

On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical Aspects

Ezequiel Alejandro Di Paolo

Submitted for the degree of DPhil

University of Sussex

January, 1999

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

During the period I spent in Sussex I have learnt a lot, explored a lot and met a lot of interesting people who were always keen to take each other seriously in a relaxed environment. Many of the ideas presented in this thesis have matured within this context.

I would like to thank my supervisor during the last three years, Phil Husbands, for providing guidance, help and being always open to new ideas no matter how strange they may have sounded initially. I would also like to thank Inman Harvey, who often played a role of supervisor himself by being always available to discuss problems and make useful suggestions which often were strangely radical and down-to-earth at the same time.

Seth Bullock and Jason Noble contributed a lot to this thesis. The many opportunities in which we discussed our work and the work of others have often been key moments in the development of my research. I am very grateful to them.

I must also thank the help I got from other people inside and outside COGS in the form of comments on my work, discussions and suggestions: Guillermo Abramson, Ron Chrisley, Dave Cliff, Kerstin Dautenhahn, Joe Faith, Ronald Lemmen, Matt Quinn, Lars Risan, Anil Seth, Oli Sharpe and Mike Wheeler. I am also grateful for comments and suggestions made by anonymous referees on my journal and conference publications.

By being passionate about their own work, some people inspired me to approach ideas of which I was completely ignorant when I first came to Sussex. In this category I must especially thank Inman, Ronald and Mike.

This research could not have been carried out without the support of my family in Argentina, my parents, my brothers, my sister and my in-laws. This support has so many dimensions I can hardly begin to explore this space but I must especially thank their love and help.

Among my friends I cannot avoid mentioning the support received from Alberto Baruj whose selfless help in taking care of our house in Bariloche provided the fundamental peace of mind I needed to carry on. Thanks, Beto.

Thanks also to those two undending sources of canine motivation: Psycho and Trash. Without them life would have been much easier and extremely boring.

Finally, this thesis could not have even started were it not for the encouragement, love, confidence, intellectual support and unique character of a single person: Marina, my wife. She has made a lot of sacrifices which I will never be able to repay entirely although I hope she takes this thesis as a first installment.

During the last three years I received financial support from different sources. I am grateful to the Argentine Ministry of Education, the National Atomic Agency and the *Consejo de Investigaciones Científicas y Técnicas de la República Argentina*. I am also grateful for an Overseas Research Student Award I received during 1997/1998 from the Committee of Vice-Chancellors and Principals. My wife contributed enormously to our financial viability as a result of her personal effort; for this I will be permanently in debt with her.

Preface

Some of the chapters within this thesis have benefited from comments made by anonymous reviewers as a result of their being published elsewhere.

In particular, parts of chapter 6 have appeared as (Di Paolo, 1997b) and (Di Paolo, 1997a). Shorter versions of chapters 7 and 9 were previously published as (Di Paolo, 1998a) and (Di Paolo, 1998b) respectively. Also, parts of chapter 9 will appear as (Di Paolo, 1999).

The author is extremely thankful to all the reviewers for their criticisms and comments on how to improve the manuscripts.

On the Evolutionary and Behavioral Dynamics of Social Coordination: Models and Theoretical Aspects

Ezequiel Alejandro Di Paolo

Abstract

An exploration is presented of the interplay between the situated activity of embodied autonomous organisms and the social dynamics they constitute in interaction, with special emphasis on evolutionary, ecological and behavioral aspects. The thesis offers a series of theoretical and methodological criticisms of recent investigations on the biology of social behavior and animal communication. An alternative theoretical framework, based on a systemic theory of biological autonomy, is provided to meet these criticisms and the elaboration of the corresponding theoretical arguments is supported by the construction and analysis of mathematical and computational models.

A game of action coordination is studied by a series of game-theoretic, ecological and computational models which, by means of systematic comparisons, permit the identification of the evolutionary relevance of different factors like finite populations, ecological and genetic constraints, spatial patterns, discreteness and stochasticity. Only in an individual-based model is it found that cooperative action coordination is evolutionarily stable. This is due to the emergence of spatial clusters in the spatial distribution of players which break many of the in-built symmetries of the game and act as invariants of the dynamics constraining the path of viable evolution.

An extension to this model explores other structuring effects by adding the possibility of parental influences on phenotypic development. The result is a further stabilization of cooperative coordination which is explained by the presence of self-promoting networks of developmental relationships which enslave the evolutionary dynamics.

The behavioral aspects involved in the attainment of a coordinated state between autonomous systems are studied in a simulated model of embodied agents coupled through an acoustic medium. Agents must locate and approach each other only by means of continuous acoustic signals. The results show the emergence of synchronized rhythmic signalling patterns that resemble turn-taking which is accompanied by coherent patterns of movement. It is demonstrated that coordination results from the achievement of structural congruence between the agents during interaction.

Submitted for the degree of DPhil

University of Sussex

January, 1999

Contents

1	Introduction	1
1.1	Overview of the thesis	3
1.1.1	Towards a systemic view of social behavior	3
1.1.2	Historical processes	4
1.1.3	Methodological issues	5
1.1.4	Modelling the evolution of cooperative coordination	6
1.1.5	An evolutionary role for social influences on development	8
1.1.6	Coordination without non-shared information	9
1.1.7	Coordination, structural congruence and embodiment	10
1.2	Summary of contributions	11
2	Towards a biology of social behavior: a systemic perspective	14
2.1	The pragmatic need for a biological grounding	15
2.2	Functional and operational approaches	16
2.3	Confusing the domains of discourse: the case of animal communication	20
2.3.1	Selective advantages as a defining feature of communication	20
2.3.2	Some ill-defined terms	21
2.3.3	Practical consequences for research	24
2.4	A systemic perspective on social behavior	26
2.4.1	Basic concepts	26
2.4.2	Coordination	29
2.4.3	Changing the practical issues of interest: animal communication again	32
2.5	Summary	36
3	Historical processes	38
3.1	From homogeneous time to historical time through self-organization	38
3.1.1	Constraints	39
3.1.2	Invariants	41
3.2	Different manifestations of history	43
3.2.1	Trails on grass and Pask's artificial ear	43
3.2.2	Polya's urn scheme	44
3.2.3	Evolution	45
3.2.4	Embodiment, culture and <i>habitus</i>	49
3.3	Explanatory value of spontaneous invariants	51
3.4	Summary	52

4	Methodological issues in the use of formal models and computer simulations	53
4.1	A pragmatic approach to formal modelling	54
4.2	A scientific use for computer simulations	59
4.2.1	Science and instruments	61
4.2.2	Tools in scientific activity	61
4.2.3	Thought experiments	64
4.2.4	Simulations as tools	65
4.2.5	Integrating simulations in biology	67
4.2.6	Exploring, experimenting, explaining	70
4.3	Summary	72
5	A game of action coordination: mathematical models	75
5.1	The game	76
5.2	Game-theoretic analysis	78
5.2.1	Strategies	79
5.2.2	Infinite population: static analysis	81
5.2.3	Infinite population: dynamical analysis	82
5.2.4	Finite population: dynamical analysis	85
5.3	Ecological coupling	88
5.4	Genetic constraints	89
5.5	Continuous spatial model with finite locality	91
5.5.1	The model	93
5.5.2	Pattern formation	96
5.5.3	Dynamics	99
5.5.4	Effect of density thresholds, move toward discreteness.	102
5.5.5	Summary for continuous spatial model	106
5.6	Summary: broken symmetries	108
6	A game of action coordination: individual-based model.	110
6.1	The model	111
6.2	Results	114
6.3	Spatial patterns	116
6.4	Correlation between cluster size and coordination level	119
6.5	The role of kin selection	120
6.6	Clusters as structuring invariants of the evolutionary dynamics	124
6.6.1	Fluctuations and phenotypic aspects of spatial position	125
6.6.2	Broken symmetries within clusters	127
6.6.3	Coordination explained	128
6.6.4	Correlation explained	131
6.7	Conclusions and analogies	132

7	Social influences on development and coordinated behavior: an assessment	137
7.1	Social influences on development: evidence and consequences	138
7.2	A model of socially influenced development	142
7.3	Computer model	145
7.4	No development	148
7.4.1	No Space.	148
7.4.2	Space.	148
7.5	Results with social development	150
7.5.1	No Space.	150
7.5.2	Space.	150
7.6	Networks of social influences on development	151
7.7	Discussion	155
8	Coordination without information	159
8.1	Computer models of communication and the informational bias	159
8.2	Coordination of recursive actions	162
8.3	Discussion	166
9	Rhythm, entrainment and congruence in acoustically coupled agents	168
9.1	Acoustic coupling	170
9.2	An example of coordination through acoustic interactions	171
9.3	Methods	172
9.4	The model	173
9.4.1	Sound	173
9.4.2	Bodies	173
9.4.3	Genetic algorithm	176
9.5	Results	177
9.5.1	Approaching behavior	178
9.5.2	Entrainment, turn-taking and structural congruence	181
9.6	Discussion	185
10	Conclusions	190
10.1	Recapitulation and contributions	191
10.2	The future	195
	Bibliography	197

List of Figures

2.1	Illustration of the concept of coordination.	30
2.2	A special case of coordination through interaction	31
3.1	Probability of drawing a black ball from Polya's urn	45
4.1	Direct explanation.	70
4.2	Indirect explanation.	71
5.1	Cooperation between strategies	80
5.2	Tetrahedron diagram	83
5.3	Trajectories for dynamical analysis	84
5.4	Non-zero eigenvalues for infinite dynamics	85
5.5	Trajectories for perturbed dynamics	86
5.6	Trajectory for finite population	87
5.7	Trajectory for ecological model	89
5.8	Trajectory for ecological model with genetic constraints	91
5.9	Stability of homogeneous solution	96
5.10	Dispersion relation	97
5.11	Total number of players in spatial model	100
5.12	Spatial patterns	101
5.13	Clusters and difference in energy intake for each role for strategy 1.	103
5.14	Invasion patterns.	104
5.15	Invasion patterns for different density thresholds.	105
6.1	Position of agents in the arena	112
6.2	Evolution of successful action	115
6.3	Coordination level (R_{12}) for different values of c	116
6.4	Formation of clusters.	117
6.5	Qualitative distribution of agents and energy in a cluster	118
6.6	Coordination vs. cluster size: examples	119
6.7	Average relatedness vs. c	123
6.8	Relatedness vs. cluster size	124
6.9	Probability of remaining in cluster after 5 generations	125
6.10	Position of ancestors	126
6.11	Differences in role frequency within a cluster	127
6.12	Critical frequency for extra gain to be equal to fluctuations	129
6.13	Differences in role frequency.	131
7.1	Two extreme cases of developmental relationships	141

7.2	Scheme of abstract developmental process.	143
7.3	Position of agents in the arena	146
7.4	Cooperative relations between strategies.	147
7.5	Population constitution, no space and no development	149
7.6	Population constitution for spatial game with no development, $c = 0.6$	150
7.7	Population constitution, no space with development	151
7.8	Coordination level vs. c for spatial games with and without development.	152
7.9	Population constitution for spatial game with development, $c = 0.6$	153
7.10	A self-promoting developmental network	153
7.11	Population constitution for spatial game with development, $c = 0.8$	155
7.12	Developmental structures	156
7.13	Cooperative relationships in developmental structures	157
7.14	Time variation for cooperative and non-cooperative strategies in secondary network	158
8.1	Examples of the evolution of dialogic coordinated activity	164
8.2	Dialogic activity vs. c	165
9.1	Body of agents and paths of acoustic signals.	174
9.2	Trajectories of approaching agents	177
9.3	Neuronal structure	178
9.4	Emitted signal and sensory gain	179
9.5	Power spectra for signal, orientation and intensity difference in each sensor	180
9.6	Trajectory for constant source with a fixed position relative to the agent	181
9.7	Signalling behavior of interacting agents and power spectra.	182
9.8	Motion of agents during period of coordination	183
9.9	Signalling behavior in presence of periodic beacon.	184
9.10	Signalling behavior in presence of imitative beacon.	185
9.11	Embedded time-delayed dynamics of inter-neurons	186
9.12	Achievement, loss and regaining of coordination	187

“Pourquoi la synergie n’existerait-elle pas entre différents organismes, si elle est possible à l’intérieur de chacun?”

Maurice Merleau-Ponty (1964, p. 187)

Chapter 1

Introduction

In the late nineteen-forties and early fifties, William Grey Walter was pioneering a particular style of research into adaptive behavior using artificial systems. Many of today's lines of research in situated robotics and computer simulations of adaptation have partly inherited their interests and methods from his work. In two articles in the journal *Scientific American* (Walter, 1950, 1951) and later in a book called *The living brain* (Walter, 1953), he described how life-like behavior could be realized in rather simple electro-mechanical devices which operated autonomously in a real environment. These robots, or "tortoises" as he called them, would propel themselves describing cycloidal paths on a flat surface in search of moderate sources of light. Once their photoelectric sensors detected a source of light, their behavior would change and they would move towards the location of the source. When the light became too intense their mode of behavior would change again, this time to an avoidance pattern. The robots would also avoid heavy obstacles and steep vertical gradients, (see also Holland, 1996). W. Grey Walter vividly described how these basic patterns were simple consequences of the involved mechanisms in combination with properties of the environment and how they could give rise to more complicated behaviors like avoiding becoming stuck in the presence of conflicting stimuli and even a rudimentary form of energetic stability by entering an illuminated recharging station when batteries were running low. The overall life-like behavior and curious "personality" of the robots make them deserve the name *Machina Speculatrix*.

Without any other apparent aim than the possibility of being able to check when the robots changed from one mode of behavior to another, W. Grey Walter wired a headlamp into their steering circuit. The light bulb would turn on while the robot was performing its exploratory behavior and turn off when the robot's photoelectric sensor detected a moderate light. The robot would not be able to sense its own light directly, but it could do so indirectly through an interaction with its environment; for instance, approaching a mirror or white surface. Such interactions elicited very interesting behavioral patterns.

But perhaps the most fascinating experiment of all occurred when two of the robots (Elmer and Elsie) were placed together on the same floor in the absence of any other sources of light. The interaction between the different modes of behaviors and the headlamps when these were mutually detectable would cause the robots to perform curious dancing patterns. A robot in the

exploratory mode would have its headlamp on until its sensors detected the light coming from its partner. Confronted with each other's headlamp each robot would engage in phototactic behavior and, as a consequence, they would extinguish the sources of each other's attraction thus provoking a return to the previous exploratory mode. The robots would approach toward each other and stop in an intermittent pattern while their headlamps would flicker until the light of one of them would move beyond the range of the other one's sensor and so they would "lose interest".

This simple experiment, explainable in terms of wires, light bulbs and mechanical bodies, poses some nontrivial questions about the nature of social behavior. Is it possible to classify Elmer's and Elsie's behavior, simple as it is, as social? When *is* a behavior social? Can social patterns of behavior be explained in the same way as any other case of adaptive behavior or is a different vocabulary required for doing the job? Can any knowledge be gained by studying social behavior using artificial systems like robots or simulated agents? Can these systems be themselves social? Today, many inheritors of W. Grey Walter's tradition still find these questions pressing.

Issues like these have inspired the work presented in this thesis. Its main motivation is to address some basic questions about social behavior using modelling techniques which, much as W. Grey Walter's experiments did, will work as tools of theoretical inquiry by trying to instantiate artificially the very patterns in need of explanation. In contrast with W. Grey Walter's robots, the techniques used in this thesis are mainly formal models and computer simulations.

The thesis will be limited to the study of basic forms of social behavior such as those that may be observed in non-human animals and, even in this case, only very simple ones. Many issues regarding the complexities of *human* social behavior will be sidestepped, although some of the conclusions arrived at will surely offer significant insights into human social activity and hopefully open a door for further investigation using similar methods.

Clearly, it makes no sense to try to answer questions about social behavior in a theoretical vacuum. A perspective, a point of departure, some form of grounding must be provided in order to explore these issues. It is possible to find such a theoretical standpoint in biology. After all, social phenomena can be found, as yet, exclusively in living systems.

But biology is not a timeless discipline. It has its own concerns and forces its own issues upon the researcher interested in the nature of social behavior. Thus, if one seeks a guiding framework in contemporary biology some additional questions about social behavior will become unavoidable. How does the evolution of social behavior fit within the current theoretical orthodoxy? What is the adaptive function of social behavior? Do social interactions highlight some aspects of an evolutionary process which are often simplified when studying other forms of behavior?

At the same time, current biology also tends to impose its own way of thinking, and while its subject is deemed appropriate for grounding an investigation about the nature of social behavior, it is possible that its current methods may not be so. Thus some care must be taken not to embrace a theoretical stance without questioning its potential and its limitations.

The project described in this thesis can be placed at the intersection between biology and cognitive science. It addresses issues of interest for both disciplines. It makes use of relatively novel research methodologies which are currently being adopted by researchers in both fields such as theoretical models instantiated in complex computer simulations and used in conjunction with more traditional analytical techniques. But at the same time, the project breaks with some

traditions in biology and cognitive science which are not totally unrelated. The project abandons explanatory monism, methodological individualism and functional analysis and embraces explanatory complementarity, situatedness, complexity and dynamics in a pragmatic approach. One of the main purposes of this thesis is to argue in favour of this break with tradition and to show its practical consequences for research in the form of concrete investigations about social behavior.

These investigations address issues like the evolution of cooperative coordination, the relevance of ecological situatedness, the dynamics of historical processes, the constraints imposed on behavior and on evolution by social influences on development, the relevance of embodiment in shaping patterns of behavior and the emergent features of coordination between dynamical systems in interaction. Although each investigation is meant as an individual contribution to the issues it addresses, they are vertically connected in the sense that they focus on the different levels affecting social behavior from global evolutionary patterns to the operational details of the mechanism giving rise to social coordination. The investigations are also thematically connected in subtle ways, since all of them will arrive at similar conclusions regarding the value of studying complex processes in a situated manner. Evolution does not occur as a clean universal process but is embedded in a web of ecological and historical dynamics. Coordination does not happen between preprogrammed general problem solvers, but arises from the messy constraints imposed by embodiment and situatedness as well as history.

In addition, the investigations are also methodologically connected since they use a relatively novel way of building computer simulations and exploring them to draw conclusions about theoretical issues. These techniques are aimed at understanding how theories work in artificial instantiations of the phenomena they try to explain. Understanding what goes on in an artificial world can provide information about the extent and suitability of theories about the natural world. In this way, the methodologies used in this thesis are the direct descendants of the research methodologies used by W. Grey Walter who, by understanding how his robots operated, was able to think more clearly about the behavior of natural organisms.

1.1 Overview of the thesis

The thesis is divided into two main parts. The next three chapters are intended to provide a conceptual and methodological background within which the investigations presented in the following five chapters can be framed.

1.1.1 Towards a systemic view of social behavior

Many of the above questions regarding the nature of social behavior are too general to deserve straightforward answers. As said before, it makes little sense to try to address these issues in the lack of a well-defined theoretical perspective. Chapter 2 sets the basis for a systemic approach to the biology of social behavior in autonomous systems. The pragmatic need for grounding a theory of social behavior in biology is established without implying in any way a form of reductionism of social phenomena. The intention is to constrain the vocabulary used in a theory of social behavior to what is biologically possible and avoid ungrounded theorizing from the start.

Here is where the project faces its first problem. Is the type of scientific discourse used currently in biology adequate to perform such a constraining function? In order to answer this question a distinction between operational and functional modes of scientific descriptions and explanations is introduced and the relation between the two modes examined. It is argued that none of these modes is intrinsically better than the other since they serve different purposes and can well complement each other. However, the distinction should be kept in order to avoid slipping carelessly from one type of discourse into the other. In order to reveal the dangers of this possibility the case of animal communication is examined. It is shown that the traditional biological approach relies on functional considerations, such as selective advantages, in order to define communication but that it also intends to use the definition in a pseudo-operational manner. As a consequence, communication remains poorly defined and full of internal contradictions, some of which will be carefully exposed.

Functional discourse may provide hints regarding at which level of description an operational description may be pitched, but operational discourse relates differently to functional statements by constraining them to what is operationally viable. In view of the reasons for wanting to ground social behavior in biology, an operational approach is therefore deemed as the most appropriate starting point.

Fortunately, such a perspective can be found in a fairly developed form in the theories of autopoiesis and biological autonomy, (Maturana & Varela, 1980; Varela, 1979). The relevant terminology is presented and discussed. It is shown that phenomena like communication can be much better characterized in this way by introducing key concepts such as autonomy, structural coupling, coordination and structural congruence. The change of perspective introduces important shifts of interest by questioning issues which were ignored by the traditional approach. At the same time, many of the questions traditionally addressed from a functional perspective (such as the evolutionary stability of honest signals, the nature of conflicting interactions, etc.) can still be studied from the systemic point of view perhaps after some reformulation.

1.1.2 Historical processes

The basic theoretical perspective for framing questions regarding socially coordinated behavior is complemented in chapter 3 with a general description of what constitutes a historical process. The idea behind this exposition is to prepare the reader for some of the issues addressed in the investigations of the second part of the thesis.

A process is described as historical if some form of temporal heterogeneity is introduced during its unfolding. It is shown that this is possible if the constraints of the process change as a consequence of its internal dynamics or as a consequence of coupling with other processes. The meaning of the term “constraint” is clarified as well as a way of distinguishing variations in the constraints from the process itself by specifying what is meant by an invariant of the process.

In this way, the word “historical” is used in a special sense. It does not suffice for a process to be subject to contingencies or to be susceptible to be explained in historical terms for it to be historical in the sense intended. Rather, historical processes lie roughly at the intersection between these classes and the class of processes which exhibit some form of self-organization. The reason for this is that a systematic coupling between the dynamics of a process and variations of its

own conditions of realization often leads to the formation of spontaneous organizational invariants which tend to perpetuate themselves. These invariants are constituted by the dynamics of the process and are at the same time constituting factors for the future dynamics.

Some examples of historical processes are discussed at length, including the important case of biological evolution which is relevant for the studies presented in chapters 5, 6 and 7.

Spontaneous invariants may provide the researcher with a powerful explanatory tool for understanding what goes on in a complex historical process. They set a frame of reference from which global dynamical patterns may be understood as well as their influence on local or microscopic patterns. This power is demonstrated in the models of the second part which address evolutionary, ecological and social dynamics and which are hard to explain otherwise.

1.1.3 Methodological issues

The usefulness of the theoretical issues presented in chapters 2 and 3 has yet to be demonstrated practically. The intention of the second part of the thesis is to present models which address specific questions regarding social coordination. On the one hand, these models are interesting in themselves and are intended as contributions to the scientific issues they address, on the other hand, the models are also interesting as demonstrations of how the theoretical ideas work.

Due to the complexity of the issues studied, a combination of mathematical models and computer simulations will be used. There is a long-standing tradition of formal modelling in biology and other sciences but the advent of new modelling techniques, such as individual-based computer simulations, may call for a re-thinking of this tradition, especially from a methodological point of view. Is a computer simulation just a complicated mathematical model or is it something else? Does it play a similar scientific role? What is the value of a computer simulation? How does it relate to the rest of scientific practice? How should a simulation be built/used?

These are important questions that a thesis like this cannot afford to ignore. Computer simulations are perhaps the most flexible tool of theoretical inquiry currently available. However, this does not mean that they are necessarily easier to use or to understand than other tools. In principle, very complex situations may be built into a computer model and many of these complexities cannot be subject to treatment using a purely mathematical approach. But this same flexibility may play to the disadvantage of computer simulations as an acceptable medium of scientific research. This is mainly because researchers often tend to build into a computer model situations that are radically different from those previously studied and also because the range of applicability of the conclusions of very complicated models tends to be rather narrow. The first reason introduces a communication rupture between researchers who use mathematical models exclusively and researchers interested in addressing the same problems using computer simulations. The work of the latter may be too complex for direct comparison with previous models and so for its value to be clearly assessed. So much so, that in certain cases it may not even be easy to understand *what* has actually been modelled into a computer simulation and what conclusions to derive from it. The second reason indicates that even if complex situations can be replicated in a computer simulation, there is no guarantee that this replication will render the phenomenon more understandable. Mathematical models can *explain* and *communicate* by abstracting and simplifying complex phenomena and offering a fairly universal language with which to express their conclusions. Can the

same be said about computer simulations?

These issues will constitute the theme of chapter 4 which is intended to raise concern about the use of formal models in general and computer simulations in particular. The chapter, however, also offers positive contributions towards the development of these issues. In the general case, a pragmatic attitude for the use of formal models is proposed. This attitude stipulates that, when modelling complex phenomena, 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 (which, by the way, may not even be available). In this manner, the use of formal models of different granularity is recommended as a way of understanding the role played by relevant factors following the performance of systematic comparisons between models.

The chapter proceeds to addressing the question of how to use computer simulations in a way that they become integrated into the rest of scientific practice. Of particular interest will be the question of how simulations relate to scientific activities such as formulating and testing hypotheses and explanations. It is proposed that the best perspective for dealing with this issue is given by a framework that describes the role of tools and techniques in scientific activity. A description of such a framework is offered based mainly on Deweyan and Heideggerian perspectives on the role of tools in skilful activity. Science is characterized as a network of practices which follow certain core conditions of normativity. The question of how to use simulations becomes the question of how to integrate a new tool into an existing network of practices without changing it so much that people would not be willing to keep calling it science. In view of this, of special relevance are the so-called conceptual or theoretical tools, in particular thought experiments.

Previous attempts at addressing similar issues will be discussed and criticized and a new possibility for integrating simulations in scientific practice will be offered. This possibility does not constitute a methodology in the sense of a set of rules that specify how simulations should be built. It rather describes the working stages that take place *after* the simulation has run successfully such as formulating hypotheses and performing crucial experiments in order to explain the observed patterns *within* the model and then relate this explanatory structure metaphorically to analogous explanatory structures in the case of natural phenomena. In this way, simulations are used as thought experiments aimed at questioning existing theories rather than aimed at reproducing natural phenomena directly.

The two positive contributions of this chapter are implicitly demonstrated in the concrete investigations that follow after it. Chapters 5, 6 and 7 show the comparative approach at work by studying a basic evolutionary scenario using increasingly complex models. Chapters 6, 7, 8 and 9 present different models using computer simulations where the proposed usage is put to the test.

1.1.4 Modelling the evolution of cooperative coordination

As implied above, the set of questions regarding social behavior is enormous and includes issues that are interesting to evolutionary biologists, ecologists, ethologists, cognitive scientists, economists, sociologists, psychologists, etc. The investigations presented in this thesis will address only a fraction of these issues starting from the evolutionary end of the spectrum, where behavioral complexity is sacrificed in order to be able to address large-scale evolutionary and ecological patterns, and then moving towards the behavioral or cognitive end of the spectrum where

behavioral dynamics and mechanisms are studied in detail.

In the first investigation, social behavior is studied from an evolutionary point of view. Many of the complex theoretical issues regarding social coordination are consequently simplified in order to address evolutionary questions. One such question in particular has received much attention in theoretical biology and it cannot be ignored in the present context. This is the issue of cooperative social behavior in situations where individual interests are in conflict.

Much interesting discussion has been generated lately regarding this issue, especially since it becomes problematic from the point of view of an evolutionary theory that rejects group selection in favour of maximization of individual reproductive success. Why should an individual organism cooperate with another if it is not in its best interest? Why should an established system of communication not be hijacked by manipulative individuals who cheat to obtain a personal benefit?

Indeed, much of the theoretical approach presented in chapter 2 regarding the systemic view on social behaviors can be criticized because, to the degree in which it is presented in that chapter, it remains silent about functional issues and therefore does not touch on the question of whether coordination would not demand behavioral patterns that would tend to be selected out because they tend reduce reproductive potential.

In response to this, different “solutions” have been offered that remain within the theoretical orthodoxy of neo-Darwinism. All these solutions take the form of a negation of the premises of the problem. For instance, Krebs and Dawkins (1984) suggest that in fact all communication *is* a form of manipulation and that if certain organisms seem to let themselves be manipulated by responding to signals or stimuli in a certain way, they do so because most of the time it pays them to respond in that way. Other proposals specify that cheating is prevented because of the inherent costs implied in the actual production of a cheating signal, (Zahavi, 1975, 1977). The only affordable way to behave as a cheater is by being in fact truthful. Another possibility suggests that when an individual behaves towards the benefit of another one but towards its own detriment, this is because there is a genetic similarity between them over and above the average genetic constitution of the population such that by behaving altruistically an individual helps in the propagation of copies of its own genetic material, which means that it does not *really* behave altruistically, (Hamilton, 1964; Michod, 1982).

But do these possibilities cover all cases? Is it not possible that an ecologically situated evolutionary process may be coupled with other processes so that its own dynamics are constrained along certain paths? Could not the dynamics of the totality of the processes involved explain cooperative coordination in the presence of conflict of interest as a natural consequence of situated evolution? Such a possibility would not deny the *premises* of the problem but it would question the universal and de-contextualized applicability of the theory that gives rise to the problem. It would, in other words, question the *logic* of the problem.

In order to investigate these issues, a simple action-response game with conflict of interest is used as a basic evolutionary scenario. The degree of conflict can be adjusted using a continuous parameter. Under what circumstances would one expect cooperative coordination of action and responses to evolve? In chapters 5 and 6, this question is approached using the pragmatic methodology for the use of formal models described in chapter 4.

Chapter 5 presents a series of mathematical models, including a game-theoretic approach and

extensions to it. These models study the differences introduced in the results when factors such as finite populations, discreteness, ecological dynamics and genetic constraints are included. A more complex model is used for relaxing the assumption of unstructured interactions between players by introducing a spatial dimension. This continuous model with finite spatial locality shows how patterns may be formed in space and how many of the inbuilt symmetries of the game, which were out of the focus of the initial inquiry, are broken naturally because of these patterns. The role of introducing minimum density thresholds in the spatial distribution of players is also examined.

These studies are preparatory for the model presented in chapter 6. This model is based not on a mathematical formulation of collective behavior but on the iterated resolution of individual games which is directly instantiated in a computer simulation that includes all the factors previously studied as well as others. While in none of the mathematical models is cooperative coordination an evolutionarily stable solution, in the individual-based model it is under certain circumstances. These circumstances have to do with the value of the parameter which measures the degree of conflict between players. For small values of this parameter (moderate conflict) cooperation ensues. As the degree of conflict is increased the global proportion of cooperative strategies decreases in a more or less linear fashion.

Based on previous results, it is possible to pinpoint the factors that play key roles in this change of outcome. First, other possible explanations are considered and eliminated such as the case of kin-selection. Then the role of structuring invariants, such as the emergent spatial aggregations of agents, is assessed. These invariants break many of the symmetries of the game by introducing qualitative and quantitative differences according to the position occupied by an agent within one of these clusters. This situation combines with the key role played by a treatment of agents as discrete entities subject to stochastic fluctuations to explain why coordination evolves as well as other observed correlations.

Understanding the individual-based model permits the formulation of interesting analogies with natural situations in which the role played by spatial structures may actually be played by other invariants which may constrain the patterns of interaction in similar ways.

1.1.5 An evolutionary role for social influences on development

The model of ecologically situated evolution studied in chapter 6 can be understood as a historical process since the activity of agents makes them constitute invariant structures that constrain their future evolutionary dynamics. In view of this and of the analogies drawn with other natural phenomena, another factor that may also play a similar constraining role in evolution is studied in chapter 7. The aim of this chapter is to assess the evolutionary significance of a process of social influences on phenotypic development.

There is much evidence from the animal world on how parents affect the development of their offspring, especially regarding the acquisition of behavioral patterns. This evidence is briefly reviewed and two possible evolutionary roles for socially influenced development identified.

The idea is to extend the same situation modelled in chapter 6 (thus facilitating comparisons) by including the possibility of slight phenotypic changes in newborn agents depending on the parental phenotypes. Thus, socially influenced development is simplified to the point of being an abstract and instantaneous process. Nevertheless, some key features of natural development are

retained such as the interplay of genetic, environmental and formal factors in determining an adult phenotype.

This simple extension is enough to introduce important qualitative differences in the evolutionary outcome. If the introduction of space and discreteness results in a stabilization of cooperative coordination for certain values of the parameter that measures the degree of conflict, the addition of social influences on development shows that the range of stabilization is not only enlarged but the type of stability is also different. Without development the global level of cooperative coordination decreases linearly as the conflict parameter is increased. However, for the case of socially influenced development, the level of coordination changes abruptly at a critical value from full cooperation to lack of cooperation.

An examination of the constitution of the population reveals the presence of self-promoting networks of developmental relations which remain invariant during evolution and enslave the dynamics by disallowing phenotypes which do not conform to the logic they impose. The operation of these networks is studied in detail. For a critical degree of conflict the developmental networks do not appear. Interestingly, the form of these self-promoting invariants can be interpreted as an ecological explanation of patterns that otherwise would be classified as imitation or cultural transmission of behaviors.

Again, the historical and situated character of evolution is highlighted as the constitution of a population now depends not only on the behavioral strategies of the agents but also on developmental constraints which are at the same time formal (due to the developmental relations that are allowed) and historical (due to the actual developmental invariants constituted by the agents).

1.1.6 Coordination without non-shared information

Unlike the other investigations in this thesis, chapter 8 presents a simple model which is intended mainly as a proof of concept. Typically, when researchers try to devise an evolutionary scenario for studying the evolution of communication, they think in terms of situations in which (semantic) information is unevenly shared by the participants. This, it is concluded, is a necessary condition for communicative behaviors to evolve. If communication makes evolutionary sense, there must be something to communicate about. Nowhere is this assumption less questioned than in computer models on the evolution of signalling systems. The language used in these models to justify this assumption suggests that communication would be impossible otherwise.

The theoretical perspective presented in chapter 2 insists that functional and operational aspects of communication should remain separate, and that communicative behaviors may well occur in the absence of a reasonable way of interpreting them in terms of an exchange of information. Is it possible that, in spite of this, all imaginable scenarios for studying the evolution of communication need to be framed in terms of transmission of information anyway? Or can other scenarios be imagined in which behavioral coordination can evolve nontrivially and yet information is not transmitted?

A simple modification to the model presented in chapter 6 is used to respond affirmatively to the last question. In this model the action-response game is extended so that instead of a single action and a single response, the participants must engage in the production of an alternating sequence of actions and responses which depends on the type of food they are trying to access.

Payoffs are allocated according not to individual actions but according to the correctness of the sequence that corresponds to the food source. Food types are equally “perceivable” to both players and act as constant environmental stimulus during the production of the sequence. Because of this, there is no information to be transmitted between the players, neither in terms of environmental features nor in terms of changing internal states since they are simply modelled as state-less machines. And because agents are so modelled the production of a sequence in the face of an unchanging stimulus (the food type) makes the task nontrivial.

The computer model shows that interesting behavioral coordination evolves in this game even in the case of conflict of interest thus reproducing a result analogous to the models in the previous chapters.

The main point of the model is to show how deeply can preconceptions affect research. An interesting, non-denotational mode of communication can evolve even if the scenario determines that information is equally shared by everyone.

1.1.7 Coordination, structural congruence and embodiment

In some of the investigations presented in this thesis, the complexity of social coordination has to be somehow sacrificed in order to study evolutionary and ecological dynamical patterns. This, however, does not mean that coordination cannot be studied in more detail from a behavioral perspective using similar methods. Chapter 9 presents such a study for the case of embodied agents who interact in an acoustic medium.

One of the aims of this model is to test the applicability of the theoretical framework described in chapter 2 which, up to this point, has served mainly as guidance for approaching different issues. However, its usefulness for understanding concrete questions has not been directly assessed. Another aim is to try to understand the relations between embodiment and adaptive social behavior as well as the relations between social and non-social behavior, albeit at a very simple level.

The choice of sound as a channel of interaction is not incidental. The physical constraints of continuous acoustic coupling pose nontrivial problems like discriminating self from non-self and structuring production of signals in time so as to minimize interference. An equally relevant factor is given by the mechanisms of binaural perception. For instance, the intensity perceived by each ear in the case of high frequencies depends on the angular orientation of the body. Factors such as this may encourage active perception as an integral part of behavior.

In this model, a pair of “blind” agents moving in an unstructured arena must locate one another and remain within a short distance from each other for as long as possible using noisy continuous acoustic interaction. The dynamics of the agents is given by the operation of an internal continuous-time recurrent neural network which couples sensors, motors and perceptual regulators. The various parameters associated with these networks form a vast space of possible structures many of which may be viable for performing the desired behavior. This space is searched using a genetic algorithm and viable solutions are studied in detail.

Careful observation of the most frequently evolved behavioral strategy shows that interacting agents perform rhythmic signals leading to the coordination of patterns of movement. During coordination signals become entrained in an anti-phase mode that resembles turn-taking. It is shown that rhythm is an emergent feature of the interaction since it is absent in the individual dynamics.

The role of active perception in the generation of rhythm is uncovered. Both the temporal structure of the interaction and the coherent patterns of movement are evidence of coordination being achieved by means of the acoustic coupling.

Perturbation techniques show that signalling behavior not only performs an external function but it is also integrated into the movement of the producing agent, thus showing the difficulty of separating behavior into social and non-social categories.

The achievement of a state of structural congruence between agents is demonstrated by exploring their internal dynamics. Coordination can only be understood as a result of the recurring mutual perturbations that the agents undergo. In other words, agents *co-adapt* to each other.

In principle, it is possible to imagine a different situation: that of agents who are able to adapt *individually* to externally imposed dynamics. In order to test whether this is the case, the capability of agents to synchronize their signals with “artificial” beacons that reproduce the signals of coordinating agents is tested. It is found that agents are incapable of achieving entrainment in the presence of the non-plastic beacons, even if the beacon reproduces the exact same signal of a previous co-participant. This demonstrates that the *performance* of coordinated patterns of behaviors does not entail the individual *competence* for performing similar patterns.

The value of the concepts of coordination and structural congruence is demonstrated as well as the explanatory power of studying adaptive behavior from an embodied and situated perspective.

1.2 Summary of contributions

This thesis aims at the exploration of basic social behavior in autonomous systems at different levels ranging from the evolutionary to the behavioral. It also aims at establishing the value of theoretical and methodological frameworks which manifest themselves to different degrees in each individual investigation.

Social behavior is viewed as a natural phenomenon which is constrained, but not necessarily determined by biology. The theoretical terms used to describe and explain social behavior must be grounded in a systemic theory of biological autonomy since only operational discourse can be used in this way. This move puts questions of functionality between brackets, but does not discard them as useless. However, caution must be exerted when switching between the different modes of scientific discourse as shown by the criticism of how evolutionary biology has addressed the phenomenon of animal communication.

The systemic view is based on the concepts of autonomy, structural coupling, coordination and structural congruence. This basis is enough for describing social behavior in operational terms. As a consequence of adopting this perspective many issues that were previously ignored become relevant and worthy of investigation. Some extensions to the basic framework are briefly discussed, such as the relevance of coordinated behavior between unevenly plastic systems for explaining themes related to social learning and social bonding.

This theoretical basis is complemented by a description of historical processes as well as some examples. By influencing their own conditions of realization, these processes are able to introduce inhomogeneities in time. This discussion is relevant for understanding how the situated dynamics of processes like evolution work when idealizations such as the total independence of an environment from the organisms that evolve in it are relaxed. Assumptions like these become especially

problematic when dealing with the evolution of social behavior.

Historical processes may spontaneously form organizational invariants which constrain the dynamics that gives rise to them in ways that will tend towards their own perpetuation. This dynamical phenomenon may fulfill a highly important explanatory role.

The pragmatic value of these issues is explored in the form of concrete investigations. Since these investigations are based on novel simulation techniques, some issues related to how they are used and how they relate to other scientific activity are explored. A pragmatic approach to the use of formal models in general is defended. This approach aims at finding the relevance of different factors and hidden assumptions by comparing systematically different models of similar phenomena. Especially interesting for this task are models which are not mere extensions to previous models but which radically break with them. A comparison between models which bear little or no genealogical relation has higher chances of bringing into light hidden assumptions and their relevance but at the same time the comparison may not be easy to perform.

This leads to the question of how computer simulations should be used. The best way of formulating and answering this question is by adopting a view of science as a human activity which is defined by a network of practices. The integration of a new tool to this network is the issue to be addressed. After examining some of the problems that may arise from not considering methodological questions with the seriousness they deserve as well as previous attempts at solving some of these problems, a modest proposal is offered. This is not a prescriptive methodology but it aims at describing the necessary steps to achieve an understanding of what goes on in a simulation and then how to relate this understanding via analogies and metaphors with theories of natural phenomena. Computer simulations used in this way fulfil a scientific role analogous to that of thought experiments.

This methodology is exemplified in the investigations. The first one deals with the issue of the evolution of cooperative coordination in the face of conflicting interests. By designing a simple evolutionary scenario, game theoretic models are built to study the expected evolutionary outcome. Extensions to this basic model reveal the relevance of different factors. In none of these cases is cooperative coordination the expected result. However, in an individual-based computer simulation in which agents are modelled as discrete entities in space, coordination does evolve. This is a consequence of considering evolution as a situated process in coupling with ecological and historical factors. The role of spatial structuring invariants that arise in this historical process is fundamental for understanding what goes on.

The addition of social influences on development also shows how other factors coupled with differential reproduction transform evolution into a historical process where organizational invariants are responsible for the patterns observed which would not be explainable otherwise.

The thesis then moves towards issues more related to the behavioral aspects of social coordination. A simple proof of concept shows that coordination can evolve in a significant manner without necessarily assuming as necessary a situation in which information is kept from some of the participants.

Finally, the usefulness of the conceptual structure presented in the initial chapters is put to the test in a more direct manner in a model of embodied behavioral coordination through an acoustic medium. It is shown that emergent features like rhythm of interaction, entrainment, turn-taking,

active perception, sensory regulation, coordinated movement and structural congruence, can be understood from a systemic point of view as arising from the constraints imposed by the embodied relations of agents and their environment.

The thesis concludes with a review of the main contributions and some speculations about possible future avenues of research.

Chapter 2

Towards a biology of social behavior: a systemic perspective

The aim of the present and the following two chapters is not merely to provide a necessary theoretical and methodological background against which to place the experimental studies presented in the remainder of the thesis. If these studies seem to fall short of achieving a thorough exploration of all the potential avenues of research implied in the theoretical exposition that follows, this is partly owing to the richness and openness of the subject. The objective, therefore, is to present a rough map of the territory and its potential, both thematic and methodological, for scientific research without necessarily fulfilling this potential in its entirety in the actual studies presented afterwards. These concrete studies are intended firstly as contributions in themselves to the particular subjects they address but also as examples of how particular methodologies and sets of ideas can be applied within the general theoretical framework depicted here.

In the present chapter a review will be made of some issues regarding the biology of social behavior, starting by specifying the need for a biological grounding of the terms used to discuss social behavior and by analysing the distinction between functional and operational statements as belonging to two different domains of scientific discourse. Comments will be made on how these domains are related and some of the methodological problems that may arise when the distinction is blurred. The chapter will then proceed with a detailed analysis of an example where it is considered that problems of this kind have arisen in biology. Such is the case of animal communication as studied by evolutionary biologists and behavioral ecologists. The corresponding section will discuss some of the problems that originate in characterizing communication only in functional terms according to selective advantages and information transfer and how this perspective has shaped the set of questions that make sense to investigate. This is followed by the theoretical core of the chapter (section 2.4) where a systemic perspective on social behavior in autonomous entities is presented based mainly on the works of Humberto Maturana and Francisco Varela (although some points will be criticized or further elaborated). The specific case of animal communication is then re-assessed from this new perspective and some pragmatic consequences derived and compared with the more traditional point of view.

2.1 The pragmatic need for a biological grounding

Any good interpretation of a map, if some use is expected out of it, must be guided by some principle so as to avoid, at the initial stages, spending too much time on details and forgetting the bigger picture. The guiding rule that is to be used here is that all the phenomena under consideration in the context of this thesis should be able to be linked in a continuous fashion to phenomena describable in the language of dynamical systems. In particular, descriptions of biological and cognitive phenomena will be expected to be able to be made continuous with the consequences of the basic logic of the autonomy of living systems as material entities.

It is important to emphasize that the aim is not to give supporting arguments for such approaches in themselves [e.g. the use of dynamical systems in cognitive science instead of a more traditional computational approach, (Van Gelder & Port, 1995; Van Gelder, 1999)]; rather, the point is to *use* these approaches mainly as conceptual tools that will enable the research to view the issues of interest from a certain systemic perspective that is believed to be fruitful (and in some cases badly needed) for approaching a scientific understanding.

Some readers may argue that the use of a conceptual framework implicitly endorses such a framework as more acceptable than others. Indeed, this could be said. However, it is preferable to say that the use of a tool indicates an opinion mainly about its suitability and nothing else. In certain parts of this thesis it will be useful to apply game theory to evolutionary problems, simply because it is the right tool for the job. This does not mean that evolutionary game theory is endorsed as the most acceptable approach or as providing any sense of truth or privileged explanatory power.

The hypothesis of continuity of cognitive and social phenomena with biology (and eventually with physical phenomena¹) demands that all descriptions be traceable to biological, physical or dynamical substrates, without in any way implying that these substrates are all one needs to *know* to generate from first principles the concrete phenomena of interest. Such a form of reductionism would be at odds with the pragmatic stance that is intended to serve as guidance since it would mean that certain distinctions are privileged over others ignoring the fact that such distinctions are made in the first place in order to fulfil a purpose and not as ontological commitments.

The purpose of tracing the phenomena of interest to a biological and dynamical substrate is not to be able to describe how the latter *determines* the former but how the former is *constrained* by the latter. In this way, if one is able to trace social phenomena to biology, it does not mean that social phenomena cannot have a logic of their own, impossible to arrive at from a strictly biological plane. In fact, what will be discovered, is that social phenomena, while traceable to biology, are usually underdetermined by it. Finding a trace of continuity between two domains of description is not meant as a reduction but rather as the identification of how one domain can constrain the other. It is not practical to spend too much time with theories that simply do not fulfill these constraints.

¹On this issue see for instance Wheeler (1997), Stewart (1996) and in general the works of Maturana and Varela as discussed in section 2.4.

2.2 Functional and operational approaches

Not all descriptions or explanations in science are equivalent. It is often convenient to make a distinction between two main forms of scientific discourse. In one case, an explanation or description is formulated in terms of a set of elements all pitched at a same descriptive level and also in terms of a set of law-like relationships between these elements so that an account can be given of how the phenomena are generated. Explanations and descriptions formulated in this way will be called “operational”. Typical examples of these are mechanistic explanations. In a different domain of discourse, it is possible to choose the relevant elements to be used much more freely and to focus only on certain type of regularities not necessarily apparent at the operational level while ignoring the rest (or putting it as part of a less specified context). One may choose to do this because it serves the purpose of facilitating an understanding of the relation between phenomena that operate at different levels or timescales or whose operational relationships are not completely understood or could not be completely understood. Such is the case of intentional, teleological or, in general, functional explanations or descriptions.

This distinction will be elaborated a little further followed by an exploration of the relationship between the two domains of discourse and why they should be kept distinct. Varela (1979) describes the distinction between the operational and the functional (or symbolic) modes of explaining in the following way:

“In both cases the recorded phenomena are reformulated or reproduced in conceptual terms that are deemed appropriate. The difference lies in the fact that in an operational explanation, the terms of such reformulation and the categories used are assumed to belong to the domains in which the systems that generate the phenomena operate. In a symbolic explanation, the terms of the reformulation are deemed to belong to a more encompassing context, in which the observer provides links and nexuses not supposed to operate in the domain in which the systems that generate the phenomena operate.” (Varela, 1979, p. 66)

Consequently, explaining a phenomenon operationally entails making a distinction of a set of elements or components which relate following a corresponding set of laws and relations, all of which play a generative role in the realization of the phenomenon.

“A characteristic feature of an operational explanation is that it proposes conceptual (or concrete) systems and components that can reproduce the recorded phenomena. This can happen through the specification of the organization and structure of a system [...] This is so because the organization of a machine [...] only states relations between components and rules for their interactions and transformations, in a manner that specifies the conditions of emergence of the different states of the machine, which then arise as a necessary outcome whenever such conditions occur.” (Varela, 1979, p. 66)

In contrast, a functional or symbolic explanation relaxes the condition that all its elements must be related through law-like links. The observer may choose to connect different elements only in terms of the correlations she finds interesting or practical between the system and a contextual set of other relevant phenomena. This allows for terms to participate in these explanations that do not necessarily refer to any concrete or observable components of the systems involved and

get their meaning from regularities between different observations made by the observer (e.g. the terms “fitness” or “selective history” in evolutionary biology). Thus, functional explanations always *escape* the boundaries of the systems that give rise to the phenomenon being explained by including contextual elements, either concrete or relational and, as a consequence, they can never be equated with an operational account of that phenomenon.

In spite of this, it is all too easy to slip from a functional type of talk into an operational one simply because certain terms in functional accounts can also be identified as concrete components of the system of interest. Consider the DNA molecule. It is possible to assign to it different functional roles depending of the context and regularities that the observer chooses to pay attention to. It could be said that this molecule encodes for sequences of aminoacids that will form proteins *given* the context of the molecular machinery that operates inside the cell (and this could be affirmed even if such a context is poorly understood). Or it could be stated that the DNA molecule encodes for heritable characteristics *given* the context of the reliability of its replication across a history of serial reproductive events. But it is wrong to state that DNA instructs the cell on how to differentiate or what chemicals to secrete in order to produce an organismic trait because, when looked as part of an operation of the internal processes of a cell, the DNA molecule is just as important as any other molecule. It is made up of atoms like all the others. There is nothing special about it. If DNA has any special status at all, it is not an operational one and not even an explanatory one *in general*. The only relevance of DNA is a functional relevance for *certain* types of symbolic explanation and *certain* contexts of discourse such as biological evolution and probably other conditions (see Clark & Wheeler, 1998).

Notice that the error in assigning operational power to functional statements does not depend on the criterion by which the functional statement is justified. Therefore, any arguments defending a context of legitimation for certain functional accounts instead of others [for instance, Millikan’s proper functions as those that are such in virtue of their selective history (Millikan, 1984)] is beside the point. No functional statement should be confused with an operational one or used in its derivation.

Given the distinction between the two modes of discourse, the question arises of whether one mode should be intrinsically preferable to the other. The answer is negative. Varela (1979, ch. 9) proposes explanatory complementarity as a pragmatic possibility. One can never hope to give a full operational account of phenomena as complex as human cognition even if it is certain that such phenomena could be logically derived from available operational statements. Such is the case of operational concept of autopoiesis [see Maturana & Varela (1980) and section 2.4] which, according to Varela, is capable of logically generating all the phenomenology of living systems but it is incapable, *per se*, of providing a cognitively satisfying explanation of specific instances of this phenomenology. An observer would have to follow an enormous body of historical contingencies (probably unknowable) to give a full operational account of a specific pattern of behavior in a given species. In contrast, symbolic abstraction may provide such a source of understanding by ignoring much of this historical detail and following a set of simpler principles.

In a recent article, Faith (1997) seems to be saying something similar to this when he also points to the practical impossibility of giving a full operational account of a cognitive system purely in terms of dynamics, state variables and attractors. One may agree with Faith in that, for

understanding systems as complex as living cognitive entities, functional analysis is pragmatically unavoidable. However, Faith downplays the explanatory role of a dynamical systems approach to the study of adaptive behavior (and by extension one would add of other operational accounts of life, cognition and social phenomena) by equating this approach with Laplacean determinism and, consequently, impossible to achieve. This goes against the pragmatic spirit upheld in this work. What seems to be implied is that the *only* role an operational description or explanation can have is one where the *whole* issue of interest is addressed purely in operational terms. If one accepts the explanatory complementarity suggested by Varela, then an operational account need not be a complete or final one. It may be sufficient to derive from some basic operational givens a set of constraints which may be used in different ways: in the formulation of an operationally informed functional story, as proof that certain statements go against what is operationally possible, etc.

Therefore, it is hardly necessary, *contra* Faith, for an operational story to be useful that it be complete and final. For instance, if it is accepted as an operational given that the nervous system is a system that exhibits a closed organization (Maturana & Varela, 1980), it follows logically that it cannot operate with representations of the medium external to the organism, even if one does not know the first thing about the operational complexity of the nervous system of the particular organism in question. The statement, derived from an operational description, is telling the researcher something useful about the pragmatic value of functional explanations that postulate such representations. It will only be possible to do a limited number of things with them since there is going to be a guaranteed conflict as one starts testing their usefulness in domains where the actual mechanisms become more relevant (e.g. when building a biologically inspired autonomous agent).

This less restrictive use of operational statements is already indicating something about a possible useful relationship between the two domains of discourse. If a symbolic explanation can be formulated as an abbreviation of a network of lawlike operational relations which pertain both to the system being explained and to a more encompassing context, then *preferable functional explanations should be those which are constrained by a lack of contradiction with existing operational descriptions* (i.e. with a description of those law-like links which are abbreviated). And such constraining power may also arise from operational statements that are not aimed at providing a complete account of the phenomenon.

This kind of constraint may also be used to fill what appears to be a gap in Varela's proposal for the complementarity of both modes of explaining. There is an apparent contradiction between saying that functional explanations are symptomatic of the lack of appropriate knowledge of the context that provides the background to the system of interest (Varela, 1979, p. 65) and saying that functional explanations (in particular those of the teleonomic kind) are abbreviations of nomic relationships (Varela, 1979, pp. 67 and 73). How is it possible to abbreviate nomic relationships which are not known? Clearly, functional explanations will be especially helpful when trying to understand phenomena too complex to formulate in operational terms, at least with the tools available at the time. However, this is not to say that function cannot be constrained by existing operational knowledge even if it is not complete. Consequently, while it is possible to *postulate* functional explanations as being abbreviations of (not fully known) operational relationships in a greater context this by no means implies that the functional explanations need to be *constructed*

as such. A functional account can be formulated by an observer by simply bringing forward the correlations in which she is interested. Even if one may say that the resulting functional account abbreviates nomic links there is no need to know those operational relationships beforehand. Most of the time it may be sufficient to have an incomplete operational knowledge in order to restrict, as suggested, the type of functional explanations that make sense to formulate in the chosen context. Such is the spirit with which operational accounts will be given in the present chapter and in general in the whole thesis.

Another issue that is sidestepped in Varela's account is the influence existing functional explanations can have in the formulation of new operational ones. Clearly, the latter may act as constraints for the former but it makes little sense to say that the same happens the other way around. Operational discourse must follow the logic of the elements that constitute the domain in which the phenomena of interest is assumed to be generated and the lawlike relationships between them. The laws that operate in a chemical system cannot change depending on the function assigned to that system in different contexts.

One would like to suggest that functional explanations may indeed exert an influence on how operational sense is made of a system, albeit in a much subtler way. A functional explanation may provide hints as to which is an adequate operational level of description for that system in order to explain how the associated phenomenology is generated. And so, different functional perspectives may suggest different operational levels of description.

To take an example close to the themes of this thesis, if the process of language acquisition in humans were taken to be the result of an individual capability for learning an existing system of rules and representations then one could try to formulate adequate operational descriptions in the domain of individual mechanisms which may give rise to such a capability (i.e., in terms of neurons, genes, etc.). In contrast, if the ability to participate in linguistic behavior were to be taken as arising from the interplay between individual behavior and the dynamics of social interactions (using functional terms like internalization, parental scaffolding, etc.) then one could choose a level of description extended beyond the individual as the appropriate one to give an adequate operational account. In the first case the system is bounded by the individual organism and the rest is the (relatively static) context. (This would probably give rise to a Chomskyan-type view of language acquisition). In the other possibility, the system is extended to include part of that context. (The inclusion of social dynamics as part of the system could perhaps give rise to a Vygotskian-style operational description)².

Different functional understandings will provide different hints as to how to choose the boundaries and composition of the system whose phenomenology one would like to explain operationally. This is not to say, however, that this is the only criterion for preferring an operational explanation instead of another and in the example given above there may be additional reasons (besides the preferred functional understanding) for deciding for one of the cases (e.g. other operational statements which one would not want to contradict).

Some of the issues discussed in this section would deserve a further exploration which cannot

²For an explanation of Chomsky's theory of language learning see for instance chapter 1 of (Chomsky, 1965). For Vygotsky's theories of cognitive development see (Vygotsky, 1962, 1978; Wertsch, 1979, 1991). In none of these cases it is claimed that these theories are operational in themselves, rather that operational theories similar to them could perhaps be formulated depending on the type of preferred functional perspective.

be provided here without losing track of the main issues of the chapter. The distinction was made between two modes of scientific discourse emphasizing the point that they can be used complementarily although they should remain distinct. It has also been suggested how the two modes could be usefully related: functional discourse could be constrained by operational discourse by not contradicting it and at the same time it could suggest which is an appropriate level of description where the elements of further operational accounts should be distinguished. One important point to recall in the following sections is that insisting on addressing a phenomenon operationally should not be equated with a reductionist attitude nor should the usefulness of operational descriptions be measured exclusively by their degree of completeness.

2.3 Confusing the domains of discourse: the case of animal communication

That the two explanatory domains, the functional and the operational, may coexist in scientific practice does not imply, as insisted, a lifting of the distinction. They both serve different purposes and can be related but should not be confused. Not to fall into this trap may not be as easy as it sounds. To illustrate this point, a discussion will be presented of how theoretical evolutionary biology has addressed a paradigmatic example of animal social behavior: communication.

2.3.1 Selective advantages as a defining feature of communication

Despite the intensive attention received by biologists, animal communication remains a confused subject with “leading theorists even disagreeing as to what should be properly called a ‘signal’ or ‘communication’ ” (Dawkins, 1995, p. 72). Here it will be suggested that much of this confusion is rooted on the misapplication of a functional concepts (selective advantages) into an operational domain.

Consider the following selection of relatively recent definitions of animal communication given by behavioral ecologists and evolutionary biologists.

Wilson (1975) defines communication as the altering by one organism of the probability pattern of behavior in another organism in a manner adaptive to either one of them or to both. Lewis and Gower (1980, p. 2) define communication as “the transmission of signals between two or more organisms where selection has favoured both the production and reception of the signal(s)”. Krebs and Davies (1993, p. 349) define it as the “process in which actors use specially designed signals or displays to modify the behaviour or a reactor”. They later make clear that they understand “specially designed signals” as those that have been favoured by natural selection. Burghardt defines it as a behavior that is “likely to influence the receiver in a way that benefits, in a probabilistic manner, the signaller or some group of which it is a member” (Burghardt, 1970; MacLennan & Burghardt, 1994, p. 163). Maynard-Smith and Harper (1995, p. 306) define a signal “as an action or structure that increases the fitness of an individual by altering the behaviours of other organisms detecting it, and that has characteristics that have evolved because they have that effect”.

Although these definitions are not equivalent and are carefully phrased in order to emphasize differences in what the authors consider to be the key aspects of communication, they all share a common feature: communication is characterized in terms of selective advantages. A communicative event between organisms is such only in virtue of a history of selection of similar patterns in the behavior of their ancestors. Accordingly, any complex interaction between organisms, no

matter how ritualized or similar to known cases of communication, cannot be considered to be an instance of communication until its adaptive value has been figured out or its selective history postulated. Surely, this is not a common situation. Even in evolutionary biology, where contribution to Darwinian fitness is a functional property of paramount importance in any trait, adaptive value rarely constitutes a *defining* feature of that trait. It is possible to define wings in terms of their structure and even in terms of functions relative to the observed behavior of a winged animal³, but there is no need to say that they provide advantages for survival in order to say what they are even if one appeals to those advantages to explain why they are there.

Marian Stamp Dawkins (1995) comes close to a recognition of the confusion whose roots can be traced to a requirement that signals have been especially designed to fulfil a communicative effect. According to this, one would expect all signals to be clearly defined and even exaggerated so that their conspicuousness is augmented. But such a requirement leaves as non-communicative lots of “more subtle and unritualized ways in which animals influence each other’s behaviors” (Dawkins, 1995, p. 75) such as coordinated flock movements. However, she fails to stress the inherent methodological problems created by the requirement of selective advantages.

All the above definitions could easily be questioned on methodological grounds as they present a (sometimes dangerous) mixture of definition and explanation. Communication is characterized in terms of a possible (and admittedly plausible) explanation of it. But what is “it” exactly? The danger evidently resides in the potential for tautological situations in which all instances of communication can only be explained in terms of selective advantages and if they cannot then they are not really instances of communication or if they are suspected to be, then the selective advantages must be there to be “discovered”. While not all circularities are necessarily vicious this one is certainly problematic because the phenomenon remains poorly characterized and one must always rely on intuitions or informal notions such as signals or information to tell whether communication is happening or not. Some evolutionary biologists might respond to this by saying that everybody intuitively knows what a signal is. Such a statement, however, would not seem to be justified by the constant succession of refinements in the biological literature to the way communication should be defined.

2.3.2 Some ill-defined terms

Interestingly, in none of the modifications just mentioned has the requirement of selective advantage been abandoned in favour of an operational definition. On the contrary, it could be argued that the history of refinements to definitions of communication has followed the path of using the property of communication as conferring selective advantages as the mould other more diffuse terms (such as “signal” and “information”) should fit. Given the fact that terms like these are needed to provide an intuitive grounding which cannot be provided by selective advantages alone, trying to tie less fuzzy versions of these terms to the very property that they are supposed to complement makes the potential for vicious circularity a very real possibility.

Notice, for instance, the emphasis on “specially designed signals” in the definition by Krebs and Davies (1993). Is it not the case that, according to their own logic, if signals were not specially

³Although it makes little difference to the point being made, it is debatable if such non-selective functions will make sense; see for instance the work of Ruth Millikan (1984). Thanks to Mike Wheeler for pointing this out.

designed by natural selection they would not be signals (in the sense of communicative interactions) at all since they would not participate in any type of communicative behavior? Selective advantage plays the normative role of specifying which actions constitute a signal and which actions do not but at the same time selective advantage is used to distinguish communication from non-communication. In the end, what a signal is remains a mystery.

Consider another logical consequence of defining communication as a self-benefiting activity that strips the word “signal” of all its intuitive meaning. Can it be said that camouflage, the resemblance of an organism to a part of the background in order to make itself difficult to be detected by a predator, is a signal from the former to the latter? According to Hasson (1994) this would be the case. Camouflage would be a signal that “reduces” or “hides” information from the receiver. Maynard-Smith and Harper (1995, p. 307) admit that this is an odd situation but they support Hasson by saying that “the camouflage is an evolved adaptation which changes an animal that once did not signal to one that signals ‘I am a leaf, or other piece of background, and not edible’.” As they say, if camouflage were not a signal, then they would be faced with some hard distinctions as there is no logical dividing line between camouflage and what they consider to be more intuitive instances of signals such as warning coloration. But, at the same time, if camouflage is accepted as a signal an internal contradiction becomes apparent. A structure or behavior that supposedly has evolved in order to modify the behavior of other organisms becomes effective when it “fails” to do so. The predator “acknowledges” the signal by *not* perceiving it and acting in exactly the same way as if the signaller was not there. Communication occurs without interaction! Of course, a possible answer to this criticism would state that the predator’s behavior has changed from a situation in which the prey was easy to detect to a situation in which the prey is camouflaged, and *in that sense* camouflage can be considered a signal. But these two situations can only be compared on an evolutionary time-scale, from a time where camouflage was less than a poor resemblance to the background to a time where the prey actually exploits specific features of the background to become very hard to detect. Consequently, this answer only makes the situation worse because the meaning of a signal becomes even more diluted as it is supposed to be an action or structure that operates on a behavioral time-scale and not on an evolutionary one.

Equally confusing is the meaning given to the word “information” in its semantic sense. Such a notion remains very diffuse and can be taken to signify different things for different authors. However, the concept has remained intuitively important given the reigning functional framework in which animal communication is studied. Accordingly, Maynard-Smith and Harper have recently presented a taxonomy of different types of signals according to the type of information they confer. They make clear that they refer to the semantic, “colloquial” meaning of the term rather than the technical one, (Maynard-Smith & Harper, 1995, p. 305).

Indeed, the mathematical concept of information as reduction of uncertainty over noisy channels between largely congruent systems (Shannon & Weaver, 1949) can play no semantic role (see for instance von Foerster, 1980; Oyama, 1985, p. 65). Shannon information is contentless and therefore of no use for the functional type of explanations usually sought after in evolutionary biology. It will be seen later (section 2.4.3) that it is also of little use for an operational definition of communication since once it has been explained how to get to a situation in which one may legitimately speak of Shannon information being transferred between two communicating organisms,

all that needs to be explained has already been explained⁴.

It is the semantic use of the word “information” that is problematic, especially when it is intended to be applied operationally. It should be clear that information in this sense lacks operational efficacy simply because it can never play a generative role in the phenomenon of communication. It is not to be found as a part of the systems involved (if these are taken to be the communicating organisms). Semantic information always refers to a nexus made by an observer between the communicative behaviors and a more general context in which those behaviors take place but which does not form part of the behaviors themselves. It is, therefore, an inherently functional concept.

So, why is this term still used in a pseudo-operational tone for making statements like “[a] signal alters the information available to the receiver” (Maynard-Smith & Harper, 1995, p. 307)? It should not be surprising to see that the answer is again related to the fact that communication has been defined in terms of selective advantages. The information gained in a communicative act is important because it has fitness value (the location of the piece of food, the presence of the predator). If it has fitness value then it is relevant to actually confirm whether, according to the conventional definition, the act was communicative or not. Information is used as a common currency to compare changes in fitness in different situations. To see to what extent this is so, one may point to what could be called “Hasson’s convertibility rule”, i.e. his suggestion of how changes of information should be measured in terms of changes of fitness (Hasson, 1994). According to Maynard-Smith & Harper (1995, p. 307), who endorse Hasson’s criterion, a “signal alters the information available to the receiver. The change of information may be positive, in which case the fitness of the receiver is increased, or it may be negative, so that the fitness of the receiver is reduced”. In fact, for Hasson the entailment is double (Hasson, 1994, p. 225); his suggestion is that changes of fitness should be easier to measure or deduce and then the corresponding change of information estimated. Such a refinement to the idea of semantic information (now a quasi-measurable quantity) follows, as argued before, the notion of selective advantages as the mould to fit.

Of course, having this use for semantic information has afforded evolutionary biologists with ammunition to argue about what sort of signals could be expected to be found under selective pressure for increased individual fitness. The pressing question becomes, for whom should information have fitness value? Given the neo-Darwinian emphasis on individual fitness, in principle this should be for the originator of the signal⁵ because the originator supposedly has the choice (in the evolutionary sense) not to send the signal if it has detrimental effects. But, in contrast, the receiver also has a choice to ignore signals that carry no valuable information.

Based on this argument, Krebs and Dawkins (1984) criticize the use of semantic information in describing communicative behaviors. According to them it makes little evolutionary sense to im-

⁴Shannon information, however, may be put to practical use as a tool in ethological studies, animal communication being a particular case. The application usually involves dividing the behavioral space into a number of discrete options and estimating from observation the amount of transmitted information by correlating discrete events and subsequent reduction in entropy, (Krebs & Dawkins, 1984; Dawkins, 1995). This method, however, is not without problems, in particular those arising from the use of different choices in the partition of behavioral spaces.

⁵Notice, by the way, how easy it is, when an operational characterization is lacking, to fall into the simplistic assumption that signals are discrete behavioral events which necessarily *have* a single originator. Think of howling in wolves. See chapter 9 for other natural examples and a dynamical model where statements of this kind cannot be so easily made.

prove the efficacy of a signal in conveying information if the originator is not going to benefit from the production of that signal. For them, it makes more sense to dispose with the notion of information altogether and consider communication as actions with which individuals manipulate others by making them behave in ways which are convenient to the manipulator. Forgetting for a moment that this approach may have problems of its own⁶, one may ask instead if the notion of information is completely disposed of. Initially this would seem to be the case. But when the complementary aspect of manipulation is considered, which Krebs and Dawkins call “mind-reading”, information slips back in. Mind-reading is intended to account for the capacity to predict the future behavior of an animal by observation of its present behavior and therefore it suggests an informational interpretation. Mind-readers take advantage of the manipulative signals of other animals and act according to their own benefit. But this is tantamount to saying that they assign an informational content to those signals; a content which is clearly contextual on the current pool of behaviors in the population but which is informational nonetheless. The ensuing co-evolutionary arms-races that may occur according to Krebs and Dawkins are indeed possible because the reliability of the information gained through mind-reading is under the selective control of the manipulators. Consequently, from the mind-reader perspective, manipulation is about the transfer of information or, rather, the distortion of such information⁷.

Here too, as with Hasson’s convertibility rule, the understanding of information, a notion originally invoked to intuitively complement a definition of communication poorly characterized by selective advantages alone, is made dependent upon selective advantages in a way that it almost acquires operational efficacy in contradiction with its functional origins. As with the notion of signal, information becomes so tied to the idea of adaptive value that it makes it even more difficult to act as a grounding for the concept of communication.

2.3.3 Practical consequences for research

What does it mean to study the evolution of communication under this framework? For the traditional standpoint the ability to communicate is an established fact. This is paradoxically so even when the issue of interest is the *origin* or *evolutionary change* of a communicative system. What is meant by this is that the issue that is usually addressed in conventional research practice is the evolution and stability of *content* in communication. Or, more specifically, how a selective scenario is capable of imposing meaning on a set of behaviors which initially have none. In most models on the evolution of animal communication the organisms/agents always have the ability to send a signal, make a move in the game, produce a warning call, etc. The form and structure of communication are always pre-defined and pre-existent. What these models seek to find out, be they verbal, mathematical or computational, is what kind of signal, which move or warning call the organisms/agents will produce in different situations after many generations of selective pressure and not how the communication system ever came to be.

In spite of the inherent paradox, such a methodology is far from futile since according to the

⁶For instance, distinguishing manipulation from cooperative communication depends on how one defines a conflict of interest, which may be problematic in itself.

⁷Notice that the argument does not really depend on the occurrence of an actual arms-race. As Maynard-Smith and Harper note (1995, p. 305), “it is not evolutionarily stable for the receiver to alter its behaviour unless, on average, the signal carries information of value to it. The fitness of the receiver need not be increased on every occasion that it responds to a signal”.

results the researcher could derive interesting conclusions about other issues. As an example, the evolutionary stabilization of altruistic signalling can tell the researcher under what circumstances cooperative behaviors will obtain (e.g. when agents communicate with relatives or when signals are costly, etc.). And so the approach can be useful for answering these questions. But it is clear that one is missing the chance of investigating other important issues, such as why should an agent respond to a signal in the first place, why should a particular game be played at all, how can novel games appear in evolutionary history, are there any constraints in the form and structure of communicative interactions, etc.

What it is being suggested is that the main reason for questions like these being rarely asked in research practice has to do with the influences of the traditional theoretical understanding of communication. As a response, it could be argued that this is not the case. Maybe most researchers find the question of the origin of communicative behavior not as interesting as the question of the functional role played by an evolutionarily stable set of signals in pre-existing communicative patterns. Or maybe the current tools of research do not allow one to address issues such as why should an agent produce and “interpret” certain behaviors as signals in the first place, etc. Maybe the traditional theoretical framework is blameless after all.

It could be. Clearly, it is not possible to give a definite answer to this argument. But, if one considers what sort of phenomenon is described in the traditional view (apart from all the contradictions and problems previously pointed out) it is clear that what is highlighted of communication is the ability to provide adaptive meaning to certain types of interactions in the form of transfer of information or in the form of exploitative manipulations and that such ability must provide (generally individual) selective advantages as the definition of communication requires. If this is what is relevant about communication then the type of research carried out so far should not be unexpected because the questions being asked are precisely the natural questions that the framework suggests, i.e. questions about informational content, about function and about selective stability of communicative behaviors⁸.

As for other practical consequences of this type of definition of communication, it is clear that providing one possible explanatory factor with the credentials of being the defining feature of a phenomenon introduces a class distinction in the domain of explanations. Accordingly, any factor other than natural selection which may be relevant for explaining the operation and evolution of a communicative system becomes *ad hoc* and second rate. If this were just the criticism of a purist, then maybe the situation would not be pragmatically very damaging since, as some biologists would probably believe, those additional explanatory factors could indeed be second rate or *ad hoc* in the face of natural selection. But the evidence points in the opposite direction. The assertion that behaviors with a strong interactive component (such as social behaviors) can only be explained in terms of individual selective advantages is constantly being challenged and other factors arise at the level of the dynamics of dyads, groups and whole populations that constrain the action of natural selection in ways that they become elements of equal or greater importance in understanding the final evolutionary outcome⁹.

⁸More concrete evidence of how the informational exchange metaphor has affected the construction of many computational models is given in chapter 8. The basic design specification in these models is that information should be unequally shared because otherwise there is no reason for communication to arise. The model presented in that chapter is meant to serve as a counterexample in which this requirement is abandoned.

⁹Examples of how such factors can be manifested are given in the models presented in chapters 6, 7 and 9 where

Surely, if these problems are to be avoided it is necessary to look for ways of operationalizing the definition of communication, thus dodging the circularity of defining it in terms of functional properties which depend on what communication is. Once this is done, functional accounts and selective stories will certainly become key ingredients in understanding features of specific communicative systems. But if the operational definition is lacking these functional accounts will remain grounded only on intuitions and will provide a poor notion of what communication is in terms of what it (supposedly) is for.

2.4 A systemic perspective on social behavior

The previous section shows an actual and quite relevant example of the kind of problems that may arise when the operational and functional domains of discourse are confused. For this reason, it is the purpose of the present section to introduce a purely operational perspective on social behavior in autonomous systems and then come back to the issue of animal communication.

Before proceeding, the reader should be reminded that the purpose of offering an operational account is not necessarily to act as a total replacement of functional statements. This may be too narrow a point of view given the earlier discussion regarding the purpose of an operational explanation as acting as a constraint to, rather than a replacement for, possible functional interpretations. For this reason, it is a worthwhile enterprise to try to identify such operational constraints using dynamical systems theory and other systemic concepts even if a complete dynamical description cannot be given in concrete cases. This is particularly so in the light of the pragmatic claim that a theory of social behavior should be biologically grounded. The purpose of that claim was not for the researcher to be able to provide ways in which social behavior could in general be reduced to biology, but rather to constrain a theory of social behavior within the realm of biological plausibility. Therefore, if biological grounding is to serve a constraining purpose, it makes sense that this should be done using an operational account.

2.4.1 Basic concepts

When one thinks about social behavior the first thing that comes to mind is some notion of coordinated activity between two or more autonomous entities. In order to understand what it is meant by this it is necessary to describe the meaning of autonomy, interaction and coordination.

When speaking about a system in general a distinction will be made between its structure and its organization following (Maturana & Varela, 1980, p. 77). By *organization* of a system it is understood the set of relations that define that system “as a unity, and determine the dynamics of interactions and transformations which it may undergo as such a unity”. In contrast, *structure* refers to “[t]he actual relations which hold among the components which integrate a concrete machine [or system] in a given space”. The organization, therefore, does not specify the properties of the specific components that realize the system, it only defines the relations that these components must hold in order for the system (or machine) to belong to a given class. Thus, the organization of a car may be described as being of a certain size, with four wheels placed in a certain way, and engine that provides power for moving the wheels, a set of controls, and so on. The structure of

the influences of spatio-temporal dynamics, social development and embodiment are investigated respectively.

a particular car refers to the actual physical properties of the components that realize it as an instance of the class of cars. It will include details regarding bodywork, state of the engine, mileage, plate number, etc. It follows that many concrete systems may exhibit different structures while preserving a same organization¹⁰.

The word “structural” originating from this distinction is well defined if one is dealing with a material system. This is how the term will be used in this chapter; in particular, when speaking of an organism’s or an agent’s structure, this will refer to the set of components that constitute the whole of its body including its nervous system or controller. In contrast, the term would not be as clearly defined if one is dealing with non-physical systems which are presumed autonomous (say, a financial market), which for this reason are left out of the discussion.

It will be assumed that, at any moment, the state of a well defined dynamical system is completely determined by a previous state and by its structure at that moment. This is the hypothesis of *structural determinism* (Maturana & Varela, 1987, pp. 96 - 97). If this is not so, then either the system is not well defined and its boundaries must be extended or reduced, or the observer is not in a domain where it is possible to speak of the regularities of the operation of the system as something to be found in the properties of the system (for instance, if the operation of a system could be explained by supernatural causes). The latter option would make a scientific approach very difficult if not simply impossible, (see Maturana, 1978, p. 34).

Autonomy can then be defined in non-functional terms as the property of the organization of a system of being operationally closed¹¹. This does not mean that the system does not interact with its external environment. It means that the organization of the system is defined by a network of internal processes and that the operation of this network is sufficient for the those constituting processes to be generated and sustained (constituted) without any of them being driven from outside the system. However, the autonomous system is never isolated from its context. This would be a thermodynamic absurdity which would result only in the trivial case where the system does not change at all. Hence the emphasis on the term “process” which is meant to highlight the dynamical character of the definition and which places autonomous system always far from their thermodynamic equilibrium. In what sense can then a network of processes be operationally closed? In the sense that although the system relates to its medium through other processes (perturbations, exchanges of matter and energy, etc.) these do not constitute part of the organization of the system.

The *identity* of the system is, at the same time, defined as long as it remains operationally closed, (Varela, 1979, p. 57). This is possible because autonomy becomes an invariant of the system, i.e. a way of being able to point unequivocally to the same system or the same class of systems. As soon as autonomy is lost, the operation of the system is “opened” by allowing processes that were once external to the closed network to enter into the internal operation. Consequently, its identity becomes ambiguous.

This characterization of autonomy is general, living systems constituting a particular case where the processes involved are processes of production, transformation and destruction of components in the molecular space including a boundary as a special component, in which case the organizational property is called *autopoiesis* (Maturana & Varela, 1980, pp. 78 - 79).

What sort of relation can an autonomous system have with its environment in order to remain

¹⁰See (Varela, 1979, pp. 8 - 12) for further discussion on the duality between structure and organization.

¹¹In fact, this is a thesis rather than a definition, (see Varela, 1979, p. 55).

autonomous? It is clear that as soon as this relation is one where the closure of the internal organization of the system is disrupted from the outside, autonomy will be almost certainly lost¹². Preservation of autonomy divides the space of possible interactions into those that are allowed and those that are not, and this space is obviously contingent on the present state of the system. Allowed interactions will be manifested as perturbations to the system that do not break its operational closure and not as instructions of the dynamical path that the system must follow. A process whereby the system interacting with its environment undergoes a succession of allowed perturbations (resulting in changes in its structure) without losing its autonomy is called a process of *structural coupling*, (Maturana & Varela, 1980, p. xx). As long as the organization of a structure-determined system is conserved during its coupling with the medium, the system remains distinguished from the medium and operates independently (i.e. its dynamics is always a consequence of its own structure at a given moment). Obviously, this conservation of organization does not entail a conservation of structure since, as noted earlier, different structures may realize a given organization. This is why, during structural coupling, the structure of the systems involved does indeed suffer changes.

As long as the present structure can be realized without a disruption of autonomy, the system is said to be *adapted* to its medium. Structural coupling, then, is equated with conservation of adaptation and loss of adaptation with destruction of the autonomous system or transformation into a different one¹³.

Concepts similar to the ideas of structural coupling and adaptation were introduced in the middle years of cybernetics, (Ashby, 1960; Pask, 1961). Ashby arrived at a similar definition of adaptation in terms of stability and homeostasis: "... a form of behaviour is adaptive if it maintains the essential variables ... within physiological limits" (Ashby, 1960, p. 58). And a generalized version of structural coupling is used in physics to describe coupled systems governed by equations, where it can be easily translated into the rule that distinct systems can affect each other by perturbing parameters but not variables in a direct way (otherwise the systems are not really distinct)¹⁴.

¹²Consider the nervous system as a candidate operationally closed system where the internal constituting process is the recursive transformation of relative states of electrical activity of the neurons. If one wants a dog to perform a trick and one imagines that this can be achieved by inserting electric currents in specific locations of its brain, then in that case the dog's nervous system will cease to act as operationally closed. The same happens if the system is prevented from operating normally, for instance if an animal is anaesthetized and clamped in order to measure the response of its neurons to certain stimuli.

¹³As with other relevant concepts, such as operational closure, adaptation seems to be an "all or nothing" feature of the system (in this case, in relation to its medium). This characteristic plays an important role in the re-interpretation of evolutionary change as "natural drift", see chapter 3 and (Maturana & Varela, 1987, ch. 5) and (Varela, Thompson, & Rosch, 1991, ch. 9). Similar "all or nothing" interpretations of adaptation have recently been used to explain the dynamics of speciation as percolation in a two-state (viable and non-viable) genetic hypercube (Gavrilets & Gravner, 1996; Gavrilets, Li, & Vose, 1998). However, these clear-cut distinction can introduce some problems when one wants to account for processes where organizational properties can be said to change *gradually*. Although it remains clear that, according to the above definitions, a system is either autonomous or not, is either adapted to its medium or not, one would like to be able to speak of changes that are certain to drive a system, for instance, towards the loss of its autonomy (e.g. a serious wound). Such language is sometimes used by the Maturana and Varela on different occasions. They refer to "facilitation" of autopoiesis (Maturana & Varela, 1980, p. 108) and Varela finds it useful to speak of "minor" and "major" breakdowns in autopoiesis (Varela, 1997, p. 80). This seems to be a blind spot in the theory but not an unrecoverable one. However, since the purpose here is to present the basic concepts, the existence of problems like this will be merely pointed out without attempting to come to a resolution within the current context.

¹⁴This is by no means the only restriction, see (Ashby, 1960) particularly chapters 19 and 21 for discussions on state-determined systems and parameters.

2.4.2 Coordination

Structural coupling occurs between an unity and its medium which may include other autonomous unities, in which case one can speak of an *interaction* between them. However, mere interaction between autonomous entities does not seem to be sufficient to describe the resulting behavior as social even if it happens to have an adaptive function. There is something lacking in two animals just bumping into one another while trying to escape from a predator to call *that* a social interaction. What is being looked for is a concept that will allow a description of the complex patterns of social behavior observed in humans and other species. This is the idea of *coordination* or *orientation*¹⁵:

“An organism can modify the behavior of another organism in two basic ways:

(a) By interaction with it in a manner that directs both organisms toward each other in such a way that the ensuing behavior of each of them depends strictly on the following behavior of the other, e.g.: courtship and fight. A chain of interlocked behavior can thus be generated by the two organisms.

(b) By orienting the behavior of the other organism to some part of its domain of interactions different from the present interaction, but comparable to the orientation of that of the orienting organism. This can take place only if the domains of interactions of the two organisms are widely coincident; in this case no interlocked chain of behaviour is elicited because the subsequent conduct of the two organisms depends on the outcome of independent, although parallel, interactions.

In the first case it can be said that the two organisms interact, in the second case that they communicate.” (Maturana & Varela, 1980, p. 27 – 28)

According to Maturana and Varela, a behavior can be considered communicative when a pattern of interactions elicits some form of coordination between the participants. But what does this mean exactly? Coordination is a subtle concept. In one interpretation it involves the fact that many organisms can have a complex behavioral repertoire that allows what, for an observer, seems to be a simultaneous instantiation of different behaviors, (for instance, walking and visually scanning for the activity of a moving con-specific). When two or more organisms are interacting only a part of this behavioral space may be occupied directly in the interactive activity (say, keeping mutual visual contact). However, if a coherence is observed between behaviors *not* involved directly in the interaction (say, walking together), then one is in the presence of coordination. Another, equally valid, interpretation would *not* require that the different behavioral domains, the one in which the interaction occurs and the one in which coordination is elicited, be *simultaneously* instantiated but coherence in the latter still needs to show dependency on the outcome of the interactive activity in the former. Such would be the case of the coordinated response of a group to an alarm call.

By coherence it is meant an observable agreement between behaviors of different organisms. Such an agreement may range from simple instances of synchronized activity or other types of temporal consistency (such as the group response to an alarm call just mentioned) to more complex cases such as the patterns of large prey hunting by members of a wolf pack or the approaching behavior and maintenance of the pair bond in monogamous species of tropical birds by means of antiphonal duetting (see section 9.2). One may ask why should there be any relation between

¹⁵The term “orientation” will be reserved, when possible, to refer in the model presented in chapter 9 to the *angular* orientation of moving agents.

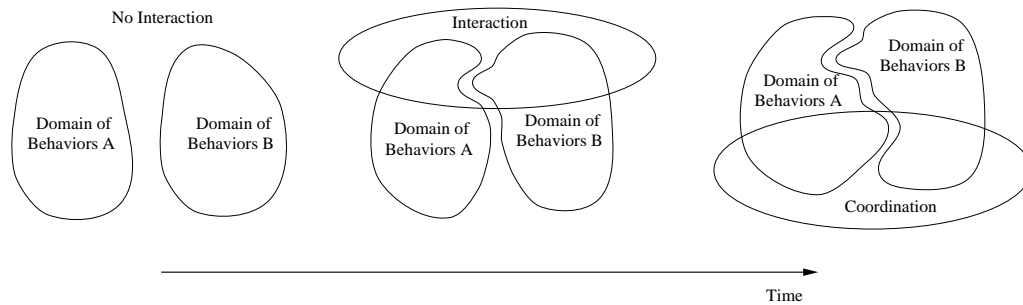


Figure 2.1: Illustration of the concept of coordination.

the coordinated behavior (in the last example, song synchronization and approaching) and the interaction (singing), unless both organisms were somehow congruent enough in structural terms so that 1) the coordinated behavior is *possible* for both of them, and 2) their structures are such that the coordinated behavior is somehow related operationally to the fact that they are undergoing a specific pattern of interaction.

See figure 2.1 for an illustration where the state of the behavioral domains of two organisms in interaction is shown through time. Each closed curve is meant to represent the behavioral domain of an autonomous system. Coherence is depicted by correspondence in shape between the two curves. Interaction is shown as a single activity in which both organism engage through structural coupling (top). Coordination is shown as additional coherence which depends (operationally) on the *existence* of interaction and the internal operation of each system (bottom part).

It would be preferable to avoid a possible reading of the above quotation which is considered misleading. Paragraph (b) in the quotation could be taken as saying that during coordination there are necessarily two different roles, the orienter and the orientee, and that the former orients the latter towards a interactional mode *previously specified* by it. This would entail an asymmetry in the coordinating systems, which need not be the general case. When the behavior of the orientee is changed towards a form of interaction which is comparable to the orientation of the orienter, this should generally be interpreted as the orientation of the orienter *once coordination has been achieved* and not necessarily before. The last case, although possible, is not general and could be wrongly interpreted as an intention on the part of the orienter in modifying the behavior of the orientee.

This is why in figure 2.1 the attainment of coordination is represented as a coherence in the behavioral domains of the participants in the “area” of their behavioral space (bottom) where a predefined disposition is not necessarily there in the first place. Figure 2.2, in contrast, shows a *particular* case where one of the participants already has a strong behavioral disposition towards a specific sub-domain of its behavioral repertoire and the other participant coordinates its own behaviors accordingly. Such asymmetries, although not the general case, are not uncommon either. A comment will be made below on one possible cause for asymmetries of this kind being a difference in ontogenetic stages between the coordinating organisms and their corresponding difference in susceptibility to plastic changes (e.g. in the case of mother and offspring).

Each perturbation that an autonomous system undergoes during structural coupling induces structural changes in it and some of these changes may be plastic. Plastic changes occur when

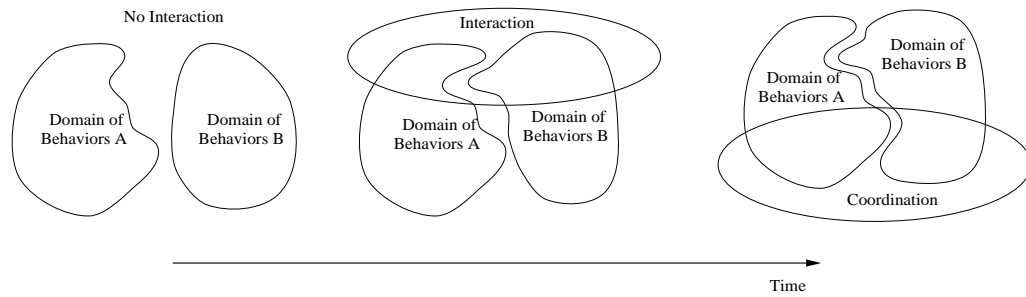


Figure 2.2: A special case of coordination through interaction where system B is oriented towards system A.

the structure of a system undergoes an alteration from which it does not recover within the same time-scale with which the change happened but with a much longer one. Clearly, some of these changes may be permanent.

The domain of coordinated behaviors established by two or more autonomous unities in structural coupling, during the course of which the systems mutually trigger in each other plastic structural changes, is sometimes referred to as a *consensual domain* (Maturana & Varela, 1980, p. 120) to emphasize the fact that although the operation of each system always depends only on its own structure, the observable behaviors arise as a consequence of an interlocked history of mutual structural perturbations between the systems; and so coordinated behaviors seem to depend on (or connote) this history. *Communication* as defined in the last quotation is coordinated behavior in a consensual domain.

The process of coordination in autonomous systems is, at the same time, a process of mutual selection of plastic changes in their respective structures, so that not only the ensuing behaviors result in a coherent pattern but also the corresponding structural changes may show some degree of coherence. The resulting relation between the structures of the coupled systems is known as *structural congruence* and it is to be found particularly between organisms that engage in interaction repeatedly and recursively. Sustained patterns of interaction tend to become embodied in the participants in the form of a history of plastic structural changes during each individual's lifetime. As a result, following encounters may be affected either in ways that facilitate the reproduction of the pattern of interaction or in ways that do not. If facilitation of future encounters is the result of certain patterns of interaction, it is clear that those patterns will tend to be conserved. For certain plastic systems this process could constitute the basis of social affinity, (an example is presented in chapter 9). In such cases quite complicated patterns of coordinated behavior can follow from very simple initial interactions just because the systems are already “tuned” to one another.

Of special interest is the case in which structural congruence is achieved between unevenly plastic organisms, as in the case of parent/offspring social interaction mentioned before. If structural congruence is understood as the meeting of two distinct, though not completely dissimilar, structures in some common ground, it is clear that those interaction patterns that facilitate their own reproduction by inducing structural changes, will tend to produce what for an observer would look like a *directed* structural change in the more plastic organism towards a structure that is congruent with the less plastic organism. This phenomenon and its opposition (the structural “rejection”

of patterns that make their own reproduction difficult) could be used as an operational basis for explaining many instances of social learning. This matter, however, will not be developed in the current context

It should be stressed that there is nothing magical about coordination. Consider one of its possible manifestation in rhythmic types of behavior: entrainment or synchrony. There is a growing literature on synchronization of coupled oscillators in biology and chemistry (see Winfree, 1980; Kuramoto, 1984, for “classic” introductions). The striking fact is that under a vast set of conditions synchronization is the expected result. Additionally, coordination in rhythmic behavior is not just manifested in phase-locking, but more remarkably in tendencies to correct phase deviations. Kelso calls this phenomenon *relative coordination*, (Kelso, 1995, p. 100 and ff.). Entrained behavior may be difficult to maintain in the case where the interacting systems are not very similar. However, the systems, under certain conditions, may manage to compensate phase slippage. Such is the case of an adult and child walking together at the same speed in spite of differences in their individual “natural” speeds. In spite of being able to provide an operational description, in many natural cases an observer would tend to *interpret* this compensating coordination as if the systems involved had an “interest” in maintaining certain types of interaction.

2.4.3 Changing the practical issues of interest: animal communication again

How much of a difference will the adoption of a theoretical framework such as the one described above make in practice? Is it possible that, after becoming convinced by arguments that the above is the “right” theoretical stance to adopt, both research interests and methods utilized will remain unchanged? These are important questions because they will help in deciding whether this is a purely argumentative exercise or whether it has pragmatic consequences for the way research is conducted. Some of these consequences will take the form of dissolution of problems which arose as part of a confusion between operational and functional discourse. But, perhaps the most interesting repercussion of adopting the systemic perspective is the opening of new issues that can be investigated within this framework but were traditionally by-passed because of the use of functional primitives in a pseudo-operational manner.

In order to explore some of these issues more specifically, one may concentrate again on animal communication. The given description of the phenomenology of social behavior and its explanation from a systemic perspective is fairly general and should be able to account for all types of social phenomena in autonomous structure-determined systems. To compare some of the practical consequences of this perspective with those of more traditional ones it is necessary to describe in more detail what is highlighted by the term “communication” within this framework and in what aspects it differs from the conventional understanding. This will be helpful for distinguishing some important shifts of interests.

Consider what would be the natural questions to ask when communication is described from the systemic standpoint in comparison to the discussion given in section 2.3.3. As mentioned before, a characterization of communication in terms of transfer of information is problematic in itself and especially so under the systemic perspective where autonomy is equated with operational closure and therefore interactions of the instructive kind (an animal telling the other one what to do) are not allowed if autonomy is to be preserved. So how is it possible to account for instances

in which it seems right to say that an animal is telling the other what to do and the other one is responding accordingly (say, a mother signalling for an infant to come closer)? Despite being a legitimate problem in the traditional view, this question was rarely asked because it never became apparent that there was a need to answer it. In the systemic perspective such observations demand a deeper investigations of the behaviors involved.

Another related issue that is highlighted is the issue of denotation. In the informational view, information is a stuff that is intended to play an pseudo-operational role. Such a role is given by the (to this point magical) ability of communicative behaviors to denote; so that it can be said that such and such a signal denotes an object, an emotion or a state of affairs. But if information is to be deprived of this pseudo-operational role, as the systemic view insists, then denotation, as a primitive operation, becomes mysterious. How can it be accounted for? Again, from the traditional standpoint there was, in all appearance, never a need to formulate this question. Both denotation and the apparent instructive character of certain types of communication were never evident as problems in virtue of the mixture of operationality and functionality of primitive terms like information. But once such use for this and other terms has been criticized in the systemic view, these issues become pressing .

Consider denotation. What sort of operation/behavior is it? Indeed, it is possible to understand it in systemic terms but this involves the recognition that it is not a primitive operation but requires an agreement (consensus) which can only be attained in an existing consensual domain, (Maturana, 1978, p. 50). An operational understanding of how such a consensual situation is achieved becomes *the* issue of interest. This point is repeated by Maturana and Varela in a comment on the evolution of language:

The understanding of the evolutionary origin of natural languages requires the recognition in them of a basic biological function which, properly selected, could originate them. So far this understanding has been impossible because language has been considered as a denotative symbolic system for the transmission of information. In fact, if such were the biological function of language, its evolutionary origin would demand the pre-existence of the function of denotation as necessary to develop the symbolic system for the transmission of information, but this function is the very one whose evolutionary origin should be explained.” (Maturana & Varela, 1980, p. 30)¹⁶

Similar themes reappear if one tries to make the case for a more formal use of the term information as reduction of uncertainty, as mentioned earlier. If one grants that certain communicative systems can be thought of as constituted by a sender and a receiver who are connected via a channel and for every change in the state of the sender, a single new state is generated in the receiver, then it must be assumed that both sender and receiver possess some degree of operational congruence because otherwise the correspondence between states could only be a consequence of instructive interactions (the emitter specifying the new state in the receiver instead of this new state being specified autonomously by the receiver’s structure) which must be discarded if autonomy is to be preserved. Two extremely dissimilar systems cannot be in this situation. Maturana (1978, p. 54)

¹⁶Maturana and Varela are able to explain denotation in systemic terms. That explanation, though, is beside the point that is intended here and too complicated to treat fairly. Maturana, in his solo writings and in a few collaborations, goes on to extend this explanation to account for the capability of linguistic creatures to become observers and self-conscious. The reader is referred to (Maturana, 1978, 1988a, 1988b; Maturana, Mpodozis, & Letelier, 1995).

describes the situation as a homomorphic relation between the domains of possible states in the sender and receiver. Such homomorphism must be assumed to exist by the observer unless she designed the system herself.

If such homomorphism, in fact, exists then any interaction can be trivially considered as an instance of communication and information makes sense as the degree of reduction in the uncertainty produced by the interaction under noisy circumstances.

“What is not trivial, however, is what takes place in the process of attaining communication through the establishment of ontogenic structural coupling and the shaping of the consensual domain. During this process there is no behavioral homomorphism between the interacting organisms and, although individually they operate strictly as structure-determined systems, everything that takes place through their interactions is novel, anticom communicative, in the system that they then constitute together, even if they otherwise participate in other consensual domains. If this process leads to a consensual domain, it is, in the strict sense, a conversation, a turning around together in such a manner that all participants undergo nontrivial structural changes until a behavioral homomorphism is established and communication takes place. These pre-communicative or anticom communicative interactions that take place during a conversation, then, are creative interactions that lead to novel behavior. The conditions under which a conversation takes place (common interest, spatial confinement, friendship, love, or whatever keeps the organisms together), and which determine that the organisms should continue to interact until a consensual domain is established, constitute the domain in which selection for the ontogenic structural coupling takes place. Without them, a consensual domain could never be established, and communication, as the coordination of noncreative ontogenically acquired modes of behavior, would never take place.” (Maturana, 1978, pp. 54 - 55)

Paradoxically (but, in fact, not so), it seems that one of the most interesting issues in the ontogeny and evolution of communicative behaviors is what happens *before* they become established. What sort of interactions can lead to coordinated behavior? How are they related to the other activities of the organism? How do these behaviors influence future encounters (for instance by promoting spatial proximity or by triggering structural changes that favour the repetition of the same behaviors)? Related to this, it is also possible to ask questions regarding the occurrence of novelty in existing communicative systems. It would be of interest to inquire about the nature of those interactions that somehow manage to “escape” from an existing consensual domain and produce a disruption which may eventually create a new domain of communication.

The question of origin becomes the question of the establishment, either in evolutionary or in ontogenetic terms, of the consensual situation that allows for communicative behaviors to exist at all. The original functional questions may still be asked, but the answers will be partly grounded on how the consensual domain is established in each case, i.e. giving ample space for issues related to embodiment and natural and cultural history to enter naturally in the picture. Similarly, the study of evolutionary change is enriched by the introduction of a mechanism that allows for novelty and creative interactions to take place and play a role that does not exist in the traditional framework. Pre-communicative behavior is historically antecedent to the establishment of the consensual domain and it is at the same time a constraint and a medium of realization of the structure of communicative behaviors. As evolutionary change can also occur in these pre-communicative stage, novelty is not restricted to a change in the repertoire of signals an animal uses, but it is

expanded to include non-communicative behaviors which can become part of a communicative event¹⁷.

So far this shift of interest seems to provide with positive additions to the set of questions that are relevant to investigate. This is particularly so because the original questions regarding the functional aspects of communicative phenomena and their selective contribution need not be discarded, although some of these questions will surely need reformulation. But is it also possible to identify what may look like “negative” pragmatic consequences in the adoption of the systemic framework? It is clear that a focus on the achievement of structural congruence and behavioral homomorphism through sustained structural coupling posits a greater emphasis on communication as a species-specific activity while many inter-specific cases of communication in the traditional view would probably need to be recast as non-communicative.

The emphasis in the systemic literature seems to be on the attainment of communicative situations purely at an ontogenetic time-scale, i.e. through sustained interactions in the life-history of individual organisms belonging to a group. It is because of this that the systemic perspective seems more focused on species-specific, ontogenetically attainable, forms of communication. However, it is possible to extend the operational arguments to account for similar phenomena occurring at an evolutionary time-scale although the transition is not trivial. This transition involves the fact that certain patterns of species-specific behavior become stable through successive generations by being selected or by influencing the likelihood of being themselves reproduced (see previous section). These patterns become identifiable (at least in principle) for an observer that is able to witness (or deduce) their resurgence generation after generation. As such, these patterns are therefore subject to changes which can be traced through evolution. However, the origin of those changes remains the same as in the ontogenetic scale. In the case of communicative systems, the origin of evolutionary novelty in communicative patterns is precisely the same pre-communicative interactions plus (re)-attainment of consensuality that generates novelty on the ontogenetic time-scale. This need not be interpreted as a form of Lamarckism because it is not being said that all ontogenetic novelty will necessarily re-appear in future generations, but that evolutionary novelty in communicative behaviors must first appear on the ontogenetic and behavioral time-scales and when they do, they do so in the form of pre-communicative interactions as described in the last quotation by Maturana.

In principle, similar arguments could be advanced for the case inter-specific patterns of communication. But a general treatment here is more difficult in virtue of the inherently dissimilar structures frequently presented by organisms of different species. The domains of consensuality must rely on the degree of potential congruence between the participants in communication. To achieve this across species requires at least some pre-existing dimension where these domains can be attained. It is indeed possible to observe this across species sharing a common habitat and a joint history. An example close to home is the ability of human beings to establish communicative domains with other animals such as dogs whose own social codes are easily interpretable even by small children and for whom many human moods and intentions seem to be quite transparent.

¹⁷The closest that the traditional framework comes to a treatment of these issues is in the discussions regarding signal ritualization (Krebs & Dawkins, 1984) as originating in pre-existing non-signalling behaviors thus following the principle of *derived activities* (Tinbergen, 1952). Accordingly, some present day signals in birds are derived from non-signalling activities such as preening or feather-settling (Krebs & Dawkins, 1984).

2.5 Summary

The reader may have noticed that many of the presented theoretical concepts regarding the logic of autonomous systems would benefit from a development deeper than the one it was possible to offer here. Moreover, much of this conceptual framework has found a more sophisticated expression in recent years, especially in connection with what is known as the enactive approach to the cognitive sciences (Varela et al., 1991) as well as in specific areas of neuroscience and theoretical immunology¹⁸. Enaction combines the main precepts of biological autonomy with another fundamental ingredient: embodied experience. A whole new angle is thus opened for the study of cognition, which breaks with the traditional computational approach. In spite of the lack of an extensive treatment in these pages, these ideas should not be interpreted as exceeding the thematic interest of the thesis. Quite on the contrary, there is a whole area for fertile development of issues regarding social behavior, cultural history, language and intersubjectivity within the framework provided by the biology of autonomous systems and the enactive approach. These issues, however, do exceed the practical limits of the thesis and therefore they are only mentioned here.

It was practical concern as well which prompted the claim that any framework used for studying social behaviors should be provided with a proper biological grounding intended to act as a constraining rather than as a reductive factor. What kind of biological grounding became clear after the introduction of the distinction between operational and functional discourse. Only after understanding the possible relations between these two modes of describing and explaining it was apparent that operational statements cannot in general work as a replacement of functional ones, but they can indeed be used as constraints to the latter by requiring that a functional account be in accord with what is operationally known about the systems giving rise to the phenomenon of interest. Functional explanations and descriptions, in contrast, can only exert a much subtler influence on operational discourse, by suggesting appropriate levels for distinguishing the systems and components which generate the phenomenon. It is because of this asymmetry that only operational explanations are deemed appropriate to fulfill the task of grounding social behavior in biology.

The actual need for providing this grounding was exemplified by the current state of confusion surrounding the issue of animal communication as viewed from the classical framework. Some methodological and conceptual problems of not properly distinguishing the two modes of discourse were examined and it was found that the main defining feature of communication was not an operational one but a functional one: selective advantages. The behavior proper was only intuitively characterized by making use of terms like signals and information. Ironically, much of this intuitive grounding is lost when an attempt is made to formalize these concepts *in terms of* selective functionality. This results in logical contradictions as well as in a poor characterization of communicative behaviors. The consequences of this go beyond the lack of a consistent theoretical framework, something quite grave in itself, and are manifested in the way the origins and evolution of communication are addressed in practice and the sort of questions that are asked and the sort of questions that are avoided. Thus, the evolution of communication is mainly viewed as

¹⁸See for instance (Varela & Stewart, 1990; Varela & Coutinho, 1991; Varela, Stewart, & Coutinho, 1993, and others) for an approach to immune networks based on autonomy and dynamical closure, (Thompson, Palacios, & Varela, 1992) for an enactive approach to color perception and (Neuenschwander & Varela, 1993) for an example in the neurosciences.

the evolution of an appropriate selection of a repertoire of actions or signals within a pre-existent behavioral context whose origins are rarely questioned.

Fortunately, the systemic perspective on the biology of autonomous systems has provided an appropriate body of operational theory that serves the intended purpose of grounding social behaviors biologically. After explaining key concepts such as the difference between structure and organization, autonomy, structural coupling and adaptation, it has been possible to introduce the notion of coordination as the achievement through interaction of a consensual state of orientations in the behavioral domains of the systems involved. Thus, systems undergoing sustained structural coupling are able to mutually trigger changes in their respective structures in ways that certain homologous areas in their spaces of interactions become coherent in spite of not being necessarily involved in direct coupling between the systems. Such coordination observed in the behavioral domain is accompanied by a state of congruence in the structural domain of the systems.

Communication, defined as behavior in a consensual domain established through coordination, can therefore be characterized in operational terms. An interesting consequence for research of adopting this definition is that it highlights a set of questions which were not actively pursued in classical investigations while, at the same time, many of those questions that were pursued can still be asked, perhaps after being re-formulated. Thus, the focus of interest is shifted from questions regarding what sort of action or move in a game is going to be selected to fit an existing behavioral slot, to an inquiry into how do those slots come to exist in the first place, be that in behavioral, developmental or evolutionary terms.

The presentation of the systemic view on social behavior, as well as the justification of its scientific potential, is intended partly as a goal in itself. This view will also be used as a guide for the concrete investigations offered later in this thesis. However, some of these investigations are not necessarily rooted directly on the systemic framework. An actual direct test of its suitability for research will have to wait until chapter 9 where the notions of coordination and structural congruence will play a fundamental explanatory role.

The rest of the concrete models will explore other derived aspects of the general systemic framework, which have to do with the dynamics of complex systems, historical processes, and self-organization, issues which are further developed in the following chapter. Those ideas will be specifically applied to evolutionary, ecological and developmental aspects of social coordination in simple games. Some (but not all) of the questions asked within those models could well be fitted within a classical framework, which is not a proof that the systemic view has been abandoned but evidence of its more encompassing scope.

Chapter 3

Historical processes

This chapter introduces a series of ideas which will together constitute one of the recurring themes of this thesis. This theme will be manifested in the studies presented in chapters 6, 7 and 9. It will be found that a common feature in all these cases is a mode of explaining the results which appeals to certain characteristics of the dynamics of the processes involved, namely that these are historical processes.

But what does it mean for a process to be historical? The aim of the present chapter is precisely to explore this question and evaluate the relevance of this concept in a series of examples, some of which are directly relevant to the dynamics of social coordination.

To this aim, the idea of what constitutes a constraint to a process will be examined, as well as how it relates to the dynamics of the process in operational as well as in explanatory terms. This will permit a specialization of the word “historical” to processes which are not merely contingent, or susceptible to be explained in genetic or historical terms, but to processes which, in addition, are able to introduce some temporal heterogeneity in their own unfolding due to the interplay of variations at different time-scales.

It will also be shown that one special feature of these processes is the presence of *organizational invariants* which are at the same time constituted by the previous dynamics and constituting of the future dynamics. Such features can provide a powerful explanatory tool when dealing with complex historical processes.

Some of the concepts presented in this chapter are related to the ideas of scientists like Waddington, Bohm and Prigogine who were in turn influenced by Whitehead’s metaphysics of process, (Waddington, 1969, 1975; Bohm, 1980; Prigogine & Stengers, 1984). However, the purpose of the chapter is to make a basic presentation of these concepts in order to facilitate their subsequent practical use and not to provide a review and comparative exposition of the philosophical and scientific extend of these ideas.

3.1 From homogeneous time to historical time through self-organization

There are many senses in which the word “historical” may be applied to a dynamical process. For instance, a process may be so called if its unfolding involves a set of contingencies which cannot

be predicted until the moment they occur. In that sense, it could be argued that all real dynamical processes deserve the name historical since all of them involve at least a set of initial conditions which are not specified in the laws of evolution of the process, i.e. they cannot be predicted by these laws. While the word “historical” would perhaps be inappropriate in such a general case, its use would probably be accepted if contingent and unpredictable factors were at play not just in the initial conditions but also at different stages during the unfolding of the process. Such factors could take the form of discrete events (as in the case of founder effects or catastrophes in biological evolution) or they could operate with constancy, in which case their effects may become manifested over long periods of time (e.g. random fixation of alleles in the gene pools of finite populations due to genetic drift).

Another related criterion would consider adequate to apply the name “historical” to a process if an explanation of what goes on in it and how its current state has been attained would be best given in historical terms. Such genetic or historical explanations [see (Nagel, 1961, pp. 25 - 26) and (Gould, 1989b, pp. 283 - 284)] would account for a state or event in a process in terms of previous *key* states or events. A chain of these events would be understandable if it is possible to understand the connection between one link and the next. It is possible that, due to some of these events being contingent, the chain could have been different. As a consequence, a genetic explanation of this type can mainly account for a given “instantiation” of a class of processes. It is possible, however, to use genetic explanation referring to a class of processes if they can be thought to reliably undergo a series of stages (e.g. in the case of development of species-specific behavioral patterns). The value of this type of explanation in understanding the genesis of intelligent behavior has been recently highlighted by Hendriks-Jansen (1996, ch. 6).

The word “historical”, in the current context, is not intended strictly in any of the above senses. Rather, a historical process would lie at the intersection between the cases just mentioned (i.e. contingent or noisy processes and processes explainable in historical terms) and the set of processes which are sometimes characterized as *self-organizing*. Such historical processes are indeed irreversible, contingent and probably most of them afford historical explanations. However, the key feature to be highlighted is the capability of the process to dynamically influence its own constraints and thus to introduce an interplay between dynamics at different time-scales which may result in open temporal inhomogeneities. In order to explain this capability the concept of constraint needs to be expanded.

3.1.1 Constraints

All observable events and processes are underdetermined by the fixed universal laws which presumably are at play in them. This is so even in the idealized case in which all randomness is taken away, as well as any confusion about which laws are relevant. The trivial reason for this is that such laws can only be universal because they are disembodied and refer to no actual system or structure in particular. In order to apply them to the understanding of a real process a description must be provided of how these laws are constrained by the actual structures and conditions that make up the process of interest.

There are two senses in which the word “constraint” may be used. One is the specific sense of the collection of relations between a system and other structures that place limitations to the

unfolding of a system's dynamics. The other more general sense includes these limitations *and* a description of the structure of the systems and their medium as well.

Consider a physical pendulum. A finite mass of finite size is hung from the ceiling by a piece of string. After some idealizations, a description of this physical system could be offered that would permit the application of an universal law, namely Newton's second law. Thus, considerations about the size of the mass, the elasticity of the string, the atmospheric conditions in the room, the stability of the ceiling, etc., would permit a description in terms of a zero-dimensional particle hanged from a fixed point by an inelastic string under the exclusive influence of gravity, etc. In mechanical terms the word constraint is reserved to those relations that place direct limitations to the variation of the coordinates with which the system is described, (see for instance Goldstein, 1980). For the pendulum, such a constraint is found in the position of the particle which must, at all times, conserve its distance to the point in the ceiling from which it is hung.

In the second, more general sense, the word "constraint" can be applied not just to these relations but to the set of parameters and other relations that make it possible to embody a universal law into a description of an actual system. If the system remains ideally isolated and such contextual factors remain fixed, it seems that calling these factors constraints would be unnecessary. However, the meaning of the word is recovered when one considers that the system may participate in coupling with other systems with their own dynamics which, through their effect in such contextual factors, may influence the system's behavior. Thus, vibrations of the ceiling which originate independently of the movement of the pendulum may affect its trajectory and so they would have to enter into the description as an additional time-dependent boundary condition. Or, after a number of oscillations, the length and elasticity of the string may change depending on the accumulated tension, i.e. on the previous trajectories; then the system would have to be re-described to take into account this effect of wear.

It is clear though, that any addition of new boundary conditions or any re-description will end up with a new fixedly defined system and a known relation to its environment. Such a tendency for re-describing actual systems is obviously limited since future changes in the contextual (and internal) conditions need not be predictable either because of random factors or because of unexpected effects of the dynamics on the conditions which granted validity to the initial idealizations.

In view of this, it makes sense to associate all these contextual factors and a description of the internal structures of the systems involved in a process under the single name of constraint. In this more general sense, the word "constraint" intends to indicate any factor which may exert some influence on the unfolding of a process as described by some generalized dynamical principle.

In a way, this meaning of constraint is a generalization of the meaning favored by Gould (1989a) for the case of evolution. According to him, a constraint is "theory-bound term for causes of change and evolutionary direction by principles and forces outside an explanatory orthodoxy", (Gould, 1989a, p. 519). Thus, any source of change apart from the general explanatory framework for the type of process in question would qualify as a constraint. In particular, being Darwinian theory functionalist in nature, it is not surprising that constraints to evolution should belong to the complementary historicist and formalist worldviews.

Readers who are familiar with the work of Howard H. Pattee will also have noticed certain similarity between his idea of constraint as an alternative description of a process and the concept

as presented here, (see for instance Pattee, 1970, 1971, 1972). Pattee uses the term “to characterize whatever auxiliary conditions must be appended to the fundamental equations of motion in order to predict more easily how a system will behave”, (Pattee, 1971, p. 264).

The term loses the negative connotation of the more formal notion of constraint as limitation and acquires a more encompassing meaning which may include the senses of direction or canalization, (see also Gould, 1989a, p. 518). The word will be used in this general sense in what follows.

3.1.2 Invariants

Although, as seen above, constraints are not necessarily fixed, one could tentatively distinguish their variations from the actual process by one of the following two conditions: these variations are independent of the operation of the system or, if they vary dependently, they do so at a much slower time-scale so that, at the scale in which the changes of state of the system occur, constraints may effectively be considered fixed. It can easily be seen how these conditions are qualitative rather than strict.

In the first case, influence on the process is exerted through coupling with other processes which operate independently. But such coupling may also reflect how those contextual processes were in turn previously influenced by the process in question. Thus, the oscillations of the pendulum may have produced structural changes in the ceiling which have altered its own pattern of oscillations (which is presumably generated as a result of a different process), and this may be reflected back into the movement of the pendulum.

In the second case, when variations in the constraints depend directly on the dynamics of the process, one could question what is exactly meant by a much slower time-scale and why are not such changes included as part of the actual process.

It is necessary to have a more formal criterion. In a way, this issue is a manifestation of a bigger problem which has been avoided so far. If the dynamics of a process may alter the constraints that define the structure of the systems involved and their contextual relations, is it not possible that these systems could change so much that they would effectively become different systems with different dynamics? In such case, with what right can one speak of unique and well-defined process? A fixed set of constraints used to do the job of assuring that the systems remained the same from one moment to the other and therefore one was able to speak of their dynamics as a single process. Such rigidity, however, entailed that no process involving some sort of innovation could be so described.

But if the constraints can also change there must be something else that one can point to in order to be able to say that one is referring to a same process. There must be an *organizational invariant* of the process which maintains certain relations fixed. Such invariants were already discussed in chapter 2 when it was mentioned that living systems are characterized by an autopoietic organization and a relation of adaptation to their environments which remain constant as long as the system exists as a living entity. This is so even when the structural changes undergone can be so radical as the changes from an embryo into an adult organism. The living entity has an identity which is given by its invariant organizational features.

This idea could be extended to processes in general. A process can then be *defined* as the

dynamics of a set of systems whose actual structures, rules or laws of operation and coupling relationships conserve some global organizational feature unchanged. For instance, one could describe the set of basic organizational features in the example of the pendulum. Among these one could include the applicability of Newton's second law in a gravitational field, the relative positions between hanging mass, string, and ceiling, the very existence of these components, etc. One could imagine that if the string is chemically unstable in air it will break at a certain point. When this happens, the process, *as defined by the above invariants*, has ceased.

There is clearly certain freedom of choice on the part of the observer regarding what is to be called a process. That freedom is in the distinction of the relevant invariants. Thus, if the only invariant in the case of the pendulum is the mass that hangs and the process is the variation in position of this mass, then it does not matter if the string breaks, this is just a change of constraints, the process goes on with the free fall dynamics, the bouncing on the floor, etc. *This* particular process would cease only if the mass disintegrates.

The identity of a process is defined by the organizational invariants chosen by the observer. This identity may differ from the individual identities of the participating systems. These systems may change structurally while conserving their own organization but they may also change in ways that their individual organization is lost, thus being destroyed or transformed into something else. In this case, one would be entitled to speak of a single process only if the defining invariant were some global organizational feature in the set of systems which remains unchanged. Thus, processes involving living systems could be identified by the same organizational invariants used to identify these systems individually but, then again, they may not. In the case of an ecosystem, in which organisms are born and die and species may even become extinct, other organizational relations must be conserved between the constituting organisms to be able to identify an ecological process as a well defined one.

These comments apply to processes in general, but they hold a special significance for historical processes, as these are the only processes in which, besides the basic invariants distinguished by the observer, the interplay between process and constraints may lead to the *spontaneous* formation of new invariants.

Such spontaneous invariants can be manifested in actual structures and relations which influence the dynamics of the process. At the same time, these invariants are constituted by an interplay between the dynamics and the constraints of the process. The form assumed by the new invariants can be explained as emerging from the influences of the dynamics on the constraints and as constituting in themselves further constraints to the process. In such cases, one may speak of the process as exhibiting some sort of spontaneous organization. Due to the amplification of the effects of fluctuations and the breaking of in-built symmetries, complex processes in which many systems interact non-linearly may exhibit non-equilibrium transitions to highly ordered (lower-dimensional) dynamics. This transition is manifested in the emergence of a coherent regime which is not pre-specified in the initial definition of the process nor externally imposed. Such processes are sometimes also called self-organizing, (Haken, 1978; Prigogine & Stengers, 1984; Kelso, 1995)¹.

¹This word must be used carefully. As noted by Ashby (1962) a self-organizing system can only be so if it has been badly described in the first place. Maturana also suggests that the term is contradictory if it is to be used in its strict sense. A system can never be *self*-organized if its identity is precisely defined by its organization. A more relaxed use

After a transitional phase, a self-organizing historical process arrives at a situation in which new organizational relations have been formed. The state of order is manifested in the form of patterns in the dynamics and in relations within its constraints. These patterns, in turn, influence the dynamics of the process in ways that tend to their own perpetuation. In addition, these organizational features may also exert an influence over other aspects of the process which need not be directly involved in the conservation of the invariant, (detailed examples will be provided in chapters 6 and 7). From this perspective, it is possible to say that a self-organized invariant may be able to “explain itself”.

3.2 Different manifestations of history

The above considerations give a rough idea of how to differentiate historical processes from processes which are non-historical or merely contingent. A historical process is a process subject to fluctuations whose dynamics affects its constraints either directly or through recurrent coupling with other processes. In such a process new invariant features may be spontaneously manifested as organizational properties which are both constituted by, and constituting of the dynamics of the process.

It does not suffice for a process to be contingent in order to be historical in the sense intended. To return to the original example, the pendulum, it is possible to describe a process which would be the former but not the latter. If the pendulum is perturbed by random horizontal impulses but everything else remains the same, the process would depend on contingent factors and a given state would have to be explained (in the conservative case) by a combination of dynamical laws, constraints, and historical events. In contrast, a historical process is such in virtue of the temporal inhomogeneity that it is able to introduce by affecting its own constraints and consequently its own dynamics. One could imagine a situation in which a massive pendulum influences the vibration pattern of the ceiling so that both of them become entrained in a well-defined oscillation mode which is not the pendulum’s or the ceiling’s natural one but a property of the whole process, i.e. a spontaneously formed invariant.

In order to make the meaning of these concepts clearer it will be helpful to consider some examples of historical processes.

3.2.1 Trails on grass and Pask’s artificial ear

Consider the trails made naturally by pedestrians on areas which are covered with grass. These trails are made by the action of walking which makes it difficult for grass to grow on zones which are frequently trodden upon. The lack of grass makes walking along the trail easier and people tend to use the trail rather than cut across the grass, even if this implies a small deviation from the optimal route to their destination. Trail formation has been recently studied using a very simple and powerful model, (Helbing, Keltsch, & Molnár, 1997). The process is self-reinforcing and, in the bigger picture, it is also a historical process.

Let the process be the set of individual pedestrian trajectories within a piece of land covered with grass (say a square) with a few preferred entry and exit points. Walkers are driven by two preferences: they want to arrive at their destination cutting across the square and they prefer to

of the word is admissible, if one wants to describe a process in which some degree of order emerges spontaneously.

walk where the grass is not grown. Initially, no path is marked on the grass and walkers choose a direct route to their destinations. As time passes, and for a certain frequencies of crossings, the effect of the initial trajectories will begin to be manifested in areas where the grass is worn. In the most used trajectories the effect of wear will be so much that the grass will not be able to compensate by growing again before the path is re-used. Thus, trails are formed and maintained in a dynamical equilibrium. The process can be quite complex since the different trails may “interact” during the process. For instance, it will be more common to observe a single exit point halfway between two frequently used and relatively close destinations than an exit point corresponding to each one of them, which means that two trails may have converged.

Once a pattern of trails is formed the history of the process has become partially embodied in it and walkers are constrained by its shape (and their preferences) to walk along the trails. Thus, the pattern acts as a constraint to the dynamics of the process but, at the same time, it is constantly being constituted by the process as trails can only be maintained if enough people use them.

The pattern of trails is the structural manifestation of an spontaneous organizational relation within the process which arises from the combination of a fast and slow dynamics (respectively: the trajectories of the walkers and the growth of the grass). Notice, that the existence of trails is evidence of this set of relations that remains invariant (relations between time-scales, form and connectivity of trails, etc.) while the actual spatial pattern may change in response to contingent events such as the appearance of a new destination.

A very similar process was used by cyberneticist Gordon Pask for the construction of artificial sensors and effectors out of an initially undifferentiated physical system, (Pask, 1959, 1960). The system consists of a network of amplifiers and associated electrodes which are not directly connected but are submerged in a solution of ferrous sulfate. The electrodes may act as sources or sinks of direct electrical current depending of the activity of the system. Crucially, if direct current is passed from a source to a sink, a metallic thread of very low resistance is formed in the ferrous solution which, as the trails on grass, will be much easier to use if current is to pass again between the same electrodes. In contrast, if the thread is not re-used, it will gradually dissolve because of local acidity. After some time, a network of threads may be formed and maintained dynamically.

The system can be “trained” to respond in a desired way to different sorts of couplings. The method of training consists simply in increasing the available energy for forming and reinforcing threads if the system’s performance is close to the desired one. Such a scheme is translated into a growth and pruning dynamics at the level of the network of threads. Interestingly, being a physical system, there are many ways in which the process of thread formation could be affected: mechanical, thermal, chemical and electrical. In particular, Pask was successful in training the system to respond to acoustic vibrations of a specific frequency. The system responded by growing a network of threads around the vibrating regions of the apparatus.

3.2.2 Polya’s urn scheme

Consider the following stochastic process known as Polya’s urn scheme (Blackwell & Kendall, 1964; Port, 1994). Put two balls in an urn, one red and the other one black. Extract one of the balls, observe its color and then replace it and put another ball of the same color into the urn. Repeat indefinitely. What is the expected probability for extracting a black ball after a large




Figure 3.1: Probability of drawing a black ball (Pb) from Polya's urn after vs. number of iterations for 10 different runs starting with one black ball and one red ball.

number of iterations?

This process was originally proposed as a model of epidemics. Interestingly, it can be shown that the probability of extracting a black ball will converge to a specific value which can be *any* number between 0 and 1. Figure 3.1 shows this convergence for 10 different numerical instantiations of the process.

The process may be understood as historical if its dynamics are taken to be the extraction, observation and double replacement of balls in a repeated manner. At any moment, the probability of extracting a black ball depends on the number and color of the balls present in the urn. This is taken as the context or constraint of the dynamics. Such context is itself affected by the same process that it constrains. After many iterations, this interplay between dynamics and constraints reaches a point of equilibrium. This is because the addition of a new ball, whatever its color, will not affect too much the existing distribution within the urn and the accumulated set of added balls will tend to reflect this distribution over a number of iterations. The actual equilibrium, however, is strongly dependent on the history of the process. In particular, much weight is given to the initial steps; figure 3.1 shows how the variations can be very extreme in the initial 10 iterations, then more moderate in the next 100 iterations and from then less and less significant.

This example shows that historical dynamics may be instantiated in processes which are relatively simple.

3.2.3 Evolution

Biological morphogenesis and evolution are historical processes *par excellence*. However, giving a full account of why this is so for each case would take too much space. For this reason, only some key aspects of why evolution is a historical process will be commented on. For important accounts of the historical and self-organizing aspects of ontogenetic processes the reader is referred to (Oyama, 1985; Goodwin, 1993, 1994; Webster & Goodwin, 1996).

The historical nature of evolution is rarely denied, although there is a tendency to think of

evolution as historical only in the sense of being a process subject to contingencies. These may take the form of “frozen accidents” or they may indeed be the result of the accumulation of small events, as mentioned earlier.

However, the current neo-Darwinian perspective tends to downplay the role of historical factors to that of contextual or initial conditions in a process subject to an universal law of fitness maximization². The process in itself is not considered historical in the sense given here to that term. As argued below, this is a simplification. But how is this simplification achieved?

This is possible thanks to two key assumptions. First, that there is a direct correspondence between the material an organism inherits from its parents during the physical process of reproduction and the phenotypic features of the organism that will grant it a good chance of survival and reproductive success. Second, that, through a process resembling trial-and-error, random changes in this inherited material will affect differentially the match between organismic and environmental features so that some variations will be selected as better adapted to the environment than others.

Even though the first assumption can be attacked from different flanks, it is the second assumption that allows the historical nature of evolution to be relegated to independent contextual factors. In effect, evolution by natural selection becomes a process of optimizing the adaptation to a pre-existent environment which does not depend significantly of the evolving organisms.

The crucial error in this view, as pointed out on many occasions (Bohm, 1969a, 1969b; Lewontin, 1982, 1983; Levins & Lewontin, 1985; Odling-Smee, 1988; Maturana & Varela, 1987; Varela et al., 1991, among others), is that the key environmental features that are significant for the reproductive success of an organism are not independent of the organism itself. “The world external to a given organism can be partitioned into *a priori* ecological niches in a non-denumerable infinity of ways. Yet only some niches are occupied by the organisms. How can we know, in the absence of the organisms that already occupy them, which of the partitions of the world are niches?”, (Lewontin, 1982, pp. 159 - 160). A similar point was nicely made by Jacob von Uexküll (1934) in his descriptions of how similar scenes represent different worlds of perception and action from the points of view of different organisms. “Perceptual and effector worlds together form a closed unit, the *Umwelt*”, (Uexküll, 1934, p. 6). Thus the tick responds only to a few features of the environment; the butyric acid that emanates from a passing mammal prompts the tick to drop itself from a position at the top of a branch into the animal’s fur, this releases a tactile cue which prompts movement until heat is detected which prompts boring. Of all the possible sources of environmental events and of all the possible actions that the tick could perform, only these basic stimulations and actions constitute the tick’s *Umwelt* which does not pre-exist the tick.

In addition to “choosing”, rather than adapting to, their own niches, organisms also alter their medium, and that of other organisms, in significant ways, (Lewontin, 1982, 1983; Odling-Smee, 1988; Jones, Lawton, & Schachak, 1994). Birds and social insects build nests and other structures, rabbits and rats dig tunnels, beavers create ponds and alter local water levels, leaves accumulate under high plants, etc³. These alterations may have both short and long term effects. The cases just mentioned would qualify mainly as short term. An often cited example of a long term effect is the presence of oxygen in the atmosphere and the seas which, due to its chemical instability,

²For classical expositions of neo-Darwinism, see (Maynard-Smith, 1969, 1993; Dawkins, 1982).

³A longer list with additional references can be found in (Jones et al., 1994).

cannot be explained unless it is constantly being renewed (by algae, plankton, trees, etc.). As a byproduct of life, a high concentration of oxygen is currently a basic condition for the existence of most animal species.

In spite of the mutual inter-dependence between organism and environment, both in defining each other and in affecting each other, neo-Darwinism has been able to approximate evolution as a non-historical process by introducing a separation between them and sweeping all contingent factors under the carpet of independent contextual variation to which organisms must adapt. This amounts exactly to a separation of the dynamics of the process and its constraints which is characteristic of non-historical processes. It is, therefore, not surprising that the problems related with innovation in biological evolution cannot be accounted for in the neo-Darwinian orthodoxy, (Crutchfield, 1994a, 1994b; Fontana, Wagner, & Buss, 1994; Fontana & Buss, 1996; Kauffman, 1993), since such innovation can only take place in historical processes.

In practice, most neo-Darwinian explanations represent the environment as either fixed or independently varying (usually in uniformly random ways, Bohm, 1969b). However, the neo-Darwinian perspective is not as naive as it may sound. Indeed, it allows a partial relaxation of this scheme, for instance, when studying co-evolutionary dynamics (prey-predator or host-parasite arms races, symbiosis, sexual selection, etc.) or frequency dependent effects. Even in these cases, however, the guiding principle remains the same. For instance, a game-theoretic analysis of different foraging strategies will probably find a strategy which is evolutionarily stable in the face of any other combination of strategies, (see chapter 5). For calculating the value of playing a given strategy one must take into account the current distribution of competing strategies played by con-specifics; at the same time, such a distribution will be modified by the actual success of each strategy. The study, however, will be unlikely to care about variations in the distribution and patchiness of food sources which is presumably going to be affected by which strategy is used and will in turn affect the benefits of using one strategy instead of another one.

While neo-Darwinism has a problematic relation with the historical aspects of evolution, the evolutionary perspective which can be derived from the systemic view presented in chapter 2 incorporates history in a natural way. Organism and environment, while distinguished as systems, mutually influence each other by means of perturbations in a process of structural coupling. As long as the relation of adaptation is conserved, i.e. as long as the organism does not enter in a relation that will cause its autopoiesis to be lost, both the environmental and organismic structures may change. These accumulated changes in a chain of autopoietic systems linked by reproduction and in their corresponding environments perform a historical trajectory without a pre-defined direction in the space of allowed structures. Evolution is equated to a process of *natural drift* (Maturana & Varela, 1987; Varela et al., 1991). From this viewpoint, the fit between the structural (and behavioral) characteristics of an organism and features of its environment is an effect of historical evolution and not its cause.

Selection acts negatively when the relation of adaptation is lost, but it also acts through the process of structural coupling. Notice, however, that this viewpoint is more general than the traditional understanding of evolution as guided mainly by differential reproduction, as it allows evolution to proceed by different sorts of dynamical couplings between organisms and environment. Thus, if similar species must evolve in a confined region without much variety of resources,

evolution by differential reproduction (i.e. fitness maximization) will be the result of natural drift since all the evolving organisms have similar niches and all of them are directly coupled to a finite local resource. However, if more variety of resources is introduced, or if the habitat is not spatially confined, then the evolutionary outcome may well be avoidance of competition and branching of the originally similar niches in dietary or spatial terms.

The outcome of natural drift may take different dynamical forms under different circumstances, as long as the relation of adaptation between organisms and environment remains invariant. Additionally, on an ecological scale, other invariants may also appear, such as the conservation of ways of behaving, social structures, affiliative networks, food webs, etc. Some examples are described in more detail in section 7.1 for the case of social influences on development. Factors such as these may give the evolution of social traits a strong historical dynamics as the particular examples and model presented in chapter 7 will show.

The relation of co-determination or co-definition between evolving organisms and their environments acquires additional significance when the subject of study are social behaviors. There is a degree of truth in the metaphor that social traits and behaviors constitute, to a certain extent, their “own environments”. This is true in the sense that the significance of a social behavior depends on the existing pool of social behaviors in the group in which that behavior is executed. A mutation that allows an individual rabbit to understand and speak fluent Italian is devoid of biological significance, unless some fellow rabbits are at least able to respond in some way to orders, requests or serenades in Italian.

On a more serious note, a conclusion that may follow from this appreciation is that, given that, at least partly, both “dynamics” and “constraints” have become so intertwined in the evolution of social behaviors as to be represented in the same social traits, it seems natural to expect the resulting process to exhibit much richer dynamical characteristics, probably in the form of fast changes in comparison with the evolution of non-social traits. This is not necessarily so. Consider the situation in which a group of organisms live in a condition where food is scarce and energy conservation is a priority in their behaviors. Suppose that these organisms have evolved a warning call *A* by which they alert each other of the presence of a predator. *A* is adequate for its purpose, but it is very inefficient in energetic terms and, under the condition of energy scarcity, the organisms would benefit from using *B* as an alternative warning call. In these circumstances, if the call were a non-social trait, it would suffice with the appearance of a mutation in single individual, so that it would tend to make it produce *B* instead of *A*, for *B* to become the new signal after some generations. Given that it *is* a social trait, a single mutation may be insufficient, since the call is only effective in terms of the responses it elicits in other individuals. For these individuals, *B* may be completely insignificant or, even worse, misleading. In general, any change in a social trait must be more parsimonious with the existing situation than changes in non-social traits. For these changes to be conserved, they must be at least quasi-neutral; a new warning call *A'* which is a bit more efficient than *A* and which is interpreted most of the time as *A* will have a better chance of becoming fixed than the more efficient *B*. The resulting evolutionary trajectory may never come near the production of *B* even if it happens to represent an optimum of some sort.

In conclusion, the recognition of the historical nature of evolution opens a wider dynamical perspective that permits the understanding of specific instantiations of evolution not as the combi-

nation of some universal principle and a set of contingent circumstances, but as processes which exhibit their own dynamical characteristics which depend on their history and the history of other concurrent processes. In the words of Lewontin: “[w]e do not further our understanding of evolution by general appeals to ‘laws of nature’ to which all life must bend. Rather, we must ask how, within the general constraints of the laws of nature, organisms have constructed environments that are the conditions for their further evolution and reconstruction of nature into new environments”, (Lewontin, 1982, p. 163).

3.2.4 Embodiment, culture and *habitus*

Of all the historical processes perhaps the most obvious, and the one that should be approached with greatest care, is the process of human history. There is little doubt that cultural and historical changes can be understood as a process in which the lower level dynamics constitute and, at the same time, are constituted by the habits, institutions and traditions inherited from previous cultural activity. An initial search for candidate invariants of this process would find two possibilities which correspond roughly to two important trends by which social practices have tended to be objectified, namely the structuralist view according to which social behaviors are regulated by the logic of existing global structures and the nativist view of the social agent as a rational actor “without inertia” (Bourdieu, 1990, p. 46) by which the logic of a social system is but the combination of the logic of a large number of rational optimizers. These two extremes run into problems because they forget about history being a process and prefer to describe it in term of achieved universal factors. In the words of Pierre Bourdieu, objectivism, in any of its extreme forms, concentrates on the *opus operandum* rather than on the *modus operandi*. The philosophy and social anthropology of Bourdieu, following acknowledged influences by Merleau-Ponty and Wittgenstein⁴, steers a careful middle course by recognizing the inherent dynamics of human practices, (Bourdieu, 1977, 1990).

Bourdieu has even coined a name for the organizational invariant which results from the way practices become embodied both in actual bodies and in institutions: *habitus*,

“systems of durable, transposable dispositions, structured structures predisposed to function as structuring structures, that is, as principles which generate and organize practices and representations that can be objectively adapted to their outcomes without presupposing a conscious aiming at ends or an express mastery of the operations necessary in order to attain them. Objectively ‘regulated’ and ‘regular’ without being in any way the product of obedience to rules, they can be collectively orchestrated without being the product of the organizing action of a conductor.” (Bourdieu, 1990, p. 53)

The *habitus* is a product of history. It is the accumulation of past experiences which are actively present in the form of embodied habits which tend to be perpetuated by helping to re-activate the same habits in future behavior in a non-mechanical way, but as a consequence of an ongoing process of social practice. In this sense, the *habitus* is the main motor of its own perpetuation.

“ This system of dispositions – [...] – is the principle of the continuity and regularity which objectivism sees in social practices without being able to account for it; and

⁴See (Dreyfus & Rabinow, 1993; Taylor, 1993; Bourdieu, 1989, pp. 5 & 9)

also of the regulated transformations that cannot be explained either by the extrinsic, instantaneous determinism of mechanistic sociologism or by the purely internal but equally instantaneous determination of spontaneist subjectivism. Overriding the spurious opposition between the forces inscribed in an earlier state of the system, outside the body, and the internal forces arising instantaneously as motivations springing from free will, the internal dispositions – the internalization of externality – enable the external forces to exert themselves, but in accordance with the specific logic of the organisms in which they are incorporated, i.e. in a durable, systematic and non-mechanical way.” (Bourdieu, 1990, pp. 54 - 55)

It is evident that while the *habitus* is an organizational invariant of a historical process, its structural manifestations do not remain unchanged but constitute in themselves a condition for openness by constraining, but not specifying, practices in ways which may result in further changes in constraints.

Interestingly, the idea of the *habitus* articulates the issue of social coordination in a manner that can be directly linked with the biology of embodied autonomous systems presented in chapter 2 and which complements it. Coordination is always the result of previous or ongoing interaction. But a focus on single instances of interactive behavior makes it difficult to explain the patterns of coordinated activity exhibited by humans who belong to a same culture. Two persons who never met may show the same table manners even if there is no ongoing interaction between them. How can such a behavioral coherence be explained in the absence of some form of mutual perturbation leading to structural congruence? Bourdieu answers this question by showing that in cases like this structural congruence already pre-exists the encounter thanks to previous interaction with similar bodies of behavioral patterns, i.e. with a same *habitus*, (Bourdieu, 1990, p. 59).

The conceptual power of the *habitus* lies in its potential for biological grounding and for expansion. There are strong reasons to suspect that this concept will be able to unify in a single dynamical scheme, albeit not necessarily in a smooth manner, themes like Vygotsky’s theory of cognitive development as a socio-historical process, (Vygotsky, 1962, 1978), Wittgenstein’s “forms of life” (Wittgenstein, 1963; Margolis, 1987, pp. 436 - 440) and Dewey’s (1929a, pp. 279 - 281) and Merleau-Ponty’s (1962, pp. 142 - 147) accounts of habit formation, interaction between habits and their relation to the social. All these themes afford interpretations as parts of open historical processes.

A thorough justification of this claim would exceed the limits of this chapter, (but see Dreyfus & Rabinow, 1993; Taylor, 1993; Margolis, 1987, pp. 441 - 444). As an example, however, the following passage by Dewey is offered:

“Apart from communication, habit-forming wears grooves; behavior is confined to channels established by prior behavior. In so far the tendency is towards monotonous regularity. The very operation of learning sets a limit to itself, and makes subsequent learning more difficult. But this holds only of *a* habit, a habit in isolation, a non-communicating habit. Communication not only increases the number and variety of habits, but tends to link them subtly together, and eventually to subject habit-forming in a particular case to the habit of recognizing that new modes of association will exact a new use of it. Thus habit is formed in view of possible future changes and does not harden so readily. As soon as a child secretes from others the manifestation of a habit there is proof that he is practically aware that he forms *a* habit subject to the

requirements of others as to his further habit formations.” (Dewey, 1929a, pp. 280 - 281)

In this quotation many of the ideas that make up the notion of a historical process are clearly manifested together with a microscopic view of how the *habitus* operates. The dual nature of a constraint as limiting and as enabling is also clearly expressed. Habit-forming, like the formation of trails on grass, is a process that “sets limits to itself”, a historical process. At the same time, a different historical process obtains as a result of the embeddedness of embodied habit-forming in a social world. Habits are the active embodiment of personal history, but they are also the ground on which social history grows. Thus habits are formed in response to the behavior of others, and the behaviors they elicit contribute to influencing the formation of habits in others. Personal constraints are no longer isolated but they interact, conflict and combine with the constraints of others.

3.3 Explanatory value of spontaneous invariants

The above list of examples of historical processes, in no way comprehensive, is intended to demonstrate the diversity and importance of phenomena that can be interpreted as historical and, at the same time, that there is some gain in giving expression to the unifying themes implied by grouping these phenomena together. History can take many forms and there probably will not be an unique way to approach all its instances, but there are some recurring ideas which may help by providing starting points from which the understanding of a historical process can be approached.

An important notion of this sort is that of spontaneous invariants of the dynamics. Once an invariant is constituted, understanding the dynamical relations that permit it to remain invariant can provide a powerful frame of reference for addressing specific questions of what goes on in a complex historical process. It allows the researcher to understand why certain things can change while others remain the same.

Contained within a spontaneous invariant lies an explanation of its own perpetuation. Even if the properties of the process in need of explanation are not directly related with its maintenance, the invariant sets conditions to how these properties can change usually by limiting a high dimensional space of possibilities into a few ordered modes. This may occasionally be enough to explain the patterns observed. To give an example, in social mammals high-ranking individuals in a dominance hierarchy tend to locate themselves spatially in central positions with respect to the rest of the group. This can be explained as a consequence of the self-maintenance of the hierarchy and how its structure conditions the patterns of interactions and movement between members of the group, (Hemelrijk, 1998).

On occasions, it may suffice with a broad description of how an organizational invariant is constituted and what effects it has on the rest of the dynamics. In other cases, the process may have to be modelled more formally and the model attacked with mathematical tools like those used for studying self-organizing processes, or alternatively, by a combination of analytical techniques and computer simulations. In all cases, however, the actual recognition of an organizational invariant involves a phase of discovery of order in the relevant patterns. Sometimes this order may be quite evident but it may often require a careful investigation.

3.4 Summary

In this chapter the concept of historical process has been presented. In the sense intended, it does not suffice that a process be subject to contingencies to be able to say that it is a historical process. Nor is it sufficient to say that because a historical explanation is suitable, the process should be called historical. Rather, the term is restricted to processes in which contingencies occur and genetic explanations may often be useful, but at the same time, the processes present a special relation between their dynamics and changes in their own context or constraints. The operation of the systems involved in a historical process affects its constraints in ways that are often manifested as changes occurring at a slower time-scale or as occurring due to the recurrent interaction with other processes. Any process that is able to alter its own conditions of realization is a historical process.

The term “constraint” has been used in a more general sense than its classical meaning of a limiting relation in order to point to the contextual conditions by which general dynamical principles are embodied in an actual process. These constraints may change during the unfolding of the process and these changes remain distinguished from the dynamics of the process itself because the process is identified by an observer as maintaining certain organizational relationships unchanged.

In addition, due to the interplay between the dynamics of the process and its conditions of realization, new spontaneous organizational relations may appear which also remain unchanged. Understanding how these spontaneous invariants are constituted can provide a powerful explanatory tool, since they specify the conditions for their own perpetuation. As such, they are sources of order which may also affect other aspects of the process.

Concrete examples of such an explanatory power will be presented in the studies that will follow. But before that, since these studies comprise mainly a combination of mathematical models and computer simulations, it will be appropriate to address some issues regarding how these models are going to be used as tools of scientific inquiry and what knowledge should one expect to extract from them. These and other methodological questions will be the subject of the next chapter.

Chapter 4

Methodological issues in the use of formal models and computer simulations

The previous two chapters have described the theoretical landscape in which concrete questions about social coordination can be framed. These questions will be investigated by a series of mathematical models and computer simulations. But before proceeding to those specific models, some comments will be made on the way formal models in general and simulations in particular will be used in this thesis.

This insistence on methodological issues is not capricious, particularly when computer models are involved. There is a growing concern in the research communities in which these models are used about how simulations should be built, how should they be applied to scientific questions and how do they relate to other scientific activities like formulating hypotheses and proposing explanations, (Taylor & Jefferson, 1995; Boneabeau & Theraulaz, 1995; Miller, 1995; Di Paolo, 1996b; Jackson, 1996; Kitano, Hamahashi, Kitazawa, Takao, & Imai, 1997; Robson & Beshers, 1997; Noble, 1997; Bullock, 1998).

In building and using the computer models in this thesis, questions of this sort have been asked and answered. This is reflected in the models themselves. However, addressing these questions was not a trivial task; it involved a process of elaboration of methodological issues. In order to find answers to these questions it was necessary first to formulate them clearly and within an appropriate context.

There are three main reasons for presenting this process of elaboration in this chapter. First, because it is intended to raise awareness of the kind of methodological pitfalls that may often be encountered when using simulations in a scientific context. Second, because it is also intended as a contribution to these issues in itself. And finally, because it will help in highlighting how the proposed methodologies have been used in the rest of the thesis. At the same time, it will become clear that when these methodologies included novel factors these have been applied successfully to concrete research questions.

In the following section, an argument will be presented in favour of a comparative approach to the use of formal models in science. This argument is based on a pragmatic view of science as a culturally situated activity. Performing systematic comparison between formal models of different

granularity helps in bringing into light the role played by different factors. However, there is no need to suppose that some of these models are closer to reality than others since that is not the purpose of their use. Rather, they are used pragmatically and without ontological commitment to help bring the inter-relation between the different factors into a clearer understanding.

After that, section 4.2 will specifically address the issue of computer simulations and their use. An adequate framework is proposed for studying the question of how to integrate simulations into the rest of scientific activity. The positive contribution of this section will be based precisely on the issue of how to relate the organization between the patterns observed in a simulation to an existing explanatory framework that connects analogous patterns in natural phenomena. For this purpose, a brief background description will be given of the role played by tools in scientific practice with special attention paid to conceptual tools, such as thought experiments. Also a brief description will be offered of what constitutes a computer simulation followed by a criticism of previous attempts to integrate computational modelling into the rest of scientific activity. Finally, yet another attempt will be explored in terms of the different working phases that take place *after* the simulation has run successfully, the last one of which is associated with a role similar, but not identical, to thought experiments.

4.1 A pragmatic approach to formal modelling

At a time when scientific activity is becoming increasingly specialized and disciplines ramify into watertight compartments, it is unusual to find practising scientists worrying too much about how science is related to other human activities and what implications for their own work may follow from such an awareness. In fact, an insistence on methodological issues and the embeddedness of science within culture is likely to be met with contempt and accusations of destabilizing intentions.

Given the strong undermining of foundationalist and realist views of scientific activity articulated by many positions in today's philosophy of science¹, it is not surprising that scientists would prefer to avoid conversations where words like "relativism" or "constructivism" are likely to come up.

At first sight, it seems quite likely that, even if one regards non-realist views of science with a sympathetic eye, very little practical good can come from exposing scientists to these views. Kuhn himself expressed on many occasions that, in order to proceed somewhere, science within a paradigm cannot avoid being blind to the assumptions of the paradigm and sweep under the carpet those anomalies that the paradigm cannot explain². In contrast, Feyerabend suggests that the logical conclusion of any relativist position is that "anything goes" and therefore existing methods and traditions need not be followed since the chances of producing new knowledge are not necessarily increased by their conservation.

But is there a middle term between these extremes? Must a scientist be either a blind follower of a paradigm or a constant questioner of its assumptions? One may ask whether a scientist whose

¹Consider for instance the Kuhnian view of scientific progress as constituted by historically punctuated, incommensurable and paradigm-oriented epochs, (Kuhn, 1970), the more anarchist view of Feyerabend who proposes no intrinsic methodological difference between science and any other human activity, (Feyerabend, 1988), constructivist views, (von Glasersfeld, 1984), social constructionist views, (Gergen, 1994) and views based on the biology of cognitive systems, (Maturana, 1988b, 1991).

²"To reject a paradigm without simultaneously substituting another is to reject science itself", (Kuhn, 1970, p. 79). See also (Kuhn, 1977a)

work is informed of the cultural and historical situatedness of her activity could not perhaps orient this activity in ways that are practically significant for her interests and that of her community. Among the different criticisms of foundationalism in science, only the *pragmatic* position appears to allow for this possibility. Being aware of the limitations of a paradigm can be aimed at making a concrete practical difference in the way science is practiced. According to Rorty, pragmatism takes the Deweyan or Baconian view of what sort of activity science is:

“... Baconians will call a cultural achievement ‘science’ only if they can trace some technological advance, some increase in our ability to predict and control, back to that development. [...] This pragmatic view that science is whatever gives us this particular sort of power will be welcome if one has developed doubts about traditional philosophical inquiries into scientific method and into the relations of science to reality. For it lets us avoid conundrums like ‘what method is common to paleontology and particle physics?’ or ‘what relation to reality is shared by topology and entomology?’ while still explaining why we use the word ‘science’ to cover all four disciplines.” (Rorty, 1988)

Rorty admits that this Baconian view is no less fuzzy than the notions of *prediction* or *control*, but that nevertheless any other demarcation of science as following special methods or having a special relation to the world will not be able to avoid the criticisms of being foundationalist or transcendental. Perhaps a more encompassing view would replace *prediction* and *control* with the more general idea of *understanding*, which in its practical guise would include the former. This may turn out to be an unavoidable generalization particularly in view of complex processes such as those described in the previous chapter.

Thus, pragmatism recognizes science as a human cultural endeavour that strives to understand the relations with the phenomenal world and human actions within it by following particular languages, more or less strict rules, technological constraints and general assumptions that may originate in different domains other than the strictly scientific³. In short, science is embedded in the cultural practices of the time and many scientists choose either to ignore or to accept this fact without fully recognizing its consequences. It would seem that a critical approach to claims of ontological priority for scientific vocabulary (which are notably more often made by philosophers than by scientists themselves), while illuminating and welcome, cannot, by itself, help one practise science in a better way.

Instead of dwelling in the implications of recognizing the relativity of scientific methods and vocabularies, a pragmatic scientist chooses to be aware of them but also to carry on doing science since, as implied in the original criticism, no other human practice can escape being in the similar situation and, in most areas of scientific endeavour, the proven record of successful application of the methods of the discipline is enough incentive at least for an initial belief in the likelihood of future successes. There are two possible ways in which a scientist may choose to do this. Either she chooses to throw herself back into the network of concepts, norms and practices that define her discipline without further questioning them and pondering about the preconceptions upon which such a network is built only occasionally (as a Kuhnian scientist would do), or she chooses to reflexively make use of her awareness in order to redirect her own work and see how different perspectives and practices can add to its unfolding. The benefits of the first path are clear (science

³A modern exposition of the pragmatic view on science can be found in (Margolis, 1986, 1987).

proceeds as normal, problems are solved, etc.) while the second path, although it may sound as an interesting possibility in theory, is of dubious practical relevance, and it would probably not be encouraged by fellow scientists. Nevertheless, the claim is made that in this second path lies a way of doing science that is both legitimate, practicable and with a good potential for achieving both deeper understanding and practical success.

If a practising scientist is to make use of a relaxed attitude towards the lack of a foundational framework of legitimation for scientific activity and approach this activity with a Baconian or pragmatic stance, then she must be prepared to criticize foundationalism but at the same time avoiding a reflexive paralysis which can hinder actual practice. In short, if a vocabulary, an explanation, a tool or a method work adequately for the job that they are intended to do, then there is little point in constantly questioning their appropriateness based on their not being the ultimate truth about the subject, because following that logic, scientific practice should cease. This form of reflexive paralysis is a symptom of not letting go, of wanting to reach a complete understanding of the subject at hand. *Prima facie*, it looks as if it is based on critical assessment of the place of science in culture, but in fact it does not fully understand the implications of this embeddedness and ends up dismissing every scientific practice as pointless because of being without support. But continually dismissing practices in this way can only imply the existence of an unspoken belief that somewhere a practice will found, or a method, or a set of basic assumptions that will be, in some undefined way, the correct ones. Such a belief is incompatible with a true understanding of what the relativity of scientific knowledge implies.

So the pragmatic scientist must be prepared to recognize the historical and contingent nature of scientific activity, without this recognition necessarily being equated with a requirement for abandoning the actual practices as if there were something inherently wrong with them. Contingent practicality, rather than a privileged view on the world, becomes the guiding standard⁴.

But is there anything else that a pragmatic scientist could do? Obviously, there cannot be a *method* for finding more appropriate vocabularies, explanations or tools. This would go against the cultural situatedness of practical concern. There are, however, methodological attitudes that may help to expose the practical scope of different theoretical approaches to similar areas of inquiry. In fact, it is possible to think of a modest example which is relevant to the use of formal models in biology and has broadly guided the way in which those models have been used in this thesis.

Formal models may have different uses within a scientific framework. They may be used to emulate as accurately as possible a given natural phenomenon, with the aim of presenting a set of hypothesis about the generative mechanisms responsible for the observed evidence and being able to formulate extrapolations and predictions which may be tested empirically. In contrast, formal models may also be used exclusively as theoretical tools whose basic goal is not to emulate a natural phenomenon, but to discover the workings of theoretical frameworks, their consistency and suitability for explaining a range of phenomena. As such, the relation of this kind of formal modelling to empirical evidence is indirect and its use is comparable to the use of thought experiments (see section 4.2.3).

⁴Interestingly, many of the most creative scientists, without necessarily calling themselves relativists, would not be completely uncomfortable in a position which recognizes theoretical terms as pragmatic constructs. Talking about the reality of electrons, Richard Feynman said that “the electron is a *theory* that we use; it is so useful in understanding the way nature works that we can almost call it real”, (Feynman, 1986, p. 70).

Despite insisting on their character as tools of research, there is a tendency to equate the formality of some of these models with a logical demonstration of the truthfulness of their conclusions, thus forgetting that, even if the theory they are supposed to be instantiating is widely accepted as the most adequate, the actual construction and expression of these models rely on a set of explicit and implicit assumptions which allow their rendering in a mathematical or logical form. This process of abstraction is a necessary precondition for “securing the ability to deal with affairs which are complex”, (Dewey, 1929b, p. 173). But at the same time, abstraction breeds familiarity and the illusion of secured knowledge (*ibid.* p. 174). How reasonable the corresponding assumptions are is a contingent, rather than an absolute, fact which may depend on issues as varied as the availability of appropriate techniques or questions of elegance and clarity of the resulting model.

It is necessary, then, to emphasize the obvious conclusion that never should a formal model be taken to say more than what it says, both explicitly and implicitly, and that whatever it says is mainly about an existing theoretical structure and only indirectly about observable phenomena. The question then arises of whether the use of these models can be justified at all given their vulnerability to being contradicted once a seemingly innocuous assumption is exposed to be playing an unexpectedly important role. The answer is affirmative if the use of these models is accompanied by an attitude aimed at testing the importance of such assumptions.

The basic idea consists in using simultaneously models of different degrees of granularity aimed at addressing a set of similar questions. This is done in the hope that, by means of systematic comparisons, the relevance of the different assumptions will be highlighted. There seems to be nothing new about this. On the contrary, it is the usual way in which this type of scientific activity is carried out. However there is a change of concern from the traditional practice of formal modelling. When two models are compared this is often done with the aim of deciding which one approaches the real situation better. The idea being proposed here is that that should not be considered the key concern in the practice of formal modelling. Rather, comparisons should be used to bring about an awareness of the role played by different factors and thus help the researcher to better understand the consequences of a theoretical framework and even modify it. In this context, even a model which bears no resemblance to the real situation may be helpful.

This attitude could be specially fruitful in view of the introduction of novel modelling techniques. When radically different modelling approaches are compared, although such a comparison may not be easy to perform, the chances of uncovering the relevance of implicit assumption are greater. In contrast, the advantage of extending a current model by a parsimonious relaxation of only a few assumptions is that the task of comparing results becomes a rather simple one. The main disadvantage, however, is that most of those implicit assumptions will be tested in this form only with a very low probability and consequently they will be inherited by the new model. This is mainly so because many of these assumptions are made in order to be able to formulate the model in the first place, i.e. for technical, rather than conceptual reasons, and therefore, the formulation of a parsimonious extended version of the model will find it very hard to revise them.

To give an example, it will almost certainly be possible to extend a mathematical model based on differential equations in order to include different initial and boundary conditions or different coupling between the variables, but it will be very hard to extend the same model in order to treat the continuous variables as discrete entities or to allow those entities to form new variables whose

dynamics is not initially specified but will depend contextually on state of the system, (Fontana et al., 1994; Crutchfield, 1994a, 1994b). A model which is able to address these issues will have to be built not as an extension to the set of differential equations but from scratch, in a way that may differ radically from the approach used in the initial model. It is in the comparison between two radically different models where one would expect the importance of implicit assumptions to be highlighted.

In short, by using different formal approaches comparatively an attempt is made to achieve a series of goals:

1. To highlight which conclusions are dependent on the type of model chosen and which are not.
2. To pinpoint what are the relevant features and assumptions that make certain results different from one model to the other.
3. To question the meaning of these assumptions and whether they were explicitly identified as relevant from the start in those models that make them.
4. To assess the plausibility of these assumptions in natural cases and thus put some weight on the conclusions of the corresponding model.
5. To question the reasons why these assumptions were made in the first place, whether for technical reasons or because they were assumed to be reasonable or not to present much of a difference in the expected results. In the second case in particular, a subsequent reordering of the relative importance of the theoretical entities involved in the original assumptions can be proposed along with ways to test the pragmatic value of the resulting modified conceptual structure.

This methodological attitude has been successfully applied in this thesis, in particular, within the context of the issues examined in chapters 5, 6 and 7. Chapter 5 offers a series of different mathematical models based on extensions to a basic game-theoretic approach to an evolutionary question. These comparisons are useful and revealing but it is only when a radical departure is made in the form of an individual-based computer simulation aimed at addressing the same question (chapter 6), that many non-obvious assumptions of the mathematical approach are highlighted resulting in different conclusions both about the actual expected results and about the role played by those implicit assumptions in the broader theoretical framework of evolutionary theory. Similarly, the model presented in chapter 7 is in a way an extension of the individual-based model of chapter 6 which facilitates the understanding of how a factor that was not contemplated in the original version (social influences on behavioral development) can also introduce fundamental differences in the expected results.

It is clear then, how in the quite specific domain of formal modelling a pragmatic approach does not question every single questionable choice in an anarchistic manner but neither does it accept every choice without question. The assumptions that will end up being questioned will be those whose significance is highlighted by the comparative use of formal models among themselves.

4.2 A scientific use for computer simulations

Given the relative flexibility of computer simulations, their use seems particularly adequate in the light of the comparative methodology described above. But their use also introduces new questions. How can a computer simulation be used in a scientific context? Is it possible to gain new understanding by building and observing a simulation? Isn't a simulation just a computer program from which one cannot gain more information than what is already in it from the beginning?

These are legitimate questions that anyone who intends to use a computer simulation as a research tool should ask and hopefully answer. Before doing that, however, it will be helpful to try to narrow the scope of what is meant by a computer simulation in the present context and then proceed to the description of a background in which these questions can be framed more clearly.

First, it should be made clear that the concern of this chapter (as that of the whole thesis) is with the use of simulations in a scientific context. More precisely, simulations will be used here only as tools of theoretical inquiry. These restrictions already serve for narrowing the scope of what is usually understood by the term computer simulations to a rather small subset.

Simulations can be used for engineering or training purposes or to control or predict the behavior of particular systems such as aircrafts, nuclear power plants, etc. This sort of simulations require a high degree of detail and granularity, and the implementational choices have to be made so that as good a match as possible with the actual system is achieved. Usually for these simulations, the greater the sophistication, the better. This is not the business of the simulations one is concerned with here.

There are scientific uses of simulations which are not included within the subset of interest of this thesis. These are simulations which attempt to replicate as faithfully as possible specific natural phenomena in order to be able to make predictions of future observations and perhaps test assumptions by direct empirical comparison with the outside world. For these simulations quantitative results are important and therefore lots of implementational details and parameters are restricted by the requirement to find a good match with actual measurements. While some of the issues discussed below will apply to this type of simulation, many of the methodological problems are solved by this comparative use which establishes a direct empirical link with the natural world.

In contrast, the computer simulations of interest need not be directly applied to empirical data, or if they are, only a qualitative match is what is looked for. These simulations are rather used to explore ideas, to query the limitations and consequences of hypotheses and test the soundness of theories. Therefore, many of the requirements which are essential for other types of simulations, can be ignored. In fact, an abundance of details and a high degree of sophistication can be a hindrance for this type of simulations since they may make ultimate understanding of the ideas under test something very difficult to achieve. This use of computer simulations has become increasingly frequent in science (for instance, in physics and ecology) and particularly so in some relatively novel research disciplines such as Artificial Life and the study of adaptive behavior. Disciplines like the physics of systems far from equilibrium, biology, economy, linguistics, cognitive science and the social sciences have a lot to gain from this type of computer-based approach since in all these disciplines the limitations of other forms of modelling are constantly being exposed.

In a sense, these simulations⁵ seem to be much less restricted than other classes. The researcher is, in principle, free to explore many different possibilities and include as many mechanisms as desired. Such freedom is partly responsible for the potential of simulations but also partly responsible for the problems that arise in their use. At least, when the intention is to test the results of a simulation against empirical data, the construction and usage of a simulation are constrained by this aim. What does constrain the use of a simulation which is only intended to explore theoretical ideas? Who says which choice of implementational details is an adequate one?

To add to these problems one could ask what would be a legitimate use of these simulations. How does work using simulations relate to the rest of scientific activities such as the formulation and testing of hypotheses, explanations, etc.? Some people hold the implicit view that a simulation that meets certain criteria, so that it can be said that a specific phenomenon is replicated by it, is in some way an *explanation* of that phenomenon, (Boneabeau & Theraulaz, 1995). But is it always (or ever) so? Is it inconceivable that after successfully replicating certain patterns that approximate a phenomenon of interest this phenomenon should remain as obscure as before? As will be seen below a simulation does nothing by itself. It must be observed, and this is a far from trivial operation. Not only does it imply choosing observables, but also finding the relationships between them: “You can look at the computer screen and be amused and amazed at what you see, but what precisely is connected with what? These “whats” need to be comprehensible, and not just some nice dancing patterns”, (Jackson, 1996, p. 20).

There is also a tendency to use simulations as proofs of feasibility in a theoretical vacuum that come short of being proofs of concept. A proof of concept should have a concept to prove beforehand and this is not always the case for a considerable proportion of the works in the relatively novel disciplines mentioned above. Comparisons and contrasts between artificial and real patterns are not of much scientific interest *in themselves*. According to Bullock (1997, p. 457): “Unless such parallels were previously hypothesised to exist, they are either merely accidental (and thus not interesting), or merely purposed (and thus not interesting)”. This should not be taken as meaning that a simulation cannot be surprising and that after running it, the researcher may not answer questions different from the originally intended ones. An accidental parallel is always welcome and maybe Bullock is being a bit too harsh. Nevertheless, his main point should be taken seriously and a reformulation of the project is called for if the accidental result is to be presented as a piece of scientific work. Purely exploratory experimentation, in contrast with “theory-driven” research, may be an important part of scientific practice (especially when an adequate theoretical framework is lacking, see Steinle, 1996), but in the case of computer models it cannot, by itself, give a piece of work sufficient scientific legitimization⁶.

It could be said that there is no fully accepted methodology at present to deal with these issues despite various clever attempts to address them in more or less general terms, see (Taylor & Jefferson, 1995; Boneabeau & Theraulaz, 1995; Kitano et al., 1997; Noble, 1997; Miller,

⁵Henceforth the term “simulation” will be understood only in the restricted sense of interest unless otherwise specified.

⁶A possible exception to this would be the rather restricted case in which a class of formal systems is in itself an object of study which, for some reason, it makes sense to approach in a more or less empirical manner. For instance, the effects of synchronous updating may be empirically studied in a class of cellular automata (Ingerson & Buvel, 1984; Bersini & Detours, 1994), the class of coupled-map lattices (Abramson & Zanette, 1998a; Rolf, Bohr, & Jensen, 1998) or in the class of random Boolean networks (Harvey & Bossomaier, 1997).

1995). Part of the reason for this lies in the lack of a clear formulation of the problem. The following paragraphs are an attempt to describe a framework in which this formulation can be clearly expressed.

4.2.1 Science and instruments

It is suggested here that the best framework for addressing the issues raised in the previous section is a pragmatic one and that the methodological problems can be expressed in terms of how well a new tool may be integrated into existing scientific practices. Such a pragmatic formulation cannot take the form “What method of use for computer simulations will be the “right” one for doing science?”, but instead: “*What method will integrate computer simulations into an existing body of practices called “science” without changing it so radically that one would not be willing to keep calling it “science” while, at the same time, allowing for the potentialities of the new tool not to be so suppressed that the new tool becomes a trivialized case of older ones?*”. The answer to this question will be necessarily contingent on present views about what constitutes scientific practice.

For this purpose, a brief introduction is offered to the subject of how tools contribute to shape scientific activity. The word “tool” is used in a general sense and includes things such as measuring devices, formal techniques, sample preparation methods, etc. From a more encompassing point of view, communicational strategies, conceptual structures, argumentative techniques, etc. can also be embraced under this concept. However, those instruments that, because of their complexity, require the coordinated action of many individuals (for instance, particle accelerators) will be kept out of the discussion for reasons that will become clear later.

4.2.2 Tools in scientific activity

The prevalence of mediation is a characteristic factor in all human activities, science being no exception. As much as scientists and philosophers well into the beginnings of the twentieth century have tried to separate those aspects of science that relate to a realm of pure knowledge (*theoria*) from the contingent set of instrumentalities used to obtain such knowledge (*praxis*), such separation has been impossible to achieve and, more importantly, it has shown little pragmatic value. John Dewey was one of the most fervent advocates of the point of view that knowledge and action bear not just a historical relation, but are in fact irremediably intertwined. In *The Quest for Certainty* (Dewey, 1929b) he shows that although scientific knowledge differs from “mere” traditions or beliefs in the need for a methodological framework, it is no different from them in its concern with the practical, both as a means as well as an end. Much of the following is rooted both in his views on scientific practice and in Heideggerian considerations of what constitutes skillful activity (Heidegger, 1962; Dreyfus, 1991; Ihde, 1979). Human activities are characterized by specific frameworks of normativity brought about by the use of tools and from which tools receive their character. By tools it is meant any form of means used towards some end. When one deals with tools one does not deal with their immediate properties; these are subordinated to their use as agencies towards some ulterior result. One does not make a distinction of mediational means out of the background of one’s experience nor, in fact, does one distinguish them from one’s own body as long as the engagement in some sort of skillful activity continues, see for instance (Heidegger, 1962, pp. H69 - H62, 98 - 102) and (Merleau-Ponty, 1962, p. 143). This is why, for certain uses,

tools can be considered transparent. If a scientist is examining a blood sample under a microscope, her concern is with the sample, not with the microscope nor with her own eye. But for this to be the case it is necessary for tools to be integrated into a framework of accepted normativity that allows the user to forget about them, in this case the correct preparation of samples, calibration, conditions of use and the observation of other norms⁷. The role played by tools tends to be ignored precisely because it is obscured in the light of the concluding event towards which the activity is aimed. One only becomes aware of them when they “stand in our way” (for instance, when one is inexperienced in their use, when they cease to work properly or when they are found to be inappropriate for the job).

Tools may typically play two basic roles in science. The first role is corporeal enhancement. This is the most obvious role and it is readily associated with measuring devices and observation techniques, but it may also be associated with formal methods, data processing, etc. The second, less obvious role is what allows science to be a social activity. Tools “objectify”, in the sense that they provide meaningful inter-subjective vocabulary to talk about observations of natural phenomena and actions within the world. By following the same procedure, two scientists can be certain that they are giving compatible meanings to their terms because the procedure itself has taken care of selecting a finite amount of relevant observables out of the virtually infinite number of possible variables that may be extracted from an observed process. The procedure has also canalized the results of the observations into a finite number of (usually quantifiable) outcomes. In this way tools modulate the meaning of the terms used by scientists and, therefore, they have communicational value. This communicational capability can reach extreme refinement as in the obvious example of mathematical models. These models constitute both a tool for exploring, deducing and arriving at conclusions and, at the same time, the communicational channel through which these activities are made public. This is not generally the case with computer simulations: they require additional effort in order to render them communicable, therefore their results may not be initially “objective” in the same pragmatic sense that a mathematical proof is.

Although, for certain uses, tools can be considered transparent, not all the relations to technology are necessarily of this sort. This caveat applies especially to complex uses of technology in science. Consider the different modes of engaging in skillful activity which are revealed in the transition from the use of a magnifying glass to a microscope to a telescope to a spatial probe, (Ihde, 1979, pp. 29 - 33). Consider in each case whether it is possible for the tool to become embodied in a loop of action and perception. Clearly the magnifying glass presents little problem. One easily becomes accustomed to its use and is able to find the appropriate spatial relations to obtain the desired perceptual effect. The optic microscope is a bit more complex, but there is a sense in which, after some training, one can position the sample and adjust the focus without concerning oneself too much with these tasks. An electron microscope is a more complex. In the old ones, the result of the observation was a photographic plate so that the experience of a continuous flow of action and perception became interrupted. In modern ones, however, the scientist is able to manipulate the image in real time and the “feel” is very similar to that of an optic microscope. Similar comments could be made for the uses of small optical telescopes, bigger ones and, say,

⁷Tools can also be associated with more theoretical activities, for instance when a physicist linearizes a dynamical system around a fixed point to determine its stability or, in a much more obvious way, when we use equipment for writing, see (Heidegger, 1962, p. H358, 409).

radio-telescopes. And in the extreme case of a spatial probe (or a particle accelerator or experimental nuclear reactor) the tool has become so thematized that its use requires the collaborative and specialized efforts of many scientists and technicians, and a great deal of those efforts are concerned with the actual workings of the tool rather than directly with the ultimate experiment. All these specialized tasks, however, may indeed fall individually into the format of embodied skillful activity.

This last example suggests that, in general, the schema “means/ends” is a rather simplified one, as one is constantly engaged in some sort of mediated activity and there is no clear-cut distinction (unless it is arbitrarily drawn) between an end and a state or condition for further mediated activity. Thus, an examination of blood samples may be done *in order to* prove the presence of a contaminating agent within the context of an autopsy. The set of constraints that allow one to forget about the tool being used, is not just “local” in the sense that it is only concerned with a specific instrument or technique, but it is affected by a web of norms in which the immediate activity is embedded. This can be said both of simple tools which directly become embodied in the scientist (a magnifying glass or even a pen) and of much more complex tools like particle accelerators. In the latter case, the framework of normativity does not only surround the use of the tool but it also acquires a complex and sometimes rigid “internal” structure in the form of rules and procedures which must be strictly followed. So, in general, it can be said that all legitimate uses of tools follow norms and that these norms are bound to become more explicit as soon as the tool becomes more complex and its use requires greater care.

It is a particular framework of normativity that defines whether a practice can be considered scientific; a framework which is not disassociated with the subject matter of the discipline. *And it is the degree of integration into this framework that provides a criterion for deciding whether such activity can be called scientific (within the corresponding discipline).* This is why finding bacteria in a meteorite sample which has been carefully treated to avoid contamination is considered a legitimate scientific result while the same finding in a sample which was kept in the kitchen or was used as a toy by a dog, is not. Scientists have been using this a criterion in an implicit form for a long time.

Although a scientific discipline can be characterized as a set of activities that comply with certain norms this does not mean that any such framework of normativity can be called science. This important issue is deliberately avoided here since it would take too long a detour to confront it. So it must be accepted, for practical reasons, that there may be more than one view of what core features a scientific framework of normativity should have, a view that will depend on personal beliefs about science. Still, there is a stronger reason for avoiding this issue and this is basically that it is of no fundamental relevance to the argument presented here: all that is needed is the acceptance of the above criterion.

Given such a criterion for evaluating whether an activity may be called scientific, a particular problem now appears when this activity involves the use of new tools or of existing ones in novel ways. The reason for this is that the successful integration of new tools into an existing framework of normativity may imply, at least in principle, an alteration to such a framework, and therefore a modification to what is regarded as legitimate scientific practice. In many cases this integration is successful because the new tool is accompanied by a body of theory that validates its use (another

example of the close relation between *theoria* and *praxis*). Thus, for instance, game theory was successfully integrated in theoretical studies of evolutionary biology not because there are justifiable ways of saying that organisms take rational decisions, but because there is an *implicit theory*, already accepted before the introduction of the new tool, that states that optimization will result from a process of natural selection (Maynard-Smith, 1982). Such a successful introduction of a new tool can lead, as this example has shown, to the opening of a new ground of research that was previously unavailable partly because the prior web of constraints made it difficult to explore. It may also lead to a careful conceptual re-examination. But, of course, its use will at the same time introduce new constraints, the most trivial of which will be the implicit acceptance of the theory that validated the introduction of the tool. So, in the case of game theory applied to the study of evolution one will not expect to use it to question the assumptions that gave it validation.

It may also happen that if the new activity is validated with a theory of use which was arrived at utilizing other tools or techniques with their own theories of use, then the new activity may partly “inherit” those validating theories in its own validation. The importance of this possibility will become clear later in the discussion of previous attempts to integrate computer simulations in biology.

Before proceeding, the following clarifications should be made. First, not all the normativity accompanying the use of a tool can be spelled out in the form of a concrete theory, although probably part of it can. In fact, the validation of a new tool by means of theory of use is only one possibility among many. Other tools may be validated in different ways or directly accepted as legitimate once enough people have been convinced they are so. Second, the separation of the theory of use of a tool and the area of inquiry in which the tool is used is an *ideal* objective, not necessarily attained in all cases and certainly not always when tools are introduced for the first time: “Before he could construct his equipment and make measurements with it, Coulomb had to employ electrical theory to determine how this equipment should be built. The consequence of his measurements was a refinement in that theory”, (Kuhn, 1970, pp. 33-34). It should not be surprising then that novel tools will be used before enough is known about them to say how their use integrates into existing scientific practice and that that knowledge may indeed come from the results obtained by using these tools.

4.2.3 Thought experiments

As emphasized above, a tool is not necessarily a physical “thing”. It is possible to talk about procedures, techniques, etc., as tools because they share common features with more obvious examples of tools. Particular attention will be paid here to what can loosely be called *conceptual tools*. These are tools for dealing with conceptual structures, though they may be applied for “practical” purposes. Thus, mathematics can be used for formulating and working with complex conceptual relationships and logic may be used for studying the internal consistency of a theory.

One particular conceptual tool of interest in the present context are *thought experiments*. Work using computer simulations has often been compared with “prosthetic” thought experimental excursions into questions too complex to deal with using naked cognitive capabilities. Such excursions could indeed catalyze an avalanche of material for philosophical activity (Dennett, 1995). Although formulated with the aim of addressing mainly philosophical issues, this is an interesting

point of view which can find similar applicability in a scientific domain and therefore deserves further exploration.

Thought experiments are difficult to define and classify, although their role in the history of science is undeniable. There are many views on the nature of thought experiments; here follows a brief description of an account by Kuhn which seems the most appropriate in the present context.

Kuhn (1977b, p. 241) poses the following questions about thought experiments. First, to what conditions of verisimilitude must a thought experiment be subject? Second, given that a successful thought experiment involves the use of prior information which is not itself being questioned how, then, “can a thought experiment lead to new knowledge”? Finally, “what sort of knowledge can be so produced”? Kuhn says that these questions have a set of rather straightforward answers which are important but “not quite right”. These answers suggest that all the understanding that can be gained from thought experiments will be understanding about the researcher’s *conceptual apparatus*; for example, in eliminating previous confusions or showing inconsistencies within a theory. If this is the use that one makes of thought experiments the only requirement of verisimilitude that must be fulfilled is that “the imagined situation must be one to which the scientist can apply his concepts in the way he has normally employed them before”, (Kuhn, 1977b, p. 242).

Kuhn describes a well known thought experiment in which Galileo shows the Aristotelian concept of speed (something similar to the present day idea of average speed) to be paradoxical. An immediate interpretation of this thought experiment can be given along the lines mentioned above. However, by carefully analysing how the Aristotelian concept of speed was used by Aristotle and his followers, Kuhn finds that it “displayed no *intrinsic* confusion” (Kuhn, 1977b, p. 261). Contradictions arise when the scientist tries to apply this concept to previously unassimilated experience, as it happened to the “corrected” Galilean concept of speed itself when it was later confronted with situations where its application showed inappropriate (such as the additivity of velocities of electromagnetic waves). In this way, Kuhn argues, the thought experiment is indirectly saying something about nature and has a historical role similar to empirical observation. But how is this possible, when it was assumed that no new empirical information was “fed” into the thought experiment? Kuhn answers: “If the two can have such similar roles, that must be because, on occasions, thought experiments give the scientist access to information which is simultaneously at hand and yet somehow inaccessible to him”, (Kuhn, 1977b, p. 261). Scientists decide to pay attention to “problems defined by the conceptual and instrumental techniques already at hand” (Kuhn, 1977b, p. 262). Therefore, some facts, although known, are pushed to the periphery of scientific investigation, either because they are thought not to be relevant, or because their study would demand unavailable techniques. A thought experiment will, on occasions, bring the relevance of these facts into focus, and therefore catalyse a reconceptualization which may involve an undramatic re-organization of relationships between existing concepts, or indeed a scientific revolution.

4.2.4 Simulations as tools

In order to assess the possibility of relating computer simulations to thought experiments it is convenient to explore the nature of simulations in more detail.

Rasmussen and Barrett (1995) have given a general description of what constitutes a simu-

lation. A simulation is an iterated mapping performed on a set of objects $\{S_i\}$, each of them characterized by an internal state, a local time coordinate, a dynamical mapping that specifies state transitions, and a set of relations to other objects. This global mapping is performed by an update functional U applied over the entire set. For simplicity Rasmussen and Barrett assume the following form:

$$\{S_i(t+1)\} = U(\{S_i(t)\}), \quad i = 1 \dots n.$$

However, this expression implies that all objects are updated synchronously (notice the common time parameter) and the set of objects remains invariant (notice the common index). A more general form that allows for asynchronous and constructible objects would be:

$$\{S_j\}_{(t+1)} = U(\{S_i\}_t), \quad i = 1 \dots n, \quad j = 1 \dots m,$$

where the update of the whole set does not necessarily imply that all the objects are actualized on the same “time step” and the number of objects need not remain constant (m is not necessarily equal to n , although a direct one-to-one mapping exists between the indexes i and j with the exception of objects that disappear and objects that are created).

If the definition stopped here, it would seem that there is no difference between simulations and other kinds of formal systems, and it would be very hard to explain how a computer simulation can give rise to emergent patterns. Rasmussen and Barrett very rightly point to the relevance of the inspection methods which are applied to the data generated by a simulation, a point also made by Baas (1994) and Crutchfield (1994a). An observation function Obs may permit the distinction of higher level properties of the basic objects, (for instance, in a spatial model of interacting objects, an observer may conclude that they aggregate into more or less well-defined clusters or flocking patterns), and these properties may indeed point to higher level structures with their own range of phenomena, (e.g. the flocks themselves, which may show complex patterns of movement and interaction with other flocks). These new structures could, in turn, be subject to further observation (Obs^2) leading recursively to a potential hierarchy of levels.

While it is clear that the iterated mapping is a formal, analytical process, the observation functions need not be so (for instance, it may be computationally undecidable whether a particular object at a given time belongs to one flock or another if both flocks are approaching and beginning to fuse⁸). It is the observer who discovers the relations which are interesting and is able to express them in a language that may be different from the terms used in the actual construction of the simulation. Thus, the observer is entitled to make statements like “flocks travel around obstacles in such and such a way” and use these statements to derive conclusions about the subject of interest. Such statements, however, cannot be validated by a careful inspection of how the model is built and instantiated since nothing resembling a “flock” will be found in it. It is in this sense that, in the current context, statements like the above are said to refer to emergent properties of the model.

⁸To give a more formal example also cited in (Baas, 1994) and (Rasmussen & Barrett, 1995), consider the iterated map $z(t+1) = z^2(t) + c$ with z and c complex numbers. It can be shown that the location in the complex plane of the boundary of the set of initial points z_0 which do not result in a divergence of the map for large values of t , the so-called *Julia set*, is non-computable for most values of c when computation is defined over the real numbers. It cannot be formally decided whether a given point of the plane is a member of the Julia set, although it is possible to make a decision by direct observation. Rasmussen and Barrett (1995) conclude that a system may be simulatable and present non-computable emergent properties.

In view of the themes discussed in chapter 3, a case of interest occurs when objects at a certain level of description incorporate ways of responding to constraints imposed by the higher level structures/properties they give rise to as, for instance, motorists who try to avoid heavy traffic routes following a radio report. Such a case has sometimes been termed *downward causation* (Rasmussen & Barrett, 1995) or *intrinsic emergence* (Crutchfield, 1994a).

This issue runs parallel with problems that suggests that there is room for other interpretations of the term “emergence”. For instance, the problem of how could the “observational capabilities” of basic constituted objects be themselves constructed during the running of a simulation (Cariani, 1991; Pattee, 1996). These issues are fascinating but beyond the scope of this chapter and would deserve a separate treatment.

It is the possibility offered by computer simulations for allowing the exploration of emergent properties of a model that makes them simultaneously novel, interesting to use and difficult to integrate into traditional scientific frameworks because of the lack of other tools of theoretical inquiry sharing this property. Even in physics, where there is a strong tradition of computer modelling in the form of cellular automata, lattice gases, etc. which originated in the pioneering numerical experiments by Fermi, Pasta and Ulam (1965), emergent structures and properties are only limitedly exploited. Models that allow for object construction, creation of structures, intrinsic emergence, etc., are not so common partly because “there are no physical principles that define and dictate how to measure natural structure”, (Crutchfield, 1994a, p. 18). It is not surprising, in contrast, that “... physics does have the tools for detecting and measuring complete order [...] and ideal randomness”, (Crutchfield, 1994a, p. 18).

If one wants to move into the domain of biology, the questions of innovation or interactions between dynamical structures at different levels become essential. Take for example the problems of how new biological structures may appear (Fontana et al., 1994; Fontana & Buss, 1996), or the interaction between emergent dynamical structures and more common processes involving differential reproduction (Boerlijst & Hogeweg, 1991b, and chapters 6 and 7). Or studies where systems must obey non-holonomic constraints (especially non-limiting ones), as is the general condition for adaptation in living organisms to be conserved⁹.

In cases like these, there are strong reasons to suspect that traditional mathematical tools are insufficient as well as computer simulation which are directly based on them. This is an important fact to be taken into consideration at the time of examining other proposals for the use of computer simulations, particularly in biology.

4.2.5 Integrating simulations in biology

A fairly straightforward way of integrating computer simulations within biological research has been recently suggested by Kitano *et al.* (1997). They propose a cycle of research very much

⁹A non-holonomic constraint is a relation between the coordinates of a system and its environment which does not follow an integrable law. A non-limiting constraint is a special case in which the relation is not even given between higher order derivatives of the coordinates. For example, a disc rolling in a plane surface is constrained non-holonomically by a relation between angular and translational velocity but not between angular coordinate and position. A particle moving “freely” within a box is constrained by its walls but it is not limited in the movement it can perform between collisions. This a non-limiting case. See for instance (Goldstein, 1980). In order to evaluate the importance that these constraints may have for adaptive behavior consider the all-or-nothing nature of adaptation as described in section 2.4.

like the one used in other sciences, with the addition of a *virtual* phase in order to generate predictions and explore the soundness of a hypothesis. In this cycle, hypotheses are generated from the observation of natural phenomena, and they are built into simulations which may point to new observables which, in turn, can be used to design new experiments and discard or support the initial hypotheses. Indeed, there is not much difference between this scheme and a traditional use of formal models and, at first glance, there is no reason why there should be. The only difference would be that in this case the modelling phase implies the use of a simulation. The authors argue that simulations, instead of purely mathematical models, will be needed to test theoretical soundness given the inherent complexity of biological phenomena in which systems are expected to be constituted by a large number of interacting elements which are difficult to describe in global terms.

There are two problems with this approach. First, it is assumed that there is a methodologically clear way of translating hypotheses into a simulation. Much of the methodological problems in the use of simulation arise because of confusions regarding what should guide implementational choices and how to interpret results. These issues are not dispelled in this approach but remain problematic. A second and more specific objection is that this scheme relies heavily on the availability of empirical data in order to derive hypotheses and to contrast the result of the simulation. This will surely be the case in a number of cases, like well studied problems in morphogenesis for instance. Inevitably, though, the chances of success will be reduced when the empirical evidence is scarcer as happens to be the case in many issues of interest, such as the origin of life, the evolution of communication, etc. It's difficult to see how any methodological "cycle" could ever be completed in such cases.

A different approach has been proposed by Miller (1995). He suggests that a good starting point for a simulation that aims at addressing a scientific issue in theoretical biology can be found in existing formal models. The idea is to replicate the results of a formal model in a minimal simulation. Once this step is achieved, some of the assumptions of the formal model can be gradually relaxed and the consequences explored. Here, there is an implicit reliance on the assumption that a computer simulation will constitute a more powerful medium with which to explore the subject than the original formal model, which is usually the case, and that an observer will be able to understand what goes on in the simulation as compared to the expectations of the formal model, which is *not* necessarily the case, but since assumptions are relaxed systematically, it is possible to formulate very strong guesses. This proposal is not incompatible with the approach advanced by Kitano *et al.*, they are in fact complementary. Whereas Kitano *et al.* try to integrate simulations into an existing background of other scientific activities, Miller tries to relate them to existing tools with an already defined role within that background.

Miller's proposal can be criticized for being perhaps too conservative for a general methodological solution to the integration problem. No doubt, this method of integrating computer simulations with an existing body of theoretical research can prove to be very fruitful in specific cases¹⁰. But, from the viewpoint of the instrumental analysis presented above, to rely on existing formal models, themselves tools with a valid range of applicability, and expand such a range by relaxing some assumptions of these models does not preclude the (partial) systematic inheritance

¹⁰See for instance (Bullock, 1997).

of conceptual constraints and theory of use of the original tools.

In this way, to return to an example given before, if a game-theoretic model is applied to a problem in evolution one is assuming the implicit theory of use of this tool as valid for the problem in question, namely that evolution unfolds by a mechanism of natural selection of individual phenotypic traits directly mapped from their inheritable genotypes which leads to a process of (constrained) optimization. One does not question the implicit theory of use of a tool if it is being used *qua* tool because if one does, the very use of the tool loses its scientific validation. If it happens that the potential domain of applicability of a new tool has an important intersection with the implicit theory of use of another tool, then it seems that one would be unnecessarily restricting the use of the new tool if one requires, as part of a general methodology, for it to partially adopt the theory of use of the latter. Such inheritance is what should be expected if the methodology prescribes specifically that simulations should emulate and extend existing formal models. The simulation and the initial formal model will be genealogically related which, in a way, will limit the potential of the simulations perhaps unnecessarily. Following the above example, since a large number of formal models in theoretical biology make use of game-theoretic techniques, it would seem unwise to directly derive simulations from them in order to study questions like “Is natural selection the only factor in evolution?”, “How relevant is development for evolution?” or “Do the dynamics of emergent group structures influence evolution in unexpected ways?”, etc.

It seems that building simulations based on existing models in theoretical biology will work successfully in those particular cases in which there is little overlap between the object of investigation of the simulation and the basic framework of validation of those models. In the general case, this approach cannot be used as a methodological recipe for integration. Moreover, it was already seen that many of the potential areas of application of computer simulations are those that escape traditional methods in present day biology. These methods must assume simplifying conditions about phenomena in these areas (such as “genotypes map directly into phenotypes”). Fontana and his colleagues have repeatedly pointed out that biology lacks good theories of object construction and biological organization precisely because of the lack of appropriate techniques to approach these issues, (Fontana et al., 1994; Fontana & Buss, 1996). Such techniques are perhaps within the potential capability of computer simulations. Consequently, it can be said that an important degree of intersection between areas of applicability of this new tool and the theory of use of the old ones is not at all unthinkable for many cases of interest.

This point is reinforced if one intends to use formal models comparatively following the pragmatic approach described in section 4.1. Although a similar philosophy can be found in Miller’s suggestion of comparing simulations with the models on which they are based by relaxing assumptions, it is clear that the advantages gained in the ease of comparison could be overshadowed by the amount of assumptions which will be hard to identify and relax because of being implicit and deeply buried in the initial construction of the formal model. There is a better chance of bringing those assumptions to light by comparing two very distinct models which bear no genealogical relation, (although, clearly, the resulting task could be more difficult).

At the risk of overemphasizing the point, it must be said that there is nothing wrong with Miller’s proposal in itself as a guidance for people with specific projects for which it makes sense. Its problem is that it does not constitute a methodological solution to the general case, and this is




Figure 4.1: Direct explanation.

because the issue of *building* a model for use in a simulation is side-stepped by it. If one follows Miller systematically, the user of a computer simulation will be passing the job of building models and formulating theories exclusively to theoretical biologists. The use of simulation will become just a collection of interesting techniques with little say in the generation of novel understanding. In more general terms, this is the same warning expressed by Fontana, Wagner and Buss (1994, p. 224); they remark that “to inform biology means not simply imitating it, but understanding it” and that the presumption must be avoided “that all the needed concepts [are] available to be imported from the other disciplines”.

4.2.6 Exploring, experimenting, explaining

Supposing that one is in a case in which a mathematical formulation is not practicable and there is not a vast amount of empirical evidence to use as a guidance in the construction and interpretation of a model, how should one proceed?

There is no general answer to this question. However, since a defense was offered of the view that an appropriate framework for addressing this type of questions is in terms of tools and their use, the following paragraphs describe a *possible* way of using simulations as tools in such an integrated manner. This description is no prescription. In particular, not much is said about how a simulation *should* be built, or when is it adequate for the job. Rather, some landmarks in its *use* are being pointed to which ultimately will help in assuring that the simulation plays a scientific role without undermining its potential. Nor is there anything particularly new about this description. This sort of methodology has been applied successfully in many instances (Boerlijst & Hogeweg, 1991b; Fontana & Buss, 1996; Hemelrijk, 1997, 1998, and others; see also chapters 6, 7 and 9).

The first preconception that must be changed is the idea that all that is required from a computer model is the choice of a plausible mechanism and the replication of a certain pattern in order to claim that an explanation of a similar natural pattern has been achieved. This viewpoint is based on the premise that successful replication implies understanding of how the pattern arises in the simulation (figure 4.1). It does not always work like that. Only in certain cases will the researcher be concerned with *just* the basic mechanisms of the model, most of which will be when one wants to present a proof of concept of the type: “it is commonly thought that M is needed to generate P , but here is a model in which M' , which is simpler (more plausible, nicer, etc.) than M , reproduces something that looks like P ”.

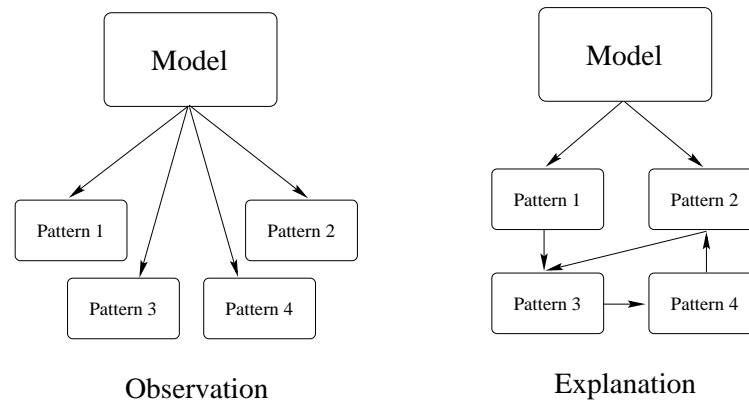


Figure 4.2: Indirect explanation.

While clearly an important use of computer models, proofs of concept are just part of what can be done with a simulation. In many cases it is not directly obvious how the patterns of interest arise or which aspects of the model are involved and which are inconsequential. Nobody dictates that an explanation must be uniquely constituted by the entities already present at the microscopic level of the model itself. Simulations must be explored, decisions have to be made about choosing relevant observables and these may lead to the discovery of emergent patterns, some of which may not have been suspected initially. Some of these observations can be explained in terms of the basic model but others may have to be explained in terms of higher order structures and patterns. Consequently, different observations have to be related through an *explanatory organization* which can be more complex than that shown in figure 4.1 but more like the one depicted in figure 4.2. This explanatory organization (which explains what is observed *within* the simulation) must then be related to the corresponding theoretical terms which describe analogous phenomena in the natural world.

It is possible to distinguish three different phases for achieving this:

1. *Exploratory phase*: Run different cases of interest, define observables, explore patterns, re-define observables or model if necessary.
2. *Experimental phase*: Formulate hypotheses that organize observations. Run crucial experiments to test these hypotheses. Explain what goes on in the simulation in these terms.
3. *Explanatory phase*: Relate the organization of observations to the *theories* about natural phenomena and the hypotheses that motivated the construction of the model in the first place.

The first two phases concern the simulation itself. Here the practitioner is dealing with her own created universe. The observations organized by the second phase prevent random fact-gathering and provide a theoretical perspective proper to the simulation that can be meaningfully compared with existing theories or hypotheses about natural phenomena (third phase). This comparison implies a backward “metaphorical step”. The first use of metaphors is found when the model is built. Entities in the model represent theoretical entities metaphorically or analogically. However, nothing guarantees that this same set of metaphors will be sufficient when one wants to project the observations made after running of the simulation back into existing theoretical entities about the natural world. This may be a trivial step if the observed patterns, or relationship between patterns,

have a corresponding counterpart in the existing theory. But it is possible to come across relationships between observations that do not translate into *existing* theoretical terms. This is a tricky but interesting situation, because lack of translation may mean that the model is not modelling what it is supposed to or that the model is pointing to a genuinely new phenomenon, which perhaps deserves a new name. The organizing theory of what it observed in the simulation can provide a new perspective for looking at the analogous theory of natural phenomena. Equally, the existing theory about natural phenomena may prompt a re-consideration of the theory of what is observed in the simulation, prompting the researcher to explore where the crucial differences lie and make a possibly useful conclusion about them¹¹.

By using the relationships between patterns in the simulation in order to explore the relationships between the theoretical terms that correspond to those patterns via metaphor, simulations are integrated as a sort of thought experimental machinery. By means of computer aided thought experiments, some theoretical terms may be shown to hold different relationships than previously thought. This is an unusual kind of thought experiment, though, given that one is not dealing with theory in a direct manner, but with computer simulations which must be observed and understood and then translated back into theoretical terms. An additional difference lies in the fact that it may indeed be possible to make a stronger case than with a “naked” thought experiment since the simulation also provides evidence that could not be arrived at by thinking alone.

As with traditional thought experiments, the information “fed” into the computer model may not be controversial but, in the end, the researcher may be forced to focus on facts/processes that were at the periphery of her conceptual structure and place them in a different relationship with other theoretical terms.

4.3 Summary

A pragmatic approach to the use of formal models in theoretical inquiry has been proposed in this chapter. This approach suggests that instead of looking for a good replication of natural phenomena, a formal model should be used in comparison with other similar models in order to expose the relevance of their differences and help in this way to understand the theoretical roles played by the corresponding factors.

A researcher practicing this method is less concerned with looking for the best or most complete model than with looking for models with enough potential to yield different results. Parsimony between the compared models can facilitate the understanding of how the small differences can affect the overall result but, for the same reason, they are bound to explore only a small subset of the possible implicit factors. In contrast, comparison between models with little genealogical connection is bound to reveal much more information but this information should be expected to be more difficult to obtain.

In the light of this, the use of mathematical models in comparison with computer simulations is particularly interesting. Simulations may be able to instantiate models which are radically different from mathematical ones, especially due to the relative ease of construction and lack of constraints which are needed to make the mathematical approach tractable. But this appreciation introduces

¹¹See (Hesse, 1980) for a description of the use of metaphors in science and particularly of the two-way conceptual dynamics which is generated when two domains are related metaphorically.

the question of how exactly should those computer models be built and used.

The question has been more carefully framed from the perspective of how new tools can be integrated into an existing network of practices. This required a brief description of the role played by tools and techniques in shaping scientific activity and how new tools can be adopted if their use is practically advantageous and validated in some form; for instance, by an acceptable theory of how the new tool works.

A description has then been offered of a way for computer simulations to be integrated into the rest of scientific activity following a pattern similar to that of a thought experiment. However, the activity demanded by the use of the tool in itself in order to reach this point also implies other methods usually associated with “experimental” tools: observing patterns in a simulation, postulating and testing relationships between them, etc., plus an additional step of metaphorical translation into existing theoretical terms, that may not be a trivial inversion of the initial metaphors that gave shape to the model.

In terms of the activities required under this scheme, computer simulations do have an experimental flavour, which is manifested in the commonly adopted phrase “computer experiment”. Indeed, this phrase is not inappropriate since much of what of the required work involves understanding the simulation in itself, something for which observations and experiments are needed since the simulation will in general be too complex for achieving a direct understanding in terms of its micro-structure. Yet integration into the rest scientific practice is done from a more theoretical side, i.e. as a conceptual tool.

The scheme proposed describes the phases that take place after the simulation has been built. A scientist that intends to follow this scheme should have all these phases in mind from the start, although obviously not in detail. Since it is the last phase where the integration into the rest of scientific activity is actualized, this last phase will constrain the decisions taken in the other phases even back to the modelling phase. It is this last phase that, in addition to the set of initial theoretical questions, provides the norm that implementational issues, statistics collected, choice of vocabulary, etc. must satisfy. Since many details will become known only after all the phases have been completed, it will not be uncommon for many iterative “sweeps” to be required before some form of “convergence” is achieved.

This scheme is not intended to become a generalized methodology. Many interesting cases could be addressed, perhaps more simply, following the proposals by Kitano *et al.* or by Miller. A general methodology cannot be prescribed. Instead, a background which is deemed appropriate for addressing the methodological problem has been proposed. A consensual methodology will emerge if there are serious efforts to fully integrate computer simulation into existing scientific practice. A *possible* way of doing this has been discussed (and concrete examples will follow), yet the actual choice will be done by a community of researchers in their contributions.

Such a convergence will depend on how well people understand each other. One of the problems that comes with the use of simulations is the relative obscurity of underlying projects, methods and conclusions. It is not usually obvious what the model is supposed to represent, why can it not be made simpler or more detailed, what motivated the implementational choices and how hard is it for someone else to replicate the results. Much progress towards a consensual methodology will be achieved by improving the communicability of computational models. Computer

simulations are at a disadvantage in this issue compared with mathematical models. Anyone with a certain degree of mathematical literacy can follow the assumptions and results of a mathematical model and spot those things that make sense and those that do not. This is not the general case with computer simulations, even when computer code is shared (not an universal practice, by the way). In this sense, all the proposals discussed in this chapter aim at improving communicability in one way or another, simply because integration constrains both the projects and the methods used so as to make computer models more understandable to other scientists.

The actual practical value of the methodological proposals made in this chapter will be tested in the rest of the thesis. As mentioned before, the comparative approach to formal modelling will be used in chapters 5, 6 and 7. Also, chapters 6, 7, 8 and 9 are intended to demonstrate how explanations can be built based upon the careful study of different patterns and their interrelations (as in figure 4.2) and not simply as proof that the basic building blocks are responsible for the generation of these global patterns (as in figure 4.1).

Chapter 5

A game of action coordination: mathematical models

Up to this point some important topics have been presented and discussed. These topics belong to what can be considered a promising theoretical and methodological perspective for addressing the issues of interest in this thesis. From the present chapter up to chapter 9 a series of models will be advanced in order to explore some of the themes that arose in the previous exposition in more concrete manifestations as well as other more specific issues.

As was commented in chapter 4, it is not the concern of this thesis to make accurate models of natural phenomena and then test these models against actual measurements. Rather, the motivation behind the use of formal models lies in their utilization as tools for exploring conceptual relations and gaining insight into complex phenomena.

This series of models will start in this chapter with a study of some dynamical issues in the evolution of social behavior. Of particular interest will be the presentation of an evolutionary scenario approachable by the use of different formal models so that knowledge can be gained by performing systematic comparisons between them. The study will be motivated in questions regarding the evolutionary and ecological dynamics of social coordination in situations in which the interest of the participants may be said to be in conflict. As mentioned in the Introduction, the systemic view on social behavior presented in chapter 2 remains silent about functional issues like the conditions for cooperative behavior to evolve in the case of conflict of interest. Given the relevance of this topic in current evolutionary biology, it cannot be ignored by this thesis.

In order to facilitate the construction of the models it will be necessary to sacrifice quite a lot of the behavioral complexity involved in the concept of coordination between autonomous entities. The basic scenario will be limited to a simple and evolutionary fixed action-response game. Such games, frequently used in theoretical evolutionary biology (Hurd, 1995; Bullock, 1997, and others), are quite simple to formulate although, as will be seen later, the resulting behavior may be far from trivial. A treatment of more behaviorally justifiable models of social coordination will be delayed until chapter 9.

The basic action-response game is described in the following section. The next sections will present some key concepts of evolutionary game-theory and use them to analyse the proposed evolutionary scenario. The chapter will then proceed by relaxing some of the assumptions of the basic game-theoretic treatment in order to include an investigation of the evolutionary dynamics both in

the cases of infinite and finite populations. The influence of ecological and genetic constraints will also be addressed. Finally, a continuous model will be presented in which the evolving population of agents is spatially distributed and interactions occur within a finite locality surrounding each agent. The existence of spatial instabilities leading to pattern formation will be demonstrated and some of the questions arising from abandoning a continuous population variable in favour of a discrete one will be explored.

All the models presented in this chapter are mathematical in nature. They involve the formulation of a set of equations which specify the evolution of global variables of the system together with a set of constraints in the form of additional rules. Some results are shown analytically while others are obtained by solving the equations numerically¹. An effort was made to explain in detail the formulation of the models and the mathematical concepts utilized in analysing them where it was considered necessary. This has been complemented in some cases with references to appropriate sources where the reader may find extended treatment of the methods involved. The significance of main results have been highlighted in the text so that they can be understood even if the formal steps are not followed in all their detail.

The answer to some of the questions posed in this chapter will have to wait until the following one where a similar evolutionary scenario is investigated, this time directly in the form of an individual-based computer simulation. The relevance of many of the issues presented in chapter 3 will become evident in the next chapter. The models presented here will be of great importance for understanding what goes on in the individual-based computer model. Consequently, both chapters should be read as a single study.

5.1 The game

The evolutionary dynamics of simple action-response games [see for instance (Hurd, 1995)] can be investigated by comparing different formal approaches more or less systematically. The investigation will be restricted to cases in which the game presents a conflict of interest between the participants (Trivers, 1974) since these cases are considered to cover the more general ground and do not assume already established conditions of cooperative behavioral coordination which would probably require additional explanation.

Formally, the game consists of a simple tournament between two players, after which the outcome of the game is computed and each player receives a certain payoff but, in order to keep in mind the sort of situations that the game is intended to model, it is helpful to describe this game in metaphorical terms within the background of an ecological situation in which two organisms try to extract energy from a single shareable food source.

The inspiration is drawn from a simple scenario in which an adequate chain of interactions between players may coordinate ensuing behaviors cooperatively in a situation of conflict of interest. Suppose that an animal has found a piece of food and has to decide whether to advertise its position or to hide it from another, nearby con-specific. The first individual will try to access the piece of food, which requires an appropriate action such as digging the ground or climbing a tree. Such actions will be accompanied by different manifestations such as noises, movements,

¹In all cases numerical integration was done using either a second order Runge-Kutta method or an Euler method (Press, Teukolsky, Vetterling, & Flannery, 1992).

etc. which may be perceivable by other individuals and influence their behavior. The actions of the players are said to be coordinated cooperatively when the outcome of the interaction is the joint exploitation of the food source.

Suppose that the first individual is able to “choose” whether or not to fake these external manifestations so that he can access the food source and at the same time try to divert the other individual towards another activity. The spectator organism, in contrast, is able to “interpret” what it perceives and “choose” an action accordingly. Of course, by allowing for this type of freedom a rather simplistic abstraction is made of both physical constraints and those imposed by bodily dispositions. The noises and movements associated with climbing a tree cannot be “produced” when one is digging the ground nor can an aggressive gesture be derived from an indifferent body activity. Consequently, a study of games like this one can reveal nothing about issues like signal ritualization or character exaggeration. The relevance of these sort of physical and bodily constraints will be explored in more detail in chapter 9.

By following the above metaphor, a realistic payoff structure for the game can be worked out. It is clear that an action whose manifestation ensues the approach of another individual will result in the first participant receiving a lesser payoff (say, in units of energy) since the food will have to be shared. Another reasonable supposition is that energy extraction from a food source can be said to be generally more efficient if more than one individual participates so that, if the energy contained in the resource is E , then each participant receives $E/2$ as a result of coordination. If, on the contrary, their actions are not coordinated, the second player gets no payoff (because the first player has succeeded in distracting his attention from the food source) and the first one gets a proportion c of E which is greater than half but less than 100%. The rest of the energy remains in the food source. The parameter c serves as a way of tuning the strength of the conflict.

For the time being, it will be assumed in these models that evolution proceeds mainly by differential reproduction². Now, if cooperation, or the lack of it, depended totally on either one of the participants in an individual fashion, the preferred outcome of the game would be clear and uninteresting: the first participant would choose not to cooperate and the second to cooperate, [this is what is meant by a conflict of interest in these circumstances (Trivers, 1974)]. However, cooperation is not an individual behavior, it is a relation between the behaviors of more than one individual and their consequences. All an individual can do is either initiate an action or respond to an action initiated by another. In this sense, whether a particular action or response is cooperative or not is a contextual fact, and depends on the pool of different actions present in the population of players at that time.

Notice also that this game has a defined time-structure (which is why it is possible to talk about an action being “initiated” or “responded” to) and that the differences between roles (both in available actions and payoffs) makes this game asymmetric in this respect, (Maynard-Smith & Parker, 1976; Maynard-Smith, 1982).

Given this situation, the question of interest is how a population of players of this game, who can play either one of the two roles on different occasions and who can have different behavioral

²It was mentioned in section 3.2.3 that the natural drift perspective on evolution allowed different dynamics depending on the form of coupling between agents and environment. For instance evolution driven by differential reproduction is characteristic of organisms that share a niche. As the models in this and the following chapters get richer, it will become clear what other forms evolutionary dynamics may take.

strategies, will evolve.

5.2 Game-theoretic analysis

Introduced as an analytical tool for studying evolutionary problems where the fitness of an individual trait is frequency dependent (Maynard-Smith & Price, 1973), game-theory has proven to be a remarkably flexible and practical technique in evolutionary biology. It has been applied to problems such as the evolution of parental investment, sex ratios, parent-offspring conflict, prey-predator strategies, animal signalling, mimicry, etc., (Maynard-Smith, 1982). Inspired by the mathematical theory of games advanced by Von Neumann and Morgenstern in the mid-forties (Von Neumann & Morgenstern, 1944), evolutionary game theory translates the assumption of a rational economic agent, who decides which is the best strategy to play, into the world of populations of organisms evolving under natural selection and their characters and behaviors. It could be argued that this is a domain where the rationality assumption is generally more justifiable than in economic decision making where players, due to different factors, do not necessarily behave rationally. Such validation is given by assuming the existence of a unique force that drives the evolution of individual strategies towards a situation where each individual will behave optimally: natural selection.

However, as will be seen later, game-theory also has its limitations which, from the methodological perspective of this thesis, is not a reason to abandon its use but something to take advantage of in order to highlight the role played by different assumptions in the case under study.

Successful game-theoretic analyses often rely upon a set of far from innocent assumptions, for instance, about the space of possible strategies that may be adopted by an individual, the relative accessibility of those strategies (in evolutionary or developmental terms) and the relation of the corresponding behavior to the fitness of organism. The importance of some of these issues has previously been addressed in the literature. For instance, regarding the appropriate choice of strategies Maynard-Smith (1982, p. 6) comments that the predictions of a game-theoretic model by Brockmann, Grafen and Dawkins (1979) on the competition between female digger wasps were incorrect because the available strategies were initially interpreted incorrectly. The question of accessibility, and particularly the influence of developmental constraints (Gould & Lewontin, 1979) is similarly treated by Maynard-Smith (1982, p. 5) as highlighting the need for a careful definition of the space of available strategies. However, this can only be done in those cases where development as a process unfolding in time with its own internal logic and external influences can be successfully decoupled from the evolutionary process³. The very possibility of such an approximation is still an issue very much under discussion. As to other types of constraints, such as those arising from genetic mechanisms, these are usually recognized and subsequently ignored by claiming that a game-theoretic approach takes the “phenotypic gambit”, i.e. the assumption that those constraints are not relevant, [see Grafen (1991) for a full discussion and section 5.3 below].

The question of the relation of the payoffs given by a particular strategy to the Darwinian fitness of the individuals using it is also a subtle one. Most game-theoretic models deal in terms of payoffs assigned to the players after the game. These payoffs, which are often interpreted as

³This issue will be addressed in chapter 7.

energy, resources, territory, mating opportunities, etc., are directly translated into fitness points. In general such a translation obeys the simple rule of “the-more-the-better”, which implicitly relies on a symmetric situation where “all other things are equal” such as the frequencies with which an individual engages in games, the relevance of the payoff received to the individual’s historical, developmental, social and ecological situation, etc.

The relative importance of some of these limitations will be explored by presenting a series of game-theoretic models with different levels of granularity and comparing the results obtained in each case. Other issues, as already mentioned, will be explored in subsequent chapters.

5.2.1 Strategies

The range of possible individual behaviors is divided into a discrete set of pure strategies which specify the actions performed by a player corresponding to each role. Suppose that the available types of food sources can be separated into two categories, those that require an action “A” and those that require “B”. Suppose also that performance of action “A” can be accompanied by one of two possible perceivable manifestations α and δ and performance of “B” can be accompanied either by β or γ . In order to define the strategy of a player in the *first* role it is enough to specify which manifestation accompanies each action. The analysis is simplified, without losing too much generality in the final results, by also requiring that, in a same individual, manifestation α for action “A” occurs in concordance with manifestation β for action “B” and correspondingly for δ and γ . By thinking of external manifestations as signals it is possible to say that α -emitters will also be β -emitters and δ -emitters will also be γ -emitters.

In order to define the strategy for the *second* role it is necessary to specify how a player responds to the different perceivable manifestations of the actions of the first player. In keeping with the above simplification some players will perform actions “A” or “B” when perceiving α or β and a non-relevant action in other cases and others will perform “A” or “B” when detecting δ or γ . The combination of the strategies for the first and second roles results in four possible overall strategies or types of player. Tables 5.1 and 5.2 show the behavior of these four types. Table 5.1 lists the external manifestation or signal corresponding to the actions “A” and “B” according to the type of agent. This table determines which signal accompanies the action of the agent playing the first role. The signal is “interpreted” by the second player according to table 5.2, where “-” means a non-relevant action.

Action	Type 1	Type 2	Type 3	Type 4
A	α	δ	δ	α
B	β	γ	γ	β

Table 5.1: External manifestations associated with actions “A” and “B” performed by each type of player when playing the first role.

Coordinated behavior occurs when both players end up performing the same action (“A” or “B”) which must correspond to the action required by the food source. In other cases behavior is not coordinated. By examining tables 5.1 and 5.2 it is possible to work out the cooperative relationships between the four strategies. These are shown in figure 5.1 where each arrow represents

Signal	Type 1	Type 2	Type 3	Type 4
α	A	A	-	-
β	B	B	-	-
δ	-	-	A	A
γ	-	-	B	B

Table 5.2: Action (“A”, “B”) performed by each type of player when playing the second role, depending on the perceived external manifestation (“Signal”); “-” means a non-relevant action.




Figure 5.1: Cooperative relationships between the four types of player.

a game that culminates in coordination when the strategy at the origin of the arrow is played in the first role and the strategy at the end is played in the second. A cyclical relation between strategies becomes apparent. This relation is similar to the (symmetric) Rock-Scissors-Paper game (Maynard-Smith, 1982)⁴. Notice also that players of types 1 and 3 are able to coordinate actions with players of their own type. These will be called “coordinating” or cooperative types.

From the payoff assignment scheme described above (section 5.1) and assuming $E = 1$ the payoff matrices for the first and second role ($M1$ and $M2$) are:

$$M1 = \begin{pmatrix} 1/2 & 1/2 & c & c \\ c & c & 1/2 & 1/2 \\ c & c & 1/2 & 1/2 \\ 1/2 & 1/2 & c & c \end{pmatrix} \quad M2 = \begin{pmatrix} 1/2 & 0 & 0 & 1/2 \\ 1/2 & 0 & 0 & 1/2 \\ 0 & 1/2 & 1/2 & 0 \\ 0 & 1/2 & 1/2 & 0 \end{pmatrix}$$

where $(M1)_{ij}$ is the payoff obtained by a player of type i when playing the first role against type j and $(M2)_{ij}$ is the payoff obtained by a player of type i when playing the second role against type j .

Players can play either one of the two roles on different occasions. If the frequency with which a player plays one role or the other does not depend on the role or on the strategy played, then this asymmetric game can be transformed into a formally equivalent symmetric one with the following payoff matrix:

⁴An example of an asymmetric (role-dependent) cyclic case is given by Dawkins (1976) for the “battle of the sexes” conflict of parental investment. A dynamical study of this example has shown the presence of oscillation in the space of strategies, (Schuster & Sigmund, 1981).

$$M = \frac{1}{2} \begin{pmatrix} 1 & 1/2 & c & c+1/2 \\ c+1/2 & c & 1/2 & 1 \\ c & c+1/2 & 1 & 1/2 \\ 1/2 & 1 & c+1/2 & c \end{pmatrix}$$

where $(M)_{ij} = (M1 + M2)_{ij}/2$ is the average payoff obtained by a player of type i over a large number of games against players of type j .

5.2.2 Infinite population: static analysis

One of the key theoretical tools in evolutionary game-theory is the concept of an Evolutionarily Stable Strategy or ESS (Maynard-Smith & Price, 1973; Maynard-Smith, 1982, ch. 2). A trait or behavioral strategy whose fitness depends on the distribution within a population of other traits or behavioral strategies is said to be an ESS if in the case that it is adopted by the majority of the population it will not be invaded by descendants of individuals who bear a different trait or use a different strategy. This resistance to invasion is a consequence of the fact that, by definition, the fitness of any trait or strategy other than the ESS is always lower than that of the ESS or at most equal to it. And, in the latter case, the alternative strategy will eventually be invaded by the ESS. This definition can be formalized in the following manner.

Let $p = (p_1, p_2, p_3, p_4)$ be a vector describing the proportion of players of types 1, 2, 3, and 4 in a polymorphic population which is considered infinite⁵. The sum of the components of p is always equal to one, which means that p belongs to the simplex S :

$$p \in S = \{x = (x_1, x_2, x_3, x_4) \in \mathbb{R}_4 : \sum_j x_j = 1, x_i \geq 0, \forall i\}.$$

An ESS, if it exists, is given by a distribution p^* when:

$$p^* \cdot Mp^* \geq q \cdot Mp^*, \quad \forall q \neq p^* \quad (5.1)$$

and

$$p^* \cdot Mq > q \cdot Mq \quad \text{if} \quad p^* \cdot Mp^* = q \cdot Mp^*, \quad (5.2)$$

where M is the payoff matrix and $x \cdot My$ is the average payoff of a sub-population described by distribution x within a population that plays with a distribution y . The above conditions can be interpreted as saying that, if p^* is an ESS, in a population described by distribution p^* a small group described by a different distribution q will have at most an equal average payoff when playing against the rest of the population and in that extreme case the rest of the population will do better when playing against players in the small group than these players against themselves.

⁵A polymorphic population is a population in which more than one pure strategy can be present. Under certain circumstances results obtained for a polymorphic distribution p are equivalent to a case with mixed strategies where each player plays strategy i with probability p_i , however this is not generally the case, see (Maynard-Smith, 1982, Appendix D). The cases studied here are restricted only to pure strategies.

One way of finding an ESS involves looking for the extreme points of the function $F(x) = x \cdot Mx$ subject to the constraint that $G(x) = \sum_i x_i - 1 = 0$, and then verify that the solutions obtained fulfill the ESS conditions. Using the method of Lagrange multipliers⁶, the extremes points can be found by solving simultaneously the system:

$$\nabla F(x) = \lambda \nabla G(x), \quad (5.3)$$

$$G(x) = 0. \quad (5.4)$$

A solution set for this system is given by:

$$P = \{p_a = (a, \frac{1}{2} - a, a, \frac{1}{2} - a); 0 \leq a \leq \frac{1}{2}\}.$$

The average payoff for a distribution of strategies given by p_a is:

$$p_a \cdot Mp_a = \frac{1}{2}(c + 1), \quad (5.5)$$

which means that all strategies in P are neutral because the average payoff does not depend on the parameter a . Therefore, none of them can be an ESS since conditions (5.1) and (5.2) cannot hold between strategies characterized by different values of a . Consequently, the game does not possess an ESS. However, this information does not reveal much about how an evolutionary system like this should be expected to behave. This question is addressed in the following section.

5.2.3 Infinite population: dynamical analysis

By definition the ESS concept is a static one. It refers to a situation of stable equilibrium. However, the same kind of game theoretic analysis that leads to the notion of ESS has been successfully applied to model evolutionary dynamics by simply “assuming the hypothesis that the growth rate of those playing each strategy is proportional to the advantage of that strategy”, (Zeeman, 1979, p. 471). Dynamic approaches to evolutionary game theory were first presented by Taylor and Jonker (1978) and further developed by Zeeman (1979, 1981) and others (Hofbauer, Schuster, & Sigmund, 1979; Schuster & Sigmund, 1981) and, in general, the resulting dynamical equations are similar in form to those studied in macromolecular evolution, (Eigen & Schuster, 1979; Schuster, Sigmund, & Wolff, 1978, and others).

The rate of growth of the proportion of individuals playing a particular strategy i is proportional to the current number of i -players and to the relative advantage of those players with respect to average fitness of the whole population. This can be denoted in the following way:

$$\frac{dp_i}{dt} = p_i(e_i \cdot Mp - p \cdot Mp), \quad \forall i \quad (5.6)$$

where e_i is the vector whose components are all 0 except the i th which is 1. The first term in the parenthesis indicates the average payoff for strategy i when playing against a distribution of strategies p and the second term, as before, is the average payoff for the whole population.

Zeeman (1979) puts the ESS concept in perspective by studying its relation to the attractors of the dynamics. He presents proof that an ESS is a fixed point attractor although a fixed point

⁶Most introductory texts on multivariate calculus discuss this method, see for instance (Finney & Thomas, 1990; Marsden & Tromba, 1996).




Figure 5.2: Representation of a vector of four components of fixed sum as an internal point of a tetrahedron (right) as an extension of the triangle diagram for representation of three variables of fixed sum (left).

attractor is not necessarily an ESS. If there is an ESS in the interior of S (denoted by $\overset{\circ}{S}$), this is an unique attractor and its basin of attraction is the whole of $\overset{\circ}{S}$. However, there may be non-ESS attractors in $\overset{\circ}{S}$ [they do not satisfy conditions (5.1) and (5.2)] which are also unique in $\overset{\circ}{S}$ but do not drain the whole of S and co-exist with other attractors in the frontier of S . A population situated in such internal attractor will remain stable against invasions in the long run. According to Zeeman this is “where the notion of ESS fails” since it cannot account for the possibility that the stable population may be initially invaded by mutants (i.e. mutant strategies having initially a positive net fitness advantage over the rest of the population) but after an increase in the mutant population this advantage is lost and the stable situation is restored. In the particular game presented here no ESS was found but since this does not mean that there may not be other attractors it is necessary to proceed with the dynamical analysis.

First, the fixed points of the dynamics are found by equating the time derivatives in (5.6) to zero. The solutions found in S are $(1,0,0,0)$, $(0,1,0,0)$, $(0,0,1,0)$, $(0,0,0,1)$ and again the set $P = \{p_a = (a, 1/2 - a, a, 1/2 - a); 0 \leq a \leq 1/2\}$ described previously. The first four are saddle points (i.e. certain trajectories move toward them and others move away from them) and the points in P are point-wise fixed so that movement along P is neutral.

So what kind of dynamics is to be expected in this game? Equations (5.6) are solved numerically. In order to represent the trajectories of the four population variables in a single plot a tetrahedron diagram is used (figure 5.2). This is an extension to four variables of the triangle diagrams for representing relative compositions in systems of three variables. Each component of $p = (p_1, p_2, p_3, p_4)$ is represented as the height from the face of the tetrahedron opposing the vertex label with the same number. The sum of all heights is always equal to 1 which makes it possible to represent the four variables at the same time in a 3-dimensional space thanks to the fact that their sum is also always 1. Each vertex indicates a coordinate where the corresponding type constitutes the whole population.

The resolution of equations 5.6 is shown in figure 5.3. As expected from the relationship between strategies (see figure 5.1), there is an oscillatory attractor in which each strategy becomes successively dominant for a period of time, [figures 5.3 (a) and (b); the reader should notice how in the long term dynamics the trajectory passes near each vertex of the tetrahedron.

The lack of strong stable points does not mean that all fixed points are necessarily unstable.




Figure 5.3: Trajectories in strategy space for the action-response game under equations (5.6). Initial conditions, (a): $p_0 = (0.27, 0.23, 0.24, 0.26)$, (b): $p_0 = (0.81, 0.09, 0.01, 0.09)$.

In fact, by moving the initial conditions continuously along the set P it is found that some points in P are weakly stable, i.e. they are neutral with respect to translations to neighboring points in P , but stable to other perturbations. This can be proved by performing a linear stability analysis⁷. The three eigenvalues⁸ for the linearized equations (5.6) are calculated around a general point p_a in P . One of the eigenvalues, as expected, is always zero (indicating neutral displacement within the neighbouring points in P) and the others are given by:

$$\lambda_{1,2} = ac + \frac{1}{4} - \frac{1}{4}c - a \pm \frac{1}{4} \sqrt{16c^2a^2 - 8c^2a + 1 - 2c + c^2}$$

Figure 5.4 shows that the real part of the eigenvalues, $Re(\lambda_{1,2})$, is negative for $1/4 < a < 1/2$. In this region there is a higher proportions of strategies T2 and T4 than T1 and T3. Points in this region are expected to behave as weakly stable and correspondingly points in the region given by $0 < a < 1/4$ to behave as unstable. This is shown by adding a perturbation term to (5.6) in the form of a small random fluctuation (white noise with range 0.01). The results can be seen in figures 5.5 (a) and (b). In the latter the initial condition is on the weakly stable subset of P and so the movement is confined only to nearby points. In contrast, the initial condition for figure 5.5 (a) is on the unstable subset of P and quickly gets attracted towards the oscillatory dynamics until, by chance, it enters into the weakly stable region and remains there.

What all this means is that, even though there is not stable point in the infinite population dynamics, there is a region of the strategy space where the population is most expected to be found after some time. This is the quasi-stable region where strategies 2 and 4 are equally abundant and

⁷This is one way of analysing the stability of fixed points. It consists of making a linear approximation for the nonlinear dynamics in the vicinity of the fixed point and evaluating whether small perturbations tend to move away from the fixed point or to return to it. This is achieved by calculating the eigenvalues for the temporal dynamics and examining the sign of its real part. If it is greater than zero it means the perturbation has a tendency to increase the distance away from the fixed point in the corresponding direction. The fixed point is therefore classified as unstable. If it is negative the perturbation will tend to decrease and to restore the equilibrium which is then classified as stable. For an introduction and examination of the possible conditions and other methods see (Haken, 1978; Murray, 1989) and others.

⁸Only 3 of the 4 equations (5.6) are independent, since $\sum_i p_i = 1$.

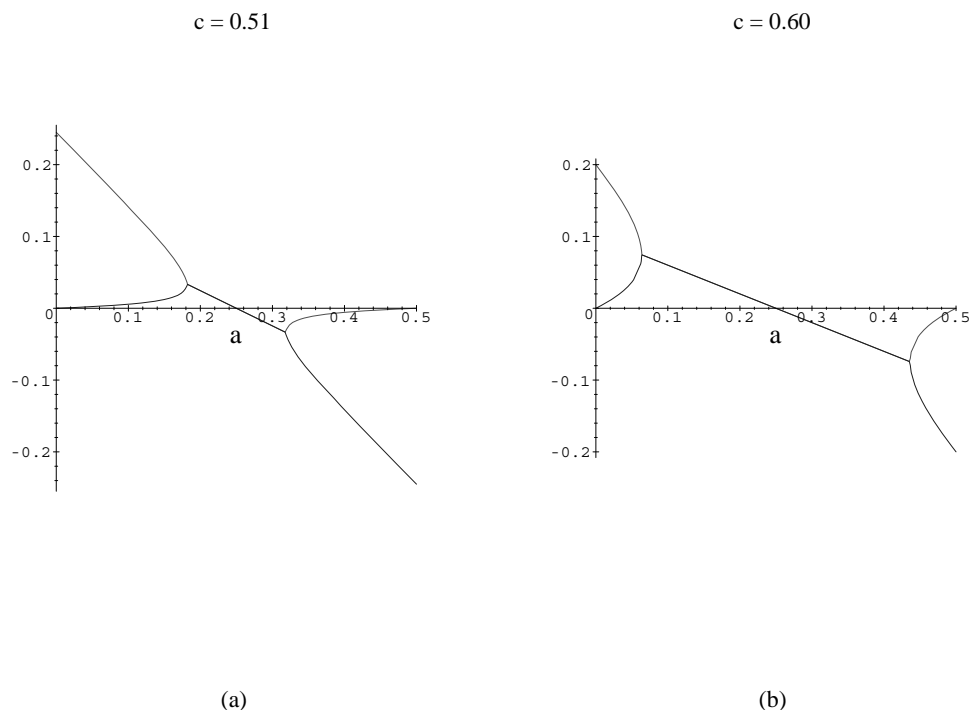


Figure 5.4: Real parts of non-zero eigenvalues (vertical axes) for points in point-wise fixed set P as a function of parameter a for two values of c . Each plot shows two eigenvalues as a function of a . One eigenvalue is always greater than the other except in the linear region in the middle where they are equal. Both have negative real part for $1/4 < a < 1/2$.

are present in a larger proportion than strategies 1 and 3. The symmetrically opposed region is unstable. In other words, the quasi-stable region is that of non-coordination.

In conclusion, two different kinds of behaviors can be expected in the dynamical infinite model, either oscillations in the case without noise or quasi-stability of non-coordinating strategies when noise is added.

5.2.4 Finite population: dynamical analysis

At present it is accepted that the infinite population limit in Darwinian selection can be a strong simplification whose consequences are not merely quantitative. With the use of mainly statistical and stochastic techniques the role of finite populations in evolution has been highlighted as a limiting factor for allele substitution under natural selection (Kimura, 1995) as well as an order production factor even in the total absence of selection (Derrida & Peliti, 1991). And from the point of view of evolutionary dynamics finite populations have been shown to induce metastability, i.e. epochs of stasis punctuated by rapid changes in global fitness, (van Nimwegen, Crutchfield, & Mitchell, 1997a, 1997b).

In the case of evolutionary game theory, however, the main effect of considering finite populations is given by the fact that different players “perceive” different distribution of strategies. In the ESS conditions (5.1) and (5.2) as well as in the replicator dynamics given by (5.6) the quantity




Figure 5.5: Trajectories for the perturbed dynamics. Initial conditions, (a): $p_0 = (0.3, 0.2, 0.3, 0.2)$, unstable region of P , (b): $p_0 = (0.2, 0.3, 0.2, 0.3)$, stable region of P .

p_i plays the dual role of density of individuals playing strategy i on the one hand and probability of one player randomly selecting another player of strategy i on the other. For these magnitudes to be equal the population must be infinite so that the fact that a player cannot play against itself has no effect. Other factors introduced by finite populations, such as random fluctuations in the way players choose their partners can be safely ignored in deterministic models under the assumption that players play a large number of games whose fitness effect is cumulative.

As a consequence, it is found that, under generalized equilibrium conditions, strategies that were ESS for the infinite population case may not be so for finite populations, (Riley, 1979; Vickery, 1987; Schaffer, 1988). Much of the discussion concerning this issue has been focused on the general stability of mixed strategies (Vickery, 1987; Maynard-Smith, 1988; Vickery, 1988) but, as will be seen in the present case, for a polymorphism of pure strategies the consequences can also be important.

The effects of finite population will only be explored for the dynamical analysis presented above. First, equations (5.6), which are formulated for an infinite population, should be modified. The change is straightforward after noticing that in (5.6) the first term in the parenthesis ($e_i \cdot Mp$) represents the average fitness W_i for playing strategy i , and the second term ($p \cdot Mp$) represents the average payoff (\bar{W}) for the existing distribution of strategies. In the finite population case the population is characterized by the number of players P_i for each strategy. The total number of players N can be assumed as a constant given, for instance, by ecological factors [see sections 5.3 and 5.5]. In a purely random situation the probability for a player of strategy i of finding a player of a different strategy j will be given by $P_j/(N-1)$ and the probability of finding a partner playing the same strategy will be given by $(P_i-1)/(N-1)$. Then the average fitness for playing strategy i will be:

$$W_i = \frac{1}{N-1} \left(\sum_j M_{ij} P_j - M_{ii} \right) \quad (5.7)$$




Figure 5.6: Trajectory for perturbed finite population dynamics with $N = 100$. Initial condition, P^5 . (a) continuous evolution, (b) quasi-discrete evolution, (see text).

and the average fitness for the whole population will be given by:

$$\bar{W} = \frac{1}{N} \sum_k P_k W_k. \quad (5.8)$$

And so (5.6) are transformed into:

$$\frac{dP_i}{dt} = P_i(W_i - \bar{W}), \quad \forall i. \quad (5.9)$$

As in the case of equations (5.6), it is easy to see that the sum over i of all the derivatives is null and therefore $\sum_i P_i$ is constant, as required.

A search is made for relevant fixed points in (5.9) (i.e. those that fulfil the condition $\sum_i P_i = N$) finding the following:

$$\begin{aligned} P^1 &= (N, 0, 0, 0) & P^3 &= (0, 0, N, 0) & P^5 &= (N/2, 0, N/2, 0) \\ P^2 &= (0, N, 0, 0) & P^4 &= (0, 0, 0, N) & P^6 &= (0, N/2, 0, N/2) \end{aligned}$$

By performing a linear stability analysis P^6 is found to be the only stable attractor. There is already an interesting difference with the infinite case. Instead of a set of point-wise fixed points only the two extremes of that set (P^5 and P^6) remain fixed. Figure 5.6 (a) shows the numerical resolution for $N = 100$ and initial condition P^5 where a small random term has been added to test stability. In this tetrahedron diagram P_i/N is plotted as the density of i -strategists in the population⁹. The trajectory is drawn towards P^6 .

Although equations (5.9) are treated as continuous over the P_i , these in fact represent discrete quantities (whose minimum variation is 1 individual) and therefore linear stability may not guarantee that a fixed point will be stable against discrete perturbations. In order to test if this is so it is helpful to perform a discrete perturbation analysis on the above fixed points consisting of the addition of one individual to one of the sub-populations and a compensating subtraction from another in all possible combinations. Examining the time derivatives it can be determined whether

⁹To simplify the notation, the vertices of the tetrahedron are label just with the corresponding vector P^i .

there is a tendency to return to the fixed point or move away from it. In all cases the linear stability diagnosis has been confirmed, although in some, for instance P^6 , it was necessary to follow the dynamics for many steps using a “quasi-discrete” approach to the numerical resolution of equations (5.9). This is done in the following manner: instead of approaching variations in P_i in the integration method as

$$\Delta P_i \sim \frac{dP_i}{dt} \Delta t$$

the following approximation is used:

$$\Delta P_i \sim INT\left(\frac{dP_i}{dt} \Delta t\right)$$

where $INT(\cdot)$ means “integer part of”, and the remainder is accumulated for the next integration step. Quasi-discrete resolution can introduce important differences in the overall dynamics, (see discussion in section 5.5.4 below), as can be observed in figure 5.6 (b) which shows a trajectory for a quasi-discrete resolution starting from the same initial condition as figure 5.6 (a).

5.3 Ecological coupling

In the simplified ecological setting assumed so far the distribution of strategies in a population has no effect in the payoff structure of the game. However, more detailed ecological considerations would require an account of how differently constituted populations exploit the available finite resources. Going back to the original metaphor that inspired the game it is easy to realize that, if the energy contained in food sources is renewed more or less independently of their use by players, then the average energy level $E(t)$ in any given food source will depend on the exploitation by an average player, i.e. on the population constitution. Until now $E(t)$ has been considered constant and equal to 1.

The expression of payoffs in energetic and fitness terms was taken as equivalent in the previous models. However, an explicit treatment of the environmental energy as an extra variable permits the expression of the evolutionary dynamics purely in energetic terms and lets the effects of relative fitness between different sub-populations be a consequence of the “environmental coupling”. This means that the optimization assumption which is explicit in equations (5.6) and (5.9) can be relaxed. It can be noticed in these equations that the time derivatives are proportional to the difference between the fitness of a given strategy and the average fitness of the population. Such an assumption is meant to provide a competitive coupling between individuals in order for natural selection to act. Instead of this, the time derivatives are now made proportional to the net energy gained by a given sub-population, which will be given by the difference between the incoming energy due to payoffs obtained from playing the game and the outgoing energy due to fixed self-maintenance costs. In the end, in the case where environmental coupling involves a single energy variable accessible to all players, optimization will be recovered as a result of the dynamics. This is hardly surprising since it is a classic case of competition for shared resources. However, as will become clear later, this need not be so if resources are distributed and accessibility is not global.

Suppose that players can access a number F_Q of food sources, each with an average energy of E/F_Q . E is renewed at a rate R per unit of time with a saturation time constant given by γ_e . The




Figure 5.7: Trajectory for ecological model with added perturbation. (Range of added white noise: 0.1 per time step, 3×10^6 time steps).

coupled equations describing the dynamics of the different populations and environmental energy can be then be written as:

$$\begin{aligned} \frac{dP_i}{dt} &= BP_i(\beta W_i E - L_o) \\ \frac{dE}{dt} &= R - \gamma_e E - \beta \bar{W} \frac{N}{F_Q} E \end{aligned} \quad (5.10)$$

with W_i and \bar{W} as defined above [(5.7) and (5.8)], now interpreted as the average proportion of energy E taken in by a player of strategy i per game and the population average of the same proportion for all strategies respectively. β is a constant, added mainly for dimensional reasons, which specifies the number of games played by one individual per unit of time ($\beta = 1$ is assumed). L_o is the individual energy cost per unit of time due to self-maintenance and B specifies how many individuals are born per unit of energy gained. For simplicity, players do not accumulate energy individually but instead all the energy obtained by a sub-population is transformed into new individuals of the same population. The same constant B is also used to specify how many individuals die when energy is lost at the sub-population level. All the energy used by the players is discounted in the last term on the right-hand side of the energy equation, where the total energy intake is calculated by multiplying the number of players (N) by the average energy intake per player ($\beta \bar{W} E / F_Q$).

Figure 5.7 shows a trajectory for the numerical solution of (5.10) where, as before, small random perturbations were added, ($R = 0.05$, $\gamma_e = 0.01$, $c = 0.6$, $B = 0.01$, $L_o = 0.25$, $F_Q = 10^4$). Without the perturbations the system evolves towards the cyclic attractor described before. However, this attractor is unstable and the system ends up in the same stable fixed point attractor as the previous model.

The total number of players N is no longer a constant, but it is a result of the ecological interplay between the average payoff (i.e. population constitution) and the energy variable. Otherwise, the ecological level of description does not introduce important differences at this stage.

5.4 Genetic constraints

Another issue that remains unexplored is the influence of genetic constraints. As with most game-theoretic cases this analysis has made the so-called “phenotypic gambit” (Grafen, 1991), that is,

the calculations proceeded as if all strategies were freely realizable in any given population and their growth depended only on their payoff. It is clear, though, that if behavioral strategies, as other phenotypic traits, somehow depend on genetic constitution, then their realization may be constrained by factors other than the strategy payoff. In the most obvious example, the best strategy could be realized by a heterozygote in a diploid system. This would mean that a proportion of the non-optimal strategies realized by the corresponding homozygotes would exist at equilibrium while in the purely game-theoretic analysis they would disappear, (Grafen, 1991).

The purpose of this section, however, is not to describe a complex scenario, such as a diploid genetic system realizing the four strategies, in order to study the effects of genetic constraints since there is no evidence of how such a genetic system would be constituted. Instead, the simpler question of how a trivial haploid system in the presence of mutations would affect the evolutionary dynamics will be addressed. The main reason for investigating this question is because such a scheme will be included in the computer model described in the next chapter.

Players have two sets of choices, one set for each role, with two possibilities in each set. These choices are assumed to be controlled genetically by two independent loci in a haploid genotype each one with two possible alleles. The first locus specifies which set of signals the player uses when acting on the first role (see table 5.1) and the corresponding alleles are denoted by **A1** and **a1**. The second locus specifies how the player interprets perceived signals when acting on the second role (see table 5.2), the corresponding alleles being **A2** and **a2**. Table 5.3 shows each strategy with its corresponding genotype.

Strategy	Genotype
1	A1 A2
2	A1 a2
3	a1 a2
4	a1 A2

Table 5.3: Haploid genotype for each of the four strategies.

The addition of a point mutation operation to the dynamics establishes an additional metric over the strategy space. If the chances of a mutation are independent and the same for each locus a mutant will be more likely to end up playing one of the strategies “nearest” to the parental one, a change requiring just one mutation, than a strategy requiring a change of two alleles. If the probability of a mutation occurring is μ , ($\mu \ll 1$), then the chances of two mutation events leading to a “non-neighbor” mutant strategy is given, in this model, by μ^2 .

This mechanism is incorporated into the equations describing the ecological dynamics. Before modifying equations (5.10) it is important to notice that it only makes sense for mutations to occur if a sub-population is increasing in number, i.e. if its derivative is positive. Consequently, the following distinction is made:

$$Q_i = \begin{cases} BP_i(\beta W_i E - Lo), & \text{if } \beta W_i E - Lo \geq 0 \\ 0 & \text{otherwise} \end{cases}$$




Figure 5.8: Trajectory for ecological model with genetic constraints and added perturbation. Same initial conditions as in figure 5.7. (Range of added white noise: 0.1 per time step, 3×10^6 time steps).

$$M_i = \begin{cases} BP_i(\beta W_i E - Lo), & \text{if } \beta W_i E - Lo < 0 \\ 0 & \text{otherwise} \end{cases}$$

and so equations (5.10) can be re-written to include mutation:

$$\begin{aligned} \frac{dP_i}{dt} &= Q_i(1 - \mu_i) + M_i + \sum_{j \neq i} Q_j \mu_{ji} \\ \frac{dE}{dt} &= R - \gamma_e E - \beta \bar{W} \frac{N}{F_Q} E \end{aligned} \quad (5.11)$$

where μ_{ij} is the rate of mutation of newborn individuals of kind i to a different kind j and $\mu_i = \sum_{j \neq i} \mu_{ij}$.

Figure 5.8 (a) shows the trajectory for the numerical resolution of (5.11) for a mutation rate of $\mu = 0.01$ per loci per unit of time. The initial conditions are the same as in the case shown in figure 5.7 and so is the added random perturbation. As in that case, the system ends up in the same fixed point attractor where the whole population is constituted by strategies 2 and 4 in equal proportions. However, it is clear that the existence of genetic constraints makes it more difficult to abandon the cyclic attractor than in the case with no mutation. Figure 5.8 (b) shows the case for a higher mutation rate ($\mu = 0.05$). Here it can be observed that, for the same range of noise, the system does not abandon the cyclic attractor at all.

5.5 Continuous spatial model with finite locality

So far the investigation has proceeded under the assumption that the population of players presents no structure with regard to the chances of two individuals playing against each other¹⁰. All players are equally accessible and this permits the mean-field approximation needed to formulate the

¹⁰Or, as will be shown later, the chances of playing one role instead of the other.

above models. This situation would change if players were distributed in space and were only allowed to choose their partners locally. The probability of two individuals playing against each other would then be modelled as a decreasing function of the distance between the two.

Spatial models have become a hot topic in biology in recent years, partly because of the non-intuitive results they yield and partly due to the possibility of implementing them relatively easily in computer simulations. Many issues come to the front when dealing with a spatially distributed population. One is the obvious fact that interactions become correlated with locality and this may provide conditions which, for instance, favour self-reciprocating strategies against parasitic ones by forming accumulations where the chances of meeting a “cheater” are very small (Axelrod, 1984; Krakauer & Pagel, 1995). Alternatively, local interactions may increase the chances of playing with relatives, if reproduction is also a local process, thus favouring a process of kin selection (Hamilton, 1964; Hines & Maynard-Smith, 1979; Michod, 1982; Ackley & Littman, 1994, and see also section 6.5).

Other spatial effects may depend on the local density of players. These may be related to the ecological effects discussed in the previous section. For instance, the actual payoffs of the game may vary with the local density of players¹¹ and/or local density of resources. Also, other density-dependent effects, not directly related to the game itself, may affect the absolute local rates of population growth or other factors like mobility. For instance, Cressman and Vickers (1997) compare two spatial games that reduce to a same non-spatial case when interactions are global. One of them assumes an additional fitness component which depends on the variable density of players and the other does not. The results are qualitatively different proving the point that spatial extensions to non-spatial games are far from unique.

A different factor that may influence the outcome of a spatial game is the formation of spatio-temporal patterns which impose certain regularity on the interactions between players. One of the most interesting recent examples is the case of models giving rise to spiral dynamical patterns in molecular evolution (Boerlijst & Hogeweg, 1991a, 1991b, 1995a, 1995b, and others) and similarly in host-parasitoid systems (Boerlijst, Lamers, & Hogeweg, 1993; Rohani & Miramontes, 1995; Comins & Hassell, 1996; Savill, Rohani, & Hogeweg, 1997, and others). The interesting situation found in these cases is that spatio-temporal patterns may present regularities of their own which enslave the dynamics at the micro level. For instance, Boerlijst and Hogeweg (1991a, 1991b) speak of competition between spiral waves that result in the exclusion of parasites which would take on the population if the dynamics were not spatial. In their model, rotating spiral waves are an emerging feature of the lower level dynamics. Their speed of rotation is related to the “fitness” of the constituting species. However, certain spirals of lower fitness rotate faster than other higher fitness spirals resulting in more rapid growth for their constituting species, even though this species would actually be less fit in a non-spatial model¹².

In general, a deterministic approach to spatial games involves some sort of reaction-diffusion or similar model¹³ able to account for the effects of density dependence and pattern formation

¹¹The stability of such density-dependent ESS (DDESS) has been analysed by Cressman (1990).

¹²In a model of the same system but based on partial differential equations instead of cellular automata, Cronhjort and Blomberg (1994) did not find the same stability against parasites. This is an interesting result because it points to the recurring theme of continuous versus discrete models. The relevance of this distinction will become evident in chapter 6; see especially section 6.6.

¹³See (Murray, 1989) for an introduction to the use of reaction-diffusion equations in biology.

as well as the effects of local interactions mentioned above. Such models may be formulated by simply adding a diffusion term to the basic dynamics given by equations (5.6) in order to account for spatial changes in densities, [see for instance (Vickers, 1989)] or they may be more complex (Cressman & Vickers, 1997; Durrett & Levin, 1994a). Typical reaction-diffusion equations are generally intended to model local interaction between infinitesimal components and they ignore long range effects. This is appropriate for contact processes such as chemical reactions but in the present case it would mean that the range of interactions of players would be very small (just the nearest neighbor). An usual answer to this situation is to “abandon the continuum of spatial scales ... in favour of a subdivision into patches” (Durrett & Levin, 1994a, p. 364). These patches may or may not have additional spatial structure. Durrett and Levin (1994a) present very interesting comparisons between spatial and non-spatial, patchy and continuous models. Unfortunately, models involving patches of population are difficult to explore deterministically which is the aim in the present section. So an alternative should be found in the form of a spatial model with finite locality as discussed below.

The remainder of section 5.5 will present the continuous spatial model in detail together with proof of its capability for the formation of spatial patterns, numerical explorations and a study of the effect of density thresholds. Some passages in this section are a bit more mathematically oriented than the preceding sections of this chapter. An effort was made to provide the readers not familiar with the techniques used in this section with enough guidance to facilitate understanding of how the analysis proceeds. However, given that a point by point understanding of this analysis is not fundamental for assessing the relevance of this model in the context of the rest of the thesis, the results will be summarized in section 5.5.5, in order to highlight the main points. If the reader so wishes, the summary may be read directly and the rest of the section ignored without much loss.

5.5.1 The model

The model is intended to be described using continuous variables but at the same time there is an interest in taking into account the possibility of long range effects instead of just hyper-localized interactions¹⁴. Therefore, a model with *finite* locality is presented. In this model players are distributed in space in positions that do not change with time. Each player can only interact with other individuals located within a neighborhood of finite radius L centered at the player’s position x . The neighborhood is denoted by $Loc(x)$. A model like this can be seen as a middle case between mean-field, non-spatial approaches and contact approaches where interactions occur only at infinitesimally small scales. By varying the parameter L both ends of the spectrum can be reproduced. Large values of L approximate global accessibility while small values approximate a contact process.

A sub-population of players of strategy i is described by its *spatial* density $p_i(x, t)$ where x is the spatial position. At any position, the total density of players is obtained by summing over all the strategies: $p_T(x, t) = \sum_i p_i(x, t)$. Energy resources are also distributed continuously in space; the amount of available energy at a given position and time is denoted as $E(x, t)$. All players play the first role in the game with a frequency which is proportional to the chance of finding a partner

¹⁴For an introduction to models that include long-range effects see chapter 16 of Murray (1989).

in the neighborhood who will play the second role. Each player per unit of space “intends” to play the first role with a uniform frequency F_0 and they make N_a attempts at finding a partner. If the probability of succeeding in locating a partner is $P_S(x, t)$ then the frequency for playing the first role will be given by:

$$F^1(x, t) = F_0 \{1 - [1 - P_S(x, t)]^{N_a}\}. \quad (5.12)$$

where the factor in curly brackets indicates the probability of having located a partner after N_a trials.

The probability $P_S(x, t)$ is simply taken as proportional to the total number of players $P^L(x, t)$ in $Loc(x)$, the neighborhood centered at x :

$$P_S(x, t) = \frac{P^L(x, t)}{\int_{\Omega} P^L(u, t) du}, \quad (5.13)$$

where a normalization has been made by dividing by the total number of players in the whole space Ω . The number of players in $Loc(x)$ is the sum of the total number for each sub-population, $P^L(x, t) = \sum_i P_i^L(x, t)$, and each term in the sum is calculated as:

$$P_i^L(x, t) = \int_{Loc(x)} p_i(u, t) du$$

Following the initial metaphor (section 5.1), it is supposed that energy is extracted from the location of the first role player. As in equations (5.10), the rate of change in the number of individuals of a given strategy will be proportional to the net energy accumulated by those individuals. This net energy is calculated by summing the contributions of the energy an individual gains when playing the first role and the energy it gains when playing the second role minus its individual energy cost for self-maintenance.

The energy received by an individual of type i situated at x as a result of its participation in games as a first role player is given by:

$$E_i^1(x, t) = F^1(x, t) E(x, t) \sum_j M1_{ij} \frac{P_j^L(x, t)}{P^L(x, t)} \quad (5.14)$$

where the matrix $M1$, defined in section 5.2.1, indicates the payoff received by a first role player, [$M1_{ij}$ is used to denote $(M1)_{ij}$]. The factor $P_j^L(x, t)/P^L(x, t)$ represents the local average density of players of strategy j within the neighborhood $Loc(x)$. Given that energy is extracted from the location of the first role player, $E_i^1(x, t)$ must be proportional to the available energy in that position and to the frequency with which an individual plays the first role in the game; hence the two first factors in (5.14). The remaining factor is simply the average payoff that a player of strategy i will receive given the number of players of each strategy present in the neighborhood.

The energy received by the same individual of type i due to its involvement in games as a second role player depends on the number of contests in which other players in the neighborhood have selected this individual as a partner and on the strategies played by those neighbors. The probability that the partner found by an individual located in position u be located in position x is simply one divided by the total number of players in $Loc(u)$, that is $[P^L(u, t)]^{-1}$. Integrating over

all the positions u in $Loc(x)$, the total energy received by a player of strategy i located at x for playing the second role is:

$$E_i^2(x, t) = \int_{Loc(x)} F^1(u, t) [P^L(u, t)]^{-1} E(u, t) \sum_j [p_j(u, t) M_{2ji}] du \quad (5.15)$$

The matrix M_2 , given in section 5.2.1, specifies the payoff obtained by a second role player. The integrand is readily interpretable as the proportion of all games originating at u which find a partner located at x times the average energy payoff received by that partner if it is of type i for its participation as a second role player.

Dropping momentarily the dependence with x and t , the coupled differential equations that define the system are given by:

$$\begin{aligned} \frac{\partial p_i}{\partial t} &= B p_i (E_i^1 + E_i^2 - L_0) + D \nabla^2 p_i \\ \frac{\partial E}{\partial t} &= R - \gamma_e E - 2 F^1 E \sum_{i,j} p_i M_{ij} P_j^L / P^L. \end{aligned} \quad (5.16)$$

A diffusive term was added to the first equation to account for mobility effects. The diffusion coefficient D is constant and the same for all strategies. In a way similar to equations (5.10), B gives the number of newborn players per unit of space for each unit of energy gained and L_0 indicates the energy costs per unit of time for a density value of one due to self-maintenance. R and γ_e represent the rate of energy renewal per time step and per unit of spatial dimension and the energy saturation time constant respectively. The matrix M is the same payoff matrix used in the previous sections, i.e. half the sum of M_1 and M_2 . The coupling term on the right-hand side of the energy equation can be interpreted as the number of games originating in position x per unit of time (F^1) multiplied by the average use of resources per game considering the contributions of both roles together.

In order to reduce the number of parameters in the model, equations (5.16) are non-dimensionalized by performing the following scaling operations:

$$t^* = \gamma_e t, \quad x^* = x/L, \quad E^* = E \gamma_e / R, \quad F^{1*} = F^1 / F_0, \quad p_i^* = p_i F_0 / \gamma_e.$$

The non-dimensional system can be written as (dropping the asterisks):

$$\begin{aligned} \frac{\partial p_i}{\partial t} &= p_i [A_1 (E_i^1 + E_i^2) - A_2] + D' \nabla^2 p_i \\ \frac{\partial E}{\partial t} &= 1 - E - 2 F^1 E \sum_{i,j} p_i M_{ij} P_j^L / P^L. \end{aligned} \quad (5.17)$$

where

$$A_1 = \frac{B F_0 R}{\gamma_e^2}, \quad A_2 = \frac{B L_0}{\gamma_e}, \quad D' = \frac{D}{\gamma_e L^2}.$$

The non-dimensional parameter A_1 gives an idea of the strength in the coupling between energy and density of players, it can be interpreted as a comparison between the energy needed to “produce” a typical number of players per unit of space $[(B F_0 / \gamma_e)^{-1}]$ and the maximum value for




Figure 5.9: Stability of homogeneous solution as a function of the non-dimensional parameters A_1 and A_2 under homogeneous perturbations. The vertical axis shows the real part of the growth factor λ for $k = 0$.

the environmental energy (R/γ_e). A_2 can be interpreted as a comparison between the characteristic saturation times for the energy (γ_e^{-1}) and the density of players $[(BL_0)^{-1}]$. D' is simply the scaled diffusion coefficient.

In general, in order to solve (5.16) or (5.17) it is necessary to specify an initial condition for each variable and a set of boundary conditions. A finite one-dimensional spatial domain will be used with two possible boundary conditions for all the variables: either Neumann conditions, i.e. all flux must be zero at the boundary (which for the 1-D case studied means that all spatial derivatives must be zero at the boundary) or the rather artificial, but useful, periodic conditions which specify, for 1-D domains, that the values adopted by the variables at one extreme of the domain must equal those at the other. The applicability of this last sort of boundary condition can be justified if the size of the domain L_D is much larger than the typical length of the system (L).

5.5.2 Pattern formation

It can be noticed that the system (5.16) [(5.17)] is symmetrical to transformations of the type $x \rightarrow x + \delta$ which suggests the existence of spatially homogeneous solutions. However, the system is not symmetric to scaling transformations of the type $x \rightarrow vx$ due to the finite size of neighborhoods. It would be interesting to explore the possibility of spatial patterns arising from the long range coupling between strategies and energy variable. In order to do this an analysis is performed of the stability of a spatially homogeneous solution for different spatial scales.

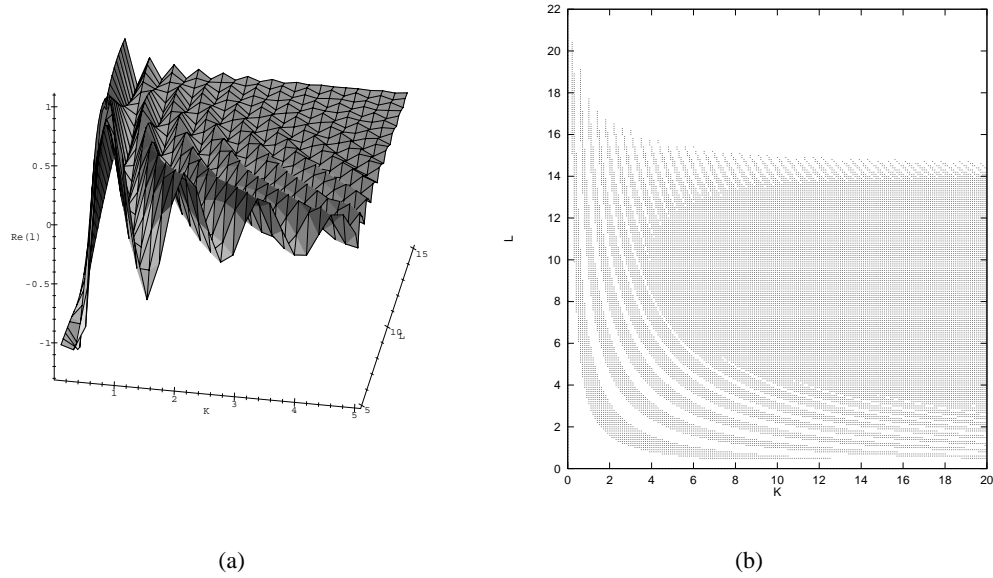


Figure 5.10: Dispersion relation. (a): the vertical axis shows the real part of the growth factor λ . (b): unstable combinations of L and K ,

This method is described in detail by Murray (1989) in the context of proving the existence of spatial patterns by testing the inhomogeneous instabilities of spatially homogeneous solutions, i.e. by testing the system's capacity for pattern formation. A finite domain of only one spatial dimension is assumed and the following steps are followed. First, it is necessary to find a homogeneous solution for (5.16) describing the same equilibrium state for all x . Then the linear stability of this solution is tested by adding to each variable a small in-homogeneous perturbation (the solution is assumed to be stable for homogeneous perturbations). In the linear limit the evolution of the perturbed variables will be given by a sum of terms with the form $\exp(\lambda t + ikx)$ where λ is the growth factor, k the wavenumber and $i = \sqrt{-1}$. This means that the perturbed state of the system can be described by a sum of wavelike spatial modes each with its own growth factor. If the homogeneous solution is stable this will be manifested with negative growth factors for all possible values of k . However, if for certain wavenumbers $\text{Re}(\lambda) > 0$ then the corresponding spatial modes will grow indefinitely until a point where the linear analysis is no longer valid and non-linear effects may take control. The fastest growing modes will end up originating spatial patterns.

In principle, the whole spatial domain is finite and this implies that the spectrum of wavenumbers k will be discrete and will depend on the boundary conditions. In order to simplify the analysis, it will be considered that the actual size of the spatial domain is much larger than the typical size of the system, i.e. $r \ll 1$ where $r = 2L/L_D$. This will allow a treatment of the wave spectrum as continuous, an often used simplification (Murray, 1989, p. 386).

As this is a proof of existence, it suffices to address particular cases. Consequently, other

simplifications can be made. First, the diffusion term is dropped since it is not the intention to test for diffusion driven instabilities. The proof will be restricted to the case where the whole population is constituted by players of strategy 2 (or, equivalently, 4). This will considerably simplify the analysis. More general cases can be studied with similar results. Pattern formation in the most general case will only be shown in the numerical resolution examples.

Taking into account that $M_{22} = 0$, $M_{12} = c$ and $M_{21} = c/2$, the simplified system obtained is:

$$\begin{aligned}\frac{\partial p_2}{\partial t} &= p_2(A_1 F^1 E c - A_2) \\ \frac{\partial E}{\partial t} &= 1 - E - F^1 E p_2 c.\end{aligned}\tag{5.18}$$

And the only non-trivial homogeneous solution is:

$$\begin{aligned}E^h &= \frac{A_2}{A_1 F^h c} \\ p_2^h &= \frac{A_1}{A_2} - \frac{1}{F^h}\end{aligned}\tag{5.19}$$

where $F^h = [1 - (1 - P_S)^{N_a}]$ is the non-dimensional frequency of games given by (5.12). P_S , given by (5.13), can be easily calculated since the distribution of players is homogeneous. In a 1-D domain of total size L_D , $P_S = r = 2L/L_D$.

The stability of solutions of the form $p_2(x, t) = p_2^h + \epsilon_p \exp(\lambda t + ikx)$ and $E(x, t) = E^h + \epsilon_E \exp(\lambda t + ikx)$ is now tested by replacing these expressions back into (5.18) and keeping only those terms that are linear in ϵ_p and ϵ_E .

After some calculations the factor $F^1(x, t)$ can be written as:

$$F^1(x, t) = F^h + \epsilon_p G(k) \exp(\lambda t + ikx) + O(\epsilon_p^2)$$

with

$$G(k) = \frac{r N_a (1 - r)^{(N_a - 1)} \sin(k)}{p_2^h k}$$

where again it is considered $r \ll 1$ so that the integral over Ω in the calculation of $F^1(x, t)$ can be approached as a constant [see (5.12) and (5.13)].

The following linear system is obtained for ϵ_E and ϵ_p :

$$\begin{aligned}[\lambda + 1 + c F^h p_2^h] \epsilon_E + [c E^h F^h + c E^h p_2^h G(k)] \epsilon_p &= 0 \\ [-c A_1 F^h p_2^h] \epsilon_E + [\lambda + A_2 - A_1 c E^h F^h + A_1 c E^h p_2^h G(k)] \epsilon_p &= 0.\end{aligned}\tag{5.20}$$

The condition for non-trivial solutions to exist is that the two equations not be independent, i.e. the determinant of the corresponding matrix must be null. By performing this operation a relation $\lambda(k)$ is obtained which tells what value of growth factor corresponds to each spatial mode defined by the wavenumber k . This is called the dispersion relation (Murray, 1989; Haken, 1983). Using this relation it is possible to study the temporal behavior of the different spatial modes, and in particular if any of them is unstable, that is if $Re[\lambda(k)] > 0$.

By equating the determinant of system (5.20) to zero two values for λ are obtained. One of them always has negative real part, so it can be ignored. But before analysing the other value it is

necessary to specify the value of some of the parameters. No exhaustive exploration of parameter space will be made since the aim is to prove that spatial patterns exist for reasonable parameter values. The following values have been used: $c = 0.6$, $N_a = 2$ and $L_D = 100$. Before specifying A_1 and A_2 it is important to recall that the homogeneous solution should be stable under *homogeneous* perturbations, that is $Re[\lambda(0)] < 0$. In order to see for which values of the non-dimensional parameters this is the case, a plot is made of $Re[\lambda(0)]$ vs. A_1 and A_2 for values of $r = 0.1$ and $r = 0.02$, (figure 5.9).

It is apparent that not for all values of A_1 and A_2 is the homogeneous solution stable under an homogeneous perturbation. In particular, it seems that for small values of A_2 the growth factor has a positive real part. To be on the safe side $A_1 = 1$ and $A_2 = 0.8$ have been chosen, which roughly means that the values given by the ecological parameters and by the population parameters for typical energies and times are of comparable magnitude. $Re[\lambda(k)]$ is now plotted as a function of the neighborhood radius L and as a function of the dimensionalized wavenumber $K = k/L$, (figure 5.10). In this figure it can be observed that certain combinations of neighborhood size and wavenumber are clearly unstable. In figure 5.10 (b) those combinations are plotted in the (K, L) plane (grey areas). For moderate values of L (remember that $L_D = 100$) there is a whole open region of instability. As L increases, the dynamics approach the case of the non-spatial game and patterns cannot form and this is manifested as a cut-off of the unstable region for $L \cong 20$. For small values of L both stable and unstable combinations are found in a complex pattern. And for $L \ll 1$ stability is recovered. The importance of having a parameterizable model of finite locality can here be appreciated. The expected behavior for the two extreme cases, non-spatiality and contact process without diffusion is recovered, but the behavior for intermediate locality sizes is very different.

The conclusion of this analysis is that the system described by (5.16) [(5.17)] is capable of pattern formation in the spatial domain. There is still no information about what kind of patterns will be formed, though. As will be seen, for some parameter values, the resulting patterns are stable accumulations (“lumps” or “clusters”) of players surrounded by areas of zero density.

It is important to point out that the mechanism of pattern formation is given by the coupling between the energy and population variables and not by diffusion mechanisms, which have been ignored in the analysis. The main factor responsible for the existence of spatial inhomogeneities is F^1 , the frequency of games which depends on the availability of partners in the finite neighborhood. If the density of players is too small, the opportunities for playing the game will be scarce and the players will tend not to be able to compensate their self-maintenance costs. In contrast, if the density is high, these opportunities are plenty and players will tend to accumulate in those areas. Such accumulation will be limited by nonlinear factors, such as the rate of energy exploitation.

5.5.3 Dynamics

The dynamics of the model are studied by with the numerical resolution¹⁵ of (5.17) in a finite one-dimensional domain of size $L_D = 200$ and with neighborhood radius $L = 10$ spatial steps of size

¹⁵The spatial domain is regularly partitioned and all spatial derivatives are calculated using finite differences. The update of the density values for each discrete position is performed synchronously once all the new values have been calculated.




Figure 5.11: Total (non-dimensional) number of players for spatial model. $A_1 = 0.25$, $A_2 = 0.05$, $D' = 10^{-4}$. (a) without random perturbation, (b) with random perturbations.

0.1. Periodic boundary conditions have been used as well as random uniform distribution for the initial conditions in the population densities (with range 2.0 per unit of space) and uniform initial value of 1 for the non-dimensional energy. The non-dimensional parameters are set to¹⁶ $c = 0.6$, $A_1 = 0.25$, $A_2 = 0.05$, $D' = 10^{-4}$. Figure 5.11 (a) shows the evolution of the four sub-population as a function of time. Oscillations in number of players can be seen for each sub-population. These oscillations compensate each other to give a constant total number of players in the whole domain after the initial transient. It is also apparent from this figure that two of the strategies (1 and 3) seem to have a smaller mean number of players than the other two. This difference, however, is not a generally observed result and it will be understood in this case once the patterns in the spatial domain have been examined.

Figure 5.12 shows the spatial distribution of players at different times, starting from a random initial condition. Players arrange themselves into two clusters. (The use of a non-zero diffusion term introduces little changes in the results, mainly by smoothing the shape of the clusters which are also found if diffusion is removed.) The relative composition of these clusters changes over time, but the behaviors are different. The composition of the cluster on the left settles, after an initial variation, into a stable state (at about $t = 5500$) where the densities of strategies 1 and 3 have the same variation with x and so do the densities of strategies 2 and 4. Integrating the densities over the cluster it is found that the cooperative strategies 1 and 3 form 43.1 % of the total population of the cluster. The cluster to the right does not settle into a stable state. It is this cluster that provides the oscillatory dynamics observed at the population level in figure 5.11 (a) while the other cluster provides the difference in mean value for the two pairs of strategies. This is a nice example because it shows, at the same time, the two possible situations for a numerical resolution in the absence of noise and because it illustrates a trivial but important property of separation between clusters: independent evolution. However, the more frequently observed behavior is given by the cluster on the right.

¹⁶These values were chosen for convenience in numerical resolution and similarity to the computer model presented in next chapter.

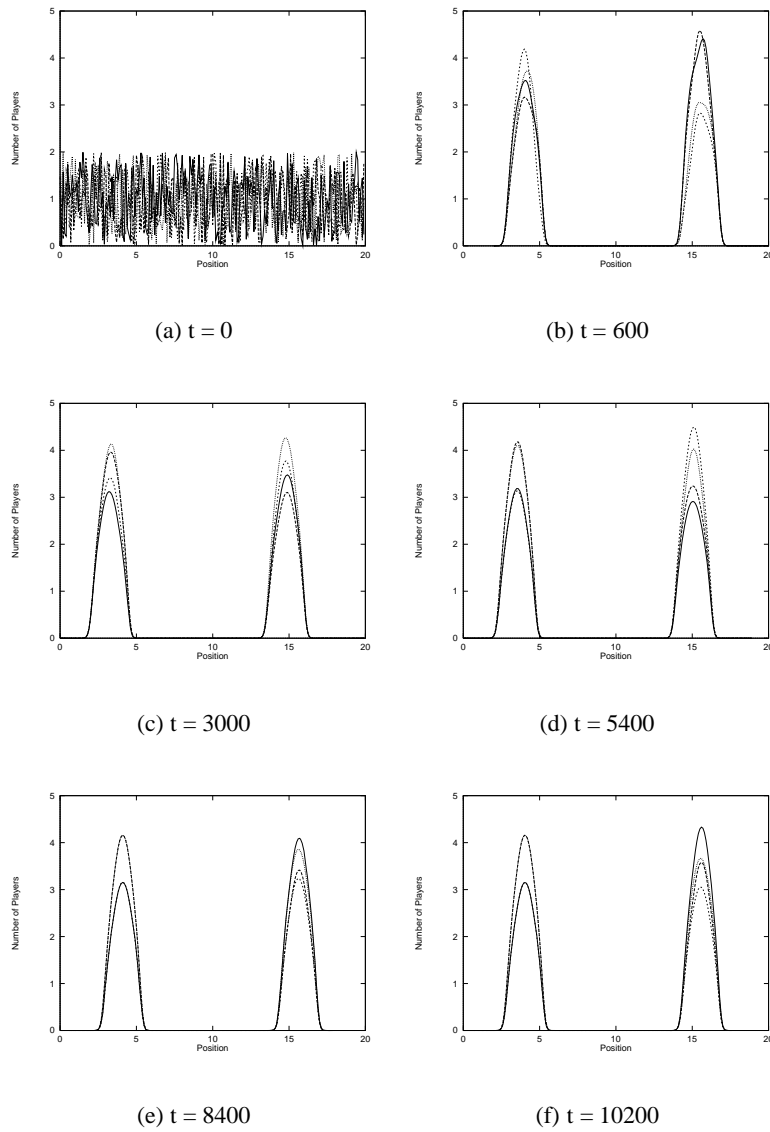


Figure 5.12: Spatial patterns for numerical resolution of (5.17), $c = 0.6$, $A_1 = 0.25$, $A_2 = 0.05$, $D' = 10^{-4}$.

It is interesting to compare these results with the non-spatial models, where it was concluded that there was either a weakly stable region or a stable attractor where strategies 1 and 3 would be present in a small quantity or even disappear. Only in the *absence* of noise a similar cyclic dynamics was also established in those cases¹⁷. The question is then whether a similar stabilization to a non-cyclic state is to be expected if noise is added to the spatial dynamics.

In order to test if this is so the equations are solved with the same conditions adding uniform noise in the density variables (mean: 0, range: 0.01 per unit of space and per unit of time). The time evolution is shown in figure 5.11 (b), where it can be seen that oscillations dominate the dynamics. (In this particular run there is only one cluster of players.) In the many independent runs that were studied, stable non-cyclic solution were not observed when perturbations were

¹⁷The exception is the case of high mutation rates, as shown in figure 5.8, where the cycle becomes stable even with the addition of noise.

added. In conclusion, oscillations seem to be the long term mode of behavior for the continuous spatial model.

An interesting feature that this model allows to study is the relative contribution to a player's net energy increase for the first and second roles, a quantity that will be important in explaining some results presented in chapter 6. For each strategy i the quantity $\Delta E_i = E_i^1 - E_i^2$ can be measured as a function of (x, t) . The resulting form depends on the composition of the cluster and the distribution of sub-populations within it. As an example, two different cases are observed in figure 5.13 where the cluster composition (and variations in environmental energy) and the corresponding variation for ΔE_1 for strategy 1 have been plotted in different runs. In a case where all strategies are present in similar quantities [(a) and (b)], ΔE_1 is always greater than zero, and therefore players of strategy 1 receive most of their payoff from games in which they engage as *first* role players. In contrast, for a cluster that is composed mainly of players of strategy 1, ΔE_1 can assume negative values near the center of the cluster, which means that individuals in that area are receiving most of their payoff for participating in games as *second* role players.

It is clear that the introduction of a spatial dimension establishes a difference in the frequencies with which an individual plays the first and second roles. In non-spatial models these were taken to be equal and now they depend on the spatial position of a player relative to the rest of the cluster¹⁸. Additionally, the spatial variation of energy also affects the payoff a player receives depending on which role it plays. For a population mainly constituted by a cooperative strategy such as 1, as in figure 5.13 (c) and (d), the payoff for each player is always one half of the energy in the food source (which is located at the same position of the first role player). According to the energy distribution shown in the figure, in a contest between a player at the center of the cluster and one at the periphery the first player will receive a greater absolute payoff if it plays the second role than if it plays the first one just because energy resources of different quality are at stake.

In conjunction, both these effects can contribute to a greater energy intake at the center of the cluster when playing the second role than when playing the first. In other, more general cases, this may not be so, but it will be important to keep in mind that this is a possible situation for the discussion presented in next chapter for the individual-based model.

5.5.4 Effect of density thresholds, move toward discreteness.

Additional numerical tests were performed in order to investigate the stability of initial distributions that favour one strategy over the others. In the most extreme case, the system was initialized with a random spatial distribution for only one of the strategies as before (uniform with range 2.0 per unit of space) and a very small uniform distribution for the other strategies (range 0.1 per unit of space). These distributions roughly mean an allocation within the spatial domain of about 20 individuals of one type and 1 individual for each of the remaining three types. The aim is to observe whether some of the strategies in small numbers will invade the population.

A series of tests were performed solving (5.16) numerically as above and using a mutation scheme as described in section 5.3. This is achieved simply by adding in each resolution step a re-distribution stage for the newborn players after the normal numerical resolution. Accordingly,

¹⁸This is true, however, only if the number of trials an agent makes in order to find a partner (N_a) is greater than one trial per unit of time. For $N_a = 1$ both frequencies are always the same.




Figure 5.13: Spatial distribution and corresponding difference in energy intake between roles for strategy 1 (ΔE_1). (a) and (b) a distribution with all strategies present, (c) and (d) a distribution where strategy 1 is predominant. Variations in energy levels are also plotted.

some of the newborn players of one strategy are mutated into the other strategies in a proportion μ for neighboring strategies and μ^2 for non-neighboring strategies, (as done in section 5.4). Tests using mutation are interesting in this context since the stability of single strategy distributions is being examined and, therefore, it is necessary to allow for the possibility of their disruption and invasion by mutants.

Figure 5.14 shows the population composition for a run in which, initially, the population is composed mainly by strategy 1. It is possible to observe that this distribution is not stable and it is invaded by sub-population 2 (which is afterwards invaded by 3, and so on. until the cyclic behavior is restored).

There is, however, something suspicious about the postulated situation, since it is not clear what is the meaning of an initial distribution where all but one, of the strategies are present in very small densities. A careful observation of figure 5.14 shows that some sub-populations are sometimes present in very small numbers (< 1) and after that follows a period where the same sub-population is dominating. Considering that players are discrete entities (although they have not been modelled as such) it could be asked whether this is a reasonable behavior for the model.




Figure 5.14: Invasion of a population initially constituted roughly by players 20 of type 1 and 1 player for each of the remaining strategies.

In a spirit similar to that of the quasi-discrete stability analysis presented in section 5.2.4, it would be interesting to ask what would be the effect of passing from a continuous to discrete analysis of spatial densities. After all, if the spatial volume of interest is of size V a density value less than V^{-1} make no sense unless it is zero.

This is hardly a matter of little consequence. Abramson and Zanette (1998b) have studied the statistics of extinction in Lotka-Volterra systems with many species introducing minimum density thresholds and finding interesting qualitative differences in the resulting dynamics with respect to purely continuous treatment such as, for instance, the transformation of unstable equilibria into stable ones¹⁹. This is more than just a theoretical or technical point. Similar conclusions have also been reached in molecular evolution experiments on the invasion rates of highly fit RNA virus variants (de la Torre & Holland, 1990). Highly competitive clones invaded their progenitor population only when seeded above a *threshold* level otherwise they were suppressed by the existing quasispecies dynamics. Similarly, the effects of discreteness in combination with stochasticity, have been recently studied in models of the evolution of altruism and kin-selection in viscous populations, (Goodnight, 1992; van Baalen & Rand, 1998). It is found that, in contrast with continuous models where an infinitesimal amount of non-altruists suffices to take over the population in the long term, treating the population as discrete shows that altruism can be stabilized due to the inter-group advantages it provides and the un-invadability of altruist groups with less than one non-altruist.

Following the mentioned mathematical models, a density cut-off value p_0 is introduced for the spatial densities of players and the invasion experiments are re-run. A reasonable upper bound for this threshold may be obtained as the inverse of the volume of a whole cluster which, in these experiments, is approximately equal to 0.4. Figure 5.15 shows the evolution of the different sub-populations for different values of p_0 where it is clear that as the threshold value increases it takes

¹⁹For other mathematical models including density thresholds see (Tsimring, Levine, & Kessler, 1996; Kepler & Perelson, 1995).

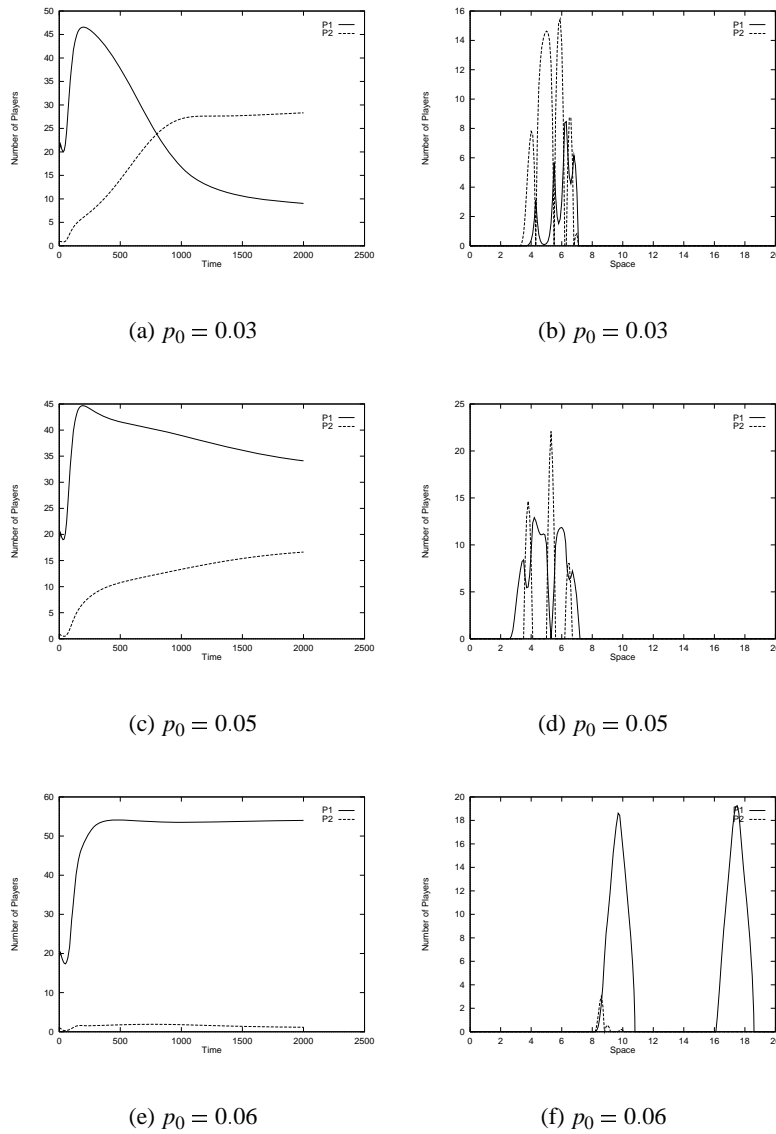


Figure 5.15: Invasion patterns for a mutation rate $\mu = 0.01$ and different values of density threshold p_0 . The resulting spatial patterns are plotted on the right for $t = 2000$.

longer for the invading population to surpass the initial one until finally for $p_0 \cong 0.055$ invasion does not occur at all.

The relevance of density thresholds points to an interesting aspect of the current model. Since the continuous approximation for a discrete variable does not hold in the evident case of very small density values in finite domains, it would be legitimate to ask in what other aspects this approximation is introducing spurious results (Durrett & Levin, 1994a, 1994b). One way to find out is to explore different approaches in which the variables are treated as discrete, but still choosing as a starting point the continuous model. A possible way of doing it is to allow only for discrete values in the numerical resolution of (5.16). A comment on how to do this was already given in section 5.3, and it consists simply in allowing only discrete increments in the densities variables and accumulating the remainder of the continuous increment for the next integration step. This is still not a discrete model, but an approximation in that direction. The partition to choose is again related to

the choice of the characteristic spatial volume where it may be safely said that one individual can be added. If the typical volume is V only density changes in units of V^{-1} should be allowed.

However, this direction of inquiry will not be pursued since the same issue will be explored more thoroughly in the following chapter where the continuous base for modelling the variables is abandoned from the start.

5.5.5 Summary for continuous spatial model

The aim of the spatial model described above has been to provide a first order approach to understanding the possible roles that a spatial distribution may have on the evolutionary dynamics. In order to be able to use the other results of the chapter, the model was based on the ecological equations presented in section 5.3, to which it was necessary to add a spatial dimension. There are different ways of doing this, some as simple as adding diffusion terms for the population density equations. However, it was considered important to include further ecological considerations regarding the way agents should be able to play with their neighbors.

Many reaction-diffusion approaches intend to model so-called contact processes, which means that interactions are considered to occur at infinitesimal distances. This approach is suitable for phenomena like chemical reactions, where the actual size of the interaction range of individual molecules is not much greater than the molecule itself. However, in the ecological case addressed here, this is not a suitable assumption. It is not reasonable to suppose that a player will only interact with a couple of nearest neighbors in its whole lifetime. It is because of this that a model with finite locality was used instead.

In this model each player is surrounded by a neighborhood of finite size where all its interactions will take place. This neighborhood is intended to represent the player's local habitat. There are, in principle, a variable number of neighbors for each player. It is this sort of long-range coupling that provides one of the key factors in understanding the model capability for pattern formation. This is so, because the actual frequency with which a player plays the game is dependent on the probability of finding a partner, which in turn depends on the number of available players in the neighborhood.

Consequently, the frequency for initiating a game as a first role player is proportional to a fixed frequency for all players times the chance of finding a local partner after a pre-specified number of trials.

Instead of working directly with the number of players, the main variables in the model are the continuous spatial densities of players according to their strategy. Assuming that energy is also locally and continuously distributed, it becomes clear how to calculate the average energy obtained by an individual per unit of time. Such energy has two main contributions: the energy the individual gains as a first role player and the energy gained as a second role player. It is assumed, according to the original metaphor, that the energy used as a payoff for the game is extracted from the location of the first role or initiating player. This makes the calculation of the energy gain as a first role player quite straightforward; it is simply the product of the frequency with which that individual plays the first role times the local energy times the average payoff obtained. This latter factor is calculated according to the current average of strategies surrounding the individual in the neighborhood. The calculation of the energy obtained as a second role player

is a bit more complicated. It involves calculating for all the other players in the neighborhood, the proportion of times in which they select the original individual as a second role participant, and then integrating this quantity over all the positions in the neighborhood. Once these energy contributions are calculated, it is easy to write the equations for the evolution of the system.

The system is then tested for its capability to form spatial patterns, that is, non-homogeneous long-term solutions. This is done using a special form of linear stability analysis which consists on finding homogeneous solutions and then testing how stable they are under non-homogeneous perturbations. If some non-homogeneous variations to the perturbed solution tend to grow, then the dynamics is expected to proceed to the formation of a pattern in space. What sort of pattern cannot be actually determined by the analysis because the analysis is concerned with the linear limit. As soon as some spatial modes begin to grow the linear limit is lost and non-linear effects begin to take place which in the end will determine the long term patterns.

As an analytical proof of existence for pattern formation, it is possible to apply this method only in a particular case described by the model. After doing the corresponding simplifications and defining parameters it is found that the models will produce spatial patterns under many circumstances. The actual “driving force” behind pattern formation is the coupling term in the equations which links an individual with its neighbours and consequently the capability for pattern formation depends on the neighborhood size. Tentatively, it can be said that players with many neighbors will tend to gain more energy partly because they are involved in more games than players which are surrounded only by few neighbors. In the first case the local density will tend to increase providing positive feedback, while in the other case it will be more difficult for new players to be born and therefore isolation will tend to increase. Such effect is clearly manifested when the actual model is explored numerically.

By solving the equations, after specifying a set of boundary conditions and random initial distributions in space for all the strategies, it is found that the population tends, in the long-term, to be arranged into clusters which are quite stable in shape and separated from one another. Within these clusters, though, the composition of the population may change in an oscillatory fashion not unlike the original oscillations of the non-spatial model. Alternatively, stable compositions may also exist within a cluster, but once noise is introduced in the system the cyclic solution becomes the only one observed. An interesting point is that once the clusters are separated the corresponding dynamics are independent.

The internal structure of the cluster introduces asymmetries between two given positions as well as asymmetries between the two roles at a given position. This can easily be seen by studying the differences in the energy contribution between the two roles. This difference varies from one position to the other and may also change signs under certain circumstances indicating that in a same cluster there are players who receive more energy for playing a different role than the role played by other agents in the cluster. Such asymmetries will become a key explanatory element for the model presented in next chapter.

Another such key factor has also been preliminarily studied here. This is the case of the role of discreteness. By performing simple invasion tests, in which one population is initialized in large numbers and the others in very small ones. Adding a mutation stage to the resolution of the system, the stability of single strategies is thus studied. It is not surprising to see that the initial

strategy is invaded by another and after that the cycle is initiated. However, according to both theoretical and empirical studies, it is not reasonable to allow for the density of players to reach very small values because it becomes a meaningless variable. After all this density intends to model the spatial distribution of discrete individuals. Because of this, a minimal density threshold is introduced. This threshold indicates the density value below which all other values are so small that they should be considered equal to zero. Interestingly, as the value of the threshold is increased, invasions become more difficult, until a point in which they do not occur and the single strategy becomes a stable solution.

This result points to discreteness as a factor which may radically change the dynamics of a model. If the inclusion of a minimum density threshold is enough to introduce important changes, one may ask what other influences would become apparent if the model were further pushed toward a discrete situation. This question is investigated in the following chapter.

5.6 Summary: broken symmetries

The motivation behind this chapter has been to define an evolutionary scenario where the evolution of a very simple game of action coordination could be studied using a traditional game-theoretic approach and some extensions to it.

Once the game was defined, it was found that there were no ESS's when the investigation proceeded using traditional game-theoretic analysis. However, after extending this basic stability analysis to include a model of evolutionary dynamics, it was found that there was a weakly stable region of the strategy space where the population was expected to be found. This region (not yet a single point) was defined by specifying that strategies 2 and 4 should be present in equal quantities and in greater proportion than strategies 1 and 3. This can be interpreted as saying that a combination of non-coordinating strategies is an expected solution. This result has been confirmed both analytically and with numerical resolution of the equations. Oscillations in strategy space have also been observed. However, these are suppressed with the addition of noise.

By allowing population size to be finite, it becomes clear that only one point in the weakly stable region is in fact stable. This attractor indicates that a population constituted by equal amounts of players of the non-coordinating strategies 2 and 4 will not be invaded.

The addition of an ecological setting which allows the assumption of optimization which is implicit in the equations to be dropped, shows no major difference in results. Under perturbations, the result is still a single stable point in which strategies 2 and 4 were dominating in equal proportion. A global tendency towards optimizing individual fitness is dynamically recovered due to the existence of a global energy variable with which all players are coupled. Distributing the strategies in a haploid genetic space shows that high mutation rates can have a role in stabilizing the oscillatory solution.

As the depth of the analysis is increased, the sort of factors that can be expected to have important effects are highlighted. An interesting tendency is already beginning to show. This tendency indicates that important variations in the results should be expected as the inbuilt symmetries of the game are gradually broken. The first broken symmetry is the fact that in the infinite population case, all players "perceive" the same field, which is not the case for finite sizes where the perceived field depends on the strategy played although, clearly, as the population size increases

this dependency becomes less important.

More evidence of the importance of broken symmetries can be obtained from the continuous spatial model. The most obvious evidence is given by the formation of discrete spatial patterns of characteristic size which achieve a steady evolutionary state. These clusters divide the population into independent sub-groups. Within a cluster the dynamics remains oscillatory if random noise is added, but evolution may reach a stable state if this noise is absent. Clusters present a definite spatial structure manifested in the shape of the population and energy distributions. Such a structure introduces differences between any two given positions within the cluster in terms of the frequencies for playing each role and in terms of the local amount of resources. Some examples were observed in figure 5.13 where the difference between in energy gained according to role played shows a dependence with position and may even change signs in certain circumstances. A change of sign indicates that some players may receive more energy per unit of time because of their involvement as second role players, something which is not possible in the non-spatial game.

Moving from a continuous model toward a quasi-discrete one introduces new factors in the dynamics especially when studying invasions of an initially dominant strategy by mutants. Inspired both by recent theoretical models and empirical evidence, minimum density thresholds were introduced below which all density values were considered to be zero. An increase in the value of this threshold (well within reasonable limits) results in invasions not occurring and, consequently, in a single strategy constituting a stable solution.

As progress was made through the chapter, a series of relevant factors that may introduce important differences in the expected results has been identified. However, there is still a sense of the investigation not being finished. The issues of discreteness and space have merely been touched upon. What would happen if the model were made discrete from the start? Would the spatial patterns also exist? In order to investigate these questions a move will be made from the realm of models that can be formulated mathematically into an individual-based model directly instantiated in computer simulations. Some of the results of this computer model will be radically different from what has been shown so far. Understanding the reasons for this is not straightforward. However, the different influences identified for the various factors studied so far will provide crucial help in formulating explanations of what happens when many of these factors are combined in the individual-based model.

Chapter 6

A game of action coordination: individual-based model.

As the importance of the simplifying assumptions of each model is assessed by building new ones that do not include them, arriving at a purely mathematical formulation becomes increasingly difficult. An exploration has been made of the effects of including dynamics into a purely game-theoretic analysis followed by an investigation of the consequences of not assuming infinite populations and of situating evolution ecologically and spatially in cases where density-dependent effects and pattern formation in space could appear. It has been shown that, as a consequence of some of these factors, general results may be altered. Yet there is still reason to suspect that some of the features which have not been so thoroughly investigated may also fulfill crucial roles. After bringing to light the fact that populations are composed by discrete entities subject to random fluctuations, it seems plausible that this fact may be one such crucial feature. From what has been shown in chapter 5 (especially sections 5.2.4 and 5.5.4) it seems that the effects of discreteness on ecological and evolutionary stability may not be trivial. Another unexplored and related issue concerns the nature of reproduction, not as an instantaneous, population level phenomenon (where one may collect all the energy accumulated in a certain period and estimate the number of new players) but as attained by individuals who have to accumulate energy and find an opportunity to mate, all of which involves stochastic time delays.

In order to be able to explore these issues without resorting to complex and cumbersome mathematics, a choice was made to formulate and analyse an individual-based model implemented in a computer simulation. The potential of this type of models has been discussed in chapter 4 along with some methodological issues concerning their use. On the one hand, it is often found that it is relatively easy to attain different degrees of complexity in this kind of models and therefore to investigate factors that are necessarily over-simplified in mathematical models but, on the other hand, without the previous theoretical insight provided by those less complex models, it is actually very hard to interpret the relevance of the many dimensions in which an individual-based simulation can be explored. It is for this reason that an effort has been made to explore different analytical routes *before* embarking in a complex computational approach.

6.1 The model

In this model a population of agents play the same game as described in chapter 5 in a shared spatial environment represented as a toroidal arena of size X_{max} by Y_{max} (typically 100 x 100). As in the continuous spatial model of section 5.5, agents are situated in a position which remains unchanged during their lifetime. The whole spatial environment is partitioned as a regular grid (typically of 10^4 positions) and each agent occupies one of these positions. At the beginning of the simulation agents are distributed randomly. The normal spatial habitat of an agent during its lifetime is represented by a neighborhood around its position. This neighborhood is simply defined as a square of a fixed size (usually 10 x 10) centered around an agent's coordinates (see figure 6.1)¹. Agents can interact with other agents and access food sources only within the limits of their neighborhoods. A number, F_Q , of energy repositories, or food sources, are also distributed within the environment, (in general, $F_Q = X_{max} \cdot Y_{max}$ has been used). Food sources can be uniformly or randomly spaced and this has proven not to be relevant to the results. These sources are created at the beginning of the simulation run and their position and associated food type remain unchanged. A food source cannot be created or destroyed during the rest of the simulation, but its energy content changes due to consumption by the agents and stochastic renewal. Food types are denoted as "F0", "F1", etc.

One of the reasons for modelling spatial relations using a finite locality, apart from those discussed in the previous chapter (section 5.5), is to avoid possible spurious effects caused by the spatial regularity imposed in lattice-like models. Most models of this kind are intended to study local contact processes (Durrett & Levin, 1994a) and rely on exclusion principles to simulate ecological density regulation. The space is regularly subdivided into cells and each cell can contain just one particle or agent which is able to interact with its nearest neighbors and move (or send offspring or propagules) only to nearby empty cells. As a result of this scheme unlimited overcrowding is prevented. However, it is necessary to be aware of the sort of spatial regularities imposed by this scheme. This is a situation analogous to the question of synchronous updating in the temporal domain. It has been shown that interesting results in this sort of models can sometimes be due to the spurious effects occasioned by updating all cells at the same time². Might not the grid-like regularity of spatial relations be an ordering factor which provokes similar spurious effects? Studies on the matter are lacking at the moment. But, to be on the safe side, it would be desirable for the present model to depend as little as possible on such choices. To achieve this, an ecological approach has been chosen instead of a cell exclusion principle in order to naturally include features like density regulation as demonstrated in the continuous spatial model in the previous chapter. Overcrowding is avoided through coupling with a space dependent energy variable which is renewed at a fixed finite rate. At the same time, even though spatial position is grid-like in this particular model, it is almost certain that spurious structuring effects are overcome by the fluctuations within an agent's neighborhood caused by random choice of partners and asynchronous updating. More than one player may occupy the same position since in this model these

¹In the model presented in the chapter 7 experiments were made with circular neighborhoods and with continuous spatial coordinates without any significant changes in results.

²See for instance (Nowak & May, 1992), criticisms by Huberman & Glance (1993) and a response (May, Bonhoeffer, & Nowak, 1995). On the question of asynchronous updating and its effects see for instance (Ingerson & Buvel, 1984; Bersini & Detours, 1994; Harvey & Bossomaier, 1997; Abramson & Zanette, 1998a).




Figure 6.1: An example of agents' positions and their square neighborhoods. "F" represents a food source within reach of agent "A1" who has selected "A2" as a partner in the game. Notice that although both "A2" and "A3" are "within reach" of "A1", they cannot interact with each other.

coordinates do not represent the *actual* physical position of the agent, but the center of a region where the agent lives.

Each agent possesses an internal energy counter that specifies the energy available for self-maintenance and reproduction. Every time an agent is born its energy level is set at a value chosen following a Gaussian distribution $G(e, \sigma_e)$ centered around e . (Typical values are $e = 150$, $\sigma_e = 10$, negative values cut off.) Energy flow is strictly accounted for. The whole environment is subject to a constant average energy flux R (500 per time step) which is manifested in a stochastic increase of the energy stored in the F_Q food sources. Agents derive their initial energy from their parents and, during their lifetime, from the food sources. This ensures energy conservation. Agents also spend a fixed amount of energy Lo (0.25 per time step) every time they are selected to act (independently of the outcome of the action) which does not return to the environmental resources but is dissipated, so as to make sure that an equivalent to the second law of thermodynamics is also observed.

At each time step a number of agents equal to the size of the current population are randomly selected to try to find a partner and perform the first role in the game. As mentioned before, the updating is performed asynchronously and, on average, each agent will always be selected once per time step. A second player and a food source, both within the neighborhood of the first player, are randomly selected. If no agent is found after a finite number of attempts N_a (typically 10), the first player loses its chance to play the game and the energy cost is discounted anyway. Each time an agent is selected its age is incremented by one unit regardless of the role it plays.

Payoffs are assigned according to the rules of the game as described in section 5.1 and con-

sidering the fact that their absolute value will depend on the energy contained in the selected food source. When an agent has accumulated enough net energy it is able to reproduce. The corresponding threshold level of energy is selected at each time step according to the distribution $G(2e, \sigma_e)$, i.e. a similar Gaussian distribution to that used in the assignment of initial energy values but this time centered at twice the initial energy value e . After each game, the energy level of each agent is compared with a number extracted from this distribution. If the energy of either one of them exceeds the selected value, then that agent initiates a reproduction event. The use of a stochastic threshold for reproduction is intended to introduce noise in the dynamics. Reproduction is sexual and selection of mate is based only on the locality condition. This has been implemented by choosing the last co-participant in a game as the second parent. The offspring occupies a randomly chosen position within the first parent's neighborhood. As neighborhoods do not overlap completely in the general case, gene flow is not totally prevented. The initial energy of the offspring is supplied by the parent whose high energy level triggered the reproduction event in the first place. This energy value is again selected stochastically using a Gaussian distribution centered around e .

Those agents unable to sustain their costs will eventually run out of energy and die. Costs are the same for all agents, with the exception of very "old" ones. After a certain age has been reached (typical values range between 1000 and 4000), costs are mildly incremented linearly with age. This is the only "developmental" feature included in this model. This is done in order to avoid the presence of "immortals" which could render the results on evolutionary dynamics questionable for not allowing enough variation in the population constitution. This mechanism replaces the lack of a good model of developmental rules with which it could be expected that typical lifespan would result from an interplay between population dynamics, availability of resources and evolution of those rules.

The structure of agents is that of a state-less machine. According to the role, actions can depend on the perceived food type or on the perceived external manifestation of the partner in the game. These actions may be represented in a look-up table that defines an agent's behavior (an example is shown in table 6.1)³.

1st role			2nd role	
Food Type	Action	Signal	Signal	Action
"F0"	B	γ	α	D
"F1"	C	β	β	C
\vdots	\vdots	\vdots	\vdots	\vdots
"Fn"	B	α	v	A

Table 6.1: Example of a look-up table that determines the behavior of an agent. Capital letters represent effective components and Greek letters external manifestations of actions (signals).

The behavioral look-up table is encoded in a haploid genome represented by a binary bit-string.

³Notice that it is possible to drop the simplifying requirement that the signals used by a player must belong to a definite subset of all the possible signals, see section 5.2.1. While this may introduce additional complications since agents would be able to play a restricted number of mixed strategies, the general results were not different from those cases where the restriction was kept.

Typically, if four food types, four actions and four signals are possible, as in the experiments described in the following section, then each action and signal is represented by two bits and the total size of the (active⁴) genotype is 12. Newborn agents receive their genome as the result of a uniform crossover operation on their parents genotypes plus certain probability of mutation μ per locus. The result of uniform crossover is a genotype in which each place in the bit-string is selected from the corresponding place from one of the parental genotypes with equal probability (see Syswerda, 1989; Mitchell, 1996, pp. 58 and 172). A typical value for the mutation rate was $\mu = 0.005$ per locus in the genotype and per reproduction event. This guarantees that, after crossover, the probability of a genotype to remain unchanged after mutation is: $(1 - \mu)^{12} \cong 0.942$.

6.2 Results

Results discussed in this section were obtained using two sets of simulations in which only four different types of food were included in the environment. In one set of simulations each food type had a different associated action and in the other a given action corresponded to two different food types and, therefore, only two actions were relevant even if agents were able to perform more actions, (for example, action “A” may correspond to food types “F0” and “F1” and “B” to “F2” and “F3”, while actions “C” and “D” may not be relevant). Only results on this last set will be described since it approaches the cases studied with the mathematical models. Simulations with more food types and actions were also carried out and the results were similar, though much more expensive to obtain and analyse.

In contrast with the mathematical models, agents are initialized totally at random. In particular, this means that they will also have to evolve an association between food types and the effective components of actions when they play that first role. Previously this was assumed to be always correct. The player that “found” the piece of food was supposed to know what to do with it. The evolution of such a task is consistently and very rapidly achieved in all the simulations, and its explanation in selective terms does not present any problems given the static character of the association involved.

Population size is variable. In all simulations the average population size stabilizes, after a transient period, at a value between 600 and 800 agents.

A baseline case was run in order to understand what kind of activity emerged when agents were not able to perceive each other reliably. In this run, the external manifestation of the first player was replaced by a random signal when perceived by the second player. The instantaneous proportional average of cases in which both players acted correctly simultaneously is called the *coordination ratio*. This ratio stabilized at 50 % for random signals. Since no food type is predominant, the best guess a second player can make, given that there is no correlation between the “signal” it perceives and the particular type of food in question, is to perform any of the two relevant actions out of the four possible ones. The achievement of this level of coordination is independent of the parameter c . This means that when signals are not random any success in coordination of behaviors will be manifested as a greater coordination ratio than that observed in the baseline case.

⁴The look-up table 6.1 is in fact implemented as a double entry matrix where for each food type and perceived signal corresponds an action/signal pair. This matrix is not used entirely. The reason for this implementation was to facilitate the experiments presented in chapter 8 in which there is a *sequence* of actions and responses and the whole matrix is used.




Figure 6.2: Proportional success in action for the first player (top line, R_1), second player (R_2), and both simultaneously (R_{12}) in a typical simulation run with $c = 0.55$.

Figure 6.2 shows the evolution of the success ratios for the whole population in a typical run with $c = 0.55$, (for a same set of parameters variations between different runs were not significant). These quantities should not be confused with any measure of fitness of the population. R_1 indicates the average proportion of successful actions for first players which, as predicted, climbs rapidly up to a level near 100 %, (perfect success is not necessarily achieved, due to the effects of random mutations). R_2 shows the same quantity for second players and R_{12} for both roles simultaneously (the coordination ratio). Although it is not very clear in the figure, R_1 and R_2 start near a value of 25 % and R_{12} near 12.5 % which corresponds to a randomly initialized population. Note that the proportion of coordinated games reaches a higher than baseline level of about 90 % and is maintained in a relatively steady state.

By observing the resulting behaviors once a stable state has been achieved it is possible to determine that signalling behaviors present a strong correlation with actions and not with food types. Therefore agents evolve a “two-signal code” for the two relevant actions that can be performed. However, as shown later, such a code is not at all universal within the population.

A study of the effect of parameter c on R_{12} was performed in order to see how the extra payoff against coordination affected the level of simultaneous success. The intuitive expectation is that this level should decrease as c increases from 0.5 to 1.0 because coordination becomes more costly as the degree of conflict is augmented. Figure 6.3 shows this effect on R_{12} for different values of c . It is apparent that the proportion of coordinated activity for the whole population decreases as c is incremented from 0.5 to 0.65 but still remaining above the baseline case. The level of coordination goes from almost total coordination for $c = 0.5$ (neutral cost) to coordination by guessing when $c \geq 0.65$. This result is quite different from the results obtained with the game-theoretical approach and the extended mathematical models and it requires an explanation.




Figure 6.3: Coordination level for different values of c . (b) shows the average value for the coordination level for different values of c . Each point is the average of 5 simulation runs. The line represent a linear regression, (linear correlation coefficient: -0.982). (a) shows some examples of the time evolution of R_{12} for some values of c .

6.3 Spatial patterns

It was shown in section 5.5 that in a spatially extended model, if the coupling between the density of players and the available environmental energy depended positively on a term reflecting the density of players in a finite volume, spatially heterogeneous solutions may exist. In particular, these solutions tend to be aggregated “lumps” or clusters of players separated by uninhabited regions. In this section a study is made of the spatial distribution of agents as a function of time to see whether this is also happening in these simulations.

Figure 6.4 shows the first stages in the spatial distribution of the agents. Initially these are distributed randomly across the whole environment, (border effects have been avoided by the use of periodic boundary conditions). Since agents do not move, any change in the spatial distribution is due only to birth and death processes. One can observe how the initial symmetry is rapidly broken and how agents show a tendency to aggregate into clusters. As in section 5.5, symmetry breaking is caused by minor differences in the distribution of initial positions and also by the nature of the updating rules. Some agents will reproduce more than others just because they have a few more neighbors to interact with or have been called to act a few more times and, therefore, they have a slightly better chance of accumulating enough energy for reproduction. As reproduction is also a local process, the effect is self-reinforcing. Relatively isolated agents will have less chance of interaction and will tend to die sooner also contributing positively to increase the isolation of agents in similar positions.

Clusters remain relatively stable, at least if compared with the time scales of an agent’s life-time. Although they rarely break or fuse, they have a tendency to move slightly in random directions⁵. This was also seen in the model presented in section 5.5 and, although it was not done explicitly then, it is possible to use that model to support the claim that clusters do not change significantly in their general form. Here, a more “physical” description of why this is so will be

⁵Or probably following slight gradients in environmental energy.







(a) $t = 0$ (b) $t = 2000$ (c) $t = 4000$ (d) $t = 8000$ (e) $t = 16000$ (f) $t = 28000$

Figure 6.4: Formation of clusters. Each figure shows the distribution of agents in the toroidal arena at different times during the initial transient.




Figure 6.5: Qualitative energy and population distribution inside and around a cluster (circle) in one spatial dimension.

given.

Tendency to expand. Figure 6.5 shows qualitatively the spatial distribution of energy and density of agents inside and in the vicinity of a cluster (compare with figure 5.12). It is easy to see that resources will be more frequently used in the more populated areas toward the center of the cluster than on peripheral areas in which the population is sparser so that the amount of available energy will, on average, decrease towards the center as shown in the figure. Agents living in the periphery will have access to resources of greater quality and the average energy gained per game played will be greater than that of agents in the center region. So, at first glance, peripheral agents would seem to be better off and have a better chance of generating more offspring; therefore the cluster would experience a tendency to increase the peripheral population and consequently a tendency toward *expansion*.

Tendency to aggregate. In contrast, while all agents have the same probability of being *chosen* to play the first role in the game, the frequency with which an agent *actually* plays the game depends on the surrounding density of players. The more densely populated the area, the higher is the chance of finding a partner in that area with which to play. For this reason, agents living near the center of a cluster will play the game more frequently. So, in principle, they will stand a better chance of receiving (by coordination or by guessing) more energy per unit of time. This provides agents in the cluster with a tendency to *aggregate*.

Equilibrium The equilibrium of both these tendencies determines the size of the quasi-stable cluster. At a certain size the extra energy gained by the peripheral sub-population will match the energy lost by their being more sparsely distributed than the center sub-population and similarly the extra gain due to higher frequency of games in the center region will be compensated by a smaller payoff per game. Resulting clusters have been observed to have a typical area of 2 up to 8 neighborhood sizes with populations ranging from 30 to 200 agents.




Figure 6.6: Examples of the correlation between coordination and cluster size (estimated as an average diameter) for different values of c . Each point represents the state of a cluster in a typical simulation run after 3×10^5 time steps. Lines depict linear regressions. In these examples the coefficients of linear correlation are 0.977 for $c = 0.52$, 0.9998 for $c = 0.55$ and 0.882 for $c = 0.65$.

6.4 Correlation between cluster size and coordination level

The above description can provide an idea of the tendencies that give origin to clusters as stable spatial patterns. However, this explanation may give a wrong impression and lead the reader to think that cluster size, as determined by the interplay between opposing “forces”, depends exclusively on the values given to the simulation parameters, and that any variance in the resulting sizes is attributable only to the random contingencies of the particular run. If this were the case, it would not be reasonable to expect any correlation between cluster size and any other variable in the simulation for a given set of parameters. However, a consistent correlation has been observed between the level of coordination achieved within a cluster and its size and number of agents.

In a given simulation run the average level of cooperation (i.e. the proportion of games that result in coordination) for each cluster admits certain variation from one cluster to the other. This is due to the fact that each cluster evolves for long periods in a more or less independent manner (which is also the reason for signalling “codes” used by the agents being different between clusters). By looking for geometric characteristics that will help to differentiate clusters according to their level of cooperation it was found that larger clusters tend to be associated with higher levels. This correlation can be observed in figure 6.6, where the coordination level within a cluster is plotted as a function of its size for three simulation runs using different values of c (error bars represent the standard deviation).

How can this correlation be explained? It is clear that a particular cluster being large or small can depend only on the initial conditions and the random fluctuations suffered by it. However, there seems to be a tendency in the dynamics toward establishing some order in the end distribution with respect to the level of coordination in each particular cluster. The implication of this is that any explanation for this correlation must provide a mechanism that links the corresponding variables, either causally to one another or operationally to the rest of the process. Such an

explanation will be provided in section 6.6.

First, it will be important to consider a plausible candidate for explaining the stabilization of coordination in the first place and see if it is able to provide an explanation for the size/coordination correlation at the same time. Such a candidate is kin selection.

6.5 The role of kin selection

The importance of a spatial dimension in evolutionary games with conflict of interest, whereby organisms are constrained both in their reproductive patterns and possible choice of interaction partners, has recently been put into focus. In particular, it is often found in these cases that cooperative interactions are stabilized under circumstances where simpler selective arguments would predict the contrary⁶. This has been demonstrated both by mathematical considerations and computer models in the case of the Prisoner's Dilemma (Axelrod, 1984; Oliphant, 1994; Nakamaru, Matsuda, & Iwasa, 1997) and communication games (Krakauer & Pagel, 1995; Ackley & Littman, 1994; Oliphant, 1996). Depending on the particular features of the model there may be more than one possible explanation for this phenomenon. For instance, Axelrod argues that spatial clustering favours reciprocity in the case of the TIT-FOR-TAT strategy (Axelrod, 1984, p. 65 - 69). A different explanation is favoured, in their respective models, by Ackley and Littman (1994), Nettle and Dunbar (1997) and also by Oliphant (1996). Such explanation involves the concept of kin selection, (Hamilton, 1964; Hines & Maynard-Smith, 1979; Michod, 1982; Grafen, 1991). It is argued that, since the spatial regions where mating partners are chosen from and offspring allocated into tend to coincide with areas where game co-participants are selected from, this ensures that players will be highly related so that a cooperative player will tend to increase, on average, the frequency of genes identical to its own in other players, i.e. their inclusive fitness.

If fitness is related to the degree to which an organism can pass copies of its own genes to the next generation, inclusive fitness would be a generalization of this concept so that it can also account for traits that help the transmission of copies of the same genes that happen to be located in other individuals. As the well known example goes, a behavior that is detrimental for an individual but which produces benefits of a same magnitude for more than two siblings will tend, under certain circumstances, to increase the individual's inclusive fitness and therefore it will be selected for. This is so because the probability of finding the same given gene in a sibling is, in diploid genetic systems, equal to 1/2. Thus, the general condition for a trait or behavior to be selected for is not that it has a positive fitness effect for an individual W_{ind}^i but that it has a positive inclusive fitness effect W_{Inc}^i :

$$W_{Inc}^i = W_{ind}^i + \sum_j r_{ij} W_{ind}^j > 0 \quad (6.1)$$

where r_{ij} is the degree of relatedness with another individual j (defined more carefully below) whose own individual fitness is affected by a quantity W_{ind}^j by the actions of i . The sum is extended over all the other individuals who may be affected by the trait or behavior in question. Even in

⁶And, more generally, a growing body of ecological literature shows that spatial extension introduces non trivial results which sometimes are unexpected when compared with non-spatial models. See section 5.5 and literature cited in the references therein.

cases where $W_{ind}^i < 0$ for individual i , its inclusive fitness can still be positive, as long as the remaining term in the sum is large enough.

Here it is clear why space can play a relevant role. If interaction and reproduction are local processes, i.e. if co-participants as well as reproductive mates are selected from an individual's vicinity and offspring are allocated in the same vicinity, then the average value for the r_{ij} 's can be expected to be high. Therefore, cooperative behaviors at the expense of one individual but which benefit others in the local vicinity could at the same time tend to be beneficial to individuals that are highly related; a situation which may result in $W_{inc}^i > 0$.

Such is indeed the kind of argument advanced by Ackley and Littman; Oliphant; Nettle and Dunbar, and others. Unfortunately, none of these authors actually *shows* that this is the case in their models. They just content themselves with finding a good candidate explanation without verifying if the condition (6.1) is fulfilled, despite the fact that such a calculation should be easy enough to perform in the computer models involved.

In the case presented in this chapter, kin selection arguments are not straightforward in the sense that the actions and responses of the players cannot be said to be intrinsically cooperative or non-cooperative, but they depend on the context of the strategies used by the rest of the players. However, under the assumption of weak selection pressure it is possible to postulate a situation of quasi-equilibrium in which the context is fixed and then actions and responses could be seen as cooperative in themselves. This is by no means the general situation in the present model but this assumption will be maintained in order to see that even if this situation were true, kin selection would not be enough to explain the obtained results.

In order to test the plausibility of kin selection as a valid explanation for the evolution of cooperative coordination in the present model, a calculation is performed of the degree of relatedness between individuals and their average partners in the game. Genetic similarity can be caused by descent but also by other factors such as convergence, founder effects, etc. However, relatedness does not intend to measure just genetic similarity, otherwise all sorts of intra-specific conflicts and competition would be unexplainable. In order to take into account this subtlety, relatedness is estimated following Grafen (1991) as the degree of genetic similarity between two individuals over and above the average similarity within the population in which the individual interacts. In this way, if the difference in genetic constitution between two individuals is zero then their relatedness is equal to one, which means that from the point of view of gene frequencies for an individual to help the other is the same as to help itself. If the difference between the genetic constitution between the players is the same as the difference between one of them and the average genetic constitution in the population, then for that individual relatedness is zero, since in cooperating with the other player, it is not contributing to an increase in the frequency of genes similar to its own⁷.

A calculation of relatedness in the current model is straightforward by taking the perspective of the individual player and keeping track of the average relatedness with the partners it encounters throughout its lifetime. For each game that is played the Hamming distance⁸ d_{ij} between the binary genotypes of the participants is calculated as well as the distance between each genotype

⁷Relatedness is not clearly defined when all players share the exactly the same genetic constitution, but this is very rare case, practically impossible in the presence of mutations.

⁸The Hamming distance is defined as the number of different bits between two binary strings of the same length.

and the population average genetic constitution⁹, d_i^{avg} and d_j^{avg} . And then, for each individual, if the distance between its own genotype and its partner's is less than the distance between its own genotype and the population average, then the relatedness is simply:

$$r_{ij} = 1 - \frac{d_{ij}}{d_i^{avg}}, \quad (6.2)$$

otherwise $r_{ij} = 0$. Notice that while $d_{ij} = d_{ji}$, in general $r_{ij} \neq r_{ji}$ since for each individual relatedness is defined with respect to its own distance to the population average constitution.

In the present model, if the mean genetic constitution is calculated using the *whole* population instead of just the population which is bound to be affected by the actions of a given individual, the resulting measure of relatedness will be a generous one in the sense that it will yield greater values than, for instance, if the mean genetic constitution is calculated for each cluster. It must be remembered that clusters can have a fairly independent evolutionary history, as was shown for the case of the mathematical model in figure 5.12. To be fair, the degree of “extra” similarity that relatedness intends to measure should be in terms of the mean population with which the player or its offspring has a chance to interact, and so an average over the cluster for the mean genetic constitution would yield a more accurate estimation of relatedness. Since this average is hard to perform computationally due to the costs implied in defining the boundaries of a cluster algorithmically, a more conservative estimation is added to the above measure of relatedness for the purpose of comparison. Whereas the average constitution used for the first measure involves the whole population, in the new estimation the average genetic constitution used for the calculation of relatedness is taken as the one “seen” by each player during its lifetime. For each agent, the genetic constitution of all its partners in the game is averaged. Now, this measure of relatedness underestimates slightly the real quantity since given the typical cluster size, no agent will play with all the other agents in the cluster, although it can be expected to play with a good representative portion of them. The first estimation (using the whole population) is called “relatedness-1” and the second (using the constitution as seen by each player) “relatedness-2”. It is assumed that the measure of interest lies somewhere between the two and probably closer to the second one.

Figure 6.7 shows the variations in the two estimations of relatedness for different values of c , each point obtained by averaging the temporal mean after transients of 5 simulation runs with identical parameters. It is observed, as expected, that relatedness-1 is greater than relatedness-2, but also that neither of them varies much with c . This behavior can be understood considering that relatedness is given mainly by the spatial relations between players of different generations which should not be affected too much by c .

By calculating the corresponding inclusive fitness per unit of energy according to (6.1) it is found that for a first player (i) the payoff for no cooperation is $W_{Inc}^{NC,i} = c$, and the payoff for cooperation is $W_{Inc}^{C,i} = (1 + r_{ij})/2$. For a second player (j) the corresponding payoffs are $W_{Inc}^{NC,j} =$

⁹The average genetic constitution is simply a string calculated as the sum of all the genotypes in the population and divided by the number of agents. Since there are only two possible alleles per locus which are represented by 0's and 1's, the result of this calculation is that every component in the average string gives an indication of the proportion of each allele in that particular locus. In order to calculate the distance from an individual genotype to this real-valued string, a measure was used which reduces to the Hamming distance when applied between two binary strings. This measure involves summing, for each component of the string, the absolute distance between the individual allele and the corresponding population average. For instance, the distance between the string (1,0,1,0) and the string (0.7,0.3,0.9,0.9) is $d = 0.3 + 0.3 + 0.1 + 0.9 = 1.6$.




Figure 6.7: Average estimations of relatedness vs. c . Lines depict linear regressions.

$r_{ji}c$ and $W_{Inc}^{C,j} = (1 + r_{ji})/2$. Taking the measure that most favours an explanation in terms of kin selection (relatedness-1), it is found that in all cases the relatedness factor is about 0.4, which means that if $c < 0.7$ cooperation should be the favoured outcome. This conclusion is reached by asking for the conditions in which $W_{Inc}^{NC} < W_{Inc}^C$ is true for both roles. Alternatively, taking relatedness-2 as the estimation ($r \cong 0.11$) the same can be concluded for $c < 0.55$, which does not seem to be enough to explain the results. Since it is expected that this is a better estimation of relatedness it can be concluded that kin selection is insufficient for explaining the results.

But even so, if the assumptions regarding weak selection pressures are stretched and a generous estimation is used for calculating relatedness, then kin selection would *seem* to suffice for explaining the evolution or cooperative coordination in the present model. However, the expected result, if such were the case, would be a *constant* high level of cooperation for any value of c between 0.5 and 0.7, and a step-wise change to no cooperation for $c > 0.7$ which is *not* what is observed. As seen in figure 6.3 (b), the level of cooperative coordination is a linearly *decreasing* function of c , even in ranges where relatedness does not vary. The condition for cooperation is given by establishing the sign of the inequality between the inclusive fitness for cooperation and for no cooperation, both calculated using equation (6.1). There are two parameters in this equation: relatedness and c . A linear variation of the degree of cooperation with c could only be explained by some spatial variation in relatedness resulting in a linear decrease of the proportion of players for whom cooperation implies the best increase in inclusive fitness. Not only is this variation not observed, but it also would have to be manifested as a decrease in the global average of relatedness. This observation suggests that, whatever the mechanism responsible for the stabilization of cooperation, it must not work equally well for all values of c below 0.7 as kin selection, in these circumstances, does.

Another reason for not supporting an explanation purely in terms of kin selection is given by the coordination/size correlation discussed above (6.4). It seems logical that if the determining factor of the level of cooperation within a cluster is given by an optimization of inclusive fitness, then any variation in this level from one cluster to another must be given by variations between




Figure 6.8: Relatedness vs. cluster size for different two of c . Lines depict linear regressions. “R1” and “R2” refer to relatedness-1 and relatedness-2 respectively.

the average values of relatedness within those clusters because the parameter c is constant for the whole population. Then, the relation between changes in relatedness and cluster size should be examined. For two typical simulations with different values of c the results are plotted in figure 6.8. An increase in relatedness with cluster size is not observed as would be the tendency to expect given the correlation shown in figure 6.6. On the contrary, if any tendency may be drawn from this figure, it is a slight tendency towards increased relatedness for *smaller* clusters, which is not very surprising since such clusters have smaller interbreeding populations and the chances of meeting unrelated individuals are less than in larger clusters.

In conclusion, even if it is accepted that kin selection could play a role in the most favourable case, it is certainly not the whole story and further investigation is needed in order to explain the variation of the average coordination level with c and the correlation between this average coordination and cluster size for a fixed value of c . As will turn out, it will be possible to discard kin selection altogether when the structuring effects provided by the clusters are examined in detail.

6.6 Clusters as structuring invariants of the evolutionary dynamics

In this section a description will be provided of the two most important factors that arise from the presence of clusters as stable or quasi-stable patterns in space. One of these factors is purely geometrical and it is related to the fact that players are discrete entities, connected by a relation of genealogy along which the positions of successive generations describe a random walk of variable step size. The other factor, already discussed in section 5.5.3 is ecological as well as geometrical, and it is related to the breaking of in-built symmetries in the game and how agents are subject to different environmental conditions depending on their position within a cluster.




Figure 6.9: Probability of an entire genealogy remaining within cluster (delimited between positions 0 and 30) after 5 generations as a function of the ancestral position (x_0) within the cluster. $L = 10$.

6.6.1 Fluctuations and phenotypic aspects of spatial position

Section 6.3 described how clusters are constituted as a consequence of the interaction between agents and how they form a coherent quasi-stable spatial structure due to the equilibrium between a tendency by the agents to exploit more resources by expanding the size of the cluster and a tendency to accumulate in zones of high density so as to increment the chances of playing the game. Apparently, this equilibrium would seem to establish a degree of spatial neutrality in the sense that spatial position does not matter for the rates of energy intake and offspring production. Players at the center play the game more often than those at the periphery but they do so for poorer resources. If the rates of net energy intake were different, the cluster would not be at equilibrium. Such a homogeneous rate of energy consumption is indeed what is observed in the simulations. All players within a stable cluster have the same average lifespan and number of offspring. However, it is *not* true that spatial position is neutral in evolutionary terms.

Imagine a stable cluster of a definite size, and suppose that a new player is born into this cluster. As a rule of the game, offspring are allocated near parents, so the position of the new player depends on the position its parents in the cluster¹⁰. Now, if the new player happens to be allocated “outside” the cluster, for instance if its parents live near the periphery, its chances of survival will be small since it will be able to play only with a few other players and the majority of agents in the cluster will be too far to reach.

So, at least in this “all-or-none” sense, spatial position is playing a phenotypic¹¹ role which is not genetically based: if you are born outside the cluster your chances of survival are reduced.

¹⁰Actually, it depends on the position of just one parent, the one whose high energy level triggered the reproduction event, (see section 6.1).

¹¹The word “phenotypic” is used to refer to those attributes of an individual which have an effect on that individual’s fitness. A phenotype is usually associated with bodily and behavioral characteristics of an organisms, but the concept can be extended to characteristics outside the body of the organisms which affect variations in genetic frequencies within a population, (Dawkins, 1982). Phenotypic properties may trivially be unrelated to genetic constitution, e.g. social position in a dominance hierarchy, etc.




Figure 6.10: Spatial distribution of ancestors (a) and current population (b) after 10 generations.

Would it be possible to retain a weaker concept of spatial neutrality and say that at least, the positions *within* a cluster are evolutionarily neutral? A careful consideration of the above example indicates that this is not possible. A player at the periphery may generate the same average number of offspring as a player at the center, but the chances of some of these dying are greater because of its own position. In general, by retracing the genealogy of the current population of players in the cluster, it would be expected that, as one goes backwards in time, *ancestors will be more likely to be found near the central region of the cluster.*

This is simply a geometrical consequence of the stochastic and local character of the process of offspring allocation. This point can be illustrated with a one-dimensional example. Suppose that all agents receive the same amount of energy per unit of time and reproduce at the same rate. The cluster occupies the space between position 0 and position L_c . Each new offspring is allocated at a position $x_{g+1} = x_g + \Delta$, where x_g is the parental position at generation g and Δ is a random variable uniformly taken from the interval $[-L, L]$, (L being, as before, the neighborhood radius). This is a random walk process with stochastic step size and non-recoverable boundaries, i.e. if for any generation g the corresponding position x_g is greater than L_c or less than zero the genealogy is cut for all generations greater than g . Let $P(x_0, g)$ be the probability that no player in a lineage starting in x_0 at generation 0 has escaped from the cluster after g generations. How does $P(x_0, g)$ vary with x_0 ? This calculation is done using a Monte-Carlo method¹². As shown in figure 6.9 for a particular case of $g = 5$, $L_c = 30$ and $L = 10$, this probability has a peak right at the center of the cluster and decreases towards the borders.

In order to find out whether this is the situation in the present model a record was performed of the spatial position for the ancestors of the players constituting a cluster. The theoretical expectations are confirmed. Figure 6.10 shows the spatial distribution for the current population and for the corresponding ancestors of ten generations ago. Ancestor clusters are indeed much less disperse suggesting that the ancestors occupied the central region of their corresponding cluster. The size of the ancestral cluster decreases as one goes further backwards in time (also showing slight spatial displacements in the cluster's position).

¹²See for instance, (Hammersley & Handscomb, 1964; Rubinstein, 1981).




Figure 6.11: Differences in role frequency within a cluster. Agents in central regions (triangles) have played on average the second role more often than the first.

It is reasonable to conclude that position within a cluster has a phenotypic role with bearing on the fitness of the players but a role which is not under genetic control. Genealogies initiated in the central region will tend to be longer than those initiated in the peripheral region and, at any moment, the stable cluster will be constituted by players whose ancestors were most probably inhabitants of the central region.

6.6.2 Broken symmetries within clusters

It has already been mentioned in section 5.5.3 that the frequencies with which an agent plays the first and second roles are a function of the agent's position within a cluster. In a non-spatial situation this does not happen, but in the presence of spatial patterns this symmetry may be broken. As a result, some intuitions about the game may be challenged. For instance, it is clear that in a symmetric situation playing the first role is always more effective in terms of payoff than playing the second role, since a first role player is bound to receive at least half the energy in the food source, and maybe more depending on the strategies played, while the second role player may receive *at most* half the energy. However, as demonstrated in the continuous spatial model (see figure 5.13), when the symmetry in relative frequencies is broken there may be circumstances in which a player receives a greater average payoff when playing the second role than when playing the first (for instance, if the cluster is constituted entirely by a unique self-cooperating strategy).

Figure 6.11 shows, in a particular simulation run, the spatial distribution of players according to which role they play more often. It is clear from this figure that there is a spatial correlation according to which those agents that play the first role more often than the second tend to occupy the periphery of the clusters and those that play the second role more often, the central region.

Of what consequence is this manifest asymmetry? One may advance the hypothesis that for values of c sufficiently closed to 0.5, there will be an inner region in each cluster where players will be receiving a higher contribution to their payoff per unit of time due to their involvement as second role players and, given the spatial proximity of players in this region, they would evolve towards a self-cooperating strategy such as 1 or 3, since any growth of a non-cooperative mutants would rapidly be overcome by a compensation of the inner population towards the other self-cooperating strategy. Given that, as discussed in the previous section, all players in a cluster come

from a genealogy originated at the cluster center, then one would expect a level of coordination above the baseline level to evolve for the whole cluster. As c increases this level would drop since the inner region of agents being selected mainly for being good “interpreters” would become smaller. This would be an explanation which is able to account for the observed variation of the level of coordination with c , [figure 6.3 (b)].

However, this hypothesis is only tentative and may not be very convincing. In order to make the situation clearer a more careful analysis is needed.

6.6.3 Coordination explained

Consider first what would be the expected outcome if the model were to behave like one of the mathematical models presented in chapter 5. In those models it was shown that a stable attractor of the non-spatial dynamics was, for any value of c , a balanced combination of the non-cooperating strategies (strategies 2 and 4) and no individuals, save for fluctuations, using cooperating strategies (1 and 3). Could a similar situation occur in the spatial model?

It is clear that this stable solution “works” in the mathematical models because of the fine global balance between the strategies. A consideration of the cooperative/conflictive relationships between the different types (figure 5.1) shows that strategy 2 is susceptible to invasion by strategy 3 which can be “restrained” from growing by strategy 4. Similarly, strategy 4 is susceptible to be invaded by strategy 1 which is “restrained” by strategy 2. As a result, the two non-cooperating strategies are able to “look after each other” but only if they are present in equal numbers and accessibility between players is global. In the spatial and stochastic case there is no way that this uniformity can be maintained, especially since spatial symmetry is rapidly broken, and there will always be a possibility of an “un-constrained” invasion in areas where one non-cooperating strategy forms a pocket and the other is absent.

The other typical solutions, both for spatial and non-spatial continuous models, were oscillations. In the present model these are not observed although, for high values of c the dynamics do not seem to settle to a steady state and the proportion of coordination fluctuates around a mean value of 50 %. This could be understood as the effect of cyclic local invasions between strategies, which occur without synchronization due to the relative independence of clusters. However, it is important to consider that one factor that may make oscillations more difficult to occur in this model is the fact that agents have a finite life spans, thus adding in combination with spatial relations a certain degree of inertia into the dynamics¹³.

For values of c near 0.5, the model presents a much more stable dynamics (figure 6.2), and an investigation of the composition of the populations at different times suggests convergence into a stable state. In order to explain this it is necessary to consider some features of the central region of a stable cluster. It will be important to remember the result of section 5.5.4 according to which an invasion can only occur if the invaders achieve a threshold concentration. This will happen if the extra gain obtained by the invaders is significantly greater than the payoff of the established strategy. Due to the stochastic nature of the process it is important that this extra payoff be greater than the fluctuations in payoff experienced by the population, otherwise the effects of the extra

¹³In the similar model presented in chapter 7 a test was performed where the effects of the spatial dimension were removed by making local neighborhoods as big as the whole environment. In this case oscillations were recovered.




Figure 6.12: Critical frequency in which an invader's extra gain equals the fluctuations in payoff in an initially cooperating cluster. The horizontal line is the typical central frequency for playing the second role as measured in the simulations. Also according to measures, the standard deviation of fluctuations in payoff are about 9 % of the mean payoff, see text.

gain will be overshadowed by noise¹⁴.

1. *The central region can be modelled as a patch.* Because of its typical size of about 1 and 2 neighborhood areas, the central region of the cluster can be approximated as a single patch where only one strategy will be represented in high proportions and the rest of the strategies will be present only in small quantities. If one of these strategies were able to invade, this strategy would be expected to dominate in the long term. What cannot be expected is a similarly weighted stable distribution for two or more strategies in such a small area. Thus, the center of the cluster will be considered, for explanatory purposes, as a patch in which only one strategy can be dominant surrounded by a field (which is not necessarily treatable as patch in itself).
2. *Invasions occur only after a threshold number has been achieved.* In general, independently of the invading and invaded strategies, for invasions to occur a minimal threshold density must be achieved, since the population is constituted by discrete entities subject to fluctuations. This was already demonstrated in the spatial model presented in the previous chapter, as in other models and empirical studies, (see section 5.5.4 and references therein). Such a minimal threshold is needed even if the mean payoff is greater for the invader. A good estimation of the circumstances in which this threshold will be achieved is when the extra payoff obtained by an invader rises above the level of fluctuations in payoff. If the extra gain is smaller than the level of fluctuations it is unlikely that the minimal threshold will be achieved since the invader can be said to be doing as well as the established population.
3. *Non-cooperative strategies are invaded for all values of c .* Let f denote the mean frequency with which players in the central patch play the second role and $(1 - f)$ the frequency for playing the first role¹⁵. Consider a whole cluster constituted by a non-cooperating strategy (2 or 4). In each contest the first player receives a proportion c of the energy at stake and the second player receives nothing. Therefore, the (time) average payoff for a player in the

¹⁴For an introduction to the theoretical basis of some of observations that follow, see chapter 5 of (May, 1973)

¹⁵Typical values for f were 0.53, with individual peaks at about 0.62.

central patch will be:

$$\langle W_o \rangle = [(1-f)c]\langle E \rangle,$$

where $\langle E \rangle$ is the average energy extracted per unit of time. If a cooperative individual capable of invasion is placed in the patch, its average payoff will be

$$\langle W_i \rangle = [f/2 + (1-f)c]\langle E \rangle.$$

This is so because all games for this new agent are against the established strategy; therefore it will receive a payoff of half the energy when playing the second role and a payoff of c when playing the first.

Fluctuations in payoff are considered to be given mainly by fluctuations in energy (i.e. fluctuations in frequencies are ignored). The standard deviation in payoff is estimated from observations as

$$\sigma_W \cong 0.09\langle W \rangle.$$

The net difference in payoff for the invading strategy will be

$$\langle W_i \rangle - \langle W_o \rangle = f/2\langle E \rangle$$

which is of the same order of magnitude as $\langle W_o \rangle$. Therefore, this difference, which does not depend on c , will always be larger than the observed fluctuations and as a result the number of cooperative invaders will grow steadily.

4. *Invasion of cooperative strategies depends on c and f .* Now consider a whole cluster constituted by a self-cooperating strategy such as 1 or 3. The average payoff for a player within the central patch is

$$\langle W_o \rangle = f/2\langle E \rangle,$$

while the payoff for an individual non-cooperative invader will be

$$\langle W_i \rangle = [f/2 + (1-f)c]\langle E \rangle.$$

The extra average payoff for the invader is

$$\langle W_i \rangle - \langle W_o \rangle = (1-f)(c - 1/2)\langle E \rangle.$$

It is clear that for certain combinations of c and f this difference will be smaller than the size of the fluctuations in $\langle W \rangle$, so that the invader will be indistinguishable from the cooperating type. In fact, if one keeps the estimates that the standard deviation of the payoff is about 9 % of the average payoff and $f \cong 0.53$ and plots the critical value of f where the extra payoff equals the standard deviation, it is possible to see (figure 6.12) that for c up to 0.6 the extra gain of the invader will be shadowed by the fluctuations. For these values of c the number of invaders will be contained below the threshold of invasion. However, statistically this number is bound to increase (without necessarily invading) as c approaches the critical value, which can be seen by noting that if one takes a mean of the extra payoff among all the potential invaders, the size of the fluctuations for this averaged quantity decreases with the number of individuals.

This seems reasonable enough. For certain values of c non-cooperating invaders in the central region do not make enough extra gain to overcome the fluctuation effects and achieve threshold concentration. But why should this affect the cluster as a whole? After all the conditions in the




Figure 6.13: Differences in role frequency. Triangles indicate individuals that play the second role more than 105% (a) and 110% (b) more often than the first and dots indicate the rest of the agents.

peripheral region are the opposite in terms of role frequencies, so that invaders should be more favoured in this region than in the center. As discussed in section 6.6.1, fluctuations in fitness are more pronounced near the border due to a higher risk of placing offspring “outside” the cluster. This is especially manifested in the lengths of genealogies initiated in the peripheral region which tend to be short. In fact, at any given moment, the whole population of the cluster can be traced back to a group of ancestors in the *central* region. So that whatever happens at the center of the cluster, although it does not completely determine what happens with the cluster as a whole, it can strongly influence its final composition.

Consequently, as a result of the internal structure of the clusters, both in terms of relative frequencies for playing each role and in terms of the directionality of the lines of genealogy, and also as a result of fluctuations due to the discrete nature of the agents and the stochastic nature of their interactions, for certain values of c cooperative coordination is the stable evolutionary outcome instead of oscillations or combinations of non-coordinating strategies.

The above explanation is able to account for the observed behavior of the mean coordination level as a function of c , [figure 6.3 (b)]. Even though total invasion may not occur, the number of potential invaders maintained at a quasi-neutral evolutionary state does increase as c increases, thus reducing the average coordination level for the whole population. This was one of the features that could not be accounted for in a kin selection explanation. Another such feature was the correlation between cluster size and coordination level within clusters. Since an important test of any scientific explanation is whether it is able to account for more phenomena than the specific questions that gave origin to it, it would be important to see if the coordination/cluster size correlation can also be explained using the same argument.

6.6.4 Correlation explained

In general, it is observed that bigger clusters either grow from smaller ones, or are relatively big at the formation stage, so that there are two phenomena in need of explanation: 1) why do initially bigger clusters favour more coordinated activity, and 2) why an increment in coordinated activity

can provide a tendency towards cluster growth.

Explaining the second phenomenon is rather straightforward if one recalls the opposing tendencies that determine cluster size. Suppose that a cluster is at equilibrium with a certain level of coordinated activity. If, for any reason, this level is slightly incremented, this means that the agents in the cluster will exploit the resources more efficiently since the average fraction of the energy extracted per game will also increase. This is an unstable situation resulting in a number of players greater than the allowed by the energy rate of renewal in the area occupied by the cluster. The possibilities are either a compensatory growth in cluster size, a restoration of the previous level of coordination, or a split of the cluster into two or more sub-clusters. Assuming only a small increment in the coordination level there is little reason for the third alternative to occur (although it has been observed, especially in large clusters where the energy level in the central region is very close to zero). So it could be said that if the increment in the level of coordinated activity remains fixed, then the cluster will have a tendency to grow.

In order to understand the other alternative, i.e. why bigger clusters may favour a higher coordination level, it is necessary to recall that for certain combinations of f and c a cluster constituted by cooperating agents will not be invaded in its entirety although a number of non-cooperating players may exist within it. Such a number will also depend on the combination of c and f . Since c is fixed during the simulation, variations in the level of coordination of uninvaded clusters will be given mainly by variations in f , the mean frequency for playing the second role in the central region. The higher the value of f the smaller the extra payoff obtained by non-cooperating invaders (see above) and consequently, the higher the chances that these invaders will be “shadowed” by fluctuations. So, for fixed c , the level of coordination within the cluster will be a decreasing function of f .

Now it is necessary to show that bigger clusters will tend to have higher values of f than smaller ones. This can be seen empirically by plotting in a way similar to figure 6.11 the position of players who play the second role not just more often than the first one, but more often than 105 % or 110 % the number of times they play the first one (figure 6.13). Here it is observed that these agents tend to form a more compact subcluster in the central areas. Comparing figures 6.11 with figure 6.13 (a) and (b) it is apparent that big clusters admit higher values for the second role frequency (f) than small ones. And therefore invaders in the central patch will tend to have a lesser extra payoff.

Consequently, the structuring effects of clusters on the evolutionary dynamics are sufficient to explain why coordination is stabilized, its dependence on c and, additionally, why coordination level and cluster size are positively correlated.

6.7 Conclusions and analogies

The investigation pursued in this chapter affords the derivation of a series of conclusions in different contexts of relevance for this thesis. First, there is the methodological context which has to do with the issues discussed in chapter 4. Then, there is the specific interpretation of the results in terms of historical processes as discussed in chapter 3. The role of spatial clusters as structuring invariants of the evolutionary dynamics allows a series of comparisons with similar processes and a re-evaluation of the importance of different dynamical factors in natural evolution. Such a role

will suggest one possible avenue of further investigation which will be pursued in chapter 7. Finally, interesting analogies may be drawn between some aspects of this model and more concrete natural phenomena. These analogies may serve as a basis for further speculation regarding the usefulness of the explanatory mechanisms utilized here for other domains.

On the methodological side, the use of an individual-based model is at the same time a continuation of the series of comparisons between models of different depth which was initiated in chapter 5 and a discontinuation in terms of the techniques involved for building the model and analysing the results. The comparative approach has facilitated the identification of relevant factors and their consequences. While, in order to use such a comparative approach, the present model can be considered as an extension of the previous mathematical models, the actual construction and analysis of this model has followed the methodological guidelines expressed in chapter 4 for the use of computer simulations. Accordingly, the model was not built simply to reproduce the mathematical results and permit further relaxation of assumptions. On the contrary, the whole situation under study was directly modelled in a computational form thus allowing a radical break with a set of underlying assumptions whose effects were not necessarily evident in the mathematical models but were brought to light at the time of comparing results. Such is the case, for instance, of the coupling between the stochastic nature of agent reproduction and spatial position resulting in a better understanding of the phenotypic role of space. This phenomenon was not even suspected in the mathematical approach because of the nature of the way reproduction was modelled as a population based instantaneous phenomenon. However, in no way should this be interpreted as saying that the individual-based approach is better or more important than the mathematical models since, without the latter, the former would not have been understandable¹⁶.

The study of the results also followed the guidelines given in chapter 4. An explanation of the global patterns observed in a complex computer model does not necessarily have to be based on the lower level phenomena which are directly expressed in the model itself, but may be achieved in terms of other global phenomena to which the model gives rise. Accordingly, an explanation of why cooperative coordination is stabilized in certain circumstances was primarily based on the structuring effects of clusters on the evolutionary dynamics. These effects, mainly to do with the assignation of a meaningful phenotypic dimension to spatial position and a breaking of the in-built symmetries of the game, are in themselves global phenomena which arise from the basic rules of the model and which are partly dependent on each other. Thus, the explanation of the observed results is better described with a diagram like figure 4.2 (indirect explanations) than with a diagram like figure 4.1 (direct explanation) which is the way that is usually suspected that a computer model should be used.

In conclusion, the model studied in this chapter is a concrete proof of the suitability of the methodological guidelines proposed in chapter 4 for the use of computer modelling in scientific research.

As with many complex phenomena, it is usually preferable to try to find explanatory principles which will unify the type of evolutionary processes investigated here and other evolutionary processes under the same framework. Indeed, the more contexts in which similar explanations are

¹⁶Unfortunately, the presentation of the last two chapters has tried to proceed in logical, rather than historical, steps and therefore it may give the reader the impression of a linear unfolding of the research process while, in actuality, this process has been a constant re-working and comparative analysis of all the models until the last moment.

applicable the more powerful the explanatory principle. That unifying role has often been given in evolutionary biology to natural selection. However, when natural selection was intended to be used as the main explanatory factor in the current model, it failed. Such failure became evident when the best candidate explanation in the present context, kin selection, was intended to account for the observed results. In fact, kin selection proved to be quite weak as an explanation given that a lot of concessions regarding actual observations had to be made even to be able to apply it and then not only did it not account for the observed stabilization of cooperative strategies but it also was contradicted by other phenomena like the correlation between cluster size and coordination level.

In contrast, the actual explanation provided in this chapter, while still relying heavily on differential reproduction, is not constructed in terms of any unifying and unsituated principles but it is rather based on the organizational properties of some invariant features of the actual dynamics. As such, in this explanation some historical aspects of the process are brought forward. These aspects are manifested clearly in the co-dependence of many of the relevant explanatory factors such as cluster formation, cluster stability, cluster size, symmetry breaking and behavioral strategies selected. While this model lacks the richness that one tends to associate with historical processes in general, the main point in the explanation of the stability of cooperative coordination is the fact that some features which emerge from the whole ecological and evolutionary dynamics become the constraining and organizational conditions which influence the path that the dynamics will follow thereon. That is the nature of a historical process as discussed in chapter 3. In addition, those same features tend to select dynamical paths that will result in their own stability which, consequently, gives them their invariant character. Such is clearly the role played by the spatial clusters of interacting agents in this model.

The actual explanation of the results starts with an understanding of the dynamical “forces” giving rise to cluster formation and stability. Once these are understood, other aspects become evident such as the differences in environmental conditions between the central and peripheral regions of a cluster. While cluster stability implies an uniform level of individual rate of reproduction for all the agents, it soon becomes apparent that the phenotypic effects of spatial position are not neutral but are manifested in the length of the genealogies initiated in different regions of the cluster. Accordingly, after a few generations most of the current inhabitants of a cluster will have ancestors who lived in the central region. This is a consequence of the stochastic nature of local reproduction and the discrete nature of the agents in combination with the forces that maintain the stability of the cluster.

Because of its typical size and the local nature of sexual reproduction, the central region of a cluster can be considered as a patch composed mainly by a single strategy. One cannot expect to find a stable combination of different strategies in similar quantities in such a small area. If the main central strategy can be invaded by another one, then the new strategy becomes the central one. However, due to the discrete nature of the agents and due to the stochasticity involved in reproduction events, for an invasion to occur it is necessary for the invading strategy to achieve a minimal threshold in the number of players. Otherwise, the beneficial effects of the invading strategy are overshadowed by fluctuations.

A combination of these facts with the broken symmetries induced by the cluster’s structure

in the frequency for playing each role in the game, results in different conditions for achieving the invasion threshold depending on the nature of the strategies involved. A cooperating strategy which has an advantage over a non-cooperating one may always achieve the invasion threshold. In contrast, if the invasion is to be achieved by non-cooperating agents who have an advantage over a cooperating strategy at the center, then the invasion will depend both on the frequency with which the second role is played at the center (f) and on the parameter c . For certain combinations of these quantities invasion of the central patch will not occur. Since, after a few generations, most individuals in the cluster will be descendants of the agents in the central patch, the central strategy will always tend to be represented in a high proportion in the rest of the cluster even if the competing strategies do better there.

This explains how cooperative coordination can be stabilized under certain circumstances. Additionally, as c is incremented the number of individuals bearing a competing strategy will rise (even if they do not invade the cluster completely) causing a dependency of the global level of coordination with c in accordance with what is observed in figure 6.3 (b). Finally, the same factors can be used in explaining the correlation between cluster size and the level of coordination.

Besides the main structuring factors (like symmetry breaking), several seemingly innocuous factors are also brought to light in this explanation, like discreteness, fluctuations, etc. It can be concluded that use of novel techniques, within the proposed methodology, allows a re-evaluation of the relative importance of different explanatory factors in complex evolutionary phenomena. Similar conclusions can be derived from other lines of research mentioned elsewhere in this thesis such as the work of Hogeweg, Boerlijst and colleagues on spiral waves and hypercycle dynamics (see section 5.5) or the organizational investigations of artificial chemistries by Fontana, Buss and colleagues (section 7.7). If this tendency affords any general reading, it surely must be that no single explanatory principle can do the whole job of explaining complex situated evolutionary and historical processes.

In addition to this general comparison between the results of this investigation and other models in which the structuring effects of invariants become key explanatory factors, it is possible to draw other broad analogies with natural phenomena if the spatial structures of this models are interpreted as metaphors for other factors which may have similar constraining effects on the dynamics. Such could be the case of group hierarchies, social networks, etc. It must be remembered that the current investigation is not intended to provide accurate models of natural phenomena but to gain a theoretical understanding of the many factors that are assumed to be involved in the explanation of natural phenomena. Therefore, these analogies are intended to be suggestive and perhaps open the door for further speculation.

As a particular case, the positive correlation between cluster size and level of cooperative coordination within the cluster has interesting parallels in the animal world. In primate societies a similar correlation has been postulated between group size and relative neocortical size (Sawaguchi & Kudo, 1990; Dunbar, 1992). It has been suggested that the evolution of complex forms of communication leading to human language has been influenced by the growth in group size in humans as compared with other primates, and a functional explanation has been advanced in terms of the role of language as a bonding mechanism in the maintenance of stability in large groups (Dunbar, 1993). The way that the analogous correlation has been explained in the present

study was rather different, more in terms of dynamics and feedback mechanisms than in terms of functional adaptations. Nevertheless, the parallel is worth noticing because some aspects of the operational account offered in this chapter could probably be translated into the specific domain of group size and communicative capabilities in primates.

Another important issue discussed in this chapter is the fitness value of a spatial position within a cluster. Interesting parallels may be drawn between this situation and the fitness value of spatial position within a natural group of animals. In general, individuals positioned at the edge of a group may suffer greater risks of predation or parasites than individuals located at the center. On the other hand, while edge individuals incur into greater energy expenditure, in some cases they benefit from a greater energy intake due to more effective foraging. Dominance hierarchies also seem to be correlated with spatial position, where dominating individuals occupy central positions in a group. Krause (1994) reviews a number of studies on the subject. The theoretical basis for most of them is Hamilton's "selfish herd" theory (Hamilton, 1971), which stipulates that individuals seek a position within a group following their self-interest even if the end situation results in a lower average fitness for the whole. Unfortunately, not all studies revised by Krause confirm this hypothesis; some even contradict it. A recent computational study by Hemelrijk (1998) also challenges the idea that dominant individuals have a *preference* for spatial position, showing that they will tend to occupy a central position anyway, as a feature of the self-organizing dynamics of interaction. If the analogy is valid, the way in which spatial position plays a phenotypic role in the current model would seem to support an explanation of the latter kind.

All these analogies suggest that it may be fruitful to speculate about these phenomena using what has been learnt from the current model. In addition to that, the same analogies suggest ways in which it would be interesting to extend the current investigation. This is what is done in the following chapter.

Chapter 7

Social influences on development and coordinated behavior: an assessment

This chapter presents an extension to the models studied in chapters 5 and 6 which is still concerned with the question of social coordination from an evolutionary and ecological point of view rather than from a behavioral one. In particular, the present chapter will continue to explore the dynamical role played by processes that could be roughly situated at an intermediate time scale between behavior and evolution.

The dynamics of spatial organization, as described at the end of chapter 6, can be read as one example of such a process. In conjunction with natural selection, spatial structures play a fundamental explanatory role in the understanding of the evolution of cooperative coordination between interacting agents. In order to explore this theme further a search will be made for natural candidates that may (potentially) fulfill similar organizing roles as the spatial dynamics did in the previous model.

One possible family of candidates comes to mind when one recapitulates how spatial structures could be interpreted as metaphors for other types of structures such as social hierarchies, affiliative networks, etc. All these examples have as a common feature a similar way of influencing the evolutionary process by constraining the patterns of realizable interactions, in that particular case this is manifested in the form of restrictions to the available partners in the game and symmetry breaking in role frequency.

This chapter will concentrate on the particular case of parental influences on development. To support the hypothesis that a process such as this can affect the evolution of social interactions, the relevant evidence from studies of animal behavior will be reviewed. The specific question to be asked will be whether such an influence can introduce qualitatively important changes in the evolutionary dynamics, and what sort of changes these may be. As a point of departure, the same action-response game studied in previous chapters will be used in a simplified version. This will facilitate understanding by means of systematic comparisons.

As before, the model does not attempt to address specific behaviors observed in natural cases, but rather to extend the understanding of the relationships between the processes involved in order to help in the formulation of concrete hypotheses of what to expect in natural cases. For this

reason, the model will be kept at a rather abstract level, but specifying what aspects of the relevant natural phenomena are intended to be included.

7.1 Social influences on development: evidence and consequences

As discussed in chapter 3, a tacit supposition in many evolutionary models is that those factors that remain outside the influence of natural selection are either fixed or vary at random. The evolution of a trait is often studied under the assumption that the relevant environmental factors change so slowly that they may be considered to be static (e.g. the distribution of food patches) or their variation is dependent upon so many contingencies that, for large populations, they may be considered to change randomly (for instance, mating opportunities). These are fruitful assumptions under which the evolution of the particular trait of interest becomes, as a dynamical process, decoupled from a myriad of other processes which, if considered, would make the treatment of the problem a very difficult task¹. However, it is apparent that in many cases this is an over-simplification, even for the limited understanding that is intended to be obtained from such models. The distribution and quality of food patches will partly depend on the efficiency of foraging strategies, and mating opportunities will depend, for instance, on space-time constraints and mate preferences.

This is a criticism that has been repeated many times (Bohm, 1969a, 1969b; Lewontin, 1982; Odling-Smee, 1988, among others)² and should be taken seriously if an understanding of evolutionary processes in general is to be achieved. On the other hand, it is clearly impractical to model a whole ecosystem if the intention is to understand the changes of a single trait in a particular species. If a compromise is to be found between accuracy, realizability and explanatory value, the relevance of those contextual factors which are not completely independent of the evolutionary process itself should be fully justified before including them in the model.

The aim of this chapter is to accomplish such a justification in the particular case where the evolution of cooperative social coordination is coupled with a process of social influences on individual development.

The assertion that social influences on development can play an important role in the evolution of social behavior is far from surprising. Indeed, the literature on social effects on early life-history in animals is vast and diverse. However, there seems to be a degree of conflicting terminology (in particular for cases where the influences are on behavior rather than on other physical traits like in the case of nutritional maternal effects), so that it is not uncommon for terms like “social learning”, “imitation”, “teaching”, “social facilitation”, “encouragement”, “scaffolding” and others to be used more or less in similar contexts but with varying meanings.

Caro and Hauser (1992) review the evidence for *teaching* in non-human animals. They also point to the existing terminological confusion and provide their own definition of teaching:

An individual actor **A** can be said to teach if it modifies its behavior only in the presence of a naive observer **B**, at some cost or at least without obtaining an immediate benefit for itself. **A**’s behavior thereby encourages or punishes **B**’s behavior, or provides **B** with experience, or sets and example for **B**. As a result, **B** acquires knowledge

¹In problems involving co-evolution assumptions like these are weakly relaxed so that the mutual coupling between two or more “single” evolutionary processes can be accounted for.

²See also section 3.2.3.

or learns a skill earlier in life or more rapidly or efficiently that it might otherwise do, or that it would not learn at all. (Caro & Hauser, 1992, p. 153)

It is clear from this definition that social influences on development include a wider range of phenomena than just teaching³. Thus, one could also include as an instance of social influence cases where the behavior of **A** cannot be said to be either encouraging or punishing the behavior of **B**, but just affecting it, without obvious immediate consequences to any of them. Or cases in which **A** is not actually “modifying its behavior”, but just behaving normally in the presence of **B**.

An example that fits very well with the above definition of teaching (but not perfectly) is given by the teaching of predatory behavior in felines, especially in domestic cats (*Felix catus*) (Leyhausen, 1979; Baerends-van Roon & Baerends, 1979; Caro, 1980), but also in free-living tigers (*Panthera tigris*) (Schaller, 1966) and lions (*P. leo*) (Schenkel, 1966) although the data is mostly anecdotal. Similar observations have been made on other species, including cetaceans like the killer-whale (*Orcinus orca*) in the Argentinian South Atlantic coast (Lopez & Lopez, 1985). In the presence of their offspring, mothers have been observed to modify their predatory patterns, which are usually quite efficient, in ways that will allow the former to gain some experience in hunting. Thus, an adult female cat does not kill her prey immediately but tries to catch the attention of her kittens by giving characteristic calls, letting the kittens play with the prey (which is immediately recaptured if it escapes), and so on.

In primates, examples of maternal “teaching-like” effects include the assistance of mother chimpanzees (*Pan troglodytes*) in the nut-cracking activities of their offspring (Boesch, 1991) and the encouragement of the development of infant locomotion in pigtail macaques (*Macaca Nemestrina*), (Maestripieri, 1996). This last report presents quantitative, less anecdotal data where mothers are seen to increase the distance of separation from their infant and display certain gestures that encourage the infant to come closer.

Complementing the observation of the actual maternal influences of behavior, some hypotheses regarding social learning have been supported by studies on the effects of mother-infant separation; Harlow (1958) and Hinde & Spencer-Booth (1971) are two classical examples in the case of primates. Another case is related in (Maturana & Varela, 1987, p. 127) where even a separation of a few hours of a newborn goat from its mother can induce long term detrimental effects on its adult social patterns.

As this last example suggests, maternal influences can also affect social relationships and social behavior, (see Chalmers, 1983). For instance, Holmes & Mateo (1998) have studied the maternal influence on the development of litter-mate preferences in Belding’s ground squirrels. They found that in the absence of their mothers, juveniles do not play preferentially with litter-mates while if mothers are present, but unable to intervene, juveniles develop a preference for litter-mates which they will tend to maintain throughout their lifetime. This suggests that the mere presence of a mother, and not her active intervention, is enough to influence changes in social behavior.

Mothers can also influence the future social status of their offspring. For instance studies of “rank transmission” from mother to daughter in rhesus monkeys (*Macaca mulatta*) suggest that the infant’s social networks resemble its mother’s not just during the first weeks of life, when

³Caro and Hauser conclude that none of the cases reviewed fulfill their definition in all respects.

the mother has total control over the infant's movements, but also when the infant starts to spend more time separated from its mother, (Berman, 1982a, 1982b). A more careful study on this phenomenon has been presented by de Waal (1996). Among the possible mechanisms suggested for explaining this perpetuation of social networks is *early familiarity* which has also been suggested as a plausible mechanism for kin-recognition (Frederickson & Sackett, 1984). According to this mechanism, affiliative preferences develop as a simple consequence of frequency of exposure which, during the early stages of an infant's life, depends strongly on the existing maternal affiliative network. Such a mechanism can be naturally regarded as potentially instantiating a historical process in the sense discussed in chapter 3.

Similar cases that would fall under the rubric of social influences on development, would be filial or sexual imprinting, whereby parents influence the mate preferences of their offspring. The relevant literature is vast and the phenomenon will only be mentioned here⁴.

One salient characteristic of all the above examples is that social influences on development do not necessarily imply the complex cognitive capabilities that one may initially tend to associate with terms like imitation, teaching, etc. As suggested by Ewer (1969), the setting of the right environmental situation by the parent is enough for the developing animal to be influenced; others mechanisms may be simpler, as in the case of early familiarity. Nor, for that matter, are all of them classifiable as being cases where a behavior or trait is "transmitted" or "copied" via non-Mendelian inheritance. Some of the examples may indeed look as cases of non-genetic (or cultural) transmission (as the case of predatory patterns in felines) and some mileage may be obtained from thinking in those terms. But this is not the general case and therefore this assumption will not be made in the model⁵.

Is it possible to speak in general of the evolutionary consequences of socially influenced development given the variety of cases discussed? Of course, from a functional perspective it is possible to argue for the selective value of some of these processes. For instance, imprinting of sexual preferences may represent an optimal compromise between inbreeding and outbreeding (Bateson, 1978, 1980). Other adaptive stories may relate the effects of the directionality of maternal effects in affecting plastic development. However, not all the effects may be functionally adaptive. Kirkpatrick and Lande, (1989, 1992) present quantitative models that show that maternal effects may produce time-lags in adaptation or even the evolution of maladaptive characters. Laland (1994) presents a model which predicts similar results for sexual imprinting whereby the probability of the spread of novel adaptive traits is reduced due to the preferential resistance to novel characters⁶.

In the context of this thesis particular interest is given to the operational or dynamical aspects introduced in an evolutionary process by allowing social influences on development. Regarding the question about the conditions for social influences to evolve in themselves, this will not be addressed here (although some simple modifications to the model would probably suffice to study it). The aim is rather to compare evolutionary processes without socially influenced development

⁴For reviews on imprinting and its evolutionary consequences, see for instance (Ridley, 1987; Hess, 1973). For a computational model on speciation induced by imprinting and sexual selection see (Todd & Miller, 1993).

⁵This will allow to arrive at some interesting conclusions concerning this issue.

⁶For another relevant evolutionary model incorporating maternal effects, see (Cowley & Atchley, 1992). Falconer (1965) presents a pioneering experiment on the selective consequences of maternal effects and Bernardo (1996) discusses their ecological significance.




Figure 7.1: Two extreme cases of developmental relationships between parent and offspring. (a) can give rise to plastic evolutionary effects (e.g. genetic assimilation), (b) can only give rise to plastic and other dynamical effects (e.g. perpetuation of developmental relationships). See text.

with processes in which social influences are unavoidable, although the nature of these influences may be subject to evolutionary changes.

In order to picture what kind of dynamical consequences can be expected, a distinction will be introduced between two aspects of the process of socially influenced development: 1) the influence on the development of a certain behavior and 2) the influence on the propensity to influence the development of future generations. Probably in most natural cases this distinction will not be very sharp, but it is intended to provide an idea of the different evolutionary roles that these processes may have. For instance, mother encouragement of the development of locomotion in macaques does not seem to be a behavior that would affect its *own* probability of being found in a future generation (although locomotion can be said to be influenced by mothers, it is not clear that the ability to influence locomotion in the next generation is affected by the mother's behavior). In other cases social influence could probably be seen as influencing their own chances of being found in future generations, as in the perpetuation of social networks⁷.

The distinction is illustrated in figure 7.1 where the two extreme cases are depicted. In this figure, phenotypes are represented by ovals divided into two sets of traits or behaviors (A and B). Arrows represent the influence on development from one generation to the next. Each arrow starts from the set of traits that is responsible for the influence on development in the next generation (A in all cases), and ends in the set of traits whose development is influenced. Figure 7.1 (a) shows a case where social influences affect certain behaviors (B) but which do not reflect on the capacity of influencing the next generation (A). In contrast, figure 7.1 (b) shows the case where the behavior (or trait) affected (A) also affects the chances of the process repeating itself in the next

⁷Adoption in non-human animals has been proposed as a possible case of a self-perpetuating social influence, (Avital & Jablonka, 1994). See (Hansen, 1996) and (Avital & Jablonka, 1996) for critical comments on this hypothesis and a response.

generation. It is clear, that the ecological and evolutionary consequences of these cases are not the same. In both circumstances it is possible to expect effects related to developmental plasticity in general [for instance, the so called Baldwin effect (Baldwin, 1896; Hinton & Nowlan, 1987), or the time lags or momentum effects discussed in (Kirkpatrick & Lande, 1989, 1992)], but the second case is more complex since the modified behavior includes the probability of modifying future generations and so it can be viewed as a process in which the dynamics affects its own conditions of realizability, i.e. a historical process. As such, it could give rise to self-sustaining or self-promoting structures (for instance, in the space of developmental relations).

It will be of interest to explore what kind of dynamical effects can result from the latter case. In order to do so behaviors and the ability to influence the development of offspring will be fused in a single phenotypic space. This is done ignoring many of the complexities involved by using a simple but justifiable model.

7.2 A model of socially influenced development

The purpose of this investigation is to make a *first order* assessment of the relevance that a socially influenced developmental process may have for understanding the evolution of cooperative action coordination, and also to try to identify what sort of constraints such a process may impose on the evolutionary dynamics. This justifies the choice of an abstract, and indeed crude, model of development under social conditions. In spite of this simplicity, the model is intended to account for the following relevant aspects:

1. Phenotypes are the result of a dynamical process which is affected both by genetic constitution and coupling between the developing organism and its medium.
2. The nature of the process itself is not totally determined by genetic constitution nor by environmental coupling but there are also formal constraints (for instance, the number of possible attractors a developmental system can have depends on the specific transition dynamics and may not be easily changed during evolution).
3. In most natural cases of socially influenced development, the most relevant coupling is between developing organisms and their parent(s).

The third point is supported by the examples discussed above while the first two can be derived from a Waddingtonian perspective on development as a dynamical process⁸. Figure 7.2 provides a qualitative picture of the kind of process considered. An offspring's genotype is determined by the genotype of its parents in the usual way (in this model using uniform crossover and point mutation). Its adult phenotype is determined both by its genetic constitution and by parental influences which depend on the adult phenotypes of the parents.

Put more formally, suppose that the *adult phenotype* of an individual i (denoted by P_a^i) is totally determined by its *phenotype at birth* P_b^i with the exception of certain changes which will depend both on P_b^i and the adult phenotype of i 's parents $P_a^{P^1, P^2}$. P_b^i depends only on the genetic constitution of i and it is obtained by a universal genotype-phenotype mapping.

⁸For a collection of Waddington's key papers on development and evolution see (Waddington, 1975). Also of relevance is the work of Brian Goodwin and Gerry Webster, for instance (Goodwin, 1993, 1994; Webster & Goodwin, 1996).

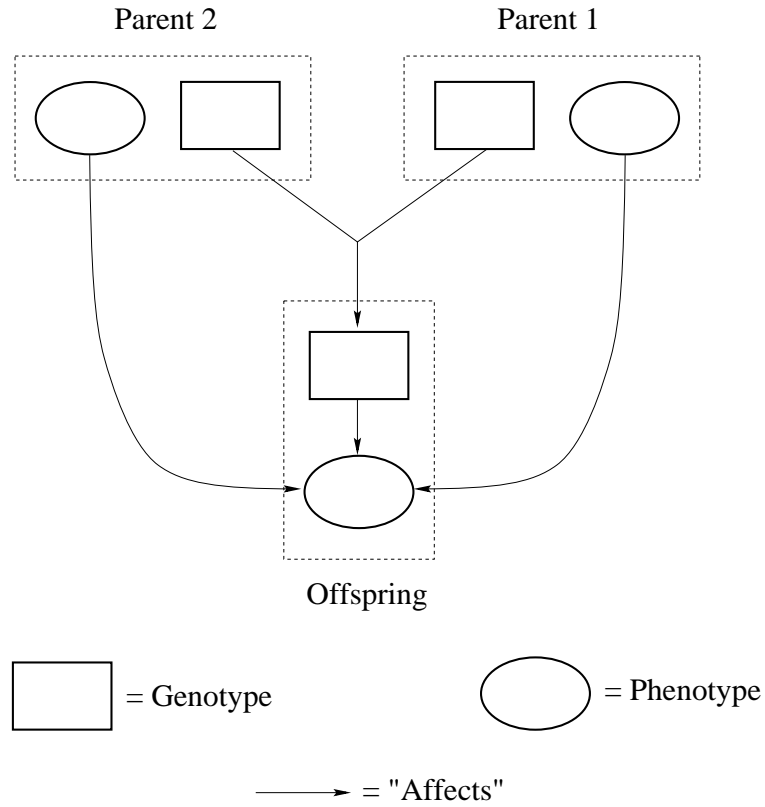


Figure 7.2: Scheme of abstract developmental process.

Of course, it is desirable that parental influence on development should reflect some formal relationships between adult phenotype and that of the undeveloped offspring. However, these formal rules will not be specified explicitly, but rather the possibility will be allowed for many of them to coexist in the same simulation run.

A set of general constraints is first defined for these developmental rules. The model will consist of a discrete and finite set of behavioral strategies and a correspondingly discrete and finite set of possible genotypes. Additionally, a discrete and finite set of phenotypes will also be defined. This space will have its own *phenotypic metric* which, in principle, need not be the same as the genotypic metric in terms of Hamming distances nor the same as a “behavioral” metric in terms of possible strategies and their relations. The reason for this is that the metric may also take into account other traits besides those being studied. As a process, development will be manifested on these “generalized” phenotypes and so it will clearly influence behavioral traits as well as following a logic of its own.

Under such a metric two phenotypes P and Q may be said to be close or apart and the distance between them is denoted by $dist(P, Q)$. Parental influences will be taken into consideration only if the following condition holds:

$$dist(P_b^i, P_a^{p_j}) \leq D, \quad j = 1 \text{ or } 2,$$

for D a small distance and the influence will be such that:

$$dist(P_b^i, P_a^i) \leq D,$$

where $P_a^i = Dev(P_b^i, P_a^{P_b^i})$ is the adult phenotype of individual i and the function $Dev(.)$ specifies the developmental mapping under parental influence.

This means that social influences on development will be manifested when parents are not too distant from the offspring's phenotype at birth and they will be such that the resulting adult phenotype will remain within the same distance⁹. If the above conditions are not fulfilled then $P_a^i = P_b^i$.

The above abstraction of the process of social influences on development follows a scheme similar to those studied in models of chemical networks of catalysis, especially, as will be apparent in the definition of the function $Dev(.)$ below, random reaction networks (Stadler, Fontana, & Miller, 1993). In such networks the interaction of two elements in a population of objects of different types results in the production of another object (product) which is added to the population. These models can contain as special cases models of catalyzed replication of individual objects (replicator equations, hypercycles) or of self-sustaining or self-replicating finite subsets of elements. The latter cases have been explored recently in the work of Walter Fontana, Leo Buss and colleagues (Fontana et al., 1994; Fontana & Buss, 1994a, 1994b, 1996) and some of their findings will be helpful in explaining what goes on in the present model (see section 7.7).

However, the model presented here is a bit more complicated, since entities are not only related through the developmental process, but also through genetic inheritance, and their relative abundance, as in the earlier model, will additionally depend both on selection pressures and spatial structuring processes.

The space of possible behavioral strategies will be confined to four different types which are discussed in the following section. However, in order not to trivialize the issue of development, the model will include the larger number of 256 phenotypes represented by integer numbers from 0 to 255. As previously mentioned, these phenotypes are thought of as including both the behaviorally relevant features (the strategies) and other features which are behaviorally (but not developmentally) neutral. All four types of behaviors are represented in equal proportions in the set of possible phenotypes¹⁰.

The metric over the phenotypic space will be the traditional metric over the integers between 0 and 255 but "wrapped around at the edges":

$$dist(P, Q) = \min(|P - Q|, 256 - |P - Q|)$$

The adult phenotype of newborn agents is obtained as a function of P_b^i and $P_a^{P_b^i}$, observing the conditions discussed above. For each value of P_b^i its developmental function $Dev(.)$ is represented by an integer symmetric matrix $M^{P_b^i}$ of size $(2D + 1) \times (2D + 1)$ which is defined at the beginning of the simulation run and is not modified afterwards (D , as defined above, is the maximum distance between phenotypes for a developmental influence to take place). The entries in this matrix are integers randomly selected from the interval $I = [P_b^i - D, P_b^i + D]$. The indexes of the rows and columns correspond to the same interval, so that the first element of the first row will have indexes

⁹The object of this restriction is to avoid making genetic constitution an irrelevant factor on development. If any adult phenotype can influence development then the relevance of genetic constitution is diluted while if only a small group "close" to the phenotype at birth can exert influences it is possible to guarantee some degree of correlation between phenotypic and genotypic proximity.

¹⁰Different mappings between developmental and behavioral space have been tried with no evident difference in the results. See next section.

$(P_b^i - D, P_b^i - D)$ and the last element of the last row $(P_b^i + D, P_b^i + D)$. If the adult phenotypes of both parents fall within I the adult phenotype for the offspring will be:

$$P_a^i = M^{P_b}(P_a^{P_1}, P_a^{P_2}).$$

That is, the adult phenotype will be the element of the matrix corresponding to the indexes specified by the parental phenotypes. If only one of the parental phenotypes falls within I the remaining index is taken to be equal to P_b^i :

$$P_a^i = M^{P_b}(P_a^{P_1}, P_b^i).$$

and if none of the parental phenotypes falls within I the adult phenotype is taken to be equal to the phenotype at birth. The matrices M^{P_b} for all the 256 possible phenotypes at birth define a *constrained random reaction-network*. Table 7.1 shows an example of M^{P_b} for $P_b = 50$ and $D = 2$. All entries are chosen randomly from the interval $[48, 52]$. For example, if the parental adult phenotypes are 51 and 52, the resulting phenotype will be changed from 50 to 49. (Note that the symmetry of the matrix means that it does not matter which parent defines the row and which one defines the column.)

Table 7.1: Developmental matrix for $P_b = 50$.

	48	49	50	51	52
48	50	50	52	51	53
49	50	50	51	52	53
50	52	51	48	50	49
51	51	52	50	50	49
52	53	53	49	49	53

In this abstract model the process of development is instantaneous. As soon as a new agent is born its adult phenotype is calculated and this will be its phenotype for the rest of its lifetime. This means that the model ignores the complex effects that may arise due to developmental dynamics and their coupling with other processes.

7.3 Computer model

The abstract reaction-network model described above can easily be implemented in the context of an action-response game equivalent to those presented in chapters 5 and 6. A population of agents are distributed in a two dimensional spatial environment. During its lifetime an individual agent does not leave its finite local habitat which is defined as a circle of radius R_0 around a fixed position with real-valued coordinates¹¹, and where all the interactions involving that agent take place (figure 7.3). This is a different definition of neighborhood from the one presented in chapter 6 (cf. figure 6.1).

As in chapter 6, at each time step agents are chosen asynchronously and with equal probability to try to play the game. An agent selected in this way will try to randomly locate a partner within

¹¹Typically, $R_0 = 8$ in a 200×200 arena.




Figure 7.3: An example of agents' position with circular neighborhoods. "F" represents a food source within reach of agent "A1" who has selected "A2" as a partner in the game. Notice that although both "A2" and "A3" are "within reach" of "A1", they cannot interact with each other. Compare with figure 6.1

its neighborhood with a probability of success proportional to the local density of agents ¹². At the same time a food repository will be randomly selected from those located in the neighborhood. The first player performs an action A and the second player performs a response R . If A and R are complementary in the sense described below, the two players are said to be coordinating their behavior. In such a case both agents share equally the payoff which is an amount of energy extracted from the local repository. If they do not coordinate the first agent receives a proportion c of the energy in the repository where, as in the previous chapters, $0.5 < c \leq 1$, and the other agent receives nothing.

In this simplified version A and R are fixed for all circumstances in a given individual: they do not depend on food types nor on a history of interactions. They are represented each by an integer number between 0 and 15. The condition for A and R to be complementary is simply that $parity(A) = 1 - parity(R)$ where $parity(.)$ returns 0 for even numbers and 1 for odd numbers. From this condition it is possible to distinguish four possible behavioral types depending on the combination of parities for A and R . Actions and responses may be odd (O) or even (E), therefore the four possible strategies are (EO), (OO), (OE) and (EE) where the first place indicates the parity of A and the second the parity of R . The relations between these strategies are exactly the same as the relations held by the four strategies used in the mathematical analysis of the coordination game presented in chapter 5. These are shown in figure 7.4 (cf. figure 5.1). In this figure, arrows

¹²As before this is implemented as a fixed number of trials in which the agent selects a player at random from the whole population. See section 5.5.

indicate a cooperative relationship between first player (arrow's starting point) and second player (arrow's end). Absent connections "indicate" non-cooperative relationships.




Figure 7.4: Cooperative relations between strategies. (EO) and (OE) are called “self-cooperating” or simply “cooperating” strategies.

As in chapter 5, two of these strategies, (EO) and (OE), can be called self-cooperating since agents playing these strategies will coordinate when playing against each other. Other cooperative outcomes will depend on the roles played by the agents. The mathematical analysis presented in chapter 5 applies directly to this game.

In one version of this game A and R are encoded in an 8-locus binary haploid genome using binary encoding. A is encoded in the top 4 bits and R in the lower 4 bits. In both cases only 1 bit is behaviorally significant: the one that determines the parity of the corresponding integer. However, the whole genotype determines the phenotypic number that is used in the developmental process; P_b , the phenotype at birth, is obtained by transforming the whole genome into an integer number. As an example, consider the string 00010100. The corresponding values for action and response will be $A = 1$ and $R = 4$. The phenotype at birth will be $P_b = 20$. After applying the developmental rules, a new adult phenotype is obtained (say 21) which is transformed back into a binary string (00010101) in order to obtain the corresponding new values of A and R (1 and 5 respectively).

It is clear that the correspondence between phenotypic number and the values of A and R will follow a strict rule determined by the above process. Although this rule does not favour any of the four behavioral types over the others, it may be the case that the results may depend strongly on its application, which after all does not intend to model any specific phenomenon, but is rather an implementational choice. To test whether this is the case, a different set of experiments was carried out using a random mapping between phenotypic and behavioral space. In this case, the 8-loci genome also encodes the phenotype at birth, but the corresponding behavioral strategy for that phenotype is assigned randomly at the beginning of the simulation. Although, the corresponding association remains fixed during the simulation run, any spurious effects should disappear after comparing many different runs. In all the cases, the results for both types of encoding were the same. A report will be made only on cases using the first type of encoding.

In all the experiments reported here the distance of allowed parental influence (D) is equal to 2 (modulo 256). Other values were tested between 1 and 5. Results did not seem to differ much, although the exploration was not as methodical as in the case of $D = 2$.

In a way similar to the model presented in chapter 6, agents store the energy gained from the

games they play and use it for self-maintenance (a fixed energy cost is discounted at each time step) and for reproducing once a certain amount has been accumulated. Selection of mating partner is based only on locality constraints and the offspring is placed within the neighborhood of the first parent (the one whose high energy level triggered the reproductive event) who also provides the offspring with its initial value of internal energy. Neighborhoods do not overlap completely so that gene flow is not prevented. The offspring's genotype results from a uniform crossover operation on the genotypes of the parents plus random point mutations with a fixed probability $\mu = 0.002$ per loci (which in the present case means roughly 3 mutants every 200 newborn agents).

At the beginning of the simulation the 256 developmental matrices (M^{P_b}) are defined and an initial population (generally about 800) with randomly chosen genotypes is distributed at random in the arena. The simulation is run for a fixed number of time steps. As before, many parameters, especially those linked to energy values, are chosen so as to avoid extinction and maintain a reasonably high average number of players (about 700). These parameters are fixed for all simulations.

7.4 No development

Based on previous experience with similar analytical and individual-based models one already knows how to expect the model to behave if no social development is included.

7.4.1 No Space.

Since the game is formally equivalent to the 4-strategy action-response game studied using game-theoretic approaches in chapter 5, all the conclusions made in that chapter apply. In general, two different attractors could be expected. A fixed point attractor where the population is constituted in equal proportions by strategies (EE) and (OO) and a cyclic attractor where each strategy becomes dominant for a finite period and then is replaced by another one in the order shown by the straight arrows in figure 7.4.

In the current model, the lack of spatial relationships can be implemented by setting the local neighborhoods to a size greater than that of the whole arena. Thus, each agent has equal access to any other agent and the situation is that of a mixed medium. Figure 7.5 shows the population constitution as a function of time where it is clear that the oscillatory solution dominates the dynamics.

7.4.2 Space.

As in chapter 6, the introduction of space and stochasticity radically changes the outcome of the evolutionary dynamics. In a way practically identical to the model presented in that chapter, despite differences in parameters and the shape of the local neighborhoods, the results of the spatial game showed that for certain values of c the level of cooperative coordination remains stable above the baseline level of 50 %. As c is increased the overall level of coordination decreases in a quasi-linear fashion, from almost 100 % for $c = 0.5$ to the baseline level for $c \cong 0.7$ (see figure 7.8).

An inspection of the spatial distribution of the players shows that they form stable clusters which present the same phenomenology as those discussed in chapter 6. This suggests that the




Figure 7.5: Population constitution, no space and no development.

same type of processes are at play in this model and, consequently, that the stabilization of cooperative coordination can be explained in a similar way.

It has been observed that populations tend to be constituted largely by a mixture of strategy (EO) and a minor amount of (OO) or, conversely, by (OE) and a minor amount of (EE). Figure 7.6 (a) shows an example of such a case, for a typical run of 1.5×10^5 time steps (roughly 400 generations). The population is rapidly invaded by a self-cooperating strategy (OE) and a small number of agents of type (EE) is maintained since they take advantage of the (OE) strategy. In contrast with figure 7.5 the spatial dynamics prevents an invasion of strategy (EE), (other strategies disappeared early in the simulation). Figure 7.6 (b) shows, for the same run, the phenotypic constitution of the population through time. Since there is not developmental process, adult phenotypes are identical to phenotypes at birth, and they are completely determined by genetic constitution. In this figure, the horizontal axis represents the phenotypic number and the vertical axis represents time (running downwards). Darker areas represent a higher number of the corresponding phenotype¹³.

As c increases beyond 0.6 this two-strategy constitution ceases to be stable and all strategies

¹³This figure does not reveal any apparent structure in phenotypic space other than the fact that the greater part of population is located in a “cloud” surrounded by smaller “clouds” with shorter life span. Since most of the transformations in phenotypic space are effectively neutral (remember that only 2 of the 8 bits in the genome are relevant for defining the behavioral strategy), one could ask why the population constitution does not perform something like a “wandering” movement in phenotypic space rather than remaining in the main “cloud”. This is not one of the issues treated in this thesis but, nevertheless, the hypothesis can be advanced that such “wandering” movement would effectively be observed if the only operation in genetic space were mutation. The present model includes mutation *and* recombination in a finite population. Thus, while many mutations would produce neutral phenotypes, in order to effectively increase the number of players bearing those new phenotypes, a threshold concentration would have to be achieved. Such concentration is difficult to achieve if most of the reproductive events in which the new players will engage will be with partners drawn from the main “cloud”, which means that the resulting offspring will be more likely to be of a different phenotype than if the only genetic operation were point mutations. On the subject of the role of neutrality in evolution see for instance (Huynen, Stadler, & Fontana, 1996; Barnett, 1997, 1998) and literature cited therein. A related issue concerning the role of crossover in allopatric speciation without environmental barriers has been discussed in (Di Paolo, 1996a).




Figure 7.6: Population constitution for spatial game with no development, $c = 0.6$.

participate in a strong dynamics.

7.5 Results with social development

When the evolutionary process incorporates development under social conditions some interesting qualitative differences become apparent. All simulations were run using the same parameters and for values of c between 0.6 and 0.95.

7.5.1 No Space.

As concluded in chapter 6, the imposition of spatial constraints changes the evolution of cooperation from an oscillatory pattern into a stable level. Is it possible that a process of social development could, by itself, constrain the evolutionary dynamics in a similar way? The answer is negative. The dynamics remain oscillatory, although qualitatively different in that strategies become phase-locked by pairs (figure 7.7).

7.5.2 Space.

When a spatial dimension is added, results at the population level are similar to the case without development except that the range of c 's where cooperation is stabilized is larger: $0.5 \leq c \leq 0.8$. More interestingly, almost for the totality of this interval, the resulting population is constituted *purely* by a single cooperative strategy [either (OE) or (EO)], and only for $c \cong 0.8$ do polymorphic populations appear. In contrast with the case without development, the global level of cooperative coordination does not decrease linearly as c increases, but is maintained at a level close to 100 % and then changes rather abruptly for $c \cong 0.8$. In figure 7.8 one can compare the variation of the level of coordination with c both for spatial games with and without development, (each point is taken as the average of the long term coordination levels of 5 simulation runs). Such a qualitative difference is an indicator that the process of socially influenced development is introducing an




Figure 7.7: Population constitution, no space with development.

additional constraining factor in the evolutionary dynamics. To reveal this factor the situation should be analyzed in more detail.

7.6 Networks of social influences on development

By studying the population constitution it is found that, within the range of c 's where cooperation is stable, the population tends to be constituted mainly by a single adult phenotype either of type (EO) or (OE). This can be observed in figure 7.9. In this figure the population constitution is shown for a typical simulation run with $c = 0.6$, (compare with figure 7.6). After a transient the whole population is constituted by strategy (OE) and the other strategies disappear almost entirely [figure 7.9 (a)]. More remarkably, figure 7.9 (b) shows that the population is constituted mainly by a *single* adult phenotype: $P_a = 176$. In contrast with the case without development, where for certain values of c the competing strategy was prevented from invading, when development is added it seems that the competing strategy is prevented even from existing.

By examining the genetic constitution of the agents it is discovered that the whole population has converged (save for small fluctuations) on a single genotype (or equivalently a single phenotype at birth: $P_b = 174$). A study of the developmental table for this phenotype at birth (table 7.2) reveals that adult phenotype 176 not only is the most abundant in the table but it also has the effect that if both parents have this adult phenotype their offspring (except in the case of mutations) will also develop into the same adult phenotype. The situation is depicted in figure 7.10 where the number in the square represents the phenotype at birth, the number in the oval the adult phenotype, the solid arrow can be read as “develops into” and the dashed arrow as “influences”.

A set of developmental relationships which somehow perpetuates itself will be called a *developmental network*. It is so distinguishable precisely because its organization remains invariant in time. In the particular case shown in figure 7.10 it is apparent that this network does not just per-




Figure 7.8: Coordination level vs. c for spatial games with and without development.

Table 7.2: Developmental matrix for $P_b = 174$.

	172	173	174	175	176
172	176	175	176	172	172
173	175	172	175	176	176
174	176	175	172	173	173
175	172	176	173	173	173
176	172	176	173	173	176

petuates itself but it may also be called “self-promoting” since the product plays a “self-catalytic” role in its own (future) production. As will be shown later, other type of networks can be obtained by increasing the value of c . But, for the moment, one could ask if it is possible to use this self-promoting network to account for the observed results.

An explanation is required for 1) why the presence of these networks prevents the existence of competing phenotypes for a certain range of c ’s, and 2) why the transition to no cooperation is abrupt rather than smooth as in the case without social influences on development.

The first question is easy to explain if the evolutionary process is momentarily imagined as a search for viable phenotypes. If the population is constituted in its totality by a self-cooperating strategy like (EO) then individuals using strategy (EE) will constitute good alternative solutions in the sense that they will get, in principle, more payoff than those using (EO). “Finding” this competing strategy is not a major problem for the case with no development, since it is “located” just one mutation away in genetic space (all that is needed is to flip an appropriate bit in the genome). In the spatial game, although individuals bearing this mutant strategy can easily be found, the corresponding invasion is prevented by the spatial dynamics, as explained in chapter 6. But the inclusion of socially influenced development changes the situation. Now not only must the mutant be easily found by simple genetic operations, but it also must comply with the logic of the dominant developmental network.




Figure 7.9: Population constitution for spatial game with development, $c = 0.6$.




Figure 7.10: A self-promoting developmental network corresponding to the same simulation run as in figure 7.9.

In the example shown in figure 7.10, suppose that a mutation changes the phenotype at birth from 174 to another value P_b . This value can fall into one of two cases: either its distance to the parental phenotype (176) is less than or equal to D or it is greater than D . In both cases one must look for the probability that the new adult phenotype will be able to perpetuate itself by generating offspring who will also result in the same adult phenotype. An estimation is made of the upper bound for the probability of this happening in each case.

If the mutant phenotype at birth cannot be affected by the parental phenotype because its phenotypic distance is greater than D , the corresponding adult phenotype will be equal to the phenotype at birth, (section 7.2). Supposing that this mutant is able to produce an offspring with the same genetic constitution, the adult phenotype of this offspring will be calculated as the entry in the developmental matrix which corresponds to the phenotype at birth of the original mutant.

According to the developmental rule, if new phenotype at birth is Q the adult phenotype of any future offspring, $M^Q(Q, Q)$, should also be equal to Q in order to compete by establishing a self-promoting network of its own. The probability of finding a matrix with this property is $1/5$ (for $D = 2$ each entry is filled by one of five possible numbers). The probability of the resulting adult phenotype being one that is able to compete with the dominant one is $1/4$ (the four behavioral strategies are equally distributed along the 256 phenotypes). And the probability of picking up a genotype at random and that this genotype is within *one* mutation from the dominant one is less than $8/256$ since some of these mutations will map into phenotypes which are closer than D . So the probability of finding a mutation that will produce a competing mutant able to develop into

itself and maintain this developmental structure in a self-promoting way is less than 1.562×10^{-3} .

The other case is a bit more complex. If the new phenotype at birth may be affected by its parents, then it has to develop into a competing adult phenotype. But it must do so both when the parents belong to the dominant structure, and when they belong to the mutant population. Suppose the new phenotype at birth is 175, which happens to develop into 177 if the parents are both 176. As the number of agents with adult phenotype 177 increases, many newborn agents will have one or both of their parents bearing this new phenotype. So, in order to compete with the dominant network, it is also required that phenotype at birth 175 and adult phenotype 177 establish a self-promoting network of their own. Again, it is possible to look for an upper bound for the chances of this happening. In the example, the condition that must be met is that the entries in the matrix corresponding to 175 for indexes (176,176), (176,177), (177,176) and (177,177) all be equal to 177. These are actually three conditions since the matrix is symmetric, so the chance of this happening in the general case is $1/125$. Again, only one fourth of these cases will map into a competing phenotype, and the independent probability of one of those being accessible by one mutation is, as before, less than $8/256$. An upper bound for the probability is then 6.2×10^{-5} .

It is clear that the existence of a self-promoting developmental network poses much harder constraints for evolution to find viable alternatives. These are not impossible to find, but much more rare. In a very concrete way social development introduces changes in the phenotypic landscape, so that even if two phenotypes are behaviorally neutral, they are no longer evolutionary so, simply because they are not developmentally neutral. And this explains why, in the presence of a self-promoting network of the type shown in figure 7.10 other phenotypes are prevented from existing.

It remains to be explained what happens when the value of c is increased up to the point where coordination is lost. Intuitively, it is possible to think that the selection “pressure” for finding the viable alternatives increases as c is incremented, but it still must be the case that, *if* a self-promoting network exists, then the competing phenotypes must comply to the developmental logic as imposed by that network. For high values of c these networks have not been observed and so it is concluded that they are too costly, in behavioral terms, to exist. Nor have any evident alternative developmental networks been observed in this case. This explains the sudden drop in coordination for values of c higher than the critical value of 0.8. Interestingly, although the drop is quite abrupt it is still possible to find a certain level of coordination above the baseline for values near 0.8, [figure 7.8 (b)]. This case will be examined in more detail.

What is observed as the value of $c = 0.8$ is approached from below, is that the increased “pressure” to find alternative phenotypes in the presence of a dominant self-promoting structure results in complex co-existing developmental networks with interesting properties. Figure 7.11 (a) shows the constitution of the population for such a value of c in a particular simulation run. Similar cases were observed for different runs. The figure shows the main cooperating strategy constituting most of the population (solid line). In figure 7.11 (b) the variation of the population constitution is shown. Besides the single main phenotypic component as in figure 7.9 (b), other phenotypes can also be seen, (these are seen to the borders of the figure, it must be remembered that the space is wrapped around and 255 and 0 are neighbour phenotypes). The main phenotype is associated with the self-promoting developmental network shown in figure 7.12 (a). In addition,




Figure 7.11: Population constitution for spatial game with development, $c = 0.8$

by examining the population, the discover is made that some agents relate through a more complex developmental structure [figure 7.12 (b)].

Both networks shown in figure 7.12 originate from three different genotypes. The Hamming distance between ‘238’ and ‘254’, and between ‘254’ and ‘255’ is 1 meaning that uniform crossover between these pairs will not introduce new genotypes. However, the two developmental structures are separate since developmental influences exist only if adult and newborn phenotypes differ in less than $D = 2$. From the behavioral perspective it is seen that two of the adult phenotypes produced in the second structure are of type (OO) which, under ideal circumstances, would tend to invade a population constituted mainly of type (EO). However, as can be seen in figure 7.12 (b), these adult phenotypes are not “free” to grow because they are constrained by the amount of adult phenotypes that *favour* their development, some of which *compete* with them on the behavioral plane. In spite of what figure 7.12 (b) may suggest all the phenotypes in this developmental structure remain bounded in number. This may be easier to understand by considering the cooperative relationships between these phenotypes also as a network (figure 7.13).

In fact, a study of the time evolution of cooperating and non-cooperating phenotypes in the secondary network [figure 7.12 (b)] shows a very strong correlation in their variation as if the overall effect tended towards a form of behavioral neutrality which suggests that some complex self-regulation is taking place. This time evolution is shown in figure 7.14. The dotted line indicates the number of players using non-cooperative strategies (phenotypes 0, 253, 255) and the solid line the number of players using cooperative strategies (phenotypes 1, 252, 254). The linear correlation between the two time series has a coefficient of 0.8841.

7.7 Discussion

As said earlier, the purpose of this study is mainly exploratory, hence the abstractions made in this early model. The relevance of social development for understanding the evolution of cooperative coordination could well be supported independently by theoretical arguments. A choice was made

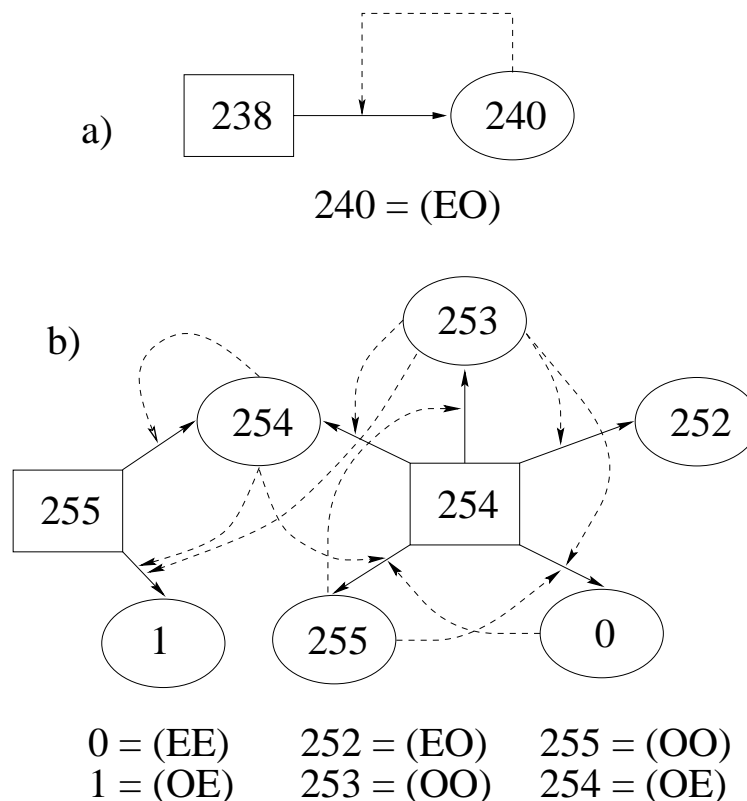


Figure 7.12: Developmental structures (see text) and corresponding behavioral types for a simulation with $c = 0.8$. a) Main self-promoting structure, b) secondary structure

to propose a minimal model where the nature of such influence could be explored.

In addition to the stabilization provided by spatial organization, cooperative coordination receives further stabilization by the presence of developmental networks which constrain the range of viable competing strategies and act as organizational invariants.

Further understanding of the dynamics of these structures could be obtained from other studies using reaction-networks. By the introduction of a chemical analogy, it is possible to use notions such as catalysis, hypercycles, etc., (Eigen & Schuster, 1979), whereby developmental relations can be seen as favouring the production of certain group of phenotypes in a self-sustaining manner. The work of Fontana and colleagues on artificial chemistries (Fontana et al., 1994; Fontana & Buss, 1994a, 1994b, 1996) could be particularly relevant in developing this analogy. In this line of work the dynamical organization of randomly reacting artificial objects is investigated in a minimal chemistry paradigm using λ -calculus expressions as the structure (and functionality) of the reactants. Reactions follow the logic of deterministic construction of new entities from the reacting elements and their subsequent substitution. Starting from random initial conditions the whole system converges after a number of interaction into a set of λ -expressions which, after application upon an identical expression, ends up producing a copy of itself. These self-copying functions are analogous to the self-promoting developmental networks that emerge in the present model. Due to its simplicity, Fontana and colleagues have labeled this kind of invariant regularities a Level 0 organization. By the introduction of syntactical constraints the action of copy functions can be restricted and more complex self-maintaining organizations emerge (Level 1) which in




Figure 7.13: Cooperative relationships in the developmental structures shown in figure 7.12.

certain aspects resemble the developmental networks that obtain when c is incremented to the critical value for which coordination begins to disappear.

However, the link between the two types of invariant structures is, for the moment, purely phenomenological, and some care should be taken in drawing hasty conclusions as the differences are also apparent, especially in terms of the other two “dimensions” in which the developmental networks of the present model exist: the behavioral and the genetic. One can have an idea of the added complications by considering the complex relations between the same set of phenotypes exposed by figures 7.12 and 7.13 in the developmental and behavioral planes¹⁴.

Another interesting aspect of the model in relation to issues mentioned at the beginning of the chapter is that the self-promoting developmental networks that dominate the dynamics for a large set of parameter values happen to look, within the restriction of the model, a lot like cases of imitative learning or behavioral transmission. Although none of the minimal cognitive mechanisms that would seem to be necessary in order to speak of such a phenomenon are present in the agents, at least from an ecological perspective it seems that of all the ways that parents can affect their offspring phenotype, those that most resemble teaching or imitation, are the most likely to be encountered. Such an observation should be taken as a warning when the intention is to try to understand the evolutionary role of social learning or imitative learning. It may not be directly related to selected advantages to individual organisms at all, but it may impose its own phenotypic significance simply by propagating more efficiently than other possible developmental relations. This would explain the apparent maladaptiveness that can be derived as a consequence

¹⁴Interestingly, in the case where only self-promoting structures like that shown in figure 7.10 exists, the corresponding network of cooperating relations is also rather simple since the adult phenotype happens to be of the self-cooperating type, [(EO) or (OE)], and so the analogy can be directly applied. But this may not happen in more complex cases.




Figure 7.14: Time variation for cooperative and non-cooperative strategies in secondary developmental network.

of other formal models such as Laland's model on sexual imprinting (Laland, 1994) and the model on maternal effects by Kirkpatrick and Lande (Kirkpatrick & Lande, 1989, 1992).

Finally, it would be important to emphasize that this study has looked at the issue of social development from an abstract perspective, that of an instantaneous process coupled with, but not completely governed by differential reproduction. This is a strong simplification of the issue of development under social conditions, especially if its behavioral and dynamical aspects become the focus of interest. Whatever the complexities of natural cases one can at least be certain that those features abstracted in the formulation of the model have relevant consequences for the study of cooperative action coordination and should be taken into account in future studies.

Chapter 8

Coordination without information

The model presented in this chapter is intended to act as a simple proof of concept against the preconception that information must be unevenly shared in order for communication to evolve. At the same time, it is meant to provide an example of the way in which a move can be made from the very simplified interpretation of coordination needed for studying some evolutionary and ecological issues towards a richer behavioral approach where other themes can be explored. This tendency will culminate in the strictly behavioral study presented in chapter 9.

As discussed in chapter 2, the preconception regarding information unavailability as a requisite for the evolution of communication arises from deriving an operational requirement from a functional *a posteriori* interpretation of the phenomenon in a selective context. In the argument presented in that chapter it was emphasized that such a derivation is not just equivocal from a theoretical perspective but that it also has specific pragmatic consequences in the form of a whole range of phenomena that is left unexplored because the models (in particular computational ones) are unnecessarily restricted by a conceptual bias.

The model of social coordination presented in this chapter will extend the models presented in chapters 5 and 6 and will not assume hidden information as those other models did. The intention is not to address specific scientific questions. Rather, the interest is in presenting it as a very simple example of the sort of issues that can be explored by relaxing a problematic assumption. For this reason, results will not be analysed with as much detail as the with other models in this thesis. At the same time, it is expected that this model will suggest other ways of extending the very basic action-response game used so far to games where the term “coordination” can be applied more in accordance to the meaning given in chapter 2.

8.1 Computer models of communication and the informational bias

“A first prerequisite for communication is that some organisms have access to information (knowledge) that others do not, for if they all have access to the same information, no communication is necessary. The non-shared information could be about the organism’s own internal state (e.g., hunger), or it could be about features of the external state of the environment that cannot be directly perceived, or as well perceived, by the other organisms”, (MacLennan & Burghardt, 1994, p. 165).

Such is the logic that, for quite some time, has been taken as a design dictum in many computer based explorations on the origins and evolution of communication. As argued in chapter 2, this view is the logical consequence of a monetary understanding of communication as currency exchange; a currency whose value is converted into fitness points and therefore renders communication tractable in simple cost-benefit terms¹. Information (understood in the colloquial or semantic sense) is the obvious candidate for such an unifying currency. In spite of the lack of clear definitions of the term, it must be recognized that this perspective on communication is attractive partly due to the abundance of imaginable scenarios that could fit naturally into an interpretation in terms of non-shared information. An agreement with the usefulness of this point of view should be evident from the way the models in chapters 5 and 6 were built and the situation represented therein interpreted. The criticism made in this chapter is not aimed at the specific situations presented in many computer models as being unrealistic, uninteresting or in any other way inaccurate. The criticism is aimed at the language used in justifying those models, a typical example of which is the above quotation. What this language manifests can only be considered, from the theoretical perspective of this thesis, as defining a conceptually and methodologically biased (rather than theoretically informed) research program².

How is this bias translated into computer simulations? The most common trend may be labeled “denotational”. This trend equates the evolution of communicative behaviors with the evolution of a shared code of tokens (usually plainly symbolic) that successfully map to environmental or internal states or the categories the agents construct from perceiving environmental or internal states. The corresponding models can be very straightforward: agents must simply converge into a conventionalised set of symbols to denote abstract states with predefined range of possibilities and which are privately accessible but must be somehow shared for effective action (MacLennan & Burghardt, 1994; Oliphant, 1996). Other models present more complex variations on this idea such as the more ecologically inspired model by Cangelosi & Parisi (1998), where agents must learn to classify mushrooms into poisonous or edible (a feature which can be only perceived proximately although the mushroom itself can be seen from afar) and pass on this information to other agents to save them the trouble of going for the wrong type of food. And a more complex and integral version is presented in studies by Luc Steels and colleagues where agents engage in language-games such as pointing to environmental features and naming them, and where words in a shared lexicon converge on perceptually grounded categories which can themselves be subject to change in a coupled process, (Steels, 1996, 1997a, 1997b, and others).

Some of the more biologically inspired models of this sort present an interesting potential from a research perspective on evolutionary issues, because they permit the study of questions regarding the benefits and costs of signalling, particularly in cases where the game posits a scenario where interests are in conflict; see for instance (Bullock, 1997; Noble, 1998a) and the models presented in chapters 5 and 6. Related to the latter, although not as clearly belonging to the “denotational trend”, are simulation studies of aggressive signalling such as those by de Bourcier and Wheeler (de Bourcier & Wheeler, 1994, 1997) and Noble (1998b). These models address issues regarding

¹See, for instance, Hasson (1994) and Maynard-Smith & Harper (1995) for definitions of information “augmentation” and “reduction” in terms of increased and decreased fitness. See also section 2.3.

²Clearly, this would not be so viewed from a theoretical standpoint which equates communication with information exchange and therefore rejects, or somehow manages to address, the broader theoretical issues raised in chapter 2.

the conditions for reliable signalling of aggression to evolve and become fixed in a population of agents. Wheeler and de Bourcier study territorial signalling in an ecological context where agents may have to compete for local resources. In its basic incarnation, the model shows that honest signalling of aggressive intent becomes stable only when the cost of signalling is high enough, thus supporting Zahavi's handicap principle (Zahavi, 1975, 1977) which states that reliability must be tied to signalling cost in order to make signals of high quality too costly for low quality individuals to produce (i.e. to bluff). Noble studies the evolution of intention movements: displays that could signal aggressive intent (such as the baring of fangs). None of these models assume a denotational view of signalling. However, the emphasis on signal *reliability* suggests that the type of communication that they are meant to address is indeed of the informational kind³.

In spite of the fact that an interest in whether signal reliability only makes sense from an information-centered point of view, in the studies on honest signalling under natural selection just mentioned one cannot speak of a *methodological* or *conceptual* bias in favour of a denotational perspective on communication since the possibility of *non-trivial* uninformative uses of signals, although rarely realized, is at least contemplated within these models⁴.

A different and less frequent trend is exemplified by the work of Werner and Dyer (1992). Here a scenario is constructed so as to avoid exerting direct pressure on a signalling system to evolve, (although whether the actual model is successful in achieving this is a matter of debate). A population of agents controlled by recurrent neural networks must evolve efficient mating strategies. Males can move blindly in a two-dimensional arena. Females can perceive nearby males and, although immobile, they are able to signal to the male who may choose to deviate on its course thus facilitating the meeting and reproduction event. After many generations the behavior of the males (initially finding mating partners at random) can be correlated with the signals emitted by the nearby females which can be interpreted as constituting a "directional code" which "tells" the blind males in which direction to move. Although, the results of this model could perhaps be interpreted in terms of behavioral coordination, it is clear that the authors choose to make an informational interpretation of the situation at the moment of designing their model. "To put evolutionary pressure on the animals to communicate, we needed to design animals in an environment such that some animals would have information that other animals needed to know but were not capable of finding out for themselves. The animals with this valuable information would have to communicate it to the other animals. The relevant information in this particular simulation is the location of the female animals relative to the male animals.", (Werner & Dyer, 1992, p. 661). Again, as in the denotational trend, non-shared information is taken as an a priori requisite for communication to evolve. Consequently, this model and others like it fall in the same category being criticized in this chapter.

Although the habit of assuming hidden information as a requisite for communication to evolve

³Unfakeable signals such as the aggressive roaring of the red deer fall trivially into a "denotational" perspective on animal communication and so do signals which are honest in virtue of Zahavi's handicap principle since these signals relate in a unique way to costs and, consequently, to quality (or need). Hence, the methodological necessity in these models to *allow* for non-denotative or dishonest signalling (although the honest-dishonest distinction remains informational in the sense discussed in chapter 2).

⁴For instance, both of Noble's models (1998a, 1998b) allow for the possibility of manipulative bluffing as well as receiver's counter-response to co-evolve in an arms-race fashion, although this is not actually observed. Consequently, these models cannot be said to have a denotational bias. However, as it has been argued in chapter 2, they may still rely on a much subtler informational perspective, that of a combination of manipulation and mind-reading. See section 2.3

is quite generalized in computational studies, there are fortunately a few exceptions.

In a recent model of the influences of sexual selection in song diversity Werner and Todd (1997), argue that linking symbols to states of the world ties the variability of signals. Werner and Todd connect the functional aspects of a signal (song) to the receptivity of this signal by a partner (and potential mate) whose taste is itself affected by evolution. Interestingly, in contrast with other models of sexual displays, in this one it is more difficult to speak of signals carrying information about the emitter since they are not linked with the emitter's quality nor do elaborate signals make the emitters incur additional costs which would put the game in a condition where Zahavi's handicap principle would apply. In such cases it would be possible to say that signals or displays carry information about the quality of their producers which the potential mate is in a position to assess.

In a different approach, Ficici and Pollack (1998) attempt to evolve communicative behaviors in a Pursuer-Evader game instantiated in a "bit-string space" rather than in the traditional 2-dimensional arena. At each time step, the evader generates a new string which the pursuer must simultaneously predict making use of a limited amount of history of the previous bits. Pursuers get scores for predicting the actions of the generator who get scores for making prediction difficult. Such a scheme may easily degenerate in the production of evaders who are extremely hard to predict, so a variation is introduced. Evaders are tested against the pursuers who they must evade, and also against "friendly" pursuers or partners for whom they must be easily predictable. The situation forces generators to be "predictable to someone". According to Ficici and Pollack, "[i]n terms of communication, this arrangement can be viewed as a competitive game between the two agent populations (generators and partners) trying to evolve a proprietary behavioral convention and the third agent population (pursuers) trying to crack the convention". This is actually observed using operational measures based on Shannon information. Although, this type of behavior does not fit exactly within the definition of coordination offered in chapter 2, it is certainly an interesting scenario with much potential for studying communication since a few simple extensions, like dropping the differentiation between the types of agents and allowing mutual observation and prediction, would probably put this model in a situation where the systemic definition would apply.

8.2 Coordination of recursive actions

As mentioned earlier, in the model presented in chapters 5 and 6 an interpretation was given in terms of information unavailability in order to make a link between the game played by the agents and a plausible natural situation. The player in the first role is able to perceive a piece of food of which there may be different types requiring different actions. The player in the second role is not able to perceive the food type but must act upon the external manifestation of the action of the first player. Although such a way of interpreting the game was not essential for the results obtained, as is shown in the fact that it is possible to remove all differences between food sources in the formally equivalent version presented in chapter 7, one can be certain that the metaphor guided many of the choices made while building the model. In chapter 7 the game is interpreted directly in terms of coordination as certain actions are "naturally" coupled (un-coupled) with certain responses. Here there is no information about food types (food sources are "neutral" for the outcome of the game). It could be argued that the action performed by the first player carries information about

its strategy but since the second player's response is fixed and does not depend on the first player's action, this information has no receiver.

It can be concluded that the game presented in chapter 7 does not assume information unavailability. However, blind one-step action coordination, as presented in this model, does not reflect the idea of ongoing mutual orientation of behaviors that is implied by the view of communication as arising from the activity in a consensual domain. Despite the insights that this model provides from an evolutionary and ecological point of view, it can be said that the behavioral aspects of communication are too simplified to make this a very strong case as an example of the superfluosity of information unavailability.

The game of chapters 5 and 6 can be extended in order to *approach* a situation in which it would seem appropriate to speak of such a consensual domain. Strictly speaking, this will not be possible as long as the structure of the agents remains non-plastic. Another issue that should arouse suspicions is the fact that (both in those models and in the extension proposed below) the patterns of interaction possess a pre-defined temporal structure which already implies the existence of a lower layer of coordinated activity⁵. These issues, along with the question of autonomy, will be explored in chapter 9. For the moment, the following modifications to the game will take it nearer a situation where one can speak of coordination in a less artificial way.

A return is now made to the case in which there are different types of food sources. However, instead of requiring a unique action by each player, access to energy in the food sources requires the performance of a specific sequence of alternating actions by both agents. Energy is released partially depending on the actions being correct at the required steps of the sequence.

In this case the actions of *both* players depend on the perceived food type and the perceived external manifestation of the co-participant in the previous step. This means that the second agent also has access to information about the food type. If, for instance, food type "F" requires the sequence of actions "A, B, C, D" the first player is expected to produce action "A" and a signal, the second player to produce action "B" and a signal followed by the first player acting again and producing "C" and a signal and finally the second player producing "D". Table 8.1 shows an example of two agents successfully coordinating their behaviors in this game. In this table A_i denotes the "behavioral operator" of each agent: a function that for each pair (food type, external manifestation) returns a pair (effective component of action, external manifestation). The signals chosen for this table are arbitrary and the symbol \emptyset indicates an initially assumed external manifestation for the first move (see below).

Time	Sequence for food type "F"		
	$A_1("F", \emptyset) \rightarrow (A, \alpha)$	A	
↓		B	$(B, \beta) \leftarrow A_2("F", \alpha)$
	$A_1("F", \beta) \rightarrow (C, \delta)$	C	
		D	$(D, \gamma) \leftarrow A_2("F", \delta)$

Table 8.1: An example of dialogic coordination.

This game has a dialogic structure and it resembles the kind of interaction described by von

⁵Such is the case for the majority of simulation models on the evolution of communication (and many other areas). A few exceptions can be found in Noble (1998b), Saunders & Pollack (1996) and in chapter 9.




Figure 8.1: Examples of the evolution of dialogic coordinated activity for two values of c in two typical runs.

Foerster's examples of two machines interacting and reaching their eigen-behaviors (von Foerster, 1977, 1980). The performance of a whole correct sequence may be called *dialogic coordination*. The situation is somehow more complicated than the simple action-response game because not only must the external manifestations of one agent and effective components of the other become correlated, but also the external manifestations of both of them must become correlated through the whole sequence. Not only must the second player rely on the signal of the first player to act correctly, but this signal must also trigger in the second player the production of another signal that in turn will help in selecting the correct behavior in the first player for the next step in the game. This is a recursive task much more in accordance with the nature of communication as defined in chapter 2. As before, agents are modelled as state-less machines.

It can be seen that a description of this task in terms of semantic notions of information is useless. If by information it is meant information about features of the environment, these are equally accessible to both participants, if it is meant information about the changing state or intention of the agents, they have no changing state or intention. In spite of this, the coordinating task that is required from the agents is not trivial in view of how they are constituted.

A simulation was run with four different types of food, two of them requiring a sequence of actions such as "A, B, C, D" wherein the first player must perform "A, C" alternating with the second player who must perform "B, D" and the two other types requiring the sequence "B, A, D, C" which means that each agent must invert its own actions with respect to the role played. Payoffs are allocated after the first two actions, and then again after the last two in a manner similar to that described in the chapters 5 and 6. The amount of energy in the food source is divided in halves, one to be assigned to each subsequence of action-response. If the actions of the two agents correspond to the actions required by the first half of the sequence each agent gets an equal part of the first half of the energy, if only one agent performs a correct action it gets a proportion c of that half ($0.5 \leq c \leq 1$) and the other gets nothing. The process is repeated for the next half of the sequence using the other half of the energy.




Figure 8.2: Dialogic activity in the steady state for different values of c . Each point is the average of 5 simulation runs. The line represents a linear regression.

Exactly as in chapter 6, agents are distributed in space and select their partners and food sources from their local neighborhoods. Reproduction is also sexual and local. The spatial, energy and genetic parameters were kept equal to those used in chapter 6. For details, see section 6.1.

Figure 8.1 shows the evolution of the average instantaneous success in dialogue over time for two typical runs with different values of c . (Variations between runs with the same set of parameters were not qualitatively significant.) This quantity is similar to R_{12} in the chapter 6, the difference being that it is defined as the proportion of games in which the *whole* sequence of actions is performed correctly. A comparison with figure 6.2 reveals that (for the same set of parameters) the achievement of a high level of dialogic coordination is a much more difficult task than for the case of one-step coordination. After 2×10^5 time steps the level of dialogic activity is starting to stabilize while in the case of one-step coordination transients were of the order of 5×10^4 time steps.

Analogous considerations to those mentioned for the model in chapter 6 apply here as well. The level of dialogic success depends on the parameter c in a similar fashion as the level of coordinated activity did for the previous case. This dependency is shown in figure 8.2. As in the case of the level of coordination in chapter 6, dialogic activity decreases linearly with c reaching a baseline level for $c \cong 0.65$. The spatial organization of the population is also similar. Agents form quasi-stable clusters with a same internal structure and these clusters can differ in their own levels of dialogic success in a manner positively correlated with their size. All this evidence suggests that the evolution of dialogic activity may be explained in a similar manner as the evolution of coordination in chapter 6.

It is possible to investigate the resulting behaviors by analysing the evolved behavioral structures. For example, consider the evolved behavioral matrix within a highly successful cluster (above 90 % of dialogic activity) shown in table 8.2. As in chapter 6, these genetically encoded matrices define what action the agent will perform for a given food type and perceived signal. The matrix specifies which action and external manifestation will be produced by the agent for each

Food Type	Perceived Signal			
	α	β	δ	γ
“F0”	A γ	C α	C γ	B δ
“F1”	A δ	D γ	B γ	C β
“F2”	B β	A δ	D γ	C γ
“F3”	B δ	D γ	A β	C α

Table 8.2: Behavioral matrix of an evolved player of the recursive coordination game.

food type and perceived signal from the other agent. In this particular case, agents must produce the sequence “A, B, C, D” in the presence of food types “F0” and “F1” and the sequence “B, A, D, C” in the presence of food types “F2” and “F3”. The column corresponding to the signal α encodes the initiating action of the first player for each food type ($\emptyset = \alpha$). The food type and α are the starting point of the interaction. In contrast with the game in chapter 6, the entire matrix is used and not just one column and one row (see section 6.1). This may explain why it takes longer for stable state of coordination to be achieved.

Assuming that the particular agent shown in table 8.2 belongs to a rather homogeneous cluster, as it does, it is possible to explore the result of the interaction with a structurally identical agent. It is easy to see that in such an interaction the agent will achieve an average level of dialogic coordination of 93.75 % just by counting for each food type the success levels that this agent would achieve in a dialogue with an identical agent starting for the entry determined by the food type and signal α and then moving to the entry determined by the new signal and the same food type. For instance, if the food type is “F3” the initiating action of the first player (corresponding to column α) is “B” and its signal is δ . For a second player of identical constitution perception of “F3” and signal δ will elicit the production of action “A” and signal β and so on. The actual level of dialogic coordination achieved in the simulation by this agent was in fact 96.2 %, probably because other agents in its neighborhood are slightly different, or because there is a relatively smaller number of food sources of type “F0” surrounding it. (Success when the food is of type “F0” is only 75 %, note that the effective component “D” is missing in the corresponding row and this prevents the production of the whole required sequence.)

These results show that agents coordinating their activity are able to perform tasks beyond their individual capabilities. This particular game requires that both agents perform a sequence of different actions in the presence of an external environmental feature that remains unchanged over time. Given that agents are stateless machines, such a behavior is impossible for them to implement individually. However, pairs of interacting agents can achieve an important level of success in this task, each player taking advantage of the presence of the other, and “using” their interactions as the internal states they lack.

8.3 Discussion

The aim of this model has been to provide a simple counter-example. Contrary to what is commonly thought, it is not necessary that a game should present players with unequal access to information for such a game to give rise to an evolutionary situation in which some form of com-

munication may result. The type of dialogic coordination which evolves in this model for certain parameters cannot be readily associated with a conventional signal exchange since, from the traditional perspective, signals must have a referent.

The problem with this kind of thinking was already addressed in chapter 2. If the perspective presented there is taken seriously, not only is dialogic coordination in the present model as good an example of an artificial communication system as other denotational models presented in computational studies, but one can also say that it is certainly nearer the idea of communication as behavioral coordination expressed in systemic terms.

Near, but not quite there yet. After all, behavioral coordination entails some sort of structural coupling between autonomous systems followed by structural changes in those systems which result in pure interactions eliciting coordinated behaviors. Agents in this game cannot perform such a task, since they are simple look-up tables. The sense in which they can be said to be acting autonomously (i.e. as operationally closed systems) is at best very restricted. However, it will be seen that it is possible to provide the necessary structural richness to build a model in which behavioral coordination does occur in a more strict sense.

Besides pointing to the possibility of exploring via computer models non-informational forms of coordination, this chapter has also highlighted, albeit in a simplistic way, one of the themes that motivates the study of social behavior from the point of view of cognitive science. Such is the case of the relation of logical implication between performance and competence which exists in most methodological individualist positions. Accordingly, an instance of a performance of certain behavior by an individual entails the capacity in that individual to instantiate such behavior (competence). This is not necessarily the case if the behaviors are instantiated by individuals acting in coordination as is clearly demonstrated in this model by the capability of the agents to produce sequences of actions when presented with a static piece of food. The lack of internal states makes this task impossible at the individual level, but two agents in coordination may perform it. Clearly, the simplicity of this example means that it should be taken as an illustration rather than as actual evidence for the possibility of performance to exist without individual competence. However, the illustration is suggestive and more serious evidence will be presented in the following chapter.

Chapter 9

Rhythm, entrainment and congruence in acoustically coupled agents

Up to this point, the studies presented in this thesis have explored the role of different kinds of ecological situatedness in the context of the evolution of certain modes of social coordination. Accordingly, it has been shown that when questions about the evolution of social coordination are not formulated in a generalized, unsituated manner, but are instead put into the context of an ecological and historical situation, the conclusions can differ radically from what is initially expected. This is hardly surprising in the light of the recent general realization in the fields of cognitive science, Artificial Intelligence and Artificial Life, that adaptive behavior should be addressed, not in terms of generalized all-purpose features, but in terms of whole embodied agents inhabiting their specific environments. In view of this general trend, it is only natural that this thesis should also address that other aspect of situatedness which has been neglected in previous models: embodiment.

The conception that adaptive behavior in general (and social behavior by implication) has less to do with general cognitive mechanisms instantiated in some form of computational processing within the brain and more to do with the workings of a whole body, with its own plasticity, habits and manners of making its way in the world, is slowly, but firmly, gaining the recognition it deserves. This recognition is reflected in the behavior-based approach to robotics (Brooks, 1991), in the field of evolutionary robotics (Harvey, Husbands, Cliff, Thompson, & Jakobi, 1997, and others), in simulated approaches to the study of adaptive behavior in whole agents (Beer & Gallagher, 1992; Beer, 1995), in issues regarding cognitive development (Rutkowska, 1994; Thelen & Smith, 1994; Hendriks-Jansen, 1996) and also in contemporary philosophy of mind (Lakoff, 1986; Varela et al., 1991; Clark, 1997; Wheeler, 1995, 1996; Lemmen, 1998).

Most of the inspiration that drives these lines of thought have been around for quite some time. On the philosophical side one can mention the earlier works of Martin Heidegger (1962; Dreyfus, 1991; Wheeler, 1996) and the phenomenological (and metaphysical) approach of Maurice Merleau-Ponty (1962, 1968; Lemmen, 1998) as well as the works of John Dewey (1929a). Also, inspiration has been drawn from the works of W. Ross Ashby and W. Grey Walter as well as other cyberneticists (Ashby, 1960; Walter, 1953). Other currents, not so well established in the cognitive sciences so far, are also reflected in the fields of anthropological linguistics (Foley,

1997), and social anthropology (Bourdieu, 1977, 1990).

Considering this general background, the goal of this chapter is to present a model in which it will be possible to explore the value of the theoretical concepts introduced in chapter 2 in a direct manner as well as issues related to the role of embodiment in a concrete example. More specifically, the present chapter sets out to explore, by means of simulations, some of the coherent behavioral patterns that can arise in the behavior of embodied autonomous systems as a result of sustained interaction through an acoustic medium. It will be shown that this coherence can be explained in terms of the systemic concept of *structural congruence* (section 2.4.2), the attainment of which is to be expected under a variety of circumstances. This concept attempts to describe one important consequence of sustained coordination between plastic autonomous systems. As discussed in chapter 2, a process of coordination is, at the same time, a process of mutual selection of plastic changes in the structures of the participating systems in such a way that not only it is possible to witness a coherence in the domain of their behaviors, but also in the domain of their structures. If the result of structural coupling is coordination, then such structural congruence can indeed be deduced operationally, through the principle of structural determinism, from the coherence which is manifested at the behavioral level of description.

The choice of sound as a medium of interaction is not incidental. Acoustic interactions are conspicuous in nature and some of their physical features, as well as those corresponding to auditory perception, make this “channel” a particularly interesting one from the point of view of embodied adaptive behavior. For instance, organisms using sound as a means of achieving behavioral coordination must face nontrivial problems like distinguishing their own production from those of others and avoiding mutual interference of signals. As will be seen later, sound perception also affords an active component given by the effects of body orientation and movement on the actual sound perceived. Additionally, other forms of perceptual regulation may play an important role in behaviors which are constrained by the temporal features of the medium and by physiological restrictions. As an example, one may consider the regulation of perceptual gain when intense sound is emitted and attention must be directed toward other sounds of lesser intensity [for instance, in the case of echolocation in bats, (Sales & Pye, 1974; Dawkins, 1995)]. Consequently, due to its physical characteristics, acoustic coupling is a medium in which the richness of embodied adaptive behavior can be particularly highlighted¹.

The present approach differs from the models presented in previous chapters as well as from other computer models addressing the evolution of communication in that the focus of interest is on behavioral issues rather than on evolutionary ones. In order to make positive contributions to evolutionary questions those models must assume too much (in terms of how the interactions between individuals are structured) to be, at the same time, useful for understanding purely behavioral questions. Here, in contrast, the aim will be to make fewer assumptions about the nature of the interaction and the behavioral building blocks incorporated into the agents.

In the following section some basic physical aspects of sound as a channel of interaction are identified and section 9.2 presents an example of complex acoustic coordination in nature. Section

¹An interesting line of research dealing with acoustic interactions using robots has been the modelling of the behavior of female crickets in the presence of the songs produced by males, (see for instance Lund, Webb, & Hallam, 1997). So far this work has only been focused on reproducing the behavior of females in responding to songs (produced by *real* crickets) and not in song production by males, where issues like self-stimulation could arise. However, a detailed understanding of the mechanisms involved shows the importance and opportunistic nature of embodied perception.

9.3 discusses the scope and methods used for building and studying the model presented in section 9.4. The rest of the chapter presents and discusses the results obtained.

9.1 Acoustic coupling

The use of sound is conspicuous in the animal world where it is associated with a variety of behaviors, particularly, but not exclusively, with social behaviors. The following list describes some of the relevant physical characteristics of interactions that rely on an auditory channel. In the general case, these features should be taken as givens whenever these interactions are present.

Finite locality and directionality. Sound intensity decays with the square of the distance to the source and it does not linger after it has been produced. It is also affected by other factors like wind direction and speed as well as the acoustics and filtering characteristics of the environment. Perception of sound is therefore a reliable measure of proximity at the moment of perception. Due to this feature the most basic behaviors relying on (voluntary and involuntary) production of sounds are related to localization functions (detection of predator/prey or mates, etc.).

Localizability. Binaural perception allows for spatial discrimination of the sound source. The mechanisms involved are varied, but rely mainly on temporal and intensity differences between the sound perceived in each ear. Differences in time of arrival of an acoustic signal can be used to pinpoint its source if the duration is short. For continuous tones, this mechanism is only effective at low frequencies (less than 1400 Hz in humans), since it relies on a discrimination of the wave phase (see Rozenweig, 1954; King & Carlile, 1995; Kandel, Schwartz, & Jessell, 1991, p. 493). For higher frequencies this information becomes ambiguous as more than one cycle may occur within the distance that separates the ears. Another mechanism provides discrimination based on differences of intensities between the sound perceived by each ear. Here, the actual difference due to intensity decay of the sound wave is not as important as the shadowing effects of the relative angular position of the head with respect to the source. For short wavelengths the head casts a considerable shadow in the travelling sound waves. In this case, differences in intensity are much more accentuated if the body comes between the source and one of the ears. In humans this difference can reach up to 20 dB for a continuous tone of 6000 Hz (Feddersen, Sandel, Teas, & Jeffress, 1957). This mechanism facilitates active discrimination involving movement of the body as was first observed by Venturi's experiments in the 1790's (see Rozenweig, 1961)².

Sound affects many individuals at the same time. Although sound production can be directed, in the general case, sound is broadcast within its local range, and can affect more than one individual organism. Acoustic signals necessarily influence the originator unless specific mechanisms prevent this from happening (e.g., synchronized sensory inhibition in bats during emission periods in echolocation). The role of this latter feature is not to be downplayed for various reasons. A discrimination between externally and internally generated sound (when necessary) poses the non-trivial problem of distinguishing between self and non-self.

Sound is continuous. Another feature that should be considered as a given in acoustic interactions is the fact that sound signals are inherently continuous, although they can be made discrete

²More specific mechanisms, such as the auditory apparatus of the cricket, combine frequency discrimination with enhancement of intensity difference by means of a particular set of delays and filters, (cf. Lund et al., 1997, and literature cited therein).

by controlling the temporal structure of their production. In combination with the above, this feature introduces the problem of how interactions are structured temporally and with respect to the number of participants so as to avoid interference between simultaneous productions. One to one interaction and turn-taking patterns require a certain degree of existing behavioral coordination, which in some cases may result from acoustic interaction itself, but they may also involve different sorts of physical coupling like body movements, touch, direction of gaze, etc. Purely gestural interactions, in contrast, may take place with certain degree of simultaneity.

So far auditory interactions have not received the attention they deserve, especially from researchers in the adaptive behavior community engaged in the study of social behavior. For instance, many studies concerned with the evolution of communication (e.g. MacLennan & Burghardt, 1994; Werner & Dyer, 1992) already assume discreteness, turn-taking or some other structured regulation of participation as the basic substrate upon which communicative behaviors evolve. The same may be said of the other models presented in this thesis. However, to be fair, none of these models is aimed at addressing the continuous/discrete transition or the problem of how participants structure a pattern of interaction in time. Another common assumption is the one-to-one nature of interaction which, as mentioned before, is not a given in the physical properties of the sound channel. These assumptions have also been criticized by Saunders and Pollack (1996). They present a model of communication over continuous channels where many emitters can affect at the same time a given agent in varying degree depending on the corresponding distances. The physical features of their signalling channels are indeed inspired by acoustic interaction, however they explicitly exclude the effects of self-stimulation and its associated problems.

The model presented in this chapter does not intend to address all these issues. For instance, the issue of self-stimulation over a continuous channel will be addressed, but restricting interaction to the pairwise case which is in itself a severe limitation. The motive for this restriction, apart from some technical difficulties, is the (reasonable but not entirely justified) suspicion that this will be a useful preliminary step for understanding more complex models with many agents interacting simultaneously.

9.2 An example of coordination through acoustic interactions

In order to illustrate the concepts introduced in the preceding section and some of the ideas discussed in chapter 2, a brief description will be made of a natural case of social coordination and structural congruence via an acoustic channel³.

In many monogamous species of tropical birds singing is not limited to the male but both male and female sing together performing in some cases antiphonal duets, i.e. alternation of different note patterns (Farabaugh, 1982)⁴. Antiphonal duetting has been studied in a number of East African species, particularly in certain shrikes (*Laniarius*) (Thorpe, 1972; Hooker & Hooker,

³In humans such coordination is apparent. There is much evidence of coordination between speech and body movements of both the speaker and the listener, both in adults and infants (Condon & Sander, 1974; Condon, 1979; Port, Tajima, & Cummins, 1998). Even if a conversation is not performed via an acoustic channel, one still is in the presence of a case that raises most of the issues discussed, (including self-stimulation via kinaesthetic feedback in the case of sign-language). However, a different example may help to put certain distance from specific human behavior and avoid issues which, at this stage, are difficult to address using the present methodology (such as linguistic interactions).

⁴Duetting has also been observed in many monogamous primate species. See (Haimoff, 1986) for a description of the evolutionary convergence of many aspects of duetting in primates of different taxa.

1969; Wickler & Seibt, 1979). One of these species, the bou-bou shrike (*Laniarius Aethiopicus*) lives in dense forests and produces a flute-like sound. Duets are constituted by patterns of notes so precisely alternated that they can be confused with the performance of a single bird. Each bird has its own part and they are not interchangeable, although some rare records were made of birds that completed their partner's part when alone. Each pair has a variety of different patterns, some of which are exclusively their own. Performance of duets can serve both the purpose of localization and demarcation of territory in the dense foliage as well as maintenance of the pair bond (Wickler, 1980; Wickler & Seibt, 1980). Hooker and Hooker (1969) observe that there is no signal other than the production of the first note for the duet to start, and that the tendency to respond can be very strong, sometimes forcing the interruption of preening or a response through a beakful of live food. They also report the lack of observation of periods of "practising" in young shrikes, which suggests that the particular features of the duetting pattern are something birds acquire by interacting within the pair.

Duetting thus serves as an example of a type of acoustic interaction which not only requires a high degree of coordination in itself but can also coordinate other behaviors like approaching. The fact that specific pairs of duetting birds develop a repertoire that reflects their own particular history of interactions and, partly as a consequence of this, they will pair for life, can be taken as evidence of the role of acoustic coupling in the achievement of ontogenetic structural congruence.

9.3 Methods

The rest of the chapter describes a simulation model which will be helpful for exploring some aspects of the concepts discussed above. Inspired by the example of duetting birds, a simulation will be introduced in which fully embodied agents are able to interact through an acoustic channel. As mentioned earlier, duetting in shrikes can elicit localization behavior in dense foliage. Based on this observation, a model is proposed for studying how mobile agents which are unable to use visual clues can approach each other by the use of acoustic interaction (i.e. signal production and phonotactic behavior). As explained in section 9.1, there are some basic givens in this type of coupling that have to be assumed from the start, like continuity and locality of the sound channel.

Continuous time recurrent neural networks will be used as the internal architecture of the agents since it is desirable that they produce behaviors that are continuous in time and this type of network has proven adequate for generating such behaviors. In the spirit of reducing the initial set of assumptions, agents will not be built directly, although certain parameters (such as body size) will be specified as constant and given⁵. Instead, an evolutionary search algorithm will be used to explore the space of possible structures.

A word must be said about the use of a genetic algorithm (GA) in the context of this work. Since it is not the aim of the model to provide an evolutionary scenario wherein the behaviors under study are a plausible outcome, the application of evolutionary techniques will be restricted purely to the task of searching a complex design space. Therefore, no conclusion will be derived about the likelihood of evolutionary histories that may lead to the behaviors of interest. Examples

⁵An exploration of the influence of each of the many parameters in the model is not its purpose. The analysis will concentrate on some instances of behaviors which are viable under the imposed constraints. The significance for natural cases will be given by the validity of the choices made.

of natural organisms that interact through an acoustic channel are sufficiently abundant to provide evidence that such cases are not at all evolutionarily implausible. Instead, plausibility restrictions are imposed on the constraints that provide the context of the search process both in the form of the physical laws fed into the model and in the form of performance evaluators for viable structures/behaviors.

It is fair, then, to inquire about the reasons for using a GA instead of other equally, or even more, efficient search techniques. Given the number of successful cases where this method has been applied in recent years in both simulated agents and actual robots, and the “interface” between the two (see for example Beer & Gallagher, 1992; Harvey et al., 1997; Jakobi, 1997, respectively), one reason for using a GA may be attributed simply to the proven adequacy of the method for similar search tasks.

The study of the resulting behaviors will follow more traditional techniques of observation and analysis of interaction patterns and internal dynamics. In addition, perturbation and disruption of normal modes of behavior will also be used in order to try to understand how such behavior is integrated.

9.4 The model

9.4.1 Sound

Sound is modelled as an instantaneous, additive field of single frequency with time-varying intensity which decreases with the square of the distance from the source. The effects of time-delays will be explicitly ignored as well as the effects of differences in frequencies of sound production, (i.e. no Doppler effect, differential filtering etc.). This coarseness of modelling will fit with mechanism modelled for spatial discrimination described below.

9.4.2 Bodies

Each agent is modelled as a circular body of radius $R_0 = 4$ with two diametrically opposed motors and two sound sensors symmetrically placed at 45 degrees to the motors (see figure 9.1). The position of the sensors was chosen in order to introduce a back/front asymmetry (although which is which is not specified) as it is considered important to be able to evaluate angular effects and coordination of physical orientation. A sound organ is located at the center of the body and regulates the intensity of the sound produced by the agent. The motors can drive the agent backwards and forwards in a 2-D unstructured and unlimited plane. In this simple model each agent is thought of as a rigid body, small in size and having a very small mass, so that the motor output is directly the tangential velocity at the point of the body where the motor is located. The translational movement of the whole agent is calculated using the velocity of its center of mass (the vectorial average of the motor velocities) and the rotational movement by calculating the angular speed (the difference of the tangential velocities divided by the body diameter). There is no inertial resistance to either type of movement.

Agents move freely in the arena except when they collide with each other. Collisions are modelled as point elastic, i.e. no energy loss and no effect in the angular velocity of the bodies. While undergoing a collision, an agent may move in a direction which is not specified by its motor output but which corresponds to a displacement which conserves the momentum of the whole




Figure 9.1: Body of agents and paths of acoustic signals.

system. The bodies of both agents are taken as identical so that the result of an elastic collision is the instantaneous “exchange” of the vectors of velocity at the center of mass. However, due to the lack of inertia, agents recover control of their movement immediately after the collision. The body circumference is taken as frictionless so that the angular velocities do not change during collisions.

Since the task the agents must perform involves some sort of spatial discrimination this must be provided by the relative activity of the sensors. These are physically separated so that in general their activity will be influenced by different external intensities, however such a difference provides poor discrimination especially if background noise is added. A natural mechanism, mentioned in section 9.1, involves the attenuation of intensity as high-frequency sound is shadowed by the body. The degree of attenuation is linked to the angular position and movement of the agent except in the case of sound produced by itself. This “self-shadowing” mechanism is modelled as a linear attenuation without internal refraction which is proportional to the distance travelled by the signal within the body, D_{sh} . This distance is given by:

$$D_{sh} = D_{sen}(1 - A), \quad 0 \leq A < 1,$$

$$A = \frac{D^2 - R_0^2}{D_{sen}^2},$$

where D_{sen} is the distance between the source and the sensor and D is the distance between the source and the center of the body (figure 9.1). If $A \geq 1$, there is a direct line between source and sensor⁶ and $D_{sh} = 0$. The maximum value of D_{sh} is given when the sensor is directly opposed to the external source ($D_{sh} = 2R_0$). The attenuated signal is calculated by first calculating the intensity of the signal at the position of the sensor in the usual way (i.e. applying the inverse square law without attenuation) and then multiplying by an attenuating factor which goes linearly from 1 when $D_{sh} = 0$ to 0.1 when $D_{sh} = 2R_0$. The process is repeated for the other sensor.

The agent’s controller is composed of a network of four dynamical inter-neurons and an arrangement of sensors and effectors, each one controlled by one neuron. The inter-neuron network

⁶For $A = 1$ the sensor, the center of the body and the external source form a right-angled triangle.

is fully connected (including self-connections). Additionally, each inter-neuron receives one incoming synapse from each sensory node (of which there are only one for each auditory sensor)⁷. Each effector node (one for each motor, one for signal production and two gain regulators as detailed below) receives one incoming synapse from each inter-neuron. There are no direct connections between sensors and effectors. This kind of dynamical neural networks can serve as an adequate basis for a fully embodied operationally closed mechanism and so they are a good tool for studying adaptive behavior in simple autonomous agents as some successful cases have shown, (Beer & Gallagher, 1992; Beer, 1996, and others). This is especially so when dealing appropriately with time constraints becomes an essential part of adaptation (Beer, 1995; Harvey, 1996). Inter-neurons and effectors neurons obey the following law:

$$\tau_i \dot{y}_i = -y_i + \sum_j w_{ji} z_j; \quad z_j = \frac{1}{1 + \exp[-(y_j + b_j)]}.$$

while sensory neurons obey:

$$\tau_i \dot{y}_i = -y_i + I_i,$$

where, using terms derived from an analogy with real neurons, y_i represents the cell potential, τ_i the decay constant, b_i the bias, z_i the firing frequency, w_{ij} the strength of synaptic connection from node i to node j and I_i the degree of sensory perturbation on the sensory node (modelled here as an incoming current).

In similar models sensors can be directly regulated by their participation in the network dynamics (e.g. by incoming synapses). A decision was made not to model direct synapses from the inter-neuron network into the sensory neurons and instead an effector was added that directly regulates the sensory gain in a multiplicative way. Such regulation allows the agent to have the possibility of extra control on perception. The gain of effectors can be regulated as well. In all cases presented here only two regulating neurons were used, one for the gain of both auditory sensors (in a symmetric way), and another one for the gain of the sound production organ. In each sensor a transduction step occurs which transforms external stimuli into a degree of perturbation (current). Analogously for effectors, a transduction step transforms the neuron firing frequency into the effector's output. These transduction steps are simply modelled as linear scalings (i.e. a multiplication of the firing rate by a gain). In sensors and effectors with regulated gain, the activation of the regulating neuron defines at each time step the scaling value by transforming linearly its own firing rate (between 0 and 1) into the gain value (see next section for ranges).

In order to constrain the production of sound to a realistic behavior neurons are allowed to “burn up” if the cell potential exceeds certain limits due to intense stimulation. In sensory neurons this may happen in the presence of intense sounds. In real auditory systems the destruction of hair cells in the cochlea occurs for mechanical reasons and not due to intense incoming currents into the nerve cells. This is the reason for stressing that the meaning of I_i should be taken as the degree of perturbation or stimulation to the sensory cell and not necessarily as a literal current. The

⁷A different sensor type that suggests itself would be a touch sensor which would activate in the event of contact with another agent (or a wall, etc.). In the simulations reported here no use was made of such sensors. The main reason for this is because it is of interest to explore long-term behavior, even after the agents have successfully approached one another. It has been observed in simulations that the introduction of touch sensors can have as a consequence the termination of (translational) movement in the agents after the first collision (i.e. agents find each other, collide and then rotate on the spot in nearby positions).

neuronal structures which satisfy this “physiological” restriction can be seen as approaching natural cases where viable behavioral trajectories are characterized by a certain equilibrium between the autonomy of the nervous system and the autonomy of the individual cells. This mechanism also provides additional significance to the sort of interactions an autonomous agent may engage in. Evolved agents should be expected to “take care” regarding the intensity of their own sound production and/or use their sensory gain regulation accordingly.

9.4.3 Genetic algorithm

A form of rank based selection genetic algorithm has been used as a search technique with a fixed population of 90 agents evolving for up to 1000 generations, (after a few hundred generation the search converged). Each agent is selected an average of ten times (five guaranteed⁸) to play with a different agent in the population which is introduced in the arena at a random time after the first one. The aim of this delay is to avoid cases in which agents may undergo similar dynamics in an artifactual way simply because they start their operation at the same time and from similar initial conditions⁹. The second agent is placed at a random distance no smaller than 50 from the current position of the first agent. The initial state of the agents is reset at the start of each trial; the cell potential of each neuron y_i is set to a randomly chosen small value taken uniformly from the interval $[-0.1, 0.1]$.

Fitness values are averaged over all the trials. Fitness is allocated in terms of how much the agents have approached each other at the end of the run, ($F_A = 1 - D_{Final}/D_{Initial}$), and what proportion of the interaction time they have spent within a distance of 4 body radii of each other (F_D). Additionally, the weighted sum of these proportions is modulated by a mild exponential term that decreases with the integrated energy used. This was done in order to compensate for the lack of an adequate model of energy consumption. Agents making excessive use of motors are therefore penalized, although very mildly. The individual fitness F for a given trial (usually lasting 200 time steps¹⁰) can be expressed as:

$$F = (a_d F_D + a_a F_A) \exp(-a_e \int_{t_i}^{t_f} \sqrt{V_R^2 + V_L^2} dt),$$

where $a_a = 0.75$ and $a_d = 0.25$ are the weighting factors for the approaching and maintenance of proximity tasks respectively, $a_e = 0.005$ scales the modulation of the exponential and V_R and V_L are the translational speeds of the right and left motors respectively. A final factor affecting fitness is cell death (see above). If at the end of the trial run an agent has burnt up one of its sensory or effector neurons, its total fitness in that run is reduced to zero.

Numerical integration of the model is done using the Euler method. This is a second order method and therefore not very accurate, however it is fast enough to study many evolutionary runs. To compensate for the lack of accuracy the integration step is chosen in such a way that similar results are obtained (in one typical run) by using an order four Runge-Kutta method¹¹ with

⁸This is achieved selecting each agent deterministically to play five times with a randomly selected partner other than itself. On average, the same agent will be selected by a different one another five times.

⁹It must be remembered that after a few generations a large proportion of the population will have converged to similar structures.

¹⁰This number is determined by the choice of other parameters, such as motor gain and initial distance of separation, in order to make the approaching task possible.

¹¹See (Press et al., 1992) for description of these methods.




Figure 9.2: Trajectories of approaching agents, (a). Distance between agents as a function of time, (b). The second agent is introduced in the arena at $t = 10$.

a time step of half the minimum neuronal decay constant. The resulting integration steps for the Euler method used is 0.1 (compare with decay constants below).

All network parameters (weights, gains and biases) are encoded in a real-valued vector of fixed dimension. Each component has a value in the interval $[0,1]$. This value is scaled linearly to an appropriate interval in order to specify a parameter of the network. Transduction gains are chosen from the interval $[0.05,10]$, biases from $[-3,3]$, weights from $[-8,8]$ and decay constants from $[0.4,2]$. An agent with N inter-neurons and N_{SE} sensors/effectors would have a genome size of $(N + N_{SE})(N + 2) + N_{SE}$. Symmetry between left and right is enforced only for biases and gains but not for weights and some of the gain parameters are directly regulated by the agent (see above), so the resulting genome size is in fact less than the above quantity.

After ranking the population according to fitness the next generation is built by making 2 copies of each individual in the top third of the current population and one copy of each individual in the middle third. No crossover operator is used and mutation consists in perturbing the genome vector \mathbf{G} with probability $\mu = 0.005$ in a random direction by adding a normalized vector \mathbf{p} multiplied by a distance m chosen uniformly in the interval $[0,1]$: $\mathbf{G} \rightarrow \mathbf{G} + m\mathbf{p}$.

Uniform noise is added to all the transduction steps affecting sensors and effectors, (range = 0.1, mean = 0 or 0.05 if the transduction implies a non-negative value). These values of noise are also scaled by the gain of the corresponding sensor/effector.

9.5 Results

Highly fit agents evolve reliably after a few hundred generations. About 16 different evolutionary searches were run. In 11 of them a same qualitative type of behavioral strategy was observed with small quantitative variations. With less frequency, other behavioral strategies were also observed but were not studied thoroughly. The most frequent strategy was found to be interesting




Figure 9.3: Neuronal structure. N0 and N1: active inter-neurons, MR and ML: motors right and left, SR and SL: sensors right and left, SenG: sensory gain regulation, SIG: amplitude of emitted sound, SigG: gain regulation for sound effector.

enough to deserve a detailed examination. The reported results correspond to a single instance of this strategy, but the same conclusions apply in all 11 instances where the same behaviors were observed.

9.5.1 Approaching behavior

Agents successfully approach each other when two of them are in the same arena (figure 9.2 shows an example of the trajectories performed by the agents and the distance between them as a function of time). The structure of the evolved neural network can be seen in figure 9.3. Only 2 active inter-neurons are “used” by the agents; the other two have a constant activation of 1 or 0 (not shown). Examination of the sensory activation shows that the signal perceived by one agent at the moment when the other one is placed in the arena is very faint in comparison with the agent’s own signal production, and even in comparison with noise levels. Agents engage in a mode of search behavior that relies on the fact that self-shadowing is a mechanism that can be exploited actively by movement with a strong angular component, (figure 9.2). This helps to discriminate external sounds since perception of an agent’s own signal does not depend on the orientation. Sensors act as leaky integrators and connect differentially to the inter-neurons with similar absolute weights (figure 9.3, values not shown), which means that the basic strategy would seem to involve rotation while moving, integration of sensed intensities and evaluation of the difference between left and right sensors (i.e. filtering out of own contribution to perceived intensity since this contribution is

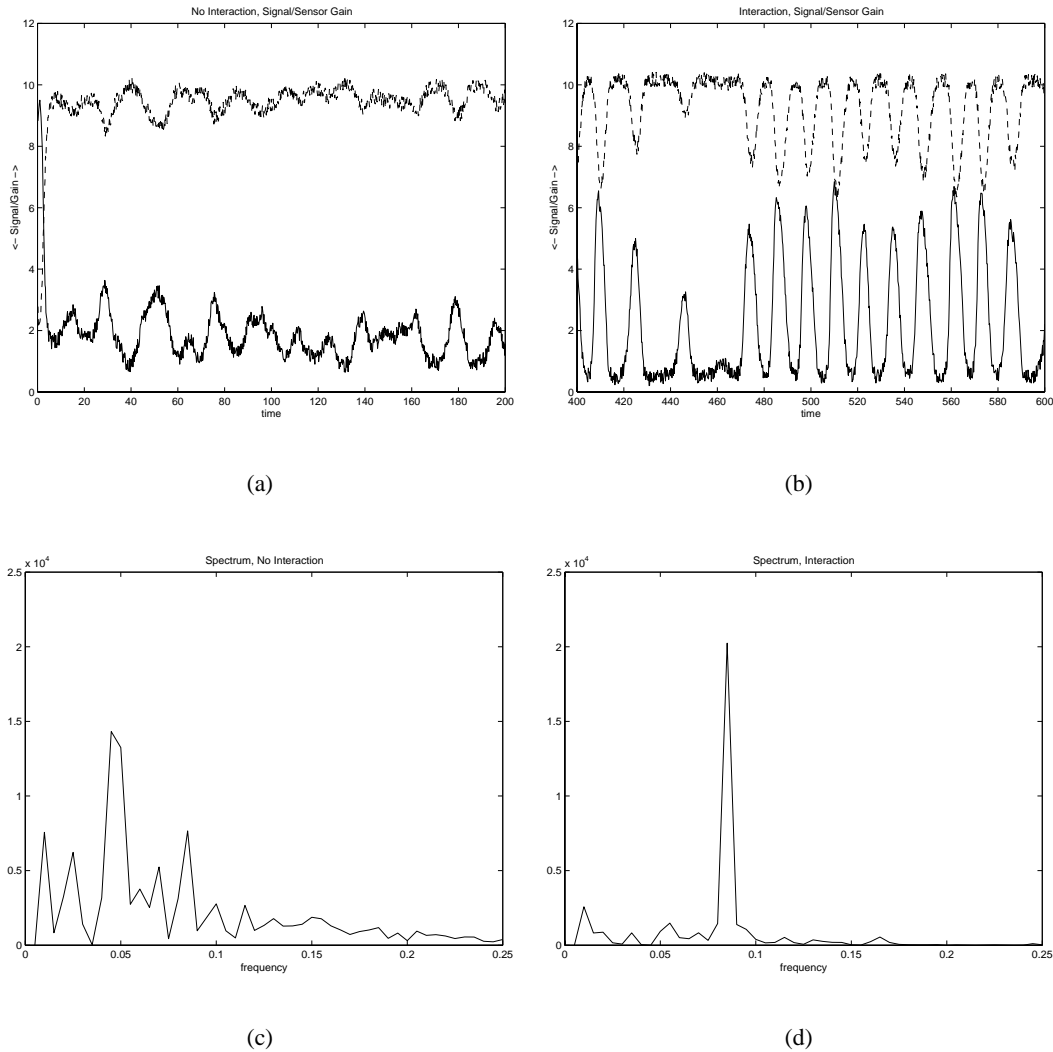


Figure 9.4: Emitted signal (full line) and sensory gain both for an agent that is interacting and an agent by itself (a and b). Fourier power spectra (c and d).

the same for both sensors). This is an efficient way of discriminating faint external signals¹².

However, this is not entirely right. It is observed that self-stimulation is also integrated into the production of movement, as is evident from the fact that if one reduces progressively the capacity to hear their own production, the behavior of the agents degenerates very rapidly into a rotation on the spot. Perturbation can be done by altering the degree of self-stimulation in absolute terms (multiplying the contribution to perceived intensity due to the agent's own signal by a factor between 0 and 1) or by introducing delays between own sound production and perception¹³. It

¹²Notice also that agents choose to place their sensors “on the back” with respect to the direction of movement. This cannot be explained in a straightforward manner, but it may be related to an enhancement of the active discrimination strategy at the initial stages. If both sensors were placed in the direction of the external source it would be unlikely that the self-shadowing mechanism will enter into action.

¹³The latter method has been used in humans to show the dependence of speech production on acoustic self-perception in subjects with normal hearing. Delayed hearing of own spoken words results in degradation of the speech produced. Deaf people can of course produce speech relying purely on kinaesthetic perception showing, therefore, no degradation. This method can be used to identify subjects who fake deafness, (von Békésy, 1972).




Figure 9.5: Fourier power spectra for signal (solid line), agent orientation relative to the other agent's position (dotted line) and difference of intensity at sensor position (dashed line).

is possible to conclude that the agents are not merely acting on external cues as was suggested in the previous paragraph, but the ability to hear themselves is also integrated into the rest of the behavior. From another point of view, this is also evidence that a functional characterization of signalling behavior as purely conveying information of position, or even of changes of position, is not possible, nor is it possible to decompose movement into active sensing and approaching.

Figure 9.4 shows the signal produced and the regulated value of the sensory gain both for an agent on its own and in interaction with another agent together with the corresponding Fourier power spectra¹⁴. As would be expected from the fact that sensors can “burn up” due to intense activation, when the agent is emitting an intense signal sensory gain is reduced. It is also apparent that signalling behavior has a marked rhythm when agents are interacting. What is the origin of this rhythm? It cannot rely entirely upon internal mechanisms since it does not appear when the agent is by itself [plots (a) and (c) in figure 9.4], although the corresponding power spectrum does indicate the presence of other frequencies.

Also from the analysis of frequency spectra it can be concluded that rhythm in signalling behavior is directly linked to angular behavior. This is achieved by comparing for one of the agents the frequency of its signal with the frequency of the variation in angular orientation relative to the line connecting both agents and finally with the frequency obtained from the “objective” difference of sound intensity at the position of the sensors, (figure 9.5)¹⁵. All three spectra show a sharp peak for the same value of frequency.

Additional evidence of a connection between signalling and angular movement is obtained

¹⁴All power spectra in this and the following figures are calculated by first normalizing the signal to a value between 0 and 1, then subtracting its mean value and calculating the square of the absolute value of the Fast Fourier Transform.

¹⁵Since the signal produced by the agent reaches its own sensors symmetrically, performing this difference in the value of intensity will also provide information about the angular movement of the agent relative to the external source.




Figure 9.6: Trajectory for an agent with a constant source with a fixed position relative to the agent (a). The source position is represented by a small circle on the side of the agent. Signal emitted by the agent, (b).

from the observation that if a source of sound is placed at a fixed distance and angular position with respect to a moving agent (i.e. movement has no influence on perceived intensities), all rhythm in signalling behavior disappears. This is shown in figure 9.6.

The previous evidence seems to suggest that rhythmic signals originate entirely as a consequence of the angular movement of the agents. However, one must be cautious with this conclusion since, as discussed above, movement is not independent of signal production. From the observations made, it could as well be argued that angular movement depends on rhythmic signalling and that rhythm in signals originates somehow within the dynamics of interaction. The first explanation seems more plausible but it has not been possible to rule out the second one so far. This difficulty in itself points to the fact that behavior is quite integrated and makes functional decomposition hard. It gets harder if one consider the global picture of both agents in interaction as will be shown next.

9.5.2 Entrainment, turn-taking and structural congruence

Since patterns of joint activity seem to be relevant for understanding the behavior of individual agents an analysis is now made of these patterns in cases of prolonged interactions (about 1000 time steps). Figure 9.7 shows the signalling behavior of two interacting agents after having approached one another. It is observed that for long periods signals seem to be phase-locked at some value near perfect anti-phase. Although agents are similar they are not identical and their “natural” power spectra [i.e. when acting on their own, figure 9.4 (c)] are indeed different. This suggests that the observed entrainment must be somehow related to the coupling between the agents. Figure 9.7 (a) (see also figure 9.12) shows that this entrainment can be lost momentarily only to be regained later. This phenomenon is similar to what has been called relative coordination: the




Figure 9.7: Signalling behavior of interacting agents and power spectra.

tendency to actively correct for phase-randomizing factors such as fluctuations or differences in natural behaviors or physical properties.

The anti-phase locking of signals can be interpreted as a basic form of turn-taking. Since agents have no other way of knowing of the presence of the other but through acoustic coupling an efficient way of doing this is by alternating the production of signals and so minimizing overlap.

Movement during periods of coordination is also highly organized (figure 9.8). Agents perform almost perfectly synchronized “dancing” patterns alternating their positions on the inside and outside of a curved trajectory while varying their angle of orientation in an almost identical manner. Agents collide slightly with each other on certain occasions and this also seems to be an ordered phenomenon (compare the smooth portions of trajectory on the top of figure 9.8 where the agents are not in contact with the portions at the bottom where collisions occur).

It is proposed that the highly ordered patterns shown by agents both in their movement and signalling behavior are evidence of the achievement of dynamical structural congruence through acoustic coupling. Coordination is therefore a coherent state that agents achieve jointly after undergoing specific transient paths in their respective dynamical trajectories in the presence of mutual perturbation. As much as coordination cannot be reduced to the behavior of a single agent, the specific ordered patterns observed during coordination cannot be explained by the activity of individual agents if they are indeed the consequence of the structural congruence attained between them. In principle this may seem strange, after all coordination could be thought of as the individual adaptation to the behavior of one’s partner instead of a co-adaptation.

In order to prove that coordination in these agents does not originate from an *individual* capacity for adapting an examination will be made of how agents behave in the presence of beacons that produce sound signals. A beacon is placed in the arena in a fixed position and one agent at a random angle, orientation and distance from it. Beacons can produce a variety of signals. In all cases agents approach the beacon successfully but their signalling behavior differs from the case of two interacting agents. Figure 9.9 shows this signalling behavior when the beacon produces a periodic signal with a period chosen to be equal to the one shown by coordinating agents.




Figure 9.8: Motion of agents during period of coordination. Agents are shown at two time steps.

If the agent had an individual capacity to adapt its signalling behavior to the sound it perceives one would expect it to show a similar signalling behavior in the presence of the beacon. It does not happen. It can be argued that the frequency used by the beacon is not exactly right. A long simulation in which a slow “frequency sweep” in the beacon’s signal is performed for the whole range of frequencies of interest shows no difference in the signalling behavior of the agent. The possibility still remains that the particular distribution of frequencies may matter. In order to test this, the following experiment is performed. A normal simulation with agents in interaction was run and the signalling behavior of one of the agents during coordination was saved. The “taped” signal was broadcast from the beacon to the other agent now by itself. The result (see figure 9.10) shows that phase locking does not occur.

All this evidence points to the importance of the presence of a history of mutual induction of changes in the dynamics of each agent. Beacons are completely non-plastic and therefore their “behavior” cannot be influenced by the approaching agent. The lack of a transient period of mutual triggering of changes of state results in no structural congruence and consequently in no entrainment. In other words, the process ceases to be historical in the sense proposed in chapter 3.

So far only some external manifestations of structural congruence have been considered. Figure 9.11 presents further evidence, this time from internal dynamics. This figure shows the embedded time-delayed plots for the activation of the same inter-neuron in two agents under different circumstances. The two plots at the bottom correspond to the agents acting on their own, i.e. no coupling. Here it is again observed that although their structures are similar (both neural networks present the same architecture shown in figure 9.3), they are not identical (parameter values differ slightly) and, consequently, their dynamics present different attractors. The top four plots show the same embedded dynamics, now for the agents in interaction. The two plots at the top are taken from a period of coordination. The striking similarity between the two attractors becomes evident as well as the difference between them and the respective natural dynamics. The plots at the center of the figure show agents interacting but during a period when coordination has been lost. These plots also show an interesting qualitative difference with respect to the other cases. Agents do




Figure 9.9: Signalling behavior in presence of periodic beacon.

not return to their “natural” dynamics when coordination is lost but to a different, uncoordinated state from which it is possible for coordination to be regained and which suggests that, during interaction, the structure of each agent has changed from its natural one.

In each case it is possible to calculate the time correlation in neuron activation between the agents:

$$\rho = \frac{Cov(\mathbf{N}_{A1}, \mathbf{N}_{A2})}{S(\mathbf{N}_{A1})S(\mathbf{N}_{A2})},$$

where $Cov(\mathbf{x}, \mathbf{y})$ is the covariance between variables \mathbf{x} and \mathbf{y} , $S(\mathbf{x})$ the standard deviation of \mathbf{x} and \mathbf{N}_{Ai} is the vector composed by the time series of activation values for the inter-neuron in agent i in the period of interest. In the case of coordination $\rho = -0.8443$, for the case of interaction but no coordination $\rho = -0.3750$, and for the non-interacting agents¹⁶ $\rho = -0.0403$. A strong anti-correlation between coordinating agents is in accordance with their signalling behavior.

The transition from coordinated towards un-coordinated state can be induced both by fluctuations (noise) or by an instability due to internal differences in the respective dynamics. However, the transition from un-coordinated state into coordination can only be understood in the presence of an organizing coupling between the two systems since fluctuations, on their own, will tend, on average, towards the loss of entrainment. Figure 9.12 shows how different variables are affected when these transitions occur. The four plots are taken from a simulation run with the agents interacting. All show certain variables as a function of time. The plot at the top shows the distance between the agents. Coordination periods are conspicuous since they present a much smaller range of variation in distance. During these periods agents perform the ordered patterns of movement shown in figure 9.8. These regions are marked “b” and “d”. In contrast, periods of no coordination (regions “a”, “c” and “e”) show greater variation in distance. The second plot shows the relative angular orientation between the agents which remains near zero degrees during coordination and is uncorrelated the rest of the time. The third plot shows signalling behavior and the plot at the bottom shows a continuous estimation of the relative phase between signals, (the horizontal lines

¹⁶This last value is only illustrative of the different “natural” behavior of the agents.




Figure 9.10: Signalling behavior in presence of imitative beacon.

mark 180 and -180 degrees)¹⁷. During coordination and even a bit before (compare the beginning of region “d” in the top two plots with the same area in the other plots) signals are produced near the anti-phase region. The fact that this correlation is manifested before the attainment of minimum distance suggests that signal coordination may come first and coordination of movement may be its consequence. This makes sense if one thinks that signals can begin to be coordinated from a certain distance greater than the minimum.

9.6 Discussion

The basic model shows some interesting phenomena like turn-taking and organized movement arising from basic features of the physical nature of acoustic coupling between embodied agents. In interpreting how these phenomena arise, use has been made of theoretical concepts concerning social behavior under a systemic, operational perspective (as presented in chapter 2) as well as other concepts taken from the dynamics of coupled oscillators. It is thought that much is to be gained from this perspective and that the present example has not exhausted its potential.

The model shows how agents interact acoustically and through this interaction they mutually coordinate their patterns of movement and they structure their acoustic coupling into a form of alternated production that resembles turn-taking. It is not inconceivable that this result could have been different, (for instance, fairly constant signalling behavior, totally decoupled from movement and perception). Initially one may be surprised that agents organize their interactions in the way they do but it has been shown that there is nothing magical about this organization if the process is analysed operationally.

However, providing a thorough operational explanation can be hard. Apart from the potential complexity of such an account, one of the main difficulties lies in the fact that many operational aspects of the system act concurrently, so it is not always possible to speak in terms of causality,

¹⁷This estimation is obtained by a continuous normalization of the signal and its derivative to the unit circle after filtering out the noise and then calculating the phase difference as a function of time. This is shown in the range between -360 and 360 degrees to aid visualization, i.e. the horizontal lines indicate a same phase value.

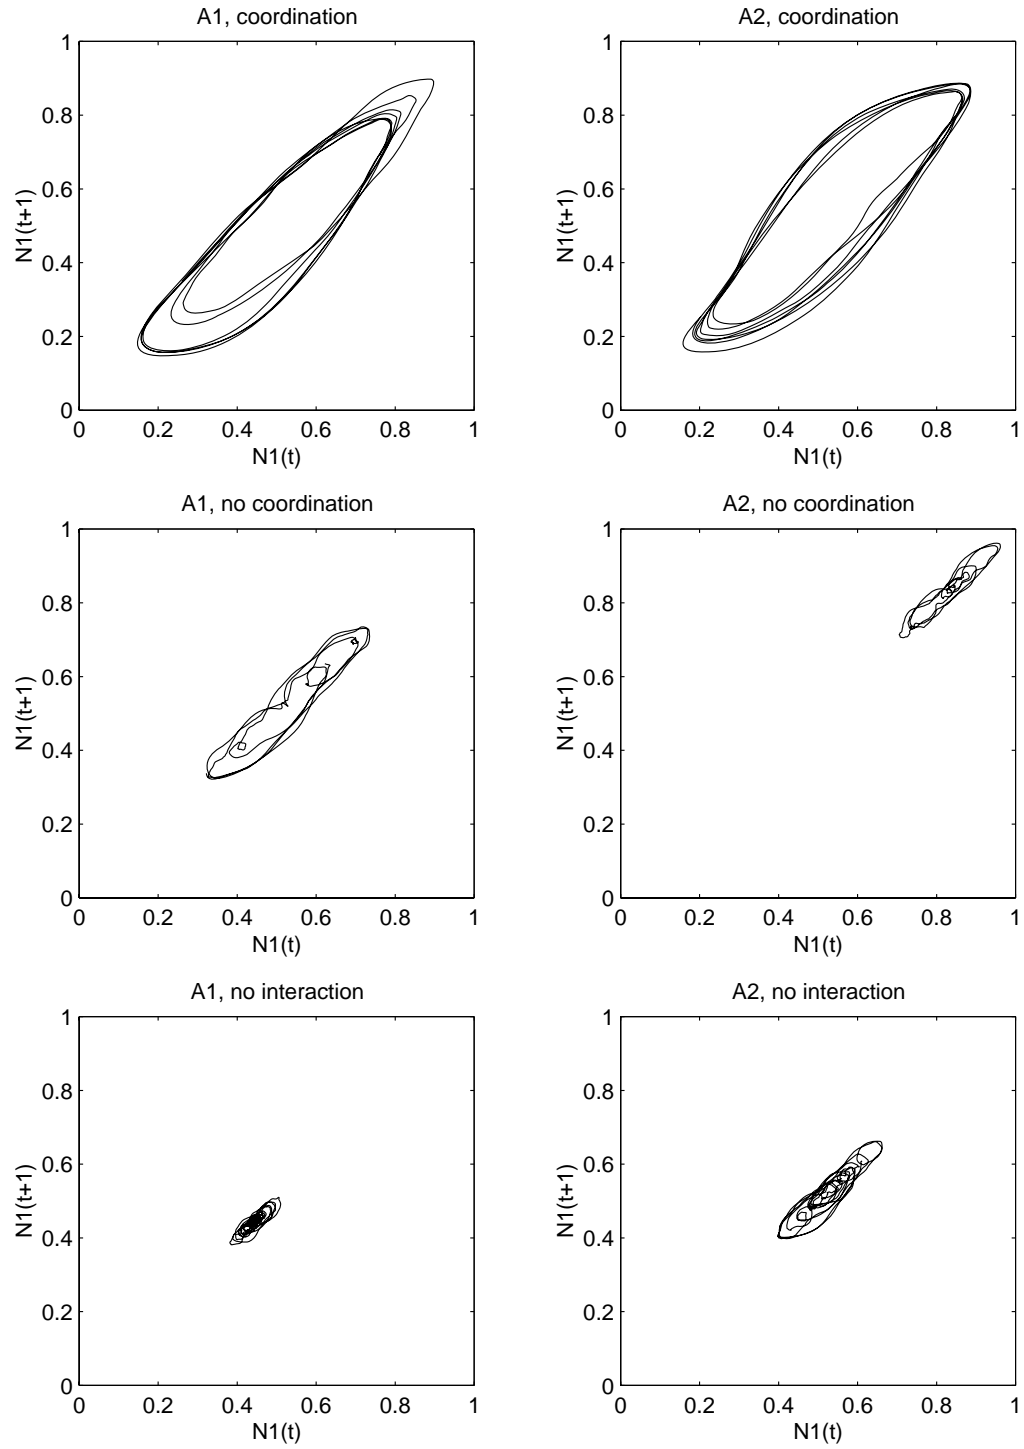


Figure 9.11: Embedded time-delayed dynamics in the same inter-neuron for two agents (A1 and A2) in different situations. The selected inter-neuron was chosen for clarity in the plots.

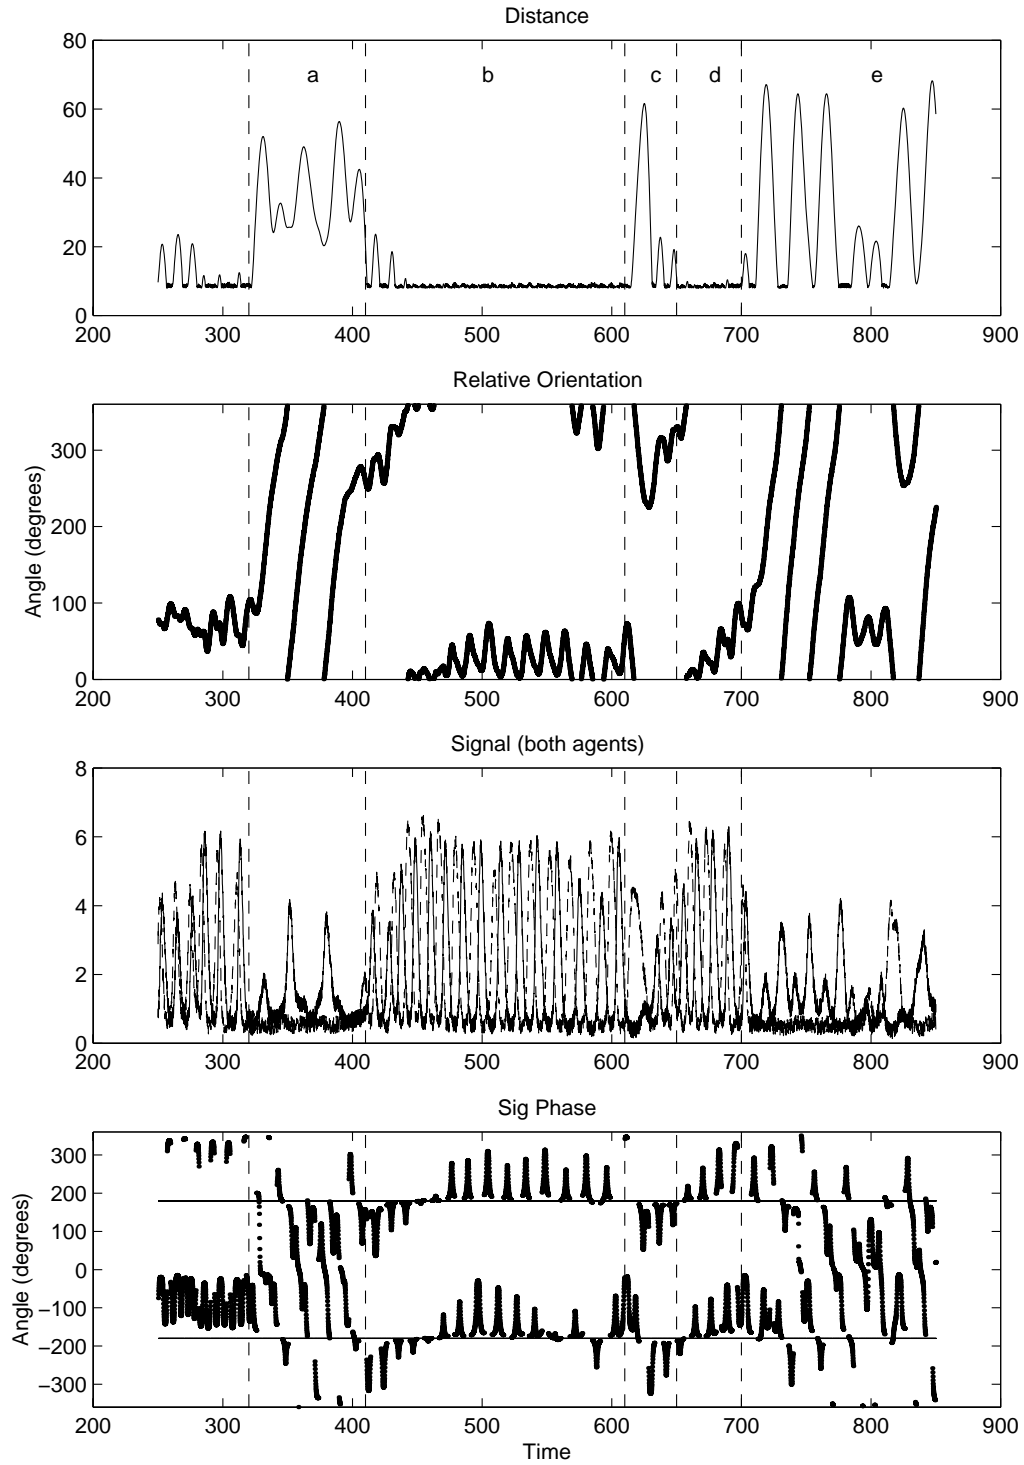


Figure 9.12: Achievement, loss and regaining of coordination. From top to bottom: distance, relative orientation, signals and estimation of relative signal phase as functions of time.

as implicitly noticed in the discussion on the origin of rhythm in signal patterns. Nevertheless, a tentative operational route can be formulated in order to understand what goes on in the model. Such route starts from the physical aspects of sound production/perception as operational givens. Embodied agents actively exploit self-shadowing as a localization mechanism by producing cycloidal movement as a search strategy. Angular movement introduces rhythm in perception which is also manifested in signalling behavior. Rhythmic acoustic signallers become entrained through mutual perturbation by driving one another to a state of structural congruence which is manifested both in the resulting internal and external dynamics. This process is similar to cases in which other coupled oscillators also entrain their dynamics under a variety of circumstances, even in the presence of fluctuations and individual structural differences (Winfree, 1980; Kuramoto, 1984; Kelso, 1995). Finally, coherent signalling behavior helps to maintain the patterns of movement in an ordered state which may be lost due to fluctuations and regained later.

If the best that can be done is to give an incomplete operational account of what happens in an artificial world one may justifiably ask what is the scientific value of this methodology. Although, as discussed in section 9.3, the model says nothing about the evolutionary aspects of social coordination, it is possible to consider some implications of the analysis as to how certain evolutionary questions could be framed.

The study of social behavior and its relevance to the evolution and development of human capabilities has often been approached from a purely functional angle. To repeat, there is nothing inherently wrong with this. However, functional considerations should be grounded on what is known about the operation (at different levels) of the systems concerned. As discussed in chapter 2, functional explanation can be derived from the abbreviation of certain nomic relationships. But without a sufficient exploration of those relationships one runs the risk of missing alternative ways of building functional links as well as the risk of building functional links that disregard operational constraints (section 2.2). For instance, the received wisdom has been in recent years that social life is important for understanding the evolution of human intelligence because social life can be very complex and our ancestors needed to be good predictors of the outcomes of social interactions¹⁸. Here social life and individual capabilities are distinguished from each other and a functional bridge is built to join them back together. And this distinction can serve its purpose, as long as it does not become reified into the operation of the various systems involved. Because from an operational standpoint, even in this simple model, the artificiality of this separation becomes evident. In particular, it was found that the use of the same channel for self-stimulation and social interaction makes it difficult to decompose the behaviors of an embodied agent into social/non-social categories. And this shows that there can be direct *operational* links between what is seen as purely social and what is seen as purely individual.

A related aspect that shows how the divide between social and individual capabilities tends to permeate into operational requirements is the re-discovered fact that behaviors arising from co-adaptation do not necessarily imply an individual capability for performing similar tasks. This has been clearly shown in the model when a demonstration was made of the fact that agents are incapable of entering into coherent signalling behavior with beacons that produce “artificial” and temporally homogeneous signals as they do with other agents. The successful performance of

¹⁸The so-called Machiavellian intelligence hypothesis, see (Humphrey, 1976) and the collections (Byrne & Whiten, 1988; Whitten & Byrne, 1997)

certain behaviors need not be within the realm of competence of the individual organism if the tasks are performed socially in a coordinated fashion, a fact that should serve as a warning whenever one tries to extrapolate operational features (how an organism should work) from functional interpretations of observed evidence (what it does)¹⁹.

To show the potential of the present methodology some brief comments will be added about how this model could be improved and what else could be studied with it. In chapter 2 (section 2.4) some ideas were briefly touched upon which could be used to explain how social forms of learning can be understood from a systemic perspective as a form of directed structural change due to the achievement of structural congruence between unevenly plastic organisms. The feasibility of this hypothesis could be easily tested by extending the model to include richer forms of plastic change, some of which would become gradually “solidified” during the lifetime of the agent. Some cases of imitative learning could possibly be explained in these terms. But the idea is more powerful still, as it could also explain other phenomena such as the bonding observed between duetting pairs in *Laniarius* (section 9.2). Once some of the structural changes undergone due to mutual perturbation have become “frozen” in ways that favour following encounters with the same individual(s), social affinity is the unsurprising outcome.

In the general context of the thesis, this chapter has shown the validity of the operational approach and its applicability to the study of the mechanisms involved in generating coordinated behavior. The direct analysis of dynamics as well as the utilization of more indirect perturbation techniques should be highlighted as particularly useful. The model has also shown that concepts derived from the biology of autonomous systems can have direct application in describing and explaining concrete examples of adaptive behavior, thus proving the practical value of the vocabulary and framework introduced in chapter 2.

¹⁹See also the model presented in chapter 8 for a simpler example.

Chapter 10

Conclusions

This thesis has originated from a combination of interests in some biological and cognitive aspects of social behavior. It has offered a theoretical and methodological perspective on how to approach research on social adaptive behavior in autonomous systems as well as concrete examples of this perspective at work in the form of specific investigations concerned with the evolutionary, ecological and behavioral aspects of basic social coordination.

In general, the importance of understanding evolutionary and behavioral dynamics as part of situated historical processes has been highlighted. This has been a recurrent theme of the thesis. Evolution cannot be considered to proceed by the combination of universal principles which bear the greatest explanatory responsibility and a set of contingencies that complement them. This is especially so in the case of social behaviors where the boundaries that separate evolving traits and environment become blurred as a consequence of ecological situatedness. Analogously, the achievement of social coordination between two autonomous systems cannot be understood, from the operational point of view, as the work of individual problem solvers capable of dealing with a variety of situations, but instead it must be approached as the result of a history of mutual perturbations between embodied systems in what constitutes also a situated process.

To take situatedness seriously sometimes implies renouncing the simplicity of explanatory monism and embracing the complexity of the multiple concurrent and interdependent factors that form a historical process. This may not always be an easy task. However, once a situated process is properly understood many issues that are initially mysterious can be explained in a natural manner. Thus, cooperative coordination in games with conflict of interest can be explained by a combination of Darwinian selection and ecological situatedness in the form of spatial and/or developmental processes that operate concurrently with evolution and in an interdependent manner. In the situation described in chapter 9 rhythm and turn-taking are not *problems* to be solved by individual adaptive agents but arise naturally from the physical constraints implied by acoustic interaction and the nature of a process of structural coupling that leads to coordination. In short, situatedness cannot be avoided if adaptive behavior is to be approached seriously.

With this general conclusion in mind, this final chapter will summarize the main contributions of the thesis and will also present some speculations about interesting future paths along which the investigations could be extended.

10.1 Recapitulation and contributions

Chapter 2 articulated the need to ground the theoretical framework used to deal with social adaptive behavior on an operational theory of biological autonomy. This choice was made both for pragmatic reasons and because of the concern raised about the way traditional biology has treated social behavior exclusively from a functional point of view and, as a consequence, has failed to account properly for paradigmatic cases such as animal communication. The logical contradictions of the misapplication of terms like signals and information and the use of selective advantage as the sole defining norm of communication were carefully exposed. The situation within traditional theoretical biology is such that, even if it is possible to theorize about the evolutionary conditions for the appearance of what is intuitively considered as communication, the phenomenon itself remains undefined beyond this intuitive level.

In response to this situation, the choice was made to ground the work of this thesis on an operational theory of biology, namely a theory of biological autonomy based on systemic concepts and given expression in the works of Humberto Maturana and Francisco Varela. According to this theory, the autonomy of a system, defined as the operational closure of its internal dynamics, restricts the modes in which such a system may engage in interaction with its environment. Those interactions that do not break its autonomy can only influence the system by means of structural perturbations which may trigger new states in the system but cannot specify these states directly. Two autonomous systems undergoing such a process of structural coupling are said to be interacting. If the systems present a degree of initial operational congruence (for instance, if they are two animals of a same species), interaction may lead to a coherence of behaviors which are not themselves directly engaged in the interaction but which would not arise if the interaction was not there (or had not been there). If this happens, the systems are said to be coordinating their behaviors. Sustained or repeated coordination of this kind constitutes a domain of behaviors which appear to depend on one another in a consensual manner. Communication is then defined as those behaviors that occur in a consensual domain. The behavioral coherence that arises as a result of sustained or repeated coordination implies, due to the structural-determinism and autonomy of the interacting systems, an analogous coherence of their structures which is referred to as structural congruence.

Complementing this theoretical basis, chapter 3 provided a general description of what is meant by a historical process. This idea proved to be useful as an aid for interpreting what to expect from certain complex situations and what to look for in order to make sense out of them. A process is historical if it is able to alter its own conditions of realization by exerting an influence on its own constraints either directly or indirectly through coupling with other processes. As such, a process is historical not only because it is contingent but also because time ceases to be homogeneous as variations in constraints change the form in which general dynamical principles are embodied as the process unfolds. This may lead to the formation of spontaneous invariants of the process which, by influencing the dynamics, specify the conditions for their own perpetuation, i.e. for their own invariant character. This peculiar relation between dynamics and constraints provides a perspective from which complex processes may be approached and much of the observed order explained.

From this rather general basis it is possible to go to many places. In this thesis, however, only a limited number of issues have been explored. Social coordination has been studied from

an evolutionary and ecological point of view in order to show that functional questions need not be abandoned (even if, in the end, functional explanations happen to be abandoned). It can be argued that, due to the simplifications made to the behavioral properties of the agents so as to make evolutionary problems tractable, the systemic perspective has not been put to a very hard test in these studies. In order to compensate for this, it seemed appropriate also to study how coordination is achieved from the behavioral point of view so as to see the theory of biological autonomy at work in a concrete case.

In all studies a similar general methodology was applied to the use of formal models and computer simulations. This methodology was explained in chapter 4. Making a model of a natural phenomenon always implies some degree of abstraction. From a pragmatic point of view finding out how much is left out by this abstraction is as important as finding out how much is captured by it. Hence, a pragmatic approach to formal modelling is more concerned with comparing different models so as to reveal the role played by different factors than with how well a model approximates the natural case. Following this approach, models of different granularity and complexity have been used in this way in the thesis. Among those, the most flexible and encompassing were individual-based computer simulations. Some problems regarding their use were discussed in chapter 4 as well as a way of formulating these problems clearly as a case of integrating a new tool into an existing framework of practices.

A solution for these problems was proposed and followed in the investigations. This solution states that a simulation must be observed and understood in itself but not necessarily as a case of the basic model giving rise to the global patterns. Most of the times the situation may not be so simple. Understanding what goes on in a simulation may be quite a complex task in itself and here is where the use of comparative models may come to the rescue by providing an angle from which to observe the simulation and compare it with the expected results. Once an explanation of what goes on in the simulation is proposed, this must be tested, just as in the case of a hypothesis about natural phenomena, by means of specific observations and crucial experiments. Finally, the integration with the rest of scientific practices is achieved by comparing the explanatory structure of the simulation with analog theoretical structures for natural cases. In this way, simulations work in a way similar to thought experiments.

Chapters 5, 6 and 7 can be read as one long investigation on the role of situatedness in the evolution of cooperative social coordination for a simple game with conflict of interest. In chapter 5 the basic action-response game is proposed and analysed from a game-theoretic perspective for which the only “driving force” behind the evolutionary process is differential reproduction. Systematic comparisons reveal how factors like noise, finite populations, discreteness, coupling with a global ecological variable and genetic constraints influence the result. However, in none of these cases cooperative coordination evolves; the results are either an oscillation in strategy space or a combination of non-cooperative strategies.

The model is then extended to include spatial interactions within a finite local domain. This continuous model is based on a long-range coupling between variables which makes homogeneous solutions unstable for non-homogeneous perturbations. The capability of the model for forming spatial patterns is demonstrated analytically and by means of numerical resolutions. Interestingly, the patterns that result are stable clusters of players that divide the population into non-interacting

groups which evolve independently. The introduction of minimal density thresholds for the player densities shows how many apparently unstable solutions may become stable if an approximation is made in the direction of treating the population variables as discrete. This supports observations made empirically and in other models that conclude that it is not sufficient for a single invader to be fitter than its surrounding population but that a minimum threshold in the number of invaders must be achieved for invasion to occur.

Following on this track, an individual-based spatial model, in which agents are discrete entities from the start, is presented in chapter 6. The results of this model could not have been predicted from the previous models. Cooperative coordination is observed to evolve for a moderate degree of conflict between players. Kin selection was examined as a possible candidate for explaining this observations since it is often claimed that spatial situatedness may enhance the chances of interacting with kin and so they may favour altruistic behavior. Unfortunately, it is unusual for researchers to confirm this claim by a proper investigation of relatedness in their computer models. Here, such an investigation was carried out and it was found that kin selection was insufficient for explaining the observed results.

An explanation was finally found after examining the structuring roles of clusters as invariants of the dynamics. Clusters introduce asymmetries in spatial position between the central and peripheral regions. These asymmetries are varied and include the greater length of genealogies that originate in the central region and the differences of frequencies for playing each role. These asymmetries follow from geometrical factors and from an understanding of the forces that maintain the stability of a cluster.

In combination with the discrete nature of the agents and reproductive events and with stochastic fluctuations, the above structuring effects suffice for explaining the evolution of cooperative coordination and its dependence on the degree of conflict of interest. Additionally, the same factors can account for other observations such as the positive correlation between cluster size and average coordination level within a given cluster.

Chapter 7 extends this investigation in a new direction. In view of the invariant structures that emerge from a coupling between evolutionary and ecological process it becomes relevant to question what other factors may introduce similar ordering effects. An assessment of the evolutionary role of one such factor, socially influenced development, is presented. In this model, adult phenotypes depend on genetic constitution and parental influences following certain formal constraints. The simplified developmental scheme that results resembles a random reaction network. The constitution of the population is the outcome of a process which combines selection, ecological factors and developmental relations, all of which are interdependent. The results show that, in the case of space and social development, not only is the range of viability of cooperation enlarged, but the dependence on the degree of conflict changes from a linear decrement to a stepwise descent from total cooperation to no cooperation. An explanation of this change of behavior is found when the developmental relations between agents are analysed. Self-promoting developmental networks constraint the population usually to a single viable phenotype in spite of the variability introduced by mutations. This occurs simply because it is very hard to find new phenotypes that conform to the developmental logic imposed by the dominating invariant network and which, at the same time, participate in a competing invariant structure of their own.

The biological implications of these models are clear. Evolution, particularly evolution of social behaviors, cannot be studied purely as a process de-coupled from its ecological context. As any other natural process, evolution is a situated process and as described in chapter 3 also a historical process which, in combination with ecological and developmental processes, may lead to the formation of invariants that constrain its dynamics to certain paths. Both spatial clusters and self-perpetuating developmental networks are examples of such invariants.

In chapter 8, the same evolutionary scheme is used to provide a simple proof of concept. Contrary to what is usually claimed, computer models of communication need not include the assumption of hidden information for interesting behaviors to evolve. The model shows this in a very simple way by extending the interactions between agents to the production of an alternating sequence of actions and responses that depends on the food type that both of them can see. The task is not trivial since the agents must produce a sequence while being modelled as state-less machines. Therefore, any non-constant response to a constant environmental stimulus must come as a result of their interaction. This simple counterexample also hints to the direction in which coordination could be studied from a behavioral point of view.

Chapter 9 follows in that direction. A model of acoustically coupled agents is presented in which coordinated behavior must follow from continuous interaction. Agents must locate, approach and remain within a certain distance of one another only by producing and sensing sound signals. Sound constitutes a rich medium for studying the role of embodiment in adaptive social behavior. This is clearly demonstrated at the time of studying the behaviors of the agents. Active perception takes advantage of the self-shadowing mechanism by exploiting the differential sensing introduced as a result of varying the relative orientation of the body. Accordingly, agents move in a cycloidal paths that introduce rhythm into the overall dynamics even though no evidence of rhythm is found in the individual agents acting on their own.

Interaction leads to coordination as agents are able to approach one another and remain close. However, further (and unexpected) coordination is observed in the temporal structure of the rhythmic signals and on the patterns of movement of the agents after having approached. Signals are entrained in an anti-phase mode that resembles a rudimentary case of turn-taking and agents move coherently in space by coordinating changes in position and in angular orientation. More interestingly, even if coordination can be lost due to fluctuations, the overall tendency is toward regaining the coordinated state.

It is here that the conceptual structure presented in chapter 2 becomes vital for understanding why agents behave as they do. It is proposed that coordination is achieved due to a process by which agents drive one another to a state of structural congruence as a result of mutual perturbations. In order to confirm that this is the case and that agents are not simply very good individual adaptors, a crucial test is made to see if agents can exhibit a similar signalling behavior in the presence of non-plastic beacons that reproduce the coordinated signalling patterns of a single agent. Failure to observe an entrainment of signals in this case suggests that agents are indeed co-adapting to one another in a historical process.

This model has some nice implications besides the emergence of coherence from embodied interactions and the applicability of the systemic perspective to a concrete case. The fact that co-adaptation need not imply an individual capability to adapt casts some doubts on the applica-

bility of the Chomskyan dictum (reproduced in most forms of methodological individualism) that competence is a precondition for performance (Chomsky, 1965, p. 25), especially in the case of social behavior (which is the domain in which it was originally meant to be applied). It is seen that agents can engage in coordinated behavior with other agents but not individually with similar structures of their environment.

Another implication is concerned with the possibility of decomposing behaviors into well-defined functional modules. This may not always be easy and it is particularly difficult in the presence of social interaction. As shown in this model, the production of signals does not fulfil the purely “social” role of interacting with the other agent, but self-hearing plays an important operational role in the generation of cycloidal approaching movement as is confirmed when this movement collapses as a result of the self-hearing capability being perturbed.

10.2 The future

Although the investigations presented in this thesis are considered to have reached a stage of completion, the reader should, hopefully, be left with the feeling that there is still much to be done along similar lines of research. The thesis will end with a speculative description of some directions in which the work presented here could be expanded.

The approach presented in chapter 9 seems particularly fruitful for extension. One obviously interesting way of extending it would be to try to realize a similar situation in actually embodied systems. This path is certainly difficult, but possible and surely quite fruitful although a clear statement of objectives would have to be made clear from the start.

Still using simulations, it would be interesting, for instance, to explore under what circumstances the coordinated behavior of the acoustically coupled agents would exhibit greater complexity, perhaps approaching something like the duetting patterns in *Laniarius* (section 9.2). One possibility for doing this would be to make the situation more complex in itself by requiring that agents do not just approach one another but that they also do something else. To achieve this, it is almost certain that the environment would have to be modified by the introduction of additional structure. An alternative to making the task more complex would be to enlarge the dimensionality with which the agents may behave. This could be achieved by making their *bodies* more complex and perhaps subject to evolution themselves, or by enlarging the variety of modes of interaction (e.g. sound could be modelled as a field of varying intensity *and* frequency).

Another alternative would be to introduce some sort of dimorphism, like two sexes, and let interactions occur between opposite sexes. Related to this possibility, a potentially more interesting approach would be a two species scenario in which intra-specific interaction requires approaching and inter-specific interaction requires pursuit and evasion. This possibility would be inspired in the approach used by Ficici and Pollack (1998) for the evolution of generators and predictors of binary strings (see section 8.1). Two agents in interactions would be in a situation of having to discriminate if they belong to the same or to different species. This scenario would probably hold a good potential for investigating evolutionary questions as well as behavioral ones such as the dynamics of competitive co-evolution between the species.

A different branch of investigation rooted in the same model would not necessarily be so concerned with complexity of coordinated patterns but with the possibility of testing the hypothesis

that structural congruence between unevenly plastic organisms may be the basis for explaining social affinity and some forms of social learning (see section 2.4). In order to study this hypothesis, the plastic dimensions of the agents would have to be extended as well as a way of letting the history of past interactions be actively reflected in their present structures. One way of doing this is by extending the dynamics of the neural networks to include processes that occur on much slower time-scales. An interesting possibility under consideration is the use of architectures that incorporate additional mechanisms like neuron modulation through emission and diffusion of “gas” (Husbands, Smith, Jakobi, & O’Shea, 1998; Husbands, 1998). These dimensions of added plasticity would facilitate experiments on social affinity such as comparing the co-adaptation between individuals in a “fresh” state with the co-adaptation between individuals who have undergone a long history of co-adaptation with different partners. Similarly, the possibility of social “learning” could be tested by allowing a group composed of genetically identical individuals which are already able to coordinate to undergo two independent evolutionary processes in which genetic constitution is kept constant but the possibility of historical drift in the patterns of behaviors is left open by the recurrent achievement of structural congruence between different generations in the presence of noisy fluctuations.

This last possibility can be “re-connected” with the themes of investigation of the other models in the thesis which in a way can be thought of as approaching a similar issue from the evolutionary and ecological angle. While a full-blown ecological model of acoustically coupled agents would perhaps be too complex to build and probably also to understand, it is possible to make some further abstractions that would make agents behave in simpler ways using perhaps simpler architectures but retaining some interesting behavioral potential beyond mere look-up tables. This would allow to add the possibility of a richer type of history than that exhibited for instance in the study of social influences on development by coupling behavior and developmental changes in a more direct manner.

Bibliography

- Abramson, G., & Zanette, D. H. (1998a). Globally coupled maps with asynchronous upating. *Physical Review E*, 58, 4454 – 4460.
- Abramson, G., & Zanette, D. H. (1998b). Statistics of extinction and survival in Lotka-Volterra systems. *Physical Review E*, 57, 4572.
- Ackley, D. H., & Littman, M. L. (1994). Altruism in the evolution of communication. In Brooks, R. A., & Maes, P. (Eds.), *Artificial Life IV*. MIT Press, Cambridge, Mass.
- Ashby, W. R. (1960). *Design for a Brain: The origin of adaptive behaviour* (Second edition). Chapman and Hall, London.
- Ashby, W. R. (1962). Principles of the self-organizing system. In von Foerster, H., & Zopf, G. (Eds.), *Principles of Self-Organization* (Second edition), pp. 255 – 278. Pergammon Press, New York.
- Avital, E., & Jablonka, E. (1994). Social learning and the evolution of behaviour. *Animal Behaviour*, 48, 1195 – 1199.
- Avital, E., & Jablonka, E. (1996). Adoption, memes and the Oedipus complex: A reply to Hansen. *Animal Behaviour*, 51, 476 – 477.
- Axelrod, R. (1984). *The evolution of cooperation*. Basic Books, New York.
- Baas, N. A. (1994). Emergence, hierarchies, and hyperstructures. In Langton, C. G. (Ed.), *Artificial Life III — studies in the sciences of complexity*, Vol. XVII, pp. 515–537 New York. Addison-Wesley, Reading, MA.
- Baerends-van Roon, J. M., & Baerends, G. P. (1979). *The morphogenesis of behavior of the domestic cat, with special emphasis on the development of prey catching*. North Holland, Amsterdam.
- Baldwin, J. M. (1896). A new factor in evolution. *The American Naturalist*, 30, 441 – 451.
- Barnett, L. (1997). Tangled webs: Evolutionary dynamics on fitness landscapes with neutrality. Master's thesis, School of Cognitive and Computing Sciences, University of Sussex, Brighton, UK.
- Barnett, L. (1998). Ruggedness and neutrality - the NKp family of fitness landscapes. In Adami, C., Belew, R. K., Kitano, H., & Taylor, C. E. (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*. MIT Press, Cambridge, Mass.
- Bateson, P. P. G. (1978). Sexual imprinting and optimal outbreeding. *Nature*, 273, 659 – 660.
- Bateson, P. P. G. (1980). Optimal outbreeding and the development of sexual preferences in Japanese quail. *Zeitschrift für Tierpsychologie*, 53, 236 – 237.
- Beer, R. D. (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. 121–147. MIT Press, Cambridge, Mass.

- Beer, R. D. (1996). Toward the evolution of dynamical neural networks for minimally cognitive behavior. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J. B., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 421 – 429. MIT Press, Cambridge, Mass.
- Beer, R. D., & Gallagher, J. C. (1992). Evolving dynamical neural networks for adaptive behavior. *Adaptive Behavior*, 1:1, 91 – 122.
- Berman, C. M. (1982a). The ontogeny of social relationships with group companions among free-ranging infant rhesus monkeys I. Social networks and differentiation. *Animal Behaviour*, 30, 149 – 162.
- Berman, C. M. (1982b). The ontogeny of social relationships with group companions among free-ranging infant rhesus monkeys II. Differentiation and attractiveness. *Animal Behaviour*, 30, 163 – 170.
- Bernardo, J. (1996). Maternal effects in animal ecology. *American Zoologist*, 36, 83 – 105.
- Bersini, H., & Detours, V. (1994). Asynchrony induces stability in cellular automata based models. In Brooks, R., & Maes, P. (Eds.), *Artificial Life IV, Proceedings of the Fourth International Conference on Artificial Life*. MIT Press, Cambridge, Mass.
- Blackwell, D., & Kendall, D. (1964). The Martin boundary for Polya's urn scheme and an application to stochastic population growth. *Journal of Applied Probability*, 1, 284.
- Boerlijst, M. C., & Hogeweg, P. (1991a). Self-structuring and selection: Spiral waves as a substrate for evolution. In Langton, C. G. (Ed.), *Artificial Life II*, pp. 255 – 276 Redwood City, CA. Addison-Wesley, Reading, MA.
- Boerlijst, M. C., & Hogeweg, P. (1991b). Spiral wave structure in pre-biotic evolution: Hypercycles stable against parasites. *Physica D*, 48, 17 – 28.
- Boerlijst, M. C., & Hogeweg, P. (1995a). Attractors and spatial patterns in hypercycles with negative interactions. *Journal of Theoretical Biology*, 176, 199 – 210.
- Boerlijst, M. C., & Hogeweg, P. (1995b). Spatial gradients enhance persistence of hypercycles. *Physica D*, 88, 29 – 39.
- Boerlijst, M. C., Lamers, M. E., & Hogeweg, P. (1993). Evolutionary consequences of spiral waves in a host-parasitoid system. *Proc. R. Soc. Lond. B*, 253, 15 – 18.
- Boesch, C. (1991). Teaching among wild chimpanzees. *Animal Behaviour*, 41, 530 – 532.
- Bohm, D. (1969a). Addendum on order and neo-Darwinism. In Waddington, C. H. (Ed.), *Towards a Theoretical Biology 2, Sketches*, pp. 90 – 93. Edinburgh University Press.
- Bohm, D. (1969b). Some comments on Maynard-Smith's contributions. In Waddington, C. H. (Ed.), *Towards a Theoretical Biology 2, Sketches*, pp. 98 – 105. Edinburgh University Press.
- Bohm, D. (1980). *Wholeness and the implicate order*. Routledge.
- Boneabeau, E. W., & Theraulaz, G. (1995). Why do we need Artificial Life. In Langton, C. (Ed.), *Artificial Life: An overview*. MIT Press, Cambridge, Mass.
- Bourdieu, P. (1977). *Outline of a theory of practice*. Cambridge Studies in Social Anthropology: 16. Cambridge University Press.
- Bourdieu, P. (1989). *In other words: Essays towards a reflexive sociology*. Polity, Cambridge.

- Bourdieu, P. (1990). *The logic of practice*. Polity, Cambridge.
- Brockmann, H. J., Grafen, A., & Dawkins, R. (1979). Evolutionarily stable strategy in a digger wasp. *Journal of Theoretical Biology*, 77, 473 – 496.
- Brooks, R. (1991). Intelligence without representation. *Artificial Intelligence*, 47, 139–159.
- Bullock, S. (1997). An exploration of signalling behaviour by both analytic and simulation means for both discrete and continuous models. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 454 – 463. MIT Press/Bradford Books, Cambridge, Mass.
- Bullock, S. G. (1998). *Evolutionary Simulation Models: On their character and application to problems concerning the evolution of natural signalling systems*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex.
- Burghardt, G. M. (1970). Defining ‘communication’. In Johnston Jr, J. W., Moulton, D. G., & Turk, A. (Eds.), *Communication by Chemical Signals*. Appleton-Century-Crofts, New York.
- Byrne, R. W., & Whiten, A. (Eds.). (1988). *Machiavellian intelligence: Social expertise and the evolution of intellect in monkeys, apes and humans*. Clarendon Press, Oxford.
- Cangelosi, A., & Parisi, D. (1998). The emergence of a “language” in an evolving population of neural networks. *Connection Science*, 10, 83 – 97.
- Cariani, P. (1991). Emergence and Artificial Life. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II — studies in the sciences of complexity*, Vol. X, pp. 41–91 Redwood City, California. Addison-Wesley, Reading, MA.
- Caro, T. M. (1980). Predatory behaviour in domestic cat mothers. *Behaviour*, 74, 128 – 147.
- Caro, T. M., & Hauser, M. D. (1992). Is there teaching among nonhuman animals? *The Quarterly Review of Biology*, 67, 151 – 174.
- Chalmers, N. (1983). The development of social relationships. In Halliday, T. R., & Slater, P. J. B. (Eds.), *Animal behaviour. Vol 3: Genes, development and learning*, pp. 114 – 148. Blackwell Scientific, Oxford.
- Chomsky, N. (1965). *Aspects of the theory of syntax*. MIT Press, Cambridge, Mass.
- Clark, A. (1997). *Being there: Putting brain, body, and world together again*. MIT Press, Cambridge, Mass.
- Clark, A., & Wheeler, M. (1998). Bringing representation back to life. In Pfeifer, R., Blumberg, B., & Kobayashi, H. (Eds.), *Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior, SAB’98*. MIT Press, Cambridge, Mass.
- Comins, H. N., & Hassell, M. P. (1996). Persistence of multispecies host-parasitoid interactions in spatially distributed models with local dispersal. *Journal of Theoretical Biology*, 183, 19 – 28.
- Condon, W. S. (1979). Neonatal entrainment and enculturation. In Bullowa, M. (Ed.), *Before Speech: The beginning of interpersonal communication*, pp. 131–148. Cambridge University Press, Cambridge.
- Condon, W. S., & Sander, L. W. (1974). Neonate movement is synchronized with adult speech: Interactional participation and language acquisition. *Science*, 183, 99 – 101.

- Cowley, D. E., & Atchley, W. R. (1992). Quantitative genetic models for development, epigenetic selection and phenotypic evolution. *Evolution*, 46, 495 – 518.
- Cressman, R. (1990). Strong stability and density-dependent evolutionarily stable strategies. *Journal of Theoretical Biology*, 145, 319 – 330.
- Cressman, R., & Vickers, G. T. (1997). Spatial and density effects in evolutionary game theory. *Journal of Theoretical Biology*, 184, 259 – 369.
- Cronhjort, M. B., & Blomberg, C. (1994). Hypercycles versus parasites in a two dimensional partial differential equations model. *Journal of Theoretical Biology*, 169, 31 – 49.
- Crutchfield, J. P. (1994a). The calculi of emergence: Computation, dynamics and induction. *Physica, D* 75, 11 – 54.
- Crutchfield, J. P. (1994b). Evolution of communication in artificial organisms. In Cowan, G., Pines, D., & Melzner, D. (Eds.), *Integrative themes*, Vol. XIX of *SFI studies in the sciences of complexity*. Addison-Wesley, Reading, MA.
- Dawkins, M. S. (1995). *Unravelling animal behaviour* (Second edition). Longman Scientific and Technical, Harlow, England.
- Dawkins, R. (1976). *The Selfish Gene*. Oxford University Press.
- Dawkins, R. (1982). *The Extended Phenotype*. Oxford University Press.
- de Bourcier, P., & Wheeler, M. (1994). Signalling and territorial aggression: An investigation by means of synthetical ethology. In Cliff, D., Husbands, P., Meyer, J.-A., & Wilson, S. W. (Eds.), *From Animals to Animats 3: Proceedings of the Third International Conference on Simulation of Adaptive Behavior*, pp. 463 – 472. MIT Press/Bradford Books, Cambridge, Mass.
- de Bourcier, P., & Wheeler, M. (1997). The truth is out there: The evolution of reliability in aggressive communication systems. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 444 – 453. MIT Press/Bradford Books, Cambridge, Mass.
- de la Torre, J. C., & Holland, J. J. (1990). RNA virus quasispecies populations can suppress vastly superior mutant progeny. *Journal of Virology*, 64, 6278 – 6281.
- de Wall, F. B. M. (1996). Macaque social culture: Development and perpetuation of affiliative networks. *Journal of Comparative Psychology*, 110, 147 – 154.
- Dennett, D. (1995). Artificial Life as philosophy. In Langton, C. (Ed.), *Artificial Life: An overview*. MIT Press, Cambridge, Mass.
- Derrida, B., & Peliti, L. (1991). Evolution in a flat fitness landscape. *Bulletin of Mathematical Biology*, 53, 355.
- Dewey, J. (1929a). *Experience and Nature* (Second edition). Dover, New York. (1958).
- Dewey, J. (1929b). *The Quest for Certainty: A study of the relation of knowledge and action*. John Dewey, *The Later Works 1925-1953*, V.4. Southern Illinois University Press, 1984.
- Di Paolo, E. A. (1996a). A computational model of speciation in non-uniform environments without physical barriers. Cognitive science research paper 412, School of Cognitive and Computing Sciences, University of Sussex.

- Di Paolo, E. A. (1996b). Some false starts in the construction of a research methodology for Artificial Life. In Noble, J., & Parsowith, S. (Eds.), *The 9th White House Papers: Graduate Research in Cognitive Science and Computing Sciences at Sussex*.
- Di Paolo, E. A. (1997a). An investigation into the evolution of communication. *Adaptive Behavior*, 6:2, 285 – 324.
- Di Paolo, E. A. (1997b). Social coordination and spatial organization: steps towards the evolution of communication. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 464 – 473. MIT Press, Cambridge, Mass.
- Di Paolo, E. A. (1998a). Assessing the role of social development in the evolution of cooperation. In Pfeifer, R., Blumberg, B., & Kobayashi, H. (Eds.), *Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior, SAB'98*. MIT Press, Cambridge, Mass.
- Di Paolo, E. A. (1998b). Behavioural coordination in acoustically coupled agents. In Nicklasson, L., Bodén, M., & Ziemke, T. (Eds.), *Proceedings of the Eighth International Conference on Artificial Neural Networks, ICANN'98*, pp. 1097 – 1102 Skövde, Sweden. Springer, London.
- Di Paolo, E. A. (1999). Behavioral coordination, structural congruence and entrainment in acoustically coupled agents. *Adaptive Behavior*, In Press.
- Dreyfus, H., & Rabinow, P. (1993). Can there be a science of existential structure and social meaning? In Calhoun, C., LiPuma, E., & Postone, M. (Eds.), *Bourdieu: Critical perspectives*. Polity, Cambridge.
- Dreyfus, H. L. (1991). *Being-in-the-world: A commentary on Heidegger's Being and Time, Division 1*. MIT Press, Cambridge, Mass.
- Dunbar, R. I. M. (1992). Neocortex size as a constraint on group size in primates. *Journal of Human Evolution*, 20, 469 – 493.
- Dunbar, R. I. M. (1993). Co-evolution of neocortex size, group size and language in humans. *Brain and Behavioral Sciences*, 16, 681 – 735.
- Durrett, R., & Levin, S. (1994a). The importance of being discrete (and spatial). *Theoretical Population Biology*, 46, 363 – 394.
- Durrett, R., & Levin, S. (1994b). Stochastic spatial models: A user's guide to ecological applications. *Phil. Trans. R. Soc. Lond. B*, 343, 329 – 350.
- Eigen, M., & Schuster, P. (1979). *The Hypercycle: A principle of self-organization*. Springer Verlag, Berlin.
- Ewer, R. F. (1969). The “instinct to teach”. *Nature*, 222, 698.
- Faith, J. (1997). In defence of functional analysis. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 20 – 27. MIT Press, Cambridge, Mass.
- Falconer, D. S. (1965). Maternal effects and selection response. In Geerts, S. J. (Ed.), *Proceedings of the XI International Congress on Genetics. Vol. 3*, pp. 763 – 774. Pergamon Press, Oxford.
- Farabaugh, S. M. (1982). The ecological and social significance of duetting. In Kroodsma, D. E., & Miller, E. H. (Eds.), *Acoustic Communication in Birds. Vol. 2*. Academic Press, New York.

- Feddersen, W. E., Sandel, T. T., Teas, D. C., & Jeffress, L. A. (1957). Localization of high-frequency tones. *Journal of the Acoustic Society of America*, 29, 988 – 991.
- Fermi, E., Pasta, J., & Ulam, S. M. (1965). Studies of non-linear problems, Document LA-1940, (may 1955). In *Collected papers of Enrico Fermi*, Vol. 2. University of Chicago Press.
- Feyerabend, P. K. (1988). *Against method* (Third edition). Verso, London.
- Feynman, R. P. (1986). *Surely you're joking, Mr. Feynman: adventures of a curious character*. Unwin, London.
- Ficici, S. G., & Pollack, J. B. (1998). Coevolving communicative behavior in a linear pursuer-evader game.. In Pfeifer, R., Blumberg, B., & Kobayashi, H. (Eds.), *Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior, SAB'98*. MIT Press, Cambridge, Mass.
- Finney, R. L., & Thomas, G. B. (1990). *Calculus*. Addison-Wesley, Reading, Mass.
- Foley, W. A. (1997). *Anthropological linguistics: An introduction*. Blackwell, Oxford.
- Fontana, W., & Buss, L. (1994a). The arrival of the fittest. *Bull. Math. Biol.*, 56, 1 – 64.
- Fontana, W., & Buss, L. (1994b). What would conserved if 'the tape were played twice'? *Proc. Natl. Acad. Sci. USA*, 91, 757 – 761.
- Fontana, W., & Buss, L. (1996). The barrier of objects: from dynamical systems to bounded organizations. In Casti, J., & Karlqvist, A. (Eds.), *Boundaries and Barriers*, pp. 56 – 116. Addison-Wesley, Reading, MA.
- Fontana, W., Wagner, G., & Buss, L. W. (1994). Beyond digital naturalism. *Artificial Life*, 1/2, 211 – 227.
- Frederickson, W. T., & Sackett, G. P. (1984). Kin preferences in primates (*Macaca nemestrina*): Relatedness or familiarity? *Journal of Comparative Psychology*, 98, 29 – 34.
- Gavrilets, S., & Gravner, J. (1996). Percolation on the fitness hypercube and the evolution of reproductive isolation. *Journal of Theoretical Biology*, 184, 51 – 64.
- Gavrilets, S., Li, H., & Vose, M. D. (1998). Rapid parapatric speciation on holey adaptive landscapes. *Proc. Roy. Soc. Lond. B*, 265, 1483 – 1489.
- Gergen, K. J. (1994). *Realities and relationships : Soundings in social construction*. Harvard University Press, Cambridge, Mass.
- Goldstein, H. (1980). *Classical Mechanics* (Second edition). Addison-Wesley, Reading, MA.
- Goodnight, K. F. (1992). The effect of stochastic variation on kin selection in a budding-viscous population. *American Naturalist*, 140, 1028 – 1040.
- Goodwin, B. C. (1993). Development as a robust natural process. In Stein, W. D., & Varela, F. J. (Eds.), *Thinking about biology: An invitation to current theoretical biology*. Addison-Wesley, Reading, Mass.
- Goodwin, B. C. (1994). *How the leopard changed its spots: The evolution of complexity*. Weidenfeld and Nicolson, London.
- Gould, S. J. (1989a). A developmental constraint in *Cerion*, with comments on the definition and interpretation of constraint in evolution. *Evolution*, 43(3), 516 – 539.

- Gould, S. J. (1989b). *Wonderful life: The Burgess Shale and the nature of history*. Hutchinson Radius, London.
- Gould, S. J., & Lewontin, R. C. (1979). The spandrels of San Marco and the Panglossian paradigm: A critique of the adaptationist programme. *Proceedings of the Royal Society London, B* 205, 581 – 598.
- Grafen, A. (1991). Modelling in behavioural ecology. In Krebs, J. R., & Davies, N. B. (Eds.), *Behavioural ecology: An evolutionary approach* (Third edition), chap. 1, pp. 5 – 31. Blackwell Scientific Publications, Oxford.
- Haimoff, E. H. (1986). Convergence in the duetting of monogamous old world primates. *Journal of Human Evolution*, 15, 51 – 59.
- Haken, H. (1978). *Synergetics, an introduction: Nonequilibrium phase transitions and self-organization in physics, chemistry and biology*. Springer, Berlin.
- Haken, H. (1983). *Advanced synergetics: Instability hierarchies of self-organizing systems and devices*. Springer, Berlin.
- Hamilton, W. D. (1964). The genetical evolution of social behaviour, I and II. *Journal of Theoretical Biology*, 7, 1 – 16; 17 – 32.
- Hamilton, W. D. (1971). Geometry for the selfish herd. *Journal of Theoretical Biology*, 31, 295 – 311.
- Hammersley, J. M., & Handscomb, D. C. (1964). *Monte Carlo Methods*. Methuen, London.
- Hansen, T. F. (1996). Does adoption make evolutionary sense? *Animal Behaviour*, 51, 474 – 475.
- Harlow, H. F. (1958). The nature of love. *American Psychologist*, 13, 673 – 685.
- Harvey, I. (1996). Untimed and misrepresented: connectionism and the computer metaphor. *AISB Quarterly*, pp. 20 – 27.
- Harvey, I., & Bossomaier, T. (1997). Time out of joint: Attractors in asynchronous random Boolean networks. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 67 – 75. MIT Press, Cambridge, Mass.
- Harvey, I., Husbands, P., Cliff, D., Thompson, A., & Jakobi, N. (1997). Evolutionary Robotics: the Sussex Approach. *Robotics and Autonomous Systems*, 20, 205–224.
- Hasson, O. (1994). Cheating signals. *Journal of Theoretical Biology*, 167, 223 – 238.
- Heidegger, M. (1962). *Being and Time*. Harper and Row, New York. Trans. J. Macquarrie and E. Robinson.
- Helbing, D., Keltsch, J., & Molnár, P. (1997). Modelling the evolution of human trail systems. *Nature*, 388, 45 – 50.
- Hemelrijk, C. K. (1997). Cooperation without genes, games or cognition. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 511 – 520. MIT Press, Cambridge, Mass.
- Hemelrijk, C. K. (1998). Spatial centrality of dominants without positional preference. In Adami, C., Belew, R. K., Kitano, H., & Taylor, C. E. (Eds.), *Artificial Life VI: Proceedings of the Sixth International Conference on Artificial Life*. MIT Press, Cambridge, Mass.

- Hendriks-Jansen, H. (1996). *Catching ourselves in the act: Situated activity, interactive emergence and human thought*. MIT Press, Cambridge, Mass.
- Hess, E. (1973). *Imprinting: Early experience and the developmental psychobiology of attachment*. Van Nonstrand.
- Hesse, M. B. (1980). The explanatory function of metaphor. In *Revolutions and reconstructions in the philosophy of science*. Harvester Press.
- Hinde, R. A., & Spencer-Booth, Y. (1971). Effects of brief separations from mother on rhesus monkeys. *Science*, 773, 111 – 118.
- Hines, W. G. S., & Maynard-Smith, J. (1979). Games between relatives. *Journal of Theoretical Biology*, 79, 19 – 30.
- Hinton, G. E., & Nowlan, S. J. (1987). How learning can guide evolution. *Complex Systems*, 1, 495 – 502.
- Hofbauer, J., Schuster, P., & Sigmund, K. (1979). A note on evolutionarily stable strategies and game dynamics. *Journal of Theoretical Biology*, 81, 609 – 612.
- Holland, O. (1996). Grey Walter: the pioneer of real artificial life. In *Proceedings of the ALife V Conference*, pp. 16 – 23 Nara, Japan. Oral Presentations.
- Holmes, W. G., & Mateo, J. M. (1998). How mothers influence the development of litter-mate preferences in Belding's ground squirrels. *Animal Behaviour*, 55, 1555 – 1570.
- Hooker, T., & Hooker, B. I. (1969). Duetting. In Hinde, R. A. (Ed.), *Bird vocalizations: Their relation to current problems in biology and psychology*. Cambridge University Press.
- Huberman, B. A., & Glance, N. S. (1993). Evolutionary games and computer simulations. *Proc. Natl. Acad. Sci. USA*, 90, 7715 – 7718.
- Humphrey, N. K. (1976). The social function of the intellect. In Bateson, P., & Hinde, R. (Eds.), *Growing points in ethology*. Cambridge University Press.
- Hurd, P. L. (1995). Communication in discrete action-response games. *Journal of Theoretical Biology*, 174, 217 – 222.
- Husbands, P. (1998). Evolving robot behaviours with diffusing gas networks. In Husbands, P., & Meyer, J.-A. (Eds.), *Proceedings of the First European Workshop on Evolutionary Robotics (EvoRobot98)*, pp. 123 – 136. Springer Verlag.
- Husbands, P., Smith, T., Jakobi, N., & O'Shea, M. (1998). Better living through chemistry: Evolving GasNets for robot control. *Connection Science*, 10, 185 – 210.
- Huynen, M. A., Stadler, P. F., & Fontana, W. (1996). Smoothness within ruggedness: The role of neutrality in adaptation. *Proc. Natl. Acad. Sci. USA*, 93, 397 – 401.
- Ihde, D. (1979). *Technics and Praxis*, Vol. XXIV of *Boston Studies in the Philosophy of Science*. D. Reidel Publishing Co. Dordrecht, Holland.
- Ingerson, T. E., & Buvel, R. L. (1984). Structure in asynchronous cellular automata. *Physica D*, 10, 59 – 68.
- Jackson, E. A. (1996). The second metamorphosis of science: A second view. Working paper 96-06-039, Santa Fe Institute.

- Jakobi, N. (1997). Evolutionary robotics and the radical envelope-of-noise hypothesis. *Adaptive Behavior*, 6:2, 325 – 368.
- Jones, C. G., Lawton, J. H., & Schachak, M. (1994). Organisms as ecosystem engineers. *Oikos*, 69, 373 – 386.
- Kandel, E. R., Schwartz, J. H., & Jessell, T. M. (Eds.). (1991). *Principles of Neural Science* (Third edition). Elsevier, New York.
- Kauffman, S. A. (1993). *The Origins of order*. Oxford University Press, New York.
- Kelso, J. A. S. (1995). *Dynamic patterns: The self-organization of brain and behavior*. MIT Press, Cambridge, Mass.
- Kepler, T. B., & Perelson, A. S. (1995). Modeling and optimization of populations subject to time-dependent mutation. *Proc. Natl. Acad. Sci. USA*, 92, 8219.
- Kimura, M. (1995). Limitations to Darwinian selection in a finite population. *Proc. Natl. Acad. Sci. USA*, 92, 2343 – 2344.
- King, A. J., & Carlile, S. (1995). Neural coding for auditory space. In Gazzaniga, M. S. (Ed.), *The Cognitive Neurosciences*. MIT Press, Cambridge, Mass.
- Kirkpatrick, M., & Lande, R. (1989). The evolution of maternal characters. *Evolution*, 43, 485 – 503.
- Kirkpatrick, M., & Lande, R. (1992). The evolution of maternal characters: Errata. *Evolution*, 46, 284.
- 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, (ECAL'97)*, pp. 274 – 283. MIT Press/Bradford Books, Cambridge, Mass.
- Krakauer, D. C., & Pagel, M. (1995). Spatial structure and the evolution of honest cost-free signalling. *Proc. R. Soc. Lond. B*, 260, 365 – 372.
- Krause, J. (1994). Differential fitness returns in relation to spatial position in groups. *Biological Reviews*, 69, 187 – 206.
- Krebs, J. R., & Davies, N. (1993). *An introduction to behavioural ecology* (Third edition). Blackwell, Oxford.
- Krebs, J. R., & Dawkins, R. (1984). Animal signals: Mind reading and manipulation. In Krebs, J. R., & Dawkins, R. (Eds.), *Behavioural ecology: An evolutionary approach* (Second edition), chap. 15, pp. 380–402. Blackwell, Oxford.
- Kuhn, T. (1970). *The structure of scientific revolutions*. University of Chicago Press. Second Edition.
- Kuhn, T. (1977a). The essential tension: Tradition and innovation in scientific research. In *The essential tension: Selected studies in scientific tradition and change*. Chicago University Press.
- Kuhn, T. (1977b). A function for thought experiments. In *The essential tension: Selected studies in scientific tradition and change*. Chicago University Press.
- Kuramoto, Y. (1984). *Chemical oscillations, waves and turbulence*. Springer, Berlin.

- Lakoff, G. (1986). *Women, fire and dangerous things: What categories reveal about the mind*. University of Chicago Press.
- Laland, K. H. (1994). On the evolutionary consequences of sexual imprinting. *Evolution*, 48, 477 – 489.
- Lemmen, R. (1998). *Towards a non-Cartesian cognitive science in the light of the philosophy of Merleau-Ponty*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex.
- Levins, R., & Lewontin, R. C. (1985). *The dialectical biologist*. Harvard University Press.
- Lewis, D. B., & Gower, M. D. (1980). *Biology of communication*. Blackie, Glasgow.
- Lewontin, R. (1982). Organism and environment. In Plotkin, H. C. (Ed.), *Learning, development and culture: Essays in evolutionary epistemology*. Wiley Chichester.
- Lewontin, R. C. (1983). The organism as the subject and object of evolution. *Scientia*, 118, 63 – 82.
- Leyhausen, P. (1979). *Cat behavior: The predatory and social behavior of domestic wild cats*. Garland STPM Press, New York.
- Lopez, J. C., & Lopez, D. (1985). Killer whales (*Orcinus orca*) of Patagonia and their behavior of intentional stranding while hunting near shore. *J. Mammal.*, 66, 181 – 183.
- Lund, H. H., Webb, B., & Hallam, J. (1997). A robot attracted to cricket species *Gryllus bimaculatus*. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 246 – 255. MIT Press, Cambridge, Mass.
- MacLennan, B. J., & Burghardt, G. M. (1994). Synthetic ecology and the evolution of cooperative communication. *Adaptive Behavior*, 2:2, 151–188.
- Maestriperi, D. (1996). Maternal encouragement of infant locomotion in pigtail macaques, *Macaca nemestrina*. *Animal Behaviour*, 51, 603 – 610.
- Margolis, J. (1986). *Pragmatism without foundations. The Persistence of Reality Vol 1: Reconciling realism and relativism*. Blackwell, Oxford.
- Margolis, J. (1987). *Science without unity. The persistence of reality Vol 2: Reconciling the human and natural sciences*. Blackwell, Oxford.
- Marsden, J. E., & Tromba, A. J. (1996). *Vector Calculus* (Fourth edition). Freeman, New York.
- Maturana, H. (1978). Biology of language: The epistemology of reality. In Miller, G. A., & Lennenberg, E. (Eds.), *Psychology and biology of language and thought. Essays in honor of Eric Lennenberg*, pp. 27–63. Academic Press Inc. New York.
- Maturana, H. (1988a). Brain, language and the origin of human mental functions. In *Conference workbook. Texts in cybernetics. American Society for Cybernetics Conference*. Felton, USA, 18 - 23 October, 1988.
- Maturana, H. (1988b). Reality: The search for objectivity or the quest for a compelling argument. *The Irish Journal of Psychology*, 9(1), 25 – 82.
- Maturana, H. (1991). Science and daily life. In Steier, F. (Ed.), *Research and reflexivity*. Sage, London.

- Maturana, H., Mpodozis, J., & Letelier, J. C. (1995). Brain, language and the origin of human mental functions. *Biological Research*, 28, 15 – 26.
- Maturana, H., & Varela, F. J. (1980). *Autopoiesis and cognition: The realization of the living*. D. Reidel Publishing, Dordrecht, Holland.
- Maturana, H., & Varela, F. J. (1987). *The tree of knowledge: The biological roots of human understanding*. Shambhala, Boston, Mass.
- May, R. M. (1973). *Stability and complexity in model ecosystems*. Princeton University Press, Princeton, New Jersey.
- May, R. M., Bonhoeffer, S., & Nowak, M. A. (1995). Spatial games and the evolution of co-operation. In Moran, F., Moreno, A., Merelo, J. J., & Chacon, P. (Eds.), *Proceedings of the Third European Conference on Artificial Life*, pp. 749 – 759 Granada, Spain. Springer Verlag, Berlin.
- Maynard-Smith, J. (1969). The status of neo-Darwinism. In Waddington, C. H. (Ed.), *Towards a Theoretical Biology 2, Sketches*, pp. 82 – 89. Edinburgh University Press.
- Maynard-Smith, J. (1982). *Evolution and the theory of games*. Cambridge University Press.
- Maynard-Smith, J. (1988). Can a mixed strategy be stable in a finite population? *Journal of Theoretical Biology*, 130(247 - 251).
- Maynard-Smith, J. (1993). *The theory of evolution*. Cambridge University Press.
- Maynard-Smith, J., & Harper, D. G. C. (1995). Animal signals: models and terminology. *Journal of Theoretical Biology*, 177, 305–311.
- Maynard-Smith, J., & Parker, G. A. (1976). The logic of asymmetric contests. *Animal Behaviour*, 24, 159 – 175.
- Maynard-Smith, J., & Price, G. R. (1973). The logic of animal conflict. *Nature*, 246, 15 – 18.
- Merleau-Ponty, M. (1962). *Phenomenology of perception*. Routledge and Kegan Paul.
- Merleau-Ponty, M. (1964). *Le visible et l'invisible*. Gallimard, Paris.
- Merleau-Ponty, M. (1968). *The visible and the invisible*. Northwestern University Press.
- Michod, R. E. (1982). The theory of kin selection. *Annual Review of Ecology and Systematics*, 13, 23 – 55.
- Miller, G. E. (1995). ALife as theoretical biology: How to do science with computer simulation. Cognitive science research paper 378, School of Cognitive and Computing Sciences, University of Sussex.
- Millikan, R. G. (1984). *Language, thought and other biological categories: New foundations for realism*. MIT Press, Cambridge, Mass.
- Mitchell, M. (1996). *An introduction to genetic algorithms*. Bradford Books, MIT Press, Cambridge, Mass.
- Murray, J. D. (1989). *Mathematical biology* (Second edition). Springer, Berlin.
- Nagel, E. (1961). *The structure of science: Problems in the logic of scientific explanations*. Routledge, London.

- Nakamaru, M., Matsuda, H., & Iwasa, Y. (1997). The evolution of cooperation in a lattice-structured population. *Journal of Theoretical Biology*, 184, 65 – 81.
- Nettle, D., & Dunbar, R. I. M. (1997). Social markers and the evolution of reciprocal exchange. *Current Anthropology*, 38, 93 – 99.
- Neuenschwander, S., & Varela, F. J. (1993). Visually triggered neuronal oscillations in birds: An autocorrelation study of tectal activity. *European Journal of Neuroscience*, 5, 870 – 881.
- Noble, J. (1998a). Evolved signals: Expensive hype vs. conspiratorial whispers. In Adami, C., Belew, R. K., Kitano, H., & Taylor, C. E. (Eds.), *Proceedings of the Sixth International Conference on Artificial Life (ALife VI)*.
- Noble, J. (1998b). Tough guys don't dance: Intention movements and the evolution of signalling in animal. In Pfeifer, R., Blumberg, B., & Kobayashi, H. (Eds.), *Proceedings of the Fifth International Conference on the Simulation of Adaptive Behavior, SAB'98*. MIT Press, Cambridge, Mass.
- Noble, J. (1997). The scientific status of Artificial Life. Poster presented at the Fourth European Conference on Artificial Life (ECAL'97), Brighton, UK, 28–31 July.
- Nowak, M. A., & May, R. M. (1992). Evolutionary games and spatial chaos. *Nature*, 359, 836 – 829.
- Odling-Smee, F. J. (1988). Niche-constructing phenotypes. In Plotkin, H. C. (Ed.), *The role of behavior in evolution*, pp. 73–132. MIT Press, Cambridge, Mass.
- Oliphant, M. (1994). Evolving cooperation in the non-iterated Prisoner's Dilemma: The importance of spatial organization. In Brooks, R. A., & Maes, P. (Eds.), *Artificial Life IV*. MIT Press, Cambridge, Mass.
- Oliphant, M. (1996). The dilemma of Saussurean communication. *BioSystems*, 37, 31 – 38.
- Oyama, S. (1985). *The ontogeny of information*. Cambridge University Press.
- Pask, G. (1959). Physical analogues to the growth of a concept. In *Mechanisation of thought processes: Proceedings of a symposium held at the National Physical Laboratory on 24 - 27 November 1958. Vol. II*. Her Majesty Stationary Office, London.
- Pask, G. (1960). The natural history of networks. In Yovits, M. C., & Cameron, S. (Eds.), *Self-organizing systems* (Second edition), pp. 232 – 263. Pergamon Press, New York.
- Pask, G. (1961). *An approach to cybernetics*. Hutchinson.
- Pattee, H. H. (1970). The problem of biological hierarchy. In Waddington, C. H. (Ed.), *Towards a Theoretical Biology 3, Drafts*, pp. 117 – 136. Edinburgh University Press.
- Pattee, H. H. (1971). Physical theories of biological co-ordination. *Quarterly Reviews of Biophysics*, 4, 255 – 276.
- Pattee, H. H. (1972). Laws and constraints, symbols and languages. In Waddington, C. H. (Ed.), *Towards a Theoretical Biology 4, Essays*, pp. 248 – 258. Edinburgh University Press.
- Pattee, H. H. (1996). Simulations, realizations, and theories of life. In Boden, M. A. (Ed.), *The philosophy of Artificial Life*, pp. 379–393. Oxford University Press, Oxford.
- Port, R., Tajima, K., & Cummins, F. (1998). Speech and rhythmic behavior. In Savelsburgh, G. J. P., van der Maas, H., & van Geert, P. C. L. (Eds.), *The non-linear analyses of developmental processes*. Royal Dutch Academy of Arts and Sciences, Amsterdam.

- Port, S. C. (1994). *Theoretical probability for applications*. Wiley, New York.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., & Flannery, B. P. (1992). *Numerical recipes in C: The art of scientific computing* (Second edition). Cambridge University Press.
- Prigogine, I., & Stengers, I. (1984). *Order out of chaos: Man's new dialogue with nature*. Heinemann, London.
- Rasmussen, S., & Barrett, C. L. (1995). Elements of a theory of simulation. In Moran, F., Moreno, A., Merelo, J. J., & Chacon, P. (Eds.), *Proceedings of the Third European Conference on Artificial Life*, pp. 515–537 Granada, Spain. Springer Verlag, Berlin.
- Ridley, M. (1987). Imprinting. In McFarland, D. (Ed.), *The Oxford companion to animal behavior*. Oxford University Press, Oxford.
- Riley, J. G. (1979). Evolutionary equilibrium strategies. *Journal of Theoretical Biology*, 76, 109–123.
- Robson, S. K., & Beshers, S. N. (1997). Division of labour and ‘foraging for work’: Simulating reality versus the reality of simulations. *Animal Behaviour*, 53, 214–218.
- Rohani, P., & Miramontes, O. (1995). Host-parasitoid metapopulations: The consequences of parasitoid aggregation on spatial dynamics and searching efficiency. *Proc. R. Soc. Lond. B*, 260, 335–342.
- Rolf, J., Bohr, T., & Jensen (1998). Directed percolation universality in asynchronous evolution of spatio-temporal intermittency. *Physical Review E*, 57, R2503–R2506.
- Rorty, R. (1988). Is natural science a natural kind? In *Construction and constraint: The shaping of scientific rationality*. University of Notre Dame Press.
- Rozenweig, M. R. (1954). Cortical correlates of auditory localization and of related perceptual phenomena. *Journal of Comparative and Physiological Psychology*, 47(4), 269–276.
- Rozenweig, M. R. (1961). Auditory localization. *Scientific American*, 205(4).
- Rubinstein, R. Y. (1981). *Simulation and the Monte Carlo method*. John Wiley and sons, New York.
- Rutkowska, J. C. (1994). Scaling up sensorimotor systems: Constraints from human infancy. *Adaptive Behavior*, 2, 349–373.
- Sales, G., & Pye, D. (1974). *Ultrasonic communication by animals*. Chapman and Hall, London.
- Saunders, G. M., & Pollack, J. B. (1996). The evolution of communication schemes over continuous channels. In Maes, P., Mataric, M. J., Meyer, J.-A., Pollack, J. B., & Wilson, S. W. (Eds.), *From Animals to Animats 4: Proceedings of the Fourth International Conference on Simulation of Adaptive Behavior*, pp. 580–589. MIT Press, Cambridge, Mass.
- Savill, N. J., Rohani, P., & Hogeweg, P. (1997). Self-reinforcing spatial patterns enslave evolution in a host-parasitoid system. *Journal of Theoretical Biology*, 188, 11–20.
- Sawaguchi, T., & Kudo, H. (1990). Neocortical development and social structure in primates. *Primates*, 31, 283–290.
- Schaffer, M. E. (1988). Evolutionarily stable strategies for a finite population and a variable contest size. *Journal of Theoretical Biology*, 132(469–478).

- Schaller, G. (1966). *The deer and the tiger*. The University of Chicago Press, Chicago.
- Schenkel, R. (1966). Play, exploration, and territoriality in the wild lion. *Symp. Zool. Soc. Lond.*, 18, 11 – 22.
- Schuster, P., & Sigmund, K. (1981). Coyness, philandering and stable strategies. *Animal Behaviour*, 29, 186 – 192.
- Schuster, P., Sigmund, K., & Wolff, R. (1978). Dynamical systems under constant organisation I; topological analysis of a family of non-linear differential equations. *Bull. Math. Biophysics*, 40, 743 – 769.
- Shannon, C., & Weaver, W. (1949). *The mathematical theory of communication*. University of Illinois Press.
- Stadler, P. F., Fontana, W., & Miller, J. H. (1993). Random catalytic reaction networks. *Physica D*, 63, 378 – 392.
- Steels, L. (1996). A self-organizing spatial vocabulary. *Artificial Life*, 3(2).
- Steels, L. (1997a). Constructing and sharing perceptual distinctions. In van Someren, M., & Widmer, G. (Eds.), *Proceedings of the European Conference on Machine Learning*. Springer-Verlag, Berlin.
- Steels, L. (1997b). Synthesising the origins of language and meaning using co-evolution, self-organization, and level formation.. In Hurford, J., Knight, C., & Studdert-Kennedy, M. (Eds.), *Evolution of human language*. Edinburgh University Press, Edinburgh.
- Steinle, F. (1996). Entering new fields: Exploratory uses of experimentation. *Philosophy of Science*, 64, S65 – S74.
- Stewart, J. (1996). Cognition = Life: Implications for higher-level cognition. *Behavioural Processes*, 35, 311 – 326.
- Syswerda, G. (1989). Uniform crossover in genetic algorithms. In Schaffer, J. D. (Ed.), *Proceedings of the Third International Conference on Genetic Algorithms*. Morgan Kaufmann.
- Taylor, C. (1993). To follow a rule... In Calhoun, C., LiPuma, E., & Postone, M. (Eds.), *Bourdieu: Critical perspectives*. Polity, Cambridge.
- Taylor, C., & Jefferson, D. (1995). Artificial Life as a tool for biological inquiry. In Langton, C. (Ed.), *Artificial Life: An overview*. MIT Press, Cambridge, Mass.
- Taylor, P. D., & Jonker, L. B. (1978). Evolutionarily stable strategies and game dynamics. *Math. Biosc.*, 40, 145 – 156.
- Thelen, E., & Smith, L. B. (1994). *A dynamic systems approach to the development of cognition and action*. MIT Press, Cambridge, Mass.
- Thompson, E., Palacios, A., & Varela, F. J. (1992). Ways of coloring – Comparative color vision as a case study for cognitive science. *Brain and Behavioral Sciences*, 15(1), 1 – 25.
- Thorpe, W. H. (1972). *Duetting and antiphonal song in birds: Its extent and significance*. E. J. Brill, Leiden.
- Tinbergen, N. (1952). “Derived” activities; their causation, biological significance, origin, and emancipation during evolution. *Quarterly Review of Biology*, 27(1), 1 – 32.

- Todd, P. M., & Miller, G. F. (1993). Parental guidance suggested: How parental imprinting evolves through sexual selection as an adaptive learning mechanism. *Adaptive Behavior*, 2:1, 5 – 47.
- Trivers, R. L. (1974). Parent-offspring conflict. *American Zoologist*, 14, 249 – 264.
- Tsimring, L., Levine, H., & Kessler, D. A. (1996). RNA virus evolution via a fitness-space model. *Phys. Rev. Lett*, 76, 4440 – 4444.
- Uexküll, J. (1934). A stroll through the worlds of animals and men. In Lashley, K. (Ed.), *Instinctive behavior*. International Universities Press.
- van Baalen, M., & Rand, D. A. (1998). The unit of selection in viscous populations and the evolution of altruism. *Journal of Theoretical Biology*, 193, 631 – 648.
- Van Gelder, T. (1999). The dynamical hypothesis in cognitive science. To appear in *Behavioral and Brain Sciences*.
- Van Gelder, T., & Port, R. (Eds.). (1995). *Mind as motion: Explorations in the dynamics of cognition*. MIT Press, Cambridge, Mass.
- van Nimwegen, E., Crutchfield, J. P., & Mitchell, M. (1997a). Finite populations induce metastability in evolutionary search. *Physics Letters A*, 229 (2), 144 – 150.
- van Nimwegen, E., Crutchfield, J. P., & Mitchell, M. (1997b). Statistical dynamics of the Royal Road genetic algorithm. Working paper 97-04-035, Santa Fe Institute.
- Varela, F. J. (1979). *Principles of biological autonomy*. Elsevier, North Holland, NY.
- Varela, F. J. (1997). Patterns of life: Intertwining identity and cognition. *Brain and Cognition*, 34, 72 – 87.
- Varela, F. J., & Coutinho, A. (1991). Second generation immune networks. *Immunology Today*, 12, 159 – 167.
- Varela, F. J., & Stewart, J. (1990). Dynamics of a class of immune networks. I. Global stability of idiotype interactions. *Journal of Theoretical Biology*, 144, 93 – 101.
- Varela, F. J., Stewart, J., & Coutinho, A. (1993). What is the immune system for? In Stein, W., & Varela, F. J. (Eds.), *Thinking about biology: an invitation to current theoretical biology*, Series. Addison Wesley, Reading, Mass.
- Varela, F. J., Thompson, E., & Rosch, E. (1991). *The embodied mind: Cognitive science and human experience*. MIT Press, Cambridge, Mass.
- Vickers, G. T. (1989). Spatial patterns and ESS's. *Journal of Theoretical Biology*, 140, 129 – 135.
- Vickery, W. L. (1987). How to cheat against a simple mixed strategy ESS. *Journal of Theoretical Biology*, 127(133 - 139).
- Vickery, W. L. (1988). Reply to Maynard-Smith. *Journal of Theoretical Biology*, 132, 375 – 378.
- von Békésy, G. (1972). The ear. In Held, R., & Richards, W. (Eds.), *Perception: Mechanisms and models*, Readings from Scientific American. W. H. Freeman and Co., San Francisco.
- von Foerster, H. (1977). Objects: tokens for eigenbehavior. In Inhelder, B. (Ed.), *Hommage à Jean Piaget*. Delachaux et Niestel, Neuchatel.

- von Foerster, H. (1980). The epistemology of communication. In Woodward, K. (Ed.), *The myths of information: Technology and post-industrial culture*. Routledge.
- von Glasersfeld, E. (1984). An introduction to radical constructivism. In Watzlawick, P. (Ed.), *The invented reality*. W. W. Norton, New York.
- Von Neumann, J., & Morgenstern, O. (1944). *Theory of games and economic behavior*. Princeton University Press, Princeton, NJ.
- Vygotsky, L. S. (1962). *Thought and language*. MIT Press, Cambridge, Mass.
- Vygotsky, L. S. (1978). *Mind in society: The development of higher psychological processes*. Harvard University Press.
- Waddington, C. H. (1969). Paradigm for an evolutionary process. In Waddington, C. H. (Ed.), *Towards a theoretical biology 2, Sketches*, pp. 106 – 124. Edinburgh University Press.
- Waddington, C. H. (1975). *The evolution of an evolutionist*. Edinburgh, University Press, Edinburgh.
- Walter, W. G. (1953). *The living brain*. Duckworth, London.
- Walter, W. G. (1950). An imitation of life. *Scientific American*, 182(5), 42 – 45.
- Walter, W. G. (1951). A machine that learns. *Scientific American*, 185(2), 60 – 63.
- Webster, G. C., & Goodwin, B. C. (1996). *Form and transformation: Generative and relational principles in biology*. Cambridge University Press, Cambridge.
- Werner, G. M., & Dyer, M. G. (1992). Evolution of communication in artificial organisms. In Langton, C. G., Taylor, C., Farmer, J. D., & Rasmussen, S. (Eds.), *Artificial Life II*, Vol. X of *SFI studies in the sciences of complexity*. Addison-Wesley.
- Werner, G. M., & Todd, P. M. (1997). Too many love songs: Sexual selection and the evolution of communication. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 434 – 443. MIT Press/Bradford Books, Cambridge, Mass.
- Wertsch, J. V. (1979). From social interaction to higher psychological processes: A clarification and application of vygotsky's theory. *Human Development*, 22, 1 – 22.
- Wertsch, J. V. (1991). *Voices of the mind: A sociocultural approach to mediated action*. Harvester Wheatsheaf.
- Wheeler, M. (1995). Escaping the Cartesian mind-set: Heidegger and Artificial Life. In Moran, F., Moreno, A., Merelo, J. J., & Chacon, P. (Eds.), *Proceedings of the Third European Conference on Artificial Life*, pp. 65 – 76 Granada, Spain. Springer Verlag, Berlin.
- Wheeler, M. (1996). *The philosophy of situated activity*. Ph.D. thesis, School of Cognitive and Computing Sciences, University of Sussex.
- Wheeler, M. (1997). Cognition's coming home: The reunion of life and mind. In Husbands, P., & Harvey, I. (Eds.), *Proceedings of the Fourth European Conference on Artificial Life*, pp. 10 – 19. MIT Press, Cambridge, Mass.
- Whitten, A., & Byrne, R. W. (Eds.). (1997). *Machiavellian Intelligence II: Extensions and evaluations*. Cambridge University Press.
- Wickler, W. (1980). Vocal duetting and the pair bond I: Coyness and partner commitment. A hypothesis. *Zeitschrift für Tierpsychologie*, 52, 201 – 209.

- Wickler, W., & Seibt, U. (1979). Duetting: A daily routine of *Laniarius funebris* the slate-coloured Boubou (Aves, Laniidae). *Zeitschrift für Tierpsychologie*, 51, 153 – 157.
- Wickler, W., & Seibt, U. (1980). Vocal duetting and the pair bond II: Unisono duetting in the African forest weaver, *Symplectes bicolor*. *Zeitschrift für Tierpsychologie*, 52, 217 – 226.
- Wilson, E. O. (1975). *Sociobiology: The new synthesis*. Belknap Press of Harvard University Press, Cambridge, Mass.
- Winfree, A. T. (1980). *The geometry of biological time*. Springer Verlag, New York.
- Wittgenstein, L. (1963). *Philosophical investigations* (Second edition). Oxford: Basil Blackwell and Mott.
- Zahavi, A. (1975). Mate selection - a selection for handicap. *Journal of Theoretical Biology*, 53, 205–214.
- Zahavi, A. (1977). The cost of honesty (further remarks on the handicap principle). *Journal of Theoretical Biology*, 67, 603–605.
- Zeeman, E. C. (1979). Population dynamics for game theory. In Nitecki, Z., & Robinson, E. (Eds.), *Global theory of dynamical systems*. Springer Verlag, Berlin.
- Zeeman, E. C. (1981). Dynamics of the evolution of animal conflicts. *Journal of Theoretical Biology*, 89, 249 – 270.