceedings of the Third European Conference on Artificial Life*, pp. 656–667 Berlin. Springer-Verlag.
- Schiller, C., & Lashley, K. (1957). *Instinctive behaviour: The development of a modern concept*. International University Press, New York.
- Schuster, P. (1996). How does complexity arise in evolution?. Tech. rep., TBI working paper 96-05-026.
- Seth, A. (1997). Interaction, uncertainty, and the evolution of complexity. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 521–530 Cambridge, MA. MIT Press.
- Seth, A. (1998a). The evolution of complexity and the value of variability. In Adami, C., Belew, R., Kitano, H., & Taylor, C. (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on the Simulation and Synthesis of Living Systems*, pp. 209–221 Cambridge, MA. MIT Press.
- Seth, A. (1998b). Evolving action selection and selective attention without actions, attention, or selection. In Pfeifer, R., Blumberg, B., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 5: Proceedings of the Fifth International Conference on the Simulation of Adaptive Behaviour*, pp. 139–147 Cambridge, MA. MIT Press.
- Seth, A. (1998c). Noise and the evolution of complexity: A study in evolutionary robotics. In Husbands, P., & Meyer, J. (Eds.), *Proceedings of the First European Workshop on Evolutionary Robotics*, pp. 123–137 Berlin. Springer-Verlag.

- Seth, A. (1999a). A cybernetic perspective on the role of noise in the iterated prisoner's dilemma. Tech. rep. 506, University of Sussex Cognitive Science Research Paper.
- Seth, A. (1999b). Evolving behavioural choice: An investigation of Herrnstein's matching law. In Floreano, D., Nicoud, J.-D., & Mondada, F. (Eds.), *Proceedings of the Fifth European Conference on Artificial Life*, pp. 225–236 Berlin. Springer-Verlag.
- Seth, A. (2000). Unorthodox optimal foraging theory. In Meyer, J., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.), *From animals to animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior*, pp. 471–481 Cambridge, MA. MIT Press.
- Shannon, C., & Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press, Illinois.
- Shapiro, L. (1999). Presence of mind. In Hardcastle, V. (Ed.), *Where biology meets psychology: Philosophical essays*, pp. 83–99. MIT Press, Cambridge, MA.
- Shettleworth, S. (1988). Foraging as operant behavior and operant behavior as foraging: What have we learned?. In Bower, G. (Ed.), *The psychology of learning and motivation: Advances in research and theory*, Vol. 22, pp. 1–49. Academic Press, New York.
- Sibly, R., & McFarland, D. (1976). On the fitness of behavior sequences. *American Naturalist*, 110, 610–617.
- Simon, H. (1988). *The sciences of the artificial*. MIT Press, Cambridge, MA. 3rd edition.
- Simon, H. (1956). Rational choice and the structure of the environment. *Psychological Review*, 63(2), 129–138.
- Skinner, B. (1969). *Contingencies of reinforcement*. Appleton-Century-Crofts, New York.
- Smithers, T. (1994). On why better robots make it harder. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 64–72 Cambridge, MA. MIT Press.
- Smucker, M., Stanley, E., & Ashlock, D. (1996). Analyzing social network structures in the iterated prisoner's dilemma with choice and refusal. Technical report no. 1259, University of Wisconsin-Madison, Dept of Computer Science.
- Sober, E. (1990). Putting the function back in functionalism. In Lycan, W. (Ed.), *Mind and cognition: A reader*. Basil Blackwell, Cambridge, MA.
- Sober, E. (1996). Learning from functionalism - prospects for strong artificial life. In Boden, M. (Ed.), *The philosophy of artificial life*, pp. 361–379. Oxford University Press, Oxford.
- Spencer, H. (1855). *Principles of psychology*. Longman, Brown, and Green.
- Spencer, H. ([1857] 1915). *Progress; Its law and cause*. Appleton, New York. In: *Essays: Scientific, political and speculative*.
- Spencer, H. (1904). *An autobiography*. Williams and Norgate, London. 2 volumes.
- Spier, E. (1997). *From reactive behaviour to adaptive behaviour*. Ph.D. thesis, University of Oxford.
- Spier, E., & McFarland, D. (1998). Possibly optimal decision making under self-sufficiency and autonomy. *Journal of Theoretical Biology*, 189, 317–331.



- Staddon, J. (1992). Rationality, melioration, and law-of-effect models for choice. *Psychological Science*, 3, 136–141.
- Stanley, E., Ashlock, D., & Tesfatsion, L. (1994). Iterated prisoner's dilemma with choice and refusal of partners. In Langton, C. (Ed.), *Artificial life III*, Vol. XVII of *Santa Fe Institute studies in the sciences of complexity*, pp. 131–175. Addison-Wesley, Reading, MA.
- Stearns, S., & Schmid-Hempel, P. (1987). Evolutionary insights should not be wasted. *Oikos*, 49(1), 118–125.
- Steels, L. (1990). Towards a theory of emergent functionality. In Meyer, J., & Wilson, S. (Eds.), *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pp. 451–461 Cambridge, MA. MIT Press.
- Steels, L. (1992). The PDL reference manual. Tech. rep., VUB AI Lab, Memo 92-5.
- Stein, B., & Meredith, M. (Eds.). (1994). *The merging of the senses*. MIT Press, Cambridge, MA.
- Stephens, D., & Krebs, J. (1986). *Foraging theory*. Monographs in behaviour and ecology. Princeton University Press, Princeton.
- Stillman, R., Goss-Custard, J., & Caldow, R. (1997). Modelling interference from basic foraging behaviour. *Journal of Animal Ecology*, 66, 692–703.
- Sutherland, N., & Mackintosh, N. (1971). *Mechanisms of animal discrimination learning*. Academic Press, New York.
- Sutherland, W. (1983). Aggregation and the 'ideal free' distribution. *Journal of Animal Ecology*, 52, 821–828.
- Sutherland, W., & Koene, P. (1982). Field estimates of the strength of interference between Oystercatchers *Haematopus ostralegus*. *Oecologia*, 55, 108–109.
- Sutton, R., & Barto, A. (1998). *Reinforcement learning*. MIT Press, Cambridge, MA.
- Tani, J., Yamamoto, J., & Nishi, H. (1997). Dynamical interactions between visual attention, learning, and behavior. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 309–317 Cambridge, MA. MIT Press.
- Teyke, T., & Schaerer, S. (1994). Blind Mexican cave fish *Astyanax hubsi* respond to moving visual stimuli. *Journal of Experimental Biology*, 188, 89–101.
- Thelen, E., & Smith, L. (1994). *A dynamic systems approach to the development of cognition and action*. MIT Press, Cambridge, MA.
- Thompson, W. (1917). *On growth and form*. Cambridge University Press, Cambridge.
- Thuisjman, F., Peleg, B., Amitai, M., & Shmida, A. (1995). Automata, matching, and foraging behaviour of bees. *Journal of Theoretical Biology*, 175, 305–316.
- Tinbergen, N. (1950). The hierarchical organisation of nervous mechanisms underlying instinctive behaviour. *Symposium for the society for experimental biology*, 4, 305–312.
- Tinbergen, N. (1951). *The study of instinct*. Clarendon Press, Oxford.
- Tinbergen, N. (1963). On aims and methods of ethology. *Zeitschrift für Tierpsychologie*, 20, 410–433.

- Toates, F. (1982). *Animal behaviour - a systems approach*. Cambridge University Press, Cambridge.
- Toates, F. (1986). *Motivational systems*. Cambridge University Press, Cambridge.
- Toates, F., & Jensen, P. (1990). Ethological and psychological models of motivation - towards a synthesis. In Meyer, J., & Wilson, S. (Eds.), *From animals to animats: Proceedings of the First International Conference on the Simulation of Adaptive Behavior*, pp. 194–205 Cambridge, MA. MIT Press.
- Todd, P. (Ed.). (1996). Special issue of *Adaptive Behavior* on environment structure and behavior.
- Todd, P., Wilson, S., Somayaji, A., & Yanco, H. (1994). The blind breeding the blind: Adaptive behavior without looking. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 228–237 Cambridge, MA. MIT Press.
- Todd, P., & Wilson, S. (1993). Environment structure and adaptive behavior from the ground up. In Meyer, J., Roitblat, H., & Wilson, S. (Eds.), *From animals to animats 2: Proceedings of the Second International Conference on the Simulation of Adaptive Behavior*, pp. 11–20 Cambridge, MA. MIT Press.
- Todd, P., & Yanco, H. (1996). Environmental effects on minimal behaviours in the minimat world. *Adaptive Behavior*, 4(3/4), 365–413.
- Tononi, G., Sporns, O., & Edelman, G. (1994). A measure for brain complexity: Relating functional segregation and integration in the nervous system. *Proceedings of the National Academy of Science (USA)*, 91, 5033–5037.
- Tononi, G., Sporns, O., & Edelman, G. (1996). A complexity measure for selective matching of signals by the brain. *Proceedings of the National Academy of Science (USA)*, 93, 3422–3427.
- Townsend, J., & Busemeyer, J. (1995). Dynamic representation of decision-making. In Port, R., & van Gelder, T. (Eds.), *Mind as motion: Explorations in the dynamics of cognition*, pp. 101–121. MIT Press, Cambridge, MA.
- Trengenza, T. (1994). Common misconceptions in applying the ideal free distribution. *Animal Behaviour*, 47, 485–487.
- Trengenza, T., Parker, G., & Thompson, D. (1996). Interference and the ideal free distribution: Models and tests. *Behavioral Ecology*, 7(4), 379–386.
- Treisman, A., & Gelade, G. (1980). A feature-integration theory of attention. *Cognitive Psychology*, 12, 97–136.
- Tsotsos, J., Culhane, S., Wai, W., Lai, Y., Davis, N., & Nuflo, F. (1995). Modelling visual attention via selective tuning. *Artificial Intelligence*, 78, 507–545.
- Tyler, J., & Hargrove, W. (1997). Predicting spatial distribution of foragers over large resource landscapes: A modeling analysis of the ideal free distribution. *Oikos*, 79, 376–386.
- Tyrrell, T. (1993). *Computational mechanisms for action selection*. Ph.D. thesis, University of Edinburgh.
- Tyrrell, T. (1994). An evaluation of Maes' bottom-up mechanism for behavior selection. *Adaptive Behavior*, 2(4), 307–348.

- Van der Meer, J., & Ens, B. (1997). Models of interference and their consequences for the spatial distribution of ideal and free predators. *Journal of Animal Ecology*, 66, 846–858.
- van Gelder, T. (1998). The dynamical hypothesis in cognitive science. *Behavioral and Brain Sciences*, 21, 615–624.
- Varela, F. (1979). *Principles of biological autonomy*. Elsevier, North Holland, NY.
- Varela, F. (1997). Patterns of life: Intertwining identity and cognition. *Brain and Cognition*, 34, 72–87.
- Varela, F., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge, MA.
- von Uexküll, J. (1934). A stroll through the worlds of animals and men. In Lashley, K. (Ed.), *Instinctive behavior*. International Universities Press, New York.
- Vygotsky, L. (1986). *Thought and language*. MIT Press, Cambridge, MA. Translation of 1962 edition.
- Waddington, C. (1942). Canalization and the development and inheritance of acquired characters. *Nature*, 3811, 563–565.
- Waddington, C. (1966). *Principles of development and differentiation*. Macmillan, New York.
- Walter, W. (1950). An imitation of life. *Scientific American*, 182(5), 42–45.
- Walter, W. (1953). *The living brain*. Duckworth, London.
- Weber, T. (1998). News from the realm of the ideal free distribution. *Trends in Ecology and Evolution*, 13(3), 89–90.
- Werner, G. (1994). Using second order neural connections for motivation of behavioural choice. In Cliff, D., Husbands, P., Meyer, J., & Wilson, S. (Eds.), *From animals to animats 3: Proceedings of the Third International Conference on the Simulation of Adaptive Behavior*, pp. 154–164 Cambridge, MA. MIT Press.
- Wheeler, M. (1996). From robots to Rothko: The bringing forth of worlds. In Boden, M. (Ed.), *The philosophy of artificial life*, pp. 209–237. Oxford University Press, Oxford.
- Whitehead, A. (1925). *Science and the modern world*. Macmillan, New York.
- Wiener, N. (1948). *Cybernetics*. John Wiley and Sons., New York.
- Williams, B. (1994). Reinforcement and choice. In Mackintosh, N. (Ed.), *Animal learning and cognition*, pp. 81–108. Academic Press, New York.
- Wilson, E. (1971). *The insect societies*. Harvard University Press, Cambridge, MA.
- Wilson, E. (1975). *Sociobiology: The new synthesis*. Harvard University Press, Cambridge, MA.
- Wright, L. (1973). Functions. *Philosophical Review*, 82, 139–168.
- Young, R. (1970). *Mind, brain, and adaptation in the nineteenth century*. Clarendon Press, Oxford.

## Appendix A

This appendix fills out the details of the various artificial evolution models developed in the course of the dissertation. The information provided here is *not* necessary for the interpretation of these models (sufficient details for this purpose appear in the main body of the text), but may be useful for their replication.

### Prisoner's dilemma models

Table 1 provides the various parameter values for both the compulsory model and the IPD/CR model of chapter 3. A tournament selection GA (see Mitchell, 1997) is used in both IPD/CR and compulsory models in which a new population is formed by repeatedly (and randomly) choosing two individuals (with replacement), and selecting the individual with the higher fitness to be the first parent. This process is repeated to select a second parent, after which point mutation and crossover operators are applied in order to obtain a 'child' which is then placed in the new population. This whole process is repeated until a new population is created. Point mutation involves swapping the target allele (from c to d, or vice-versa); 'splitting' and 'doubling' mutations are described in detail in chapter 3.

	Compulsory model	IPD/CR model
Population	30	30
Iterations per game	60	60
Noise levels	2%	2%
Generations	10,000	10,000
Crossover rate	0.95	0.95
All mutations	0.001	0.005
Initial memory	1	1
Cost on complexity	0	0.0025
$\gamma$	n/a	0.7
$n$	n/a	2
$\omega$	n/a	1.6

Table 1: Parameters for the compulsory and IPD/CR models (chapter 3).

The replication of the prisoner's dilemma model of Posch (1999) also uses a tournament selection GA, as does Posch himself, and to ensure the fidelity of the replication the same GA parameters are employed. Crossover is implemented with a per-allele probability of 0.04 of exchange with the corresponding allele of the other parent. Point mutation is also implemented with probability 0.04 per allele; each point mutation involves replacing the target allele with a randomly selected value from the relevant initial range ( $[-1.0,6.0]$  for  $a_0$  and  $[0.0,1.0]$  for  $\gamma$ ). Also,

	high-noise simulations	low-noise simulations
IR (range 0-1024)	$\pm 50$	$\pm 10$
Background IR (range 0-1024)	$\pm 10$	0
Ambient light (range 0-512)	$\pm 50$	$\pm 5$
Floor sensor (range 0-512)	$\pm 50$	$\pm 5$
Robot position	$\pm 0.1cm$	0
Robot orientation	$\pm 0.02rad$	0
Turning noise	$\pm 0.2rads$	0
Friction	$\pm 3cm$	0
Arena size	$\pm 5cm$	0
Charge radius	$\pm 1cm$	0
$\theta$	$\pm 0.25rad$	0

Table 2: ‘Homing navigation’ simulation noise levels. Background IR refers to IR sensor activation in the absence of stimulation from arena wall proximity. Turning noise is an additional component of orientation noise applied when the absolute difference between left and right wheel speeds is above 7.0. Friction is applied when the robot collides with a wall, and  $\theta$  refers to the angle of acceptance of the ambient light sensors.

for the ‘adaptive’ strategies the inherited aspiration rate is the initial aspiration rate  $a_0$  and *not* the aspiration rate following adaptation.

### Homing navigation models

The ‘homing navigation’ model of chapter 4 utilises a distributed GA rather than a tournament selection GA. To begin with, a population of 100 random genotypes is arranged on a 10 by 10 toroidal grid and the fitness of each assessed. A random grid position is chosen, and a ‘pool’ of nine genotypes constructed from the 3 by 3 sub-grid surrounding this position. Two of the fittest members from this pool are chosen as parents using stochastic rank-based selection (with replacement) and used to generate a new genotype by means of crossover and mutation. The new genotype is placed back into the population in place of the weakest pool member (again chosen using a stochastic rank-based scheme) and evaluated; the parents are re-evaluated with a probability of 0.8. Crossover probability is set at 0.95, with a 0.03 probability of point mutation per allele. Each point mutation shifts the value of the allele by a random amount selected from the range  $[-0.5, 0.5]$ ; if the shift transgresses the allowable range  $[-1.0, 1.0]$ , the post-mutation value is set by selecting a random value between the pre-mutation value and the bound in question. The noise levels used in both high-noise and low-noise simulations are given in table 2.

### Behaviour coordination model

The ‘behaviour coordination’ model of chapter 6 utilises a distributed GA of very similar structure to that of chapter 4. , the only differences being (1) a crossover probability of 0.5, and (b) a per-allele point mutation probability of 0.04. Each point mutation shifted the value of the allele by value selected from the (integer) range  $[-5, 5]$  according to a Gaussian distribution. Transgressions

of allowable ranges are dealt with as in the ‘homing-navigation’ model.

### **Matching and interference models**

The IOS models of chapters 7 and 8 also employ a distributed GA with a population of 100. The only differences from the GA of chapter 4 are (1) a crossover probability of 0.5 (as above), (2) each new genotype carries on average 2 point mutations, and (3) each point mutation shifts the value of the locus by an amount drawn from a Gaussian distribution of radius 0.13. Transgressions of allowable ranges are dealt with as in the ‘homing-navigation’ model.

### **The sampling strategy IFD model**

The sampling strategy IFD model of chapter 9 returns to the use of a tournament selection GA, of almost exactly the same structure as described for the prisoner’s dilemma models (above). The only difference is to do with point mutation; in this model it occurs with a per-allele probability of 0.01, with each mutation again shifting the value of the locus by an amount drawn from a Gaussian distribution of radius 0.13. Transgressions of allowable ranges are dealt with as in the ‘homing-navigation’ model.

## Appendix B

This appendix sets out the various mathematical excursions referred to in the dissertation.

### The matching law

Here we consider the distinction between matching to available resources and matching to obtained resources in the context of a direct link between response and reward.

The matching law, considered in terms of matching to *obtained* resources, can be expressed in terms of ratios,

$$\frac{B_X}{B_Y} = \frac{R_X}{R_Y},$$

where  $B_X, B_Y$  = rate of response to  $X, Y$ , and  $R_X, R_Y$  = resources obtained from  $X, Y$ . Now, assume without loss of generality that  $R_X = f(F_X, B_X)$  and  $R_Y = f(F_Y, B_Y)$ , where  $F_X, F_Y$  represent the resources *available* at  $X, Y$ , and we have:

$$\frac{B_X}{B_Y} = \frac{f(F_X, B_X)}{f(F_Y, B_Y)}.$$

Our major condition is that  $f$  represents a ‘reliably arranged contingency’, and since our object is to characterise the simplest possible case of foraging, we may take  $f$  to be linear, thus (with linear scaling factors  $\gamma, \delta$ ):

$$\frac{B_X}{B_Y} = \frac{\gamma B_X F_X}{\delta B_Y F_Y}.$$

Now, this condition can only hold true if either (a)  $B_X = 0$  or (b)  $\gamma F_X = \delta F_Y$  (ie. if  $X$  and  $Y$  are equally ‘valuable’). In contrast, there are no such restrictions if matching is considered in terms of *available* resources:

$$\frac{B_X}{B_Y} = \frac{F_X}{F_Y}.$$

The truth of this condition is open to empirical assessment, it is not prefigured by any contingency between response and reward. Note that this argument does not discount matching to obtained resources as a noteworthy empirical phenomenon; it is indeed so, and has been recorded as such by an extensive body of work in experimental psychology in which  $f$  is specially constructed to reduce, or eliminate, the contingency between response and reward (see for example Herrnstein, 1997, and also chapter 9).

### The ideal free distribution

This section describes how many of the mathematical expressions pertaining to the IFD, as described in chapter 9, are related.

I start with an expression for  $N_{eat_i}$ , the total number of prey consumed in patch  $i$ , taken from Bernstein et al. (1988, 1991):

$$N_{eat_i} = \frac{QF_iF^*N_i^{(1-m)}}{1 + QhF_iF^*N_i^{-m}}, \quad (1)$$

where  $F_iF^*$  and  $N_i$  represent the quantity of prey items and foragers respectively ( $F_i$  being the fraction of total prey  $F^*$ ),  $Q$  (the ‘quest’ constant, units of square metres) is a measure of patch-independent forager search efficiency,  $m$  (dimensionless) is the interference constant, and  $h$  (dimensionless) is the handling time required for the consumption of each prey item, expressed as a fraction of unit time. To recover the number of prey consumed per forager - the intake rate  $W_i$  - in patch  $i$ , it is necessary to divide through by  $N_i$ :

$$W_i = \frac{QF_iF^*N_i^{-m}}{1 + QhF_iF^*N_i^{-m}}.$$

Now consider two patches  $A$  and  $B$ , with foragers distributed according to the IFD such that  $W_A = W_B$ , we can write:

$$\frac{QF_AF^*N_A^{-m}}{1 + QhF_AF^*N_A^{-m}} = \frac{QF_BF^*N_B^{-m}}{1 + QhF_BF^*N_B^{-m}},$$

which, after some simplification, becomes:

$$F_AF^*N_BN_A^{1-m} - F_BF^*N_A N_B^{1-m} = QhF^* (F_BF_A N_B^{1-m} N_A^{1-m} - F_BF_A N_B^{1-m} N_A^{1-m}) = 0.$$

Some simple algebra then reveals the relationship:

$$\frac{N_A}{N_B} = \left( \frac{F_A}{F_B} \right)^{\frac{1}{m}}, \quad (2)$$

which leads directly to the ‘generalised habitat matching rule’ (Fagen, 1987, equation 9.3):

$$\log\left(\frac{N_A}{N_B}\right) = \frac{1}{m} \log\left(\frac{F_A}{F_B}\right).$$

Finally, taking the total number of foragers to be  $N_T (= N_A + N_B)$ , it is possible to predict the value of both  $N_A$  and  $N_B$  directly. From equation 2 we have:

$$N_A = (N_T - N_A) \left( \frac{F_A}{F_B} \right)^{\frac{1}{m}},$$

therefore

$$N_A \left( 1 + \left( \frac{F_A}{F_B} \right)^{\frac{1}{m}} \right) = N_T \left( \frac{F_A}{F_B} \right)^{\frac{1}{m}},$$

which leads directly to:

$$N_A = \frac{N_T}{(10^{-c} + 1)}, \quad c = \frac{\log \frac{F_A}{F_B}}{m}.$$



### The problem of P-allocation

Deriving the IFD under P-allocation (chapter 9) can be difficult. The expected intake rate for the general form of the standing stock model under P-allocation may be written as follows:

$$E(W_i) = \frac{QF_iF^*N_i^{-m}}{1 + QhF^*N_i^{-m}}.$$

Therefore, the IFD condition that (expected) intake rates are equal across patches leads to:

$$\frac{QF_AF^*N_A^{-m}}{1 + QhF^*N_A^{-m}} = \frac{QF_BF^*N_B^{-m}}{1 + QhF^*N_B^{-m}}.$$

Unfortunately, this equation does not simplify as easily as the equivalent expression for the simple standing stock form (equation 9.2), and, most importantly, it will not reduce to the generalised habitat matching law (equation 9.3). Instead, the following relationship can be derived:

$$\frac{QhF^* + N_A^m}{QhF^* + N_B^m} = \frac{F_A}{F_B}.$$

Now, if  $QhF^* \ll N_A^m, N_B^m$ , it is possible to discard the  $QhF^*$  terms from the left side of this equation, in which case the generalised habitat matching law is recovered. Otherwise, however, no analytic solution for  $\frac{N_A}{N_B}$  can be recovered and therefore IFD predictions in such situations require the use of standard numerical techniques.

## Appendix C

This appendix presents the detailed results from the use of various sets of fixed parameters in the rule-based IFD model of chapter 9. Table 3 reveals the degree to which observed population distributions differ from the predictions of the IFD. Table 4 describes the extent to which isolated agents match to obtained resources. Finally, tables 5 and 6 illustrate the divergence from strict matching - to available and obtained resources - of embedded agents.

Error in population distribution with respect to IFD predictions										
	$\varepsilon_{opt}$	$\varepsilon = \bar{\varepsilon}$	$\varepsilon = .032$	$\varepsilon = .042$	$\varepsilon = .062$	$\varepsilon = .072$	$\varepsilon = \bar{\varepsilon}$	$\varepsilon = \bar{\varepsilon}$	$\varepsilon = \bar{\varepsilon}$	$\varepsilon = \bar{\varepsilon}$
	$\omega_{opt}$	$\gamma = \bar{\gamma}$	$\gamma = \bar{\gamma}$	$\gamma = \bar{\gamma}$	$\gamma = \bar{\gamma}$	$\gamma = \bar{\gamma}$	$\gamma = .227$	$\gamma = .327$	$\gamma = .527$	$\gamma = .627$
$\varepsilon$										
$m = 1.0, CA$	0.00	0.01	0.01	0.01	0.02	0.02	0.01	0.01	0.01	0.01
$m = 0.3, CA$	0.01	0.05	0.03	0.04	0.06	0.07	0.06	0.05	0.05	0.05
$m = 1.0, PA$	0.03	0.06	0.03	0.04	0.06	0.07	0.05	0.05	0.05	0.05
$m = 0.3, PA$	0.17	0.20	0.19	0.19	0.22	0.22	0.21	0.21	0.21	0.21
$\omega$										
$m = 1.0, CA$	0.01	0.01	0.01	0.01	0.02	0.02	0.01	0.01	0.02	0.02
$m = 0.3, CA$	0.02	0.03	0.02	0.03	0.03	0.04	0.02	0.03	0.02	0.03
$m = 1.0, PA$	0.01	0.05	0.08	0.06	0.04	0.04	0.03	0.04	0.06	0.07
$m = 0.3, PA$	0.03	0.07	0.04	0.06	0.08	0.09	0.06	0.07	0.07	0.07

Table 3: Error measures indicating the extent to which equilibrium population distributions differ from those predicted by the IFD (equation 9.4). Error measures are provided for 8 conditions defined by all combinations of sampling strategy ( $\varepsilon$  or  $\omega$ ), interference level ( $m = 1.0$  or  $m = 0.3$ ), and C-allocation or P-allocation (CA or PA). Ten parameter sets are investigated; near-optimal, fixed ( $\bar{\varepsilon} = 0.052, \bar{\gamma} = 0.427$ ) and 4 variations of both  $\varepsilon$  and  $\gamma$  from the fixed set. Each error measure is calculated as the average normalised deviation from the IFD prediction over 9 resource distributions ( $F_A = 0.1$  to  $F_A = 0.9$ ), with each deviation calculated from the average of 30 population distributions.

Isolated agent matching to obtained resources										
	$\epsilon_{opt}$ $\omega_{opt}$	$\epsilon = \bar{\epsilon}$ $\gamma = \bar{\gamma}$	$\epsilon=.032$ $\gamma=\bar{\gamma}$	$\epsilon=.042$ $\gamma=\bar{\gamma}$	$\epsilon=.062$ $\gamma=\bar{\gamma}$	$\epsilon=.072$ $\gamma=\bar{\gamma}$	$\epsilon=\bar{\epsilon}$ $\gamma=.227$	$\epsilon=\bar{\epsilon}$ $\gamma=.327$	$\epsilon=\bar{\epsilon}$ $\gamma=.527$	$\epsilon=\bar{\epsilon}$ $\gamma=.627$
$\epsilon$										
$m = 1.0, CA$	0.94	0.87	0.91	0.89	0.85	0.83	0.87	0.87	0.87	0.87
$m = 0.3, CA$	0.95	0.87	0.92	0.90	0.85	0.83	0.87	0.87	0.87	0.87
$m = 1.0, PA$	0.67	0.61	0.65	0.63	0.59	0.57	0.62	0.60	0.60	0.61
$m = 0.3, PA$	0.70	0.61	0.65	0.63	0.59	0.58	0.60	0.61	0.62	0.60
$\omega$										
$m = 1.0, CA$	0.91	0.93	0.95	0.94	0.92	0.90	0.93	0.93	0.93	0.93
$m = 0.3, CA$	0.95	0.93	0.95	0.94	0.92	0.91	0.93	0.93	0.93	0.93
$m = 1.0, PA$	0.88	0.87	0.91	0.88	0.86	0.84	0.90	0.88	0.86	0.87
$m = 0.3, PA$	0.93	0.87	0.91	0.89	0.85	0.84	0.90	0.88	0.86	0.86

Table 4: Extent to which isolated agents match to obtained resources. Eight conditions and 10 parameter sets are presented as in table 3. Each entry represents the slope of the line of best fit relating response frequency to resources obtained. This measure is equivalent to the ‘bias’ of equation 9.8, with 1.0 representing strict matching. Each entry is derived from the best-fit line of a scatter plot of 270 points created by testing the agent 30 times at each of 9 resource distributions, ( $F_A = 0.1$  to  $F_A = 0.9$ ).

Error in embedded matching to available resources with respect to strict matching										
	$\epsilon_{opt}$ $\omega_{opt}$	$\epsilon = \bar{\epsilon}$ $\gamma = \bar{\gamma}$	$\epsilon=.032$ $\gamma=\bar{\gamma}$	$\epsilon=.042$ $\gamma=\bar{\gamma}$	$\epsilon=.062$ $\gamma=\bar{\gamma}$	$\epsilon=.072$ $\gamma=\bar{\gamma}$	$\epsilon=\bar{\epsilon}$ $\gamma=.227$	$\epsilon=\bar{\epsilon}$ $\gamma=.327$	$\epsilon=\bar{\epsilon}$ $\gamma=.527$	$\epsilon=\bar{\epsilon}$ $\gamma=.627$
$\epsilon$										
$m = 1.0, CA$	0.02	0.02	0.01	0.02	0.02	0.02	0.02	0.02	0.01	0.01
$m = 0.3, CA$	0.14	0.11	0.13	0.11	0.11	0.10	0.11	0.11	0.11	0.11
$m = 1.0, PA$	0.04	0.05	0.04	0.05	0.06	0.07	0.05	0.05	0.05	0.05
$m = 0.3, PA$	0.03	0.04	0.03	0.03	0.05	0.06	0.05	0.04	0.04	0.04
$\omega$										
$m = 1.0, CA$	0.05	0.01	0.02	0.01	0.01	0.01	0.01	0.01	0.01	0.01
$m = 0.3, CA$	0.15	0.14	0.14	0.14	0.14	0.13	0.14	0.14	0.14	0.14
$m = 1.0, PA$	0.03	0.05	0.07	0.06	0.03	0.04	0.03	0.04	0.05	0.07
$m = 0.3, PA$	0.12	0.09	0.11	0.10	0.08	0.07	0.10	0.09	0.09	0.09

Table 5: Error measures indicating the extent to which embedded agent behaviour differs from strict matching to available resources. Eight conditions and 10 parameter sets are presented as in table 3. Each error measure is calculated as the average normalised deviation (of a randomly selected single agent) from strict matching over 9 resource distributions ( $F_A = 0.1$  to  $F_A = 0.9$ ), with each deviation calculated from the average of 30 randomly selected agents from each population.

Embedded agent matching to obtained resources										
	$\varepsilon_{opt}$ $\omega_{opt}$	$\varepsilon = \bar{\varepsilon}$ $\gamma = \bar{\gamma}$	$\varepsilon=.032$ $\gamma=\bar{\gamma}$	$\varepsilon=.042$ $\gamma=\bar{\gamma}$	$\varepsilon=.062$ $\gamma=\bar{\gamma}$	$\varepsilon=.072$ $\gamma=\bar{\gamma}$	$\varepsilon=\bar{\varepsilon}$ $\gamma=.227$	$\varepsilon=\bar{\varepsilon}$ $\gamma=.327$	$\varepsilon=\bar{\varepsilon}$ $\gamma=.527$	$\varepsilon=\bar{\varepsilon}$ $\gamma=.627$
$\varepsilon$										
$m = 1.0, CA$	0.97	0.94	0.97	0.95	0.93	0.91	0.93	0.94	0.94	0.95
$m = 0.3, CA$	0.96	0.90	0.95	0.92	0.89	0.87	0.91	0.90	0.90	0.91
$m = 1.0, PA$	0.89	0.78	0.85	0.81	0.74	0.72	0.79	0.78	0.78	0.79
$m = 0.3, PA$	0.76	0.64	0.70	0.67	0.62	0.61	0.64	0.65	0.64	0.65
$\omega$										
$m = 1.0, CA$	0.99	1.00	1.00	1.00	1.00	0.99	0.99	0.99	1.00	1.00
$m = 0.3, CA$	0.97	0.95	0.97	0.96	0.95	0.94	0.95	0.95	0.95	0.95
$m = 1.0, PA$	0.95	1.10	1.20	1.10	1.10	1.10	1.10	1.10	1.10	1.20
$m = 0.3, PA$	0.94	0.90	0.93	0.91	0.88	0.86	0.91	0.90	0.89	0.89

Table 6: Extent to which embedded agents match to obtained resources. Eight conditions and 10 parameter sets are presented as in table 3. Each entry represents the slope of the line of best fit relating response frequency to resources obtained, calculated as in table 4, but using a randomly selected embedded agent rather than an isolated agent.