


Artificial Societies

The Computer Simulation
of Social Life



Edited by

Nigel Gilbert & Rosaria Conte

Also available as a printed book
see title verso for ISBN details

Artificial societies

Artificial societies

The computer simulation of social life

Edited by

Nigel Gilbert

University of Surrey

&

Rosaria Conte

Institute of Psychology
Italian National Research Council



© Nigel Gilbert, Rosaria Conte and contributors, 1995

This book is copyright under the Berne Convention. No reproduction without permission. All rights reserved.

First published in 1995 by UCL Press

This edition published in the Taylor & Francis e-Library, 2005.

“To purchase your own copy of this or any of Taylor & Francis or Routledge's collection of thousands of eBooks please go to www.eBookstore.tandf.co.uk.”

UCL Press Limited University College London Gower Street London WC1E 6BT

The name of University College London (UCL) is a registered trade mark used by UCL Press with the consent of the owner.

ISBN 0-203-99369-1 Master e-book ISBN

ISBN: 1-85728-305-8 (Print Edition) HB

British Library Cataloguing-in-Publication Data A CIP catalogue record for this book is available from the British Library.

Library of Congress Cataloging-in-Publication Data Artificial societies: the computer simulation of social life/edited by Nigel Gilbert and Rosaria Conte. p. cm. Includes bibliographical references and index. ISBN 1-85728-305-8 1. Social sciences—Simulation methods—Congresses. 2. Computer simulation—Congresses. I. Gilbert, G.Nigel. II. Conte, Rosaria, 1952– H61.A7 1995 94–47440 300'.1'13–dc20 CIP

Contents

Preface	vii
Notes on contributors	ix

1 Introduction: computer simulation for social theory <i>Rosaria Conte and Nigel Gilbert</i>	1
---	---

Part I The simulation of social theories

2 A model of the emergence of new political actors <i>Robert Axelrod</i>	15
3 Division of labour and social co-ordination modes: a simple simulation model <i>Massimo Egidi and Luigi Marengo</i>	33
4 Emergence of kinship structures: a multi-agent approach <i>Jean Pierre Treuil</i>	48
5 Cities can be agents too: a model for the evolution of settlement systems <i>Stéphane Bura, France Guérin-Pace, Hélène Mathian, Denise Pumain and Lena Sanders</i>	72
6 The EOS project: integrating two models of Palaeolithic social change <i>Jim Doran and Mike Palmer</i>	86
7 Genetic algorithms, teleological conservatism, and the emergence of optimal demand relations: the case of learning-by-consuming <i>Roger McCain</i>	106

Part II The evolution and emergence of societies

8 Emergence in social simulation <i>Nigel Gilbert</i>	122
9 How to invent a lexicon: the development of shared symbols in interaction <i>Edwin Hutchins and Brian Hazlehurst</i>	132
10 MANTA: new experimental results on the emergence of (artificial) ant societies <i>Alexis Drogoul, Bruno Corbara, Steffen Lalande</i>	160
11 Emergent behaviour in societies of heterogeneous, interacting agents: <i>alliances and norms</i>	180

alliances and norms
Nicholas V.Findler and R.M.Malyankar

Part III The foundations of social simulation

12	Kin-directed altruism and attachment behaviour in an evolving population of neural networks <i>Domenico Parisi, Federico Cecconi and Antonio Cerini</i>	201
13	Understanding the functions of norms in social groups through simulation <i>Rosaria Conte and Cristiano Castelfranchi</i>	213
14	A logical approach to simulating societies <i>Michael Fisher and Michael Wooldridge</i>	227
	Bibliography	240
	Index	252

Preface

This volume draws on contributions made to the second of a continuing series of international symposia on 'simulating societies'. All the chapters, except the one by Robert Axelrod, first saw the light of day at a meeting held in a converted monastery at Certosa di Pontignano, near Siena in Italy, in July 1993. Robert Axelrod was not able to be present, but contributed a chapter nevertheless. The meeting in Siena followed one held at the University of Surrey, Guildford, UK, in April 1992 (see Gilbert and Doran's *Simulating societies: the computer simulation of social phenomena*, 1994). The editors and contributors are very grateful to Cristiano Castelfranchi and Rosaria Conte for their hospitality in arranging the Siena meeting.

As we note in the Introduction (Chapter 1), this volume and its predecessor provide evidence for the growth and current liveliness of the field. The vast power of current personal computers and the sophistication of the software and the programming tools now available means that increasingly the limit on the computer simulation of societies is not the difficulty of implementing simulations, but the imagination and creativity of the researcher. That this imagination is not lacking is shown by the range of contributions to this volume, which draws on almost the whole span of the social sciences to learn more about and experiment with everything from tastes for cultural objects to patterns of Aboriginal kinship, and from the organization of ant nests to the location of cities.

The editors are grateful to Justin Vaughan and Kate Williams of UCL Press for their encouragement to embark on this series and for their patience in helping to put together the book, and to Jim Doran for his help in organizing the Siena conference. We also gratefully acknowledge the support of the Universities of Siena and Surrey.

NIGEL GILBERT ROSARIA CONTE July 1994

Notes on contributors

Robert Axelrod is the Arthur W. Bromage Distinguished University Professor of Political Science and Public Policy at the University of Michigan. He received his BA in mathematics from the University of Chicago, and his PhD in political science from Yale University. He is the author of *The evolution of cooperation* (1984) and *Conflict of interest* (1970), the editor of *Structure of decision* (1976), and the author of numerous articles including “An evolutionary approach to norms” (1986). His work on cooperation and norms has received awards from the American Association for the Advancement of Science, the American Political Science Association, the MacArthur Foundation and the National Academy of Sciences. He has been elected to membership in the American Academy of Arts and Sciences, and the National Academy of Sciences.

Stéphane Bura is writing a doctoral thesis in distributed artificial intelligence at the Paris VI University. His research is focused on the emergence of structures in complex systems. His interests include artificial life, social simulation and memetics.

Cristiano Castelfranchi is a researcher at the Institute of Psychology of the National Research Council of Italy. He has been the head of the Project for the Simulation of Social Behaviour since 1988. He is the author of several books on social and cognitive modelling. In 1992 and 1993 he was Programme Chairman of the European Workshop on Modelling Autonomous Agents in a Multi-Agent World (MAAMAW).

Federico Cecconi is with the Institute of Psychology of the National Research Council in Rome. He has a technical degree in computer science. He is interested in neural networks and genetic algorithms and in their application for both scientific and technological purposes.

Antonio Cerini did his thesis work for a “laurea” in philosophy at the Institute of Psychology of the National Research Council in Rome. He is currently with ISED Informatica, a software house in Rome.

Rosaria Conte is Researcher at the Institute of Psychology of the Italian National Research Council (CNR), Rome, and is teaching Cognitive and Social Psychology at the University of Torino. Her first degree is in philosophy and her post-graduate research training is in the cognitive modelling of social behaviour. She coordinates a research group on the simulation of social behaviour at the Institute of Psychology, with the aim of working out computational models of social phenomena both at the micro and at the macro levels. Her interests include cognitive modelling, social simulation, multi-agent systems, distributed artificial intelligence and the theory of rational action.

Bruno Corbara has a doctorate in Ethology from the University of Paris XIII. His first degree is in psychology and his doctorate is on the social organization and its genesis in the *Ectatomma ruidum* ant. He has studied many species of social animals and has published two books on social insects. He is currently writing a book on the various

building behaviours, either collective or individual, that exist in the animal world. His interests include self-organized systems, theories of behaviour, and statistical methods for ethological research.

Jim Doran read mathematics at Oxford, became a specialist in artificial intelligence at the University of Edinburgh and is now Professor at the University of Essex.

Throughout his career he has applied statistical and computer methods in archaeology. His particular interest is in the simulation of societies using the concepts and techniques of distributed artificial intelligence.

Alexis Drogoul is Assistant Professor at the University of Paris VI. His first degree is in computer science from the University of Paris VI and his doctorate is on the use of multi-agents systems in ethological simulation and problem solving. The main theme of his work concerns the ability of groups of agents to self-organize in order to produce functional collective responses, and more particularly, the emergence of these responses within groups of non-intelligent agents (also called reactive agents). His interests include biological simulation, theories of the emergence of cognition, and distributed problem solving.

Massimo Egidi is Professor of Economics at the Department of Economics of the University of Trento, Italy. He is especially interested in the theory of organization and collective decision making.

Nicholas V. Findler is Research Professor of Computer Science and the Director of the Artificial Intelligence Laboratory at Arizona State University, Tempe, Arizona, USA. He has a BE (Honors, Summa-cum-Laude) in electrical engineering and a PhD in mathematical physics from the Technical University of Budapest, Hungary. He has worked in artificial intelligence since 1957. A large proportion of his projects over the past 12 years or so have been in the area of distributed artificial intelligence, ranging from applications in future air traffic control, street traffic signal control, through allocation of moving resources to moving tasks and nationwide manufacturing control, to studies of the intelligent and heterogeneous multi-agent systems described in this book.

Michael Fisher is a Principal Lecturer in Computing Research at the Manchester Metropolitan University. His first degree was in mathematics and computer science, and for his PhD he worked on the temporal semantics of concurrent programming languages. He is well-known for his work in both the theory and implementation of temporal theorem-proving, and is one of the leading proponents of executable temporal logics as programming languages. His other research interests include DAI, formal specification and verification, object-based systems, and adult learning.

Nigel Gilbert is Professor of Sociology at the University of Surrey. His first degree is in engineering from the University of Cambridge and his doctorate is on the sociology of scientific knowledge. He is a founder member of the Social and Computer Sciences research group at Surrey, which applies insights from the social sciences to the design of information technology. His interests include social simulation, natural language interaction with computer systems, and statistical methods for social research.

France Guérin-Pace is a demographer and a statistician at the French National Institute for Demographic Studies (INED). She has written *Deux siècles de croissance urbaine* (*Two centuries of urban growth*) Anthropos Ed.

Brian Hazlehurst is a Post Doctoral Fellow in Cognitive Science at the University of California, San Diego. In 1994 he received his doctorate in cognitive science and anthropology from the University of California, San Diego for his dissertation entitled “Fishing for cognition: an ethnography of fishing practice in a community on the west coast of Sweden”. His research includes the ethnographic study of cognition and communication systems, and modelling cultural processes.

Edwin Hutchins is Professor of Cognitive Science at the University of California, San Diego. He received his doctorate in anthropology from the University of California, San Diego. In 1985 he was awarded a MacArthur Foundation fellowship. His current research concerns cognition in the cockpits of commercial airliners and computational models of cultural processes.

Steffen Lalande is a PhD student at the University of Paris VI. He is currently working on the simulation of the emergence of social order within bands of apes, in collaboration with the Laboratory of Ethology of Strasbourg. His work consists of studying the individual pre-cognitive bases required for the emergence of hierarchies and the functional role of inter-individual interactions. His interests include social simulation, methodologies of multi-agents simulation, and distributed problem solving.

Luigi Marengo is a Research Fellow at the Department of Economics, University of Trento, Italy, and at the International Institute for Applied Systems Analysis, Laxenburg, Austria. He has a doctorate from the Science Policy Research Unit, University of Sussex, United Kingdom, and is an economist interested in the theory of organization and collective decision making.

Roger A. McCain is Professor of Economics and Director of the Undergraduate Honors Program at Drexel University in Philadelphia, USA. His doctorate is from Louisiana State University. Following doctoral research on economic growth in Latin America, much of McCain’s research has focused on mathematical modelling with applications in economic growth, the economics of the arts, and the economics of labour-managed enterprises. McCain is also the author of a recent work on economic psychology, and his research now focuses on computer simulations of boundedly rational agents.

Raphael M. Malyankar is a Research Scientist in the Department of Computer Science and Engineering at Arizona State University. He earned his doctoral degree in computer science at Arizona State University. His research interests include distributed intelligent systems, societal simulation, distributed artificial intelligence and tactical command and control and decision support systems.

Hélène Mathian is a statistician in the PARIS (Pour l’Avancement de la Recherche sur l’Interaction Spaciale) team at CNRS in Paris.

Mike Palmer gained a BA in archaeology and an MSc in computer science at the University of Newcastle upon Tyne. He was a Senior Research Officer at the University of Essex, assisting in the development of distributed artificial intelligence models of the emergence of social complexity. He is now working for APV Baker Ltd. as a software engineer developing an online scheduling system for the food industry.

Domenico Parisi is director of the Institute of Psychology of the National Research Council in Rome. He has a “laurea” in philosophy from the University of Rome, a Master of Arts in Psychology from the University of Illinois, and a “libera docenza” in

Psychology. His research interests include neural network and artificial life models to study the behaviour of simple and, hopefully, more complex organisms.

Denise Pumain is a geographer and a professor at the Paris I University. She is a technical consultant at the French National Institute for Demographic Studies (INED) and director at the CNRS (National Center for Scientific Research). She has written books about urbanization and modelling. She is the publishing director of the “Villes” collection (Anthropos Ed.).

Lena Sanders is a demographer and a statistician in the PARIS (Pour l’Avancement de la Recherche sur l’Interaction Spaciale) team at CNRS. She teaches geography at the Paris VII University. She has written a book on data analysis in geography and co-wrote *Villes et auto-organisation* with D. Pumain and T.Saint-Julien.

Jean Pierre Treuil is in charge of the Laboratory of Computer Sciences at the French research institute, ORSTOM. He is a graduate of the Paris Ecole Polytechnique and Ecole Nationale de Statistique et d’Administration Economique. His doctorate is in artificial intelligence and learning processes. His interests include artificial life, methods of distributed artificial intelligence and computer simulations of social processes.

Michael Wooldridge is a Lecturer in Computing at the Manchester Metropolitan University. He obtained a first-class honours degree in computer science in 1989, and went on to study for a PhD, which he obtained from UMIST in Manchester in 1992. His doctoral research was in the area of logics for modelling distributed AI systems. His current research interests lie primarily in the area of distributed AI, and he has publications in both the theoretical and practical aspects of the field.

Chapter 1

Introduction

Computer simulation for social theory

Rosaria Conte and Nigel Gilbert

In the Introduction to this volume's predecessor, *Simulating societies: the computer simulation of social phenomena* (Gilbert & Doran 1994), the editors, Nigel Gilbert and Jim Doran, described the state of the art in research on social simulation and defined social simulation as a particular approach to the study of societies that involved "the use of precisely specified simulation models, often formulated and run on computers". This approach, said the editors, has been "fruitful in the past and...appears likely to be even more so in the future" (p. 1). Their forecast has proved to be correct even within the short time since it was made. Three different indicators suggest that the use of computer simulations for studying essential aspects of societies has gained importance:

- (a) The number of studies relying on this methodology has increased.
- (b) The number of disciplines involved has also grown. Once, studies in social simulation were drawn from the central core of the social sciences and from computer science. Now other disciplines, such as cognitive science, biology, neuroscience, and some artificial intelligence (AI) subfields, such as distributed artificial intelligence and research on multi-agent systems, have been showing a growing interest in the computer-based study of societies. The motivations that are leading people with a non-social scientific background to this common area of investigation are obviously heterogeneous. However, there are some broad questions of cross-disciplinary interest that seem to demand the use of computer simulation. These include evolutionary and dynamic models, the problem of complexity and the paradigm of emergence. Later in this chapter we shall return to these questions and to the necessity of using computer simulation to deal with them successfully.
- (c) The number of theoretical perspectives involved has increased. Even radically different approaches, such as conventional AI rule-based systems and research based on neural nets and genetic algorithms, share the use of computer simulation as their methodological grounding. The more differentiated the conceptual and theoretical framework of social simulation becomes, the more influential and significant will be its scientific role.

Doran and Gilbert (1994) argue that computer simulation is an appropriate methodology whenever a social phenomenon is not directly accessible, either because it no longer exists (as in archaeological studies) or because its structure or the effects of its structure, i.e. its behaviour, are so complex that the observer cannot directly attain a clear picture of

what is going on (as in some studies of world politics). The simulation is based on a model constructed by the researcher that is more observable than the target phenomenon itself. This raises issues immediately about which aspects of the target ought to be modelled, how the model might be validated and so on. However, these issues are not so much of an epistemological stumbling block as they might appear. Once the process of modelling has been accomplished, the model achieves a substantial degree of autonomy. It is an entity in the world and, as much as any other entity, it is worthy of investigation. Models are not only necessary instruments for research, they are themselves also legitimate objects of enquiry. Such “artificial societies” and their value in theorizing will be the concern of the first part of this chapter.

In the second part, we will consider an idea that, although central to social simulation, remains controversial: the idea of “emergence”. Emergence is also a major concern for many fields of investigation adjacent to and intertwined with social simulation; for example, the study of artificial life, biology, and solid state physics. It recurs in many chapters in this book and was addressed specifically during the second Symposium on Simulating Societies, held in Siena, Italy, in July 1993, the meeting at which initial versions of most of the chapters in this book were first presented. We shall argue that social simulation may help to reformulate the question of emergence in more specific and empirically relevant terms.

The study of “possible” societies

The study of natural societies is the original objective of social simulation. The behaviour of a model of a society can be observed “in vitro” and the underlying theory tested. For example, this use of computer simulation lies at the heart of the studies of Mayan “collapse” carried out in the 1960s and 1970s (for a review, see Doran & Gilbert 1994), and Palaeolithic social organization (see Doran et al. 1994; Doran & Palmer in this volume).

During the 1980s, a different use of computer simulation developed within the social sciences. For example, within game theory, computer simulation has been applied to study the differential dynamics of processes and to assess the mechanisms and reasons for social learning; more specifically, the spreading of different social “strategies”, or behavioural patterns. As well as testing theories of society, experimental simulations are run to observe which interactional strategy among several existing in a given population is most likely to survive and spread over time, and why.

Such studies of the dynamics of social processes demonstrate another role for computer simulation. The question here is not “what has happened?”, as in the case of the computer study of the Mayan collapse or even “what might have happened?”, but rather “what are the sufficient conditions for a given result to be obtained?”. While the first two questions are exclusively descriptive, the latter may have prescriptive consequences. It may provide hints about how to enhance or reinforce some social strategies by telling us, for example, under what conditions these strategies become stabilized. Promising studies are being carried out within the social sciences, some AI subfields, operational research, organization theory, management science and so on, with the normative aim of reinforcing socially desirable outcomes or optimizing co-operation and co-ordination.

That this also gives an indirect account of how things are, in fact, or have been, is only a secondary aspect of the work.

This alternative role for social simulation deserves attention. We will not be directly concerned with purely prescriptive studies, although for a variety of reasons the goal of optimization has had a strong impact on the computer-based study of social processes. Another, at least equally important, objective in this work is to realize, observe and experiment with “artificial societies” in order to improve our knowledge and understanding, but through exploration, rather than just through description. In this mode of research, the target is no longer a natural society, but an artificial one, existing only in the mind of the researcher. A new target (the artificial system) is created with its own structure (the architecture of the system) and behaviour (the simulation). When a simulation is run, the system operates in a certain way and displays certain behaviour. The simulation may either provide a test of the model and its underlying theory, if any, or may simply allow the experimenter to observe and record the behaviour of the target system. As the emphasis shifts from describing the behaviour of a target system in order to understand natural social systems the better to exploit the behaviour of a target for its own sake, so the objective of the research changes to the observation of and experimentation with possible social worlds. With the possibility of constructing artificial systems, a new methodology of scientific inquiry becomes possible.

The value of building artificial societies is not to create new entities for their own sake. Such an approach, although common to much research and debate about artificial life, has only moderate scientific interest. It is exemplified by some philosophical disputes about whether artificial beings may be considered as living (see Emmeche 1994, Harnad 1994, Langton 1994). These disputes are of little relevance because their resolution depends entirely on how the notion of “life” is defined.

Our stress, instead, is on a new experimental methodology consisting of observing theoretical models performing on some testbed. Such a new methodology could be defined as “exploratory simulation”. The exploratory aim synthesizes both the prescriptive and descriptive objectives: on the one hand, as with the testing of existing theories, the aim is to increase our knowledge; but on the other, as happens with studies orientated to the optimization of real life processes, the aim is not to reproduce the social world, but to create new, although not necessarily “better”, or more desirable, systems. Here lies the difference from optimization research.

Exploratory research based on social simulation can contribute typically in any of the following ways:

- (a) implicit but unknown effects can be identified. Computer simulations allow effects analytically derivable from the model but as yet unforeseen to be detected;
- (b) possible alternatives to a performance observed in nature can be found;
- (c) the functions of given social phenomena can be carefully observed (we will return to this issue later); and
- (d) “sociality”, that is, agenthood orientated to other agents, can be modelled explicitly.

In the next section, we shall show why simulation is of particular value for the development of social theory about sociality.

Artificial societies and artificial sociality

Notwithstanding the importance of its theoretical contributions, the study of society has always been plagued with some crucial methodological and empirical difficulties. For example, social theorists often fail to provide sufficiently specific models, categories of analysis and conceptual instruments for the exploration of social reality. As a result, a wide gap continues to exist between empirical research and theorizing.

In the past, sociologists used to embrace distinct approaches and combine suggestions and intuitions from other disciplines in highly comprehensive systems of thought (consider, for example, the role of biological and evolutionary models in structural-functionalism in general, and the impact of the psychological and psychoanalytic literature on Parsons' theory of action in particular). Nowadays, sociologists have all but abandoned the creation of large, comprehensive theories. However, many also seem to have relinquished ambitions to engage in any form of modelling of social reality. Dissatisfied with the poor predictive or explanatory power of theories of action, sociologists have taken refuge in claiming that social reality cannot be described scientifically because it is construed by and only accessible through the minds of heterogeneous agents. Consequently, social thinkers are said to be entitled, at most, to gather and interpret lay people's self-reports about their actions. And often the warrant for these interpretations is far from clear.

An exception to this general trend is represented by those approaches to the study of sociality that are grounded in mathematical terms, such as game theory. Within this area, there is active theorizing and empirical work. Not surprisingly, however, these approaches share a strong methodological individualistic assumption; that is, the idea that there is nothing about societies that cannot be said about the microinteractions among its members (see Chapter 8). Consequently, the specific aim of sociological inquiry, the study of societies, is essentially nullified.

Social simulation studies provide an opportunity to fill the gap between empirical research and theoretical work, while avoiding the individualist tendency of most mathematically-based approaches. In particular, social simulation provides not only a methodology for testing hypotheses, but also an observatory of social processes. It can therefore offer the basis for new efforts to devise categories of description and new analyses of social reality. In other words, social simulation can provide instruments for modelling sociality.

A couple of examples will help to clarify what is meant by the modelling of sociality as opposed to developing large comprehensive social theories, on the one hand, and atheoretical, empiricist research on the other.

Between co-operation and conflict

Usually, social action is considered to be a choice between co-operation and conflict (think of the structure of the "Prisoner's Dilemma", where each player is faced with the choice of whether to renounce part of the immediate reward to let the opponent have a share of the cake, or else get as much as possible and forget about the other players' outcomes). However, such a simple dichotomy is fallacious, for several reasons:

- (a) there are many forms of co-operation and conflict, as well as many forms of altruistic and selfish, desirable and undesirable social action;
- (b) the alternatives are not clear-cut: there are forms of apparently cooperative action that may be regarded as being conflictual. Social exchange, for example, although notionally a form of co-operation, may give rise to manipulation and produce inequality;
- (c) many concrete examples of sociality are neither advantageous nor harmful to co-operation. For example, conversation is itself neither intrinsically conflictual nor co-operative. Furthermore, many abstract categories of social action are neither pro- nor anti-social. Consider, for example, communication, which may be either co-operative or aggressive; and influence and persuasion, which may be either selfish or altruistic; and
- (d) a social action may be governed by a complex hierarchy of goals so that a given action's lower-level goals may differ from the higher-level ones. A pro-social action may be included in an anti-social plan and vice versa. Deception may be used for helping, and co-operation for cheating.

Natural interaction is undoubtedly much richer than is allowed by the basic categories of co-operation and conflict. To understand and model social action and account for the complex and varied forms of interaction means working out subtler categories of analysis than are usually employed by social thinkers. The fine-grained modelling of sociality allowed by the study of artificial society could greatly enhance our capacity to account for the complexity of interaction.

Self-sufficient agents and social interference

Often, social agents are thought of as self-sufficient beings acting in a common world. The term applied by some social scientists to describe structural social relations is "interference", essentially referring to agents accidentally hindering or facilitating one another. This view is limiting, since it does not account for the role of the structures that pre-exist interaction (Conte & Sichman, in press). It does not show that certain forms of social action are present *in nuce* in certain types of structural relations and emerge from them. A detailed description of different types of social structures would allow the relationships between social structures and social action to be modelled in more revealing ways.

Groups and coalition formation

The classical view of groups and coalition formation is based on two ideas: first, co-ordination to overcome interference and, secondly, collaboration involving social intentions and beliefs.

Co-ordination

If agents are conceived to be self-sufficient beings, once they are in a social context they will find constraints to their self-fulfilment. The common social environment and the

other agents limit their autonomy and achievements. This may be ameliorated through co-ordination of the agents' actions.

Collaboration

Co-ordination alone is insufficient for any sort of coalition or agreement to take place among rational agents. Agents also need to have some degree of awareness of others' existence, wants and habits. In other words, agents need to have some mental representation of other agents' minds. These representations can be divided into two main categories:

- (a) Strategic beliefs (in the sense of strategic rationality), which consist of individuals' mental states produced by, and taking into account, the mental states of other, possibly interfering, agents.
- (b) Group beliefs and intentions (Gilbert 1987, 1989; Tuomela 1991, 1992), which consist of the sharing of similar beliefs and intentions by an aggregate of agents. Social thinkers have been debating what exactly should be meant by "we-ness" and "we-intentions", and where precisely the core of groups and teams resides, without reaching any satisfactory agreement.

However, these two ingredients are not in themselves sufficient to account for coalition formation. They do not allow the gap between the individual level of agency and the group or collective level to be filled. Or, when some attempt in this direction is made, as happened with the notion of strategic knowledge provided by game theorists, only a simplistic, reductionist view of social mental states could be obtained. What is lacking is a fine-grained modelling of sociality that allows for the variety of social beliefs and goals to be taken into account. Agency should not be seen as a twofold, individual and societal, phenomenon, but as a multi-level one, where individual and societal levels are integrated because of the special make up of the agents—their social characterization. It is this level of sociality that throws a bridge between individual and collective action.

In sum, the fine-grained modelling of sociality is extremely important, both as an independent objective of scientific inquiry and as an instrument and a condition for the study of societies. Since sociality represents the medium between individuals and collectives, the study of various forms and levels of social organization cannot do without the study of sociality. However, modelling of sociality requires:

- (a) detailed and explicit models of social agents;
- (b) an observatory of various forms and levels of sociality, not predetermined and constrained by existing theories and assumptions;
- (c) the possibility of experimenting with the observed social phenomena; and
- (d) the possibility of observing the emergent properties of the phenomena.

None of these requirements can easily be fulfilled without computational models of social agents and without computer simulations of social processes.

A central issue: emergence

“Emergence” is one of the most interesting issues to have been addressed by computer scientists over the past few years and has also been a matter of concern in a number of other disciplines, from biology to political science. The central idea is that systems of great complexity (complexity being measured along a number of dimensions, such as differentiation, organization, size, etc.) survive and reproduce thanks to their capacity to adjust, spontaneously and gradually, to the pressures of their environments by achieving varying types and levels of self- (that is to say, emergent) organization. In particular, the question of emergence is central to the field of social simulation and seems to provide a bridging notion across its component disciplines and subfields. However, the notion of emergence, perhaps because of its significance for all these disciplines, remains vague and ill defined.

Levels and types of emergence

Emergence could represent a unifying concept allowing different phenomena to be compared and different approaches to be contrasted. However, in current social and computational studies, the notion of emergence seems to have several meanings, not all of which are related. Cristiano Castelfranchi argued (during the Siena conference on which this book is based) that at least the following senses of emergence need to be kept distinct (although they might be intertwined empirically):

- (a) *Diachronic or evolutionary*: starting from some forerunners, a phenomenon reaches maturity over time.
- (b) *Gestalt-like or layered*: reality is seen as a multi-layered set of phenomena, with different levels of complexity. At some higher level, emerging properties (be they aggregational, collective, relational or Gestalt-like) might be observed which cannot be detected at lower levels of complexity. The emerging properties that are assigned to the higher levels cannot be attributed to the formal elements belonging to the lower levels.
- (c) *Representational*: structures (e.g. social structures) may variously affect phenomena at a given level of analysis B without being represented (known, learned) at the higher level A . However, at some point later in time, agents may acquire a representation at level A . We could call emergence the process by which a given phenomenon is learned or recognized at the cognitive level (e.g. the agents becoming aware of given objective relations occurring in their social world).
- (d) *Adaptive*: emergent properties are often meant as adaptive or functional, since they increase the fitness of the overall system to which they belong. This, one of the most widely used meanings of “emergence”, raises a host of problems that resemble those once encountered by functionalist explanations. In practice, the notion of emergence is often confined to the description of “positive” effects such as self-organization, a usage that comes close to the idea of function, i.e. a feedback mechanism operating on actions and responsible for the regulation and evolution of systems.

Some problems with emergence

Existing approaches to emergence share one or more of the following assumptions or features, all of which seem unwarranted.

Sub-cognitive bias

The idea of emergence has most often been applied to systems with subcognitive units (for example, reactive systems in AI and subsymbolic systems working on a neural network base). Emergence is meant to have the same role for the subcognitive as knowledge and calculation have for the cognitive. But emergence can also be significant among collections of cognitive agents. The following effects can emerge from interacting cognitive agents (see also Conte & Castelfranchi in press):

- (a) the emergence of a “structure” of objective relationships among a set of unaware agents from their individual goals and abilities once agents are placed in a common world;
- (b) the emergence of awareness by agents of these precognitive relationships because of, for example, adaptive cognition (learning from failures);
- (c) the “evolutionary” emergence of internal goals, complex representations and minds from environmental pressures (fitness) exerted on behavioural systems;
- (d) the emergence of collective activity, teamwork and group intentions from complementarity of goals and actions, sharing of beliefs, etc.; and
- (e) the emergence of co-operation from the functional effects of the agents’ deliberate actions without the agents becoming aware of it.

The behavioural bias

Even when applied to rational agents, it is usually only at the behavioural level (the spreading of given behavioural patterns and strategies, such as co-operative or defective choices, the learning of seemingly planful routines in fundamentally reactive systems, etc.) that behavioural properties are noted; non-behavioural emergent effects (e.g. cognitive structures) are usually ignored. For example, the idea that cognitive structures (the capacity for knowledge-based reasoning, planning, decision-making, etc.) represent a highly adaptive emergent response developed by complex systems under environmental pressure has not received the attention it deserves.

The individualistic bias

The study of emergent social phenomena tends to imply an individualistic approach, tacitly adopting one of the traditional solutions to the micro-macro problem, namely that phenomena at the macro-level can be said to emerge from phenomena at the micro-level of analysis and interpretation. The micro-macro link is therefore reduced to one-way (from micro to macro) relationships. However, individual social action can be forged by macro-social forces of various sorts, for example, social norms. Of course, one may claim that a norm is only a convenient term to describe a state of affairs that has no real correspondence in the minds and actions of social agents. Alternatively, one may claim

that social norms correspond to specific phenomena that are produced spontaneously and gradually by interactional practice. Then social norms (macro effects) tend to constrain not only the actions of agents, but also their minds and beliefs, their goals and even their evaluations, expectations and attitudes (a micro effect). The formation of new mental constructs can be a by-product of such norms. Over time, these mental constructs tend to form part of the social characterization of the agents. What are they but emerging micro-properties of macro-social phenomena?

Reconsidering emergence

A unifying and better-defined notion of emergence could be obtained if the above biases were avoided. A unifying notion of emergence cannot decide to account only for properties of subcognitive systems in interaction. Nor can emergent properties be reduced to observable, behavioural effects (leaving aside internal, or mental, structures) at the micro-level, ignoring the fact that actions at the micro-level can be considered to be the emergent effects of phenomena at the macro-level.

It is also necessary to find criteria that allow arbitrary, accidental or irrelevant emergent effects to be discriminated from relevant ones. This brings us back to previous functionalist debates. The puzzles left unsolved by structural-functionalist theory may be concealed under the apparently new terminology of emergence. However, the methodology of computer simulation offers an opportunity for the social sciences to address some of these puzzles. Within such a methodology, a functional interpretation of a given phenomenon can be tested and *falsified*, since the three main questions that should be addressed within a functional explanation: (a) what the limits are within which a given effect is functional; (b) who the beneficiaries of the functional effects are; and (c) what the mechanisms of reproduction and selection are (cf. Elster 1982), through which a given effect becomes functional, can finally be answered by means of computer simulation.

To summarize, the importance of emergence is that interacting agents in a common world produce phenomena that are not necessarily intended by the agents themselves but are none the less relevant to their later behaviour and achievements. Social simulation represents a testbed for the study of emergence. A scientifically adequate exploration of emerging phenomena is not practicable in the absence of such a methodology.

Artificial societies

The role of simulation in evaluating models of both natural and artificial societies and the importance of the concept of emergence will become apparent in the remaining chapters of this book. The chapters in the first section, on the simulation of social theories, focus on building models that articulate social theories in areas of the social sciences which previously have been the preserve of more traditional methods of analysis. For example, in Chapter 2, Robert Axelrod examines how political actors can emerge from aggregations of smaller political units. He points out that the major research paradigm for formal models in politics has been game theory, but game theory takes the actors involved as a given. He shows how clusters of actors that behave like independent

political states emerge from the interactions of basic units during the simulation of a “tribute” model in which the basic units compete with each other for “wealth”.

Massimo Egidi and Luigi Marengo, in Chapter 3, also use simulation to shed new light on a topic on which much has previously been written: the division of labour in organizations. They model a simple production process: the counting of a heap of banknotes by a group of agents. Each agent is represented by a form of program called a classifier system which is able to evolve the rules it uses to determine how many banknotes it will count and where it will pass on the result. They test the simulation in a variety of conditions and show that while a centralized and hierarchical division of labour is most efficient when the agents are infallible, if the agents make mistakes in counting, a decentralized control mechanism that permits local adjustments is better.

In Chapter 4, Jean Pierre Treuil applies social simulation to a problem of social anthropology. He develops a model of the emergence and transmission of kinship structures in Australian aboriginal societies, using simulation as an alternative to the more usual mathematical models, in order to show the conditions under which the kinship system converges to an equilibrium pattern of unions. In Chapter 5, Stéphane Bura and her colleagues describe a simulation located within the conceptual framework of central place theory that aims to account for the location of human settlements. They argue that simulation has advantages over other types of modelling, such as using differential equations, because it allows a much greater variety of qualitative and quantitative factors to be brought into play. Their model is used to study the genesis, development and concentration of urban functions during the long-term evolution of settlement systems. They show that despite the highly varied local conditions in which settlements are located, there is a remarkable convergence in the pattern of settlements in the system as a whole.

While all the previous chapters have reported simulations in which the agents are intentionally very simple, having at most a few rules to determine their interactions with other agents, Jim Doran and Mike Palmer in Chapter 6 describe a simulation designed to explore the growth of social complexity in societies in the Upper Palaeolithic period in south-western Europe in which the individual agents are comparatively complex. This is because the agents’ cognitive representations are modelled explicitly. The authors argue that it is essential to model humans’ unique cognitive abilities and show that, by doing so, their simulation illustrates the emergence of leaders, these being the agents that seem to offer the best plans for action to the others.

In the final chapter of this section, Roger McCain suggests that it is possible to model consumer demand in microeconomics by using simulations based on genetic algorithms. He investigates whether the agents eventually learn optimal behaviour, that is, whether they are able to maximize their utility, thus tackling from a new perspective one of the foundational problems of microeconomics.

In the second section of the book, the contributors consider the idea of emergence from a number of different perspectives. Nigel Gilbert in Chapter 8 reviews the debate in sociology between methodological individualism, the idea that macro phenomena must be accounted for entirely by the actions of individuals, and methodological holism: the idea that individuals’ behaviour is to be entirely explained in terms of their locations in society. He suggests that a more recent position, “structuration”, is more illuminating than either of these extreme positions. Comparing the theories underlying most

simulation models of societies with structuration theory, he notes that so far simulations have failed to take account of the fact that people have the ability to discover and monitor the emergent features of their own societies. He argues that this ability has major implications for human societies and needs to be considered in developing future models.

Edwin Hutchins and Brian Hazlehurst in Chapter 9 describe a simulation that illuminates the emergence of natural language. They point out that for people to communicate, there needs to be a shared lexicon in which each item has an agreed form and an agreed meaning. The puzzle is how this lexicon could have evolved in human societies from nothing. They develop an artificial society consisting of a number of neural nets, warrant it in terms of a theory of cognition and then present a simulation in which the interacting neural nets evolve a common lexicon for describing visual scenes. The simulation is even able to reproduce the emergence of dialects among some subgroups of nets.

In Chapter 10, Alexis Drogoul and his colleagues describe the work they have been doing on the simulation of the social organization of an ant colony. They report the results of their simulation of sociogenesis, that is, the creation of a society from one individual, the queen, and compare the macro-level, emergent behaviour of an artificial ant nest with records of the growth of natural nests, showing that the two are similar. In a further set of experiments with their artificial nests, they simulate the effect of having several queens and observe that although the queens are programmed to be in competition with each other for resources, the resulting behaviour can easily be interpreted in terms of co-operation and altruism. They conclude that many behaviours viewed as being co-operative can be obtained by competitive interplay between agents.

Nicholas Findler and Raphael Malyanker in Chapter 11 return to the theme of alliance formation considered by Axelrod in Chapter 2, but from a different perspective. Their emphasis is on the development of norms and in particular on the mechanisms by which norms evolve and are implemented in international power politics. They use the techniques of distributed artificial intelligence (which is also the starting point of the work of Doran and Palmer, and Drogoul et al.) to model agents that can reason about their environment and interact with each other.

In the third and final section are three chapters that address some fundamental issues in social simulation. In all three chapters, the principal concern is not to model some natural society or even an approximation to a natural society, but to develop methods and ideas that can be applied to both artificial and natural societies. In Chapter 12, Domenico Parisi et al. describe simulations of the evolution of altruism and attachment behaviour in a population of neural networks and show that it is possible to observe the emergence of a type of behaviour that tends to keep individuals close together, thus creating what might be described as a simple social group. Rosaria Conte and Castelfranchi in Chapter 13 examine the role of norms in controlling aggression among competing agents that have some limited ability to reason. They carry out a series of simulation experiments on a simple artificial society in which there is competition for “food”, to compare the effects of individual rational strategies and societal norms. They find normative controls result in a more equitable distribution of resources in the population.

Social simulations have usually been programmed using either conventional computer languages such as C and Basic or languages developed for artificial intelligence research, such as LISP, Prolog and Smalltalk. In the final chapter, Michael Fisher and Michael

Wooldridge describe a novel language for implementing models of societies. In this language, individual agents are programmed by giving them a specification of their desired behaviour in terms of temporal logic. Fisher and Wooldridge demonstrate the use of the language through examples and show how a variety of behaviour can be specified relatively simply and concisely compared with using traditional languages.

The chapters in this book stand as an illustration of the range and vitality of current work on social simulation. They draw on and contribute to a broad sweep of disciplines and show how increasingly the most promising intellectual developments are those that span disciplinary boundaries, not only within the natural sciences and social sciences, but also across these broad divisions. As editors, we look forward to a continuing stream of exciting research arising from the work reported in this book that will deepen our understanding of societies and improve our methodological skills in simulating natural societies and creating artificial ones.

Part I
The simulation of social
theories

Chapter 2

A model of the emergence of new political actors

Robert Axelrod

How can new political actors emerge from an aggregation of smaller political actors? This chapter presents a simulation model that provides one answer. In its broadest perspective, the work can be seen as part of the study of emergent organization through “bottom-up” processes. In such “bottom-up” processes, small units interact according to locally defined rules, and the result is emergent properties of the system such as the formation of new levels of organization. Thus the work is typical of the “Santa Fe” approach to complex adaptive systems (Stein 1989; Fontana 1991; Holland 1992). The concern with increased levels of organization is also reminiscent of how biological systems succeeded in making the transition from single-celled organisms to multiple-celled organisms (Buss 1987), and how brains function by organizing individual neurons into meaningful structures (Hebb 1949; Minsky 1985).

The task at hand involves the emergence of new political actors. This is a vital question in the post-Cold War world. We are experiencing an era in which the standard unit of politics, the nation, is no longer completely stable. We see, on the one hand, that some states are disintegrating, as in the former Soviet Union and Yugoslavia. We see, on the other hand, that larger units are being organized, such as the European Union (EU), and other regional associations. The question of the aggregation and disaggregation of political actors is essential for the understanding of the future of global politics, both in terms of international security affairs and international political economy.

The question of how the world can be placed on a sustainable path of development is a particularly pressing question. The emergence of new political actors is fundamental to the question of sustainability. One of the main problems of attaining sustainability is the tragedy of the commons (Hardin 1968). The tragedy of the commons arises when many independent actors (people, villages, states, or whatever) each “over-graze” because there is no mechanism to enforce the collective interests of all against the private interests of each. This leads to resource depletion, elimination of bio-diversity, overpopulation, war, and other major social problems. A major route to the prevention of the tragedy of the commons is the emergence of a political actor based upon the organization of previously independent actors. Today we have political actors at the national level that can regulate resource use within their boundaries, but we do not yet have very effective political actors at the transnational level to regulate resource use at the global level.¹

Political scientists have access to a variety of concepts and theories to analyze the emergence of new political actors. Unfortunately, they do not have any formal models that account for this emergence endogenously. In fact, the problem is much like

biologists' interest in the emergence of multi-cellular organisms: research has tended to take for granted the existence of such complex units and therefore has not developed rigorous theories to explain how they might have come about in the first place (Buss 1987). For example, the major research paradigm for formal models of politics is game theory, and game theory takes as given exactly who the actors are in a particular setting. In contrast to the rational choice approach of game theory, the model of this paper uses techniques of complex adaptive systems. It takes as given the existence of the lower-level actors, and generates higher-level actors from the interactions among them.

Given that the goal is to account for the emergence of new political actors, it is important to have a set of criteria which can be used to identify a new actor when one emerges. Here are my criteria for identifying a new political actor as an emergent set of relationships among previously existing units:

1. Effective control over subordinates:
 - (i) little rebellion; and
 - (ii) no independent "foreign policy".
2. Collective action ("all for one and one for all"):
 - (i) paternalism (protection of the weak by the strong); and
 - (ii) joint foreign policy.
3. Recognition by others as an actor.

This list is largely inspired by historical and contemporary international law and practice concerning the recognition of new states in the world community of nations. For example, the thirteen colonies became a new nation called the United States of America when the central government was established which had:

1. Effective control over the individual states:
 - (i) with only low levels of rebellion (at least until the Civil War); and
 - (ii) with a prohibition on treaties between the individual states and other nations.
2. Collective control over some important resources:
 - (i) that allowed federal interests to dominate state interests at least in certain important domains; and
 - (ii) that allowed for a common foreign policy in matters of war and peace.
3. Recognition of its independence by other nations such as France and Britain.

While the emergence of the United States is hardly typical, it does illustrate the essential properties of how a new political actor can result from the aggregation of smaller political actors. In the American case, the aggregation was of colonies whose own autonomy was short-lived and precarious during the revolutionary period. The more typical case, and the one modelled here, is where the emergent actor is essentially an empire in which a core unit dominates its neighbours. A good example of this process is the growth of the Russian (and later Soviet) empire from a small core around Moscow. Other examples include the growth by accretion of neighbouring lands by Rome and China.

Coercion in general and extortion in particular have played a central role in the formation of states over the centuries (Tilly 1985; 1990). For this reason, at the heart of the present model is the simple dynamic of “pay or else”. An elementary actor can make a demand of a neighbour for payment of resources, with the threat that if payment is not forthcoming there will be a war.

In the tribute model wars result in changes in wealth (and thus power), but not in outright territorial conquest.² This allows each territory to maintain itself as a separate actor, thereby allowing for the possible emergence of sets of territorial actors who might form a new aggregate actor. Although no territory changes hands, war results in costs to both sides, but especially to the weaker side. Thus the heart of the model is a tribute system in which an actor can extract resources from others through tribute payments, and use these resources to extract still more resources. Alliances are also allowed so that actors can work together. Whether a set of elementary actors emerges as a stable aggregate actor will depend on the dynamics of the tribute system: whether groups of actors emerge which function as a single aggregate, and whether alliance patterns emerge which lead to stable, coordinated actions.

Unlike my earlier work on the Prisoner’s Dilemma (e.g. Axelrod 1984) the tribute model is based upon extortion rather than co-operation. Also, unlike the earlier work, it does not assume that the actors are equal in power, but instead the tribute model takes power differences as vital, and as resulting directly from the dynamics of the model. Nor does it assume that the actors interact only two at a time. Like my earlier work, it is not a rational model. The tribute model assumes that actors develop more or less strong commitments to each other based upon their prior actions. These commitments can be thought of as the result of psychological processes (e.g. to stand by those who have helped you in the past), or the result of political rules of thumb (e.g. to support those who may later help you in your time of need). Actions are based upon simple decision rules rather than game theoretic calculations of optimal choice, since rational calculations would be virtually impossible to make in such a complex setting. An important and novel feature is that the behaviour of the actors changes over time as they apply simple decision rules to data about the historical experience they have gathered through prior interactions.

The project will be successful to the extent that it can account for important phenomena that cannot be accounted for by existing formal theories, and can do so with only a few assumptions. The main result to be generated is the emergence of new political actors at a higher level of organization, and the main assumptions have to do with how elementary actors interact with each other. As we shall see, the model generates spontaneously not only emergent actors, but also other interesting phenomena that exist in the world of international politics.

The tribute model

The basic units of the model are ten actors arranged on a line. The actors can be thought of as independent political units, such as nations. The reason for having so few actors and having such a simple geography as a line (rather than a two-dimensional space) is to keep the data analysis as simple as possible. As long as no vital conceptual features are lost, simplicity is the best modelling strategy. To avoid introducing arbitrary distinctions

between actors on the ends of the line and those in the interior, the line is joined into a circle. This gives each of the ten actors only two neighbours.

There is one resource in the model, called “wealth”. Each actor is given some initial wealth. The initial wealth is chosen from a uniform distribution between 300 and 500. These parameters, like all the others in the model, are somewhat arbitrary and are selected for convenience.

The basic cycle of the model is called a “year”. In each year, three actors are chosen one after another at random to become active. An active actor, *A*, may demand tribute from one of the other actors, *B*. Initially, the target, *B*, must be a neighbour, but later this restriction will be relaxed when alliances are considered. The model is based upon a dynamic of “pay or else”. The target, *B*, then has a choice of paying tribute to the demander, or fighting.

The selection of actors to be active is based upon the notion that ambitious leaders and potential disputes arise at random. But once given an opportunity to make a demand, the activated actor need not actually make a demand if it finds that the current situation is not favourable.

If *A* does make a demand, the target *B* has a choice.

- If *B* pays, wealth is transferred directly from *B* to *A*. The amount of wealth transferred to *A* is 250 if *B* has that much. Otherwise, the amount transferred to *A* is whatever *B* has. This transfer represents tribute that *B* pays to *A* to avoid a fight.³
- If *B* fights rather than pays, each side loses 25 per cent of the *other* side’s wealth (or both lose proportionally less if either side doesn’t have that much wealth).⁴ This is a simple Lanchester attrition dynamic (Lanchester 1916; Epstein 1985). The idea is that in a fight both sides suffer, but the stronger side imposes more damage than the weaker side does.⁵

After the three activations, the yearly cycle ends with a “harvest” which increases each actor’s wealth by 20. This feeds some new wealth into the system. A typical run is 1000 years.

The next question to consider is decision rules used by the actors for making and responding to demands. It would be almost impossible for actors to develop fully rational rules in the tribute game because of its complexity over time and space. Therefore, the model uses heuristic decision rules that capture some of the key short-term considerations facing the players.

- The active actor needs to decide of whom, if anyone, to make a demand. The ideal target of a demand is weak enough so that it might choose to pay rather than fight, and so that it will not cause much damage if it does choose to fight. On the other hand, the ideal target should be strong enough to be able to afford to pay as much possible. A suitable decision rule combining both of these considerations is to choose among the potential targets the one that maximizes the product of the target’s vulnerability multiplied by its possible payment. The target’s vulnerability is $(W_A - W_T)/W_A$ where W_A and W_T are the wealths of the active actor and the target, respectively. The target’s payment is how much the other side can pay, which is the maximum of its wealth and the full tribute of 250. If no potential target has positive vulnerability (i.e. no potential target is weaker than the demander), then no demand is made.

- The decision rule used for the target is simpler: fight if and only if it would cause less damage than paying would.

So far, the model has considered only pairwise interactions. But in order to study the development of emergent actors, there has to be a way for the basic actors to work together. The key idea is that actors develop degrees of commitment to each other. These commitments are caused by their choices to pay or fight, and in turn have consequences for how they will pay or fight in the future. The basic idea is that if two elementary actors fight, another adjacent actor will join the side to which it has greater commitment. If it has equal commitment to the demander and the target, it stays neutral. If it does join one side or the other, it contributes forces (i.e. wealth) in proportion to its commitment to that side.

Initially, no one has any commitments to others, and each actor is fully committed to itself. Commitment of actor i to actor j increases when:

- (a) i pays tribute to j (subservience);
- (b) i receives tribute from j (protection); or
- (c) i fights on the same side as j (friendship).

Similarly, commitment decreases whenever:

- (d) i fights on the opposite side to j (hostility).

The motivation for these rules of commitment dynamics is simple. When one actor pays tribute to another, the first actor is also likely to be partially under the second actor's political domination, and therefore compelled to assist the second actor next time. A state typically becomes committed to helping the patron, whether by choice or necessity, as illustrated by the participation of many Latin-American states in the Second World War after Pearl Harbor. Conversely, the protection of states that have provided benefits is also commonplace, to protect future sources of revenue. The next point is that if two sides have fought together in the past, they tend partially to be committed to each other for the future, as in the case of the United States and South Korea after the Korean War. On the other hand, two states that have fought each other are less likely to support one another in the future, as in the case of Germany and France after the First World War.⁶

The model assumes that increases and decreases of commitment are in constant amounts, namely increments of 10 per cent. In addition, one actor's commitment to another can never be more than 100 per cent nor less than 0 per cent. The commitment processes described above maintain the symmetry of commitment. This is because two actors start with no commitment to each other and their commitments to each other always grow and decline in unison. For example, if two actors fight on the same side in a war, then their commitment to each other will increase by 10 per cent.

The final part of the model deals with co-ordination of actors. To keep interactions similar to land combat, co-ordinated action is assumed to require contiguity. Thus an actor is an eligible target for a demander only if everyone between them joins the demander. Others may then join the demander or the target provided that the intermediate actors also join in: for example, the actor in position 5 may make a demand on actor 8 only if actors 6 and 7 join actor 5. This requires that actors 6 and 7 are both more committed to 5 than to 8. Then, if 5 does make a demand on 8, actor 10 may join the defence of 8 only if 9 joins too.

Commitments and wealths are common knowledge.⁷ Thus when an active actor is evaluating the vulnerability of an eligible target, it can take into account the commitments and wealth of all actors who would join either side. Similarly, the target of a demand can determine the cost to itself of fighting, by calculating the damage that the attacking alliance could do, and the proportion of that damage that the target would suffer, which is the proportion of the defending alliance's wealth contributed by the defender.

The dynamics of the tribute system

As with any stochastic simulation, the tribute model needs to be run many times to explore the types of history it generates. In the case of the tribute model there are two sources of random variation: the initial wealths of the actors, and the selection of three actors each year to become active. A good way to begin to appreciate the dynamics of the model is to look at history from the perspective of a single actor, say actor 5. This reveals the sequence of events (demands, payments and fights) that it takes part in. The history also illustrates how alliances are built up from the events that cause commitments.

For example, here is the first few years of history for actor 5 in the first run: in year 9, 5 was active and made its first demand, a demand on a neighbour, actor 4. Finding payment would be more costly than fighting, 4 chose to fight. In the resulting fight each lost wealth, but the defender was weaker and thus lost more than the attacker. In the next year, 5 was again active, and again made a demand on 4. Having been weakened by the previous fight, this time 4 found it cheaper to pay tribute. The payment not only helped 5 at the expense of 4, but also led to a mutual commitment of subservience and protection between 4 and 5 at the 10 per cent level. The next time that 5 became active (year 11), it was able to make a demand of 3 (with 4's support); 3 chose to pay, making it also partially committed to 5. In year 14, 5 became active again and targeted 2. However, 2 refused to pay, so 5 fought with the help of 3 and 4, both of whom were partially committed to 5. Since 3, 4 and 5 all fought on the same side, they each became 10 per cent more committed to the others. And since they each fought on the opposite side from 2, they would have become 10 per cent less committed to 2 had they had any commitment to 2.

Having looked at some event-by-event history, it is time to examine the big picture.




Figure 2.1 Wealth of each actor over 1000 years (population 1).

Figure 2.1 shows the wealth of each of the ten actors over the long timespan of 1000 years. The figure clearly shows that three actors were




Figure 2.2 Fights (bars) and population wealth (line) in population 1.

wealthier than the others over most of the history, with actor 5 reaching the highest wealth. A close examination of the figure also shows that something dramatic happened near the end of the 1000 years: a sudden decline in all three of the wealthy actors, especially in the biggest one. Figure 2.2 shows the number of fights in every 25-year period, and also shows the total wealth of the entire population over time.

There were many fights at the start but then fewer, with a resurgence of fights in the last 100 years. The resurgence of fights corresponds to the collapse of the population’s wealth at around year 900. This kind of major collapse happened only once in five runs of the model.

Quite different histories were generated for new populations simulated under the same conditions. For example, Figures 2.3 and 2.4 show that the




Figure 2.3 Wealth of each actor over 1000 years (population 2).




Figure 2.4 Fights (bars) and population wealth (line) in population 2.

second population had four successful actors (rather than three), no fights at all after year 200 (rather than continuing fighting and a late resurgence), and a steady growth of global wealth (rather than a collapse).

Figures 2.5 and 2.6 show the third population had two wealthy actors, an uneven pattern of fights, and several moderate declines in global wealth.




Figure 2.5 Wealth of each actor over 1000 years (population 3).

The fourth and fifth populations showed still other combinations of number of wealthy actors, frequency of fights, and trends in population wealth. Thus, five runs of the same model with the same parameters give quite different histories.⁸ Different combinations of initial wealth and the order in which actors become active can evidently make a large difference




Figure 2.6 Fights (bars) and population wealth (line) in population 3.

in the development of the history of the system.

Still another way of looking at the dynamics of the model is to examine the development of patterns of commitment. Typically, the first pattern to emerge was a set of proximity relationships in which actors developed low levels of commitment to neighbours, and sometimes to neighbours of neighbours. This pattern is illustrated with the pattern of commitments in year 25 of population 2 (see Table 2.1).

Table 2.1 Commitments forming a proximity pattern (from population 2, year 25).

<i>i, j</i>	1	2	3	4	5	6	7	8	9	10
1	100	30	0	0	0	0	0	0	10	40
2	30	100	0	0	0	0	0	0	20	20
3	0	0	100	20	0	0	0	0	0	0
4	0	0	20	100	30	0	0	0	0	0
5	0	0	0	30	100	0	0	0	0	0
6	0	0	0	0	0	100	0	0	0	0
7	0	0	0	0	0	0	0	100	50	0
8	0	0	0	0	0	0	0	100	50	0
9	10	20	0	0	0	0	0	50	100	10
10	40	20	0	0	0	0	0	0	10	100

Note: Actors land 10 are adjacent.

The other common pattern of commitments develops after this, and consists of two clusters of actors. A cluster can be defined as a set of actors, all of whom are highly committed to each other (say at the 50 per cent level). Table 2.2 shows the same population 25 years later, divided neatly

Table 2.2 Commitments forming pattern of two clusters (from population 2, year 50).

<i>i, j</i>	1	2	3	4	5	6	7	8	9	10
1	100	100	70	0	0	0	0	100	100	100
2	100	100	100	0	0	0	0	100	100	100
3	70	100	100	0	0	0	0	100	90	70
4	0	0	0	100	100	60	100	0	0	0
5	0	0	0	100	100	100	100	0	0	0
6	0	0	0	60	100	100	100	0	0	0
7	0	0	0	100	100	100	100	0	0	0
8	100	100	100	0	0	0	0	100	100	100
9	100	100	90	0	0	0	0	100	100	100
10	100	100	70	0	0	0	0	100	100	100

Note: Actors 1 and 10 are adjacent.

into two clusters, with no commitments at all between members of different clusters.⁹

To see how this population can become divided so quickly into two distinct clusters, look at the sequence of payments and fights involving a single illustrative actor, actor 5. Table 2.3 shows the fights and payments involving actor 5 from year 25, when it had only one partial commitment, to year 50, when it was fully integrated into a cluster of four actors.

The events of year 35 can serve to illustrate the dynamics of commitments that lead to the development of distinct clusters. In that year, 3 and 4 fought 5, 6 and 7. This led 5 to increase its commitment to 6 and 7 while decreasing its commitment (if any) to 3 and 4. In general, as fights take place, they increase the commitments of those fighting on the same side, and decrease commitments of those fighting on opposite sides. Thus, as clusters begin to form, then tend to get even more distinct. This is because pairs of actors who fought together became even more committed to each other, and pairs on opposite sides lost whatever partial commitment they may have had to each other. By year 45, fights were taking place that involved all ten actors, leading to the pattern of two strong and distinct clusters that was shown in Table 2.2

Having seen how commitments can lead to clusters in an illustration from population 2, let us now return to the unusual case of population 1, where the strongest actor collapsed. The dynamics over the full 1000 years was shown in Figures 2.1 and 2.2. Now we can look in detail at the period of collapse that took place between years 890 and 950. This is shown in Figure 2.7.

In year 911, actor 10 targeted actor 9, resulting in a “world war” in-

Table 2.3 Development of clusters of commitments. Fights and payments involving individual 5 of population 2 between years 25 and 50.

Year	Active actor	Target actor	Roles	Commitment with actor 5 increases ^a	Commitment with actor 5 decreases ^b
30	4	5	—aAD—		3, 4
30	6	5	—PR—	6	
32	3	4	—ADd—	4	3
32	4	5	—aADd—	6	3, 4
33	4	3	-dDAa—	4	2, 3
34	5	7	—R—P—	7	
35	5	4	—dDAaa—	6, 7	3, 4
36	3	4	-aADd—	7	2, 3
36	8	7	aa-ddDAaa	6,7	1, 2, 8-10
37	6	5	—PR—	6	
38	8	7	aaa-ddDAaa	6,7	1-3, 8-10

38	2	7	aAa-ddDaaa	6, 7	1-3, 8-10
41	8	7	aaa-ddDAaa	6,7	1-3, 8-10
42	6	5	——PR——	6	
42	5	4	——PR——	4	
44	5	7	——R-P——	7	
46	7	8	dddaaaADdd	4, 6, 7	1-3, 8-10
47	7	8	dddaaaADdd	4, 6, 7	1-3, 8-10
48	7	3	ddDaaaAddd	4, 6, 7	1-3, 8-10
48	5	3	ddDaAaaddd	4, 6, 7	1-3, 8-10

Key to roles: A=attacker; a=attacker’s ally; D=defender; d=defender’s ally; P=payer; R=receiver of tribute.

a) Increases by 10%, up to 100%; b) Decreases by 10%, down to 0%.



Figure 2.7 Late collapse of a powerful actor. Payments and fights involving actor 5 (population 1, years 890–950).

volving actors 5–9 versus all the others. Actor 5, the strongest actor, joined in because although it had minimal commitment to actor 9 it had none at all to actor 10. Because 5’s commitment to 9 was only at the 10 per cent level, it contributed only 10 per cent of its forces to the fight and suffered only slightly. But later that same year, 9 again attacked 10, and this time when 5 joined in it contributed 20 per cent of its forces (having become a little more committed to 9 as well as to all the others in that alliance), and suffered a little more. Now 5 was even more committed to everyone in its emerging cluster of 5–9. In years 915–18 there were five more “world wars” with the same alignments (although different demanders and targets). As both sides became more committed to the members of their own emerging clusters, more and more forces were committed and more and

more damage was done. The result was a sharp decline in the wealths of the largest actors (5 on one side and 2 and 4 on the other).

This dramatic decline of the strongest actor is because it was dragged into fights involving weak actors to whom it had developed commitments. This is reminiscent of what happened to two dominant powers, the Hapsburgs in the seventeenth century, and Britain in the nineteenth to twentieth centuries. Paul Kennedy has an apt term for the dangers of the excess commitments that strong powers tend to take on. It is “imperial overstretch” (Kennedy 1987). This is just what is illustrated in Figure 2.7. It is striking that a simple model developed to generate new political actors produced behaviour that has been the subject of a major policy debate in contemporary international affairs (e.g. Nye 1990).

Clusters are not only settings for mutual commitment; they are also settings for extraction of tribute. The strongest actor in a cluster typically stays strong and grows by making demands and collecting tribute from the other members of its cluster. As actors slowly grow due to their annual “harvest”, the strong member makes demands upon the weak and reduces their wealth. In fact, as the strong member selects lucrative targets, its demands tend automatically to be rotated among members of the cluster. This rotation has the unplanned consequence of preventing any target from growing very strong. But sometimes a weak member of a cluster escapes demands just long enough to grow strong enough to make its own demands on still weaker members. If luck lets it be active often enough at this critical stage, it can collect sufficient tribute so that the strongest member no longer finds it advantageous to challenge the newly-wealthy actor. An example of this process occurred after year 100 in population 3 (see Figure 2.8). In this example, actor 5 has managed to grow in the shadow of




Figure 2.8 Internal conflict. Growth of a second strong actor in a cluster of actors 4–8 (population 3).

the strongest actor of its cluster, actor 7.

Demands within clusters do not always get resolved by payments for tribute. In fact, civil wars can occur, i.e. fighting among members of the same cluster. The strongest member of a cluster typically is strong enough, by taking sides, to prevent such a fight from taking place. If, however, the strongest member is equally committed to the attacker and the defender (perhaps by being totally committed to both), then it would stay neutral. This allows civil wars to occur if the active actor finds a lucrative target within the same cluster, and the target finds it more costly to pay tribute than to fight. The attacker and defender may even find allies among the others if these others are not equally committed to the two sides.

Surprisingly, initial differences in wealth do not matter for wealth in the long run. In five populations of the model there were 50 actors in all, 14 of whom were quite successful, using a criterion of wealth of over 10,000 after 1000 years. Of these 14 successful actors, half had less than average initial wealth and half had more than average initial wealth. One might expect that initial wealth would be a great advantage, since it would allow an actor to make successful demands on its neighbours at first, and thereby build up a dominant position in a cluster. But having substantial initial wealth can also make an actor a lucrative target for other strong actors. In addition, having substantial initial wealth can make one over-confident, given the limited rationality of the decision rules in effect.¹⁰

The results of the model's performance can now be summarized in terms of six characteristics:

1. *Things usually don't settle down.* Instead, the history of the model shows considerable complexity. For example, as late as year 900, one of the populations suffered a series of fights that destroyed over three-quarters of the global wealth (see Figure 2.2).
2. *Histories show considerable variability.* The combinations of wealthy actors, frequency of fights, and trends in population wealth differ considerably from one population to another. Even though each actor is using the same decision rules, the results differ greatly because of random differences in initial wealth and in the order in which actors become active.
3. *"Imperial overstretch" can bring down even the strongest actor.* As an actor becomes committed to others via tribute relationships and fighting together, it becomes exposed to the risks of whatever fights the others get into. Since actors decide on demands and responses based upon their own calculations, even a weak actor can choose to make or resist a demand for its own reasons, and thereby drag into the struggle a strong actor who is committed to it.
4. *Civil wars can occur among the smaller members of a cluster.* While the strongest member of a cluster typically can prevent a fight among members of the same cluster by taking sides, it would not do so if it had equal commitment to the two sides. Therefore, smaller members of a cluster may fight each other while the strongest member stands aside.
5. *A cluster can have more than one powerful member.* Clusters are often like empires, with one powerful actor and many weaker ones who pay tribute to it. But as we have seen in the case of Figure 2.8, it is also possible for a second actor to grow strong in the shadow of the first.
6. *Initial endowment does not guarantee or even predict success.* Before clusters of commitments are established, wealth can be as much a handicap as an asset. The

reason is that wealth makes an actor a lucrative target for other strong actors, and can make one over-confident in making demands.

Have new actors emerged?

Returning to the original goal of the project, we can now assess the extent to which the emergent clusters of mutually committed actors are really new actors at a higher level of organization. Recall that a cluster is defined as a set of actors who are each committed to the others at a level of at least 50 per cent. This definition provides a set of candidates who may or may not be emergent political actors. To see if clusters in fact behave as newly-emergent political actors, we return to the original set of criteria.

1. Effective control over subordinates:

- (i) in the simulation, clusters exhibit no rebellion. There were no fights in a cluster against the strongest member. Note that this result is not built into the model's assumptions. After all, it would certainly be possible for several members of a cluster to join in attacking the strongest member (there were, however, some fights among the secondary actors within a cluster, which were tolerated by the strongest); and
- (ii) independent "foreign policy" occurred on only rare occasions. On these rare occasions, the strongest member *was* dragged into conflict by commitment to a client ("imperial overstretch").

2. Collective action ("all for one and one for all"):

- (i) the strong did protect the weak; and
- (ii) members of a cluster did act together with respect to outsiders, both in attack and defence.

3. Recognition by others as an actor did occur, since joint action of clusters was foreseen and taken into account by outsiders. Note that the model assumes that everyone can anticipate the effect of current alliances. The fact that clusters were stable means that the short-term calculation based on alliances was tantamount to long-term recognition of emergent actors based upon stable clusters.

In sum, the tribute model did produce clusters of mutual commitment, and these clusters did exhibit to a high degree all the attributes of an emergent new political actor. In fact, even the exceptions reflected the kinds of exception present in international affairs, as illustrated by "imperial overstretch".

Variations of the model

One of the advantages of a simulation model is that it is relatively easy to answer “what if” questions about the effects of varying the parameters or the premises of the model. Here is a very brief sample of some of these questions, with answers based on further simulation runs.

1. Does the model settle down after 1000 years?

– No. Even in runs up to 10,000 years there are repeated large wars, and strong actors continue to rise and fall.

2. Are there better decision rules?

– Yes. An example is adding a constraint so that a demand will not be made of a target that would find it cheaper to fight than to pay. An individual using this revised rule does much better than the others. Moreover, if everyone is using the revised rule, any single individual who uses the old rule would do worse than the others.

3. Does it pay to give and receive commitments?

– Yes. An actor who neither gives nor receives commitment does very poorly. Moreover, if just one actor gives and receives commitments (and the others do not, except with that actor), then that actor does very well.

4. What happens if everyone can reach everyone else (as in sea power)?

– In this extreme case, a single powerful actor tends to dominate, and there are relatively dense patterns of commitment. Presumably, in a two-dimensional space, where everyone has more neighbours than in a one-dimensional space but fewer than in the “sea power” case, there might be an intermediate level of dominance and density of commitment.

5. What if wealth does not enter into anyone’s calculations?

– If the demands are made without regard to wealth, and if everyone always fights, then the result is two clear clusters, one of which has virtually all the wealth.

6. Do “islands of commitment” grow?

– Yes. If two adjacent actors initially have just 10 per cent commitment to each other (and all other commitments are zero), then these two actors develop complete commitment to each other, and both tend to prosper.

Conclusion: building stepping stones to theory

The tribute model provides an existence proof that it is possible to use simple local rules to generate higher levels of organization from elementary actors. In particular it shows that a dynamics of “pay or else” combined with mechanisms to increase and decrease commitments can lead to clusters of actors that behave largely according to the criteria for independent political states.

Like most simulation models, the tribute model generates huge amounts of historically rich data. Indeed, a major challenge in any simulation study is to find ways of analyzing

and presenting data that help us to see the forest for the trees. This chapter illustrates four different kinds of historical analysis. In effect, these are four different ways of looking at history.

1. History from the point of view of a single actor: plots of relative wealth of individual actors over time.
2. History from a global point of view: joint plots of wealth and number of fights over time.
3. History from the point of view of the emergence of new actors: commitment matrices revealing emerging clusters of actors.
4. History as “news”: event data, including chronologies of tribute payments and wars fought between alliances.

Perhaps the most useful outcome of a simulation model is to provide new ways of thinking about old problems. In the present case, the need to determine whether or not the model was successful in generating new political actors forced the specification of explicit criteria for recognizing a new political actor should one arise. This in itself is a useful exercise. In addition, the same need led to the development of different ways of viewing history, and especially to ways of analyzing the interaction among emerging clusters of actors.

In the future it would be good to use these conceptual and statistical developments to answer some new questions suggested by the model. For example, the dynamics we have seen in the tribute model suggest the following interesting questions:

1. What are the minimal conditions for a new actor to emerge?
2. What tends to promote such emergence?
3. How are the dynamics affected by the number of elementary actors?
4. What can lead to collapse of an aggregate actor?
5. How can new actors grow in the shadow of established actors?

While no one would believe that the answers to these questions given by the model itself would necessarily be accurate for the real world, a realistic hope is that the concepts developed in generating the answers for the model would be useful in providing new ways of thinking about comparable questions in the real world. Indeed, understanding the dynamics of the tribute model might also suggest specific hypotheses that could be tested with data from the real world. In addition, the study of one model (such as the tribute model) can provide the basis for insights into how even simpler models can be developed that might allow deeper insights into issues of emergent political organization.

Finally, a simulation model such as the tribute model can lead to insights into where there might be policy leverage in the real world. For example, one might be able to identify situations in which slight changes in the system could lead to substantially improved outcomes—in terms of fewer destructive wars or greater co-ordination of political action. If we knew when minor inputs could lead to major gains we would have

valuable clues to consider for policy leverage in the real world—the world where avoiding war and achieving sustainability really matter.

Acknowledgements

I thank Rob Axtell, Arthur Burks, Lars-Erik Cederman, Michael Cohen, Paul Courant, Joshua Epstein, Walter Fontana, Stephanie Forrest, Murray Gell-Mann, Nigel Gilbert, John Holland, John Miller, Melanie Mitchell, Robert Pahre, Rick Riolo, Carl Simon, and Gerard Weisbuch. For financial support I thank Project 2050 through the Santa Fe Institute, the National Science Foundation, and the LS&A Enrichment Fund of the University of Michigan.

Notes

1. Of course, the ability to regulate resource use does not guarantee that the ability will be wisely used. For example, the former Soviet Union had the power to control pollution from its factories, but in the interests of maximizing production it chose not to exercise that power.
2. This contrasts with Cusack & Stoll (1990), whose simulation model uses conquest of territory as the major dynamic.
3. Having a fixed maximum demand is arbitrary, but avoids the need for the actors to calculate what demand they will make.
4. For example, if *A* has 400 and *B* has 300, *A* loses $0.25 \cdot 300 = 75$, and *B* loses $0.25 \cdot 400 = 100$. The wealths after the fight would then be $400 - 75 = 325$ for *A*, and $300 - 100 = 200$ for *B*. Note that the disparity in wealth increases from 100 to 125. If *B* began the fight with only 50, then *B* would lose all 50 (which is half the other damage *A* is capable of doing), and *A* would lose half the maximum damage *B* is capable of causing, i.e. $0.25 \cdot 50 \cdot 50 = 12.5$.
5. The demander is assumed to carry out its implicit threat to fight if the demand is not met. This assumes, in effect, that the demander's need to maintain its reputation is strong enough to maintain the credibility of the threat.
6. For evidence on the friendship and hostility aspects, see Axelrod & Bennett (1993).
7. In international politics, wealth can usually be estimated with rough accuracy, but commitments are harder to predict. There is a rich empirical and theoretical literature on commitment. See, for example, Schelling (1960), Huth (1988), Powell (1990), Bueno de Mesquita & Lalman (1992) and Axelrod & Bennett (1993).
8. A run takes only 23 seconds on a Macintosh Quadra 700 using Pascal. Of course, the analysis can take days or weeks.
9. In many cases, there were two large clusters, with one or two actors participating in both the clusters.
10. For example, suppose an actor with wealth of 500 makes a demand on a neighbour with 400. The neighbour will fight rather than pay (losing $0.25 \cdot 500 = 125$, which is less than payment of the standard demand of 250). Since the target fought, the demander loses $0.25 \cdot 400 = 100$, and is worse off by having made the demand. Had the initial wealths been greater than 1000, this would not happen, since the cost of fighting would be greater than the standard demand of 250.

Chapter 3

Division of labour and social co-ordination modes: a simple simulation model

Massimo Egidi and Luigi Marengo

I still believe that, by what is implicit in its reasoning, economics has come nearer than any other social science to an answer to that central question of all social sciences: How can the combinations of fragments of knowledge existing in different minds bring about results which, if they were to be brought about deliberately, would require a knowledge on the part of the directing mind which no single person can possess? To show that in this sense the spontaneous actions of individuals will, under conditions which we can define, bring about a distribution of resources which can be understood as if it were made according to a single plan, although nobody has planned it, seems to me an answer to the problems which has sometimes been metaphorically described as that of the “social mind”. (von Hayek 1948, p. 54)

It is customary to consider economic organizations as social systems which make the co-ordination of individual plans and decisions possible. However, little attention has been given to the connection between forms of coordination and the process which characterizes industrial societies: the process of increasing division of labour.

The fact that co-ordination and the division of labour have been somewhat connected in their historical development is hardly possible to doubt: as economic organizations increased their degree of division of labour and knowledge, the problems of co-ordination among a growing number of increasingly inter-related producers and decision-makers became more and more complex. Co-ordination of distributed decisions by markets, firms and other economic institutions appears as the other side of the process of the increasing division of labour.

The historical evidence does not show clearly whether the two processes of co-ordination and the division of labour have co-evolved as two aspects of the same phenomenon or have in the main proceeded independently, in terms of both temporal and causal relationships, but some evidence at the microeconomic level—such as analyses of the processes of design and managerial planning in modern corporations—seem to support the hypothesis of co-evolution. In economic organizations in which planning and design are highly purposeful activities, division of labour and co-ordination are the joint result of these very activities. But can the hypothesis of co-evolution of co-ordination and division of labour be extended also to the cases in which—both in markets and in organizations—they are not the outcome of a purposeful planning and design process, but

are emergent and partly unintended properties of the interaction between distributed decision-making activities (cf. von Hayek 1948)?

Behind this question lie two different ways of conceiving the division of labour:

1. As the outcome of a single mind, which designs, divides and coordinates: this is the view behind the Marxian capitalist and the Schumpeterian entrepreneur.
2. As the emergent, unplanned and unconscious outcome of the interaction of local and incomplete decisions.

We believe that an intermediate form can be found in large business organizations, where both forms, top-down conscious design and bottom-up partially unintended adaptation, co-exist. To understand why this is so we have first to take into consideration the role played by routines and contracts; both routines and contracts are necessarily incomplete and partly tacit, and this implies that some discretion must be left to the actors. Thus incompleteness and tacitness are likely to bring about cognitive conflicts.

To clarify this point, suppose that the top management puts into place a redesign of the organization in order to react to some kind of environmental change. The implementation of the new division of labour that is required by such a change gives rise to a complex process of adaptation which is far from that anticipated by the traditional theory of planning. On the one hand, the implementation of a new organizational design requires managers and employees to rethink their jobs and revise their competencies, but on the other hand—to be effective—any new design requires local checks and adjustments, i.e. the resolution of cognitive conflicts arising from a possible mismatch between the general requirements of the project and the specific, idiosyncratic knowledge of any single agent.

These considerations lead us to a brief remark on the role of managers: for what has been said, we cannot clearly distinguish between managing and doing:

the blue collar worker on the most routinized assembly line must repeatedly make decisions about how to handle non-standard situations, and in particular when to call one to the attention of the supervisor. On the other hand, sales managers, in addition to managing their salespersons, often spend considerable amounts of time with clients, engaged in selling, and thus in 'doing'.

(Radner 1992)

“Doing”, i.e. the application of a given routine, is never a mere execution of given rules but always involves some discretionary behaviour, if only because a fully specified routine, which does not leave any space for autonomous decision-making, would require an endless list of contingent behaviours. The essence of managerial activities seems rather to lie in the capability to fill knowledge and information gaps and in the capability to design the organization of labour and “imagine how to do”.

Therefore top-down plans—which are necessarily incomplete for reasons we will discuss below—have to match pieces of knowledge and skills which are largely local and tacit. Inevitably, they therefore have a largely uncertain and unintended outcome.

Before discussing the reasons for the incompleteness of planning, it is useful to note that, when considered in connection with the process of division of labour, the very notion of co-ordination takes on a meaning different from the one implicit in neo-

classical economics. For the latter, coordination means making individual and independent decisions compatible; here the problem of co-ordination concerns the relationship between the top-down activity of designing new organizational set-ups and the adaptive, intelligent, bottom-up reactions by managers and employees, to adapt the organization to the external environment.

Division of labour, distributed knowledge and specialization

The analysis of the process of division of labour as a primary determinant in economic development goes back to Adam Smith and lies at the heart of classical economics. But few of these early analytical efforts have been continued in modern neo-classical economics, and they have been brought back to their central role only within recent non-neo-classical studies of the process of technological evolution (e.g. Dosi et al. 1988).

A possibly fruitful way of interpreting the process of division of labour is to consider it within the general framework of the theory of problem-solving. From this point of view, the division of labour derives from the decomposition of problems into subproblems to be solved independently. Direct observation of the behaviour of organizations and individuals as problem-solving activity—within the tradition initiated by March and Simon (1958)—suggests that such behaviour is a peculiar and unstable balance between two opposite situations: on the one side purely routinized behaviour, in which a series of operations are repeated again and again, and on the other an active and conscious search for solutions to new problems (or new solutions to old problems) faced by the organization.

Without entering the debate on how to represent problem-solving activities formally, some general points are worth mentioning for their relevance to the subject of this chapter.

First, problem-solving activities are characterized by a search in a problem space which generally leads to the decomposition of the original problems into at least partially independent subproblems. If such a decomposition is feasible, subproblems can be solved in parallel and subsequently co-ordinated: the original problem is therefore decomposed into a set of connected subproblems. In the language of economics and organization science, this corresponds to the expectation that the decomposition of a production process (*lato sensu*) into independent subsystems, which can be dealt with in relative isolation and in parallel, will lead to increasing returns, in spite of the co-ordination costs which the decomposition generates.

Second, the working hypothesis we propose—along the lines of Simon and March—is that the search in the problem space, based on the division of the given problem into subproblems, is a model for the division of knowledge; even if division of knowledge and division of labour are not the same process (two employees co-operating in the same organization can have the same competencies and largely overlapping knowledge bases but different jobs), the former is necessary to the latter.

Third, it must be emphasized that the search in the space of sub-problems is a highly uncertain and conjectural process. In fact, when a problem has been decomposed into a set of subproblems, in general not all these subproblems will be solvable immediately and some will have to be decomposed in turn into simpler ones. The decomposition then

recursively proceeds until all the relevant subproblems have been solved. Problem-solving by subproblem decomposition is therefore an uncertain activity, for two reasons:

1. There is no set of decomposition rules which a priori allows agents to achieve a certain result.
2. The solvability of the original problem can only be verified when all the relevant subproblems have been solved (Egidi 1992).

This implies that plans and projects within firms are not merely executed by different actors: they also necessarily require a multi-agent process of learning and adaptation that can generate cognitive conflicts among the different actors.

These features allow us to consider co-ordination within organizations as the complement to the division of knowledge and labour which follows the realization of a new project. The further the division of labour proceeds, the more the different parts require co-ordination and the more information becomes dispersed. A crucial problem then arises: can we assume that the co-evolution of division of labour and co-ordination also takes place within markets?

A first approach would claim that markets are mere co-ordinating devices (by means of price mechanisms), whereas the division of labour is performed entirely within organizations (Coase 1937). A more careful examination would suggest that the two processes take place in a complementary way in both markets and business organizations, the main difference being the degree of awareness which characterizes the actors.

The context of Schumpeterian competition is a good starting point to clarify how division of labour and co-ordination can both take place in markets: suppose that a cluster of product innovations is generated in the economic system. If the new products are substitutes for old intermediate goods, they will activate the modification of the existing productive and organizational routines and a "second generation" of adaptive innovators will come along: they in turn, by using the new products will produce either new goods or the old ones in a new manner. In the first case, we have a dynamic innovation process, in which new projects lead to the modification of other projects. This is not merely a matter of the diffusion of an innovation. On the contrary, what we are describing is an "avalanche" of innovations that are activated recursively by the original cluster of innovations. Two different "responses" could occur as a consequence of the original innovative actions. In the first case, they only give rise to local dynamic processes, which do not alter the basic structure of the division of labour. In the second case, the reaction of the system causes a new definition of the division of labour, knowledge and competencies within the economic system (often going beyond the intentions and expectations of the innovators).

This clarifies how a new division of labour and knowledge can be generated within a market. It should be noted that co-ordination within markets also involves the transmission of knowledge and competencies: prices do not convey sufficient information to support the co-ordination processes that arise from the division of labour. The importance of this point and the provisional conclusions we have reached are reinforced if we consider an intermediate organizational set-up, as does Williamson in his "fundamental transformation" (Williamson 1975). When two or more business enterprises engage in a new common project, an improvement in human and physical assets internal to each enterprise and a transfer of knowledge among the enterprises will

normally be required. The emergence of sunk costs in physical and human capital guarantees the strength and stability of the links among the businesses. Even in this case, the relationship among firms which takes place in the market will involve not only price signals but also (and fundamentally!) the transmission of pieces of knowledge and information. Conflicts and bargaining between the parties will most probably be solved by protest and dispute instead of by leaving the market (Hirschman 1970). Markets are only one of a number of devices for communication and co-ordination between organizations.

To summarize, the transfer of knowledge and competence is a fundamental aspect of co-ordination which takes place not only within organizations but also among organizations in markets. Even if we believe that a fully developed model of the properties of the division of labour and co-ordination and their co-evolution should be at the top of the research agenda, to model this kind of process is far beyond our present abilities and intentions. More modestly, we intend to pick up some of the properties discussed above and try to compare the performances of the different co-ordination forms that arise from the division of knowledge and labour in a context of boundedly rational behaviour. In the next two sections the relevant assumptions will be discussed.

Division of labour and returns to scale

A fundamental question that has been already hinted at is whether and under what conditions subproblem decomposition—which we suggested was a useful representation of the process of division of labour—can generate more efficient forms of the organization of production. This question can be answered in many different ways: one of the first analyses was provided by Simon (1962) with his parable of the two watchmakers, called *Tempus* and *Hora*. Watches are complex objects made up of a very large number of small pieces. *Tempus* assembles his watches sequentially and every time he is disturbed—for instance, by his clients' phone calls—he has to start the assembly all over again. *Hora* instead proceeds by first assembling subunits consisting of only a few components and then putting the sub-units together. When he makes mistakes or is disturbed by external events he only has to restart assembling the current subunit. His average time to complete a watch will therefore be much shorter than *Tempus*'s average.

In this example there is a clear advantage in the division of labour: when a perturbation occurs in a sequential system, it affects the whole system, when instead it occurs in a system made up of small independent and parallel subunits it will not propagate outside the affected subunit. But if we rule out the possibility of perturbations and mistakes, this advantage disappears and the parallel system seems rather less efficient because of its higher co-ordination costs. Only if we consider the positive, long-term effects that the division of labour might have on individuals' working capability, because of learning-by-doing and specialization, can we again find advantages in an increase in the degree of task decomposition. Learning therefore seems to be a key factor in explaining the division of labour. But how do individual learning processes co-ordinate in a framework in which the very definition of tasks is an emergent property?¹ In the rest of this chapter we shall investigate this question by means of a very simple model of division of labour and some simulations of different social co-ordination mechanisms.

Division of labour in the accountants' model

We shall elaborate on some of the points outlined in the previous section by means of a simple model of organization, which will be then made a little richer. The model is the accountants' model suggested by Sobel (1992) which addresses directly the problem raised by Simon's parable of the two watchmakers.

Suppose that the problem to be solved is counting an unknown and possibly very large number of items, for instance, dollar bills, and delivering a correct answer. Each accountant is characterized by a certain productive capacity, i.e. the number of bills he can count successfully in a given time interval, or, alternatively, by a certain probability of making mistakes, which plausibly we can assume to increase with the time he spends counting, because he gets more and more tired and error-prone, and/or more and more bored and likely to have his attention distracted. In the first case the accountant simply cannot work over his productive capacity; in the second he has no given limit, but the likelihood that he will make a mistake will increase with the size of his workload. If an accountant makes a mistake and does not realize it, he will jeopardize the result of the entire counting process; if he is aware of the mistake he will instead be forced to start counting again from the beginning. In all these cases splitting the task into subtasks is necessary to ensure the efficacy and efficiency of the production process.

Of course, the decomposition of the global task into subtasks must be technically possible, which is usually the case only to a certain extent, because there are often technical indivisibilities. The example of counting bills is in this sense very simple because the problem can be decomposed into any number of subtasks, limited only by the fact that the minimum subtask is counting one bill.

Suppose that the global task is to count a stack of N bills of one dollar. Suppose also that each accountant can successfully count at most k bills in a given time interval, at the end of which he issues a new bill whose face-value corresponds to the number of one-dollar bills he has counted (at most, k). Such bills are put in a new, smaller stack and can be counted by other accountants who will issue new, higher-valued bills and so on, until a stack of "fictitious" bills is obtained whose size is no bigger than k and which can be therefore counted by a single accountant who will be able to announce the total number.

Let us illustrate the problem with a trivial example: suppose that $N=10,000$ and all accountants have a counting capacity of $k=10$ bills. We will need therefore $L_1=1000$ accountants at the first level, who will count the 10,000 bills, issue 1000 bills of face value **\$10** and put them in a new stack. To count this new stack, $L_2=100$ accountants are needed, who will issue 100 bills valued **\$100** and put them in a new stack that needs $L_3=10$ accountants to count them. Finally, only one accountant ($L_4=1$) is needed to count these 10 bills of **\$1,000** and announce the result of the entire counting process. All in all, we need 1111 accountants to complete the task in 4 time units, and we obtain a perfect pyramidal subdivision of subtasks.

In general, if $N=K^n$, i.e. the size of the task is an integer power of the individual productive capacity, the optimal hierarchical structure is formed by

$$n = \log N / \log k$$

levels, each consisting of 1, k , k^2 , ..., k^{n-1} accountants.

When the ratio $\log N/\log k$ is not an integer, the number n of levels of the pyramid will be the first integer greater than or equal to this ratio. This implies that the pyramid will necessarily contain some slack resources which cannot fully be exploited.² Let us consider a general task of size N bills. If every accountant can count up to k bills, the lowest level of the organization will produce N/k bills, the second N/k^2 and so on up to level w such that $N/k^w \leq k$. If we have exactly $N/k^w = k$ all resources in the pyramid will be fully exploited, otherwise the productive capacity of the accountant at the top of the hierarchy will not be entirely exploited. Moreover the numbers $N/k, N/k^2, \dots, N/k^w$ might not all be integers: some idle resources will appear at each level for which the ratio is not an integer number.

The ratio between size of the task N and total amount M of labour required for its completion is given by

$$R(k) = N/M = k^n / (1 + k + k^2 + \dots + k^n) = (1 - k)k^n / (1 - k^{n+1}).$$

Thus $R(k)$ tends to 1 as n tends to infinity and the process exhibits asymptotically constant returns to scale.

If the number N of bills to be counted increases, the organization can respond by increasing the level of employment and/or the productive capacity k . The latter can increase also with N constant (equivalent to an increase of productivity, or Marxian exploitation).

Division of labour and co-ordination

If dividing a complex task into subtasks to be handled separately is necessary and/or more efficient than carrying it out as a unit, this process of division of labour implies a growing and non-trivial co-ordination problem. The previous section analyzed some formal properties of a toy model of the division of labour and showed under what conditions a pyramidal structure can implement an optimal subdivision of a given task. In this section we examine how such a structure could emerge in different institutional set-ups and adapt to random environmental disturbances.

Generally speaking, we can imagine at least three different ways of attaining co-ordination:

1. A central planner can use the model outlined in the previous section to design the optimal organization and adapt it to changes in the size of the task and/or changes in the workers' capabilities. To perform this task, he or she needs to know at every moment in time the exact value of every worker's productive capacity and the size of the global task.³
2. A boundedly rational central co-ordinator, instead of attempting an exhaustive plan, can adjust the organizational structure adaptively by moving workers where needed. Such a co-ordinator can operate as a sort of Walrasian auctioneer: he or she receives messages about all the flows between the different levels of the organization and about unused resources, and moves workers between adjacent levels of the hierarchy in such a way as to fill the gaps between demand and supply.

The information requirements of this boundedly rational co-ordinator are quite different from those of a fully rational central planner: while the latter needs precise information about the size of the overall task and the characteristics of each accountant, the former needs information on all the flows between different hierarchical levels. Whereas the central planner needs precise information, as even small amounts of noise will make the entire organization ineffective, the boundedly rational co-ordinator will generally need some signals (even if they are only qualitative ones) about supply-demand disequilibria. Moreover this boundedly rational co-ordinator can be replaced completely by inter-level co-ordinators who each take care only of the relationship between demand and supply at one interface between levels, regardless of what happens in other parts of the organization.

As to the kind of cognitive capabilities that are required by the central planner and the boundedly rational co-ordinator, the former has to develop a highly abstract and general decision rule. This requires that the problem has been understood in its general features and decomposed into subproblems. The boundedly rational co-ordinator, on the other hand, can implement the organizational structure adaptively, using a process of trial and error which can proceed without a general understanding of the problem (Dosi & Egidi 1991, Dosi et al. 1993).

But the process of adaptive co-ordination involves a cost, given by the loss of efficiency incurred during the process of adaptation. While the perfectly rational central planner computes the optimal organization in his mind, the boundedly rational co-ordinator carries out the design process in real time and corrects mistakes only after experiencing the loss of efficiency they cause. On the other hand, if the central planner makes mistakes, these are likely to damage the organization more persistently, because the planner is unable to process signals which detect the presence of inefficiencies and make consequential adjustments, and a thorough re-design is always needed even to cope with small perturbations.

3. Co-ordination can also be achieved with a completely decentralized mechanism, a quasi-market in which each accountant adjusts his production and/or his position according to quantity and/or price signals, which are processed independently by each individual. Each interface between levels of the organization constitutes a market where the accountants of the lower level sell the “fictitious” bills they produce and the accountants of the higher level buy them. Demand and supply in each of these markets depend on the number and productive capacity of the accountants at the two levels. Suppose, for instance, that the overall productive capacity at the lower level is not sufficient to supply the accountants at the higher level with enough bills. Some of the higher-level accountants will not be able to produce enough because of insufficient input and will therefore tend to move to other parts of the organization. These adjustments could take place through simple quantity signals or through more complex price mechanisms of Marshallian type (for instance, excess demand generates a price increase which raises the profits of sellers; and this attracts new sellers who make the supply increase and balance the initial excess demand).

A decentralized co-ordination mechanism of this kind requires only local information processing: each accountant can process disequilibrium signals autonomously according to his own local knowledge about his own characteristics and without any need to know

about other accountants' characteristics. However, such a mechanism requires a powerful system of incentives in order to fuel the agents' search for maximum profits. Finally, as in the case of a central co-ordinator, decentralized co-ordination takes place through costly adjustments in real time.

To summarize, in the first case, all the decision-making capability is concentrated in the central planner, who has developed a highly context-independent co-ordinating rule ($n = \log N / \log k$). The accountants do not even have the possibility of sending signals about inter-level disequilibria to the central planner. In the second case, the central co-ordinator follows a behavioural routine based on disequilibrium signals sent by the accountants. In the third case, there is no central co-ordination.

In the next section of the chapter we will study how co-ordination could emerge as an adaptive property in the different environments defined by centralized and decentralized co-ordination mechanisms and explore some of their properties by means of a simulation model.

Division of labour and the emergence of co-ordination: a simulation model

In this section some of the properties of the co-ordination modes outlined in the previous section will be investigated further by means of a simulation model of adaptive learning based on genetic algorithms and classifier systems (Holland 1975 and 1986, Goldberg 1989).

Let us consider a population of h accountants. Each of them is characterized by two variables: his position in the hierarchy (i.e. the integer number which identifies the hierarchical level where the agent is placed) and his productive capacity (i.e. the maximum number of banknotes he can count in a given time interval). Both variables can be encoded by their binary representations. In our examples we have used a string of six bits to represent each accountant: the first three bits for his position in the hierarchy (levels 0 to 7) and three bits for his counting capacity (from 0 to 7 banknotes per unit of time). The i th accountant is represented by the string:

$$A_i: p_1 p_2 p_3 k_1 k_2 k_3 \quad p, k \in \{0, 1\}$$

A set of h such strings provides, at each moment in time, a complete representation of the organization in which the position in the hierarchy (level) and the task of each individual (productive capacity) are specified.

Accountants can move (or be moved) across hierarchical levels and can modify their productive capacity according to signals of input-output disequilibrium either at the individual (in the case of decentralized co-ordination) or at the inter-level interface (in the case of a boundedly rational central co-ordinator). Such signals, appropriately interpreted, constitute the condition part of the adaptive learning system.

The case of the boundedly rational co-ordinator

The boundedly rational central co-ordinator is able to adapt the entire organization according to signals of disequilibrium/equilibrium at the interface between different

levels of the hierarchy. It can thus be represented by a set of condition-action rules where the condition classifies such signals and the action defines a complete organizational structure. In detail, we have:

- (a) Environmental messages (equilibrium/disequilibrium signals) are given by the concatenation of eight binary strings (one for each interlevel interface, including the interface between the last level and the “final demand”). Each one of these eight substrings is composed of two digits:

$$s_1 s_2 \in \{0, 1\}$$

where the first digit is set to 1 when there is an excess supply at that interface and is set to 0 otherwise; the second digit is set to 1 when there is an excess demand at that interface and is set to 0 otherwise.

- (b) Conditions are strings of the same length (16 bits) as the environmental messages which they classify:

$$c_{11} c_{12} c_{21} c_{22} \dots c_{81} c_{82} \in \{0, 1, \#\}$$

Each bit position may be set to 0, 1, or a “wild card” marker, #.

- (c) Action parts are binary strings of length $6h$ that define the whole organizational structure:

$$p_{11} p_{12} p_{13} k_{11} k_{12} k_{13} \dots p_{h1} p_{h2} p_{h3} k_{h1} k_{h2} k_{h3} \in \{0, 1\}$$

In this way an adaptive and boundedly rational central co-ordination can be represented by a classifier system. Here we briefly review its basic features (more details can be found in, for instance, Holland 1986, Holland et al. 1986, Goldberg 1989).

An adaptive learning system is a system of condition-action rules such as those so far described. In addition, each rule is attributed a strength coefficient which, as a first approximation, measures how successful that rule has been in the past, and a specificity coefficient (the number of bits in the condition part that are not set to the wild card #), which measures the strictness of the condition. The smaller the cardinality of the set of environmental messages that satisfy that condition, the higher is its specificity, the highest specificity belonging to rules whose conditions are satisfied by only one environmental message.

This set of rules is processed according to the following cycle throughout the simulation process:

1. Condition matching: a message is received from the environment which informs the system about the disequilibrium/equilibrium conditions at the inter-level interface. This message is compared with the conditions of all the rules, and those rules which match, i.e. those which apply to such a state of the world, enter the following step.
2. Competition among matched rules: all the rules whose condition is satisfied compete in order to select the one which is allowed to execute its action, i.e. to implement the organizational structure specified by its action part. To enter this competition each rule

makes a bid based on its strength and on its specificity. In other words, the bid of each matched rule is proportional to its past usefulness (strength) and its relevance to the present situation (specificity):

$$\text{Bid}(R_j, t) = k_1(k_2 + k_3 \text{Specificity}(R_j)) \text{Strength}(R_j, t)$$

where k_1 , k_2 and k_3 are constant coefficients. The winning rule is chosen randomly, with probability proportional to the bids.

3. Action and strength updating: the winning rule executes the action indicated by its action part. It has its own strength reduced by the amount of the bid and increased by the payoff that the action receives, given the occurrence of the “real” state of the world. If the j th rule is the winner of the competition, we have:

$$\text{Strength}(R_j, t+1) = \text{Strength}(R_j, t) + \text{Payoff}(t) - \text{Bid}(R_j, t)$$

4. Generation of new rules: the system must be able not only to select the most successful rules, but also to discover new ones. This is ensured by applying “genetic operators” which, by recombining and mutating elements of the most successful existing rules, introduce new ones which could improve the performance of the system. In this way new rules are injected constantly into the system and scope for new opportunities is always made available. The genetic operators used in our simulations (as in standard classifier systems) are crossover and mutation.

The case of decentralized co-ordination

In this case each accountant can move autonomously across the hierarchical levels of the organization and change productive capacity according to individual disequilibrium signals. Some kind of market mechanism is thus necessary to generate local signals to reflect global disequilibria. One of the simplest ways in which such a quasi-market could operate⁴ is by delivering quantity signals through rationing: if at an interface between levels there is an excess demand, some (possibly all) of the accountants of the higher level will not be able to get all the banknotes they would like and will be subject to rationing on the demand side. If instead there is excess supply, some (possibly all) of the accountants of the lower level will not be able to sell all the banknotes they would like and will be subject to rationing on the supply side.

Thus, in this case, each accountant is an independent decision-maker who can be modelled by an autonomous classifier system, and the links between the different classifier systems are given by rationing signals. Each accountant is represented by a set of condition-action rules, where the condition classifies such signals and the action defines his own position/capacity pair. In more detail we have:

- (a) environmental messages (equilibrium/disequilibrium signals) are in general different for each accountant and given by a two-digit binary string:

$$s_1 s_2 \in \{0, 1\}$$

where the first digit is set to 1 when the accountant has been rationed on the demand side (i.e. there is an excess demand at the interface between that accountant's level and the lower one) and is set to 0 otherwise; the second digit is set to 1 when the accountant has been rationed on the supply side (i.e. there is an excess supply at the interface between that accountant's level and the higher one) and is set to 0 otherwise.

- (b) Conditions are therefore strings of two bits which classify such environmental messages:

$$c_1 c_2 \in \{0, 1, \#\}$$

- (c) Action parts are binary strings, each six bits in length which define the accountant's position in the hierarchy and his productive capacity:

$$p_1 p_2 p_3 k_1 k_2 k_3 \in \{0, 1\}$$

Each classifier is then processed by exactly the same execution cycle as that already described for the centralized co-ordination case.

The two co-ordination systems have been tested on a simple problem. The task is to count 25 bills and there are 6 accountants, whose capacity can vary between 0 and 7 banknotes and whose position in the hierarchy can vary from level 0 to level 7. Level 0 is a stand-by position: accountants in this position do not enter the production process⁵. Let us examine in more detail the nature of the environmental signals which characterize the two institutional set-ups:

1. In the case of the boundedly rational central co-ordinator, the winning rule's action part is decoded in order to obtain the corresponding organizational design. The productive capacity of all the accountants who are at hierarchical level 1 is then summed to obtain the total demand for banknotes at this level. If the total demand is less than 25, the environmental message will signal an excess supply in the following iteration (the first digit will be set to 1 and the second to 0). If the total demand is greater than 25, the environmental message will signal an excess demand (the first digit will be set to 0 and the second to 1). Only if the total demand is equal to 25 will the environmental message signal an equilibrium situation (both digits set to 0). The total supply of level 1 can now be computed in this way:
 - (i) if the total demand is less than or equal to 25, all accountants at level 1 will be able to exploit their productive capacity to the full. Thus the total supply will be given by a number of banknotes equal to the number of accountants, and each note will have a face-value equal to the productive capacity of the accountant who produced it; and
 - (ii) if instead the total demand is greater than 25, some accountants (randomly chosen) will be unable to use their own productive capacity fully. Total supply will be given by a number of banknotes equal to the number of accountants at the first level who received at least one banknote and their face-values will be given by the production of the accountants who produced them.

Once the total supply at the interface between levels 1 and 2 has been so computed, we can determine the total demand as the sum of the productive capacities of all the accountants who are placed at level 2 of the hierarchy. We can then set the third and fourth digits of the environmental message according to the disequilibrium and equilibrium situations which are thus realized. The same procedure can be repeated for all the organizational levels. At the last level we will suppose the existence of a “final” demand of one banknote of face value 25. If more banknotes are offered, an excess supply will be signalled and the face-value of the only purchased banknote will determine the overall payoff to the organization: if its value is 25, the organization (i.e. the winning rule) will receive a positive payoff, otherwise it will receive a negative payoff proportional to the absolute difference between 25 and the face-value itself.

2. In the case of decentralized co-ordination, inter-level supplies and demands are computed in exactly the same way, but environmental messages are determined separately for each accountant, by means of random rationing. For instance, if at a given interface between levels, demand is higher than supply, banknotes are assigned, one by one, to a randomly chosen accountant who still has some unused productive capacity. All accountants on the supply side will therefore have sold all the banknotes they produced, and at the following iteration will receive a message whose second digit is set to 0. As to the accountants on the demand side, some (possibly none) of them will have been given all the banknotes they required and at the following iteration will receive a message whose first digit is set to 0. The others (possibly all of them) will find themselves with at least a part of their demand unmet and at the following iteration will receive a message whose second digit is set to 1. The organizational payoff is in this case distributed to all the accountants’ winning rules through a bucket-brigade mechanism (Holland 1986).

Simulations have been carried out in order to test the adaptive performance of the systems in different environmental conditions. Let us examine briefly the main results:

1. The first set of results concerns the simplest situation: an error-free and stationary world. The stack of banknotes to be counted always contains 25 notes, and no mistakes can occur in the counting process, i.e. the counting can always be performed without any disturbance and the face-value of the output always equals the amount of used productive capacity. In this situation the boundedly rational centralized co-ordination mechanism is considerably more efficient in finding the optimal structure as far as both the speed of convergence to an effective organizational design and its stability are concerned.

The structure of the payoff function was found to play a crucial role in determining the speed of convergence and the type of organization which emerges. In particular, punishments for unused productive capacities are essential for reducing slack resources, and a payoff function that takes into account the total task-completion time is necessary to obtain an equal distribution of capacities across agents.

2. In error-prone environments, accountants can make mistakes. Two different kinds of mistake are possible. In the first case accountants are aware of the mistakes they make and correct them by restarting the counting process. The payoff function contains a

penalty for the time (steps of the counting process) taken to perform the overall task. A second type of mistake is represented by a random variable (with mean 0) which might cause a deviation between the amount of productive capacity used by an accountant and the face-value of the banknote he produces. In this case accountants are not aware of the mistakes they make and they deliver a result which is not correct. The decentralized mechanism is more efficient in dealing with both types of mistake. Local adjustments therefore seem to increase efficiency in coping with local disturbances.

Conclusions and directions for further research

We have pointed out that the division of labour determines a hierarchical structure of tasks of pyramidal form. Does this structure directly relate to power relations within organizations, i.e. is the hierarchical system of task division connected to the hierarchical system of power and control? The answer is, generally speaking, no. The pyramidal structure of our example emerges only from functional relations (each agent has to use a multiplicity of inputs coming from different agents) and is independent of control and power relations. But, on the other hand, the shape of the pyramid can have some important consequences for the relationship between agents. The number of agents in fact decreases as we climb the hierarchy and thus, if the relationships between adjacent levels are managed by a market, such a market cannot be a competitive one because of its strong asymmetry. Market power will be higher for agents placed in higher hierarchical positions.

In our simple model we have supposed that movement of workers across hierarchical levels can take place instantaneously and at no cost: this is clearly a most unrealistic hypothesis. Different positions in the hierarchy require different capabilities and a knowledge of different routines. There is likely to be a trade-off in the organization between the time a worker takes to learn the routines associated with a different hierarchical level, and the increase of production time and number of mistakes resulting from an increase in individual workloads.

Moreover, in real organizations the types of knowledge required normally differ between hierarchical levels: higher levels in the hierarchy demand a broader but less precise kind of knowledge (regardless of the intellectual capabilities of workers). The division of labour usually involves a strong cognitive asymmetry between different levels of the task as a part of the general decomposition of knowledge, but this element does not appear in our accountants' model.

Finally, real economic organizations always possess some mechanism of control to spot elements which do not work properly, either because of real problems or because of opportunistic behaviour. They also include a parallel and connected system of incentives that aims to avoid the appearance of such problems. All these elements should be embodied in our model (which could become a sort of "beehive" rather than an accountants' model) to make it more realistic.

Notes

1. The problem of the co-ordination of learning processes within an economic organization has also been studied, in a framework in which the division of labour is given, in Marengo (1992).
2. The presence of such idle resources is a possible reason for the existence of economies of scale in this technology: a multiplicity of tasks can be internalized in a single organization, which can therefore reduce the amount of idle resources. In the above example, two separate tasks of \$5,000 would be handled more efficiently by a single organization rather than separately. But this issue lies outside the scope of this chapter.
3. This requirement seems rather paradoxical. If the planner knew exactly the number of bills to be counted there would be no need to carry out the counting process. This paradox also arises in less caricatured production processes, although less strikingly. Fully centralized and exhaustive planning would require perfect knowledge of every aspect of the production process: only in this case could the co-ordination problem be solved in the planner's mind without the need for local adjustments.
4. Markets with prices are much more complex and to model them correctly we should introduce a series of hypotheses on individual utility functions, which would take us far from the core issue of the chapter.
5. The possibility of leaving the organization must be allowed for if we want the system to be able to adjust when there are more accountants than the optimum number.

Chapter 4

Emergence of kinship structures: a multi-agent approach

Jean Pierre Treuil

Among the various systems of representation that form a culture, those concerning kinship and spatial relationships are obviously highly significant. These have been prime elements in human exchange and they are closely related in many societies. This chapter takes as its starting point the kinship systems of the Australian aboriginal societies studied by B.Glowczewski, an anthropologist. It reports the first part of more extensive research conducted jointly with him. The research aims at studying, by means of computer simulations, the conditions governing the observed structuring of territory in these societies.

Some methodological positions: structure and genesis

Anthropology studies what is shared by the members of a society that allows them to distinguish between “us” and “the others”. Such a definition inevitably raises questions about the genesis of this sharing. Is this genesis “outside the field of investigation”, the scientific effort having to focus on the structures and functions of the shared elements themselves? Or is their genesis the main issue whose answer is at the basis of any explanation?

Many approaches, such as that defined by Sperber (1985, 1987), which reduced anthropology to an “epidemiology of representations”, see the sharing of cultural elements as being the result of dynamic processes. These circulate numerous types of mental representations (internal) and public representations (external) within society until some prevail and become the common property of the majority. According to this view, anthropology aims at studying empirically the systems of circulation, the transformations under way and the reasons for the successes of some representations and the disappearance of others. That the methods of transmission are themselves cultural elements whose sharing must also be accounted for in the same way adds to the complexity of the mechanisms without changing their nature. The influence of modes of transmission through the concept of “cultural inheritance” is at the heart of work such as that by Boyd (Boyd & Richerson 1985).

The genesis of common kinship structures

Following this general approach, we have tried to develop a simple process leading to the sharing, by all the individuals of a society, of the same system of rules governing kinship

and marriage, starting from an original society composed of individuals with differing rules. In order to formulate such a process, the interactions between systems of rules must be capable of description. Furthermore, the systems of rules must belong to the same general pattern, each representing a specific development of it.

The process described here leads to systems of kinship rules which share a significant property with the elementary kinship structures described in the literature, that of being mathematically formalizable in terms of “permutation groups”. Our kinship structures also share another important property with the real ones: the property of exogamy. However, the real systems also exhibit properties of symmetry (see the example of Kariera society below); these properties ensure that all the matrimonial classes identified in a given kinship structure have an equivalent position and that no one has a privileged status in the system. Such properties are more difficult to generate in our simulations. Nevertheless, we think that the simple hypothesis at the basis of convergence, namely an evolution which searches permanently for a consensus between the individuals who make similar choices, can provide some interesting explanations of these cultural phenomena.

Mathematical modelling of elementary kinship structures

The work by Lévy Strauss on elementary kinship structures led to numerous studies aimed at reducing the observed matrimonial rules to a small number of axioms (Boudon 1968). One of the axiomatizations, that of Courrège (1965), identifies these structures as groups of permutations.

Courrège considers “classificatory” kinship systems. In these systems, society is divided into n matrimonial classes. The rules of marriage and filiation (that is, the relationship between parents and offspring) which are forced upon an individual are expressed in relation to the class to which he or she belongs.

The marriage rules determine the class to which the spouse of a given person must belong. Courrège’s matrimonial function, designated by ω , determines the class of the woman married to a man from a given class. The inverse function ω^{-1} determines the class of the husband. The filiation rules determine the class to which the children of a person will belong, given his or her own class and sex. The fathers of two people from the same class cannot belong to different classes. The rules are formalized in terms of two functions: the maternal and the paternal function. The maternal function, designated by μ , determines the class of the children of a woman. Similarly, the paternal function, designated by π , determines the class of the children of a man.

Coherence implies that women and men who are allowed to get married agree about their children’s class. This coherence is expressed by the relationship $\pi = \mu \omega$. The “biunivocity” of the functions ω , μ , π transforms them into permutations of the set of classes. The above relation means that the total set of functions obtained by the combination of the functions ω , μ , π is a group of permutations.

This axiomatization makes it possible to analyze the interpretation of the formal properties of the permutation groups in terms of kinship and to situate the classificatory kinship systems observed in all possible groups.

An example: the Kariera structure

According to Courrège, the Kariera structure includes four classes designated by 1, 2, 3, 4; ω , μ and π are given by:

ω	μ	π
1→3	1→2	1→4
2→4	2→1	2→3
3→1	3→4	3→2
4→2	4→3	4→1

We can represent the functions ω , μ , π as matrices:

$$\omega = \begin{vmatrix} 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \end{vmatrix} \quad \mu = \begin{vmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \end{vmatrix} \quad \pi = \begin{vmatrix} 0 & 0 & 0 & 1 \\ 0 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{vmatrix}$$

Modelling the dynamics of the generation of kinship structures

According to the ideas presented in the previous section, we assume that:

1. A society governed by the kinship systems described above is composed of individuals who get the same model, R , consisting of the rules of marriages and filiation. This model R determines, on the basis of an individual's class, the classes prescribed for his or her spouse and children.
2. This homogeneity results from a convergence starting from a heterogeneous society where an individual i has model R_i , which can differ from those of others.
3. The convergence occurs after modifications are made to the models in the course of interactions between individuals.

Specification of the processes of the generation of kinship structures

To specify the processes generating kinship structures, it is necessary to determine:

- (a) the general pattern of evolution;
- (b) the range of variation of the models R_i ; and
- (c) the rules of interactions and transformations of the models.

The general pattern of evolution

We assume that the convergence occurs during a reproduction process which allow the replacement of the population. Within this framework, we propose a continuous pattern which gives a step-by-step replacement.

Let I_t be the population at time t . The evolution of the system can be described by a sequence $[I_0], \dots, [I_t], [I_{t+1}], \dots$ where the transition from $[I_t]$ to $[I_{t+1}]$ is through the birth or the death of a single individual. At each time step, an individual is selected randomly from the population. If, at that moment, the size of the population is higher than expected, this individual disappears. Otherwise, the individual searches for a partner from the opposite sex and if he or she succeeds, the female gives birth to a child who takes her place in the population.

Range of variation of the individual models

The models of classificatory kinship systems are defined by a triplet of permutations:

$$R=(\omega, \mu, \pi)$$

These permutations are linked by the relationship

$$\pi=\mu\omega$$

To specify the model as a whole, it is therefore sufficient to determine two permutations. Here, for simple reasons of symmetry, the maternal and paternal permutations μ and π are selected. We write

$$R=(\mu, \pi)$$

There are many ways of conceiving the space \mathfrak{R} within which models R are specific points. The one chosen in this study considers the space \mathfrak{R} as composed of couples $R=(M, P)$. In these couples, M and P are no longer functions, but relations; they do not determine univocally the children's class from that of the parents', but only delimit a set of possible classes.

Mathematically, a model R is a set of elements (s, x, z) , where s is the sex (1 for women and 0 for men), and x and z are classes. Each element means that z is a possible filiation of an individual of sex s belonging to class x . M and P are the subsets of elements with $s=1$ and $s=0$, respectively.

$$M=\{\dots,(1, x, z)\dots\} \quad P=\{\dots,(0, x, z)\dots\}$$

This is expressed naturally in a matrix representation: M and P are matrices that are always composed of zeroes and ones, but they do not necessarily have to have only one 1 in every row or column. For instance, in the case of the four matrimonial classes {1, 2, 3, 4}, we could get the matrix shown overleaf.

The first column of P defines two possible connections for the children of men of class 1, class 2 and class 4. For a society with N classes, the space \mathfrak{R} of the possible models will include $2 \times N \times N$ elements.

$$M = \begin{vmatrix} 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 \end{vmatrix} \quad P = \begin{vmatrix} 0 & 0 & 0 & 1 \\ 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 0 \\ 1 & 0 & 0 & 0 \end{vmatrix}$$

The rules of interaction and transformation of the models

The rules of interaction determine the conditions and protocol of breeding and their results. In choosing a set of rules, our principle was the following: we consider that individuals make their decisions during meetings, and that they use only knowledge contained in their models and nothing else. For instance, we accept that two individuals might make a decision on the basis of a comparison between their models, but we disallow the possibility that an individual might make a decision by comparing his or her model with the average value computed over the whole population. All the decisions are “model based”.

Under this restriction, the rules are completely specified by four operators:

1. A *condition of union* Ω selects from all the couples those that are allowed to reproduce and the classes of their children.
2. A *meeting protocol* U selects from all the couples allowed by Ω those which actually meet for breeding. It also chooses the class of the child resulting from the meeting, in the range of possible classes determined by Ω .
3. A *combination operator* C determines the model R of the child based on the models of its parents after each reproduction.
4. A *mutation operator* T introduces some modifications in the model of an individual, at birth or during the life cycle.

Notations

To explain these operators, we must introduce some notation, following Courrège:

i, j, k, \dots designate individuals;

n represents the total number of individuals in the population;

N is the number of classes;

$s \in \{0, 1\}$ designates a sex value and $x, y, z, \dots \in \{0, 1, \dots, N\}$ designate classes;

X, Y, Z, \dots represent subsets of classes;

$R_i = \{M_i P_j\}$ represents the kinship model of individual i ;

$R_i[s, x]$ represents the possible classes of the children of an individual of sex s and class x : $R_i[1, x] = M_i[x]$; $R_i[0, x] = P_i[x]$;

s_i represents the sex of individual i , and x_i his or her class;

$n(s, x, z)$ represents the number of individuals who have (s, x, z) in their model; and

R_m designates the average model computed from the models R_i over the population.

R_m is the set of elements $[(s, x, z), n(s, x, z)]$. From these numbers we can compute the proximity of the model R_m to a permutation group. compute a magnitude S called average specificity, which measures the

Conditions of union Ω

The conditions of union, Ω , are based on a principle of agreement between the individuals selected from the potential couples about the classes of their children. This is a generalization of the relationship $\pi=\mu\omega$

The simplest principle only requires that the maternal part of the model of the mother and the paternal part of the father should be consistent, at least with respect to the class of their children. Let i and j be the potential parents. Ω can be written

$$R_i[s_i, x_i] \cap R_j[s_j, x_j] \neq 0$$

Then $Z_k = R_i[s_i, x_i] \cap R_j[s_j, x_j]$ will designate the possible children of the couple.

A more limiting variant of this condition requires first the separate agreement of the father's and the mother's models. Thus, under this condition, a man of class x_i should marry a woman of class x_j . This condition can be written:

$$R_i[s_i, x_i] \cap R_l[s_j, x_j] \cap R_j[s_i, x_i] \cap R_j[s_j, x_j] \neq 0$$

Meeting protocol U

We experimented with a protocol which runs as follows:

1. Individuals first choose a filiation for their future child. This choice is made from all the filiations allowed by their own model.
2. Couples are selected to breed at random from among those who made similar choices.

The procedure we used, called the "veto" procedure, requires that the choice made by an individual i from a filiation z must not be prohibited by another individual of the same sex who has already chosen this filiation z . In more detail, the process is as follows:

An individual i first chooses a filiation z at random from his or her own set $Z_i=R_i[s_i, x_i]$; then i asks individuals of the same sex who have already chosen the filiation z if they agree with this choice; if one of these individuals j does not have the rule (s_i, x_i, z) in his or her own model R_j , the individual i will search for another filiation in Z_i . Moreover, in what we call below "Lamarckian processes", it will remove the rule (s_i, x_i, z) from its model R_i , so that it will never again choose the class z .

This kind of procedure is sophisticated but it corresponds to a situation which is not unknown in human behaviour. Its application has an important consequence: it reduces the diversity (in terms of classes) of the set of individuals choosing a given filiation. This set becomes increasingly specific with time.

Combination operator C

The principle consists in establishing the maternal and paternal relations of the child, based on the parents' maternal and paternal relations, respectively.

Let i be the father and j be the mother of the child, k :

$$R_k=C[R_i, R_j] \quad P_k=C[M_i, M_j] \quad \text{and} \quad M_k=C[P_i, P_j]$$

The combination operator is based on genetic combination following techniques taken from genetic algorithms (Goldberg 1989). Each relation M or P is coded in the form of a sequence (a “genotype”) of $N \times N$ “genes” each assuming the value of 1 or 0 in accordance with the matrix representations of these relations. As a safety measure, the sequencing of the matrices on the genotype is randomly determined. Because of the length of each genotype ($N * N$ positions) the operator uses a crossover with two points.

The position in the genotype of the gene showing the presence (value 1) or the absence (value 0) of a given rule (s, x, z) is the same for all the individuals, but it is determined randomly to avoid biasing the groupings. In these conditions, the use of the combination operator amounts to building the set R_k of the child’s rules as follows:

- If (s, x, z) belongs both to R_i and R_j then it belongs to R_k .
- If (s, x, z) belongs neither to R_i nor to R_j it does not belong to R_k .
- If (s, x, z) belongs to one of them, it will belong to R_k or not according to random selection.

Mutation operator T

We experimented with two kinds of mutation operator. Both work on the “genotype representations” of the individuals’ kinship models R_i , changing the value of one gene.

1. The genetic mutation operator. In this case the gene and its new value are chosen randomly. The mutation operator runs at the birth of the individual, after the combination of the parents’ models. The genetic mutation amounts to randomly removing or adding a rule (s, x, z) in the set R_i .
2. The correcting mutation operator. This operator resets to 0 the value of a gene during the life cycle of an individual. The correcting mutation operator amounts to removing from R_i a rule (s, x, z) for instance, a rule which does not lead to a choice accepted by the other members of the community.

Analyzing the simulations

Results are shown in two windows built automatically during the simulations (Figures 4.1 & 4.2) that show a set of graphics.

Synchronic analysis

The first graphic is a table about the unions and filiations observed at any time that is used to analyze the structures of relationships between classes and their approximation to a group of permutations. The second graphic is also shown in Figure 4.1. It is used to give the mean configuration $R_m = [M_m, P_m]$ of individuals’ models $R = [M, P]$. In terms of the notations listed above, these graphics are histograms of the distribution of the numbers $n(s, x, z)$.

The third graphic shown in Figure 4.1 is used to evaluate the cycles of classes through successive generations, for women and men. A cycle is the set of the class of an individual and the classes to which that individual’s ancestors of the same sex have

belonged. The length of the cycle is the size of this set. These cycles are characteristic of the average structure of kinship models over the population.

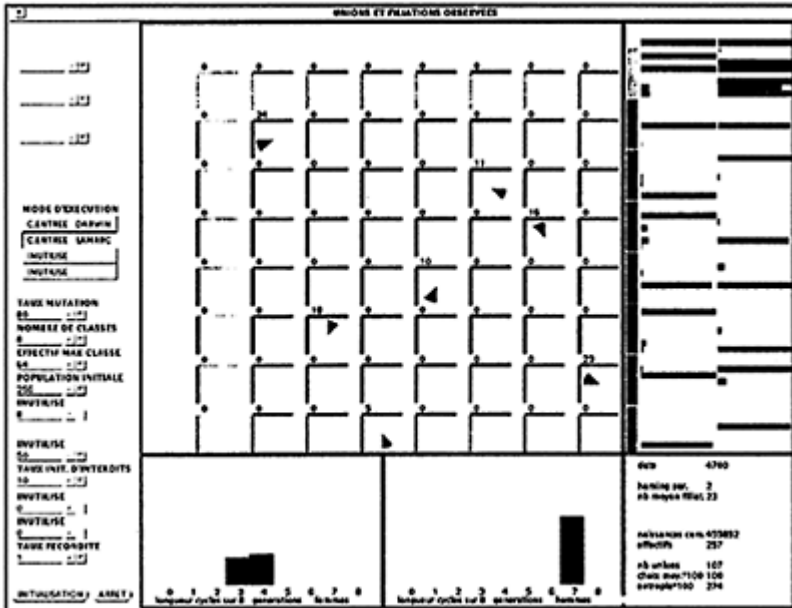


Figure 4.1 Analyzing the simulation: synchronic analysis.

Diachronic analysis

All these graphics are “snapshots” of the situation at a given time. Four further graphics analyze processes over time (see Figure 4.2). They show, from top to bottom:

1. The evolution $F(t)$ of the mean number of rules (s, x, z) in individuals’ models, i.e. the average number of possible filiations. $F(t)$ corresponds to the average number of 1s in the matrices M and P .
2. The evolution $C(t)$ of the mean number of possible classes of the children of actual unions. This graph is used to evaluate the dispersion or concentration of the population in the classes, computed every hundred time periods.
3. The evolution $E(t)$ of the entropy of the distribution of numbers per class. This is at a maximum when the distribution is regular and each class includes roughly the same number of individuals. It is equal to zero when the whole population is concentrated into a single class.
4. The evolution $S(t)$ of the average “specificity” of the classes. This is

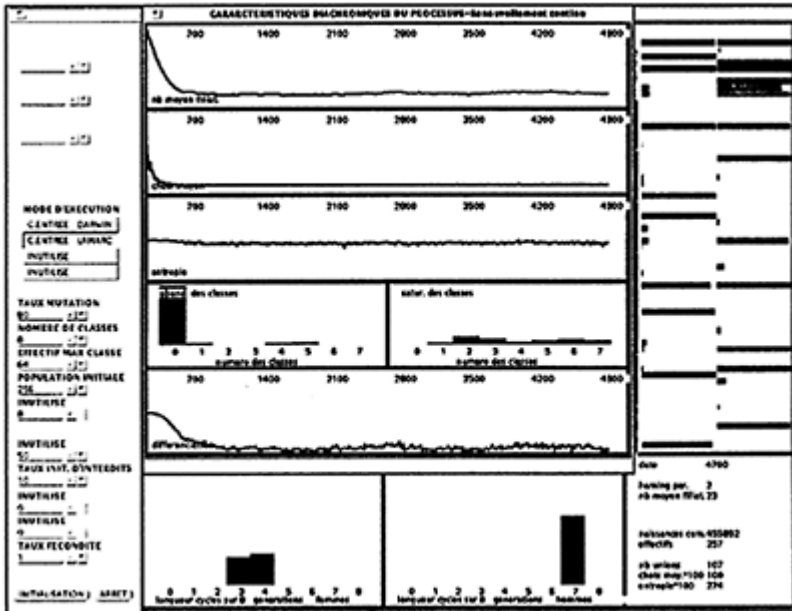


Figure 4.2 Analyzing the simulations: Diachronic analysis.

based on the conditional entropy, computed on the set of classes that have the same filiation, in the average configuration of the individuals' models. S is zero when specificity is at its maximum and this occurs when a given filiation has only one class for the parent of each sex.

The time unit t used in these graphics consists of n of the elementary steps described earlier. During one time unit, all the n individuals of the population get on average one opportunity to be randomly selected to reproduce or die. Therefore, the time unit amounts to roughly half a generation.

Results

All the experiments analyzed here are carried out with eight classes and a population of 256 individuals. The choice of eight classes is consistent with many observed structures of kinship. We have also carried out experiments with 16 classes and more numerous populations, but the results are not essentially different in their qualitative features. We discuss later the influence of the number of individuals.

Convergence

We observed the evolution of the above measures for the same periods for each kind of simulation. We analyzed the results from the values of:

1. The mean number of possible filiations in the individuals' models, $F(t)$. The maximum value of F is $2 \times 8 \times 8 = 128$ filiations, and the minimum is 0.
2. The mean number of possible classes of the children of the actual unions is $C(t)$. The maximum value of C is 8. The minimum is 1, otherwise a union would be not realized.
3. The entropy of the distribution of the population on the set of classes $E(t)$. Its range is between 0 and 3.
4. The average specificity of classes $S(t)$. Its range is also between 0 and 3.

Within any one kind of simulation, we observed that $F(t)$, $C(t)$, $E(t)$ and $S(t)$ always have similar shapes. These variations are related to discontinuities within the evolution of classes and the structure of R . The curves and observation of the process suggest that it is useful to break the simulation into three phases (see Figure 4.3).

1. During the initial phase, F and C decrease evenly at a speed linked to the mutation rate. They are related by a formula which comes from the hypergeometric law governing the statistical distribution of C when F is given: $C = F^2 / 4N^3$. During this phase, there are meetings between all classes; there are many non-zero values within the union matrix. So the sizes of all the classes change very slowly. Entropy remains close to its maximum value, 3, for all eight classes. This phase ends when C reaches a value close to its minimum, i.e. close to 1. At this point the individuals who meet to reproduce usually get only a single common possible filiation: the one they choose.
2. During the second phase, corrections and stabilization can be observed: C remains constant, slightly over its minimum value; F and S continue to decrease until they reach minimum values which are related to the mutation rate; E remains close to its maximum or falls into a lower value. Many null values appear within the union matrix. Individual kinship models become similar, but their common structure does not correspond to a permutation group because different classes may have the same filiation. The union matrix is not stable. There are

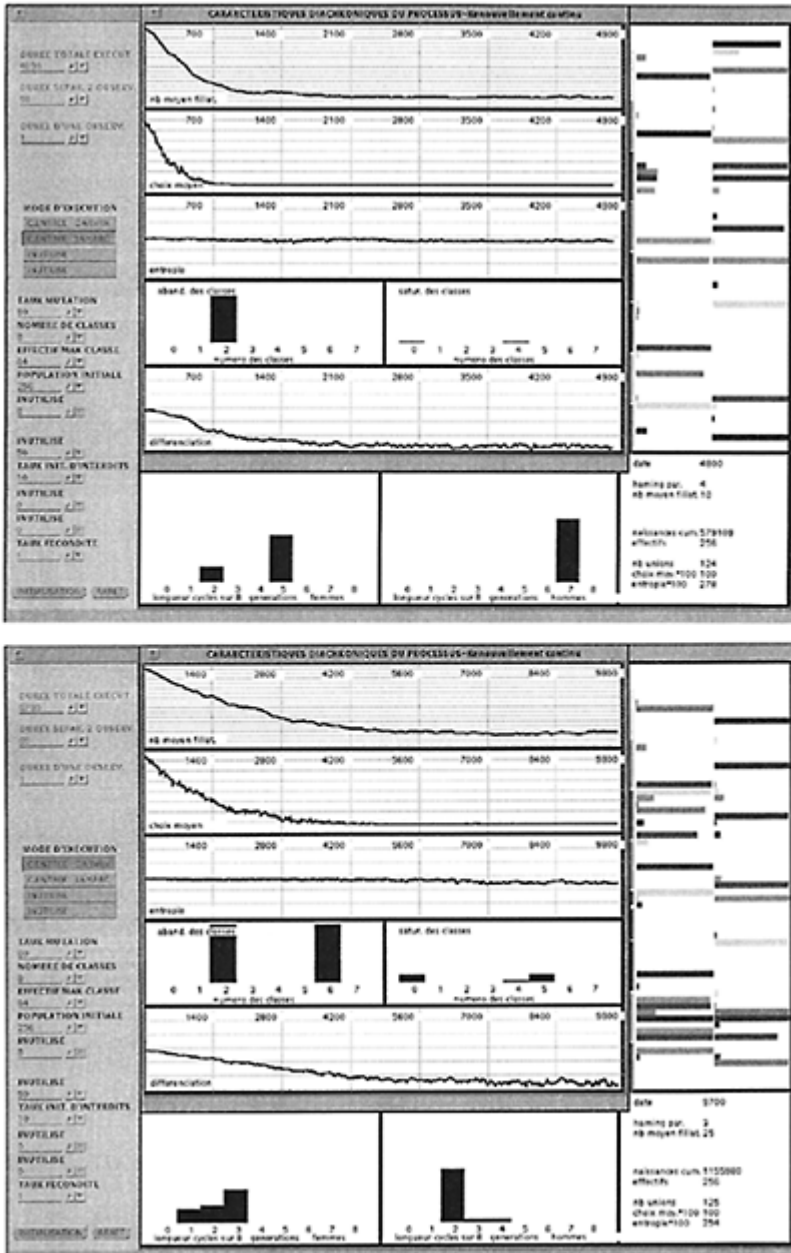


Figure 4.3 The three phases.

important changes in the cycle's structure and in the class size.

It is not always easy to determine the end of this phase precisely. $S(t)$ is the relevant indicator, but with high mutation rate, “noise” is important. The earliest point at which the phase ends is when $S(t)$ stabilizes.

3. The third phase is the equilibrium phase: C has reached its minimum value and remains constant. All the individual kinship models are approximately the same. This common model R_m corresponds to a permutation group (a bijective mapping of the set of occupied classes). The cycle structure is clear and remains constant. However, observation shows that even permutation groups are not completely stable. This residual unstability occurs in Lamarckian processes with high mutation rate, when the permutations do not cover all the initial classes.

We can observe three variations of these common features. They are all related to the second, stabilization phase:

Variation 1. Stabilization ends suddenly with a dramatic fall in the number of occupied classes (see Figure 4.4(a)). It sometimes results in the termination of the reproduction process.

Variation 2. The decline in entropy ceases and future adjustments result in the even distribution of the population over the classes (Figure 4.4(b)).

Variation 3. The situation does not stabilize. (Figure 4.5).

Table 4.1 summarizes the results of the 40 simulations.

States of the system after stabilization

The states of the systems after stabilization are shown Table 4.2. This table gives the equilibrium values of F (average number of filiations per class in the individuals' models), N^* (the number of occupied classes), and the number and lengths of the stabilized cycles for women and men. For example, the third simulation with a Lamarckian principle, a low mutation rate and limitation of the class size is shown in Figure 4.3(a).

For this simulation, in Table 4.2, we have $F=18$, $N^*=7=N-1$ (one unique class is removed). There are two cycles (lengths 2 and 5) for women and one cycle (length 7) for men.

Synthesis

Tables 4.1 and 4.2 allow us to provide a quantitative answer to several questions.

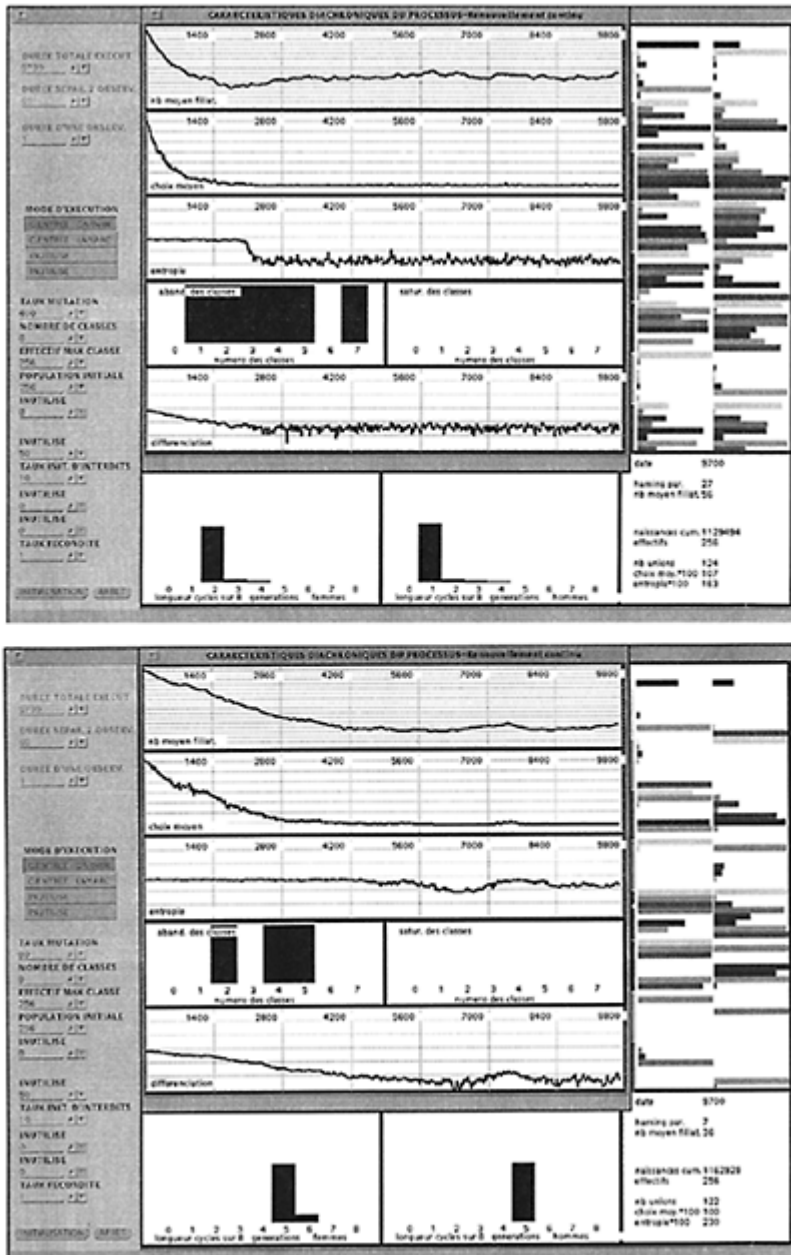


Figure 4.4 Variations of the second phase.

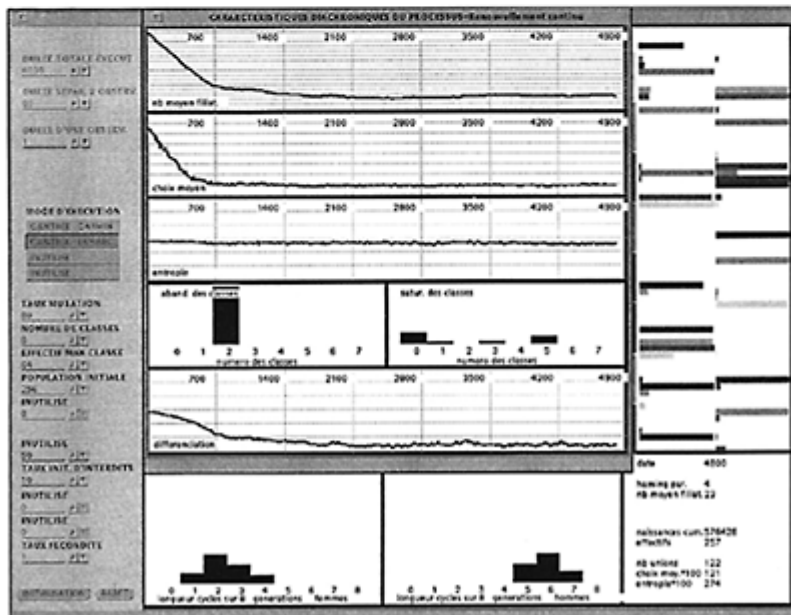
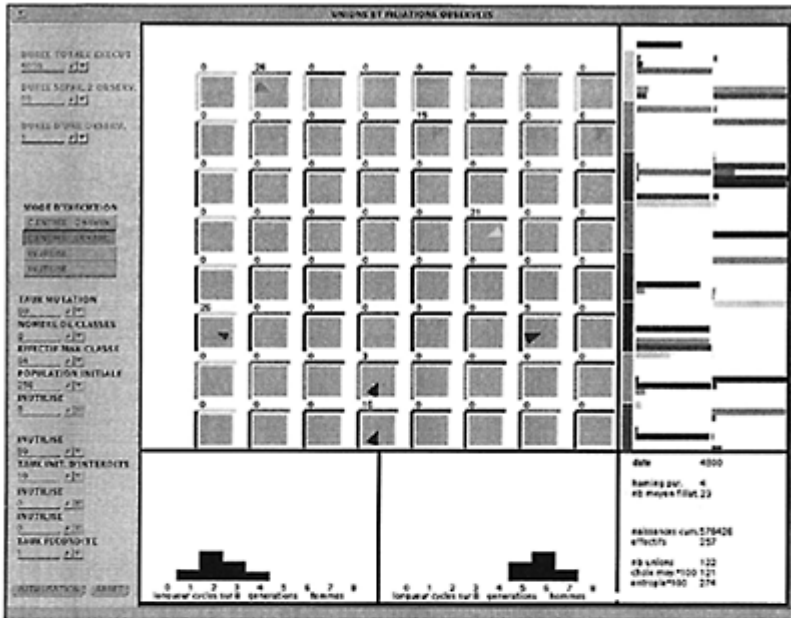


Figure 4.5 How do we determine that a phase is not stabilized?

Table 4.1 Overview of the simulations: diachronic features of the convergence.

(a) Simulations with restriction on the class size.

	Low mutation rate					High mutation rate				
DARWINIAN										
End of the init. phase	4200	3500	2800	2800	4200	1400	1000	1000	2100	2100
End of the stab. phase	7000	7000	–	–	–	7000	8400	–	–	–
End of the observ.	9700	9700	9700	9700	9700	9700	9700	9700	9700	9700
LAMARCKIAN										
End of the init. phase	700	600	700	700	700	350	500	350	350	350
End of the stab. phase	1800	1800	2500	2100	–	1400	1000	700	?	–
End of the observ.	4800	4800	4800	4800	4800	4800	4800	4800	4800	4800

(b) Simulations without restriction on the class size.

	Low mutation rate					High mutation rate				
DARWINIAN										
End of the init. phase	4200	4200	4200	3500	4000	2100	1800	900	1400	1400
End of the stab. phase	8400	8400	–	–	4000	2100	2800	1400	5300	–
End of the observ.	9700	9700	9700	9700	9700	9800	9800	9800	9800	9800
LAMARCKIAN										
End of the init. phase	700	600	700	1000	1000	350	350	350	200	350
End of the stab. phase	1800	2100	2600	2800	–	1800	?	1000	1000	?
End of the observ.	4800	4800	4800	4800	4800	4800	4800	4800	4800	4800

The future of the system: stabilization or extinction?

Stabilization occurs in 27 of the 40 simulations. In twelve other simulations, it did not seem that the process had stabilized. Reproduction stopped in one simulation. This distribution can be considered as representative: it has been confirmed by running the same series of simulations for a second time. The second series yielded a distribution close to the first, with 25 stabilizations, 11 uncertain situations and four terminations of the reproduction process.

Does the stabilized common model strongly constrain the possible filiations?

The value of the average number of filiations F shown in Table 4.2 suggests an affirmative answer. For low mutation rates, F never exceeds 36. It means that, on average, for a given sex s and a given class x , there are at most only two rules (s, x, z) in the model. But we must take into account inactive rules. Inactive rules correspond to classes which are not occupied or only rarely occupied by individuals. When we remove these rules from the model, and when we consider only the remaining rules, the choice of filiations is still more restricted, and it appears that there is often only one possible choice.

Table 4.2 Overview of the simulations: final states.

(a) Simulations with restriction on the class size.

	Low mutation rate					High mutation rate				
DARWINIAN										
Number of filiations F	21	25	26	29	30	42	42	44	40	40
Remaining classes N^*	7	6	7	7	8	4	4	4	8	8
Women's cycles	1+6	1+2+3	–	–	–	1+1+2	4	–	–	–
Men's cycles	2+5	2+2+2	–	–	–	4	4	–	–	–
Incestuous filiations	1	0	–	–	–	1	1	–	–	–
LAMARCKIAN										
Number of filiations F	17	18	18	26	23	20	23	28	33	40
Remaining classes N^*	8	8	7	5	7	8	8	6	5	4
Women's cycles	1+2+5	2+6	2+5	1+4	–	2+6	1+3+4	1+1+4	2+3	–
Men's cycles	8	4+4	7	1+1+3	–	1+2+5	1+7	1+1+4	1+4	–
Incestuous filiations	0	0	1	0	2	0	1	0	–	–

(b) Simulations without restriction on the class size.

	Low mutation rate					High mutation rate				
DARWINIAN										
Number of filiations F	32	36	22	27	–	50	56	57	61	41
Remaining classes N^*	6	5	8	6	–	3	2	1	1	4
Women's cycles	3+3	5	–	–	–	3	2	1	1	–
Men's cycles	2+4	5	–	–	–	3	1+1	1	1	–
Incestuous filiations	1	0	–	–	2	2	1	1	–	–
LAMARCKIAN										

Number of filiations F	18	18	17	20	17	24	27	37	55	47
Remaining classes N^*	8	8	8	6	8	8	7	5	3	3
Women's cycles	1+7	2+2+4	8	1+5	–	1+7	1+6	1+4	1+2	1+2
Men's cycles	3+5	8	1+3+4	2+2+2	–	1+7	2+5	1+4	3	1+2
Incestuous filiations	0	2	0	1	0	0	0	1	0	

For high mutation rates, the same observation can be made. In stabilized situations where all the initial N classes remain occupied, the level of F is also very low. The average level of filiations becomes relatively high in situations where many classes are unoccupied, and therefore where many inactive rules exist. But even in these cases, the number of possible filiations in active parts of the model remains near the minimum.

Does the number of occupied classes remain constant?

The procedure of “veto” involved in the chosen protocol of meetings tends to eliminate rules from the individuals’ models, in both Darwinian processes and Lamarckian ones. Does this removal of rules keep the number of occupied classes near the initial number? Here again, a general feature can be observed. Evolution tends to keep a high number of occupied classes, as may be observed in Table 4.3(a) which gives the relationship between the stabilized situations and the number of occupied classes, N^* . In a majority of cases, the number is greater than 5, even without any constraint on its minimum value.

Table 4.3: Synthesis of the simulations.

(a) Classes remaining after stabilization

Number of remaining classes N^*	Number of stabilized simulations having a given N^*	
	Simulations with restriction of the number of individuals per class	Simulations without restriction on the class size
1	0	2
2	0	1
3	0	3
4	2	–
5	2	2
6	2	2
7	2	1
8	4	4

(b) Exogamous rules

Situations without cycles of unit length	7
--	---

Situations with an unique cycle of unit length	10
Situations with two cycles of unit length	8
Situations with 3 or 4 cycles of unit length	2
Situations with any incestuous rules	13
Situations with one unique incestuous rule	10
Situations with 2 or 3 incestuous rules	4

Are rules of R exogamous?

Table 4.3(b) show the distribution of the stabilized situations according to two relevant characteristics. The first part of the table shows the number of cycles of unit length. These cycles correspond to cases where the class of a child is the same as the class of one of his or her parents. The second part of the table shows the number of “incestuous rules”. These rules correspond to cases where the mother and the father belong to the same class. These figures suggest that the rules of R are exogenous.

Do the resulting permutation groups correspond to real kinship structures?

The cycle structures indicate that the permutation groups obtained do not correspond well with real kinship structures. For real kinship structures, cycles define a partition of the set of individuals from one sex into groups which are equivalent in size and class number. For instance, for the aboriginal society of Warlpiri (Glowczewski 1991), the male population is divided into four “patricycles”, each with two classes, and the female population is divided into two “matricycles”, each with four classes. We did not observe such configurations during those simulations which ended with stabilization. These tended to include quite long cycles containing a greater part of the population.

We have looked for heuristics which could lead to the balanced structures found in real societies. The idea was to use a more restricted condition, W , or a more restricted meeting protocol, U , which imposed the constraint that the individuals who choose the same filiation at time t also belong to the same generation. In every case, we have obtained a faster convergence, but qualitatively similar structures.

Does the mutation rate influence the results?

From Table 4.2, it appears that convergence speed is two or three times higher, with a high mutation rate. Moreover, a high mutation rate seems to differentiate Darwinian from Lamarckian processes. With the former, the population tends to congregate in just a few classes. For a low mutation rate, results for the two processes are quite similar.

An interpretation can be given for these results. In the Darwinian process, the concentration of population is a reaction to mutations. Most mutations have no effect upon existing cycles, which contain a small number of classes (Figure 4.6). They mainly affect inactive rules. In Lamarckian processes, the corrections which are performed during the search for a filiation





Figure 4.6 A case of a concentration of the population on three classes.

offset the disturbing mutations. There is a sort of learning which leads to preserve a larger number of occupied classes.

Does limiting the class size influence the results?

When size is limited, the size of a class must not exceed twice the average size. The mean size of each class is 32; thus the saturation number is 64. The saturation of a given class i occurs when 32 additional individuals—that is, an eighth of the total size of the population—choose their filiation in this class.

Saturation tends to occur at the beginning of the stabilization phase. After this, the number of size-limited classes reduces and all disappear in the last phase if the initial number of classes is preserved.

Two observations can be made:

1. In the case of high mutation rates, the population tends to be concentrated in the smallest possible number of classes.
2. In the case of low mutation rates, limiting the size of classes has no clear influence.

These two observations lead us to propose the hypothesis that limiting the size of classes does not play a significant role in the emergence of regular structures that preserve a large number of classes. The explanation is related to the “veto” procedure. During the stabilization phase the mean configuration R_m of individuals’ models may result in them leaving a class. But, when such a class is empty, no “veto” can prevent this class from being chosen by an individual. After some time, a filiation cycle including this class can be regenerated.

Is the duration of these processes consistent with observed cultural evolution?

Obviously, the answer is “no”. In the case where evolution is the most rapid (Lamarckian processes, high mutation rates), stabilization does not occur until 1800 time units have elapsed. This corresponds to about 900 generations.

The question led us to examine the parameters, conditions W and protocols U which could give more rapid evolution, for a given class number N . First, having a smaller number of individuals has no positive influence. Secondly, more restricted conditions W or more restricted meeting protocols make the evolution faster. But the improvement is not a significant one.

Discussion

The questions that have been raised in this chapter are as follows:

1. Is it possible to build a process based on interactions between individuals chosen stochastically that gradually homogenizes their individual kinship models by making them consistent with a common model?

2. Is it possible to choose conditions for the interactions so that this common model is consistent with models of elementary kinship structures?
3. Is it possible to achieve this under conditions based on what is known about the history of such structures?

Simulations allowed us to answer the first question affirmatively. They also showed that a certain type of protocol which does not seem to be unrealistic leads to models which prescribe exactly the unions and filiations consistent with the preservation of matrimonial classes and the constant renewal of the population. On the other hand, the symmetry properties of elementary kinship structures which allow an equilibrium to be reached between the different groups and which retains their status in successive generations could not be obtained. Moreover, it is not certain that they could be obtained naturally by remaining in the same schema, that is to say, without making use of additional knowledge in individuals' models. Finally, the negative answer given to the third question about timescales leads us to think that the models' diffusion and homogenization follow courses other than those used in these simulations.

However, we think that the simulations are of interest. They show that structuration can be established from very simple conditions. The factor common to these conditions is a transfer of information. The individuals not only act according to direct information affecting their behaviour, but they also take account of the rules of the others with whom they become acquainted during meetings.

Our approach consisted of first making a generalization on the basis of observed behaviour and then testing the processes and the modes of transmission and change that lead from an original society displaying all possible behaviours to the eventual prevalence of that observed behaviour. We shall now introduce a theoretical reflection about this way of explaining cultural evolution.

Cultural evolution and computer simulations

“Individualism” and “holism”

One of the core problems of the human sciences arises from the diverse explanations given of the rules affecting social behaviour and the ways in which the systems that underpin them are represented. To understand the various approaches to this problem, it is useful to refer to the distinction between those that focus on the primacy of the individual and those that focus on the primacy of the “whole” formed by the society (Caillé 1992).

The former approach emphasizes individuals' desires and the means used to satisfy them. It is based on the assumption that individuals have goals, intentions and knowledge. This knowledge refers to the state of society, to what takes place within that society, and the way in which information is used. As a consequence, individuals make choices and form relationships with each other and their environment. This interaction produces a general social process whose characteristics are difficult to identify but which

are none the less regarded as the consequences of hypothesized mechanisms of individual behaviour.

The “holistic” approach accentuates the influence of the whole society on its constituent parts and the way society determines an individual’s degree of autonomy. The determinants that are brought into play vary with the changing trends of holistic thought. Some have tried to explain the rules of society by proposing that the rules enable the society to reproduce and survive by adapting to the changing environment in which it exists. More generally, rules can be considered as the ways used by the society to fulfil a certain general condition. For example, in family systems, marriage rules may be linked to the need to prevent consanguine groups from becoming introverted and to maintain a balance between what each group gives and gets (Levi Strauss, quoted in Mercier 1968).

Caillé (1992) argues that the individualism/holism debate (see also Chapter 8) should give way to the idea of “collective subject” because a society builds itself through choosing its environment and therefore the constraints within which it must live; through choosing the values which it will nurture, in view of these constraints; and through choosing the solutions adopted to resolve the tensions created by these constraints and values.

Tautology in simulations

The charge of tautology has been brought against both the individualistic and the holistic approach (Caillé 1992). In the former case, the tautology can be found in the link between the choice of the individuals and the goals they pursue. In the holistic approach it arises from the link between the problem to be solved and the means used to solve it. Similar criticisms are made of computer-simulated social processes: if a machine generates an artificial world, what is observed can only be the result of programming the rules which govern its existence.

Some methodological precautions must be taken to test the tautological nature of an explanation; they are essentially the same whether computer simulation is used or not. But experiments that can be carried out with the use of a computer do give the researcher an excellent chance to test the results.

In our opinion, the following methodological steps need to be applied in carrying out research by simulation:

1. Identify in the social processes general properties: satisfied constraints, objectives pursued or values to be nurtured.
2. Characterize these general properties by measurement or qualitative means so that their fulfilment can be checked in a real or simulated society.
3. Delimit a space of possible and concurrent behaviours.
4. Determine within this space those types of behaviour that are incompatible with the general properties identified, as well as those that encourage the emergence of these properties, and analyze the minimum hypotheses required for this emergence.

The problem of genesis

A common observation is that all individuals within a given society act within a common cultural setting made up of conceptual cultural tools that are shared by everyone. It is this setting that defines their particular society and distinguishes it from others. It is the existence of these shared elements and their origins that is at the heart of the problem.

The individualistic approach favours the theory that the sharing arises from a historical process, e.g. a certain cultural element, which was originally non-universal and co-existing with other competing cultural elements, may gradually become predominant and finally the unique element. Although the existence of such processes cannot be denied, their ability to cause the emergence of coherent cultural systems is debated. It is somewhat similar to Darwinian theory and the relationship between natural selection and evolution: on the one hand there exist occasional genetic mutations and the selective powers of the environment; and on the other, there is the extraordinary diversity and complexity of living forms and the coherence that each living system must develop in order to survive. The extension of this debate into the study of the genesis of cultural structures has engendered strong polemics.

It is worth reviewing the main points agreed upon between those who study the relationship between biological and cultural evolution:

1. The development of a second means of transferring information in addition to the genetic code: the use of human language and all related systems of mental representation (Gervet 1992). Biological evolution has been governed by the modes of transfer, modification and utilization of information specifically associated with genetic coding. With human language, additional mechanisms for the circulation, transformation and utilization of information are in place. They introduce new dynamics that are exclusive to the cultural evolution observed in human societies.
2. Mechanisms governing biological evolution cause the diversification of species. Similarly, mechanisms governing cultural evolution cause cultural diversity. The effects of the latter emerge in a shorter space of time and are much less influenced by initial conditions and the environment. This explains why cultural forms may legitimately be viewed as free inventions that are not determined by the requirements of the milieu in which they emerge.

The origins of such differences can be sought in the processes relating to the transfer and modification of information. It could be said, for example, that biological evolution is “Darwinian” and that cultural evolution is “Lamarckian” through the inheritance of learned characteristics. But many other aspects of cultural transfer intervene. Cultural transfer is neither exclusively nor fundamentally linked to reproduction, but is accomplished through additional channels: for example, physical encounters between individuals who do not have family ties. The transfer channels, whose identification is part of the transferred information, have in turn become objects of evolution.

Studying biological and cultural evolution in parallel leads to questions about the mechanisms by which information is modified. Although the processes of mutation and biological modification are well understood, our understanding of cultural modification is much less advanced. A similar issue is considered by Lumsden and Wilson (1985) who introduce the idea of “culture-genes” or “units of culture”. The weakness of this concept seems to be linked to the fact that cultural information is structured into interdependent

elements. It is difficult to see how a partial and random modification of such a structure could maintain its consistency.

At this stage, it seems useful to put forward the following hypothesis: human knowledge and representations are arranged recursively in a system of hierarchically organized schemata, each of them delimiting a set of possible changes. Cultural evolution would be analyzed as the result of a succession of permitted changes within schemata, at a different level of this organization. Some of the changes, occurring initially in a few individuals, could be transmitted to others through the usual modes of transmission and under some conditions, snowball to spread to the whole society.

Chapter 5

Cities can be agents too: a model for the evolution of settlement systems

*Stéphane Bura, France Guérin-Pace, Hélène Mathian, Denise Pumain
and Lena Sanders*

Models of global change sometimes take into account the population dimension, with various conjectures about demographic growth. More rarely they include hypotheses about the spatial distribution of the population, that is, the geography of settlement patterns. This neglect is regrettable because most of the relationships between societies and their environment depend on the nature and size of human settlements, at the micro-level of the individual as well as at the macro-level of regional planning. In particular, the general and powerful contemporary trend towards an increasing concentration of population within city systems in many countries of the world has to be considered seriously (Moriconi 1993). The grounds for such a process are not yet fully understood: for instance, many authors have interpreted the negative correlation between demographic growth and settlement size in the industrialized countries during the 1970s as initiating a new trend, called “counter-urbanization” (Berry 1976). The subject is still controversial (Champion 1989).

Settlement patterns have been a source of concern for geographers for a long time. However, a conceptual link is still lacking between the comprehensive theories of urbanization describing the transformation of a rural population into an urban one, and the genesis of the corresponding settlement patterns. It is not clear how a highly skewed size distribution and heavily concentrated spatial pattern of towns and cities emerges from a more homogeneous and more scattered set of rural settlements. Comparative analysis suggests that developing countries may follow a process quite similar to the one that was observed in the industrialized countries during the past two centuries, although involving different parameters values. A model simulating such a process would be useful to test hypotheses about the key parameters of the dynamics of settlement systems, and to predict the possible evolution of systems in various demographic, economic and environmental conditions.

Central place theory is the major conceptual framework that has been employed in this field but, although its dynamic aspects were mentioned in early work (Christaller 1933:86–128), more recently they have received only intermittent attention from modellers. Central place theory deals with the regularities found in the spacing of settlements according to their rank in a hierarchy of size. The theory suggests that the cities with the most important functions (e.g. administration) are spaced further apart than cities with only simple functions. Recently, several authors have proposed to simulate the development of central place systems by means of mathematical dynamic models based upon differences (White 1977, 1978) or differential equations (Allen & Sanglier 1979,

Gamagni et al. 1986). These models use a limited number of equations, defined for each state variable and the same for every spatial unit. They are based on general hypotheses, defined at the macro-geographical level, about the way urban change is occurring in the system.

In White's model, the relationships between retail activities, described by cost equations, and consumers are simulated by spatial interaction equations. He links the spatial dispersion of centres to an interaction parameter and tests the compatibility of the urban hierarchy with a rank-size distribution for the cases of one-sector (White 1977) and two-sector (White 1978) simulations. Allen and Sanglier (1979) use differential equations for simulating the more or less random appearance of central functions at higher and higher levels and their competition within a system of settlements. However, although these authors get final distributions of city sizes which satisfy the rank-size rule, the temporal developments of the spatial patterns of centres are more typical of a market progressively colonized by entrepreneurs rather than the genesis of a realistic central place system, since the models start with two central locations only, sharing the whole region as their market areas (Allen 1978).

Such models seem well adapted for simulating competition between already established urban centres. Modelling with differential equations has some advantages: the model is written in a rather concise way and it can be experimented with under specified conditions with a guarantee of repeatability. However, such models treat geographical space as an isotropic function of distance, using various synthetic spatial interaction formulations; they do not allow the mixing of quantitative and qualitative information; they have difficulty handling a large variety in both the range and scope of spatial interactions; and they cannot consider more than two different geographical scales at the same time. It would have been difficult to adapt any of these models to our project in which we consider the historical development of a set of cities over a long period of time, involving several qualitatively different functional types of city.

That is why other kinds of modelling have been tried in studying the evolution of large settlement systems. Among the first attempts was the pioneer work of Morrill (1962), using a type of Monte Carlo simulation for modelling the development of a central place system, a method first used by Hägerstrand (1952) for simulating the spatial diffusion of innovations. Progress in computer science and artificial intelligence (AI) has now opened new avenues to spatial modelling: for example, with cellular automata. A cellular automaton consists of an array of cells, each of which may be in any one of several states. At each iteration, each cell may change to another state depending on the state of neighbouring cells. Some models, like the famous "Game of life" (Berlekampe et al. 1982), are very simple, but more complex cellular automata could be useful in geography, if they involved more than two possible states and if they allowed for more sophisticated definitions of neighbourhood. Tobler (1979) first mentioned them as the "geographical type" of model, while Couclelis (1985) drew attention to their use for modelling micro-macro relationships within spatial dynamics and for deriving complex dynamics from simple rules (Couclelis 1988). White (1991) applied this formalism to simulating the evolution of land use patterns within urban areas.

Our research relies on the same basic principles as sophisticated cellular automata but has benefited from an advanced modelling method that is known in the field of AI as "multi-agent modelling" (Ferber 1989). Compared to cellular automata, multi-agent

modelling more easily takes into account a large diversity of interactions between the set of cells or “objects” under consideration. In particular, it can include definitions of neighbourhoods that are more flexible and more suitable for the study of urban systems.

In this chapter, multi-agent models will be described first, and we shall see how they can be applied to geographical simulations. A theoretical framework for the development of settlement systems will then be presented, so that objects and rules can be chosen for the design of the model.

Simulation by multi-agent systems

Multi-agent systems are a part of distributed artificial intelligence (DAI) which differs in some way from classical AI (see Chapters 6 and 10). Classically, an AI system is meant to *represent* a human being carrying out a task which needs knowledge, experience and reasoning. A DAI system is conceived as a “society” of autonomous agents working together. They communicate, co-operate and interact in order to reach a global objective. The agents together can treat problems where global control of the local interactions is impossible. The agents have different and sometimes even contradictory knowledge and aims. It is the combination of their local interactions that produces the general behaviour of the system.

Ferber (1989:249) defines an agent as:

a real or abstract entity that is able to act on itself and its environment; which has a partial representation of its environment; which can, in a multi-agent universe, communicate with other agents; and whose behaviour is the result of its observations, its knowledge and its interactions with the other agents.

As we are interested in simulating the dynamics of a settlement system, we will model each aggregate of population as one agent. This is possible because we are only interested in the properties of such a geographical entity as a whole, leaving out the details of the individual behaviour inside the aggregate. Our aim is to analyze the set of all aggregates of population as forming a system, and to model the emergence, growth and decline of the cities belonging to this system. As a first step, we will interpret Ferber’s definition of an agent in the context of simulating the dynamics of settlement systems:

- (a) an agent “*has a partial representation of its environment*”. The information related both to the spatial aggregate itself (the number of inhabitants, functional profile and so on), to its environment (e.g. soil quality), and to the other elements to which it is related (distance, transport commodities, number of inhabitants, etc.), are stored and handled at the local level, which means at the level of the geographical entity.
- (b) the agent “*can, in a multi-agent universe, communicate with other agents*”: because information is treated locally, communication is essential to the working of the system. So the agents are able to receive and send information. For example, information flows can be modelled between a city and surrounding villages, describing the city’s functions, the goods and the services that it supplies. Through this communication process, mechanisms of supply and demand between the different

units of the system can be modelled. The differences of range that depend on the level of the city in the urban hierarchy can easily be taken into account;

- (c) *“the behaviour of an agent is the result of its observations, its knowledge and its interactions with the other agents”*: the rules governing the evolution of each agent are also handled locally. These rules make it possible for each unit to define its own evolution from knowledge of its actual state and of its exchanges with other agents. This formalism makes it possible to consider the city as an evolutionary system. It receives and diffuses information; and it can create and innovate. It can also transform its qualitative functional structure: for instance, a hamlet or a village can become a city at an increasingly higher level in the urban hierarchy. As it competes with others, the reverse can also occur; and
- (d) an agent is a *“real or abstract entity that is able to act on itself and its environment”*. An aggregate of population acts on its environment of course; for example, by exploiting its resources, or building transport infrastructure. The effectiveness of its action will depend upon the size of its population and of its functions.

A multi-agent approach differs in many respects from the sets of differential equations that are often used in dynamic modelling. The differences appear both in the treatment of the data and in the working of the model. An example helps in understanding these differences. The negative effect of distance is a frequently-used component in spatial interaction models. In a differential equation approach, this effect will appear as a function of distance (for example, a negative exponential) inside a global equation that combines many effects and which describes the evolution of various state variables. In a multi-agent approach, the level of generality will be the same: the rules are defined in the same way for all agents of the same kind, but the evolution of the system's state is handled at the level of each geographical unit. In a multi-agent framework, each geographical unit is able to identify the other units of the system, their proximity and their properties by exchanging information. The volumes of migratory flows or trade exchanges between two cities will be computed at the scale of the city, using the rules that are associated with its status as a city, for example. Although the rules have a quite general definition, they are used and handled at the local level of the geographical units that are the building blocks of the model.

The most useful quality of multi-agent system modelling is its flexibility. It is easy to introduce modifications: it is possible, for example, to introduce new behaviour. This kind of change involves the construction of new rules but does not entail changes in those that already exist. Each enrichment of the model is thus made by adding rules that do not affect the global structure of the system. It is even possible to create a new class of agents (a transport network, for example) which is able to communicate with the city agents. In particular, it is possible to test the effects of individual decisions or specific local strategies without modifying the general rules. One could, for example, test the effect of the past-orientated attitude of a town council that hinders a town from growing as quickly as its potentialities might allow. The handling of special cases is very simple as each agent has local knowledge of the general rules that are associated with its status and the specific rules that correspond to its particular case. In this way, it is possible to introduce simultaneously into the model some rules that depend on general laws (e.g. distance decay, competition, supply and demand mechanisms) and others that reflect local peculiarities.

With this kind of modelling it is also easy to handle geographical units that are of various kinds and can change from one qualitative state to another (e.g. hamlet, village, cities of various levels). In this way the genesis of a city and of an urban system can be modelled.

In many respects, this way of conceiving of the evolution of a complex system seems relevant to modelling of the evolution of the urban system. Our purpose is not to study the development of one city but to analyze the evolution of the structure of the whole settlement system. In order to transfer this methodology to our problem, we have to recall the main evolutionary properties of an urban system, the purpose of the next section. The model is supposed to reproduce such an evolution from the modelling of the interactions at a meso-level. The micro-level of individual actors will not be considered here.

The dynamics of settlement systems

Settlement systems are produced by human societies to exploit and control their environment. Some very general evolutionary rules can be derived from the empirical observation of settlement systems in the very long term (for a review see, Pumain 1991). Although the most detailed statistical studies available deal only with settlement systems of the industrialized countries of the ancient world (for instance, De Vries 1984), they allow us to propose general laws of the evolution of settlement systems (see Bairoch 1985). The main regularities are:

- (a) the spatial pattern of settlements is, on the whole, stable while a sedentary agrarian economy remains established within a country: very few new settlements are created during the life of the system;
- (b) the early emergence (Fletcher 1986) and long-term persistence of a strong differentiation of settlements by size within the system according to a regular geometric progression. Hierarchical organization is characteristic of settlement systems;
- (c) a process of internal speciation occurs over time within the system, differentiating types of settlement according to functional division of labour and social representations. This process is partly sporadic, stemming from newly valorized local resources, and partly systematic, linked with the somewhat cyclical appearance and further hierarchical diffusion of bunches of innovations within the system (Pred 1977). Some of these functional or symbolic types remain attached to specific places for very long periods of time; and
- (d) the settlement system expands over time, by increasing the size, number and diversity of its elements. After various stages of limited expansion and decline in a mainly agricultural economy, a major bifurcation has been observed during the past two centuries: the industrial revolution has increased the mean size of settlements and the number and size of levels within the urban hierarchy dramatically, inducing a type of phase transition from a rural to an urban settlement system. The decisive parameters seem to be demographic and economic growth, which allowed a global expansion of the system. At the same time, the accelerating speed of communications contributed to an increase in the range of settlement sizes. As the speed of communications grew, the number of cities needed to provide functions to all settlements decreased. As a result,

only the biggest cities accumulated the functions and the hierarchy got “simpler” because of the growing differences between the cities and smaller settlements (Guérin-Pace 1993).

The historical genesis of urban systems from a dispersed pattern of rural settlement may be interpreted in terms of two alternative theoretical explanations. The first claims that towns emerged endogenously as villages accumulated wealth and developed central functions. Variants of such a “bottom-up” theory either consider that this process is economic by nature and that the emerging towns were at first rural markets exchanging surplus of agriculture within small regions (Bairoch 1985); or that the selection of specific nodes was political, as feudal lords decided to group in particular places all the craftsmen they needed for their own convenience (Duby 1984). According to the second theory, the emergence of towns and cities was linked from the very beginning with the development of long-distance exchanges. Large cities could grow without being supported locally by a rich agricultural production, because they had an especially good situation for long-distance trade (Pirenne 1925, Braudel 1967). The very general process of hierarchical diffusion of innovations within urban systems could support this second hypothesis. Both explanations, “bottom-up” and “top-down”, can be exemplified in the literature about urban history and can be tested in our model.

The main objective of our work is, however, to test the consistency of our hypothesis about the self-organized character of settlement systems: as open systems, the dynamics of their structure at the macro-level are supposed to be generated by interactions of the elements of the system at a meso-geographical level. By simulating the interactions between the local units of a set of settlements, we should be able to produce an overall evolution of the whole system that is compatible with the state of our knowledge about settlement systems dynamics, as enumerated above.

Despite the irreducible uniqueness of the historical path of development of each human settlement (Arthur 1988), it is possible to identify a few rules that may characterize the successive transitions leading from a small agricultural settlement to a large multi-functional urban metropolis. In addition, other rules describing the way settlements interact during their development are needed.

The main transition within the history of a settlement occurs between villages and towns, which are quite different as geographical entities. Villages exploit local renewable resources, levied on their own site. Their ability to sustain and to develop their wealth and their population depends mainly on the quality of the available resources; they are, of course, affected by such external events as climatic fluctuations and military invasions, but they have very little power over such perturbations. A kind of Malthusian equilibrium limits the growth of population according to agricultural production. However, villages may accumulate the surplus of production and commercialize it.

Some of those settlements may then become towns by developing a totally different way of accumulating wealth and sustaining the population. Trading ability or political power is used to capture the wealth produced on distant sites, by means of unequal exchange. Production is also diversified by inventing new products and new skills. The ability to accumulate wealth no longer relies on their own local resources but on the quality of the settlements’ situation within communication networks. The constraints on development depend mainly on competition with other cities for markets and levying taxes on other sites. Town markets are at first confined by transportation costs to a small

neighbouring area, but they enlarge over time with progress in transportation facilities. The scope of market trading also enlarges as towns grow and create more activities that are new urban functions. The “neighbourhood” is then defined in terms of topological and hierarchical distances, measured on the transportation and information networks between urban centres, instead of consisting of a restricted contiguous market area. To our knowledge, this is the first time that the effect of space-time contraction has been introduced explicitly within a model of central place dynamics.

A complete model of settlement dynamics should therefore include explicitly at least three kinds of interaction flow between places: wealth, population and information. In our model, each place, according to its natural resources, its functions and its position within the system, produces wealth. Differentials of wealth between places and the circulation of wealth through exchanges induce the development of the settlement system. The effect of wealth accumulation on population growth is formulated in a simplified way: the ability of a place to increase its population depends on its own wealth. In order to avoid the deterministic effects of that too-simple assumption, population growth rates are computed as stochastic functions of wealth (in the first version of the model, population migrations are not considered explicitly but could be integrated later). The circulation of information is another important component in the dynamic of the multi-agent system: information about available production is provided to neighbouring consumers, and information about existing demand to trading centres. During the course of development, innovations and new economic functions are created, through which places gain the means of producing a wider range of goods with a larger “value added”. Such information circulates between towns and cities according to a hierarchical definition of their neighbourhood. Therefore, the circulation of information between places allows the circulation of wealth. It also encourages competition for accumulating and producing wealth.

In a further step, population flows could be added to the model, following, for example, the suggestions made by Huff (1976) or by Fik (1988) and Fik and Mulligan (1990) about the distribution of migration flows within a central place system. A more general synergetic model of inter-urban migrations could also be included (Haag et al. 1991, Frankhauser 1990, Sanders 1992).

The model

This application has been developed in the object-oriented language, Smalltalk, because of its modularity and flexibility (Goldberg & Robson 1983). Its properties of inheritance and data encapsulation make it appropriate for our formalism (Ferber 1991). The agents are organized into a hierarchical structure of classes according to their characteristics, from the most general to the most specific. For instance, a town inherits all the properties of settlements and these in turn inherit all the properties of places.

Multi-agent models use an anthropomorphic vocabulary that is not well suited to our approach to settlement systems at a meso-scale, but which will be used for convenience. What is called here the “behaviour” of a place or a settlement is in fact the aggregated result of the individual actions of its inhabitants. Each settlement is said to be able to accomplish a set of actions: for example, to consume, to produce or to exchange as if it had autonomous behaviour.

The action space in which the agents operate is represented by grids of variable size and shape. Each cell of the grid is called a “place”. Places can be of various shapes (even irregular) but the most commonly used are squares and hexagons. Places are themselves agents, as they need to communicate to be able to deal with events or information which do not depend exclusively on settlements (e.g. the building of a road, or a natural catastrophe). Each place has a type of natural environment (e.g. a plain, mountain, sea, swamp, etc.) and may contain one or more parts of a network (e.g. a river, road, etc.). It also receives natural (i.e. agricultural, mineral, maritime) resources, according to its nature and position. These resources correspond to a potential to be exploited by the population, the productivity varying with respect to such factors as technological ability or the activity of neighbouring settlements.

Each inhabited place contains a settlement agent. Settlements are characterized by the size of their population, their wealth and the functions they possess (agricultural, economic, industrial, administrative, etc.). The nature of each agent determines its class. For instance, a settlement which possesses only the agricultural function is a hamlet; another with several economic functions is a town, its level depending on its size and the nature of these functions. We assign a part of the population to each function, although hamlets devote the whole of their population to agriculture.

The dynamics are simulated according to a set of rules which define the interactions between agents as well as the constraints imposed by the environment. Some of the rules are deterministic, depending only on constraints imposed by the environment, and others involve random processes that produce a variety of trajectories. Thus two places, although belonging to the same hierarchical level and having the same environment and a similar productivity rate, would not necessarily have the same evolution.

Every simulation proceeds by a sequence of iterations. At the end of each iteration, a balance of the exchanges and the actions performed during the period is made for each settlement. The results modify the settlements for the next step through growth, the gain or loss of functions, environmental changes and so on.

The population P of a settlement is distributed between the different activity sectors according to the functions the settlement possesses. We have the relationship

$$P = \sum_F P_f$$

where F constitutes the functions of a particular settlement. Each time a city acquires a new function, a proportion of its population is devoted to it. This is done by transferring fractions of the total population from old activities to new ones.

Settlements have an evolving demand for consuming various goods and services. The total demand is a function of the total population P , and is distributed among the different goods and services which are proposed. Let γ_f be the propensity of a settlement to consume the goods and services associated with the function f . The total demand D of a place is:

$$D = \sum D_f \text{ and } D_f = \gamma_f P \text{ where } f \in F'$$

and F' is a subset of the economic functions, F .

A settlement usually produces and supplies the maximum that its status allows. Each subpopulation associated with an economic function f has a productivity β_f which determines the value of its supply S_f :

$$S_f = \beta_f P_f$$

The places communicate the amounts of their supplies and demands by means of messages. A place first satisfies its own demand with its own production. It can exchange goods and services with other places according to its economic functions and the amount of its remaining stocks. Productivity, propensity to consume and a range which represents the maximal distance at which the goods are to be exchanged are associated with each economic sector. Trade consists of the transfer of wealth from one entity to another. The total supply of a commercial city is divided between the neighbouring units that are able to buy some of it, according to a function of their wealth and of the distance between them.

The adjustment between supply and demand is carried out iteratively by means of a large number of complex interactions between places. For each cell, we calculate the consumption of goods, i.e. the demand that is satisfied effectively. The price for each good is proportional to the distance between the buyer and the seller. For each sector, we compute an income, which represents the value of the stock of goods that has been sold.

The growth of a place depends on the previous trading balance, and particularly, on the actual profits. According to whether the balance of trade is negative or positive, the growth rate is set above or below the average growth rate of the system. More precisely, it is the result of a normally distributed random variable.

The relative growth rate of the subpopulation of each economic sector depends on the amount of unsold goods in that sector. If a place has some goods left unsold in one sector, this means that its population for that sector has grown too fast or that the market has shrunk. As a consequence, the share of this activity in the total population will be reduced. On the other hand, if a sector makes profit and its market is growing, we increase its population according to a normal random variable. If the share of the population devoted to a given function drops to zero, the settlement loses this function.

Each settlement then is tested to see if it can acquire new functions. The criteria to meet (different for each function) usually depend on size, wealth and which functions the settlement already possesses, and this can involve such factors as the existence of the function in a neighbouring settlement, or be partly related to a random process.

This short explanation of the model shows the simplicity of its inner working. All the complexity goes into the handling of the multiple local interactions between the various components of the system. The framework is very “user friendly”, because the stages of modelling follow very closely the steps of geographical reasoning mentioned above.

A simple example

To illustrate this description of the model, we can observe the results of a specific simulation. The parameters for this simulation are:

1. The environment is a hexagonal grid of about four hundred cells.
2. The length of a cycle is ten years.
3. The initial distribution of resources and population is random over a small range of possible values.
4. There are only three functions available:
 - (i) an agricultural function which is possessed by all the settlements;




Figure 5.1 The simulation after 1000 years.

- (ii) a commercial function with four levels (from basic to high-level goods); and
- (iii) an administrative function with two levels.

As time goes by, agricultural productivity increases with the introduction of new agricultural methods and fertilizers. The range of commercial functions increases over time with the development of new modes of transport and the creation of a transportation network. When a city gets to a higher functional level, it increases the range of all its activities. On the other hand, administrative functions slightly repel each other (for example they cannot appear in two neighbouring settlements).

Even with so few constraints, after the simulated passage of 1000 years, an embryonic hierarchy is clearly visible on the screen (see Figure 5.1). Each settlement is represented with a size proportional to its total population using a symbol corresponding to its functions. Empty hexagons are uninhabited (swamps or mountains).




Figure 5.2 Portions of trading population and cities. The growth is a percentage.

The relative growth rate of the population is calculated for each period (see Figure 5.2). The values fluctuate slightly but in many simulations there is first a period where the growth is close to zero, even sometimes negative, and then it increases.

For a few centuries the growth rate remains moderate, increasing around year 600 to about 5 per cent every ten years. The growth curve shows that the appearance of commercial functions impaired growth until the agricultural population could produce enough wealth to satisfy their demands, around year 600 (see Figure 5.3). In the meantime, the trading cities were better able to increase their wealth and thus grew faster, the




Figure 5.3 Statistics for the system (summing the characteristics of all the settlements).

more so, the higher their level. However, they did not all grow at the same rate because of the competitive market (larger and nearer cities' products are more attractive).

A rank-size representation of the settlement-size distributions over time shows that the process of hierarchalization is very slow (see Figure 5.4).

Figure 5.5 plots the slopes of the rank-size curve during the simulation. It was noticed from previous simulations that a permanent rate of innova-

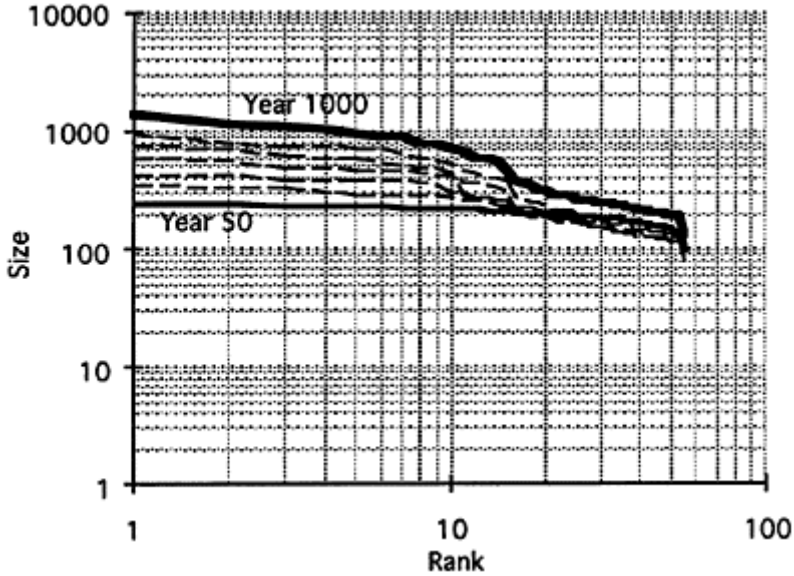


Figure 5.4 Rank-size curves through time (year 50, 200, 350, 500, 650, 800 and 1000).

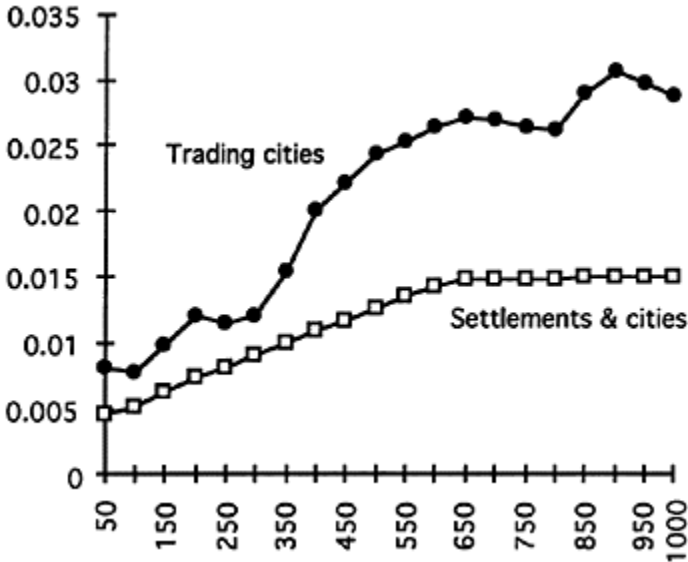


Figure 5.5 Slopes of the rank-size curves.

tion (by the creation of new urban functions) is necessary to produce plausible rank-size distributions.

As the wealth curve (Figure 5.3) shows, the system, having produced a fair amount of wealth, is now ripe for the introduction of more differentiated functions which further emphasize existing differences.

Concluding remarks

The application of tools and methods derived from AI to geography, although very promising (Couclelis 1986), is still in its infancy. We have shown how multi-agent systems could help to build simulation models for studying the evolution of complex geographical systems. They are more flexible and easier to handle than the usual mathematical formalisms and can be built step by step.

This method of modelling is a way to handle the problem of dealing with the double character of the evolution of human settlements from the perspective of global change. On the one hand, rather strong deterministic and general processes lead to a remarkable convergence in the properties of settlement systems all over the planet. On the other hand, diverse local conditions lead to a specific spatial organization in each autonomous territory and confer a unique historical trajectory to every village, town or city.

Further developments of this model should allow us to test the major current questions involved in describing settlement systems as self-organized: is it possible to identify which local properties and interactions, defined at the meso-level of spatial aggregates, can produce the global dynamics that lead from a scattered rural pattern of settlements to

a more concentrated form of human habitat? Is such an evolution, which totally changed the relationship of societies with their environment, still reversible, and if so, under what conditions?

In order to answer such questions, we would also like to use the model to test hypotheses about the driving forces which may lead to different spatial configurations of settlement systems all over the world. We think that two parameters are decisive in controlling the spatial arrangement and degree of hierarchicalization of settlement systems: the average speed of communication which prevails at the moment when the settlement system starts its major bifurcation; and the relative importance of the bottom-up and top-down processes generating an urban organization. That is why we want to test the influence of these two factors on the development and characteristics of settlement systems by means of our simulation model.

Chapter 6

The EOS project: integrating two models of Palaeolithic social change

Jim Doran and Mike Palmer

A previous report on the EOS project (Doran et al, 1994) described the development of a specialized software testbed, and the design, implementation and testing within it of a simple computational version of the Mellars model for the growth of unprecedented social complexity in the Upper Palaeolithic period of south-western Europe (Mellars 1985).

In this chapter we describe new EOS project work and new results, which additionally address Gamble's (1991) alternative model of social change in the Upper Palaeolithic period. We also discuss some important methodological issues that have arisen.

Previous EOS project work

In the EOS project we have set out to create and explore a model which is a computational interpretation of the informally specified Mellars model for the growth of social complexity in the Upper Palaeolithic period (around the time of the most recent glacial maximum) in south-western France (Mellars 1985). Mellars' model relates changing features of the natural environment to the emergence, perhaps for the first time, of centralized decision-making and other related social phenomena within the human population.

Our approach is explicitly experimental. We have endeavoured to identify the essentials of the processes explicit or implicit in Mellars' model, embody them in a working computational model using the concepts and techniques of distributed artificial intelligence (DAI) (e.g. "agents", "negotiation protocols", "reactive behaviour", "planning"), and then study the behaviour of the model in key respects.

We see little alternative to using DAI concepts and techniques if progress is to be made in understanding social dynamics by the use of scientific modelling. It seems unrealistic to model human societies without paying attention to the particular characteristics of human cognition; and to handle the complexity that computer modelling of cognition inevitably entails is impossible without substantial computer support.

The fundamental assumptions that we have made within the EOS model are closely based on Mellars' model and complementary work by Cohen (1985) who has in particular considered organizational and psychological responses to crowding stresses in prehistoric hunter-gatherer populations. The main assumptions are:

1. Macro-level behaviour (the collective behaviour of a set of agents) is best viewed as deriving from the micro level (the behaviour, including cognitive behaviour, of individual agents). An important corollary of this assumption is that extensive reliance on preconceived categorisations of macro-level behaviour is to be avoided.
2. The trajectory followed by a relatively simple human society is greatly influenced by its environmental circumstances. For example, large fluctuations in the availability of faunal or other resources mean that a society will follow one trajectory; but stability of resources means that it will follow another. In particular, it is conjectured that population concentration has a major impact.
3. When modelling human society, any deep insight must address the unique cognitive abilities that humans deploy, e.g. planning and mental representations of one another and of their own society. As noted above, it is this assumption that motivates and justifies the use of DAI and the complexity it brings.
4. A key aspect of human society is the agreement or otherwise between agents' cognitive representations and the reality within which the society exists (what might be called "the degree of socio-cognitive congruence"). Put crudely, if agents believe all the wrong things they will tend not be good survivors, and vice versa.
5. In the Upper Palaeolithic Period, the human population had already developed cognitive abilities akin to our own. The development of human cognition itself is not at issue.

Within the framework established by these assumptions, the EOS project is focusing on the need for individuals to co-operate in resource acquisition and, as a means to co-operation, to conceptualize effectively their physical and social environment. Such issues as kin relationships and organized in-ter-group conflict are not prominent in the Mellars model and therefore do not appear (explicitly) in our work.

The EOS project makes heavy use of DAI concepts and techniques: for example, agents structured as production systems, multi-agent planning, models of acquaintances, and team formation. It also models the relationship between environmental structure and the emergence of social structure (cf. Conte and Castelfranchi 1994). For an introduction to DAI see Bond and Gasser (1988a).

The EOS testbed and the EOS1 model

The main features of the EOS software testbed are:

- (a) a two-dimensional simulated environment or "landscape" with a population of mobile agents and changing resources which provide "energy" for the agents. If an agent does not acquire energy by consuming resources sufficiently regularly, its energy level falls below the target level (we may say that it becomes "hungry"), and if the level falls to zero then the agent "dies", i.e. disappears from the simulation;
- (b) agents structured as production systems (in the AI sense) with rules that "reactively" connect sensory input to action. In addition, there are rules which implement inter-agent communication and which generate, maintain, update and use simple plans and social models (that is, simple representations of other agents and of their social arrangement).

(c) a variety of adjustable parameters including those that specify the behaviour of the resources.

The testbed is implemented in Prolog.

The initial EOS1 model was developed within the testbed, primarily as a particular set of rules and a generic social model within the agents. It established a recurring process of:

1. (Simulated) concurrent and asynchronous information collection, goal setting for resource acquisition, and the proposal of multiple agent plans by individual agents. A multiple agent plan is an allocation of resource harvesting tasks to individual agents with specified timing.
2. Inter-agent communication leading to the adoption of a small number of multi-agent resource acquisition plans for execution, with consequent allocation of, in effect, “prestige” to agents whose plans are selected and executed.
3. Once an agent has acquired planning prestige, semi-permanent leader-follower relationships are created (expressed within the social models of the agents involved), giving the objective appearance of semi-permanent agent groupings.
4. Hierarchy formation (groups of groups of groups...) recursively by the same group-formation processes.

By observation of which agent first acquired each resource, agents came to recognize that some resources were “owned” by particular agents or groups. Agents then planned first for the acquisition of their own or their group’s resources. Since resources were immobile, this implied that a form of territoriality was displayed.

EOS1 experimental results

The experimental results we obtained with the EOS1 model were primarily concerned with the varying circumstances in which hierarchy formation did or did not take place, and the degree to which it took place. We have results indicating, for example, that (i) relatively greater agent awareness of nearby agents and resources (which may be identified loosely with population concentration) encourages hierarchy formation; (ii) inaccurate social models limit hierarchy formation; and (iii) flexibility of commitment of agents to leaders is good for the survival of the agent community as a whole but not for the survival of particular hierarchies. For details of these experiments, see Doran et al. (1994).

“Clobbering” and a new perspective

EOS1 had obvious deficiencies and over-simplifications. For example, the planning processes and the generic structure of the social model within agents were very simple, as were the criteria for creation of the leader-follower relationship and its effects. One limitation was that a hierarchy involving N agents could only be formed if there existed a resource which required specifically the co-operation of N agents for its acquisition. This tight linkage seems unrealistic (though not entirely so), and in the work now to be

reported we have instead linked co-operation to the agents' need to avoid "clobbering". By "clobbering" we mean that one agent may accidentally damage the execution of another's plan (e.g. when two agents plan independently to acquire the same resource, and when only one can be successful).

The need for co-ordination

We have taken as a key growth point of the EOS model the need for agents to avoid "clobbering", whatever the complexity of the resources they are harvesting. Clearly, the avoidance of clobbering is in the interest of all agents (in the absence of explicitly aggressive behaviour in the model). So, in times of scarce resources, agents should seek to plan co-ordination between all the agents of which they are aware. To make this possible, substantially more complex planning processes are needed within agents than were present in EOS1.

Contrasting models of Upper Palaeolithic social change

We have also found it possible to gain a wider perspective and to set clearer experimental objectives by considering an alternative model to Mellars': that put forward by Gamble (1991). Gamble's model, like Mellars', is informally stated and takes as a starting point Palaeolithic art and the need to understand its flourishing in the late Upper Palaeolithic period. Like Mellars, Gamble envisages resource deterioration as a consequence of climatic change around the time of the glacial maximum. But whereas Mellars envisages the consequent growth of decision hierarchies in conditions of localized reliable resources, with complex culture and cave art as a further consequence, Gamble sees an intensification of the negotiation of alliances (not hierarchies) in conditions of uncertain and distributed resources with the art bound up with the negotiation.

Although Mellars and Gamble seemingly disagree about what in fact happened, we see no immediate conflict between their views of the processes involved. They are proposing different sociocultural reactions to different environmental circumstances.

Accordingly, we set out to demonstrate that if a suitable set of assumptions is made in the EOS testbed, then:

1. Lowering of resource availability without resource concentration and with unpredictable variation in richness of individual resources leads to greater numbers of alliances.
2. Lowering of resource availability together with their spatial concentration leads to hierarchies.

In both cases there is an obvious but important condition: that there are appropriate cognitive mechanisms within the agents. Once these two proposed social trajectories are juxtaposed, it is clear that there are other questions which might be asked. For example, what happens if resources are unpredictably concentrated, i.e. if the concentrations occur in different localities at different times?

To address Gamble's model within the EOS testbed requires something of a change of scale. The area considered by Gamble is much greater than that considered by Mellars (much of Europe, rather than valleys in what is now south-western France) so that

intuitively there would be a much greater relative immobility and territoriality amongst agents. We must therefore think a little differently about the relationship between a unit of distance in the testbed and distance in antiquity. And, of course, we now need an interpretation of “alliance” within the model.

EOS2

We now describe and explain our new extended computational model, EOS2. Most extensions have been implemented by adding new rules to, or modifying existing rules in, the standard production rule set associated with agents (with its supporting function library) and by simple extensions to the generic social model and other representational structures within agents.

Agent planning

Agents in EOS2 must be able to plan the co-ordination of many agents, not to acquire complex resources (as in EOS1) but rather to avoid inter-agent clobbering. This requires them to undertake a more complex form of planning than in EOS1. Agents decide the resources to target and the number of agents needed to acquire them, and then seek to make an agreed allocation of agents to resources. A single agent and its following may be assigned several resources to harvest. Plans have explicit estimated numerical payoffs. The way in which the anticipated payoff of a plan is calculated takes the risk of clobbering into account (among other things) so that agents can “see” the advantage of co-ordination.

It is important to keep in mind that agents are planning in terms of what they believe to be true, rather than what in fact is true. Although agents’ awareness of resources and other agents is accurate (if possibly incomplete) in all the experiments we describe in this chapter, the judgements that agents make are heuristic and not tightly related to the reality of the agents’ world. In particular, an agent may wrongly believe that one plan is likely to be more rewarding than another, or may believe that a proposed plan is sufficient to meet its needs when it is not, and vice versa.

Plans are built incrementally by an agent, thus:

WHILE

- (a) the estimated payoff of a plan to the planning agent is less than the agent’s believed energy level deficit;
- (b) the size of the plan is less than the maximum plan size for the agent;
- (c) adding a new resource will increase the estimated payoff of the plan to the planning agent; and
- (d) there is a new resource which can be added

DO Add an appropriate new resource to the plan.

There is a distinction between the total expected payoff of a plan and the expected payoff of a plan to the planning agent. The latter takes into account the need to share out acquired resources amongst all the agents involved (which may or may not include followers and allies of the planning agent—see later in this chapter). Factor (a) captures a

notion of sufficiency: an agent stops planning if it has a plan which looks as if it will suffice to meet its current energy deficit. Factor (b) captures a notion of bounded rationality. Different agents have different limits on the size of plan they can build. Within factor (c) a balance is struck between the assumption that the smaller the plan, the more vulnerable it is to “clobbering”, and the larger the plan the more generally unreliable it is. Factor (d) requires the agent to refer to its model of resource availability. In general this is incomplete and is built up both by the agents’ own observation and by information coming from any followers the agent may have.

In a resource rich environment, the agents will typically believe that small “individual” plans suffice. However, as the number of resources and their richness decline relative to the number of agents, small plans will be judged to be less rewarding and no longer sufficient (both because individual resources are less rich and because the likelihood of clobbering increases) so that agents will tend to build and execute more complex plans. That in turn means that the “bounded rationality” limit on the size of plan an agent can build will become much more significant. Notice that this planning mechanism embodies a notion of “economic rationality” (e.g. Doyle, 1992).

Plan adoption and execution: the beginning of co-operation

In principle, agents all plan concurrently and asynchronously. Which agents’ plans are adopted for execution? This is a complex matter. In EOS2, agents select and invite other agents to join the plans they have created (selecting first their own followers). Agents adopt the plans that they judge to be potentially the most beneficial to themselves in terms of their own current beliefs: either they persist with their own plan, or they agree to join that of another agent. The effect is that, with some delay, the more highly rated plans are wholly or partially adopted for execution by groups of agents. One of the agents in each group is the originator of the plan and is therefore viewed by the others in the group as potentially a leader.

Plan execution is also complex. It involves the setting of a rendezvous point and time, and the pooling and sharing at the rendezvous of all the resources harvested. Further complexity is created by the possible failure of some or all of the agents to perform the tasks they have agreed to undertake, and the fact that the whole process is partly recursive: a subleader will execute its part of a plan by creating and recruiting to a subplan and so on.

The properties of this co-ordination process (which involves a form of “team formation”, cf. Kinny et al. 1992) are complex and difficult to establish theoretically. The process is not intended to be in any sense optimal—rather, it is intended to be effective and plausible in its broad characteristics.

Alliances in EOS2

Recall that any relationship instance is expressed by beliefs within the social models of the agents concerned. For example, to say that agents *X* and *Y* are in a leader-follower relationship means that Agent *X* believes that Agent *Y* is its leader, and vice versa.

In EOS1 such a leader-follower relationship came into existence whenever one agent (which might itself already be a leader) agreed to join another’s plan. This is

unattractively simplistic. Further, to address Gamble's model some notion of an alliance must be introduced. So key questions for EOS2 are: What is an alliance? And: When does an instance of an alliance or a leader-follower relationship come into existence?

The principle we have followed is that after multiple instances of co-operation between two agents (where an instance of co-operation means that one agent has agreed to take part in the other's plan) an alliance will be formed. When two agents are in an alliance, they exchange information about their needs and give priority to incorporating one another in their plans. The natural conjecture is that such an alliance will lead to rather more informed planning and hence to rather more efficient resource harvesting.

An instance of a leader-follower relationship, however, comes into being when co-operation is consistently "one-way". If agent *X* is consistently recruited to agent *Y*'s plans over a limited period of time, then both *X* and *Y* will come to see themselves as being in a leader-follower relationship, with *Y* as the leader. It follows, other things being equal, that this relationship is inherently less likely than that of an alliance. The consequences of a leader-follower relationship are the same in EOS2 as they were in EOS1: leaders tend to generate more highly valued plans and they attempt to recruit their followers to their plan first. A follower supplies information to its leader about its needs and its awareness of resources, and the follower will always wait for a plan to be proposed to it by its leader (but may then reject the leader's proposal in favour of an alternative).

A leader-follower relationship can evolve from an alliance and both types of relationship can break down if the agents involved lose contact with one another for a sufficient period of time.

The onset of dominance

How can it come about that one agent repeatedly proposes plans that other agents value highly and in which they therefore agree to participate? We have identified a range of possibilities capable of implementation within the EOS2 model. They may be classified briefly and listed as below:

Factors determining the *actual* quality of the plans an agent proposes:

- (a) Variation in agents' environmental circumstances:
 - (i) more central location as regards resources; and
 - (ii) availability to agent of "dominating actions", i.e. actions which enable it to influence strongly the payoff of another's plans.
- (b) Variation in agents' "cognitive" characteristics:
 - (i) variation in algorithms or bounded rationality; and
 - (ii) variation in knowledge (e.g. by better sensors).
- (c) Variation in agents' social circumstances:
 - (i) "ownership" of majority of the resources in question; and
 - (ii) the agent already has a considerable following.
- (d) Variation in agents' social behaviour:

- (i) tendency to be “greedy” or “altruistic” i.e. the agent proposing the plan may seek to retain more or less than its proportionate share of the proceeds.

Factors determining the *claimed* quality of the plans an agent proposes:

(a) Variation in agent’s “cognitive” characteristics:

- (i) variation in optimism (“my plans are always great”); and
- (ii) propensity to lie to other agents: for example, about the estimated value of plans.

From among the foregoing factors, we have implemented in EOS2 a simple form of “bounded rationality”, together with a simple form of “optimism”. The former is a matter of agents possibly being set to have differing limits on the size of the plans they can create. The latter is a multiplier that agents employ to scale up or down the value that they put upon plans when they propose them to others. Different agents may be set to have different multipliers. It is also the case that agents assess plans more highly (because they are estimated to be more reliable) when they intend to harvest their own resources. These factors are all parameterized so that they can be varied between experiments.

Results

We indicated earlier the collective behaviour that the Mellars (1985) and Gamble (1991) formulations suggest: as the number and richness of the resources available to the population decline *without* resource concentration and *with* unpredictability of individual resources, so (following Gamble) the number of inter-agent alliances should increase. But if there is resource decline coupled with survival of individually, stable resources in limited areas then (following Mellars) population concentration and development of complex hierarchies is predicted. In fact, our experimental findings are only partially in line with these predictions.

We find that many factors bear upon the dynamics of the agent community and the alliances and hierarchies within it. Three of the most important are:

1. *Resource availability*. Can the agents acquire sufficient resources to survive? This involves the maximum rate at which agents can harvest, as well as the actual rate of resource renewal. Inadequate uptake rate means that agents must die.
2. *Resource dispersion*. The more spatially dispersed the resources are, the more difficult it is for agents to co-ordinate.
3. *Agents’ “hunger” levels*. Is the present energy level of the agents up to their target level? If so, they will construct minimal plans and no co-ordination will occur. If not, they tend to plan more because agents “believe” that more complex plans avoid clobbering and hence yield a better harvest.

Recall that in all the experimental trials reported here, agents create and work with accurate, but possibly incomplete, social models. Additionally, all resources in these experiments are “simple” in the sense that each can be harvested by a single agent. Resources may differ, of course, in their energy yield when harvested.

We have adopted simple measures of alliance and hierarchy. The alliance measure is simply the number of alliances (between pairs of agents) divided by the total number of

agents. The hierarchy measure is the sum of the hierarchy depths of the agents in the hierarchy, again divided by the total number of agents. A leader has a hierarchy depth of zero, one of its followers has depth one, a follower's follower has depth two, and so on. The depth of a hierarchy is defined as the maximum agent depth it contains. In graph theoretic terms, a hierarchy is a tree rooted at its leader.

A "low resource period" trial

The impact of two of the factors just listed, resource availability and agent "hunger" levels, on hierarchy and alliance formation is illustrated by the following experimental trial. A group of 12 agents is subjected to a resource regime consisting of an initial period of ample available resource energy, a comparatively long period of very low available resource energy and then sustained ample resource energy again (see Table 6.1).

The resources are spatially semi-distributed: randomly scattered over an area of 500×500 distance units within a 1000×1000 environment (see Figure 6.1).

Resources renew every 30 cycles. Agents "consume" one energy unit each per cycle. They begin with energy levels of 1000 and with a target energy level of 1000. Agents' awareness range is set at 200 units of distance

Table 6.1 The resource regime for the experimental trial.

Cycles	Resources
1–200	Ample resources (12, each of yield 35)
201–1200	Severely inadequate resources (12, each of yield 1)
1201–2500	Ample resources (12, each of yield 35)

Cycle Number: 61

log file:

initialise	load	run	step	pause	present	past	past1	break	exit
------------	------	-----	------	-------	---------	------	-------	-------	------




Figure 6.1 Snapshots of the EOS environment at moments in the “low resource period” trial. Dots indicate the locations of agents, and circles containing 1s the location of resources. (a) After 61 cycles, the agents are mostly located at resources and there are no relationships between them.

and an agent’s movement in each cycle is at most 50 units.

We observe (see Figure 6.2) that in the initial phase (to cycle 200) the community survives easily without any complex activity. As the reduction in resources makes an impact, from about 750 cycles, agents begin to build more elaborate co-ordination plans

and co-ordination structures appear (i.e. alliances or hierarchies), but then (at around 1200–1500 cycles) deaths

Cycle Number: 1000

log file:

initialise load run step pause present past past1 break exit






Figure 6.1 (b) After 1000 cycles, a number of relationships between agents have been established. A solid then dotted line from one agent to another implies that the first agent is the leader of the second. A line connecting two agents but dotted throughout its length implies an alliance.


(a) Total Available Resource Energy



(b) Average Agent Energy



(c) Hierarchy Measure



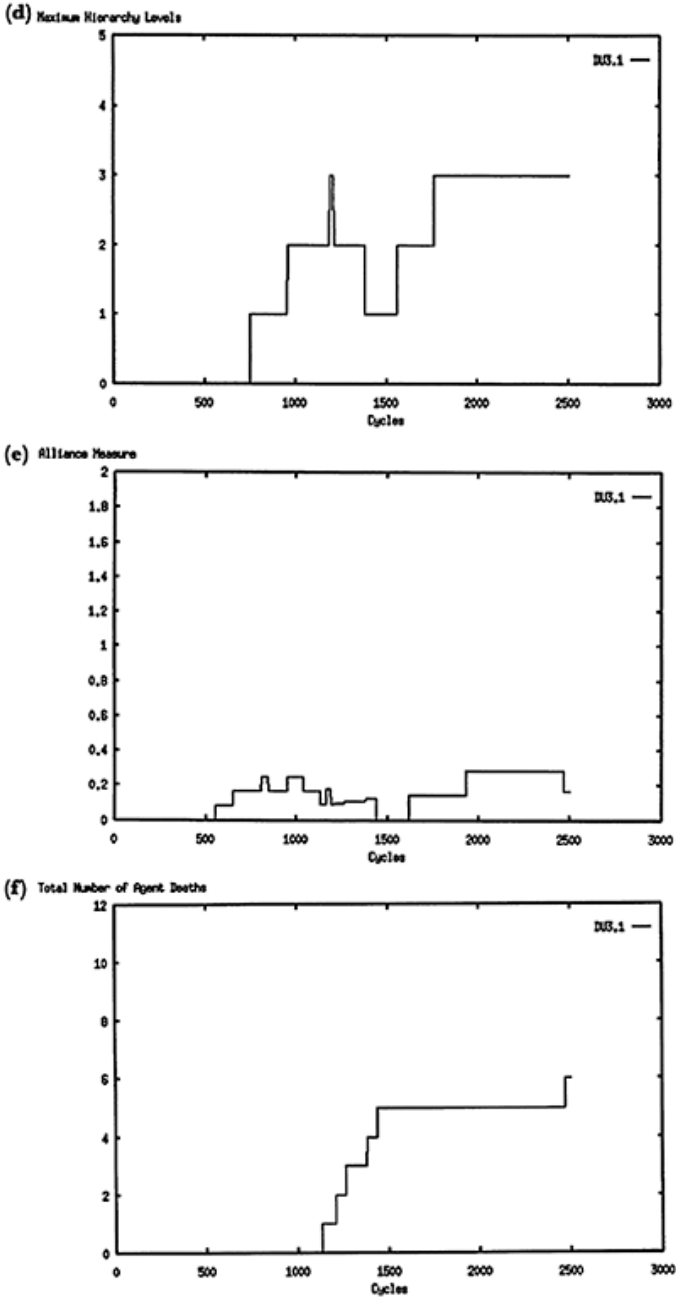


Figure 6.2 (overleaf) Graphs drawn from the “low resource period” trial: (a) the resource energy available in the

environment over the 2500 cycles of the trial. Note the very low available energy between 200 and 1200 cycles. The short-term fluctuation in levels is the result of agent harvesting; (b) the average (surviving) agent energy levels; (c) the hierarchy measure over the 2500 cycles. Note the partial collapse after about 1500 cycles; (d) maximum depth of hierarchy at each cycle; (e) the alliance measure over the 2500 cycles; (f) cumulative agent deaths.

occur and part of the co-ordination structure collapses, even after resource availability has returned to normal. After this period of partial collapse the agents continue with high levels of co-ordination planning (because their energy levels at that time have dropped well below their target level) with the result that even more substantial co-ordination structures emerge. Note that the rate of increase in average energy level when resources return to “normal” is relatively low—hierarchical co-operation is not efficient in this particular environment.

This trial is relatively simple. It involves somewhat distributed resources and therefore agents. It does not involve variation between individual agents or short-term unpredictability in resources. We have conducted further trials to investigate the impact of these additional factors.

Multi-factor trials

We now report further trials exploring the impact of additional factors upon agents’ survival and upon degree of hierarchy and alliance formation. We have addressed the following three conjectures:

1. Spatial concentration of stable resources encourages the formation of hierarchies.
2. Variation between agents (for example, some agents better able to plan than others) encourages the formation of hierarchies.
3. Short-term unpredictable fluctuation in the energy provided by individual resources encourages the formation of alliances rather than hierarchies.

All trials are set in a resource “upturn”: 12 agents start with low energy levels (100) relative to their target levels (1000) and are set in an environment with relatively plentiful resources (12, each of average yield 25, renewing every 20 cycles). As before, agents “consume” one energy unit each per cycle. Agents’ awareness range is again set at 200 units of distance and an agent’s movement in each cycle is at most 50 units. Note that to work with more than 12 agents would require a very large amount of computer processing.

Table 6.2 shows the results of 3×2×2 trials, each involving two replications with two different pseudo-random number streams. For each replication of each trial, results are presented after 1500 cycles: the number of agents surviving, the values of the alliance and hierarchy measures, and the average energy level of the surviving agents. The trials vary three factors:

1. The degree of resource concentration/dispersion (concentrated: re-

Table 6.2 Results from a set of 12 trials, each in two replications with different pseudo-random number streams. Note the much greater level of hierarchies and alliances when resources are spatially concentrated.

	Concentrated				Intermediate				Dispersed			
	Ag	Al	Hi	AE	Ag	Al	Hi	AE	Ag	Al	Hi	AE
Base case	10	1.4	1.1	463	5	0.2	0.8	294	5	0.0	0.6	271
	8	0.6	1.6	354	5	0.2	0.8	243	7	0.0	0.4	524
Varying resources (VR)	8	0.4	1.5	524	6	0.3	1.2	249	3	0.0	0.7	82
Varied agents (VA)	9	0.1	2.6	242	3	0.0	0.7	262	6	0.2	0.7	398
	11	0.5	1.3	384	5	0.0	1.2	237	6	0.2	0.7	463
VA+VR	7	0.0	1.7	334	5	0.0	1.2	257	4	0.0	0.5	302
	6	0.2	2.2	398	4	0.0	1.0	233	8	0.0	0.5	435
	12	1.0	0.4	397	5	0.0	0.6	380	6	0.0	0.5	322

Key: Ag, agents surviving; Al, measure of alliances; Hi, measure of hierarchies; AE, average energy of surviving agents. All results after 1500 cycles.

sources in a 100×100 area; intermediate: resources in a 600×600 area; dispersed: resources in a 1000×1000 area).

2. Whether there is random variation between agents (in “bounded rationality” and “optimism”).
3. Whether resources are renewed with random variation about their mean energy level.

From Table 6.2 it is apparent that agent survival is greater, and hierarchy and alliance formation are much more extensive, when resources are spatially concentrated. Somewhat unexpectedly, however, there is no clear relationship between agent variation and the formation of alliances and hierarchies, nor is there a clear relationship between random resource variation in time and alliance formation. As before, there is some suggestion (in the “intermediate” trials) that co-ordination tends to lower the average agent energy levels (i.e. it is counter-productive). There is substantial variation between replications, suggesting that the outcome of any particular trial is highly sensitive to the spatial patterning of the resources involved.

Summary and discussion of results

It would be quite wrong, of course, to put undue weight upon our results. The implications of our model in its different variants have not been fully explored, and the arbitrary nature of many of the assumptions made within it is apparent.

With that caveat, we find that our results do not quite agree with our expectations. Although our agents do form extensive hierarchies when resources are spatially concentrated relative to their awareness range and when total energy uptake falls below their target expectations, this is not the whole story. First, when total resource energy availability is low, the hierarchies may collapse quickly as a result of the deaths of constituent agents and, secondly, it is after the end of a resource “famine” that the strongest hierarchy formation occurs, when agents are still seeking to cooperate to survive, although their situation is, in fact, secure. Thus our results are in line with Mellars’ (1985) view, but suggest adding some additional complexities to it. On the other hand, we do *not* find that dispersed and unpredictable resources lead to alliances as suggested in (our interpretation of) Gamble’s (1991) informal model.

It may have come as some surprise that in these experiments co-operation is counter-productive. That is, agents harvest fewer resources if they co-operate than if they do not. But the reason is quite simple: organizing and sustaining co-operation takes time and this disadvantage outweighs the beneficial reduction in “clobbering”. Agents co-operate not because it is good for them, but because they believe it to be good for them. However, the negative effect of co-operation would almost certainly disappear were we to set complex resources in the environment, that is, resources requiring the collective effort of two or more agents to be harvested.

Extending the model

We see great potential for development of models of the EOS type. By way of illustration, we briefly discuss the emergence of leadership. In the present model the agents who become leaders are those who offer, or seem to offer, the best plans. We have not yet experimented with “greedy” or “altruistic” leaders, that is, leaders who reserve to themselves more or less than their “fair” share of a multi-agent plan’s payoff, but it would be easy to do so. And we have not addressed the issue of agents who have the ability to “damage” the interests of others (for example, by destroying resources in another agent’s territory, or by destroying the agent itself) and use that ability to establish group control by, in effect, making all plans other than their own highly unattractive.

But there are other possibilities. Boehm (1993) has drawn attention to substantial ethnographic evidence supporting a quite different view of the transition from egalitarian to ranked society. He envisages the failure of mechanisms which restrain potential leaders, rather than the occurrence of a positive opportunity. Can this conjecture be expressed within the EOS framework? In principle, the answer seems certainly to be “yes”, but a further substantial increase in complexity would be required, seemingly including the ability of agents to represent and manipulate within their social models alternative futures involving the appearance of a new dominant leader and the potentially negative consequences.

Finally, we note that there is good reason to believe that consideration of economic rationality, however comprehensively captured within the model, is not enough. It is intuitively apparent that being a leader has to do with the dynamics of “emotional energy” (Randall Collins’ phrase) as well as good decisions. A number of authors (e.g. Alexander 1989, Collins 1990, Hammond 1990) have considered recently the role played by emotions in society, including the emergence of stratification. To simulate aspects of the flow of emotional energy in a model of the EOS type would certainly be possible and might suggest rather different conditions for the emergence of leadership and hierarchies.

Methodological issues

The EOS project raises seemingly important methodological questions. Science is about finding useful and tractable models, whether mental, informal or formal. The attraction of DAI modelling in this and similar contexts is easy to explain. It enables cognition to be explicit in the model (cf. treatment of “endomorphism in agents” in Zeigler’s (1990) terminology) and encourages the use of symbolic rather than solely numerical representation. But, as always, models have to be developed in a valid way if their results are to be trusted.

A prominent feature of both the EOS1 and EOS2 models (and indeed of all DAI-based models) is the large number of adjustable parameters, numerical and symbolic, embedded within them. As these parameters are varied the behaviour of the model system varies, sometimes dramatically. Getting “interesting” behaviours, or a clear understanding of particular types of behaviour, is rarely easy. We feel that a number of important questions arise from this observation.

1. *Validity* How can any model be taken seriously when there are so many arbitrary parameters within it? Cannot any desired behaviour be obtained from a model with a sufficient number of adjustable parameters?
2. *Classification of behaviour* How can the set of model behaviours be classified?
3. *Abstractness* Why should any one particular model be used when there are clearly a range of alternatives to hand at differing levels of abstraction?
4. *Formality* Is there a need for formal support for the work?

We now attempt to give partial answers to these questions.

The validity of the EOS models

The EOS models have not been, and clearly cannot be, validated in anything like the traditional sense of the word. The many assumptions within the models have not been supported empirically, some of the claims made explicitly in connection with the agents are controversial even within the DAI community (e.g. our reference to agents’ “beliefs”), and the actual behaviour of the models has not been and cannot be checked against reality in any detailed way.

But this lack of “validity” in the traditional sense is not as damning as might at first appear. What the EOS models do is to give relatively precise interpretations to the informal models of Mellars (1985) and Gamble (1991) and enable relatively precise study

of their consequences. That is an important step forward. Both Mellars and Gamble (and others) are proposing possible processes which they feel may fit the limited available evidence. Their emphasis is on insight and plausibility, rather than a precise fit to particular events in antiquity. So is ours. There is an important sense in which our work, like that of Mellars and Gamble, is more theory building than it is modelling. The aim is to discover more about certain processes and how they interact. By so doing, the way is prepared for more traditional modelling studies in the future. Yet at the same time we do not exclude the possibility that testable hypotheses may be derived directly, as experimental exploration focuses our attention on particular relationships between structures and parameters.

It is not the case that any behaviour may be obtained where there are a multitude of structural parameters. But in such circumstances experimentation must go beyond a small number of trials, however interesting, to the formulation of a coherent and sound overview of the possible behaviours that may be obtained from the model, and in what circumstances they occur.

Understanding a model's behaviour

Even when the usefulness of models of the EOS type is accepted in principle, there remains the problem of how best to comprehend and to express so much possible variation. It is not clear whether the behaviours of these models may be relied upon to fall into simple patterns (as seems implicit in the informal models from which they are derived) or whether we are faced with something more like chaos—a multiplicity of structural parameters, with unpredictable sensitivities and no simple patterns of behaviour to be discovered. Should the latter be the case, progress may be very difficult indeed.

Choosing between differing levels of abstraction

A further issue is how best to make the choice between the many alternative models, at varying levels of abstraction, whose behaviour might be studied to address a particular problem. For example, we may imagine a model, called ABSTRACT-EOS, whose operation is as follows:

FIRST

- scatter resources over a two-dimensional region according to a multi-normal distribution and agents over the landscape according to a uniform distribution

THEN REPEATEDLY

- compute a (locally) optimal multi-agent/multi-resource plan for each agent based upon the resources and other agents within the agent's sensory range (factoring in any existing alliances and leader-follower relationships);
- compute the set of plans to be executed as that subset of the agent plans which satisfies the constraint that no agent should appear in more than one plan in the set, and which provides the maximum total payoff;

- compute the payoff from the plan set to be executed, and discard/ replace any agents that “die” (note that “clobbering” may occur); and
- insert and discard examples of alliances and dominance relationships by simple rules of the type used in EOS2,

WHILE monitoring the dynamic pattern of alliances and leader-follower relationships.

Using this model, we might expect to demonstrate that agents located in the midst of many resources become high leaders, but with the results sensitive to the size of the agents’ sensory range and the parameters of the resource distributions.

The important feature of ABSTRACT-EOS is that it bypasses both the internal cognition of agents and inter-agent communication, while retaining other essentials of the EOS model. But it loses the important potential match or mismatch between agents’ social models and social reality. The questions are, which is the more useful model, EOS or ABSTRACT-EOS, and why?

A standard modelling principle is that the complexity and level of a model should be chosen so that it answers the questions and embodies the theoretical elements we are interested in, but is otherwise as simple as possible. Complexity for its own sake, including “cognitive” complexity, is to be avoided. Agents in a model should not, for example, be programmed to do predictive planning if that complexity is not clearly needed for the modelling objectives to be achieved. Does that principle apply here, and if so, how is the judgement of need to be made? It might be argued, alternatively, that when modelling human society we should focus on the “human” level, and regard higher phenomena as emergent from that level, if only because that makes everything easier for us to understand. We feel that this issue is unresolved.

A need for formal support?

It may be suggested that the difficulties with the EOS models just discussed demonstrate the need for some more formal framework within which these models should be set (or by which they should be replaced). This argument will appeal particularly to those researchers within distributed AI studies who are taking a more logicist road to the study of teams and societies (e.g. Kinny et al. 1992, Werner 1990).

The argument is attractive but is, we believe, less persuasive than it at first appears. Such a formal framework typically will be a system of axioms and inference rules, or some partially formulated version thereof. But how can this help? Certainly, certain types of general inference could be made and abstract theorems proved. But in the context of human society, a formal framework will always be either a gross oversimplification (at best it will illuminate certain essentials) or replete with domain specific axioms and tractable in particulars only as a computer simulation of essentially the EOS type.

Unquestionably, any model should be designed in a principled way. But we must avoid imposing additional mathematical requirements which in practice compel a model to be too simple to be of value and which may even focus attention away from the essential requirements of the modelling task.

Conclusions

We feel that we have already achieved some useful and original insight into the close relationship between the Mellars and Gamble models. We find that our experimental results, as far as they go, are compatible with the Mellars (1985) view but not yet with (our interpretation of) Gamble's (1991). But much more experimental work is needed. And there are important methodological issues that have yet to be fully clarified and resolved.

Acknowledgements

We are grateful to our colleagues Nigel Gilbert and Paul Mellars for many valuable discussions related to this chapter, and to Nigel Gilbert for making available additional computing facilities at the University of Surrey. The EOS project has been supported by the UK Joint Council Initiative in Cognitive Science/HCI by grant no. SPG 8930879.

Chapter 7

Genetic algorithms, teleological conservatism, and the emergence of optimal demand relations: the case of learning-by-consuming

Roger McCain

The demand relationship (the relationship between the price of a good and the quantity of that good demanded) has long been a central construct of economic theory. Economic theory conventionally assumes that consumer demand curves are determined by the “maximization of utility” or preference. That is, the consumer demand relationship is the locus of price and quantity pairs such that the quantity demanded yields maximum utility or preference ranking at the given price.¹ However, a large body of evidence from cognitive science, experimental economics and management studies indicates that people are not capable of judging maxima in many conditions.² Instead, rationality consists of behaviour according to heuristic rules which, while not optimal, give satisfactory results in a wide range of cases.³

The consumer demand relationship might instead be considered as an heuristic rule, or as the predictable outcome of a set of heuristic rules. This would be consistent with modern cognitive science and with a number of proposals for a more cognitively realistic conception of economic rationality. Among the many questions that then arise are: (a) How do these heuristic rules arise and what determines them? If they are not “hardwired” into the human genetic constitution, then their origin and characteristics should themselves be the subject of investigation; (b) Can we characterize with any generality the heuristic decision processes that consumers use? and (c) Is it possible that these heuristic processes lead, in some cases, to approximately optimal results? If so, how good is the approximation, and what determines its quality from case to case? Of course, if it should prove that the answer to questions (b) and (c) is “yes”, and that the approximation is almost always quite good, we could then return to the neo-classical optimality theory without further concern about the specific cognitive processes involved.

One obvious possibility is that the heuristic rules “evolve” in some sense (Nelson and Winter 1982). However, “evolution” is not a uniquely specified process. A recent development in computer science, genetic algorithms (Holland 1975, Goldberg 1989, Koza 1992) provides a model of evolution that includes an analogy with sexual reproduction.⁴

This chapter reports on computer simulations of the determination of consumer demand by such heuristic processes. The chapter begins with a relatively conventional demand relationship and then proceeds to apply these processes to some more

complicated models of demand, including cultivated tastes. In each case the question is whether, and to what extent, the simulated agents learn approximately optimal behaviour.

Some simulations of simple cases

The first set of simulations were based on a consumer demand model of textbook simplicity. There are only two goods: a numeraire⁵ good and the good whose demand we investigate. The consumer has a given income to spend on the two goods and aims to maximize utility within one period. If “cognitively realistic” processes cannot find their way to an optimum in a case of this kind, it is hardly plausible that they would do so in the far more complex and ambiguous environment with which real consumers must cope.

Some results have been published elsewhere (McCain, forthcoming) and will be summarized very briefly here. Suppose that people choose a quantity to consume and, from time to time, experiment with alternative quantities, continuing to buy the experimental quantities only if, in the experiment, the different quantity leads to a higher utility. In a simulation based on this hypothesis, demand rapidly converges to the optimal quantity when the price remains stationary. However, when the price changes—either as a stationary random variate or as a random walk⁶—the average individual (aggregate) demand does not respond well to price changes. In such a case, econometric estimates of the demand relationship would underestimate the elasticity of the optimal demand curve (which could lead to bias in cost-benefit estimates) and changes in the structure of demand would be observed where there have been none.

Genetic algorithms

The simulations reported in the previous section were based on models that share a common shortcoming which is the converse of the criticism of neo-classical economics: they assume cognitive processes that are more rudimentary than real human cognitive processes. The truth presumably must be somewhere between these models and neo-classical maximization.

Suppose that demand rules “evolve” (Nelson & Winter 1982, McCain 1992b, forthcoming). That does not specify a single sort of model for theoretical work or simulation. Evolution is a process of random change and selection. Then what varies, and in what sense does it vary randomly? And how does selection come about? These are by no means trivial questions and it may be that seemingly arbitrary answers to them determine the outcomes of a simulation.⁷ A recent literature that shows some promise in this connection is organized around the idea of genetic algorithms. A genetic algorithm is an algorithm based on random variation and selection in which the random variation in the main takes the form of random recombination⁸ of given elements, rather than mutation. A large literature now exists on the use of such algorithms to find numerical maxima, fit curves, and evolve highly effective rules of procedure (Holland 1975, Goldberg 1989, Rawlins 1991).

The routines in the simulations reported here are demand rules; that is, rules that make the quantity purchased a function of price.⁹ Both the mathematical specification and the

parameters of the demand rules are coded as inputs to a genetic algorithm. The details of this coding are reserved for the Appendix on page 138. The population is “seeded” with demand rules with randomly determined parameters. Selection to “reproduce” is based on current-period utility (rather than, for example, average utility over a longer period), when the rule is applied in the context of a price that is a stationary random variable.

Thus, the demand routines of a simulated individual consumer are modelled by a string of about 160–170 zeros and ones, the consumer’s “chromosome” in the genetic algorithm. On each step of the simulation, the population (of demand routines) is reconstructed as follows. First, two simulated consumers’ demand routines are selected to “reproduce”. The mates are selected by a pseudo-random routine that assigns a higher probability to the selection of routines that yield a higher utility. The “chromosomes” are then split at a randomly selected point and the first portion of one linked to the last portion of the other, and vice versa.¹⁰ (This simulates the biological phenomenon of “crossover” in sexual reproduction, according to some accounts of the rationale of genetic algorithms.) This is repeated until there are 100 members of the new population. Before this recombination takes place, mutations occur at a 2 per cent rate¹¹; that is, two of the 100 strings are chosen at random and one of the zeros is changed to a one, or a one to a zero, at random.

Since these programs are meant to simulate the behaviour of consumers, we should be explicit as to the behaviour they are meant to simulate. It is emulation.¹² The idea is that consumers reconsider their demand routines periodically and adopt—in part—the routines of other consumers who have had greater success in matching their purchases to their needs. Thus the more successful routines would spread through the human population by a process of imperfect emulation.

The question, then, is this: how effective can such a process be in approximating optimal demand? The genetic algorithm seeks the optimum quantity demanded fairly effectively. Figure 7.1 gives an indication: it plots the simulated quantity demanded for 100 steps of a representative simulation run in which price is constant and the optimal quantity demanded is always 980.¹³

The performance is not especially impressive by comparison with simpler heuristics (McCain, forthcoming); but one would think that the heuristic demand process would have the advantage when price, and thus optimal quantity, varies, since in a case of that kind the heuristic rule would be capable of responsive behaviour.




Figure 7.1 Inverse demand curve for price fixed at 10. Optimal demand is 980.4.




Figure 7.2 Approach of simulated demand to optimal demand with inverse specification.




Figure 7.3 Approach to optimal demand with linear specification.

In fact, the performance of the genetic algorithm with a variable price can be very effective. Figures 7.2 and 7.3 show such cases. In each of them, the optimal demand curve is shown in grey, while the paths of the simulated price and quantity are shown in black.

Price is on the vertical axis and quantity demanded is on the horizontal axis, as usual for a demand relation. In the simulation reported in Figure 7.2, the optimal demand relation is an inverse function, and the simulation converges very rapidly to a consensus on that optimal form.

However, this success proved to be an artefact of the initialization conditions—in other words, a fortuitous mistake.¹⁴ When the optimal demand curve was linear, the simulation tended to choose an inverse specification in any case. When the simulation was forced to choose a linear specification, and when that specification was carefully initialized, results like those in Figure 7.3 could be attained. As we can see, the simulation does not get the slope quite right, but it moves quite rapidly to a roughly correct demand relationship. However, even when it chose the wrong specification, the genetic algorithm gave rise predictably to a very good numerical approximation to an optimal demand relationship.

To the extent that these simulations are representative of the evolutionary processes that give rise to consumer demand functions, what tentative conclusions may we draw for the rationality of consumer demand? We see that the simulations can converge rapidly to a very good approximation of the demand function. Emulation is a powerful aid in the generation of demand rules that are both responsive to price and reliable in coming close to optimal quantity demanded at the given price.

Teleological conservatism

In earlier work, McCain (1991a, forthcoming) explored similar models in which choices are determined by “simple groping”, that is, a process of random trial and error in which trials that lead to deterioration of the results are abandoned. The genetic algorithm, however, is not a generalization of groping. The simple groping model has a property the genetic algorithm lacks that makes a difference to the process of selection, a property we might call *teleological conservatism*. In a genetic algorithm (and in evolution) reproduction is, as it were, a pure lottery. The parent organisms, or code, are lost.¹⁵ The filial organisms, or codes, are selected from more fit parents, but because the code is recombined randomly, there is always some probability that the filial organism, or code, will be less fit than the parents. In a groping process, by contrast, the performance of the recombined heuristic is compared with that of the “parent” heuristics, and if the filial heuristic is less effective, it is discarded and the parent heuristics kept. Thus groping selection is not a lottery, but an option. To put it in more biological terms, it is as if the parent, having engendered a child, could reject the child if it is less fit and replace the child with a clone of the parent itself. Of course, evolution by mutation and sexual reproduction and selection of the “fittest” also lacks the property, as we might suspect, since the genetic algorithm is meant to model the evolutionary process. Thus the property of teleological conservatism serves to distinguish groping from evolution.

An advantage of genetic algorithms is that they can find the neighbourhood of the global optimum in many optimization problems in which there are multiple local optima (Rawlins 1991). It is not clear that a groping process with teleological conservatism would have this capacity. Biological evolution has given rise to manifold creativity in natural history, but a groping process may not be able to do the same.

However, our purpose here is not to model evolution or to maximize a function but to model the determination of demand, when people are imperfectly rational but are capable of learning by emulation of one another. The process is as follows. Each person determines his or her own demand by means of some heuristic routines that would cut back consumption when prices rise. A person then observes the success of some other consumer, and selects a relatively successful one to emulate. Success is measured in terms of utility achieved, but since utility is imperfectly perceived, the person chosen for emulation may not be the single most successful one. The emulation of the chosen model’s routines is also imperfect, and takes the form of random recombinations of the model’s and the imitator’s own demand routines. These recombinations form a new “population” of demand routines, which replace the old ones only to the extent that the replacement increases the utility of the imitator and which are rejected otherwise. If they could not be rejected—if the new population of routines were to replace the old, regardless—then we would have a genetic algorithm.

In what follows, genetic algorithms (without the property of teleological conservatism) will be called standard GAs, or just GAs, and modified genetic-like algorithms with the property of teleological conservatism will be called heuristic groping or teleologically conservative algorithms (TCAs).

We apply this procedure first to a conventional demand example. The results will be summarized very briefly. It appears that in these models, either GAs or TCAs can give rise to approximately optimal demand relations, predictably and reliably. One might

suppose that teleological conservatism could be helpful when individuals have differing utility functions. The danger of a standard GA in such a case is that the individual would emulate (borrow code for an heuristic routine) from another individual whose utility function is quite different from her/his own. The success of the individual emulated would not signal that the same routines would be successful for the consumer who emulates. To investigate this possibility, simulations were run based on linear demand, with one modification: utility functions are allowed to differ among the simulated agents. Again, the results will be summarized very briefly. Here, as before, the aggregate market behaviour of the simulated consumers yielded an average demand relation of the conventional type, and in the rough neighbourhood of the optimal average demand, after some delay as optimal demand patterns were “learned”. This was true both of the standard genetic algorithms and of the TCAs. However, the TCAs gave rise to a better average performance.

Learning-by-consuming

In a number of papers written in the later 1970s and early to middle 1980s (McCain 1979, 1981, 1982, 1986), I explored the implications of the cultivation of taste in otherwise fairly conventional models of the demand for works of art. All of those models assumed that consumers act as if they maximize a utility function, although in most of them ¹⁶ consumers are assumed to be short-sighted. A key prediction of my earlier, as-if-maximization models is that the demand for works of art is multi-modal.¹⁷ We shall see that this prediction extends also to some of the “cognitively realistic” models.

We shall model the demand for art. In the terms of my 1979 and subsequent papers, the hypothesis of cultivation of taste is expressed in the following way: the consumer’s utility depends on the flow of a subjective quantity we might call “enjoyment-of-art”. The subjective quantity will be denoted by Q and the utility function used in the simulations will be

$$U = \gamma + 5Q - 0.01Q^2$$

where γ represents spending on all other goods and services.

In turn, the subjective flow of enjoyment-of art depends on the objective quantity of art consumed (Stigler and Becker 1977). Greater experience leads to greater sensitivity to the benefits of art and thus to greater consumption and enjoyment. In economic terminology, this experience is a kind of immaterial capital, which we might call taste capital, and the learning process is the accumulation of taste capital. In this study the objective quantity of art consumed will be denoted by X and the sensitivity by β , and the dependency of the subjective flow on the objective flow is

$$Q = \beta X - 0.2$$

Thus β is a measure of taste capital. If the consumer were to maximize utility, the optimal demand relation (written as a demand price function) would be

$$p = 5.004\beta - 0.02\beta^2 X$$

In these simulations the taste parameter β is determined recursively by

$$\beta_t = 0.9\beta_{t-1} + 0.1X_{t-1} + 0.002$$

This might be called a “learning-by-consuming” model by analogy with “learning-by-doing” in the economics of technology and production (e.g. Arrow 1962).

In simulations of learning-by-consuming based on simple groping, the average quantity demanded and utility attained differ systematically from one simulation to the next (McCain, forthcoming). This is because random factors initially and during the groping process lead a smaller or larger proportion of the population to develop cultivated tastes. We are particularly concerned with the prediction that cultivation of taste would lead to multi-modality of demand. These simulations show the bimodal distribution of cultivation of taste previously predicted on the basis of catastrophe theory (McCain 1979).

When learning-by-consuming is modelled in a simulation based on a standard GA, the results are far less positive (McCain 1992b). Multi-modal outcomes are *not* observed. Depending on the details of the simulation, these simulations could eliminate rapidly any taste capital randomly assigned in the initialization process and converge rapidly to very low rates of consumption and taste capital, or converge equally rapidly to universal cultivation. Nor was optimal demand well approximated. Figure 7.4 shows a simulation that led to universal cultivation.

In this simulation, the highest utilities were emulated, so that those who were randomly assigned a high level of cultivation at the outset were more likely to be emulated. The grey line shows the optimal demand for a culti-




Figure 7.4 Approach to optimal demand with linear underlying specification.

vated consumer. We see that the simulation does not settle down to a systematic demand relationship, as if it were “confused” by learning-by-consuming.

The difficulty seems to be that cultivation of taste results in idiosyncratic individual differences in demand, and GAs do not handle that kind of individual difference well, as we observed in a previous section. Accordingly, let us consider some simulations of learning-by-consuming in terms of TCAs. Once again we ask: does the population tend to develop dichotomous groups of cultivated and uncultivated individuals?

The answer is that, in these simulations with modified GAs, it does. This is illustrated by Figure 7.5, which shows the distribution of β at the end of five simulation runs.




Figure 7.5 Dichotomization of the population in five simulation runs with heuristic groping, cultivation of taste, and the price treated as a stationary random variate.

In these simulations, the initial value of β was a normal pseudo-random number with mean 5 and standard deviation 3, subject to truncation at zero. Thus, at the outset, about half of the population had “sensitivity” at or above the equilibrium level, with a single mode at that level. We see in Figure 7.5 that in each case the population has become dichotomized, with almost all the population having a value of β either at zero or in the neighbourhood of 5. Despite the fact that, at the beginning, about half of the populations had a β of 5 or more, at the end more than half (from 65 per cent to nearly 90 per cent) of the population were entirely uncultivated.

Figure 7.6 shows the dichotomization in terms of utility. Notice that the five simulation runs are pattern-coded consistently: in each case the black




Figure 7.6 Dichotomization of the population in five simulation runs with heuristic groping, cultivation of taste, and the price treated as a stationary random variate. Note that the vertical axis begins with 20,000.

line stands for the first simulation and the grey line for the last and so on. Notice also that in Figure 7.6, the population is sorted as in Figure 7.5, i.e. so that the end-of-simulation β is non-decreasing. We note that utility is almost non-decreasing: the more cultivated tastes are, with very few exceptions, the better-off the consumers. (The exceptions are presumably consumers who are less skilful at approximating their optimal consumption level than are their neighbours with the same or lower β . This is possible because of bounded rationality.)

Finally, we might ask whether the simulation leads to something approximating optimal demand. Since the consumers now have different utility functions, because of their different degrees of cultivation of taste, we cannot simply compare the average consumption to a uniform optimum.¹⁸ For the uncultivated, zero is optimal consumption and most of them attain it. Figure 7.7 shows actual and optimal consumption for the cultivated group in a representative simulation. (Specifically it is the penultimate simulation, shown by the diagonally-dashed line. Thus the cultivated group in this case are about 20 per cent).

As usual, actual price-quantity movements are shown by black lines and optimal by the grey one. In this case, only the second hundred iterations are shown, so the simulation is essentially in equilibrium from beginning to end. In this case, the performance is comparable with that for the heuristic groping simulations of demand without learning-by-consuming, and is what we observe in general.

Thus we may say in summary that when consumer decisions with learn-




Figure 7.7 Approximately optimal and simulated demand for those with cultivation parameter of 3 or greater.

ing-by-consuming are simulated with a model of heuristic groping, the teleologically conservative counterpart of a genetic algorithm, the population tends to become dichotomized into two groups, one less cultivated and with lower utility at the same price, and the other more cultivated and with higher utility. Each group tends to evolve demand routines that approximate their respective optima, based on the sensitivity they have, as would be the case in the absence of learning-by-consuming. It appears that dichotomization is a consequence of teleological conservatism, rather than of learning-by-consuming *per se*. In any case, this dichotomization is an instance of an emergent phenomenon in the economics of demand.

Summary and conclusions

This chapter has considered computer simulations of the determination of consumer demand, focusing mainly on a case with learning-by-consuming and on two sorts of heuristic process: genetic algorithms and groping over heuristics (teleologically conservative algorithms). Genetic algorithms involve slightly less cognitive capacity than groping over heuristics, but the two processes differ mainly in details, in that groping over heuristics is teleologically conservative while standard GAs are not. The genetic algorithms produce an optimal demand relationship when the sensitivity coefficient is the same for all consumers, but is not adaptive to differences in utility functions among consumers and thus is less efficient than groping over heuristics in cases in which consumers differ, either intrinsically or as a consequence of cultivation of taste and different histories. Groping over heuristics is highly effective when individual differences among consumers are intrinsic, but shares the short-sightedness of simple groping when there are opportunities for cultivation of taste.

The following conclusions are drawn: if human cognitive capacity is well approximated by genetic algorithms, and especially if they are teleologically conservative, then real demands will in many cases approximate closely optimal demands. However, there are some important exceptions, and cultivation of taste is among them.

Appendix: The genetic algorithm for consumer demand routines

In the literature on curve fitting by genetic algorithms, the mathematical form of the curve to be fit is generally given a priori and only the parameters are “evolved”. In this study, however, the functional form is also allowed to evolve, with some qualifications that will be mentioned.

The computer programs used in this study simulate a population of 100 consumers. Each consumer is assumed to determine her/his demands in each period by means of routines expressible by a function from price and income to quantity demanded: the demand function. Each simulated consumer’s demand-determining routines are coded by a string of zeros and ones. Five numerical constants are coded in binary notation, with the sign indicated by the leading digit (zero if negative). Each of these binary-coded constants is expressed by a string of about thirty zeros and ones. The five binary strings are concatenated to a string of about 150 zeros and ones. A string of 15 zeros and ones is concatenated to the beginning of this string and this codes for the specification of the demand function. This 15-digit coding string works as follows:

1. The sixth through tenth digits code for transformations in the first variable:
 - (a) if the sixth digit (first of the five) is 0, there are no transformations;
 - (b) if the second digit of the five is 1, the number is transformed to logs, and no further transformations are made;
 - (c) if the third digit is 1, the variable is raised to the power of a corresponding constant, and there are no further transformations;
 - (d) if the fourth digit is 1, the variable is transformed to an inverse function, and there are no further transformations; and
 - (e) If the fifth digit is 1, the variable is exponentiated.
2. The eleventh through fifteenth digits code for the transformations of the second variable, in a similar way.
3. The first five digits code the way in which the variables (transformed or not, as the case may be) are combined, as follows:
 - (a) if the first digit is 1, there is a constant term (the first constant of the five);
 - (b) if the second digit is 1, the variables are combined additively;
 - (c) if the third digit is 1, the variables are combined multiplicatively (and these last two are, of course, exclusive; addition takes precedence and the third digit is ignored if the second is 1);
 - (d) if the fourth digit is 1, the result is transformed to its logarithm; and
 - (e) if the fifth digit is 1, the result is exponentiated. Again, the last two are exclusive.

Thus there are two ways to code for a Cobb-Douglas function, i.e. the product of power functions of the independent variables, a function often used in mathematical economics: 110011100011000 and 101001010010100.

Acknowledgements

I am indebted to participants in the Seventh International Conference on Cultural Economics, Fort Worth, Texas, October 1992; the Seminar on the Economics of the Arts, Venice, Italy, December 1992; and the Symposium on Simulating Societies, Siena, Italy, July 1993; for useful comments on some of the research reported in this chapter. Errors and omissions are, of course, attributable to the author.

Notes

1. It should be observed that this is a statement about the individual demand curve, not the aggregate demand curve. Economic theory has sometimes posited the hypothesis that, if the individual demand relationships are derived in this way, aggregate demand relationships will have certain characteristics that are sufficient for a market equilibrium to exist. This may be called the “representative man” hypothesis in demand theory. However, it is now clear that utility-maximizing demand relationships are neither necessary nor sufficient for the representative man hypothesis, and thus for the existence of market equilibrium. Nevertheless, it is basic to many applications of the demand concept, for example, in cost-benefit analysis, and conversely, evidence of responsiveness to price-like variables is often offered in the economic literature as evidence of the rationality of economic action. These facts justify a continued interest in the rationality of individual demand relationships. Moreover, issues with respect to the aggregation of demand relationships are raised, not resolved, by the finding that utility-maximization-derived individual demand functions are neither necessary nor sufficient for the representative man hypothesis. Aggregate demand relationships as described in the representative man hypothesis might emerge in any case, not as necessary but as contingent outcomes of individual decisions, at least in some circumstances. This would constitute emergence, and is precisely the issue to which this chapter is directed. In order to address it, it is necessary to avoid both of the contradictory a priori hypotheses about demand relationships common in the social science literature: the neo-classical hypothesis of utility maximization and the hypothesis that, since people are “obviously” not rational, demand decisions have nothing to do with price and utility maximization considerations. The truth may be in between the two, or partake of both hypotheses, and adopting either would then blind us to the important issue of when one or the other is reliable. It should also be noted that this chapter is limited to the preliminary issue of the rationality and aggregateability of demand, and does not address the existence or characteristics of market equilibrium, issues which are confronted by the work of Chattoe (1993).
2. For example, Camerer (1987), Kahneman and Tversky (1984), and Kahneman, Slovic, Tversky (1982).
3. As Herbert Simon (1981) has long argued. Note also Reddaway (1936) for a very early example. For a survey see McCain (1992a).
4. By contrast, such earlier studies as Nelson & Winter (1982) allowed only “mutation”.
5. That is, a good in terms of which prices and incomes are measured. For purposes of this discussion, “money is a veil” and plays no part in the model.

6. These are statistical processes commonly assumed, and often observed, in econometric work. Constant prices are of limited interest to us, since a major concern of this chapter is with the responsiveness of the simulated agents to changing prices. Thus, constant price assumptions play only a comparative role.
7. Nelson & Winter (1982), for example, find that they have to make use of the production function as part of the specification of their simulation of evolution in firms. This seems to be something of a failure—what they give us is not a model of evolution but of adaptation.
8. I cannot resist the temptation to observe that in the *Theory of economic development*, Schumpeter habitually refers to innovations as “new combinations”.
9. This, of course, imposes some “unrealism” by identifying the heuristic with the demand relationship itself. There would be merits in a more realistic heuristic, and future work should address that. However, there are pitfalls in such a research strategy. Suppose, for example, that a model were to identify a budget of total spending on a commodity as the demand heuristic that is to evolve. This might be “more realistic”, but it entails a specific demand price—quantity relationship, and one that is highly responsive to price changes—the implied relationship would have an “elasticity” of exactly -1. The assumptions made here are in fact less restrictive and allow for unresponsive demand heuristics, the very possibility we wish to confront.
10. One implication of this is that the functional form is unlikely to be changed by recombination, since the probability that the split takes place in the first 15 of 160 or so places is small. Thus it may be more accurate to say that the functional form is “selected” in this study than that it “evolves”. One way to force the functional form to “evolve” would be to let the first 15 positions play the role of a second “chromosome”, and have it recombined on every round of the simulation. Multiple-chromosome models have been proposed but have not been explored extensively in the genetic algorithms literature.
11. This rather high rate was chosen without any particular rationale, but some experiments with alternative rates suggested that the rate of mutation made little difference. For present purposes, mutation is a cognitive phenomenon, representing incorrect memory, perhaps. Thus it is not clear what rates might be “realistic”, but we may say that, in a certain sense, higher rates correspond to lower cognitive capacity.
12. For a very deep and useful discussion of the rationality of imitation or emulation, see Pingle (1992). Pingle’s paper, excellent as it is, is somewhat compromised by his adherence to the neo-classical conception of rationality, a conception which is also used in this chapter. As I pointed out in my unpublished comments on Pingle’s paper, his ideas would fit more naturally in a relative concept of rationality, such as my “linguistic” conception (McCain, 1991b, 1991c, 1992a). A referee observes that people imitate behaviour, not heuristic routines. Perhaps; and there might be much interest in a study in which individuals model the behaviour of others and consider emulating the behaviour they model (Johnson-Laird 1983). That is beyond the scope of this chapter. In any case, it does not seem clearly wrong to suppose that people have a repertoire of heuristic routines which are a common social heritage that they “recombine” in various ways as their heuristics evolve. There being some connection between heuristics and behaviour—that, of course, is the whole point of heuristics!—this would lead to imperfect emulation of heuristics, which is what the chapter assumes.
13. The simulation does not converge to as near this optimum as do the models of simple groping (McCain 1994) and remains below it. It has been noted in other studies of genetic algorithms that their convergence tends to be incomplete. Due to initialization conditions, the approach to the optimal quantity is from below. It may be that these considerations together explain the tendency of the quantity demanded to remain below the optimal.
14. That is, the linear form requires that the constant and the coefficient be of very different orders of magnitude, while the inverse specification did not. Thus, when the two constants were initially random draws from the same pool, the inverse specifications tended to perform

best in early rounds of the simulation, even when the true specification was linear. This experience indicates that some functional forms are much more sensitive to initialization conditions than others and thus may gain an early advantage in an evolutionary process that prevents the correct specification from emerging. A different selective process, designed to retain diversity in early rounds, might perform better.

15. I should remark that in practice, applied genetic algorithms have often included some element of teleological conservatism. To that extent, however, they depart from the evolutionary metaphor. This issue seems more important for simulation than for applications to function maximization—in the latter case, the correct assumptions are the ones that work; in simulation, the case is much less clear and needs explicit discussion.
16. The exception is McCain (1981).
17. This is testable econometrically. See McCain (1992a). I note in passing the econometric support for a catastrophic long-run demand model for wine in my 1979 paper, which I believe to be still the only implementation of catastrophe theory in econometrics.
18. In other words, this outcome presents an “aggregation problem” in which average demand is not the demand of a representative consumer. Indeed, in a multimodally distributed population, there can be no “representative consumer”.

Part II
**The evolution and
emergence of societies**

Chapter 8

Emergence in social simulation

Nigel Gilbert

Every discipline is based on a unique foundation of epistemological assumptions and concepts. This means that even when one discipline develops so that it begins to share its concerns with another, there may be little or no contact because the practitioners are, literally, speaking different languages. Even if contact is established, the neighbouring disciplines may still have nothing to say to each other because, while a topic may be common, the questions being asked and what count as interesting answers differ so greatly. To bridge the gap, considerable work in translation has to be done.

This situation seems to be developing in the case of artificial intelligence (AI) and the much older discipline of sociology. Both are interested ostensibly in the behaviour, structure and properties of collections of actors or agents. Yet the literature on distributed artificial intelligence (DAI) makes little or no reference to sociology, or even to its major concepts. Similarly, there are very few sociologists who have considered the relevance of DAI to their concerns.

In this chapter, I shall show first that sociologists have also struggled with one of the basic conceptual problems that those interested in simulating societies have encountered: the problem of understanding “emergence” and, especially, the relationship between the local and global properties of complex systems. Secondly, I shall indicate some ways in which DAI simulations may have oversimplified important characteristics of specifically human societies, because the actors (agents) in these societies are capable of reasoning, and do so routinely, about the emergent properties of their own societies. This adds a degree of reflexivity to action which is not present (for the most part) in societies made up of simpler agents, and in particular is not a feature of current DAI simulations.

The macro and the micro level

The issues I shall address stem from the fact that both societies and the computational systems with which distributed AI are concerned are composed of many interacting agents (also known as people, actors or members). These systems can therefore be described either in terms of the properties and behaviour of the agents, or in terms of the system as a whole. The former mode of description focuses on the “micro” level, that is, the features of individual agents and their local environment (which they can perceive directly), while the latter focuses on the “macro” level, that is, the global patterns or regularities formed by the behaviour of the agents taken as a whole. The problem is how to characterize the relationship between these two types of description, in particular when the macro properties can be said to “emerge” from the micro-level behaviour.

Sociological approaches to emergent properties

Methodological holism

The relationship between local and global properties is one which has exercised sociologists since the foundation of the discipline. For example, Emile Durkheim, one of its founding fathers, emphasized the external nature of social institutions (a macro-property of a society) and argued that they imposed themselves on individuals at the micro-level. It is worth quoting at some length from his methodological writings, because he was addressing issues very similar to those which continue to worry people who construct social simulations. After having defined a category of facts that he called “social facts”, he wrote

Another proposition has been...vigorously disputed...it is the one that states that social phenomena are external to individuals ... The states of the collective consciousness are of a different nature from the states of the individual consciousness: they are representations of another kind. The mentality of groups is not the mentality of individuals; it has its own laws... To understand the way in which a society conceives of itself and the world that surrounds it, we must consider the nature of society, not the nature of the individuals. (Durkheim 1895:65–6)

Thus, for Durkheim, there are social representations which can be examined independently of the individuals that make up the society. These social facts are

a category of facts with very special characteristics: they consist of ways of acting, thinking and feeling that are external to the individual and are endowed with a coercive power by virtue of which they exercise control over him. (Durkheim 1895:69)

This view of the relationship between an individual and society was later developed by Parsons (1952) into a “normative conception of order” which emphasized the role of internalized norms in ensuring the integration of society through shared values and obligations. Subsequent theorists have extended this tradition, which has become known as “methodological holism” (O’Neill 1973). It asserts that the social behaviour of individuals should be explained in terms of the positions or functions of these individuals within a social system and the laws which govern it.

Methodological individualism

In contrast, “methodological individualists” see macro-phenomena accounted for by the micro-level properties and behaviour of individuals. The individualists’ position demands that all the concepts in social theory are analyzable in terms of the interests, activities, etc. of individual members of society.

If social events such as inflation, political revolution, “the disappearance of the middle classes”, etc. are brought about by people, then they must be explained in terms of people; in terms of the situations people confront and the ambitions, fears and ideas which activate them. In short, large scale-social phenomena must be accounted for by the situations, dispositions and beliefs of individuals (Watkins 1955:58).

A rather sterile debate between these two camps continued for much of the 1970s and 1980s. With the benefit of hindsight, it is now possible to argue that while there was some truth in both, neither was a particularly helpful or revealing way of conceiving the relationship between global and local behaviour. It is the case, however, that most, if not all, current simulations of human societies in essence adopt one or other of these positions, often without making this explicit.

Structuration theory

Perhaps the most influential recent sociological approach to the relationship between macro and micro is the theory of structuration. The theory argues that there is a duality between society and knowledgeable human agents. Agents are seen as reproducing in action the structuring properties of society, thereby allowing social life to be reproduced over time-space (Giddens 1984). This rather dense, and possibly even opaque, statement of the theory needs some considerable explanation.

First, by “structuring properties”, Giddens means things such as institutional practices, rules and resources. These properties are the means by which social structure or “society” is produced and reproduced. Human action has these structuring properties “embedded” in it, so, as people act, they contribute to the reproduction of society. Also (“reflexively” according to Giddens), human action is both constrained and enabled by social structure, for this is the medium through which action is performed. Structure is at the same time both the outcome of knowledgeable human conduct and the medium that influences how conduct occurs (Giddens 1984:25).

Complex adaptive systems

We can now consider whether structuration theory can be linked to or has anything to offer computational views of emergence in simulated societies. The question of emergence has probably been considered most systematically by those interested in complex systems (Stein 1989, Jen 1990, Stein & Nadel 1991). Such systems have the following general characteristics (Forrest 1990):

1. The system consists of a large number of interacting agents operating within an environment. Agents act on and are influenced by their local environment.
2. There is no global control over the system. All agents are only able to influence other agents locally
3. Each agent is driven by simple mechanisms, typically condition-action rules, where the conditions are sensitive to the local environment. Usually, all agents share the same set

of rules, although because they may be in different local environments, the actions they take will differ.

A typical example of such a complex system is an array of cellular automata. Each cell can take on one of a number of internal states. State changes occur according to the results of the firing of simple condition-action rules dependent only on the states of neighbouring cells.

Some complex systems are also adaptive. In such systems, agents either learn or are subjected to a process of mutation and competitive selection, or both. In learning systems, agents are able to modify their rules according to their previous “success” in reacting to the environment. In systems involving competitive selection, agents’ rules are altered, either at random or by using some specific learning algorithm, and then those which are more successful are copied while the less successful are deleted from the system. An important characteristic is that agents adapt within an environment in which other similar agents are also adapting, so that changes in one agent may have consequences for the environment and thus the success of other agents. The process of adaptation in which many agents try to adapt simultaneously to one another has been called “co-evolution” (Kauffman 1988).

Emergence in complex systems

Because complex systems, whether adaptive or not, consist of many agents, their behaviour can be described either in terms of the actions of the individual agents or at the level of the system as a whole. In some system states the global description may be very simple (e.g. if the agents are either not interacting or interacting in repetitive cycles, the global description might be that “nothing is happening”), or exceedingly complex (e.g. if the agents are in complete disequilibrium). In some circumstances, however, it may be possible to discover a concise description of the global state of the system. It is in these latter circumstances that it becomes possible to talk about the “emergence” of regularities at the global level.

An example relevant to social simulation can be found in the work of Nowak and Latané (e.g. 1993) which aimed to model the effect of social influence. This draws on Latané’s theory of social impact which states that the impact of a group of people on an individual’s opinion is a multiplicative function of the persuasiveness of the members of the group, their social distance from the individual and the absolute number of the members. For example, in a simulated world in which there are only two mutually exclusive opinions, the opinion adopted by an individual is determined by the relative impacts on the individual of those propounding the two opinions. At any moment in time, an agent’s opinion is determined by the multiplicative rule that implements the theory of social impact, and is depend-ent only on the opinions of other agents. The simulation fits the definition of a complex system.

The outcome of Nowak and Latané’s simulations is that coherent clusters of opinion emerge and remain in dynamic equilibrium over a wide range of assumptions for the parameters of the model. The population does not move wholly to adopt one opinion, and minority views are not completely expunged. Nowak and Latané show that the clusters of individuals with the same opinion can be visualized as patches on a two-dimensional grid

in which the nodes represent the members of the population. These clusters can be said to have emerged and to form regularities.

In systems in which there is emergent behaviour, it is convenient to think of the emergent properties of the system as influencing the actions of the agents. Thus, not only do the agents' actions at the local level, when aggregated and observed at the global level, constitute the emergent behaviour, but also the global emergent behaviour can also be said to influence the local actions of the agents, in a form of feedback. For example, in Nowak and Latané's simulation, one can speak of an opinion cluster influencing the opinions of individuals outside the cluster. But it must be remembered that in the standard complex system, there is no explicit mechanism for such feedback to take place: agents are affected only by their local neighbours. The idea of feedback, or "feed-down", from global to local level is merely a convenient shorthand.

Levels of emergence

So far, for simplicity, it has been assumed that there are clearly distinct "levels": the micro and the macro. In fact, the situation is more complex than this. For example, it may be the case that individual identity is best regarded as an emergent phenomenon, where the micro-level "agents" are sub-cognitive, such as neurons. And societies are perhaps best considered as emergent phenomena arising from the interaction of social institutions. Thus it is better to consider a complex hierarchy of levels of emergence, rather than a straightforward division between micro and macro.

A continuing worry of those who are interested in emergent systems, and particularly of those who are simulating social process, is whether the emergent behaviours they observe are in some sense "programmed in" and an inevitable consequence of the way in which agents are constructed. Some criterion is required which will distinguish emergent behaviour from behaviour which is predictable from the individual characteristics of the agents. The description of complex systems suggests that a candidate criterion is that it should not be possible to derive analytically the global emergent behaviour solely from consideration of the properties of agents. In other words, emergent behaviour is that which cannot be predicted from knowledge of the properties of the agents, except as a result of simulation. This kind of definition, of course, provides an obvious explanation for the popularity of simulation as an investigative technique in the field of complex adaptive systems.

This criterion, however, is not sufficient to end the worries of those who are concerned about whether they have *really* achieved emergence. For it is always possible that at some future time analytical methods will be developed that can be used to derive global properties from local ones. In a few cases, this seems to have happened. For example, Forrest and Miller (1990) define a mapping between classifier systems (i.e. systems consisting of nodes that use the genetic algorithm to synthesise patterns of connections that collectively co-evolve to perform some function) and Boolean networks whose properties are, relative to classifier systems, much better understood. This mapping opens the way to predicting the emergent properties of classifier systems where previously predictions were impossible. The example shows that if we define emergence in terms of an inability to find an analytical solution, any particular emergent property stands the risk of being demoted from the status of emergence at some time in the future. This suggests

that emergence may be neither a stable nor an especially interesting property of complex systems: what are interesting are the systems' macro properties and the relationship of those macro properties to the micro ones.

Complex adaptive systems and human societies

The above account of complex adaptive systems is generic, in the sense that it is intended to apply equally to many disparate domains. For example, there has been work applying these ideas to the evolution of antibodies, stock market crashes, ecological dynamics, computer networks and so on (Waldrop 1992). However, I am particularly concerned with understanding human societies. Can these accounts be applied unchanged to human societies, or is there something special about these societies which would indicate that such models are completely inappropriate, or must be amended or developed in order to provide insights into their structure and functioning?

Debating whether there is a boundary between humans and other animals and, if so, its nature, is a matter of some controversy, for example with the argument that great apes deserve the same moral status as humans (Vines 1993). The question of whether human societies differ in any fundamental way from non-human societies threatens to become mired in the same moral debate. Nevertheless, it is clear that, while an exact boundary line may be difficult to draw and some characteristics of human societies are also to be found to some degree in some animal societies, human societies are significantly different from any animal society. Does this also mean that present theories of complex adaptive systems, while perhaps useful for understanding non-human societies, are lacking in some significant way as models of human societies? I believe that the answer to this question is that such theories do at present fail to model a crucial feature of human societies, and that the preceding discussion, especially the notion of structuration, can help to identify what this is and how it may be corrected.

A fundamental characteristic of humans, one that is so important that it makes the societies they form radically different from other complex systems, is that *people are routinely capable of detecting, reasoning about and acting on the macro-level properties (the emergent features) of the societies of which they form part*. The overall lesser intelligence of animals prevents this occurring in nonhuman societies. The consequences of this capability to "orientate" to macro-level properties are far-reaching.

Our discourse is permeated by references to the macro-level properties of our society. For instance, we are not only able to speak a language using a common lexicon in which words have shared¹ meaning, but we are also able to describe and comment on our lexicon; that is, as members of a human society, we can observe and reason about this macro-level property of our society.

Let us now consider another example, one where a current simulation has been examining a theory relating to human society, and which tries to provide explicitly a relatively sophisticated approach to modelling human sociality, and compare its implicit characterization of the relationship between macro and micro with sociological accounts that theorize society in terms of structuration.

Simulating human institutions

Institutions are among the most pervasive emergent features of human society. An institution is an established order comprising rule-bound and standardized behaviour patterns. Examples include the family, tribes and other collectivities, organizations, legal systems and so on. Sociological approaches to institutions can be considered in terms of the three positions described at the beginning of this paper: methodological holism, individualism and structuration theory.

In brief, the holists' position sees individuals being socialised into membership of an institution (for example, the family, or the nation) by parents and teachers. The socialization process consists of the transmittal of the rules and norms of behaviour of the institution to new members. These rules and norms then guide members in institutional action. In contrast, the individualists see institutions as arising from negotiation among the members, each pursuing his/her interests. An approach from structuration theory might focus on the way in which individual action and interaction in everyday life generates institutions through rationalizations of action, in which the institution itself becomes acknowledged by individuals as part of the process of accounting for their actions. As Bittner (1974), although himself not an advocate of structuration theory, put it with reference to a particular kind of institution, the organization: "Organisational designs are schemes of interpretation that users can invoke in as yet unknown ways whenever it suits their purposes" (Bittner 1974:77). In other words, institutions, as macro-level features, are recognized by members, who use them to warrant their own actions, thus reproducing the same features.

The EOS simulation

We can now examine an example of a simulation based on DAI principles to see whether it fits neatly into any of these theoretical perspectives on the relationship between macro and micro. I have been associated with the EOS (Emergence of Organised Society) project since its inception, although Jim Doran and Mike Palmer are the people who have done all the work (Doran et al. 1994, Doran & Palmer, Chapter 6 in this volume). The objective of the project is to investigate the growth of complexity of social institutions during the Upper Palaeolithic period in south-western France, when there was a transition from a relatively simple egalitarian hunter-gatherer society to a more complex one, with centralized decision-making, specialists and role differentiation, territoriality and ethnicity (Cohen 1985, Mellars 1985, Doran 1992). The project is intended to examine a particular theory about the causes of this transition, one which takes seriously the effect of human cognitive abilities and limitations. Consequently, the simulation that is at the heart of the EOS project is based on a computational model that incorporates a number of agents having AI capabilities that enable them to reason about and act on their (computational) environment. The agents, implemented using a production system architecture, are able to develop and store beliefs about themselves and other agents in a "social model". Doran et al. (1994:204) describe the social model thus:

The social model is particularly important. It is designed to record an agent's beliefs about the existence of groups of agents and the identities of their leaders. Associated with beliefs about leaders is other information,

for example, the size of its immediate following. Some followers may themselves be leaders. An agent may well appear in its own social model as a leader or follower in particular groups. Also held within the social model are beliefs about territories. There is no expectation that the information held in an agent's social model will be either complete or accurate.

The social model is important because it has an effect on how an agent behaves in relation to other agents. For example, an agent's membership of a semi-permanent group depends on the agent becoming aware of its membership of the group and then treating fellow members differently from non-members.

Agents start without any knowledge of groups or other agents, but with a need to obtain a continuing supply of "resources", some of which can only be secured with the co-operation of other agents. If an agent receives a request to co-operate with the plans of another agent, it will do so, or it will attempt to recruit agents to a plan of its own. Agents that successfully recruit other agents become group leaders and the other agents become followers. Such temporary groups continue to work together, with group membership recorded in the members' social models, unless other processes (such as the agent's movement out of the proximity of other members) intervene.

A number of observations can be made about this simulation. First, in terms of the definition offered previously, it is a complex system, although not an adaptive one. Secondly, group membership is an emergent property of the behaviour of the agents. Thirdly, agents pursue their own interests to obtain resources, forming "groups" as and when necessary through negotiation with other agents. These groups have no existence other than their representation in agents' social models.

In particular, the agents, even though their membership in a group is recorded in their social model, have no awareness of the group as an entity in its own right with which they can reason. Agents not in a group have no way of recognising that the group exists, or who its members are. In short, the simulation permits agents to reason about micro-properties that they perceive in their local environment, and from this reasoning and the consequent action, macro-level properties emerge, but because the agents have no means of perceiving these macro-properties, the model as a whole fails to match the capabilities of human societies.

Structuration theory begins to suggest what additional functionality could be built into the agents. It proposes that actions are sedimented into "structuring properties", such as institutional practices, rules and typifications. Thus it would be necessary for the EOS agents to *recognise* that they are members of groups and to realize what this implies about members' actions; and do so explicitly and in a way that results in a representation they can subsequently reason with. Their own actions, as group members, would reinforce ("reproduce") these explicit representations of the group and its practices.

In the present simulation, a "group" emerges from agents' actions, but it remains an implicit property of the agents' social models. As would be expected from a simulation that adopts the individualistic approach implicitly, the existence of the "group" can be observed by the researcher either by looking at the patterns of interaction at the global level or by "looking inside the heads" of the agents, but it is not visible to the agents themselves. An alternative way of implementing the simulation, rightly rejected by the

research team, would have been to build the notion of a group into the system at a global level so that the agents' actions were constrained by their group membership. In this case, which corresponds roughly to the holists' position, the group would not have been an emergent property of the simulation, but a constraining "social fact".

A third possibility is to design the simulation so that agents themselves are able to detect the emergence of groups and react ("orientate") to this structure, so that they become no less privileged in this respect than the observer. This latter course is the approach that might be advocated by structuration theorists.

Conclusion

This chapter originated from a feeling that, despite the sophistication of the present-day work, the models being used to simulate human societies are lacking in crucial aspects. Few computational models of human societies have supported features which could be argued to correspond to "meaning", "structure" or "action", although to sociologists, concepts such as these are central to understanding human society. One of the suggestions of this chapter is that more sophisticated models of human interaction are essential if we are to progress with the simulation of human societies. Human societies are not just ant hills writ large. However, pinning down clearly, and in a computationally precise way, the nature of the difference between non-human and human societies is not straightforward.

The second suggestion of this chapter is that sociological debates about the nature of human society can contribute to our understanding of these issues, even if a certain amount of effort is required to "translate" these debates into forms that are accessible to computational modelling. We have seen that one theory—structuration—provides some useful suggestions about how the relationship between macro- and micro-level properties might be conceptualized and even some broad indications about how a simulation might be implemented. It must be noted, however, that structuration theory is certainly not the only position on these issues to enjoy respect among sociologists, and that other present-day ideas may be equally, or more, fruitful.

The third suggestion is that one capacity of humans which is probably not widely shared by other animals is an ability to perceive, monitor and reason with the macro-properties of the society in which they live. It seems likely that this ability is related to, and is perhaps even a consequence of, the ability of humans to use language. Much of the lexicon of modern adults is, for example, composed of concepts about macro-properties.

The fourth and final suggestion, lurking behind the rest of the argument, is that human societies are dynamic; that is, they are (re-)produced by people as they proceed through their everyday life. Routines, typifications, structures, institutions and so on, are not reflections of some external reality but are human constructions. In particular, the "macro-properties" that are recognized by members of human societies are not fixed by the nature of the agents (as they would be in other complex systems), but are an accomplishment, the outcome of human action.

Acknowledgements

I am grateful to the EOS project team, especially Mike Palmer and Jim Doran, and the participants at SimSoc '92 and '93 for opportunities to discuss these issues.

Notes

1. I am simplifying here; meanings are indexical and negotiated in the context of use in a way which we do not yet properly understand, at least from a computational perspective. See Hutchins and Hazlehurst, Chapter 9 in this volume.

Chapter 9

How to invent a lexicon: the development of shared symbols in interaction

Edwin Hutchins and Brian Hazlehurst

Human language provides, among other things, a mechanism for distinguishing between relevant objects in the natural environment. This mechanism is made up of two components—forms and meanings—which must be shared by the community of language users. The lexicon constitutes much of a language's form, but its description as a set of shared, form-meaning resources for communication is seldom addressed by formal models of language. This chapter presents a simulation model of how shared symbols (form-meaning pairs) can emerge from the interactions of simple cognitive agents in an artificial world. The problem of creating a lexicon from scratch is solved by having cognitive agents that are capable of organizing themselves internally share their expressions of visual experience in interaction through learning to classify a world of visual phenomena. The model is seen as an instantiation of a theory of cognition which takes symbols to be a product of inter- and intra-individual organizations of behaviour, the result of cultural process.

Below, we first present a theoretical stance that has been derived (elsewhere) from empirical investigation of human cognitive phenomena. This stance leads to the articulation of some simple information-processing constraints which must hold for lexicons and their users. We present later two simulations employing communities of simple agents that allow us to model how a lexicon could come into existence or emerge from the interactions between agents in an extremely simple world of experience. The model allows us to explore issues which involve interaction between group- and individual-level properties of cognition, social organization and communication. Subsequently, we briefly review some findings from the literature on lexicons of natural language and conclude the chapter with an evaluation of the model as an instantiation of the theoretical stance.

A community of minds as a cognitive system

Recently, we have been exploring a novel approach to cognitive anthropology by trying to push the boundaries of a genuinely cognitive unit of analysis out beyond the skin of the individual. Ever since symbolic and cognitive anthropology embarked on their ideational odyssey in the 1950s they have proceeded away from the material and the social aspects of human life. Of course, many people are interested in *social cognition*, in which the social world is the content of cognition. And, in fact, there are good arguments for believing that human intelligence developed in the context of reasoning about social

situations (Byrne & Whiten 1988, Levinson in press). This kind of relationship between the social and the cognitive is important, but it is still centred on the notion of the individual as the primary unit of cognitive analysis. The social world is taken to be a set of circumstances “outside” the individual, about which the individual reasons. What we intend, instead, is to put the social and the cognitive on an equal theoretical footing by taking a *community of minds* as our unit of analysis.

This new perspective permits two things to happen that are not possible from the traditional perspective. First, it permits inquiry about the role of social organization in the cognitive architecture of the system, and allows description of the cognitive consequences of social organization at the level of the community (Hazlehurst 1991, Hutchins 1991). Second, it permits symbolic phenomena that are outside the individuals to be treated as real components of the cognitive unit of analysis (Sperber 1985, Hazlehurst 1994, Hutchins in press).

Making this move also presents an opportunity to view language in a new way. Cognitive science generally takes the existence of language as a given and concerns itself with the sorts of cognitive processes that must be involved when an individual processes language (in production and comprehension). From the perspective of the community of minds as cognitive system, a language’s information-bearing capacity, conventions of use, functionality, distribution, variability, etc., all become determinants of the cognitive properties of the community because these things are instrumental in determining where, when, and what kinds of information move through the system (cf. Freyd 1983). This attention to the movement of information in the larger system necessarily brings the material world back into play, since, having acknowledged symbols outside the head, we now must take seriously their material nature. Furthermore, we need not, indeed must not, ignore the means individuals have or develop for incorporating symbols into private as well as collective organizations of behaviour. By redefining our unit of cognitive analysis, it seems to us that progress may be made towards reuniting the social and the material with the cognitive (cf. Goody 1977).

A model of the emergence of shared language

The existence of shared language is one of the central facts of human existence. Language appears to be closely tied to most high-level cognitive activities. It mediates most of the interactions among members of the most social of all species. Once a language exists, it is not difficult to think of the means by which it could be maintained and propagated from generation to generation in a population. But without anyone to tell individuals which language to speak, how could a language ever arise? How could something structured come from that which is unstructured?

There is, of course, a vast speculative literature on the origins of language which we will not attempt to treat here. Rather, this study focuses more modestly on the development of sets of local lexical distinctions which may arise in small groups of cognitive agents engaged in shared tasks. In this chapter we outline a scheme by which shared denotational resources, which we shall call *symbols*, arise in the interactions between the members of such a community. This is certainly not a model of the development of a human language, but it does demonstrate how simple shared structures

can arise where none existed before. These structures are products of a system whose organizing dynamics are an interplay between intra-individual and inter-individual coordinations.

In the presentation, these structures will be referred to as *terms*, *descriptions*, *words*, or *patterns* of acoustic features. There is, however, no strong a priori commitment to any particular level of linguistic representation, and the structures described might just as well be thought of as patterns of denotational or even relational features. Each of the above terms takes its meaning from a claim about the function of these public representations in this artificial world. We take this stance to be an important methodological and theoretical component of this work. Representations do not get to be symbols by virtue of *us* creating them or calling them such, but rather, by our reading of what they do for the members of a community of artificial cognitive agents (Clancey 1989).

The model is based on six central theoretical assumptions, derived from a theory of *distributed cognition* (Hazlehurst 1991, 1994; Hutchins 1990, 1991, 1993, in press; Hutchins & Hazlehurst 1991; Hutchins & Klausen in press):

1. No mind can influence another except via mediating structure. (The *no telepathy* assumption.)
2. No social mind can become appropriately organized except via interaction with the products of the organization of other minds, and the shared physical environment. (The *cultural grounding of intelligence* assumption.)
3. The nature of mental representations cannot simply be assumed, they must be explained. (The *shallow symbols* assumption—in contrast to the *deep symbols* assumption which brings symbols into the language of thought as an article of faith rather than as a consequence of cultural process.)
4. Symbols always have both a material and an ideal component. The material component is what makes form, structure and differences possible. The ideal component is a function of the stance that organized individuals take toward these material forms. (The *material symbols* assumption.)
5. Cognition can be described as the propagation of representational state across representational media that may be internal to or external to individual minds. (The *distributed information-processing* assumption.)
6. The processes that account for the normal operation of the cognitive system should also account for its development through time. (The *no developmental magic* assumption.)

Below we present a computer simulation that is an implementation of these assumptions. It turns out to be a very robust procedure by which a community of individuals can develop a shared set of symbols. The simulation is *not* to be taken seriously as a model of any part of human history. It is the simplest possible scheme that captures the essential properties of the system being modelled. The simulations are not to be taken as representations of actual human cognition, language, culture or experience, but as existence proofs that a particular kind of process is capable of producing a particular sort of outcome, in this case, a community with a shared lexicon. One is certainly free to question the extent to which it is reasonable to map this process on to plausible courses of human behaviour and in the discussion section we consider ways in which the model is consistent with observations about human language.

The constraints on a shared lexicon

The central problems of inventing a lexicon can be stated in terms of a description of the outcome. Consider two individuals, *A* and *B*, and a set of visual scenes or contexts in the world numbered 1, 2, 3, ..., *m*. Let the description that an individual uses for referring to a scene be denoted by the concatenation of the letter designating the individual and the number of the scene. For example, “*B*₅” denotes the description that individual *B* uses for referring to the fifth scene.¹ Now, if the lexicon is to be *shared*, the word that *A* uses for any particular scene must be the same as that used by *B*. In notation: $A_1=B_1, A_2=B_2, \dots, A_m=B_m$. Simultaneously, if the lexicon is to be a lexicon at all, there must be differences between the material forms of the words used by each individual for different scenes. In notation: $A_1 A_2 \dots A_m$ and $B_1 B_2 \dots B_m$. (It wouldn’t do to have a lexicon for *m* scenes that were *m* homonyms.²) These two constraints must in some way be satisfied simultaneously in any process that is to develop a shared lexicon. For our purposes, *a shared lexicon is a consensus on a set of distinctions.*³

It is interesting that the same mathematical basis, although using a different computational procedure, was developed by Hinton and Becker (1989) to show how modules in the brain could discover a shared communication protocol without a supervisor to specify how to communicate. The problem of discovering a lexicon may be quite general and seems to occur at a number of levels of organization in cognitive systems. Independent of its relation to problems solved by organic cognitive systems, the procedures described here might provide a general engineering solution to problems where modules must invent a means of communication but where the organization of the communications protocol cannot be specified in advance.

Cultural process

Before turning to the simulation, we need to say a few more words about the theoretical stance. The six theoretical assumptions described previously can be assembled into a model of cultural process, the temporal characterization of distributed cognition shown in Figure 9.1.

Our inventory of representational structure includes *natural structure* in the environment, *internal structure* in the individuals, and *artefactual structure* in the environment. Artefactual structure is a bridge between internal structures. Artefacts may provide the link between internal structures in one in-




Figure 9.1 A model of cultural process.

dividual and those in another individual (as is the case in communication), or between one set of internal structures in an individual and another set of internal structures in that same individual (as is the case in using written records as a memory aid, for example). Internal structures provide bridges both between successive artefactual structures and between natural and artefactual structures. Following Sperber (1985:76) we may say that “A representation [artefactual or internal structure] is of something [natural, artefactual, or internal structure] for some information processing device [internal or artefactual structure].”

Connectionist networks as simple agents

During the past eight years, developments in computational modelling employing *connectionist networks* have made it possible to think in new ways about the relationships between structure inside a system and structure outside.⁴ A *network* is composed of two kinds of mathematical objects: *units* and *connections*. Units take on activation values, generally in the range of 0.0 to 1.0, and pass activation along one-way connections to other units. This passing of activation along a connection is modulated by a real value associated with that connection, the *connection strength* or *weight*. The passing of activation from input units to output units of the network can be treated as the implementation of a function, and viewed as the network’s behaviour in a given environment. Through modification of the network’s weights, in time, the network adapts to the shape of that environment.

Connectionist networks of a class called *autoassociators* have particularly nice properties with respect to the problem of discovering and encoding structural regularities in their environment. Autoassociator networks learn to duplicate on the output layer the identical pattern of activation pre-




Figure 9.2 Autoassociator network.

sent to the input layer. Figure 9.2 shows a simple autoassociator network. It consists of three *layers* of units.

Input units are on the left, output units on the right, and hidden units in the middle. *Targets* are real valued vectors which are structurally similar to the output and input layers but, like inputs, are thought of as information external to the network. Targets are part of the environment which the network is made to learn (see below)—they provide the teaching signal. In the case of autoassociators, the targets are simply the input patterns themselves, allowing the network to learn about the environment with no additional teaching signal.

Limitations on space make a full description of this kind of information processing system impossible. The following sentences will, we hope, convey the style of computation entailed, if not the details.

Every unit in the input layer of a network has a unique connection to every unit in the hidden layer, and every unit in the hidden layer has a unique connection to every unit in the output layer (see Figure 9.2). The strengths of these connections can be adjusted. The activations of the input layer are set by external phenomena. The activations of the other units are determined by the activations of the units from which they have connections and on the strengths of those connections. Starting from random connection strengths, the task for the network is to discover a pattern of connection strengths that will produce the desired output (i.e. the target) in response to a given set of inputs. Incremental improvement in accomplish-ing this task is referred to as *learning*. The networks modelled here use a procedure called the *back-propagation of error* to find an appropriate set of connection strengths. In this scheme, the output produced is compared to the target,⁵ and the difference between output and target is an error in the network's ability to perform this input-output mapping. The connections are then adjusted to reduce this error on future trials at this task. The problem for the network can be viewed as one of finding a set of weights which meets simultaneously the constraints imposed by all of the input-output mappings it is made to perform.

Meeting the first constraint: words must discriminate between objects in the environment

Rumelhart, Hinton & Williams (1986) have shown that under certain conditions, the activations of the hidden layer units of autoassociator networks converge on efficient encodings of the structural regularities of the input data set. That is, the connections between input and hidden units must produce activations at the hidden layer which can be used by the connections between hidden and output units to produce the target, under the constraints of the function that propagates activation. For example, given any four orthogonal input patterns and an autoassociator network with two hidden units, the hidden unit activations for the four input cases should converge on $\{(0, 0) (0, 1) (1, 0) (1, 1)\}$. This is because the network must use the activations of the hidden units to encode the four cases and the encoding scheme attempts to distinguish (optimally) among the cases.

Producing these efficient encodings is equivalent to feature extraction. That is, what the networks learn is how to classify the input data in terms of distinctive features or principal components. If we fold an autoassociator in half, and allow the hidden layer to

produce representations which become a material part of the shared environment of interaction, then the (now “public”) hidden layer encodings produce one of the properties we want in a lexicon.⁶ In Figure 9.3 we have relabelled these units “verbal input-output” units because this is the location where agents produce words and receive words for comparison with the words of others.

If we take the remaining parts of the network to be a simple visual system—capable of classifying scenes in the environment—then the verbal input-output layer is capable of generating patterns of activation in response to each visual scene encountered, and these patterns are (or become) maximally different from each other (within the boundaries set by




Figure 9.3 An agent’s network.

learning and resource limitations of the network). Regarding network A ’s descriptions for the m scenes, this satisfies the constraints: $A_1 A_2 \dots A_m$.

Meeting the second constraint: word meanings must be shared

Virtually all work in connectionist modelling today is concerned with using networks to model aspects of the cognition of *individuals*. Our theoretical stance suggests that it might be useful to consider the properties of *communities* of networks. Of particular interest here is the fact that in traditional connectionist modelling, it is the programmer who constructs the world of experience from which the networks learn. Modelling communities of networks suggests that the behaviour of *other networks* might also be an important source of structure from which each network could learn. Connectionist programmers refer to the output patterns to be learned as the *teachers* for their networks. With a community of networks, we can let an important part of the teacher be embodied in the behaviour of other networks. Thus, where traditional network modelling is concerned only with the relationship of structure in the environment to internal structure, a model of interactions in a community of networks adds the universe of communicational artefacts to the picture (see Figure 9.4).

It is easy to show that consensus between two networks (say A and B) can be achieved by making the output of one the teacher for the other. If each takes the behaviour of the

other to be the target, then consensus will result. This satisfies the constraints that $A_1=B_1$, $A_2=B_2, \dots, A_m=B_m$.




Figure 9.4 Interaction between agent networks.

Implementatio

The simulation proceeds via interactions—one interaction is one time step in the simulation. An interaction consists of the presentation of a chosen scene (from the set of m scenes) to two chosen individuals, a *speaker* and a *listener* (from the set of n individuals). The functions that do this choosing determine what we call the *interaction protocol* of the simulation. The typical functions simply implement random selection from the domains of scenes and individuals, respectively. One of the individuals chosen (say A) responds to the scene by producing a pattern of activation on its verbal output layer (A speaks). The other individual (say B) also generates a representation of what it would say in this context but, as listener, uses what A said as a target to correct its own verbal representation. The listener, B , is also engaged in a standard learning trial on the current scene, which means its own verbal representation—in addition to being a token

for comparison with *A*'s verbal representation—is *also* being used to produce a visual output by feeding activation forward to the visual output layer. The effects of this learning on *B*'s future behaviour can be stated as: (a) in this context it produces a representation at verbal output more like what was said by *A*, and (b) produces a representation at visual output more like the scene itself.⁷

Over time, by choosing interactants and scenes randomly, every individual has the opportunity to interact with all the others in both speaking and listening roles in all visual contexts. The effect to be achieved is for the population to converge on a shared set of patterns of activation on the verbal output units that makes distinctions among the m scenes. That is, we hope to see the development of a consensus on a set of distinctions.

Below we discuss the results of two different sets of simulations. Simulation 1 was based on a more complex network architecture and set of scenes. Simulation 2 used the simpler network architecture already shown in Figures 9.3 and 9.4. The scenes of Simulation 2 were the four orthogonal vectors (1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 0) and (0, 0, 0, 1). The purpose of discussing Simulation 1 is to demonstrate the qualitative effects of evolving a lexicon within a (relatively) complex system space. The purpose of discussing Simulation 2 is to explore more analytically the nature of this kind of dynamic system in a more manageable system space.

Results and analysis

Simulation 1

In this simulation, each individual is an autoassociator network consisting of 36 visual input units, 4 hidden units, 4 verbal input-output units and 36 visual output units, as shown in Figure 9.5.




Figure 9.5 An individual network for the first simulation.




Figure 9.6 The scenes to be classified for the first simulation.

Notice that an additional layer of 4 hidden units appears in these networks. These additional resources were required by the networks in this simulation to enable the community to converge on a shared lexicon.⁸ The scenes to be classified are 12 phases of a moon, represented as patterns in the 6×6 arrays shown in Figure 9.6.

Developing consensus on a set of distinctions appears to be a highly likely final stable state of this dynamic system. Since the initial connection strengths of individuals are mid-range values and randomly assigned, early verbal representations do not differentiate between the scenes. Figure 9.7 shows the activation levels of the four verbal output units in response to the twelve scenes for some typical individuals at the start of a simulation run.

It is easy to see that there is little variation in the response of any individual to the different scenes. It is also easy to see that *consensus* (defined in terms of the degree of variance in terms used by individuals to represent the same scene) is quite high. That is, individuals' responses *do not* carry information which distinguishes the scenes, and these responses *are* highly similar across individuals in the community at the start of the simulation run.⁹

Figure 9.8 shows the same individuals after an average of 2,000 interactions with each of the other five individuals in the community.




Figure 9.7 Activation levels of the four verbal output units at the start of a simulation run.




Figure 9.8 Activation levels after an average of 2000 interactions.

For the most part, individuals now respond *differently* to each of the twelve scenes and all the individuals *agree* with each other on how to respond. That is, we have a consensus on a set of distinctions. Because of the random starting weights of the networks, and the random interaction protocol functions which organize their learning experiences, there is no way to predict *which* lexicon will develop—but the procedure is robust in the sense that *some* well-formed lexicon or another develops nearly every time.¹⁰

Simulation 2

In this set of simulation runs, we attempt to map out more analytically some of the properties of a simpler system. In particular, we view each simulation run as one of a large number of dynamic systems that are possible, given different initial conditions and parameter settings (Abraham & Shaw 1987, following Thom 1972). This infinite-dimensional space D of dynamic systems can be modelled by a function F which maps the following independent variables and functions into an instance of a dynamic system:

Scenes:

m =number of scenes; and

S =set of scenes $\{s_1, s_2, \dots, s_m\}$.

Individuals:

$f\text{-arch}$ =function which determines the network architecture of an individual;

n =number of individuals in the community at the start;

\mathbf{W} =set of starting weights of individuals;

μ =learning rate; and

ψ =learning momentum.¹¹

Interaction protocol:

$f\text{-pop}$ =population control function for the community;

$f\text{-scene}$ =function which picks a scene for an interaction; and

$f\text{-ind}$ =function which picks individuals for interaction.

Instantiation of these variables and functions (by the application of F) determines a unique dynamical system which evolves in time (t =cycles of simulation or interactions). In general, we expect different instantiations of these parameters to generate qualitatively different dynamic systems in D .

The parameter settings of Simulation 2

In order to analyze a small portion of the huge space D , we can make the following simplifications during the remainder of this discussion (i.e. apply F as follows):

Fix scenes to 4 orthogonal vectors of length 4:

$m=4$; and

$\mathbf{S}=\{(1, 0, 0, 0), (0, 1, 0, 0), (0, 0, 1, 0), (0, 0, 0, 1)\}$.

Fix individuals with identical network architecture, but random values for initial weights:

$f\text{-arch}$ instantiate all individuals as shown in Figure 9.3; n to vary across experiments; and

\mathbf{W} =a set of random numbers in the range -0.5 to $+0.5$.

Fix learning parameters in time and across individuals:

$\mu=0.075$

$\psi=0.9$.

Grant individuals immortality, and do not generate new individuals. Make the interaction protocol functions random:

$f\text{-pop}$ =individuals live for ever during a simulation run, and no new individuals are introduced;

$f\text{-scene}$ =random selection from the set of scenes \mathbf{S} ; and

$f\text{-ind}$ =random selection from the community of individuals.

One benefit of establishing the three parameters \mathbf{W} , $f\text{-scene}$, and $f\text{-ind}$ as random variables is that (given statistically relevant samples of simulation runs) these parameters can be approximated as fixed, thereby isolating the dynamic systems' dependence upon the one remaining variable, namely the number of individuals in the community (n).

Measures of the emerging lexicon: Avg1 and Avg2

Finally, having established a simulation run (or, preferably, a sample of simulation runs) as an instance of a dynamical system (by setting n), we can monitor the system's evolution with two measures of the community's language through time, $\text{Avg1}(t)$ and $\text{Avg2}(t)$. Avg1 is a measure of the average difference in *each individual's* verbal

representations (i.e. it measures the average “term distinctiveness”, across all scenes as denoted by each individual, averaged across all individuals). Avg2 is a measure of the variability in the *community of individuals’* verbal representations (i.e. it measures the average term variability for each scene across individuals, averaged across all scenes).¹² Our expectation is that Avg1 will tend toward 1.0 (the maximum difference in terms used across scenes), and Avg2 will tend toward

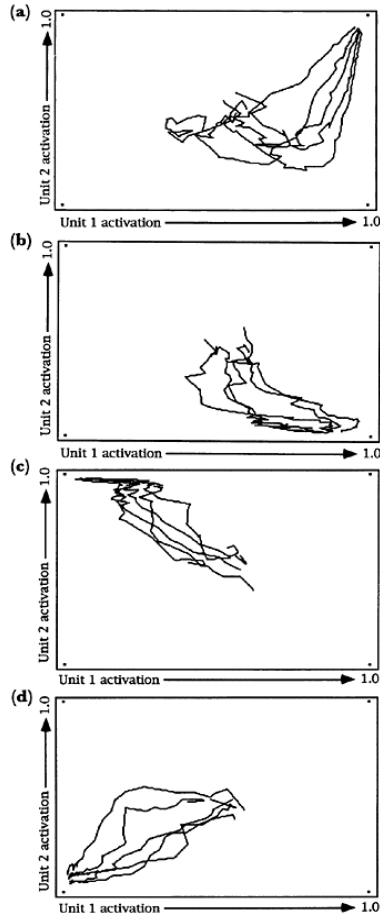


Figure 9.9 The trajectories of verbal representations for each of the four scenes in simulation 2.

0.0 (the minimum difference in terms used by the community for each scene) as the system evolves. That is, we expect to see a consensus emerge on a set of distinctions.

Figures 9.9 and 9.10 show a simulation run with a community of five individuals ($n=5$). The graphs of Figure 9.9 show the phase space of verbal representations (Unit 1

activation versus Unit 2 activation) for each of the four scenes in the world. Each trajectory on a graph represents one individual's term for that scene, parameterized by time. The trajectories all begin near the middle of each graph (verbal output activations near [.5, .5]) because, at the beginning of the simulation run, individuals are responding to the scenes with unorganized connection weights.¹³ As time proceeds, the trajectories of each graph head for one of the four corners (i.e. consensus regarding each scene increases). Furthermore, each graph's trajectories must reach a corner unoccupied by the trajectories of the other three graphs (i.e. term similarity decreases). Notice how two of the emerging lexicon's terms (Figures 9.9(a) and 9.9(b)) compete with each other for a place in the (1, 0) corner before the term representing scene (0, 1, 0, 0) of Figure 9.9(b) finally wins this competition.





Figure 9.10 The development of the measures of the emerging lexicon.

The two graphs of Figure 9.10 show how Avg1 and Avg2, plotted for the same simulation run as that shown in Figure 9.9, capture the two properties of the evolving system. Descriptions begin with a lot of consensus but lack discrimination because,

again, the networks are responding roughly the same (and uninformatively) at the beginning of the simulation. Between 1,000 and 2,000 interactions, as good representations for discriminating between scenes emerge (Avg1, the mean variability in each individual's descriptions for *different scenes*, goes up), the degree of consensus goes down (Avg2, the mean variability in descriptions representing the *same scene*, goes up). As the rate of learning to discriminate between scenes slows down (Avg1 approaches the asymptote), the representations which serve as individuals' verbal targets *change more slowly*—they become easier targets to follow, and consensus begins to emerge (Avg2 begins to descend again).

The effect of varying community size

We are now in a position to ask how the dynamic system depends on population size n , given that we have fixed all the other system parameters



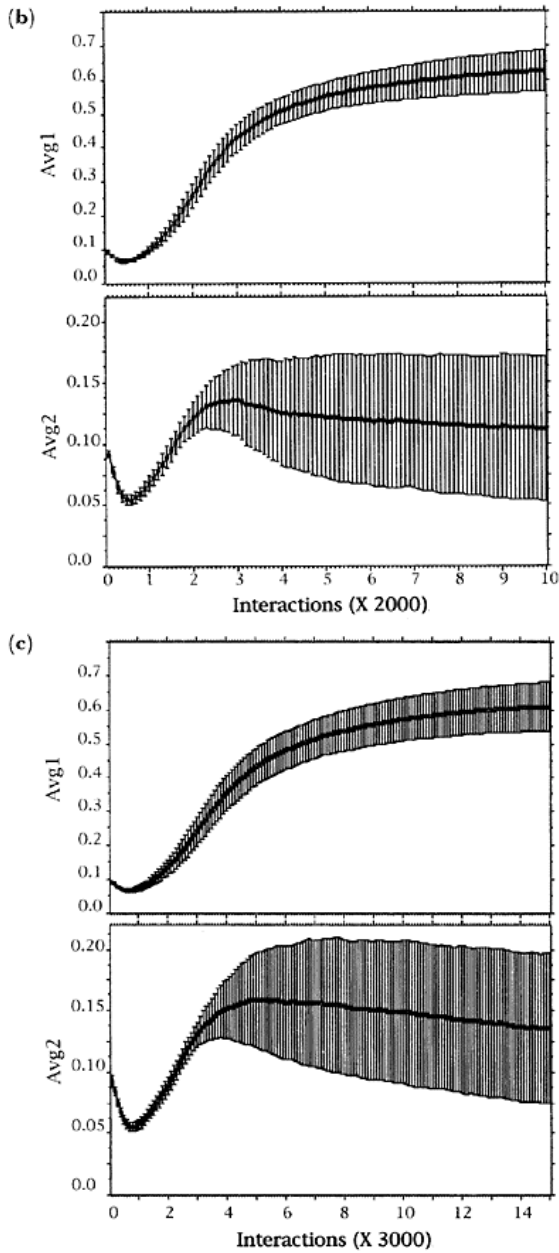


Figure 9.11 Simulations with different numbers of community members, n .

as discussed above. Figure 9.11 shows the results of three different experiments, each entailing a sample of 15 simulations.

The simulations of each experiment were run with community member sizes n of 5, 10 and 15, respectively. The means and one standard deviation error bars for the 15 observations of each experiment (sampled at regular intervals within each simulation) are shown for the two measures of lexicon structure, Avg1 and Avg2. In all three experiments, individuals have participated (on average) in the same number of interactions (namely, 2,000) by the end of the time frame shown in Figure 9.11. The general pattern is the same as that seen in the single simulation run of Figure 9.10. Of particular interest is the nature of the “decay” in the lexicon formation process shown by changes in the two measures of lexicon structure, Avg1 and Avg2, as community size gets larger. This decay is caused by the greater difficulty of organizing large communities rather than small ones. Each experiment displayed in Figure 9.11 shows that Avg1 and Avg2 of the “average community” (represented by the plotted mean values of the 15 simulations in each graph) vary smoothly and asymptotically as a function of the number of interactions. Therefore, the final steady-state of each experiment can be approximated reasonably well by the final mean value of each graph in Figure 9.11. Taking the three points so collected for each measure (Avg1 and Avg2), it appears that the decay in the ability of a community to form




Figure 9.12 The changes in the measures Avg1 and Avg2, with different sized communities.

a “good” lexicon increases exponentially by a factor of $1/n$. As community size n increases, the rate at which lexicons become less “good” slows down. This relationship, although limited to only three data points, is clearly represented in Figure 9.12.

Of course, the meaning of what “good” is (and how good is “good enough”) can only be determined by the functional properties entailed in agents *using* the lexicon, something

these simulations do not address. (But see Hutchins & Hazlehurst 1991, for how such a lexicon may be embedded in a larger context of community use.)

Discussion

The need for a critical period of language learning

We have experimented with adding new individuals with random weight-structures to communities that have already developed a lexicon. Depending on the size of the community, the addition of the new individual may have quite different effects. The behaviour of a new individual added to a large community with a highly shared lexicon will be entrained by the behaviours of the other members of the community and the newcomer will learn the shared lexicon. A new individual added to a small community may completely destroy the previously achieved solution. After such an event the community may or may not be able to recreate a well-formed lexicon with the new individual.

In running these simulations we found ourselves wishing that there was some principled way to reduce the learning rate once individuals had learned the language. In particular, one would like to reduce the learning rate at the point in the life cycle where individuals are likely to encounter many interactions with “disorganized” individuals. This would amount to implementing a critical period for language learning so that individuals learn less from the linguistic behaviour of others once they have reached sexual maturity.¹⁴ Perhaps evolution has engineered something like this into our species. We did not implement a critical period, however, because to do so seemed arbitrary and a violation of one of the core premises: the processes that account for normal operation of the system should also account for its development through time. In more complex situations, such as that of biological evolution, where adaptive searches are conducted in parallel at many levels of specification, it may be reasonable to expect violations of this premise.

The formation of dialects

Occasionally, communities fail to create a well-formed lexicon. This happens because sometimes the random initial starting points of the networks in a community are incompatible with each other, and “unlucky” choices of the interaction protocol lead to divergence in the verbal representations of these individuals, which cannot be overcome¹⁵. In this case, the kind of language that one individual is predisposed to learn is not the sort that the other is predisposed to learn, and given their starting points and learning experiences, they never find a solution that can be shared.

The fact that some pairs of initial weight configurations are more compatible than others suggested that it might be advantageous to let the individuals discover and seek out those with whom they are compatible (as demonstrated in the similarity of their behaviours, i.e. their verbal representations). A sequence of experiments was run in which the choice of a listener for each speaker was biased in favour of those who had a history of speaking (i.e. using descriptions in former interactions between the two) which

was similar to the speaker's own history of utterances. The result of implementing this interaction protocol was the formation of *dialects*, or clusters of individuals within which there is consensus on the descriptions and their referents. Since individuals interact more with "like minded" others, they are more susceptible to learning from the classification behaviours of those who act as they themselves do. We take this to be a very primitive implementation of what Sperber (1985) called "ecological patterns of psychological phenomena".

The acquisition of a coherent lexicon

Consistent with the observations reported above about dialect formation and the education of new individuals, is an interesting interpretation of the model's performance regarding the ontogenetic problem of learning form-meaning pairs that are already established in the world. In particular, it seems that the complement to the need for a critical period of language learning (which ensures that the experience of older individuals is not lost by attending to the unorganized ramblings of novices) is the need for a set of consistent models during language acquisition (which ensures that novices are exposed to a system that is coherent, and therefore learnable).

Creating a lexicon from total lack of organization (demonstrated in Simulations 1 and 2), seems to be easier than creating a functional lexicon from a system that is organized in the wrong way. That is, novices introduced into the system who must accommodate a wide range of organized variance (as with dialects) have a hard time learning any part of the form-meaning system. On the other hand, new individuals exposed to well-formed lexicons (or one well-formed dialect) gain significant leverage on the problem of learning the form-meaning system via the mediation of consistent and well-formed terms. In fact, experiments were run which showed that the strictly visual problem of classifying scenes in the simulation world is simplified by the use of consistent and well-formed terms—*individuals' abilities to learn the visual classification problem are enhanced by the existence of a coherent lexicon.*

This fact of the simulation seems to be because of the nature of the decomposition of the principal component being conducted by the autoassociator in its learning (Chauvin 1988). By grounding this process in the constraints imposed on the hidden layer by coherently organized targets, the decomposition process is accelerated significantly. There is a sense, then, in which the structure of the lexicon is an important vehicle for the learning of the visual classification problem.

In the real world, we can cite two kinds of evidence that are consistent with this aspect of our model's behaviour. The first has to do with the nature of language-acquisition environments. Children generally *do* learn language, over an extended period of time, in a relatively homogeneous population of language users. This is a fact which follows from the general nature of the social organization of our species. Second, there appears to be a good deal of empirical evidence for the fact that children *must* (or at least that they *act as though they must*) accommodate each and every form-meaning pair that they encounter (Clark 1983, 1987; Slobin 1985). That is, the language-acquisition process apparently requires the learner to assume that words contrast in meaning, and children use this as a resource in classifying the objects or events which words denote (Clark 1987).

Term contrast and consensus in natural language

Clark (1987) investigated the natural-language version of our notion of a lexicon being a *shared set of distinctions*. She provides empirical evidence consistent with the claim that natural-language lexicons exhibit two pervasive features:

1. Every possible pair of forms contrast in meaning (the principle of contrast).
2. For certain meanings, there is a conventional form that speakers expect to be used in the language community (the principle of conventionality) (Clark 1987:2).

The general claims made are that the systems which generate human lexicons are efficient (all words contrast in meaning) and conservative (established words have priority, and innovative words fill lexical gaps as opposed to replacing established words with the identical meanings). Evidence for these principles is cited from a wide range of natural-language phenomena, including the observations that:

1. Lexical domains emphasize semantic contrasts.¹⁶
2. Syntactic constructions create semantic contrast; “differences in form mark differences in meaning at both the lexical and the syntactic levels” (Clark 1987:6).
3. Well-established irregular forms in a language are maintained in the face of many resources (paradigms or patterns) for simplifying the language through regularization.
4. Innovative (new) words emerge as a consequence of a failure in an existing lexicon’s capacity for conveying the appropriate meaning (i.e. due to an inability of the existing words to establish the appropriate contrasts).

These principles of human language have parallel entailments for the acquisition of language:

1. Children rely on contrasting terms to tune their understanding of semantic fields to adult levels of specificity and generality.
2. Children assume (or at least act as though they are assuming) that new terms contrast with those that they already know.
3. Children reject (both across and within languages) terms which they understand to be synonyms with terms that they already know.
4. Children productively generate novel terms to fill expressive needs, but these terms converge toward conventional usage as their language development proceeds.

We find Clark’s analysis to be in general agreement with the phenomena modelled and suggested by our simulations, although her observations clearly exceed the range of phenomena covered by our simple simulations.

Grounding meaning in communicatory practices of the community

We find Clark’s analysis lacking in two respects: the first is made explicit by our model; the second is suggested by the general theoretical framework we have adopted, although not specifically captured in the simulations we have presented here.

First, Clark’s formulation tends to equate language use with form-meaning pairing, which lapses into a notion of meaning that is structurally derived rather than grounded in communicatory practice. This appears to stem from an analysis which takes language

from as a proximal explanation for *function*.¹⁷ In particular, without an explicit means of grounding terms in the world of experience, meaning becomes too tightly attached to form, forcing Clark into use of the *conduit theory of meaning* (Reddy 1979, Lakoff 1987).

The conduit theory of meaning takes meaning to be something transferred between language users, as if meaning is *attached to* language forms, rather than something which expresses a relationship between perceiver/ actor, context and experience as a *consequence of* situated processing of language forms. For example, terms in the lexicon only need contrast to the extent that necessary communicative functions are served in agents' uses of the lexicon. Clearly, these functions will vary across language situations, participants, and traditions—because contexts, individual experiences and the histories of use vary with each of these. This latter view of meaning becomes clear when one distinguishes between, and considers the interactions among, conventional forms of behaviour (artefactual structure), individual experience (internal structure), and the physical world (natural structure). In our own simulations, this functional grounding of meaning is evident in the ways the lexicon (artificial structure) responds to tweaking different simulation parameters (natural structure). It is also evidenced in the variability of individual network weight configurations (internal structures), which can accommodate the same lexicon (artefactual structure) given stable environments (natural structure).

Secondly, Clark's analysis undervalues the explanatory power of one of her own principles—conventionality. The principle seems to be given the role of *describing* the structure of language, but no causal role in *creating* that state of affairs. For example, in explaining the fact that children's private and idiosyncratic novel words give way to conventional expressions, Clark cites children's *efforts to contrast terms* as the mechanism responsible for con-vergence towards use of the standard forms: "It is children's discovery that two forms do *not* contrast in meaning that leads to take-over by the established term," she says (Clark 1987:18, emphasis as in original). It seems to us that Clark is ignoring a large area of communicatory functions responsible for explaining why children adopt conventional expressions. Again, structural analysis provides only proximal explanations for the mechanisms involved. Our own modelling framework keeps in focus the higher order effects of language sharing—*consensus is a functionally important property in its own right*—and seems to play a more active role in answering the question of why children conform to an established norm, than Clark posits.

In brief, a community of language users (and generations of language users) constitutes a system which enables cognitive performance that cannot be performed by individuals alone (Hutchins & Hazlehurst 1991). Children's convergence towards the use of established expressions would seem to be related importantly, not only to the individual-level problem of making meaningful distinctions in the here and now, but also to the community-level problem of constructing which distinctions are meaningful. Conventionality points to a complex cultural process—generating many language-shaping communicative properties—which only becomes clearer by taking a community of interacting language users as the unit of analysis.

The model as theory instantiation

Our model explicitly represents the interactions of the three kinds of structure discussed at the beginning of the chapter: natural, internal and artefactual. The patterns representing physical phenomena of the world (scenes) are the natural structure. The patterns of activation on the verbal input-output units are the artefactual structure. The connection strengths in the networks are the internal structure that provide co-ordination between the two kinds of external structure, and are themselves a product of artificial structure mediating experience with the world. We see this as the smallest first step towards a system in which artefactual structures invoke the experience of that which is not present in the environment.¹⁸

Let us return to the central theoretical assumptions. As we have seen, no individual can influence the internal processing of another except by putting mediating artefactual structure in the environment of the other. However, by putting particular kinds of structure in each other's environments, they all achieve a useful internal organization.¹⁹ It is possible for each individual to achieve an internal classification scheme in isolation—this is what autoassociators are known to do by themselves. But such a classification would be useless in interaction with others. That is, idiosyncratic distinctions may be useful, but not as useful as shared ones. We have noted that learning to categorize the world is easier when mediated by coherently organized verbal representations. By forcing individuals to learn from the classification behaviour of others we ensure that each individual can only become internally organized by interacting with the external products of the internal organization of others. The effects of this kind of system also enable individuals to tap the resources of an entire group (and ancestors of the group), enabling cognitive performance that is not achievable by individuals alone. This is the foundation upon which human intelligence is built (Hazlehurst 1994, Hutchins & Hazlehurst 1991, Hutchins in press).

Although this simulation is too simple to address the issue of symbolic representation directly, it suggests a way in which shared symbols that could subsequently come to serve internal functions might arise as a consequence of social interaction. Such symbols are outside the individual as pieces of organized material structure—in the behaviour of others—before they have explicit internal representations. Undoubtedly, such shared public forms can be given internal representations, as can any structural regularity in the environment, whether natural or artefactual. This perspective, in which symbols are in the world first, and only represented internally as a consequence of interaction with their physical form and social consequences, is what we mean by the *shallow symbols* hypothesis. In this view, symbols and symbolic processing may be relatively shallow cognitive phenomena, residing near the surface of the functional organizations that are the result of interacting with material structures in a cultural process.

The computations performed by the networks are well characterized in terms of the propagation of representational state. The universe of inputs is propagated through the networks and re-represented at the output. These representations, or rather the functional capacities to produce them, then become distributed across the members of the community. This general notion of computation comes from Simon (1981:153) who says, "Solving a problem simply means representing it so as to make the solution transparent." Simon may not have intended quite so broad a reading of his definition but it seems to

capture well the behaviour of this system. The structure of the natural world is fixed in this model, but the internal structures and the artefactual structures co-determine each other and co-evolve in the development of the lexicon. In the broadest sense, the solution arrived at was determined by the structure of the natural world as manifested in the phenomena encountered, in the random initial configurations of the internal states of the individuals, and in the instantiation of who actually learns from whom in the community. The process of developing a lexicon in this model is a process of propagating transformed representations of naturally occurring structure throughout a system that also contains artificial structure.

Finally, even when two networks are in complete agreement with one another about the use of the lexicon, each has a unique internal structure. Individuals in our model are able to use the lexicon *without* needing to share the internal structures that enable that use. Through learning from each other in interaction, individuals become functional equivalents, not structural replicates of each other. There is no need to posit the existence of grandmother neurons which are responsible for the like behaviours of autonomous individuals. On the other hand, behaviour *is* shaped by the constraints of autonomous individuals interacting in (problematically) shared environments.

Within this theoretical framework, we claim that meaning can be re-trieved from the unattractive positions of being equated with:

- (a) the results of a (usually innate) private language of mental symbols that stand for an uncontested, fixed and non-social objective reality (Fodor 1976; cf. Lakoff 1987); or
- (b) an unproblematically shared, static (i.e. a non-developmental, non-historical, and often non-social) semantic knowledge base (ethno-science and much of cognitive anthropology; cf. Hutchins 1980); or
- (c) a strictly public (i.e. superindividual, non-mental) symbol system (Geertz 1973; cf. Shore 1991).

Meaning in our model is an evolving property of the interaction of internal, artificial and natural structures. At any point in the evolution of such a system we take meanings to be characterizations of the functional properties of individuals *vis-à-vis* the environment (including each other). Meaning, for each individual, describes a range of possibilities for action. In a more complicated world, possibilities would be constrained by representation of the consequences of these actions, something we understand to be true of human meaning systems. In the limited world we have created, the range of possibilities for action is restricted to producing something like monolexemic utterances in the contexts of shared binary scenes. In the real world the possibilities are much greater, some of which we hope to address in future work.

Appendix

In a community of n agents who live in a world with m scenes, the community's language at time t (the recording of all agents' descriptions of the m scenes) can be written in matrix form as:

$$L_t = \begin{matrix} & S_{1,1}(t) & S_{1,2}(t) & S_{1,3}(t) & \dots & S_{1,n}(t) \\ & \vdots & \ddots & & & \vdots \\ & \vdots & & \ddots & & \vdots \\ & \vdots & & & \ddots & \vdots \\ S_{m,1}(t) & S_{m,2}(t) & S_{m,3}(t) & \dots & S_{m,n}(t) \end{matrix}$$

where $S_{ij}(t)$ is agent j 's description of the scene i at time t . Then the sequence of matrices, $\{L_0, L_2, \dots, L_t\}$ represents the evolution of the community's language from time $t=0$ to $t=\phi$. At any time t , the language L_t can be analyzed for at least two properties:

- (a) Avg1, the average difference between all pairs of descriptions of scenes for each individual; and
- (b) Avg2, the average difference between pairs of individuals' descriptions of the same scene, for all individuals and scenes.

Avg1 gives a measure of individuals' abilities to use terms to distinguish between scenes, while Avg2 gives a measure of the community's ability to agree on the terms used to identify scenes.

More formally, assume $S_{i,j}(t)$ and $S_{l,p}(t)$ are real valued vectors of length γ , and define a distance metric:

$$d(S_{i,j}(t), S_{l,p}(t)) = \sqrt{\frac{\sum_{k=1}^{\gamma} (r_{i,j}^k - r_{l,p}^k)^2}{\gamma}}$$

where r_{ij}^k is the k th real value in vector $S_{ij}(t)$.

Then,

$$Avg1(t) = \frac{\sum_{j=1}^n \left(\frac{\sum_{(i_1, i_2) \in P_2(m)} d(S_{i_1, j}(t), S_{i_2, j}(t))}{2} \right)}{n}$$

where $P_2(m)$ is the set of all pairs of integers from 1 to m , and $(m^2-m)/2$ is the size of the set. Similarly,

$$\text{Avg2}(t) = \frac{\sum_{i=1}^m \left(\frac{\sum_{(j_1, j_2) \in P_2(n)}^2 d(S_{i,j_1}(t), S_{i,j_2}(t))}{\frac{n^2-n}{2}} \right)}{m}$$

Acknowledgements

Research support was provided by grant NCC 2-591 to Donald Norman and Edwin Hutchins from the Ames Research Center of the National Aeronautics & Space Administration in the Aviation Safety/Automation Program. Everett Palmer served as technical monitor. Additional support was provided by a fellowship from the John D. and Catherine T. MacArthur Foundation.

Notes

1. In fact, a description is a vector of real values which we imagine to be articulatory features capable of generating a monolexemic word, a piece of agent-created structure, in this artificial world. The act of “referring” is a property of our construction of the world—we have built into the world the need for agents to internalize visual experiences and code that organization in externally realizable structures. This, of course, bypasses a long history of evolution which might select some of these properties as internal properties of agents. However, in granting these (on face value, plausible) assumptions we are better able to address cognitive and cultural issues with our simulations.
2. This model does not deal with homonyms or synonyms. However, note that it is a matter of current debate whether such categories in fact exist in human language when the perspective of language use and learning is taken to examine instances claimed to be synonyms or homonyms (Clark 1987, Slobin 1985).
3. See Appendix for a more formal treatment of this analysis.
4. The best background work on connectionism is by Rumelhart, McClelland et al. (1986), and McClelland et al. (1986). The behaviour of autoassociator networks is thoroughly analyzed in Chauvin (1988).
5. For an autoassociator, the target is identical to the input, thus reducing the problem to an identity mapping on the input set.
6. We thank Elizabeth Bates (personal communication, February 1991) for coining the term *public hidden units* for this construction.
7. Implementing this combination of error signals is straightforward. One error signal is computed from the difference between produced word and target word, the other error signal is the usual error back-propagated from the visual output layer. These two signals are simply added together and back-propagated to the visual input layer.
8. It is a well-known result from connectionist history that two layers of weights are required to perform mappings from input to output which are not linearly separable (Rumelhart et al. 1986). The range of verbal representations that individuals are attempting to map to in this simulation may constitute such a set, thus requiring the extra hidden layer to perform properly. We say “may”, because the mapping itself is evolving as the lexicon is being constructed. Another reason for the required extra layer has to do with the large compression

- of data from input (36 units) to verbal input/output layer (4 units). This compression tends to swamp the verbal output layer with large activation values (even for an untrained network), reducing the network's flexibility to learn. Since the learning problem is composed from the behaviours of other (now inflexible) agents, the community is unable to converge upon a shared lexicon.
9. The consensus in the starting state of the simulation is a product of the fact that random weights in the network tend to produce mid-range output values regardless of input to the network.
 10. See "The formation of dialects", section (below) for a discussion of some observed exceptions.
 11. The *learning rate* is a parameter controlling the magnitude of the effect that one learning trial has on a network, i.e. the scale of magnitude by which changes are made to weights on each learning trial. The *learning momentum* is a parameter which influences the effects that variability in the learning set has upon network learning performance. That is, the learning momentum parameter determines the scale of magnitude by which recent learning trials continue to affect the current learning trial. (See McClelland (1988), for implementation details regarding these learning parameters.) These two parameters could, conceivably, vary among individuals, perhaps also as functions of time. In the simulations of this paper we chose to fix these parameters, not letting them vary across individuals or time, in order to simplify the task of understanding the set of dynamic systems we are dealing with.
 12. See Appendix for a more formal definition of these measures.
 13. There has been no learning yet, and all individuals begin with random weight assignments to their connections, therefore all units respond at mid-range levels of activation.
 14. Of course, an alternative solution might be simply for individuals *not* to focus their learning lens too narrowly on any other single individual (or homogeneous group of individuals). The problem with this solution is that it works for organized (adult) individuals but does not work for disorganized (novice) individuals. See "The acquisition of a coherent lexicon" section.
 15. The frequency of this phenomenon is dependent upon many parameters of the simulation which determine the learning trajectories of the networks (see Simulation Two, above, for a listing of these parameters). Included here are the learning rate and learning momentum parameters of the network learning algorithm. The stability of the environment being learned also plays a critical role in determining the probability of networks becoming "stuck" in their search for a set of weights to perform the tasks required.
 16. Clark claims that true synonyms do not exist in natural languages.
 17. This fact seems to get Clark into trouble in her critique of Slobin's notion of *unifunctionality*—which denies the existence of multiple forms carrying the same meaning (something Clark agrees with), *and* denies the existence of multiple meanings being carried by the same form (something Clark disagrees with; 1987:25). Clark claims, for instance, that the English inflection *-s* is a form used productively to map on to the concepts of plurality and possession. Clark's argument here is based *solely* on evidence from the structural regularities of parts of English morphology, effectively ignoring the possible communicative and learning functions which act to create contrasts even here. These are the properties of natural language that Clark relies upon to build her own case elsewhere in the paper. Furthermore, Clark (1987:26) utilizes a logical analysis of semantic feature inheritance as an argument for rejecting Slobin's denial of true homonymy in natural language. This seems to be inconsistent with her use of language's communicative functions found elsewhere in her paper.
 18. Work on the slightly longer step of developing propositional representations and letting symbols break free of the phenomena to which they refer is now in progress (cf. Hutchins & Hazlehurst 1991).

19. Here we must acknowledge a hedge on our own claims. The model, as it stands, collapses semantic and phonological representations. There is no distinction between what an agent *conceives* of the scene and what it *says* about the scene. Similarly, what an agent says is unproblematically heard by the other agent participating in the interaction. This is, strictly speaking, a violation of the “no telepathy” assumption, and the separation of internal and artifactual structure. Having acknowledged this discrepancy we can state that the model, as it stands, gains no explanatory power from this conflation and is therefore not a violation in principle. We have experimented with architectures that produce phonological representations from semantic representations, and vice versa. Implementation of these constructions do not render invalid the general claims of the model presented here.

Chapter 10

MANTA: new experimental results on the emergence of (artificial) ant societies

Alexis Drogoul, Bruno Corbara, Steffen Lalande

The MANTA project, already described in Drogoul et al. (1992a, b) and Ferber and Drogoul (1992), is an application of the EthoModelling Framework (EMF) to the simulation of the social organization of ant colonies. EMF is based on the principles of multi-agent simulation (for a formal definition, see Ferber & Drogoul 1992), which means that each individual in a population is represented by an artificial entity whose behaviour is programmed (computationally speaking) with all the required details (for other works on multi-agent simulation see Gilbert & Doran 1994, Hogeweg & Hesper 1985, Collins & Jefferson 1991). Multi-agent simulations are of most help in modelling situations in which individuals have different complex behaviours and where the interactions involve so much non-linearity that they cannot be described easily within a mathematical framework. Our aim with MANTA is to test hypotheses about the way in which social structures, such as a division of labour, emerge as a consequence of the behaviour and interactions of individuals. In other words, we want to evaluate the minimal set of causes that has to be provided at the micro-level to observe definite structures at the macro-level.

Some early simulation experiments conducted with a preliminary version of MANTA have been presented in Drogoul et al. (1992a, b). They showed that it is possible to obtain the emergence of a division of labour within a nest of “simple” ants, i.e. ants provided with only three tasks. Nevertheless, our goal at that time was not to simulate the complexity of the social organization as it is observed in real nests but to demonstrate the relevance of our approach. This chapter presents the preliminary results of a more ambitious set of simulation experiments on the mechanisms of sociogenesis (Wilson 1985) and division of labour. These experiments have been conducted using a new version of MANTA, in which the ants are more realistically represented than in the previous one, being provided with a larger set of behaviours.

The MANTA agents’ model of behaviour

The stimulus/task model

In MANTA, the behaviour of the agents that compose a colony is programmed using the stimulus/task architecture provided by EMF. This architecture assumes that the behaviour of an agent can be completely described by a set of independent tasks, each being constructed from a sequence of elementary behaviours called *primitives*. An agent cannot

engage in more than one task at a time. The tasks are triggered by stimuli that can have variable strengths and may be either internal or provided by the environment. Stimuli are associated with pre-existing motivations, expressed in terms of tasks' weights, thresholds and activity levels.

The *weight* stands for the relative importance of the task within the agent. In addition to being a component of the basic behaviour reinforcement process¹ provided for the simulated ants, the weight is also used to accumulate the agent's previous experiences of a given task. A high weight indicates that the agent is specialized in the task. The weight is used to compute the *activation level* of a task when it is triggered by a stimulus. Its reinforcement is considered to be a form of long-term positive feedback.

The threshold is viewed as an indicator of the "motivation"² of the agent to perform a given task. This motivation is increased continuously as long as the task is not activated, and decreased whenever it becomes active. As a result, the threshold of the neglected tasks decreases and the threshold of the current task increases during the task selection process. These operations are considered to be short-term positive and negative feedback.

Finally, the *activity level* is viewed as an indicator of the agent's motivation to continue the task it is performing. The activity level is initialized with the activation level of the task and then decreased continuously as long as the task remains active. When it reaches zero, the task is stopped. Hence a task is considered to be capable of being activated when its threshold is lower than its weight multiplied by the strength of its triggering stimulus (that is, its activation level). And it is selected when this activation level surpasses the activity level of the current task. The agent then switches from the previous task to the new one.

Communication

As defined in the kernel of EMF, agents do not communicate with one another. They just drop stimuli into the environment, and these stimuli may or may not trigger the behaviours of other agents. Dropping a stimulus results in the creation of a gradient field around the emitter (by diffusing this stimulus in concentric circles around it), and in a way analogous to the chemical, visual or oral propagation of information. Some of these communications are part of the agents' behaviour (they can deliberately choose to propagate stimulus to attract or repulse other agents), and some of them are independent of their behaviour (for example, an agent will always be visible to other agents, be they friends or enemies, if they can see it). The propagation of stimuli is only stopped by obstacles such as walls.

Implementation

EMF follows an object-orientated approach, although the agents do not use message-passing as a means of communication. The agent objects are instances of classes that inherit from Ethobehaviour or one of its subclasses which implement the abstract model of behaviour they are going to use (cf. the use of objects described in Chapter 5). Each of these classes represents a species of agent and provides the agent with information about primitives, tasks, stimuli and other domain-dependent pieces of knowledge shared by all agents. This information can either be described in the agent's class, or inherited from

higher classes. The subclasses of Ethobehaviour are divided into two sets: *abstract* classes and *concrete* classes. Abstract classes (such as Ethobehaviour) cannot be instantiated. They are used to define the primitives and knowledge associated with primitives. Concrete classes, which can have instances, inherit these primitives and use them to define the tasks (behaviours) of their agents. An example of such a hierarchy is provided in Figure 10.1, in which all the classes whose names end with *Behaviour* are abstract classes.

Interfacebehaviour implements the protocols of the agents' user-interface capacities (graphical trace, inspection, etc.). *Locatedbehaviour* provides all its subinstances with the ability to be in an environment and to act on it (propagating stimuli, for example). *Curingbehaviour* implements the primitives needed by agents that will have to cure other agents or receive care from them. *Feedingbehaviour* implements primitives needed by agents that will have to feed themselves, feed another agent or be fed by one. *Maturing-*

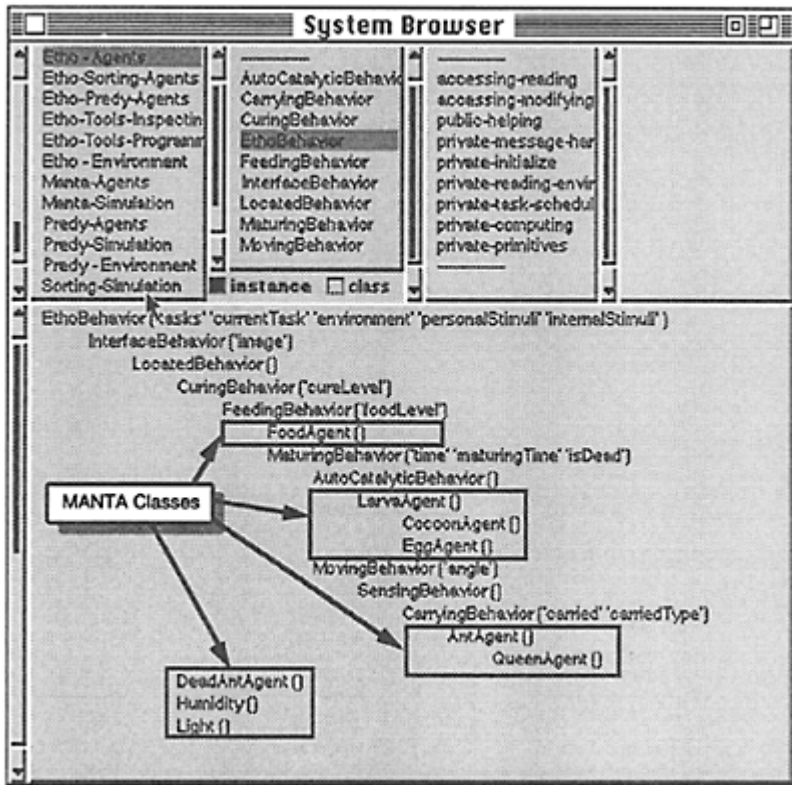


Figure 10.1 The hierarchy of classes of the MANTA application. The concrete classes are outlined. All the classes inherit from EthoBehaviour.

behaviour provides its subinstances with the capacity to grow old (and subsequently die!) and to become another agent. The notion of time is implemented in this abstract subclass, as an internal stimulus. The subclasses of these last three classes must define some domain-dependent pieces of knowledge, such as the average expectation of life of their instances, their needs for food and so on. This knowledge will be used to compute implicit stimuli that will always be propagated by these agents. These stimuli are named after the type of agent (ant, larva, etc.) prefixed by “cure-”, “hungry-” or “maturing-”. *Movingbehaviour* gives agents the possibility of moving in their environment. *Sensingbehaviour* implements two primitives for following or fleeing a gradient field. *Carryingbehaviour* implements the primitives needed for carrying other agents.

The MANTA classes

The concrete subclasses can be divided into three sets. *FoodAgent*, *DeadAntAgent*, *HumidityAgent* and *LightAgent* represent environmental agents, that is, agents that do not perform any behaviour other than that of propagating their stimuli (`#deadAnt`, `#food`, `#light` and `#humidity`). These stimuli can trigger some ants' behaviour (`#food`) or be used as guides for behaviour

		Emittor	Name	Primitive Sequence	End	Int.	Ext.
Primitives			egg	→ →			✓
	Put down...		cure Egg	→			✓
	Pick up...		larva	→ →			✓
	Has food ?		cure Larva	→			✓
	Follow/flee...		hungry Larva	→ →			✓
	Eat		maturing larva	→ →			✓
	Kill...		cocoon	→ →			✓
	Cure...		cure Cocoon	→			✓
Agents			Ant	→		✓	✓
	Eggs		hungry Ant	→ →		✓	✓
	Larvae		kill Larva	→		✓	
	Cocoon		kill Egg	→		✓	
	Ants		food	→ → →		✓	✓
	Food		light	→			✓
	Humidity						
	Light						

Figure 10.2 MANTA tasks, primitives and agents.

(#humidity or #light). *EggAgent*, *LarvaAgent* and *CocoonAgent* represent the brood agents, that is, the three stages needed for an egg to become a worker. These agents have to be cured, fed and carried by the ants during their lifetime. These needs are represented in the model by the propagation of stimuli inherited from superclasses: #cureEgg, #cureLarva, #cureCocoon, #hungryLarva, #egg, #larva, #cocoon that trigger the appropriate tasks within the ants. Finally, *AntAgent* and *QueenAgent* represent the “active” agents of the colony, namely the workers and the queen. These agents are provided with around 15 tasks, constructed from the primitives inherited from their superclasses, that cover all their needs as well as the brood’s needs (see Figure 10.2).

QueenAgent simply defines an additional task called #layEggs, which represents the only difference between queens and workers in the species we have studied. It is

triggered by an internal stimulus whose strength is a periodic function of the time (the task should be triggered every six hours).

Sociogenesis

Definition and assumptions

In the foundation process or sociogenesis (a word coined by Wilson (1985) by analogy with morphogenesis), the newly-fertilized queen alone initiates a new society. In *Ectatomma ruidum*, a tropical ant belonging to the phylogenetically primitive subfamily of the Ponerinae and whose societies are the natural counterparts of MANTA's simulated ones, the foundation is said to be semi-claustral. The founder leaves her nest to provide the food that is necessary for herself and, above all, for the larvae (in a claustral foundation the queen provides food to the larvae by using her own internal reserves derived from the degeneration of her wings' musculature). Furthermore, as we have shown elsewhere (Corbara 1991), the queen generally continues to forage after the emergence of the first workers. Sociogenesis is then a remarkable challenge for the queen: she has to face a very complex situation alone during the first stages of the society, when she has to take care of the whole brood as well as going to find food. And it is not really surprising to see that, even in laboratory conditions, only 15 per cent of sociogeneses succeed.

We chose to experiment on sociogenesis for three reasons:

1. The first is related to the phenomenon of sociogenesis *per se*. The first occurrence of emergence at the social level is to be found in the generation of the society itself. It would be unnatural to study transformations of the social organization through the impact of population increase, environmental variations or experimental manipulations without first paying attention to the early stages of the society, which obviously condition all its subsequent evolution.
2. The second reason relates to the validation of the model. Validation requires the comparison of data obtained from the observation of natural societies with data from the corresponding simulated ones. Sociogenesis provides us with at least two ways to estimate the validity of the model, as the comparison can be done with two different sets of data: demographic data, i.e. the number of each kind of brood and nestmates at each stage of the development of the societies; and behavioural data, i.e. the individual behavioural profiles of the ants and the functional groups they constitute during the societies' development. The success of natural sociogenesis under laboratory conditions can be quantified easily by looking at the evolution of the colony's demography. The analysis of behavioural data is more complex and, above all, time-consuming. Moreover, detailed results about the sociogenesis of natural ant colonies and their demographic evolution are already available (Corbara 1991; Corbara et al. 1991). It is therefore possible to evaluate the success of artificial sociogenesis by studying only the




Figure 10.3 The demographic evolution of the brood population in a natural ant colony.

demographic evolution of the colony and without paying attention to the artificial ants' behavioural profiles and their subsequent social status. As an example, the demographic evolution of the brood population in a natural colony named ERF2 is shown in Figure 10.3.

3. The third reason is that performing an artificial sociogenesis provides us with a totally “artificially built” colony, whose structure has been generated by means of a self-organizing process, rather than translated from its natural counterpart. This way, we hope to reduce to a minimum the bias usually introduced by such a translation.

Artificial sociogenesis

The simulation experiments have been conducted with a model that reproduces the laboratory conditions in which natural sociogeneses have been studied, the shape of a laboratory plaster nest, made of 9 adjoining chambers, being the same as shown in Figure 10.4.

Some humidity agents are placed on the left-hand wall, and a light agent outside the nest. Food, whose amount depends on the number of ants in the colony, is provided every day (of the simulation timescale) at the entrance to the nest, as in the case of true ants that were fed *ad libitum*. All the experiments have been run with the same parameters (tasks' initial weights, biological parameters, etc.), which means that the only source of unpredictability is to be found in the ants' random movements.

An experiment begins by putting a queen agent alone in an empty nest and leaving it to evolve. The experiment is stopped when the queen dies from starvation or more than 20 workers have been born. The first situation signifies the death of the colony and the failure of the sociogenesis process; the second is a good indicator of successful sociogenesis in natural colonies. A natural colony that reaches this stage ordinarily keeps

on growing, which means that reaching around 20 workers constitutes the end of the most difficult period in the sociogenesis process. Table 10.1 reports the results in terms of success and failure. Failure cases are clustered in seven categories, which correspond to the composition of the population when the queen dies.

In these experiments, the proportion of failures (78 per cent), appears to be close to that observed for natural colonies bred under laboratory conditions (where 86 per cent did not reach the 10 workers stage). The situations in which the foundation of the colony is likely to fail can be identified by the fact that larvae are part of the population (larvae were present in 90




Figure 10.4 The plan of the nest.

per cent of the failures). The simplest explanation of why this is so is that larvae must be cared for, carried and fed, whereas the other agents composing the brood just need to be cared for and carried. The presence of larvae generates, at the population level, a very important need for food, propagated by means of the appropriate stimuli. Therefore, the agents that can bring back food to the colony (namely, the queen and the workers) will have to do it more often, simply because their feeding behaviour will be triggered much more frequently than before. This does not cause much of a problem so long as:

- (a) many workers are available, because other tasks still have a good probability of being executed. But in the early stages of the foundation when the queen is alone, she is prevented from doing anything else, even keeping enough food to feed herself;

Table 10.1 Proportion of failures and successes in the sociogenesis experiments. The “failures with larvae” item sums all the failures obtained when larvae were part of the population, and provides the proportion of these failures with respect to the total number of failures.

Results	Composition	Number	Percentage
Failures with larvae	Eggs	8	6.06
	Eggs, larvae	16	12.12
	Larvae	27	20.45
	Eggs, larvae, cocoons	24	18.18
	Larvae, cocoons	16	12.12
	Larvae, workers	10	07.58
	Eggs, cocoons	2	1.52
Failures with larvae		93	90.29
Total number of failures		103	78.03
Total number of successes		29	21.97
Total number of experiments		132	100.00

- (b) food is not often far from the brood, which congregates near humid places so that it can be neglected for a long time; and
(c) food has not begun to run short, and so the queen is not obliged to stay outside the nest (or near its entry) until more food is supplied.

We have also conducted successful experiments. And, given the constraints stated above, we will see which “emergent strategy” is employed by the queen and the first workers to succeed in the foundation of the colony. By “emergent strategy” we mean a global long-term behaviour that *looks like* a strategy from the point of view of the observer, but which has not been coded into the behaviour of the agents.




Figure 10.5 The demographic evolution of colony #46.

We will begin by looking at the demographic evolution of colony #46, depicted in Figure 10.5, where the number of each type of agent is plotted against the time (each unit corresponds to two days of the real timescale). This colony provides us with a complete illustration of what happened in our successful experiments.

The curve in very light grey, which represents the population of eggs, is very uneven, although it is possible to detect a regularity in its sudden falls, which occur approximately every fifty days. These falls appear to be synchronized with the peaks of the second curve (in light grey) that represents the larval population. It is apparent that there is an interplay between the populations of eggs and larvae at the colony level. What is interesting, however, is that this interplay has not been coded in the system. As it could be seen as a particular side effect of a special initial configuration, we have represented the curves obtained with three other colonies (#33, 37 and 44) in Figure 10.6. These diagrams show clearly that we obtain the same kind of population fluctuations, although with small variations, in all cases.

In all these successful experiments, it is easy to detect the gain in stability provided by the increase in the population of workers. Although the frequency of the brood population variations does not change, the arrival of the workers induces a great change in their amplitudes. This is especially true for the curves representing larvae and cocoons. As the average number of eggs laid by the queen per day does not vary, we can assume that these eggs have more chance of becoming larvae when the society includes workers. The explanation is twofold: first, foraging activities are likely to involve several ants, thus increasing the amount of food arriving inside the nest. This prevents many eggs from being eaten by the ants (i.e. converted into “alimentary eggs”). Secondly, the fact that the foraging tasks are performed allows the other ants (including the queen in most cases) to be more attentive to the needs of the brood. This prevents eggs, larvae and cocoons from dying through lack of care, as was often the case in the first stages of the society when the queen had to forage and take care of the brood simultaneously.

This set of experiments allows us to underline two points:

1. The dynamics of the sociogenesis process in natural colonies can be simulated by the interactions of artificial agents provided with a small behavioural repertoire and a simple mechanism of behaviour reinforcement. It does not mean that ants really behave in this way; it just means that a stimulus/task model of behaviour at the individual level, even if it does not account for the plasticity of real ants, is able to gen-




Figure 10.6 The demographic evolution of the colonies #33, #37, and #44.

erate a complex pattern at the colony level which does account fairly well for the development of ant colonies.

2. The whole demographic evolution of the society is generated without any centralized control. The kind of “birth regulation” that emerges here is entirely due to local behaviour and interaction between ants, or between ants and the brood, and is conditioned by the constraints of the environment, as we shall see below.

Regulation of the colonies’ demography

In the previous experiments, the quantity of food provided to the colonies was incremented automatically by a fixed amount at each birth of a worker. This is what is normally done under laboratory conditions. In that way, the colony has enough food to feed all its workers. However, it does not necessarily take place this way in natural conditions. The amount of food (either prey or vegetables) is much more variable and affected by events that cannot be controlled either by the colony (climatic changes, etc.) or the experimenter. In these natural conditions, ant societies regulate their demographic evolution automatically so as to adjust the population to their environment.

In order to see whether the MANTA model could simulate this “self-control”, we have conducted a set of sociogenesis experiments in which the quantity of food provided to the colony is limited to the quantity appropriate to a fixed number of workers. The system still increments the daily food amount at each birth, but only until the limit is reached. The amount is also decremented at every death. All the experiments were conducted with a maximum fixed to the food requirements of 10 workers.

The results reveal one strong tendency, illustrated by the diagram in Figure 10.7 which depicts the demographic evolution of colony #P3_2. It could be interpreted as an inclination toward “wisdom” and is common to approximately 90 per cent of our successful experiments. The colony increases to 10 or more workers, and then decreases to five or six workers, where it stabilizes. Experiments such as #P3_2 that have run for one year of simulated timescale give the impression of stability at approximately half the theoretical maximum size of the colony.

However, it appears that this stability is the result of a dynamic equilibrium. This fact is illustrated by the demographic evolution of colony #P11_2, which clearly shows two “attempts” by the colony to increase its




Figure 10.7 The demographic evolution of society #P3_2, where the distribution of food is restricted to 10 workers (number of each type of agents against the simulation time in days).

population, immediately followed by a temporary stabilization at around seven workers (see Figure 10.8).

The stabilization or equilibrium obtained at the workers' level is the result of the dynamics at the brood level: the eggs, larval and cocoon population sizes vary quite substantially from day to day. The behavioural explanation that can be given is as follows:

- (a) when the number of workers is under the limit fixed in the experiment, the evolution of the population remains similar to that observed in the previous experiments on sociogenesis. There are no objective reasons for it to be different; and
- (b) when this number exceeds the limit, the same phenomenon as that observed in the interplay between the populations of eggs and larvae occurs. The food needs of the society increase and several eggs are killed and converted into alimentary eggs. This automatically reduces the chance of this generation of eggs of becoming larvae. As alimentary eggs do not provide enough food to maintain the number of




Figure 10.8 The demographic evolution of society #P 11_1, where the distribution of food is restricted to 10 workers (number of each type of agents against the simulation time in days).

workers, some workers die (by starving). This lowers the demand for food and stabilizes the process. But during this stage there have been savage cuts in the generation of eggs (this is particularly obvious from Figure 10.7), which constrain the society to wait for a full maturation cycle before obtaining newly-born workers. This explains the three months' stability between the two peaks in Figure 10.8. And the whole process then starts over again.

Despite the fact that these results are difficult to relate to the results of natural experimentation (this kind of experiment has not been conducted, as far as we know, on the species we study), the simulations presented here represent a very important step in the validation of the model. The simulation of sociogenesis has allowed us to obtain societies that grow in much the same way as natural ones. We now have societies of artificial ants that can adapt their size dynamically to environmental changes. Although this has to be confirmed by means of natural experimentation, we are quite sure that it is very close to reality.

From dominance relationships to hierarchy (preliminary results)

A large majority of ant species show a social structure governed by two means: a division of labour created through individual specialization, and a hierarchical organization generated by dominance relationships between individuals. This hierarchy can involve all the ants of the colony or be limited, in polygynous societies, to the reproductive females, namely the queens. In the latter case, one queen emerges during the foundation as the leader and habitually prevents the others from laying eggs, thus converting them into simple workers. This allows the society to get ahead on the schedule of the sociogenesis process, and thus reduces the chances of failure.

Ectatomma ruidum is a monogynous species and no hierarchical organization has ever been reported for it. However, we were interested in searching for the conditions in which a hierarchy could appear between queens without centralized control and seeing what our artificial queens could gain from being hierarchically organized. In the following sections the gain will be evaluated with respect to the success of the sociogenesis process.

In conducting such experiments, we abandon temporarily *ethological simulation* in order to perform *artificial ethology* (Collins & Jefferson 1991) and look at the system with the eyes of the ethologist, as if it were really a living system. We believe this approach may become a powerful tool for ethologists, a kind of “idea generator” for further natural experiments. For example, the experiments described in the next section will soon have natural counterparts³. With respect to this aspect of multi-agent simulation, we believe that what is important is not so much the answers that are given, but rather the questions that are asked.

Polygynous sociogeneses

The first set of experiments consists of beginning a foundation with four queens identical to those used previously in order to evaluate the impact on the sociogenesis process. Each individual (brood agent and worker) is labelled throughout its evolution with a lineage number (from 1 to 4) representing the queen from which it is descended. The results are surprising. First, there is a decrease in the proportion of successes, down to 6.6 per cent, that can be explained by the “group effect” of the brood. The four queens lay eggs continuously during the sociogenesis. They then have to face a brood much larger than those that characterized the previous sociogeneses, where a queen had three workers to begin with. Secondly, the successful experiments show an evolution that is exactly comparable to that of a monogynous sociogenesis (the number of eggs is roughly multiplied by 4, but the multiplying factor is around 1.5 for the other categories of the brood and for the workers—see Figure 10.9).

This is because of the very large amount of cannibalism (conversion of eggs into alimentary eggs). Thirdly, as depicted in Figure 10.10, the demo-




Figure 10.9 The demographic evolution of the whole population in colony #N8. Note that, although the number of eggs is high, the size of the other populations are comparable to those of the monogynous sociogenesis.




Figure 10.10 The demographic evolution of the four lineages of eggs plotted against the time (in days).

graphic evolution of each lineage appears to be strictly analogous to the previous ones, as if four small sociogeneses had been conducted in parallel.

In brief, a polygynous sociogenesis is not advantageous for our artificial entities. Having four queens together in the same nest is not similar to having one queen plus

three workers. What a hierarchy provides is the possibility of switching dynamically from the former to the latter. However, the lack of success of polygynous sociogenesis should not be surprising. Our artificial queens have been provided with a behavioural repertoire compiled from the observation of the queens' behaviour in a monogynous species. Even if the real queens can exhibit enough plasticity to react in a proper manner to a polygynous situation (and that is what we are going to test by performing natural experiments—see Note 4), the behaviour has not yet been studied. Hence, our simulated queens are not likely to adopt behaviour adapted to this somewhat artificial situation.

Polygynous sociogeneses with a repulsive behaviour

We can now draw together hypotheses on the behaviours that are necessary for making polygynous sociogeneses as efficient as those observed in other species. The first hypothesis comes from The raulaz et al. (1990), who explain that, in the case of bumblebees, there is a link between the rank of an individual in the society and its mobility inside the nest. According to them, a dominated bee will be less mobile than the bees that dominate it. The second hypothesis comes from observation of the spatial positioning of ants within the nest. When laying eggs, queens try to put them as close as possible to the most humid places in the nest. Because of the position of the humidity, the eggs aggregate in the top-left room. Eggs placed outside this area die more quickly than the others. The idea we had was to combine these two hypotheses by providing the queens with new behaviour, triggered directly by a stimulus they propagate, which makes them flee the gradient associated with that stimulus. The consequence is that queens will now repulse each other. Our assumption was that a queen already engaged in laying eggs would push the other queens out of the brood room, thus resulting in a differential positioning of the brood inside the nest. Being rejected from this room, the queens would then lay eggs in other places, with less chance for them to grow old⁴. This would result in turn in differential evolution of the lineages and reduce the demands of the brood, converting these queens into worker-like ants.

Despite the plausibility of this chain of reasoning, our expectations were

Table 10.2 The proportion of failures and successes of polygynous sociogeneses began with four queens. The percentages indicated for the item “Successes with...” are computed on the number of successes.

Results	Composition	Number	Percentage
Total number of failures		5	15.15
Total number of successes		28	84.85
Successes with	1 queen	20	71.43
	2 queens	3	10.71
	3 queens	3	10.71

	4 queens	2	7.14
Total number of experiments		33	100.00

not fulfilled or, more precisely, not fulfilled in the way we had imagined. Quantitatively, the experiments with these new queens were more successful than the previous ones with either monogynous and polygynous sociogeneses (see Table 10.2). The proportion of failures dropped to 15 per cent compared to 76 or 94 per cent.

However, the qualitative results were not what was expected. More than 71 per cent of the successful experiments concluded with only one queen, the other three dying. And if we look more closely at these societies, we discover that the deaths of the three queens occurred in the first stages of the society (see Figure 10.11).

We observe that:

- In the societies that end up with more than one queen, there is little spatial differentiation among the queens. A queen occupying the brood room for laying eggs or taking care of them does not always stay there. This is essentially because, being closer, she is more exposed to the larval demand for food than are the other queens and this obliges her to leave the room. While she is outside, the other queens can lay eggs or carry their eggs into the brood room. Nevertheless, the repulsive behaviour does seem to have a powerful impact (although not the one we were expecting) on the success of the sociogenesis. It seems to allow a queen to complete a task without being disturbed continuously by the other queens.
- Explaining the success of societies that end up with only one queen is more difficult. As a matter of fact, most of the simulation runs look similar to the monogynous sociogeneses. We have, however, observed that:

(a) the most difficult stage of the foundation process is when the first




Figure 10.11 The demographic evolution of the polygynous sociogenesis #N49. The sizes of the

population are plotted against the days (/2). Note the quick death of the queens 1, 2 and 4 after just one month of life.

larvae appear;

(b) in most cases (around 80 per cent), the three queens die during this stage; and

(c) around 70 per cent die because they spend their time taking care of the newly-born larvae.

We can conclude, first, that the successes arise because of the help provided by the three queens in taking care of the larvae; once this stage has been completed, i.e. when the first cocoons appear, the last queen is certain to survive (see the low proportion of failures during this stage, shown in Table 10.1); and, secondly, that the eggs laid by the three queens are always converted into alimentary eggs, which means that they constitute a supplementary source of food compared with that available in a monogynous foundation.

– We conducted 10 experiments where we provided the queens with the ability to reinforce their repulsive behaviour (by increasing its weight). Our aim was to obtain a differential reactivity to this stimulus that could be converted easily into a hierarchical pattern⁵. Although we did not obtain clear differentiation, the results were surprising: all 10 experiments succeeded and six of them finished with four queens, with several lineages clearly dominated in size by the others. These results need to be confirmed by more experiments, so we shall not develop them further here. But this approach certainly constitutes the key to obtaining emergent hierarchical structures.

Conclusion

All the results presented here are preliminary (except those on monogynous sociogeneses) and we are still conducting experiments with alternative hypotheses in order to understand the mechanisms underlying the generation of social structure within ant societies.

The results that we are waiting for are those on the division of labour within the colonies. These are still being analyzed and examined, but because they represent a huge amount of work, they will not be available for some time. We need them to begin experiments on sociotomies (Lachaud and Fresneau 1987)—the splitting of a colony into two or more subcolonies made of different functional groups—in order to observe how these subcolonies produce the specialists that are lacking. We also want to conduct studies on the relationships between social structure and spatial organization in the nest (see Fresneau et al. 1989).

We conclude with a word on the simulation we have conducted on the generation of hierarchical structures. Apart from the interest of such a simulation in itself, the underlying problem it emphasises is the apparently unnatural relationship between competition and co-operation. As it is programmed, the behaviour of the queens makes them compete with each other without any kind of collaboration. However, the results obtained with the polygynous sociogeneses (and especially those concerning the colonies

that terminate with one queen) can be interpreted easily in terms of co-operation between the queens. In most cases it is even possible to describe the evolution of these colonies by using notions such as “altruism” or “social sacrifice” that are currently employed by sociobiologists to describe similar phenomena among social insects. What we want to underline is that the notion of co-operation is a very fuzzy one and that many processes or behaviours viewed as being co-operative may be obtained by the competitive interplay between individuals and the constraints generated by their environment.

Acknowledgements

We are very grateful to Michel Denis, IPT-Université Paris XIII, for having allowed us to use 15 IBM computers which were the hosts of our simulated ant colonies for several weeks. We also wish to acknowledge the contribution of Steffen Lalande in programming the Pascal version of MANTA during the first year of his doctorate; Nicolas Czternasty for having patiently conducted most of the experiments on hierarchy; and Philippe Louise for having discovered the importance of birth control, even if it was unplanned. We thank them for their wise advice and stimulating discussions.

Notes

1. This process is detailed in (Drogoul et al. 1992a). It consists simply in incrementing the weight of a task when it has been activated.
2. In fact, the inverse of this motivation: when the threshold is low, the agent is considered to be very motivated.
3. A hundred experiments have begun on polygynous sociogenesis within a monogynous species at the Ethology Laboratory of Paris XIII.
4. It has to be noted that we do not pay attention to the phenomenon of inhibition that has been reported to occur in many species of ant. This phenomenon makes a queen inhibit the reproductive capacities of the other potential reproductive females by propagating a special pheromone. We shall probably consider it in our future experiments on hierarchical organization. The problem is that we would have to modify the model of behaviour of the agents in order to implement the inhibition of behaviours along with their activation.
5. A queen with a high weight would have been repulsed more frequently than one with a low weight and could then be considered to be dominated.

Chapter 11

Emergent behaviour in societies of heterogeneous, interacting agents: alliances and norms

Nicholas V.Findler and R.M.Malyankar

This chapter deals with alliance formation and norms of behaviour in distributed intelligent systems. An alliance is a social entity composed of intelligent agents that have made a deliberate decision to join forces to achieve alliance goals. Norms are defined as limitations on agents, such as legal constraints on individual actions. These phenomena are considered to be emergent behaviour in the context of distributed intelligent systems. Emergence is defined as the appearance of patterns of behaviour that are discernible only at higher levels of observation. Multiple levels of emergence are possible. The aim of the research is to arrive at descriptive and prescriptive theories of group behaviour.

The chapter begins by describing the basis for alliance formation in terms of the goals of individual agents. Procedures for computing the similarity of agent goals and procedures for alliance formation are informally and formally described. This is followed by a description of a mechanism for implementing norms and for the evolution of norms over time. Finally, experiments in progress using these procedures are described in the context of international power politics.

It is hoped that these studies will contribute to the understanding of human organizations, including international political action, national and international markets, industrial problems, co-ordinated efforts, and other social situations, both natural and artificial.

Introduction

Distributed Artificial Intelligence (DAI) is, in general, concerned with how the members of a group of intelligent agents perceive and reason about the environment, perform, interact, collaborate and/or compete with each other (Bond & Gasser 1988b, Durfee et al. 1989, Findler & Lo 1986, Findler 1987, 1990, 1992). This chapter describes efforts towards developing a system to study collective phenomena in the emergent behaviour of heterogeneous agents. It uses the concepts and methods of computer science and artificial intelligence as well as those of organization theory, sociology, social psychology, political science, economics and management science.

This research deals with communities, rather than collections, of intelligent agents. The agents may have different capabilities and resources, knowledge bases, goal structures, value scales, representations and viewpoints of the environment, social and organizational attitudes, methods of reasoning and decision-making, commitments and belief structures. The aim is to obtain descriptive, explanatory and prescriptive theories of group behaviour. These theories would tell us how individual members form and dissolve

non-permanent *alliances* that share some of the members' goals, act under joint constraints (principles), exchange messages, make decisions, plan actions, collaborate and compete with other alliances. The ultimate aim is to learn why the total behaviour is more than the average of the individual behaviours. The term "alliance" is used to refer to a stronger form of co-operation than generally found in the DAI literature (Davis & Smith 1983, Durfee et al. 1987, Findler & Lo 1986, 1991, 1993, Findler & Stapp 1992, Findler & Sengupta 1992).

The alliances are *social entities* subject to certain conventions and guidelines, called *norms*. The norms define the legal limits of goals and actions, and can be changed if enough members of the alliance so decide. However, there are certain constraints and limitations, called *meta-norms*, that determine in what direction and how far the norms may change.

Alliances and norms are regarded as forms of emergence in distributed intelligent systems. Emergence is defined as the appearance of patterns at higher levels that are not apparent at lower levels. Multiple levels of emergence are possible; patterns of behaviour at intermediate levels may in turn constitute emergence at higher levels. The concept of emergence adopted here is somewhat broader than the usual definition, which requires that agents be unaware of the existence of potential for emergence (see also Chapter 8). Agents in distributed intelligent systems may, if they possess sufficient intelligence, be aware of the potentiality of emergence and strive to enhance (or destroy) emergence. One example of such behaviour is a stock market, where the individual agents (brokers) are aware of the potential for cumulative consequences of many similar individual acts; for example, a rush to sell or buy stock in a particular company has effects on its price, and sometimes on the market as a whole, that are not accounted for in terms of the effects of individuals' actions. In such cases, the emergence is within the "reasoning horizon" of individuals, since they are aware of the possibility of emergence and may try to influence it. However, it is outside their control horizon, since (in most cases) agents cannot influence it by their individual actions alone. Such phenomena, which are not controllable by individuals, are also considered in this chapter to be instances of emergence. Alliances and norms are emergent phenomena given this definition, as they are (in general) not under the control of any one agent. Also, alliances exhibit (in some respects) united action, where the individuals give up some or much of their independence in return for "safety in numbers".

The concept of the alliance replaces that of a single agent as the unit actor with a single locus of control and intention. One must keep in mind the following features that distinguish alliances from other forms of co-operation:

1. Alliances may differ from one another in all possible characteristics and have no a priori motivation to collaborate.
2. If two or more alliances decide to collaborate, it is based on their momentarily shared goals, and on the accuracy and timing of their "network perception" (their perceptions of other alliances).
3. The goals of an alliance may be self-generated or induced by the environment.
4. The members of a single alliance or different alliances may assume a temporary, hierarchical or flat organizational structure if it serves their purpose and their internal power conditions.

5. The membership and the norms of an alliance may change dynamically in response to the changing environment and the individual agents' mentality.
6. The actions of an alliance may affect the others adversely if their goals are partially or totally exclusive.
7. The costs and benefits of an alliance depend on the actions of the other alliances as well. It may, however, be a non-zero-sum environment.
8. Conflicts arising between alliances may (but need not) be resolved through negotiation.
9. The malfeasance of an alliance (deliberate or accidental transgression of the norms) is to be discovered and reported by others.
10. When the goal and/or action of an alliance are found to violate the current norms, their expected value to the social entities in question is calculated. If this is greater than a certain threshold value, a suggested change to the norms (allowing the goal or action in question) is checked *vis-à-vis* the meta-norms, and, if found to be feasible, put to a vote. A negative vote by the majority (or plurality) can be followed by negotiations leading to an acceptable consensus.

The features listed above are observed in many alliances in international relations, legislation, economic behaviour, and other areas of social activity. Not all the features are necessarily present in all alliances.

Our studies are intended to enhance the understanding of human organizations, the social and organizational role of knowledge and meta-knowledge, national and international markets, industrial problems, the effect of conflict and competition, trade-offs in co-ordinated efforts, and computer tools for the control of co-operation and collaboration.

The rest of this chapter concentrates on the concepts of alliance and norms. The next few sections present formal and informal descriptions of procedures for alliance formation, dissolution and concerted action by the members of an alliance. These sections are followed by a formal definition of norms and procedures for generating, evolving and applying norms. Finally, the use of these procedures in a single social domain is described.

Goal structure

The basis and motivating force of action is the description of the goal structure of agents and alliances. The following formal description of goal structures requires the definition of a *similarity operator*, a binary operator that accepts as its operands two goal structures and returns a numerical value characterizing the degree of similarity between them. In addition, the formal definition of a goal structure must be powerful enough to allow for the description of complex goal structures, that is, complex Boolean expressions using the AND and OR connectors between clauses. Accordingly, the structure of a goal is that of an AND/OR tree. The formal definition of the goal structure follows:

$G \rightarrow X$		$(A \wedge A)$		$(A \vee A)$
$A \rightarrow X$		$(X \wedge A)$		$(X \vee A)$
$X \rightarrow \text{string_a}$		string_b		...

where G =root goal; A =subgoal; X =indivisible/atomic goal; string_i = description of atomic goal.

In the above definitions, the root goal is the original goal of the agent or alliance, and an indivisible or atomic goal is one that cannot be subdivided into constituents. A goal structure is the hierarchy of goals and subgoals that compose the root goal. It is not the same as a plan, which in this context is defined as a series of actions, partially ordered over space and time, to accomplish goals. For simplicity, a goal in a goal structure can be an AND goal (a conjunction of its subgoals), an OR goal (a disjunction of its subgoals) or an ATOMIC goal (no subgoals). This does not affect the descriptive power since a mixed expression (containing both conjunction and disjunction operators) can easily be put into normal (canonical) form and expressed as a conjunction of disjunct clauses or a disjunction of conjunct clauses.

The next section briefly describes a proposed similarity operator for computing the similarity of goal structures as defined above.

The similarity operator

The similarity operator is a binary operator that returns the degree of similarity between its two operands. The operands are goal structures and the result is a real number between 0 and 1 (inclusive). The purpose of this and the previous section is not so much to define all-encompassing goal structures and a similarity operator that works well under all circumstances, as to demonstrate by example the fact that it is possible to define fairly general goal structures and a similarity operation around those structures.

The operator works by comparing the operands using a domain-dependent comparison mechanism. If this comparison fails to return an answer at the starting level, it performs a recursive descent into its arguments. This comparison and recursive descent cycle continues until the lowest levels of the structure are reached. Similarity values of parent nodes are computed as combinations of the similarity values of the child goals. The form of the combination depends on whether the parent goal is an AND or OR combination of the child goals. The operator is informally defined below.

The similarity computation relies on the *basic_match* function, which is a domain-dependent function that accepts as arguments two goal descriptions and returns a value between 0 and 1 (inclusive) characterizing the degree of similarity between the two descriptions, within the context of the domain of interest. The computation is based on syntactical considerations as well as the semantics of the descriptions in relation to the domain. For example, in a Blocks World domain, the Blocks World *basic_match* function should return 1 when comparing goal $G1$, stated as “the red ball is on top of the blue box” to goal $G2$, also stated as “the red ball is on top of the blue box”. Comparing $G1$ (stated as previously) to $G2$, stated as “the black ball is on top of the blue box” should return a value between 0 and 1, say 0.80. The definition of the *basic_match* function is

necessarily domain-dependent, since the degree of matching between two goals depends on the semantics associated with the descriptions as well as the descriptions themselves. For example, G_3 , stated as “the red ball is beside the blue box” should match G_1 to a lower degree than G_2 , as the relationship between the two objects is different. (Here we assume that position is more important than colour. If the opposite holds, the result of the comparison changes.)

The steps in computing the similarity between two goals G_1 and G_2 are:

1. The *basic_match* function is applied to the descriptions of G_1 and G_2 . If this results in a non-zero value for the degree of match, the value is returned as the result. If the match value is 0 and both G_1 and G_2 are atomic goals (that is, have no subgoals), 0 is returned.
2. If Step 1 fails to return a value, the goals G_1 and G_2 are compared with subgoals of each other. This handles the situation where one goal is a subgoal of the other. For example, “have a 600-ship navy” may contain “expand Portsmouth Naval Shipyard” as one of its subgoals. Each of these may be a top-level goal for different Representatives in the US Congress. If this step results in a non-zero match value, the value obtained is returned.
3. If Step 2 fails to return a value, the subgoals of G_1 are matched with the subgoals of G_2 , calling the similarity computation recursively, that is, starting over with Step 1 for each pair of subgoals. The values obtained are combined to give the match value for the parent goals. The form of combination is based on whether the parent goals are AND or OR goals, that is, whether the subgoals are conjuncts or disjuncts.
4. If all steps above fail, 0 is returned.

A more formal description of the similarity computation is beyond the scope of this chapter. Its definition tries to mimic the intuitive estimates of similarity between goals in most social situations, which are based on a deep understanding of the implications of various goals and their requirements in a social environment.

The definition of this operator requires precise descriptions of individuals’ goals. These descriptions must be written to include explicitly the dependence of the goals on agents; for example, if both agent A ’s and agent B ’s goals are to maximize their individual utilities, the goals must be expressed in forms that translate as “maximize A ’s utility” and “maximize B ’s utility”, respectively, and not simply as “maximize utility”. The first formulation allows the potential conflict between the two goals to be recognized, while the second does not. The alternative is to make the *basic_match* function smart enough to recognize such ambiguities and to make it consider the identities of the owners of the goal structures when computing the similarity of goal descriptions. These restrictions mean that a situation where one agent’s utility is other agents’ disutility is taken care of.

The rationale behind this approach to defining similarity is as follows. The computation of similarity depends on two things: first, the structural similarity between two different goal structures; and secondly, the degree of semantic match between the descriptions of those goals (in an artificial language invented for the domain, if necessary). The above formulation of the operator separates the semantic content of similarity, which is necessarily domain-dependent, from the structural content, which is not. Accordingly, application of this operator to another domain requires the replacement

of the semantic content (that is, the redefinition of the *basic_match* function), while the structural content can be carried over without change. Of course, this carrying-over requires that the goals in both domains can be expressed as goal trees. This requirement is not overly restrictive as the tree representation of goals is general enough to cover most goal structures. This approach to defining the similarity between goals is, we believe, as generic as possible for the problem.

Importance assignment

The importance of a goal is a number between 0 and 1 (inclusive) that gives the relative importance of the goal to the agent or alliance to which the goal structure belongs. Importance values are assigned to top-level goals based on the agent's or alliance's internal considerations as well as the external environment. The distribution of top-level importance to subgoals is described below. Each case corresponds to one way of dividing a root or intermediate goal, as described above.

$G \rightarrow X$	$\text{importance}(X) = \text{importance}(G)$
$G \rightarrow A_1 \ A_2$	$\text{importance}(A_1) = \text{importance}(A_2) = \text{importance}(G)$
$G \rightarrow A_1 \ A_2$	$\text{importance}(A_1) / \text{importance}(A_2) = \text{size}(A_1) / \text{size}(A_2)$ $\max(\text{importance}(A_1), \text{importance}(A_2)) = \text{importance}(G)$

The first case simply expresses the importance of an atomic goal in terms of the goal from which it is derived by means of the production rules that define a goal structure. The second case means that all AND subgoals of a goal are as important as the goal itself, since all are necessary to achieve it. The third case is for OR subgoals, and does two things: first, it prioritizes the subgoals in inverse order to their size and hence their complexity; and, secondly, it specifies that the importance of the highest-importance subgoal shall be equal to the importance of the parent goal. These conditions taken together mean that achieving the simplest subgoal has the same importance as achieving the goal G , while achieving more complex goals has a lower importance, since they are alternatives to the simplest subgoal.

The importance is used in the *basic_match* function to allow fine tuning of the similarity based on the degree of importance each agent attaches to the goal or subgoal. It is also used in selecting alliance goals.

Alliance goal structure

Alliances need to form an *alliance goal structure*, as this is what differentiates an alliance from a group of co-operating agents. Agents in an alliance work together to achieve a commonly decided goal or set of goals, rather than simply aiding one another to achieve their own goals independently of the others. The essence of an alliance (as distinct from a coalition) is giving up part of each member's independence in order to gain the benefits from joint action. Accordingly, some procedure is needed to construct an alliance goal

structure out of the goal structures of individual members. This chapter defines this procedure in terms of two subprocedures: one to select the most mutually beneficial member goal to be the initial alliance goal, and another that constructs an alliance goal structure “around” the selected initial goal. These procedures are now described informally.

The procedure to select the initial alliance goal is:

1. Choose arbitrarily any two agents from the alliance and combine all their goals (top-level goals as well as subgoals) into a rank-ordered set. The rank ordering is based on the degree of similarity between goals. Goals that are similar to many other goals are ranked higher than goals that are unique. The reasoning behind this method of ordering is that alliance goals are most likely to be those goals that contribute to furthering the aims of most members of the alliance.
2. Once a rank-ordered set has been created, the goals of other members of the alliance are added to the set one by one, using the same ranking procedure. In other words, for each goal or subgoal of an alliance member, a rank is computed based on how many goals it matches, either partially or completely. Furthermore, goals already in the set have their rankings updated, based on the degree of match between them and the new goals being introduced.
3. The result of Step 2 is an ordered set of goals that is independent of the order in which member goal structures were processed. This set contains all the goals in all the goal structures of alliance members. The set is ranked by *mutual benefit*, that is, the first goal in the set is the one that will benefit the members the most, by virtue of matching many goals in the goal structures of many members. This goal is selected as the basis for constructing the alliance goal structure.

All the agents determined to join a particular alliance do so simultaneously, after which the new goal structure of the alliance is computed. Working in this way, the membership of the alliance does not depend on the order in which individual agents join it.

Whenever an alliance is formed, or when it is reorganized, an alliance goal structure must be constructed from the goal structures of individuals or coalescing alliances. This construction, which is carried out by the second procedure, starts with the most mutually beneficial goal and adds subgoals and parent goals of this goal to the alliance goal structure one by one. While a goal is being added to the alliance goal structure, the importance of the goal to the alliance is calculated. The procedure that implements these ideas is as follows:

1. Let the highest-importance alliance goal, selected by the first procedure, be denoted by g_x . Add this goal to the (initially empty) alliance goal structure. It is assigned an importance value of 1 (the highest possible).
2. For all members of the alliance:
 - (a) find the location, if any, of the selected alliance goal g_x in the agent's individual goal structure; and
 - (b) add each subgoal g_y of g_x in the agent's goal structure to the alliance goal structure, as a subgoal of g_x .

3. Repeat Step 2 for each non-atomic goal in the alliance goal structure, until there are no non-atomic subgoals left. In other words, all descendants of the original goal are added to the alliance goal structure, until the subtree rooted at g_x is fully developed.
4. Add the parent goals of the originally selected goal g_x in a similar fashion, until either there are no more goals to be added or a specified importance cut-off has been reached.
5. Assign importance values to the constituents of the alliance goal structure, starting with the original goal (already assigned an importance value of 1):
 - (a) descendant goals of the original goal are assigned importance values calculated from the type (AND/OR) of their parent goal, as discussed in the previous section;
 - (b) ancestor goals of the original goal are assigned importance values based on their ranks in the ordered set computed by the first procedure described; and
 - (c) any goals in the alliance goal structure that have not already been assigned an importance value are assigned values based on the importance values attached to their parents, as described in the previous section.

This procedure constructs a goal structure that is the composite of the goal structures of the alliance members and is based on the most mutually beneficial goal, in the sense that no other goal has a greater importance value than this goal.

It is possible that two identical higher-level goals have somewhat different goal structures. For example, within the context of legislative action, representatives from different states may have the same top-level goal (such as “build a 600-ship navy”) but different subgoals (expanding naval shipyards in their respective home states). In this case, both subgoals will be placed in the alliance goal structure, and the potential contradiction between the two subgoals will be resolved within the alliance by an accepted voting process, so that the alliance goal structure may remain conflict free. (Recall that the members deciding to stay in the alliance are not allowed to pursue their own goals.)

Joining an existing alliance

The following algorithm is used to determine whether an agent joins an existing alliance. The rationale is that an agent should join an existing alliance if (enough of) its own goals will be helped along by the alliance goals. The degree of such helpfulness can be estimated by calculating the similarity between the agent’s own goals and the alliance’s goals.

Join(A, L) \rightarrow TRUE if and only if agent joins alliance L

1. $G1 \leftarrow$ current goals of agent A
 $G2 \leftarrow$ current goals of alliance L
2. MaxUsefulness $\leftarrow 0$
3. For each goal g_x in $G1$
 Usefulness(g_x) $\leftarrow 0$


```

For each goal  $g_y$  in  $G_2$ 
    Usefulness( $g_x$ )  $\leftarrow$  MAX( $g_x \otimes g_y, facilitates(g_x, g_y)$ )+Usefulness( $g_x$ )
endfor
TotalUsefulness  $\leftarrow$  TotalUsefulness+Usefulness( $g_x$ )
MaxUsefulness  $\leftarrow$  MAX(MaxUsefulness, Usefulness( $g_x$ ))
endfor
4. if (TotalUsefulness > TotalThreshold) (MaxUsefulness > MaxThreshold)
    return TRUE
else
    return FALSE
endif

```

Steps 1 and 2 initialize the goals and the potential helpfulness (*MaxUsefulness*) of the alliance to the agent. Step 3 compares each goal in the agent's goal set to each goal in the alliance's goal set and calculates two values: the overall potential helpfulness (*TotalUsefulness*), defined as the sum of helpfulness for all goals; and the maximum potential helpfulness (*MaxUsefulness*), the maximum extent to which any one of the agent's goals is helped by the alliance. If either of these is above a threshold, the agent decides to join the alliance. This amounts to using two criteria: (i) the overall expected utility of joining the alliance (*TotalUsefulness*); and (ii) the expected utility for a single, very important goal (*MaxUsefulness*). The extent to which one goal helps along another is obtained using the *facilitates* calculation, described in the next section.

Calculation of facilitation

This calculation is used to determine the utility of one goal in facilitating another goal. A goal is said to *facilitate* another goal when achieving the first goal helps the achievement of the second. For example, achieving a subgoal facilitates the achievement of that subgoal's parent. The degree to which one goal facilitates another is measured in terms of the number of steps required to get from the first goal state to the second. The algorithm is constructed so that the *facilitation* is a value between 0 and 1, with 0 corresponding to no help at all (or insignificant help) and 1 corresponding to direct logical implication:

$facilitates(g_1, g_2) \rightarrow$ degree to which goal g_1 facilitates goal g_2

This is calculated by computing successive action-closures of g_1 (the sets of propositions that become TRUE after successive possible actions), starting with the state in which g_1 holds. The computation of action closures is ended after a limited number of steps, at which point the *facilitation* value can be assumed to be low enough to be ignored. (In principle, *facilitation* could also have negative values, corresponding to a "counter-

effective” goal pair. This ramification is not explored in this chapter.) The algorithm for calculating F , the value of facilitation, is as follows:

1. $F \leftarrow 1.0$
2. if g_1 implies g_2 by direct logical implication
 - then
 - return F
 - else
 - $F \leftarrow F/2$; let $CL \leftarrow \{g_1\}$;
- endif
3. while ($F > \text{LowThreshold}$)
 - $CL \leftarrow \text{closure}(CL, 1)$;
 - if CL contains or implies g_2
 - then
 - return F
 - else
 - $F \leftarrow F/2$
 - endif
- endwhile
4. return 0

Some explanations are needed. The set CL is the action-closure of g_1 ; that is, the set of states which can be reached from g_1 with 1, 2, 3,... successive actions. Steps 1 and 2 initialize the values and check whether g_1 in the current state directly implies g_2 (that is, without any actions needing to be taken by agents). Informally, step 3 tries to find out how much effort is needed to achieve goal g_2 starting from g_1 (and the current state of the environment). This effort is measured in terms of the number of steps required to achieve g_2 . The rationale is that the more steps are required to achieve g_2 (given g_1 in the current environment), the less g_1 facilitates g_2 . The division factor 2 by which facilitation values are reduced in successive steps and the cut-off threshold, *LowThreshold*, below which facilitation is considered to be insignificant, are both arbitrary values.

Forming a new alliance

The following procedure describes how alliances are first formed.

1. For each pair of agents, calculate the similarity between their goals.
 - This gives a *proximity matrix* for all known agents.
2. From all pairs of agents with similarity $> \text{JOIN_THRESHOLD}$

select the closest pair.

3. For each agent i not currently in an alliance
 - if the similarity to the plurality of agents in the alliance is greater than JOIN_THRESHOLD
 - add agent i to the alliance
 - endif
- endfor

This procedure is analogous to the complete linkage method of constructing a cluster (Findler 1991, Kaufman & Rousseeuw 1990). It results in a fairly tightly-knit group of agents whose goals are relatively similar. An alternative method would be to form one single alliance initially, consisting of two members whose goal structures are the closest to one another. This procedure can also be applied to the merger of alliances to form a larger alliance. In this case, the goal structures used will be alliance goal structures.

It should be noted that Step 3 does not involve construction of the alliance goal structure immediately upon the agent's joining the nascent alliance. The alliance goal structure is constructed after the entire alliance has been formed, using the procedure described above.

Costs of joining and leaving

Under certain circumstances, a penalty may be levied for leaving an alliance and a reward given for joining one. This can be implemented in those cases where agents have an account, in the sense of a bank account, or currency/means of exchange, as in a stock market scenario, by levying a suitable penalty on leaving the alliance or giving a suitable reward for joining the alliance.

Some kinds of alliance may require that a joining agent deposit some of its own resources into a common pool of alliance resources. In that case, the penalty for leaving may consist of loss of access to all or some of the resources in the alliance pool, even those deposited upon joining. Other kinds of alliance may either allow the departing agent to take with it a proportion of the resources it deposited, or a part of the resources currently in the alliance pool.

Other conditions for alliance formation

Other possible conditions for alliance formation include strategic alliances, where agents join forces for a temporary purpose such as defeating other agents, even when they do not have common goals *ab initio*. Our contention is that even strategic alliances may be explained on the basis of goal similarity testing, as the joining of forces is caused by the recognition of a common goal, given by the purpose of the alliance. The case where two agents (with different individual goals) join forces to defeat other agents can be explained

by noticing that they do in fact have a common goal: defeating the other agents. The goal-based alliance formation described in this chapter does not require that the agents' goals stay the same throughout; the individual goals are allowed to change in response to changes in the environment, actions of other agents, improved knowledge, or negotiation with other agents.

Definition and description of norms

Norms are restrictions, limitations or constraints imposed on the actions and goals of individual agents or alliances. In general, a norm is expressed as a tuple consisting of the constrained goal or action and a description of the constraint. For example, in a Blocks World situation, one norm could be the limitation:

((*Block_B* on_top *Block_A*) forbidden)

That is, *Block_B* cannot be placed on top of *Block_A*, the reason being that *Block_A* is made of jelly and *Block_B* of lead, and placing *Block_B* on top of *Block_A* would destroy it. It is not necessary for the system to know the reason for the norm, though there will be norms for which a reason might be known, especially norms generated by the agents and alliances themselves. Any known reasons for norms may be used during the process of norm evolution.

In addition, a norm may be expressed in general rather than specific terms. For example, the above norm may be expressed as

for $\forall y$ ((substance x jelly) (y on_top x) forbidden)

which means *do not place anything on the top of any block made of jelly*.

A more realistic example is seen in the filibuster in the US Senate in April 1993 over the stimulus bill proposed by the administration of President Clinton. The essential characteristic of a filibuster is the blocking of all proposed legislation through continual debate. In its basic form, a filibuster blocks *all* legislative action until either the filibuster is broken or the bill is dropped. This is indeed the way filibusters have been carried out until recently. However, several rules and conventions have been developed to reduce their negative impact. The current convention is that a part of the day is set aside for other business by agreement between both parties, even during the course of a filibuster, although this is not required by Senate rules. The norm here is that a filibuster of one piece of legislation no longer involves the blocking of all other pieces of legislation, by common consent. This norm may be expressed formally as:

for $\forall y$ ($(\neq xy)$ (filibuster x) (\neg (filibuster y)) (blocked y) forbidden)

which amounts to saying that during a filibuster of one piece of legislation, any other piece of legislation that is not being filibustered is not blocked.

The tuple form of norm description can be extended easily from describing only forbidden actions/goals to other types of limitation. For example, it can be extended so that a goal that satisfies the precondition of the norm (the first element of the tuple) could be reduced or increased in priority relative to other goals. When this generalization is

made, the norms are clearly no more than a set of rules that are applied to actions and goals, and the norm-related computation is a kind of expert system selection mechanism (filter) applied to the actions/goals (Giarratano & Riley 1989).

For the purposes of this chapter, *norms* are defined only within the context of a particular alliance, that is, a particular norm applies only to the individual members of a particular alliance. *Meta-norms* are constraints and limitations that are functionally similar to norms, but apply over the entire environment; that is, they apply uniformly to all the agents and alliances in the system. Meta-norms correspond, for example, to humanitarian, ethical and moral standards in the real world. Norms correspond to regulations aimed at improving efficiency and adherence to the meta-norms.

One significant difference between the application of norms/meta-norms and an expert system is that the satisfaction of a norm or meta-norm by a set of goals need not mean that it is in fact followed—it might be violated, or the agent might choose to release itself from the norm by leaving the alliance. The violation of a meta-norm may come about through an explicit decision by the alliance to violate it when the payoff is high enough, where the alliance is willing to risk discovery and also risk the penalty for the violation. Norms, on the other hand (which apply to members of an alliance) are not violated by the members. Instead, agents that find the costs of staying in the alliance to be too high may choose to depart from the alliance, paying the costs of departure, if any. Another difference from expert systems is that the norms apply only in certain circumstances, that is, (in general) only to members of a single alliance.

The rest of the discussion assumes only one kind of norms and meta-norms: those that forbid certain actions/goals.

The application of norms

Norms are tested (by agents) whenever an agent or alliance creates a new goal or considers a new action, most often during the formation of a plan. Testing a norm consists of matching its antecedents against the current environment, the agent/alliance goals and actions, and any other goals or actions known to the tester. Norms that are triggered during the testing phase result in the tester being informed about the anticipated violation. The agent may choose either to violate the norm or to follow it by withdrawing the action/goal that triggered the potential violation. Note that agents can only test goals known to them; that is, their own actions/goals, the goals of the alliance they belong to, and any other goals that they may have been informed of (Findler & Gao 1987, Findler & Ge 1989, 1993, Findler 1989). The actions of other agents and alliances are known only after execution and cannot be tested during this phase.

Norm generation

The purpose of norms is to enable more efficient functioning of the distributed system and to prevent harmful events. *Efficiency* and *harm* in this context are measured by criteria that are important to the particular system. For example, in a legislative context, one of the criteria for determining *efficiency* is the time taken to pass controversial

legislation. Harmful events are also defined in a context-sensitive manner. In the real world, *harm* is measured in terms of transgressions of humanitarian, ethical and moral standards. For example, the destruction of an agent (say, by reducing all its resources to zero) may be considered to be a harmful event in some distributed intelligent systems. This is implemented in practice by providing a list of harmful events (in much the same format as the description of a norm) while setting up the system. In other words, the concepts of harm and efficiency are domain-dependent. They may be measured in terms of their effects on the payoffs to individual agents. Because in the simple case considered in this chapter, the agents are *selfish individuals* that try to maximize individual payoffs, the sum of individual payoffs constitutes the global payoff.

Thus, norms must arise from experience; that is, they must be formed after a harmful event or inefficiency is detected. In other words, a norm is in a sense a *new concept* that must be learned from examples or non-examples/ counter-examples. This means that a record of actions/goals and their consequences and payoffs must be kept. Whenever a harmful event occurs, the log is examined and a norm constructed (for example, as a conjunction of the actions/goals immediately preceding the harmful event). Deciding which of the immediate precedents of a harmful event caused the event is a difficult problem in itself, and will be simplified for the purposes of this research by depending on an omniscient black box to construct a norm from the log by deriving causal relations between events and their precedents. This is adopted as a conceptual simplification; any implementation of this black box is necessarily domain-dependent. Note that in this context, omniscience does not mean infallibility. The black box may (and does) come up with norms that are wholly or partially wrong. Norms are suggested by an individual agent (this means that each agent is capable of deriving and evolving norms). A norm, once suggested, may be negotiated with other agents and subjected to a voting procedure to decide its acceptance. The only method of norm generation used at present is direct forbidding of the action in question. This method is more general than it seems, and is often seen in human affairs (for example, laws are often devised that directly forbid specific actions that are regarded as harmful to society).

The norms generated in this phase are evolved, ideally towards perfection, by the agents. Figure 11.1 shows the process whereby norms are generated.

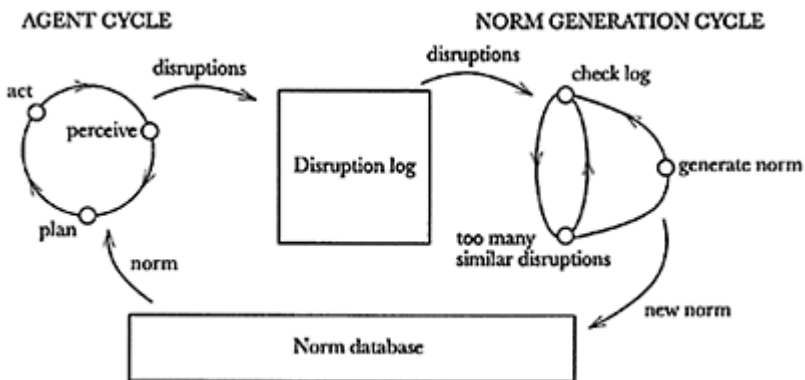


Figure 11.1 Schematic showing the generation of norms.

The evolution of norms

Three kinds of evolution are possible:

- (a) addition or removal of terms/clauses;
- (b) binding or unbinding of variables (that is, specialization or generalization); and
- (c) changing the values of numerical constants. Each change is triggered by an unpenalized violation, or by a harmful non-violation that came close to triggering a norm, or the recurrence of a harmful event that a norm was supposed to prevent. The process of changing a norm is described below.

1. The change is proposed by the discoverer/initiator to all other agents in the same alliance.
2. Weighted votes are received from other agents, ranging from $-D$ to $+D$, depending on how strongly the agent disfavours/favours the change. *Zero-votes* (*don't-care* votes) are allowed, as are *not-voting* (that is, a *non-response* is given by an agent). Both *don't-care* votes and *not-voting* are handled similarly as far as tallying votes goes. The only difference between the two is that *don't-care* votes can be processed immediately while *not-voting* must be handled by implementing a time-out.
3. The votes are tallied and the change is made or not made, depending on the outcome. All agents are notified of the outcome.
4. The changed norm is checked against the meta-norms. If it passes, it is acceptable; if not, it is rejected.

The most important part of the process is Step 2, in which an agent's vote is decided and weighted, depending on how strongly the agent favours the change. The algorithm is:
Algorithm: Calculate B , the benefit of a proposed change, and return a vote weighting.

1. $B \leftarrow 0$
2. For each element e (action or goal) in recorded history
 - if e is passed by old norm and blocked by proposed norm
 - subtract actual payoff from B
 - else
 - if e is blocked by old norm and passed by proposed norm
 - add expected payoff to B
 - else (e treated similarly by both norms)
 - ignore e
 - endif
- endif

3. return *B*

Discovering violations

Violations of norms are discovered by an agent during its *perceive* phase, by matching the most recent actions of other agents from the same alliance or actions of other alliances, to the existing norms (in the case of members of the same alliance) and meta-norms (in the case of other alliances). Penalties are levied on violators in the form of

- (a) downgrading the importance of its goals in the alliance goal structure, perhaps after holding a referendum;
- (b) expulsion from the alliance; and
- (c) retaliation directed at the offending agent or alliance.

A domain of application

The first chosen domain for experimentation is international politics, in particular, post-Second World War international relations in North Africa. Within this domain, we study the phenomena of alliance formation and dissolution related to the political alliances between nations of the region and the superpowers interested in the region. The general goals of these alliances can be identified as follows:

- 1. Deterrence of aggression by other countries/alliances.
- 2. Defence against military moves by other countries/alliances.
- 3. Alliance unity in political and military action.
- 4. Enlargement of the alliance.
- 5. Reduction in the membership of competing or hostile alliances.
- 6. Non-proliferation of nuclear and certain other weaponry.
- 7. Control of escalation in crises.
- 8. Stabilization of crises.
- 9. Stabilization of arms race in general.
- 10. Control of costs.
- 11. Elimination of the possibility of surprise attacks.
- 12. Reduction of extremism and terrorism.
- 13. Gaining economic advantage in trade.
- 14. Influencing the public opinion in the other alliances.
- 15. Adjustment of the internal political pulls within each nation to increase the influence of the group currently in power.

Of the above, some goals (deterrence, defence, etc.) are noted by Kugler (1990). Within these general goals, each alliance or agent (nation) will have its own goals, which may conflict partially with the goals of another agent or alliance.

The questions addressed are:

1. Is it feasible to model the formation and progress of these alliances as a distributed intelligent system?
2. Would these alliances form and evolve just the same as they have in reality?

Since the scope of the domain is far too large for reasonable modelling efforts, our first experiment is restricted to a small number of major events rather than modelling the entire domain over the period of interest.

An experiment

The experiment described in this chapter concerns the Soviet–Libya relationship. For the purpose of this experiment, the salient points are:

1. The goal structure of both parties.
2. The value of each party to the other.
3. Considerations determining the selection of the other party instead of other candidates.

For both parties, the goal structure consists of a main goal and auxiliary component or subgoals. For each party, the overall goal is to enhance its own status. The sub- and auxiliary goals derive from this overall goal. For the Soviet Union, the auxiliary goals deriving from its root goal (enhancing its own power) are: (a) enhance status without regard to ideology; and (b) enhance status with regard to ideology. (History has shown that these two were at times quite distinct and even in conflict.)

The first derives from superpower compulsions and the second from the ostensible ideological background of the Soviet Union. These second-level goals can be further subdivided into gaining (a) military, (b) political, and (c) economic advantage over the Western powers.

Of these, only the political advantage subgoal is considered to apply to the second-level goal of gaining advantage with regard to ideology. (This restriction is introduced for purposes of simplicity, and is not considered a general constraint.) These third-level goals are composed of the following subgoals:

1. *Gain military advantage.* This can be further subdivided into: (a) gaining military facilities; (b) improvement of military discipline among allies; and (c) getting a positive military effect on the neutral countries.
2. *Political advantage.* This can be further subdivided into: (a) gaining political friends, which includes favourable votes in international fora; (b) improving political coherence in the Soviet camp; and (c) projecting a positive political image to neutral countries.
3. *Economic advantage.* This can be further subdivided into: (a) gaining economic benefits, such as raw materials at favourable prices; (b) improving economic coherence in the Soviet camp, with a view to increasing trade under favourable

conditions; and (c) projecting a positive economic influence on neutral countries, with a view to drawing them into its sphere of political influence.

For Libya, the goal structures are similar, although the relative importance given to each component varies and the goals specific to superpower concerns are removed. The descriptions are also changed to match Libya's concerns, for example, *expanding economic influence in the Soviet camp* becomes *expanding economic influence in the Arab bloc and in African countries*.

These goal structures are encoded as in Figure 11.2. Here, each line represents a goal or subgoal. The structure of the goal tree should be clear from the indentation. The first word after the left parenthesis indicates whether the goal is an AND goal (a conjunction of its subgoals), an OR goal

(or enhance status of s-u)

(and regardless of ideology)

(or military advantage)

(atom military facilities vs western powers)

(atom improve military discipline soviet-camp)

(atom positive military effect on third world)

(or political advantage)

(atom gain political friends)

(atom improve political discipline soviet-camp)

(atom positive political effect on third world)

(or economic advantage)

(atom gain economically profit)

(atom improve economic discipline soviet-camp)

(atom positive economic effect on third world)

(and with ideological coverage)

(or political advantage)

(atom support favorable organizations)

(atom advance propaganda aims)

(atom cause factions in opposing parties)

Figure 11.2 Goal structure of the Soviet Union.

(a disjunction of its subgoals) or an ATOM goal (recall that an ATOM goal is one that is not further subdivided). The rest of each line describes the goal.

The value of each party (Soviet Union and Libya) to the other is determined by assessing relevant factors using numeric values (positive for beneficial effects, negative

for detrimental factors). These assessments were assigned arbitrarily (though realistically).

For example, one of the advantages to the Soviet Union of getting military facilities in Libya is access to military bases in the Mediterranean, another is profits from weapons sales, a third is military intelligence, and so on. These can be broken down into short-, medium-, and long-term goals. For the Soviet Union, short-term goals include economic treaties, monopolization of supplies (including weaponry), treaties of "friendship and cooperation", etc. Medium-term goals include naval and other bases and facilities, military advising, and cultural and economic exchanges. Long-term goals include the setting up of a base for exerting Soviet influence over neighbouring countries by various means, including "revolutionary wars", and so on. For Libya, the short-term goals include arms supplies, treaty benefits, and industrial supplies. Medium-term goals include transfer of technological expertise and availability of a supplier for essential materials. Long-term goals include increasing influence over other African and Arab countries.

The benefits deriving from possible alliances with other parties in the region were encoded in a similar fashion. These benefits are considered in deciding with which party an alliance should be made. The experimental set-up is constructed as an event-driven simulation. Events are read from an external file (or from user input) and processed according to the internal action rules of the agents. Norms are generated and evolved, depending on the consequences of the actions chosen by the agents. These are governed by rules that describe the evolution of the norm based on the existing norms and the consequences of actions. The system as it stands at the present time generates an alliance between the Soviet Union and Libya based on the mutual benefit to each of the other's goals (the similarity value of the individual goal structures). The norms and evolution rules are largely in the form of an embedded expert system coded in CLIPS (Giarratano & Riley 1989). At present, the only norms generated by the system relate to "hiding" provocations and invasions of neighbouring countries; that is, they relate to the military features of goals. Efforts to expand the scope of the system by increasing the number and variety of the norms and the range of their evolution are under way at the time of writing. For example, in the system, the reaction to an economic crisis in Libya is an invasion of Chad, leading to a penalization of the country by its neighbours, which in turn results in the alliance generalizing the norm to suspend invasion of neighbouring countries.

Related work

Some research on coalition forming has been described in the literature on group behaviour (for example, Mokken & Stokman 1985). In relation to coalition forming in the legislative and international domain, on the other hand, there is significantly more literature, especially concerning voting coalitions in legislative assemblies (for example Cigler & Loomis 1986, Davidson & Oleszek 1985, Mokken & Stokman 1985, Oleszek 1989). Most of the literature appears to be aimed at describing the phenomenon of coalition formation. Our research tends to be prescriptive rather than descriptive in that we attempt to arrive at a framework for the process of alliances and coalition formation in algorithmic terms rather than capture the individual features of coalitions/alliances in a particular domain.

Bond (1991) describes *commitments* as “a constraint or binding” on agent actions, states and beliefs. This concept includes the development of social relationships that lead to constraints on agent activities. The constraints include commitments to act/believe in a specified manner and also constraints on actions and states. This notion of commitments as constraints is similar to the concept of norms presented in this chapter. The fundamental difference between Bond (1991) and this chapter lies in the use made of the concept. Bond uses the concept of a commitment as a social relationship forming a basis for agent organization, where the commitments form constraints on the organizational units. The research described in this chapter considers the organizational units as alliances that have formed in order to achieve common goals. The work on norms is aimed at studying the generation and evolution of such commitments or norms, rather than regarding them as the basis of alliance formation.

Conclusion and future objectives

The proposed system provides a fairly realistic model of alliance-forming behaviour in the domain of international politics. (We may develop an encompassing computer-based theory of Soviet foreign policy that we might call “Gromyko on a microchip”.) It is hoped that such studies will provide a tool for the better understanding of social systems, including international political relationships.

Future work will include co-ordinated action by the alliances and the inclusion of referenda to revise goal structures at intervals. The effects of rewards for joining and costs of leaving an alliance are also being studied, as is the application of meta-norms to the evolution of norms. The work is expected to include extension of these ideas to the legislative domain, especially related to coalition forming in legislatures.

Acknowledgements

The authors wish to thank the anonymous referees for their comments and suggestions.

Part III
**The foundations of social
simulation**

Chapter 12

Kin-directed altruism and attachment behaviour in an evolving population of neural networks

Domenico Parisi, Federico Cecconi and Antonio Cerini

In early human societies social groups were mainly based on space. A social group was a set of individuals who tended to live close together in the same local environment. Physical nearness made it possible to interact socially and to develop social structure inside the group. It is only with the development of transportation technologies and, especially, the huge development of communication technologies in the twentieth century, that space has lost some of its importance in determining who is interacting with whom and, therefore, the nature of human groups.

A second property of early social groups was their biological basis. The social sciences have a tendency to ignore the biological infrastructure of human societies but we know that these societies have evolved from a state in which human groups were essentially families, i.e. sets of individuals who were genetically related or related because of sexual reproduction (Johnson & Earle 1987).

In his 1992 presidential address to the American Sociological Association (Coleman 1993), James S. Coleman remarks that the discipline of sociology was born at a time when Western societies were changing rapidly from what he calls “primordial societies”, based on locality and kin relations, to modern societies, in which social relationships tend to be independent of space and biology. Sociology has concentrated on modern societies and it has embraced the point of view where in order to understand X it is not necessary to reconstruct how X has come to be X , i.e. its origin and development.

Our preferences go to the opposite (genetic) viewpoint (see Parisi in press), according to which contemporary human societies cannot be understood unless one examines how they developed from earlier—in fact, even non-human primate—societies in which space and biology had a more important role. The simulation of human societies, for us, is the simulation of an evolutionary or historical process of change. More specifically, we are concerned with simulating how human groups changed from being small-sized groups largely based on space and kin relations to being progressively larger groups in which other causal factors beyond space and biology kept the group members together.

In this chapter we examine two elementary forms of social behaviour which appear to play a role in creating simple social groups in organisms, both human and non-human, i.e. attachment behaviour and kin-directed, genetically-based altruism. Attachment behaviour is the behaviour of an organism that tends to keep it near another organism. Hence, attachment behaviour is one basic mechanism for maintaining space-based social groups. Kin-directed altruism is the behaviour of an organism which increases the fitness, i.e. the reproductive chances, of a kin while decreasing the fitness of the organism

exhibiting the altruistic behaviour. Attachment behaviour and altruism are related and are, in fact, complementary behaviours. If it is the case that the organism which is the potential recipient of an altruistic act can only obtain the advantages emanating from the act if it is near the altruist, then there will be a tendency for attachment behaviour to emerge through evolution. In the simulations we shall describe below, offspring exhibit a tendency to follow their parents in their wanderings in search of food because offspring can only receive food (and survive) if they remain close to their parents, who feed them.

It is more complicated to explain how altruistic behaviour may emerge and be maintained in a population of organisms. Since altruistic behaviour by definition decreases the fitness of the altruist, individuals with altruistic tendencies will tend to have fewer descendants than do non-altruists and therefore they should in the long run disappear from the population. Kin-selection theory (Hamilton 1964, Trivers 1985) has been introduced to explain why, contrary to this expectation, genetically-based altruism is observed in many organisms. According to this theory, altruistic behaviour will be maintained in a population if it is directed towards one's own kin. The fitness of the altruistic individual will be decreased by the altruist act but the fitness of the kin individual who is the recipient of the act will be increased. Since kin individuals are genetically similar and are therefore likely to exhibit the same behavioural tendencies, altruistic behaviour present in the reproductively unsuccessful altruist will also be present in the reproductively successful kin towards which the altruistic act was directed. Hence, whenever the increase in fitness of the recipient of the altruistic act exceeds the altruist's loss of fitness, altruistic behaviour will emerge and be maintained in an evolving population.

The present chapter is divided into several sections. In the next section we describe some simple simulations whose purpose is to demonstrate that kin-directed altruism can emerge through evolution and be maintained as a stable trait in a population of artificial organisms. In the following section we describe more complex simulations of populations of artificial organisms living in a shared environment. In these simulations not only does genetic altruism of parents toward their offspring emerge, but also offspring develop a tendency to follow, i.e. to remain in proximity to, their altruistic parents. Therefore, in these simulations we observe the emergence of a simple kind of social group composed of individuals that live in the same local environment and exchange simple acts of following and giving food. In the final section, we discuss how to extend the range of phenomena that can be simulated along the same lines.

The evolution of kin-directed altruism in a population of neural networks

The purpose of the simulations to be described in this section is to demonstrate that some form of kin-directed altruism can evolve in populations of simple artificial organisms. The organisms of the simulations of this and the next section are modelled by neural networks. (Alternative, e.g. rule-based, approaches might also be used but we shall not discuss the reasons for our choice of a neural network model.) We assume that there is a population of such networks that reproduce selectively on the basis of a criterion of fitness. During its life, each network has a fixed number of social encounters in which it

must decide whether to act altruistically or egoistically with respect to its current social partner. Each network has two input units, three hidden units, and one output unit. The two input units encode the type of the current social partner which can be either a sister (10) or a non-sister (01). A sister is another network which is an offspring of the same parent of the network we are considering. (In all simulations reproduction is agamic, i.e. a single parent generates one or more offspring.) The single output unit encodes whether the network decides to give (1) or not to give (0) a piece of food in its possession to its current social partner. At the end of life the networks with the largest number of pieces of food (fitness) reproduce, while the other networks die without leaving offspring.

The initial population is composed of 100 networks divided into 20 groups of five sisters (identical networks). Each group of sisters is assigned a different set of connection weights but all the sisters in a group share the same weights. All networks live the same number of cycles. At the beginning of life each individual is given the same number of pieces of food. In each cycle the network has a different social partner randomly selected from the population. Hence, at any given cycle the social partner may be either a sister or a non-sister (with different probabilities). The network must decide whether to act altruistically by giving its current partner one of its pieces of food or to act egoistically by refusing to give food. If a piece of food is given, the donor's food counter is decreased by one unit while the social partner's counter is increased by one unit. Otherwise, nothing changes.

At the end of life, which lasts 100 cycles (i.e. 100 social encounters) for all individuals, the 20 individuals with the highest number of pieces of food reproduce by generating five copies of their matrix of connection weights, with random mutations added to some of the weights. The other individuals die without leaving descendants. The offspring networks constitute the second generation. The second generation also consists of 20 groups of five sisters each but, unlike the first generation, sisters now have similar but not identical connection weights because of the random mutations. The process continues for 50 generations.

Since individuals reproduce as a function of the number of pieces of food they possess at the end of life, one would expect individuals who always act egoistically, that is, who never give away food, to emerge and colonize the entire population. These individuals keep their own food and may occasionally receive additional food from altruistically-acting individuals. However, as predicted by kin-selection theory, what is in fact observed is a more complex type of behaviour (all results are based on five replications of each simulation, each starting from a different "seed" of a random number generator). Individuals refuse to give their food to non-sisters but they tend to give it to sisters. In other words, they act egoistically towards non-kin but altruistically toward kin. The loss of fitness that is caused by giving food to sisters is compensated by the increase in fitness of the sisters who receive food. Hence altruistic traits, i.e. connection weights which result in giving food, are maintained in the population via reproducing sisters (inclusive fitness). However, altruism is exhibited only towards sisters. Non-sisters are treated egoistically (Figure 12.1).

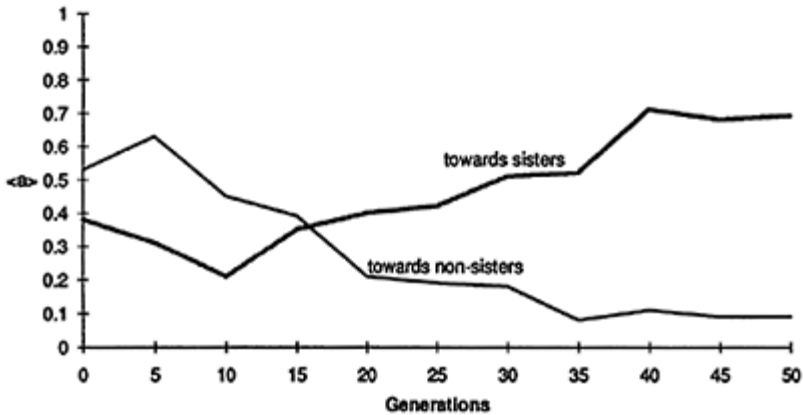


Figure 12.1 Proportions of altruists toward sisters and non-sisters across 50 generations.

The evolution of attachment behaviour and kin-directed altruism

In a new set of simulations we used conditions that were somewhat more lifelike than those used previously. Neural networks live in an environment that has its own independent structure. They are ecological networks (Parisi et al. 1990). The environment is both social and non-social. It includes randomly distributed pieces of food (non-social environment) and conspecifics (social environment). Furthermore, unlike the preceding simulations, in these simulations generations overlap. Hence, parents may share the environment with their offspring. Since food is eaten constantly and it disappears from the environment, it is reintroduced periodically. The simulation lasts for 20,000 time units (cycles) with all living organisms (networks) updated at each time unit.

Individuals have a lifespan (total number of cycles), which can vary from individual to individual. If an individual spends too many cycles without receiving energy from outside, the individual dies. The life of an individual is composed of two successive phases: young and adult. When an individual is young it cannot reproduce and can only obtain energy (and, therefore, survive) if its parent gives some of the energy it has obtained by capturing food in the environment. When an individual becomes an adult, it can only survive (and reproduce) if (a) it is able to capture food in the environment; and (b) it decides to eat the food instead of giving it to its offspring. Hence, an adult individual has a choice between eating the food (thereby increasing its chances of survival and reproduction) or giving the food to its offspring (thereby increasing its offspring's chances of surviving and reaching adulthood).

More specifically, at the beginning of life all individuals are assigned the same quantity of energy. The energy decreases by one unit at each cycle. If the energy reaches 0, the young individual dies. The only way for the young individual to increase its energy is by receiving food from its parent. If a young individual survives, at a certain age it

changes its status and it becomes an adult. At the beginning of adulthood the individual is endowed with a certain amount of energy which is equal to its parent's current energy. This energy is also decreased by one unit at each cycle. However, by capturing food in the environment an adult can increase its energy. An adult generates a new offspring at regular intervals (number of cycles) during its life. The individual is also asked regularly whether it wants to give one unit of its energy to its offspring or it wants to keep all the energy for itself. By deciding to keep its energy for itself (and not to give it to an already existing offspring) an adult not only makes its own life longer but also increases its chances of generating further offspring.

The behaviour in the environment of both adult and young individuals is controlled by a feedforward neural network. Neither the network's architecture nor the network's connection weights change when a young individual reaches adulthood. What does change is what is encoded in the network's input (sensory) units. Each network has four input units, six hidden units, and two output units. When an individual is young, two input units encode the current position of the individual's parent relative to the young. One unit encodes the angle of the parent (measured clockwise and mapped in the interval between 0 and 1) with respect to the offspring's facing direction. The other unit encodes the distance between parent and offspring (also mapped in the interval between 0 and 1). The remaining two input units are set to a fixed value of 0. When the individual reaches adulthood, the first two units are set to 0 and the other two units encode the angle and distance of the nearest piece of food. The two output units, in both young and adult, encode four possible motor actions: go one step forward (11), turn 90° left (10) or right (01), do nothing (00).

Although both young and adults move in the environment on the basis of environmental sensory information, they are informed by their senses about different aspects of the environment. Young individuals are informed about the current position of their parent. Adults are informed about the position of the currently nearest piece of food. Young and adult individuals must use this environmental information in different ways. The objective of young individuals is to receive food from their parents because this is the only way for them to survive. Their parents have part of the responsibility because if they are unable to capture food or if they decide not to give their food to their offspring, the offspring will not eat. However, another part of the responsibility for eating lies with the young themselves. The reason is that a young individual can receive food from its parent only if it is located sufficiently close to its parent. Otherwise, no giving of food from parent to offspring can take place. This implies that young individuals must use the environmental (sensory) information about the current position of their parent to generate motor actions that keep the distance between them and their parent within the required radius. Since their parent in the meantime is moving itself in the environment in order to find food, they must follow their parent.

For adults the situation is different. Adults must use environmental (sensory) information about the location of food to move in the environment so as to reach the food in as few cycles as possible. (When reached, a food element disappears from the environment.) Then they must decide whether to use the energy resulting from the captured food for themselves or to give the food (energy) to their offspring, assuming the offspring is at the appropriate distance. (If an adult has more than a single young offspring, it is the nearest offspring which receives the food.)

What decides whether a parent will act egoistically or altruistically towards its offspring? We assume that each individual inherits from its parent not only the connection weights of its network (these weights determine how good an individual is at following its parent when young and at capturing food when adult) but also a special parameter (“gene”) which determines the probability of an individual acting egoistically or altruistically as an adult. Hence, different networks are born as more or less altruistic individuals. If the inherited value of the altruism “gene” is high, when the individual becomes a parent it will tend to behave altruistically and give its energy to its offspring. If the inherited value is low, the individual will tend to act egoistically and consume all the energy derived from the captured food.

It is clear that an intermediate, optimal, value of the “gene” must be found if the population is not to become extinct. If the average altruism of the population is too pronounced, each individual will generate very few offspring and for this reason the population will be at risk of becoming extinct. If the population tends on average to act too egoistically, the average individual will generate many offspring but these offspring will tend to die (of starvation) before reaching adulthood and reproducing; in this case too the population runs the risk of extinction. An intermediate level of altruism means that the average adult will eat enough to survive and generate new offspring and, at the same time, will give enough energy to its existing offspring to allow them to survive to adulthood and reproduce.

In the simulations we are describing, the appropriate level of altruism evolves spontaneously in the population as does any other trait and does not need to be determined a priori by the researcher. Any trait can evolve in a population if (a) it varies from one individual to another; (b) it is inherited; and (c) it is subject to random mutations. All three properties characterize both the weight matrices and the value of the altruism “gene”. Hence, we can observe how the appropriate level of altruism evolves spontaneously in the population. Figure 12.2 shows how the average tendency to behave altruistically changes across the generations.

Altruism starts at around 0.5 due to the random assignment of values (between 0 and 1) to the altruism “gene” of individuals of the initial population and it quickly increases to stabilize at a value between 0.8 and 0.9.

The evolved behaviour which is of particular interest here, however, is the tendency of young individuals to follow their parents during the latter’s movements in search of food. In fact, in these simulations it is possible to observe the emergence of a type of behaviour which tends to keep different




Figure 12.2 Evolutionary change in the average value of the altruism.

individuals close together, thereby creating a primitive social group. This is the “attachment” or “following” behaviour of offspring with respect to their parents. Figure 12.3 shows the evolutionary decrease in the average distance of offspring from their parents.

Oscillations in the evolutionary curve appear to be caused by the periodical reintroduction of food. When food is periodically (“seasonally”) more abundant there is a consequent increase in population size (cf. Fig. 12.4).

Since offspring are born near their parents an increase in population size means that many offspring are located close to their parents independently of their ability to follow their parents. The evolutionary increase in the ability to follow one’s parent becomes apparent when population size is low (just before food reintroduction), and offspring must be able to follow their parents to survive. In fact, the curve of Figure 12.3 reflects two factors, i.e. the periodical changes in population size, and the increasing tendency of offspring to remain in proximity to their parents.

The “following” behaviour of offspring is directly represented in Figure 12.5, which shows the trajectories of a particular parent and one of its offspring. The parent moves in the environment approaching food and the




Figure 12.3 Evolutionary change in the average distance between offspring and their parents.




Figure 12.4 Evolutionary change in population size.




Figure 12.5 Trajectories of an adult individual and its offspring.

offspring follows its parent in an attempt to keep its distance from the parent below the threshold necessary to receive food.

The emergence of social groups in populations of artificial organisms

We have presented some simple simulations of the emergence of spatially-defined social groups that result from kin-directed altruistic behaviour and attachment behaviour. (For other simulations that relate sociality to spatial distribution in populations of neural networks, see Parisi et al. in press.) Since in our simulations offspring tend to follow their parents, there is a tendency for small groups, formed from a parent with its offspring, to emerge. The individuals that form such social groups tend to remain close together and to move as a group. Note that these groups have not been created by the researcher by inserting in the offspring explicit rules for following their parents. Rather, social groups emerge spontaneously as a consequence of the behavioural strategies that increase an individual's reproductive chances.

In our simulations, a social group is maintained by the behaviour of the offspring alone. Of course, the behaviour of the parent is also critical in that, unless parents give some of their food to their offspring, the offspring's following behaviour cannot possibly emerge. However, in terms of space-related behaviour, parents go after food, not after their offspring. Hence, spatial vicinity is maintained directly by the offspring's behaviour of following their parents, not by the parents' behaviour of searching for food. In our simulations parents cannot possibly maintain spatial vicinity with their offspring because they do not perceive their offspring at all.

This shows that social groups can be maintained by individuals who have different roles. But parents might also try actively to stay close to their offspring. In other simulations one might endow adults (parents) with an ability to perceive both food and their offspring. Adults would have two sets of functioning input units, one set informing them about the location of the nearest food and the other set informing them about the location of the nearest offspring. In these circumstances one might expect parents to develop two types of space-related behaviour. They might evolve an ability to go after food (as in the simulations of the preceding section) and, furthermore, they might evolve an ability to keep close to their offspring, thereby directly contributing to keeping the

social group together. One might expect this second ability to evolve naturally as a by-product of the present evolutionary scenario, since by keeping close to their offspring parents can increase the likelihood that their offspring will receive the energy that allows them to survive to adulthood and reproduce.

One problem which might arise in the new simulations is that a parent might find itself in situations in which it has to move in one direction to capture food and in a different direction to keep close to its offspring. Hence, in such circumstances the parent would need to decide between two conflicting courses of action. However, as Floreano (1993) has shown, this problem can be solved by neural networks similar to ours with no need, again, for explicit rules to solve the conflict. In his simulations organisms live in an environment which contains a “nest” in a particular location and a number of randomly distributed pieces of food. Fitness is only increased when an individual first captures a (single) piece of food and then goes to the nest. In other words, it is assumed that food must be captured in the open but can only be consumed in the nest. Floreano’s networks are able to evolve this type of alternating behaviour (food, nest, food, nest, etc.) provided they are informed by a dedicated set of input units whether they have a piece of food in their “hand” or not. In the same way we can expect that, rather than being caught in unresolvable doubt about what to do in a situation of conflict, our adult networks would evolve some sort of alternating behaviour of first capturing food and then taking the food to the offspring.

The social groups of the present simulations are very small. In fact, they are limited to one parent and its offspring. Two kinds of extension can easily be envisaged, one in the genetic kin line and one in the non-genetic (affinal) kin line. The extension in the genetic kin line requires that more than two successive generations live together and form social groups (families). In our simulations it sometimes happens that some offspring of a living adult reach adulthood and reproduce. Hence, individuals belonging to three successive generations may be living at the same time: offspring, parent and grandparent. However, there can be no tendency for individuals belonging to three successive generations to live together as a social group because this would require that not only pre-adult offspring follow their parents but also that adult offspring do the same with their parents. This is not possible in the present simulations because when an individual reaches adulthood it ceases to perceive its parent and starts perceiving food. Hence, adult offspring cannot follow their invisible parents.

However, we might easily allow individuals to continue to perceive their parents after reaching adulthood while at the same time perceiving and looking for food. The real problem is what adaptive function remaining near its parent would have for an adult. Adults obtain their energy by capturing food, not by being given energy by their parents. Hence, there would be no reproductive advantage in following one’s parent after reaching adulthood. Only by introducing some such advantage might one observe the evolution of larger social groups constituted by individuals belonging to three, rather than two successive generations. (Another possibility, of course, is that juvenile traits are maintained in adulthood even if they are no longer functional.)

A possible reproductive advantage for adult (or even pre-adult) individuals of staying close to one’s parent is that proximity might allow individuals to learn from their parents. This raises the whole issue of cultural transmission of behaviour through imitation or, more generally, through learning from others. Imitation can be modelled as a network

(the imitator) taking the output of another network (the model) in response to a shared input as its teaching input.¹ In fact, in the simulations described in Denaro & Parisi (1994) we have shown that neural networks evolve a genetically inherited tendency to stay close to other networks (independent of any genetic relationship among them) if this allows them to learn by imitation some useful behaviour that increases their fitness.

An increase in the size of family groups could also be caused by the emergence of altruistic behaviour with respect to other kin, e.g. cousins. Kin selection theory predicts that altruistic behaviour is a function of degree of genetic relatedness. Of course, to act altruistically towards cousins (even if less altruistically than towards offspring or sisters) requires the parallel development of an ability to recognize different types of kin (Hepper 1991).

The other extension of social groups that might be considered is in the non-genetic kin line. This extension requires that reproduction ceases to be agamic, as in our simulations, and becomes sexual, with two parents giving birth to offspring. (For sexual reproduction in ecological neural networks, see Menczer & Parisi 1992.) Sexual reproduction can have consequences for the formation of social groups at two levels. First, sexual reproduction requires that two individuals come together in space for mating to occur, hence sexually reproducing networks might evolve a tendency to form (temporary) social groups composed of two individuals. (Werner & Dyer (1991) have simulated communication between “male” and “female” networks that causes two individuals to come together for mating.) Second, sexual reproduction might cause the creation of more permanent social groups if the two parents have to live close together in order to help their offspring survive to adulthood. The two parents might live together as a by-product of living near the same offspring, or they might live together for some specific function connected with their offspring. For example, one parent might evolve a tendency to feed not only its offspring but also its mate. This behaviour might evolve if the survival of the mate is critical for the survival of the offspring. (The other parent might protect the offspring from danger.) Hence an altruistic behaviour (feeding) toward a non-genetically related individual (one’s mate) might evolve because it has an indirectly favourable effect on the reproductive chances of one’s offspring. The net result would be the emergence of permanent social groups consisting of father, mother and offspring.

If the emerging family groups can be enlarged in both the consanguineal and affinal lines (Murdock 1949), the number of individuals that constitute a social group can increase considerably. Dunbar (1993) has estimated a group size of around 150 individuals for early social humans, based on an extrapolation of the observed relationship between group size and neo-cortex size in primates. Ethnographic evidence on hunter-gatherer and other early types of society indicates a layered group structure with three levels: bands or overnight groups (30–50 individuals); clans and lineage groups (100–200 individuals); and tribes (500–2500 individuals) (cf. Dunbar 1993). Johnson & Earle (1987) distinguish between family-level groups (30–50 individuals) and local groups (300–500 individuals).

Of course, the next step in the evolution of human social groups is the emergence of groups of individuals living together when the individuals which constitute the group are not relatives. While it may be reasonable to evolve social groups constituted biologically by kin, as in our simulations, since the basis of the behaviour that leads to the formation and maintenance of kin groups appears to be mainly genetic, for non-kin social groups it

is more likely that the behaviour which creates such groups is learned from others and, therefore, culturally rather than genetically transmitted. The general tendency to learn from others can be inherited biologically in organisms—such as humans—that have developed an adaptation which assigns a critical role to cultural transmission, but the actual behaviours which create non-kin social groups are likely to be learned rather than inherited genetically.

Notes

1. We are assuming back-propagation (Rumelhart, McClelland, the PDP Group 1986) as the learning method here. For work using this neural network model of imitation see Chapter 9 and Terna 1993.

Chapter 13

Understanding the functions of norms in social groups through simulation

Rosaria Conte and Cristiano Castelfranchi

In the Artificial Intelligence (AI) literature a norm is treated as a behavioural constraint; that is, a reduction of the action repertoire and therefore of the actions physically available to the system. According to this view, norms allow co-ordination within social systems, the latter being defined as groups of jointly acting agents (Shoham & Tenneholtz 1992a, b, Moses & Tenneholtz 1992).

This view is insufficient because it accounts exclusively for what are known as norms of co-ordination (Conte & Castelfranchi 1993). Within game theory, norms of co-ordination are seen essentially as conventions; that is, behavioural conformities that do not presuppose explicit agreements among agents and emerge from their individual interests (Schelling 1960, Lewis 1969). The function of these norms is essentially that of permitting or improving co-ordination among participants. A well-known example (Lewis 1969:5) is that of two people unexpectedly cut off during a telephone conversation. They both want to continue the conversation and each of them has to choose whether to call back or wait. A convention gradually emerges from interactional practice, establishing who should do what. Norms of co-ordination, therefore, do not stem from a conflict of utility. Indeed, the solution of a problem of co-ordination is such that, if one agent succeeds, so does the other. The single agent's utility implies the joint one, and vice versa.

However, we claim that norms of co-ordination are only a subset of norms. First, norms are not necessarily conventional: they do not necessarily emerge from implicit agreement but are also imposed by explicit prescriptions, directives and commands. Secondly, norms often arise from conflicts of utilities among agents, and sometimes they provide what economists call Pareto-inferior solutions; that is, solutions which produce a state of affairs such that some agents are worse off than they were before. Hence, while norms of co-ordination provide unanimous solutions, this is not true for every type of norm. Thirdly, not all norms serve to improve co-ordination, reduce accidental collisions and other negative interferences. Some important norms have other socially desirable effects. For example, the norm of reciprocation is essentially orientated to control cheating. Other norms prevent more open forms of deliberate aggression (physical attack). A good theory of norms should stay distinct and explore their relative functionalities.

Elsewhere (Conte & Castelfranchi 1993), we have tried to account for norms as prescriptions represented in the minds of autonomous (that is, self-interested) agents in a multi-agent context. Here, instead, we are concerned with the functions of norms, independent of the agent's makeup. One specific function will be hypothesized to lie in

the control and reduction of aggression among agents in a common world; that is, in a world where one's actions have consequences for another's goal achievements.

The design of the simulation

We shall describe a simulation experiment aimed at exploring the role of norms in the control of aggression among a population of agents. In this experiment, agents are allowed little, if any, autonomy: they perform some elementary routines for surviving in a situation of food scarcity. A utilitarian strategy will be compared with a normative one and the differences will be observed and measured with some indicators (rate of aggression, average strength, and its variance).

The world

The program, implemented in C++, defines "agents" as objects moving in a two-dimensional common world (a 10×10 grid) with randomly scattered food. The term "agent", indeed, sounds lofty in such a context. However, it is used in the sense of applying some routines. Whenever an agent is said to do something, this is equivalent to saying that a given routine has been applied and that a given object (the agent) is to be considered as its performer. This use is consistent with the general use in the social simulation field. Moreover, it is by no means prejudicial to the objectives of the present work. Whether the agents are truly distinct programs, or are defined as if they were such, makes no difference. The aim is to observe the functions of some macro-social object, i.e. norms, rather than to implement a decentralized system with effectively autonomous agents. We are interested in the function of norms independent of the level of cognitive sophistication of the agents among which the norms are in force. The question raised here is, "What is the use of norms?" rather than, "How are norms implemented in autonomous agents?"

An experiment consists of a set of "matches", each of which includes a number of "games". At the beginning of a match, 50 agents and 25 food items are assigned locations at random. A location is a cell in the grid. A cell cannot contain more than one object at a time (except when an agent is eating). Locations may be filled by food items, which the agents "want" to acquire. Again, in this context, the terms "want", "goal", etc. may appear to be metaphorical. However, they are unequivocally meant to refer to performing, rather than being involved in, some routine. For example, the agents move on to the grid in search of food and may want to attack one another as a function of the scarcity and concentration of food. Concentration is obviously distinct from scarcity: while scarcity points to a food-to-agent relationship, concentration points to a food-to-cell relationship. Food may be scarce but highly concentrated in a given region of the grid. However, the agents cannot want to be attacked.

It should be noted that only the agents who perform the "eating" routine (in short, "eaters") may be attacked: no other type of aggression is allowed. When agents are about to move to the same location, one agent is given precedence at random.

The agent-to-food ratio is always maintained at 2:1. The food's "nutritional" value (20), the initial strength of each agent (40), and the costs of the action types (0 for eating,

1 times the number of steps for moving, and 4 for attacking) are also kept constant throughout the simulation. Of course, the agents' strength does change over the course of each game as a function of the amount of food eaten, the actions performed, and the attacks suffered by each agent (minus 4 for each attack, no matter who is the winner).

The routines

At present, 50 agents are included in each match. They move in search of food as soon as some food item appears within their "horizon" (extending two steps in each direction from the agents' current location). The agents move as castles on a chess board: one step either up, down, left or right, but not diagonally. Food may be "seen" within a cross-shaped area called a "territory", consisting of the four cells to which an agent can move in one step from its current location). Food can also be "smelled" within each agent's horizon. At the corners and at the edges of the grid, both territories and horizons are reduced. Even if a cell in some agent's territory is occupied by another agent the latter does not block the former's sensing of the food items (for example, in Figure 13.1, where only a region of the entire grid is shown, the food item located at cell (4; 5) may be smelt by *F*).




Figure 13.1 The grid for simulation.

There are five routines: EAT; MOVE-TO-FOOD-SEEN; MOVE-TO-FOOD-SMELT; AGGRESS; and MOVE-RANDOM. When alternative moves are allowed, a record of the existing options is kept on an agenda to facilitate selection of the most convenient move.

The algorithm is as follows:

- (a) *Search food.* If the cell that agent *x* is occupying is not supplied with food, *x* starts exploring its territory from the upper cell. If *x* finds a cell with food and unoccupied, it will MOVE-TO-FOOD-SEEN. If, instead, *x* finds one or more cells supplied with food but occupied, it will record the related moves (AGGRESS, cost 4) in the agenda and keep checking its horizon. If a cell with food is found, *x* moves there and eats (MOVE-TO-FOOD-SMELT). However, since *x* can "see" another agent only when it is close to it (one step away), that cell may turn out to be occupied. In this case, *x* will record the related move (AGGRESS, cost 4). When no cell with food is found, *x* moves to a free cell. If more than one is available, *x* chooses one at random. If no cell is free, *x* stays at its current location.

- (b) *Choose action*. This function is called by the food search. The agenda is checked and if more than one alternative is found, one is chosen at random.

Matches, games and moves

Every match is started by assigning locations to food items and agents in such a way that a cell is not assigned to more than one object at a time. When a food item has been consumed, it is restored in the next game at a randomly chosen location. All matches consist of a fixed number of games (2000), and a game includes a maximum of 50 moves.

Moves may consist of either attacking some eater, eating, changing location, or simply staying at the same location (when no free cell is available). Moves are intended to be simultaneous and time-consuming. However, eating moves begin at a given game and end up two games later unless they are interrupted by aggression. To simplify matters, the eater's strength changes only when eating has been completed. Thus, while the action of eating is gradual (to give players the chance of attacking each other), both the food's nutritional value and the eater's strength change instantaneously (the former remains at 20 during the whole of the period during which it is being eaten and only goes down to 0 when the eating has been completed; analogously, the eater's strength is unchanged during the period of eating and increases to 20 as soon as the eating has been completed).

After each game, the number of attacks that have occurred is recorded; the agents' average strength and its standard deviation are noted. The outcomes of attacks are decided in terms of the agents' respective strengths (the stronger is always the winner). When the competitors are equally strong, the winner is selected randomly. The cost of aggression is equal to the cost of being at the receiving end of aggression. Both winners and losers encounter the same costs. However, winners obtain the contested food item. Targets cannot avoid aggression: once they have been attacked, they are bound to pay the cost of aggression even if they happen to win the competition. Agents might be attacked by more than one agent at a time, in which case the victim's cost is multiplied by the number of aggressors. In case of multiple aggression, the winner is the strongest. In case of equal strength, the winner is chosen randomly.

Aggression can be repeated. Two cases have been observed to occur:

- (a) equally strong agents keep attacking each other, snatching food out of each other's hands; and
- (b) one or more weaker competitors keep attacking the eater and losing until the food supply is exhausted.

While the former case may be repeated indefinitely, until one of the competitors dies, the latter case cannot exceed two games, after which food is necessarily exhausted at the given location.

Action types

Four action types are available to the agent:

- (a) MOVE-TO

Pre-condition:

x is one step from the destination.

Effect:

x 's strength is decreased by 1; and

x 's location is changed by one step.

(b) EAT

Pre-condition:

x is at the food's location.

Effect (two games later):

the food is removed from x 's location;

x 's location is unchanged; and

x 's strength is increased by 20.

(c) ATTACK

Pre-condition:

the agent attacked (y) is eating (it is in a cell with food); and

x is one step away from y .

Effect:

either x 's or y 's location is supplied with food; or

both x 's and y 's strength are decreased by 4.

In sum, attack consists of snatching the food from someone's hands. It is the food item that, so to speak, moves from one location to another.

The competitors remain in their respective locations.

(d) STAY

Pre-condition:

all the cells in x 's territory are occupied and not supplied with food.

Effect:

x 's location is unchanged; and

x 's strength is unchanged.

The order of preference among these actions is as follows: EAT, MOVE-TO-FOOD-SEEN, MOVE-TO-FOOD-SMELT, ATTACK, MOVE-RANDOM and STAY.

Some examples

To give a flavour of the simulation, let us consider the following three examples (see Figure 13.2):

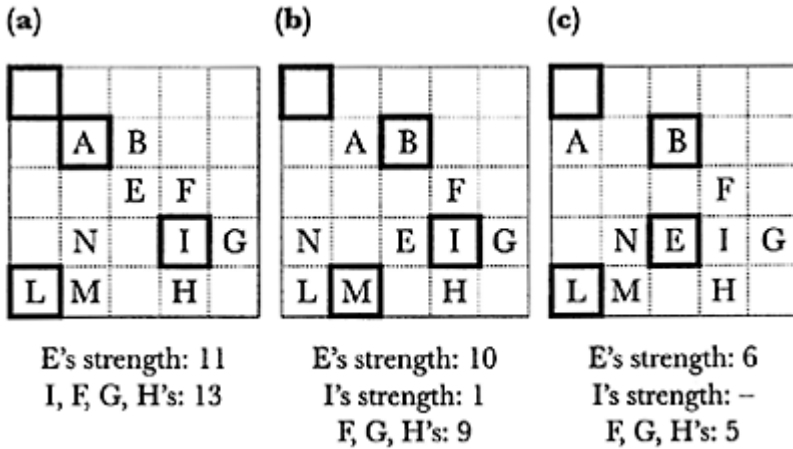


Figure 13.2 Examples of simulations.

1. An end to iterated aggression. In Figure 13.2(a), consider the relative position of *A* and *B*. Suppose the agents are equally strong and *A* is eating. Then *B* snatches *A*'s food (Figure 13.2(b)). If other food items were not around, *A* would resist the attack and the two agents would keep attacking each other until either the food was finished or one of them died. But since a less costly alternative source of food is available at cell (1, 1), *A* will turn to it (see Figure 13.2(c)).
2. Multiple aggression and unintentional favour. Look at agents *I*, *F*, *G*, and *H*. *I* is eating and is surrounded by hungry players that are as strong as *I* (Figure 13.2(a)). But *I* is lucky and holds its food tight. Later, however, *I*'s strength is considerably lowered, to 1 (Figure 13.2(b)). Next time, its attackers will succeed easily. However, the winner will not be any of the earlier competitors, since their aggression caused them to lose some of their power (their strength fell to 9 each (see Figure 13.2(b)). The agent that gets the food is *E*, although its strength in the previous scene (Figure 13.2(a)) was lower than that of the other players. Thanks to the attacks of the others, *E* can now not only easily defeat *I* but can also have the better of the other players, for their strengths are now lower than *E*'s (Figure 13.2(c)).
3. On the track of food. In our simulation, agents sometimes cannot help engaging in costly behaviour. A similar phenomenon occurs rather frequently in everyday life also. Consider the players *L*, *M* and *N* in Figure 13.2(a). *L* is holding the food, but is robbed by *M* which is equally strong (Figure 13.2(b)). Now, *M* is holding the food, but *L* attacks *M* and gets the food back. Poor *N* is always one step behind on the track of food, "smelling" it without being able to get hold of it.

The experimental conditions

The experiment was designed to investigate the role of norms in the control of aggression. To do this, situations in which agents follow norms were compared with identical situations where players follow utilitarian rules. In other words, a normative routine was compared with a strategic one. However, given the extreme naivety of the modelling of agents, norms could not be distinguished from other rules in terms of their implementation. In the context of this experiment, the difference between a normative and a non-normative strategy is in their immediate results: a normative strategy is one which may be disadvantageous to the agent that applies it, and would therefore be discarded, could that agent decide in strictly utilitarian terms.

Three conditions were compared:

- (a) *Blind or savage aggression (B)*: in which aggression is constrained only by personal utility, with no reference to the eaters. Agents attack eaters when the costs of alternatives, if any, are higher. They are aware of neither their personal nor the eaters' strengths.
- (b) *Strategic (S)*: in this condition, aggression is constrained by strategic reasoning. A confrontation of strength among competitors takes place; agents can only attack those eaters whose strength is not higher than their own. An eater's strength is visible one step away from the agent's current location. A record of each agent's strength is kept throughout the match.
- (c) *Normative (N)*: in this condition, norms are introduced. At the present stage of the experiment¹, we decided to explore the functionality of a norm of precedence to finders (a sort of finder-keeper precept): finders become possessors. At the beginning of a match, agents are randomly allocated to locations and are assigned those food items which happen to fall into their own territories. For example, in Figure 13.1, $f_{2,3}$ lies in B 's L 's, and C 's territories. Therefore, B , L and C are said to be co-possessors of $f_{2,3}$. Much as happens in real life, possession is ascribed on the grounds of spatial vicinity. However, agents maintain ascribed rights over time, even when they walk away from their possessions. Food possessed is therefore flagged and every player knows to whom it belongs. Of course, the chances are that some agents happen to be non-possessors, while others are ascribed more than one food item. Some have only shared possessions, others are unique possessors of more than one resource.

Agents cannot attack possessors eating their own food. In Figure 13.1, if B were on $loc_{2,3}$, it could be attacked by both L and C . But, in Figure 13.3, C cannot be attacked although H and A can. Non-possessors are expected to be heavily discriminated against, since they can be always attacked. However, since food items reappear at random, those that were possessors in one game may turn into non-possessors in the next, and vice versa.




Figure 13.3 Flagged food items.

Findings

Findings have been gathered from 100 matches for each condition (that is, for a total of 400 matches). For each set of 100 matches, the number of attacks (*Agg*), the average strength (*Str*), and the variance of strength (*Var*) were recorded, and the significance of the differences tested. From now on, we will consider variance of strength as a measure of inequality; the larger the variance, the more polarized, that is, the less equitable the distribution of strength is, and vice versa.

The major results are shown in Figures 13.4 and 13.5. Figures 13.6 to 13.8 show the variance of the three measures obtained (*Agg*, *Str* and *Var*) during the matches in each of the three experimental conditions, “Blind”, “Strategic”, and “Normative”.

With regard to aggression (see Figure 13.4, first three bars), even if normative and non-normative rules do well at constraining aggression (all dif-




Figure 13.4 Overall means of the number of attacks, strength and variance of strength.

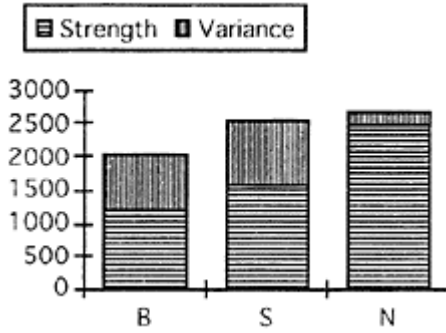


Figure 13.5 Strength and variance of strength under three conditions: blind, strategic and normative.




Figure 13.6 Variance of number of attacks (Agg) under the three conditions.




Figure 13.7 Variance of the strength (Str) under the three conditions.




Figure 13.8 Variance of inequality (Var) under the three conditions.

ferences are highly significant ($p < .001$), the normative rule does far better than the strategic one.

With regard to average strength (Figure 13.4, second set of bars), again, all differences are highly significant ($p < .001$). The findings are in the opposite direction from those for aggression. In both the *S* and *N* experimental conditions, strength is significantly higher than in *B*. The difference is mainly due to the normative condition, where the average strength is higher. Again, the difference between *S* and *N* is significant. So far, things are as expected: the lower the aggression, the higher the strength.

With variance, or inequality (Figure 13.4, third set of bars), however, the pattern of results varies. Surprisingly, the Strategic is the condition in which the polarization of strength is the widest (see the height of the second column). Conversely, the Normative condition differs significantly ($p < .001$) from both the Blind and the Strategic ones. In other words, the Normative condition is apparently the most equitable of the three. What is interesting about the variance is that it does not go in the same direction as the other measures. If variance is compared with average strength (see Figure 13.5), for *B* and *S* the variance increases with strength. However, although the difference of strength between *B* and *S* is significant, the difference of variance is not (Duncan's test). Variance, therefore, does not only vary with strength. In the Normative condition, although the average strength is significantly higher than in any other condition, variance is significantly lower. Therefore, while strength is an increasing function of the control of aggression, variance is not.

Analysis of the trends of the measures obtained during the matches (Figures 13.6–13.8) confirms the same pattern of results. We tested the significance of differences at regular time intervals (after 500, 1000, 1500 and 2000 moves) and found differences as significant as the overall ones. As for aggression (Figure 13.6), it reaches a peak around the 500th move in each condition, and then goes down for *B* and *N*, but remains steady in the Strategic condition. Strength (Figure 13.7) increases constantly in all three conditions (this is an effect caused by the amount of food being constant). As for variance (Figure 13.8), while its trend is consistent with that of strength in *B* and *S*, it is much slower in *N*. The significance of the differences in variance is because of the Normative condition (Duncan's test). At each time interval, variance differs significantly only between the Normative and each of the other two conditions, a pattern of results confirming that the Normative condition is by far, and steadily, the most equitable condition.

In sum, the control of aggression *per se* is neutral: it is neither pro-social, nor anti-social. Whether it plays a positive or a negative social role depends on the type of control employed. In the Strategic condition, all the social costs of controlling aggression are sustained by the weaker individuals. The stronger ones do not pay anything. The stronger agents are never attacked and therefore do not pay the costs of being attacked, while the weaker ones do. Stronger agents only pay the costs of undertaking advantageous aggression. Their costs are lower than their gains (cost of aggression=4, against the food item's nutritional value=20). The weaker agents sustain the costs of being attacked without obtaining anything. In other words, advantageous aggression is not constrained, while disadvantageous aggression is.

In the other conditions, although aggression is always to the advantage of the stronger agents and to the disadvantage of the weaker ones, the costs of being attacked are shared equally by both weak and strong. Furthermore, in the non-Strategic conditions, unintentional alliances among weaker agents are allowed which might prove advantageous for one (as in the example shown in Figure 13.2) or for all of them (as when a stronger eater attacked by a number of players dies during the match). An indirect effect of multiple aggression is thus the reduction of social differences.

Even the Blind condition is more equitable than the Strategic one: it allows for disadvantageous aggression, and although this is a short-term cost for weaker agents, it turns out to be a long-term benefit as well, in that it reduces the social gap.

It is not surprising that in the Normative condition the distribution is more equitable. The costs of the control of aggression are equally shared among the strong and the weak. Because food is constant and the match length is fixed, the differences between agents introduced by the norm of precedence are not maintained throughout the match. Every time a food item reappears, the status of some agents changes. Some agents, who earlier may have been non-possessors, now find themselves endowed with a new possession, while some others, who have been privileged in the past, become members of the unlucky part of the population. Thus, the norm in question is not truly a norm of property but rather a "finder-keeper" norm, since the possessions are not ascribed exclusively on the grounds of the "right of birth", but may be acquired over the course of each agent's life. The norm of finder-keeper helps the agents to find a solution to competition for the same food item, if any such competition occurs. Interestingly, this type of norm, which seems to have a biological foundation (de Waal 1982; Eibl-Eibesfeldt 1967/1978), while controlling aggression efficaciously, also reduces the variance of strength among the agents; that is, their inequality.

Speculative scenarios: the role of norms from the stronger agents' perspective

Is there any collective advantage in the control of aggression in social groups, and in particular in maintaining those norms that produce this control? Our data seem to suggest that the main effect of the normative control of aggression is not an increase in the average strength (which is high in the Strategic as well as in the Normative conditions), but reduced inequality. In other words, each agent's share of the overall advantage is different in different conditions. In the Normative condition, the weaker agents are

attacked less frequently and consequently eat more food. For the weaker agents, a life of aggression is not convenient. A distributive view of norms is suggested initially by our findings: some social laws are useful for the weaker to control the dominance of the stronger.

A problem arises immediately from this conclusion: why should the stronger obey a norm which tends to favour the weaker more than it favours them? Why should all agents accept a non-unanimous sort of solution to social conflicts? These are classical questions of moral and political philosophy that we cannot solve merely on the grounds of our simplistic experiments. However, let us at least try to formulate some answers to the questions:

- (a) *Additional costs in norm violation.* In simulating the “prescriptive” character of norms among cognitive agents, a spreading awareness of the norm, the control behaviour, and short-term and long-term punishments would be expected. Under these conditions, cheats could receive high additional costs (punishment for transgression), and the stronger are more likely than the weaker to cheat. The existence of the norm does not favour the stronger (indeed, it is a pure cost), but once the norm has been established, to observe it is advantageous, even for the stronger.
- (b) *Common interest.* Whether or not the population of agents considered does really form a “group” makes a large difference. Members do not pursue only individual interests, but also some common goals of mutual advantage. In such cases, the control of intra-group aggression becomes crucial. We may imagine three different scenarios:
 - (i) *Inter-group competition.* Suppose there is a prize for the group that obtains the highest strength score, to be distributed among its members. In this situation, it might be better for the stronger agents to “cooperate” by observing the norm and thus reducing the level of aggression.
 - (ii) *Need for mutual aid and exchange.* Suppose that the members of the population need to help one another to realize their individual goals (i.e. there are dependence relations: Castelfranchi et al. 1992; Conte et al. 1991, Cesta & Miceli 1993), and suppose that agents are able and willing to provide such help only when they have a certain amount of power. In this case, the stronger agents may have some interest in a relative decrease of their personal strength and in a more equitable distribution.
 - (iii) *The “wheel of fortune”.* A third possible explanation points to different degrees of variability in the agents’ fates. In the wheel of fortune, you cannot tell who will climb and who will sink on the social scale. Even in our artificial world, the status of the agents during the game is not constant: powerful agents can be reduced to weak, and vice versa. This is due to fortune (and misfortune) and to the effects of attacks (especially multiple attacks in which the aggressors pay just a share of the costs sustained by the target). If the stronger can lose their status and power, it is possible to ascertain to what extent and in which condition it is better for them to respect norms. For example, a norm of reciprocity might be better for the stronger to observe (be merciful and you will receive the same treatment if and when you become weaker). However, the question of whether the agents’ fates are steadier in normative scenarios than in the non-normative conditions is an open one. In our normative condition, fates are at least in part precarious because possessions change over the course of the match. However, it would be interesting to try out this hypothesis with other types of norm.

Concluding remarks

Experimental studies of the emergence of co-operation (e.g. Kreps et al. 1982, Axelrod 1984, Axelrod & Dion 1988, Lomborg in press, Bicchieri 1990, Glance & Huberman 1993, Huberman & Glance 1993) have achieved a level of sophistication which is far beyond that of the experiment presented here. Unlike those studies, however, this experiment is meant to pave the way for the study of norm functions among cognitive agents, and, consequently, to introduce both mental representations of norms and normative reasoning in the cognitive agents. Therefore, the most interesting findings are expected to be found through developing more sophisticated agent architectures, rather than experimenting with larger populations, more sophisticated dynamics, or more complex marginal utility functions.

The present experiment seems to show that:

- (a) norms may have the function of constraining aggression;
- (b) by controlling aggression, norms also have an indirect positive effect on the strength of the population (at least as effective as utilitarian constraints);
- (c) norms have a specific and positive impact on the single agents' share of the overall "cake". This seems to be a result of the specific type of aggression control allowed by normative constraints, which reduce both advantageous and disadvantageous attacks, thus redistributing the costs of controlling aggression over the population. In the non-Normative conditions, advantageous attacks were not reduced and the costs of aggression control were not fairly distributed.

Future work will aim to:

- (a) implement additional types of norm (e.g. the norm of reciprocity) and examine their consequences;
- (b) compare different strategies of aggression control within the same population, and observe their evolutionary trends;
- (c) allow some learning mechanism, and observe the spreading of normative strategies; and
- (d) allow some level of normative choice by implementing a cognitive model of norms.

Acknowledgements

We would like to thank Juri Castelfranchi and Vincenzo Catoni for their help in programming, Alessandro Laudanna and Dario Salmaso for their competent (and patient) statistical suggestions, and Fabrizio Stasolla for his assiduous and understanding collaboration.

Notes

1. In an earlier stage of the experiment, two norms of precedence were applied, namely a norm of precedence based upon possession of resources, and a convention forcing agents to give precedence to eaters to their right. Both fitted the simulated environment rather well, but while the former was grounded on "natural" and "historical" conditions, the latter was

merely formal and contractual. In that stage of the experiment, however, the simulations were run without the food supply being kept constant. As a result, the initial possession of food items introduced a difference among the players that was maintained throughout the game. In contrast, at the present stage, the food supply is kept constant, and the first norm has been changed to a finder-keeper norm: precedence is given to the agents in whose territory a given food item happens to “fall”. The difference between the first and second norms is much less evident. Therefore, at the present stage, only one type of normative condition, namely precedence to possessors, is used.

Chapter 14

A logical approach to simulating societies

Michael Fisher and Michael Wooldridge

A number of languages and testbeds for implementing and simulating artificial social systems have been described in the literature of Distributed Artificial Intelligence (DAI) and its related disciplines (see, for example, Hewitt 1977, Agha 1986, Gasser et al. 1987, Shoham 1990, Bouron et al. 1991, Doran et al. 1991, Ferber & Carle 1991). However, with the possible exception of pure Actor languages (Hewitt 1977, Agha 1986), these tools have relied on an informal, if not *ad hoc*, notion of agency, with similarly informal techniques for programming agents. As a result, it is difficult to reason, either formally or informally, about the expected behaviour of systems implemented using such tools. It is similarly difficult to derive an a posteriori explanation of such a system's behaviour. In this chapter, we present a language for simulating co-operative and competitive behaviour in groups of intelligent artificial agents. In contrast to the informal languages and testbeds referred to above, this language, called Concurrent METATEM, has a simple, well-motivated logical (and hence mathematical) foundation. In relative terms, it is easy to reason about Concurrent METATEM systems at both formal and informal levels (Fisher and Wooldridge 1993). Concurrent METATEM is, so far as we are aware, unique in DAI in having such a well developed and motivated logical foundation.

This chapter is structured as follows. In the remainder of this section, we motivate the language and discuss its origins in more detail. In the next section, we present a more detailed, though essentially non-technical, introduction to the language and its execution. We then demonstrate by example how the language may be used to model societies. Finally, we describe the current state of the language and describe our intentions with respect to future developments.

Distributed AI is a relatively young discipline which has developed its own tools for building, experimenting, and evaluating theories and applications. Over the past decade, many frameworks for building DAI systems have been reported. While some, such as the DVMT, allow experimentation only with one specific scenario (Durfee 1988) and others provide extensions to existing AI languages such as LISP (Gasser et al. 1987), comparatively few fundamentally new languages have been proposed for DAI. Arguably the most successful DAI languages have been based, at least in part, on the Actor model of computation (Hewitt 1977, Agha 1986); for example, Ferber and Carle (1991) and Bouron et al. (1991).

It is our contention that while frameworks for DAI based on extensions to existing AI languages are useful for experimentation, they will not, ultimately, be viable tools for building production DAI systems. This is because DAI systems are a subset of the class of systems known in mainstream computer science as "reactive systems". A reactive system in this sense is one whose purpose is to maintain some ongoing interaction with its environment. We therefore do not use the term "reactive system" in the way that it has

become fashionable to use it in AI, to refer to systems that respond directly to their environment, perhaps using “situation-action” rules, without reasoning in any way. All concurrent or distributed systems are reactive in the sense we intend, just as a module or agent in a concurrent system must maintain some interaction with the other components in the system (see, for example, Pnueli (1986) for a discussion of this point). It is for this reason we claim that all DAI systems are reactive.

Contemporary DAI testbeds and languages are usually built as extensions to the classic AI languages LISP and PROLOG. However, neither of these languages is well suited to building reactive systems. They are both based on a quite different view of systems, called the functional or relational view. For this reason, we suggest that an obvious development for DAI is to build languages based on a more reactive view; Concurrent METATEM is such a language.

In a 1977 paper, Pnueli proposed temporal logic as a tool for reasoning about reactive systems. When describing a reactive system, we often wish to express properties such as “if a request is sent, then a response is eventually given” and such properties are easily and elegantly expressed in temporal logic. However, proving the properties of reactive systems using temporal logic is not a trivial matter. The difficulties involved in this process led, in the early 1980s, to the idea of using temporal logic itself as a programming language (Moszkowski 1986). This idea led directly to the METATEM concept (Barringer et al. 1989). METATEM is a general framework for executing temporal logic specifications, where these specifications are expressed as a set of temporal logic “rules”. The concept of a reactive system is therefore at the very heart of METATEM. Although the original METATEM proposal did not address the issue of concurrency, the potential value of concurrently executing METATEM systems, particularly for DAI applications, was recognized immediately (Fisher and Barringer 1991). Concurrent METATEM is a simple operational framework that allows societies of METATEM processes to communicate and co-operate.

Note that although Concurrent METATEM may be regarded as a logic programming language, in that it has a well-developed and motivated logical foundation, it is quite unlike any other logic programming language with which we are familiar; in particular, it is based on a novel model for concurrency in executable logic. Whereas most previous concurrent logic paradigms are based on fine-grained AND-OR parallelism, (e.g. Clark and Gregory 1987), concurrency in Concurrent METATEM is achieved via coarse-grained computational entities called agents; each agent is a METATEM process.

Concurrent METATEM

In Concurrent METATEM (Fisher 1993), the behaviour of an agent is defined using a temporal logic formula. Temporal logic is used because, not only is it an ideal formalism in which to represent the dynamic properties of an agent, but it also contains an explicit mechanism for representing and executing goals. As well as providing a declarative description of the agent’s desired behaviour, the temporal formula can be executed directly to implement the agent (Fisher and Owens 1992). Thus, the basic behaviour of an agent consists of following a set of temporal rules representing the basic dynamics of the agent, introducing new goals, and attempting to satisfy existing goals.

Agents communicate via message-passing. As well as its temporal specification, each agent records two sets of message-types: one representing messages that it is able to send, the other representing messages that it is willing to receive. When an agent sends a message, this message is broadcast. When a message arrives at an agent, the agent will only process that message if it is one of the types of message it is “listening” for. Agents are truly autonomous; not only do they only react to messages that they want to “hear”, but they are able to dynamically change the set of message-types that they will recognize, and are able to control the number of messages consumed at any one time.

Agents are members of groups. Each agent may be a member of several groups. If an agent sends a message, then that message is broadcast to all members of its group(s), but to no other agents. Thus, the basic model of communication in Concurrent METATEM is of agents using broadcast locally, while using more selective methods for non-local communication (for example, via “eavesdropping” agents that are members of multiple groups).

Agents within groups can achieve close interaction. In particular, the use of broadcast communication allows individual agents to observe the messages that other agents send, and thus to modify their behaviour accordingly. Using this type of approach, varieties of co-operative and competitive behaviours can easily be modelled.

In the following subsections, we shall give a more detailed introduction to Concurrent METATEM. We begin with an informal introduction to the type of temporal logic that Concurrent METATEM is based on.

Temporal logic

Temporal logic is classical logic augmented by a set of modal operators through which it is possible to describe the time-varying state of the world. Although there are many different types of temporal logic (Emerson 1990), in this chapter we shall consider only one. This logic is based on a linear, discrete model of time, which is bounded in the past and infinite in the future. This means that there is only one “timeline”; that time comes in “atoms”, such that each moment in time has a unique predecessor and successor; that there was some moment in the past at which time “began”; and that time stretches infinitely into the future.

The logic we consider is based on classical first-order logic. In addition to the usual connectives and quantifiers of this language, the logic contains a number of temporal operators. First there are the future time connectives. If ϕ is a formula of the temporal language, then $\diamond\phi$ is also a formula, which will be satisfied now if ϕ is satisfied now or at some time in the future. Another unary connective is \square : the formula $\square\phi$ will be satisfied now if ϕ is satisfied now and at all times in the future; i.e. now and for ever more. Since each moment in time has a unique successor, we can introduce a “next time” connective: $O\phi$ will be satisfied now if ϕ is satisfied the next moment. There are also some binary connectives in the language: $\phi\cup\psi$ will be satisfied now if ψ is satisfied in some future time, and at all time points *until* that time, ϕ is satisfied.

The language also contains temporal operators for talking about things that have occurred in the past. The \blacklozenge operator mirrors the behaviour of \diamond in the past, so $\blacklozenge\phi$ will be satisfied now if ϕ was satisfied in some prior moment. The \blacksquare operator mirrors \square in the past, so $\blacksquare\phi$ will be satisfied now if ϕ was satisfied at all previous times. It is possible to

talk about the previous moment in time: $\bigcirc\phi$ will be satisfied if there was a previous moment in time, and ϕ was satisfied at that moment. (Note that the possibility of “now” being the first moment in time complicates the meaning of this “last time” operator somewhat, and so a special nullary operator, “**start**”, is used to denote the first moment in time.) Finally, there is a binary “since” connective: $\phi\psi$ will be satisfied now if ψ was satisfied at some prior moment in time, and at all times since then, ϕ has been satisfied.

Agents

A Concurrent METATEM system contains a number of concurrently executing agents that are able to communicate through asynchronous broadcast message passing. Each agent is programmed by giving it a temporal logic specification of the behaviour it is required to exhibit. Agents execute their specifications directly, where a specification consists of a set of “rules”, that are temporal logic formulae of the form

antecedent about past consequent about future.

Agents maintain a record of the messages they send and any (internal) actions they perform: agent execution proceeds by a process of determining continually which of the past-time antecedents of rules are satisfied by this recorded history. The instantiated consequents of any rules that do fire become commitments which the agent subsequently attempts to satisfy. It is possible that an agent’s commitments cannot all be satisfied simultaneously, in which case unsatisfied commitments are carried over into the next “moment” in time.

Inter-agent communication is managed by interfaces, which each agent possesses. An interface determines what messages an agent may send and what messages, if sent by another agent, it will accept. Whenever an agent satisfies a commitment internally, it consults its interface to see whether this commitment corresponds to a message that should be sent; if it does, then the message is broadcast. On receipt of a message, an agent will consult its interface to see whether the message is one that should be accepted; if it is, then the message is added to the history, otherwise, it is ignored.

For example, the interface for a “stack” agent might be defined as follows.

```
stack (pop, push) [popped, stackfull].
```

Here, {pop, push} is the set of environment predicates, (i.e. messages the agent recognises), while {popped, stackfull} is the set of component predicates (i.e. messages the agent itself might produce). Note that these sets need not be disjointed—an agent may broadcast messages that it also recognizes. In this case, messages sent by an agent to itself are recognized immediately.

During the execution of each agent’s rules, the two types of predicate have a specific operational interpretation, as follows:

- (a) *Environment predicates represent incoming messages.* An environment predicate can be made true if, and only if, the corresponding message has just been received. Thus, a formula containing an environment predicate, such as “*request(x)*”, where *x* is a

universally quantified variable, is only satisfied if a message of the form request (b) has just been received (for some argument b).

- (b) *Component predicates represent messages broadcast from the agent.* When a component predicate is satisfied, it has the (side-)effect of broadcasting the corresponding message to the environment. For example, if the formula “*offer(e)*” is satisfied, where *offer* is a component predicate, then the message offer (e) is broadcast.

Note that some predicates used by an agent are neither environment nor component predicates; these *internal* predicates have no external effect. They are used as part of the internal computation of the agent and, as such, do not correspond either to message sending or message reception.

Once an agent has commenced execution, it follows a cycle of reading incoming messages continuously, collecting together the rules that “fire” (i.e. whose left-hand sides are satisfied by the current history) and executing one of the disjuncts represented by the conjunction of right-hand sides of “fired” rules. Individual agents execute asynchronously, and are autonomous in that they may execute independently of incoming messages and may change their interface dynamically.

Agents may backtrack, with the proviso that an agent may not backtrack past the broadcasting of a message. Consequently, in broadcasting a message to its environment, an agent effectively *commits* the execution to that particular path. Thus the basic cycle of operation for an agent can be thought of as a period of internal execution, possibly involving backtracking, followed by appropriate broadcasts to its environment.

Finally, agents are also members of groups. Each agent may be a member of several groups. When an agent sends a message, that message is, by default, broadcast to all the members of its group(s), but to no other agents. Alternatively, an agent can choose to broadcast only to certain of the groups of which it is a member. This mechanism allows the development of complex structuring within the agent space and provides the potential for innovative applications, such as the use of groups to represent physical properties of agents. For example, if we assume that any two agents in the same group can “see” each other, then movement broadcast from one agent can be detected by the other. Similarly, if an agent moves “out of sight”, it moves out of the group and thus the agents that remain in the group “lose sight” of it. Examples such as this (which will not be explored further here) show some of the power of the “group” concept.

Examples

In this section we shall give a range of examples, starting with collections of agents acting individually and adding gradually structure and the potential for more complex interactions amongst agents and groups of agents. We will begin with a simple scenario, taken from Fisher (1993), where seven agents are each attempting to get a resource from a single agent.

Snow White and the Seven Dwarfs: a tale of eight agents

To give some idea of how Concurrent METATEM can be used, a simple example showing agents with a form of “intelligent” behaviour will be described. First, a brief outline of the properties of the leading characters in this example will be given.

The scenario

Snow White has a bag of sweets. All the dwarfs want sweets, though some want them more than others. If a dwarf asks Snow White for a sweet, she will give him what he is asking for, but perhaps not straight away Snow White is only able to give away one sweet at a time.

Snow White and the dwarfs are going to be represented as a set of agents in Concurrent METATEM. Each dwarf has a particular strategy that it uses in asking for sweets, as described below (the names and behaviours of the dwarfs differ from those in the fairy tale!).

1. *eager* initially asks for a sweet and, from then on, whenever he receives a sweet, asks for another.
2. *mimic* asks for a sweet whenever he sees *eager* asking for one.
3. *jealous* asks for a sweet whenever he sees *eager* receiving one.
4. *insistent* asks for a sweet as often as he can.
5. *courteous* asks for a sweet only when *eager*, *mimic*, *jealous* and *insistent* have all asked for one.
6. *generous* asks for a sweet only when *eager*, *mimic*, *jealous*, *insistent* and *courteous* have all received one.
7. *shy* only asks for a sweet when he sees no one else asking.
8. *snow white* can only allocate one sweet at a time. She keeps a list of outstanding requests and attempts to satisfy the oldest one first. If a new request is received, and it does not occur in the list, it is added to the end. If it does already occur in the list, it is ignored. Thus, if a dwarf asks for a sweet n times, he will eventually receive at most n , and at least 1, sweets.

This example may seem trivial, but it represents a set of agents exhibiting different behaviours, where an individual agent’s internal rules can consist of both safety (e.g. “nothing bad will happen”) and liveness (e.g. “something good will happen”) constraints, and where complex interaction can occur between autonomous agents.

The program

The Concurrent METATEM program for the scenario described above consists of the definitions of eight agents shown below. To provide a better idea of the meaning of the temporal formulae representing the internals of these agents, a brief description will be given with each agent’s definition. Requests to Snow White are in the form of an ask () message with the name of the requesting dwarf as an argument. Snow White gives a sweet to a particular dwarf by sending a give () message with the name of the dwarf as an argument. Finally, uppercase alphabetic characters, such as X and γ , represent universally quantified variables.

1. **start** \Rightarrow **ask(eager)**

○give(eager) \Rightarrow **ask(eager)**

Initially, *eager* asks for a sweet and, whenever he has just received a sweet, he asks again.

2. **○ask(eager)** \Rightarrow **ask(mimic)**

If *eager* has just asked for a sweet then *mimic* asks for one.

3. **○give(eager)** \Rightarrow **ask(jealous)**

If *eager* has just received a sweet then *jealous* asks for one.

4. **start** \Rightarrow **□ask(insistent)**

From the beginning of time, *insistent* asks for a sweet as often as he can.

5.
$$\left\{ \begin{array}{l} (\neg \text{ask}(\text{courteous}))\mathcal{S} \text{ask}(\text{eager}) \wedge \\ (\neg \text{ask}(\text{courteous}))\mathcal{S} \text{ask}(\text{mimic}) \wedge \\ (\neg \text{ask}(\text{courteous}))\mathcal{S} \text{ask}(\text{jealous}) \wedge \\ (\neg \text{ask}(\text{courteous}))\mathcal{S} \text{ask}(\text{insistent}) \end{array} \right\} \Rightarrow \text{ask}(\text{courteous})$$

If *courteous* has not asked for a sweet since *eager* asked for one; has not asked for a sweet since *mimic* asked for one; has not asked for a sweet since *jealous* asked for one; and has not asked for a sweet since *insistent* asked for one, then he will ask for a sweet.

6.
$$\left\{ \begin{array}{l} (\neg \text{ask}(\text{generous}))\mathcal{S} \text{give}(\text{eager}) \wedge \\ (\neg \text{ask}(\text{generous}))\mathcal{S} \text{give}(\text{mimic}) \wedge \\ (\neg \text{ask}(\text{generous}))\mathcal{S} \text{give}(\text{jealous}) \wedge \\ (\neg \text{ask}(\text{generous}))\mathcal{S} \text{give}(\text{insistent}) \wedge \\ (\neg \text{ask}(\text{generous}))\mathcal{S} \text{give}(\text{courteous}) \end{array} \right\} \Rightarrow \text{ask}(\text{courteous})$$

If *generous* has not asked for a sweet since *eager* received one; has not asked for a sweet since *mimic* received one; has not asked for a sweet since *jealous* received one; has not asked for a sweet since *insistent* received one; and has not asked for a sweet since *courteous* received one, then he will ask for a sweet!

7. **shy(ask) [ask] :**

start \Rightarrow **◇ask(shy)**

○ask(X) \Rightarrow **¬ask(shy)**

○ask(shy) \Rightarrow **◇ask(shy)**

Initially, *shy* wants to ask for a sweet but is prevented from doing so whenever he sees some other dwarf asking for one. Thus, he only succeeds in asking for a sweet when he sees no one else asking and, as soon as he has asked for a sweet, he wants to try to ask again.

8. **snow-white (ask) [give] :**
 $\bigcirc \text{ask}(X) \Rightarrow \diamond \text{give}(X)$
 $\text{give}(X) \wedge \text{give}(Y) \Rightarrow X = Y$

If *snow-white* has just received a request from a dwarf, a sweet will be sent to that dwarf eventually. The second rule ensures that sweets can not be sent to two dwarfs at the same time, by stating that if both *give(X)* and *give(Y)* are to be broadcast, then *X* must be equal to *Y*.

Note that, in this example, several of the dwarfs were only able to behave as required because they could observe all the *ask()* and *give()* messages that were broadcast. The dwarfs can thus be programmed to have strategies that are dependent on the behaviour of other dwarfs. Also, the power of executable temporal logic is exploited in the definition of several agents, particularly those using the “ \diamond ” operator to represent multiple goals.

We also note that, as the agents’ behaviour is represented explicitly and in a logical way, verification of properties of the system is possible. For example, given the agents’ definitions, we are able to prove that every dwarf except *shy* will eventually receive a sweet. For further work on the verification of properties of such systems, see Fisher and Wooldridge (1993).

Adding “money”

We now extend this example to incorporate the notion of some resource which the dwarfs can attempt to exchange with Snow White for sweets, i.e. money.

Bidding

Initially, we will simply change the *ask* predicate so that it takes an extra argument representing the amount the dwarf is willing to pay for a sweet. This enables dwarfs to “bid” for a sweet, rather than just asking for one. For example, *dwarf1* below asks for a sweet, bidding “2”.

```
dwarf1( ) [ask] :
start   ask (dwarf1, 2)
```

We can further modify a dwarf’s behaviour so that it does not bid more than it can afford, by introducing some record of the amount of money that the dwarf has at any one time. Thus, the main rule defining the “bidding” behaviour of a dwarf might become something like

$$\bigcirc [\text{money}(N) \wedge N \geq 2] \Rightarrow \text{ask}(\text{dwarf1}, 2)$$

Note that the behaviour of Snow White might also change so that all the bids are recorded and then a decision about which bid to accept is made based upon the bids received. Once a decision is made, *give()* is again broadcast, but this time having an extra argument showing the amount paid for the sweet. For example, if Snow White accepts the bid of “2” from *dwarf1*, then *give(dwarf 1, 2)* is broadcast.

Finally, a dwarf whose bid has been accepted, in this case *dwarf1*, must remember to record the change in finances:

$$\bullet [\text{money}(N) \wedge \text{give}(\text{dwarf1}, C)] \Rightarrow \text{money}(N - C)$$

Renewable resources

Dwarfs who keep buying sweets will eventually run out of money. Thus, we may want to add the concept of the renewal of resources, i.e. being paid. This can either happen regularly, at time intervals defined within each dwarf's rules, for example:

$$\text{start} \Rightarrow \text{money}(100) \wedge \text{paid}$$

$$\bullet [\text{money}(N) \wedge \bullet\bullet\bullet\bullet \text{paid}] \Rightarrow \text{money}(N + 100) \wedge \text{paid}$$

Or the dwarf can replenish its resources when it receives a particular message from its environment, for example:

dwarf1 (go) [ask]:

$$\text{start} \Rightarrow \text{money}(100)$$

$$\bullet [\text{money}(N) \wedge \text{go}] \Rightarrow \text{money}(N + 100)$$

Competitive bidding

As the bids that individual dwarfs make are broadcast, other dwarfs can observe the bidding activity and revise their bids accordingly. We saw earlier that the *mimic* dwarf asks for a sweet when it sees the *eager* dwarf asking for one. Similarly, *dwarf2* might watch for any bids by *dwarf1* and then bid more, for example:

$$\bullet [\text{ask}(\text{dwarf1}, B) \wedge \text{myhigh}(M) \wedge B > M] \Rightarrow \text{ask}(\text{dwarf2}, B + 1) \\ \wedge \text{myhigh}(B + 1)$$

Although we will not give further detailed examples in this vein, it is clear that a range of complex behaviours based upon observing others' bids can be defined.

Co-operation

In the previous section we showed how individual dwarfs might compete with each other for Snow White's sweets. Here, we will consider how dwarfs might *co-operate* in order to get sweets from Snow White. In particular, we consider the scenario where one dwarf on its own does not have enough money to buy a sweet and thus requires a loan from other dwarfs.

Borrowing money

In order to borrow money from other dwarfs to buy a sweet, a dwarf can broadcast a request for a certain amount. For example, if the dwarf (*dwarf3* in this case) knows that the highest amount bid for a sweet so far is X and he only has γ , he can ask to borrow $X - \gamma$, possibly as follows.

`dwarf3(lend)[borrow,ask]:`

$$\bullet [\text{highest}(X) \wedge \text{money}(Y) \wedge X > Y] \Rightarrow \text{borrow}(\text{dwarf3}, (X - Y) + 1)$$

Now, if another dwarf, say *dwarf4*, offers to lend a certain amount, say *Z*, to *dwarf3*, then another rule recording the loan must be added to *dwarf3*'s rule-set:

$$\bullet [\text{lend}(\text{dwarf4}, \text{dwarf3}, Z) \wedge \text{money}(Y)] \Rightarrow \text{money}(Y + Z) \wedge \text{owe}(\text{dwarf4}, Z)$$

Lending behaviour

Dwarfs might have various strategies for lending and borrowing money. For example, perhaps a dwarf will not lend any more money to any dwarf who still owes money. Further, a dwarf might be less likely to lend money to any dwarf who has never offered to help his previous requests. Again, a variety of strategies for lending and borrowing can be coded in Concurrent METATEM. Rather than giving further examples of this type, we next consider the use of groups in the development of structured systems of interacting agents.

Group structuring

As described earlier, as well as the notion of autonomous objects, Concurrent METATEM also provides a larger structuring mechanism through "groups". This restricts the extent of an object's communications and thus provides an extra mechanism for the development of strategies for organizations. Rather than giving detailed examples, we will outline how the group mechanism could be used in Concurrent METATEM to develop further co-operation, competition and interaction between agents.

Again, we will consider a scenario similar to that of Snow White and the Seven Dwarfs, but will assume the existence of a large number of dwarfs and possibly several Snow Whites. We will outline several examples of how the grouping of these agents can be used to represent more complex or refined behaviour.

Collective bidding

If we have a situation where dwarfs bid for sweets, we can organize co-operation within groups so that the group as a whole puts together a bid for a sweet. If successful, the group must also decide to whom to distribute the sweet. Thus a number of groups might be co-operating internally to generate bids, but competing (with other groups) to have their bid accepted.

Forming subgroups

Within a given group, various subgroups may be formed. For example, if several members of a group are unhappy with another member's behaviour, they might be able to create a new subgroup which excludes the unwanted agent within the old grouping. Note that members of the subgroup can hear the outer group's communications, while

members of the outer one cannot hear the inner group's communications. Although we have described this as a retributive act, such dynamic restructuring is natural as groups increase in size. By using a combination of individual agent strategies and of grouping agents together, we are able to form simple societies. In particular, we can represent societies where individuals co-operate with their fellow group members, but where the groups themselves compete for some global resource.

Although our examples have been based upon agents competing and co-operating in order to obtain a resource, many other types of multi-agent system can be developed in Concurrent METATEM. It is important to note that there is no explicit global control or global plan in these examples. Individual agents perform local interactions with each other and with their environment.

Related work

Concurrent METATEM is in some respects similar to Shoham's AGENT0 system (Shoham 1990). AGENT0 is a first attempt to build an *agent oriented programming* (AOP) language. AOP is a "new programming paradigm, based on a societal view of computation" (Shoham 1990:4), central to which is the idea of agents/agents as cognitive entities, whose state is best described in terms of *mentalistic* notions such as belief, choice and commitment.

Both AGENT0 and Concurrent METATEM are based on temporal logic, although these have quite different forms. In Concurrent METATEM, a tense logic approach is adopted. The language of Concurrent METATEM rules is a classical logic augmented by a set of modal operators for describing the dynamic nature of the world. In AGENT0, the "method of temporal arguments" is used. This means that the language contains terms for referring directly to time. Predicates and modal operators are then "date stamped" with the time at which they were true. Both styles of temporal logic have advantages and disadvantages and this chapter will not go into the relative merits of each. AGENT0 and Concurrent METATEM are both rule-based languages, although each makes novel use of the concept of a rule. In both languages, the rules an agent possesses determine how that agent makes *commitments*. In AOP, commitment is given a mentalistic interpretation, based on the formalisations in Shoham (1989) and Thomas et al. (1991). In contrast, Concurrent METATEM gives commitment a precise computational meaning. Despite these similarities, AGENT0 and Concurrent METATEM differ in many significant respects. An obvious distinguishing feature is the nature of rules in the two languages. In Concurrent METATEM, rules have an explicit logical semantics and are based on the separated (*past future*) form. In AGENT0, rules do not have such a well-developed logical semantics.

Another system used extensively for DAI is Georgeff and Lansky's procedural reasoning system (PRS) (Georgeff & Lansky 1987; Georgeff & Ingrand 1989), which employs some elements of the belief-desire-intention (BDI) architecture partly formalized by Rao (Rao and Georgeff 1991). The PRS also has much in common with Concurrent METATEM: both systems maintain "beliefs" about the world (in Concurrent METATEM these beliefs are past-time temporal formulae); knowledge areas in the PRS loosely resemble Concurrent METATEM rules; and PRS intentions resemble Concurrent

METATEM commitments. There are many points of difference, however. Notably, the PRS does not have a logical basis for execution that is as elegant as that in Concurrent METATEM. However, the elegance of Concurrent METATEM does not come cheap. The PRS seems to be more flexible in execution than Concurrent METATEM and is likely to have higher potential in time-critical applications.

The computational model underlying Concurrent METATEM is somewhat similar to that employed in the “autonomous agent model” of Maruichi et al. (1990). More generally, there are also some similarities with the Actor model (Hewitt 1977; Agha 1986). The key differences are the ability of agents in Concurrent METATEM to act in a non-message-driven way and the use of broadcast, rather than point-to-point message passing.

Finally, a comment on the use of mentalistic terminology. Both AGENT0 and the PRS make free with terms such as “belief”, “desire”, and “intention”. Shoham argues that such terminology provides a useful degree of abstraction for complex systems (Shoham 1990, pp. 5–6). Although Concurrent METATEM employs terminology such as “commitment”, no attempt is made to relate this terminology to a deeper theory of agency (as Shoham (1990) hopes to do in AOP).

Comments and conclusions

In this chapter, we have described the use of Concurrent METATEM, a programming language based on executable logic, in representing and simulating simple societies of artificial agents. We have given several examples of prototypical societies and have shown that Concurrent METATEM is useful, particularly for co-operative and/or competitive groups of agents. The use of a formal language for representing agents’ behaviours has the added advantage that the verification of properties of such societies becomes possible.

A full implementation of propositional Concurrent METATEM has been developed. This implementation has been used to develop and test numerous example systems, including a simulation of a railway network (Finger et al. 1993), and propositional versions of all the examples presented in this chapter (see below). Experience with this initial implementation is being used to guide an implementation of full first-order Concurrent METATEM. Several features, such as the grouping mechanism described above and the extension of individual agents’ capabilities to incorporate meta-level statements, are still under development.

Other ongoing and future work topics include: a complete formal definition of the language semantics (Fisher 1993); techniques for specifying and verifying Concurrent METATEM systems (see also Fisher and Wooldridge 1993); organizational structure in Concurrent METATEM; (more) efficient algorithms for agent execution; meta-level reasoning in METATEM and Concurrent METATEM (Barringer et al. 1991); a language for agent rules containing a belief component; and the development of a formal framework for the group structuring notion.

Bibliography

- Abraham, R. & C.Shaw 1987. Dynamics: a visual introduction. In *Self organizing systems*, F.E.Yates (ed.), 543–97. New York: Plenum.
- Agha, G. 1986. *Actors—a model for concurrent computation in distributed systems*. Cambridge, Mass.: MIT Press.
- Alexander, R.D. 1989. Evolution of the human psyche. In *The human revolution*, P. Mellars & C.Stringer (eds), 455–513.Edinburgh University Press.
- Allen, P. 1978. Dynamique des centres urbains. *Sciences et Techniques* **50**, 15–19.
- Allen, P. & M.Sanglier1979. A dynamic model of growth in a central place system. *Geographical Analysis*, **11**, 256–72.
- Arrow, K.J. 1962. The economic implications of learning by doing. *Review of Economic Studies* **29**, 155–174.
- Arthur, B. 1988. Urban system and historical path dependence. In *Cities and their vital systems: infrastructures, past, present and future*, J.H.Ausubel & R.Herman (eds), 85–97. Washington, DC: National Academy Press.
- Axelrod, R. 1984. *The evolution of cooperation*. London: Penguin.
- Axelrod, R. & D.Dion 1988. The further evolution of cooperation. *Nature* **242**, 1395–90.
- Axelrod, R. & D.S.Bennett 1993. A landscape theory of aggregation. *British journal of Political Science* **23**, 211–33.
- Bairoch, P. 1985. *De Jericho à Mexico, villes et économie dans l'histoire*. Paris: Gallimard.
- Barringer, H., M.Fisher, D.Gabbay, G.Gough, R.Owens 1989. METATEM: A framework for programming in temporal logic. In *Proceedings of REX Workshop on Stepwise refinement of distributed systems: models, formalisms, correctness*. Mook, Netherlands (published in *Lecture Notes in Computer Science*, **430**, 94–129. Berlin: Springer).
- Barringer, H., M.Fisher, D.Gabbay, A.Hunter 1991. Meta-reasoning in executable temporal logic. In J.Allen, R.Fikes, E.Sandewall (eds), *Proceedings of the international conference on principles of knowledge representation and reasoning (kr)*, Cambridge, Mass.
- Berlekamp, E.R., J.H.Conway, R.K.Guy 1982. *Winning ways for your mathematical plays*, vol. 2. New York: Academic Press.
- Berry, B.J.L. 1976. *Urbanization and Counterurbanization*. Los Angeles: Sage.
- Bicchieri, C. 1990. Norms of co-operation. *Ethics* **100**, 838–61.
- Bittner, E. 1974. The concept of organisation. In *Ethnomethodology*, R.Turner (ed). London: Penguin.
- Boehm, C. 1993. Egalitarian behavior and reverse dominance hierarchy, *Current Anthropology* **34**, 227–54.
- Bond, A.H. 1991. Commitments and projects. *Proceedings of the eighth conference of the society for the study of artificial intelligence and the simulation of behaviour*, 26–35, Leeds, UK.
- Bond, A.H. & L.Gasser 1988a. An analysis of problems and research in DAI. In *Readings in distributed artificial intelligence*, A.H.Bond & L.Gasser (eds), 3–35, San Mateo, California: Morgan Kaufmann.
- Bond, A.H. & L.Gasser (eds) 1988b. *Readings in Distributed Artificial Intelligence*. San Mateo, California: Morgan Kaufmann.
- Boudon, R.. 1968. *A quoi sert la notion de "structure"*. Paris: Les Essais, Gallimard.
- Bouron, T., J.Ferber, F.Samuel1991. MAGES: A multi-agent testbed for heterogeneous agents. In *Decentralized AI 2—Proceedings of the second European workshop on modelling autonomous*

- agents and multi-agent worlds* (MAAMAW), Y. Demazeau & J.P. Muller (eds), 195–214. Amsterdam: Elsevier/North Holland.
- Boyd, R. & P.J. Richerson 1985. *Culture and the evolutionary process*. Chicago: University of Chicago Press.
- Braudel, F. 1967. *Civilisation matérielle et capitalisme*. Paris: A. Colin.
- Bueno de Mesquita, B. & D. Lalmán 1992. *War and reason: domestic and international imperatives*. New Haven, Connecticut: Yale University Press.
- Buss, L.W. 1987. *The evolution of individuality*. Princeton, New Jersey: Princeton University Press.
- Byrne, R. & A. Whiten (eds) 1988. *Machiavellian intelligence*. New York: Oxford University Press.
- Caillé, A. 1992. Sujets individuels et sujets collectifs. *Philosophie et Anthropologie*, coll. Espace international philosophie, December, 93–114. Paris: Les séminaires du Centre Georges Pompidou.
- Camagni, R., L. Diappi, G. Leonardi 1986. Urban growth and decline in a hierarchical system: a supply-oriented dynamic approach. *Regional Science and Urban Economics* **16**, 145–60.
- Camerer, C. 1987. Do biases in probability judgement matter in markets? *American Economic Review* **77**, 981–97.
- Castelfranchi, C., M. Miceli, A. Cesta 1992. Dependence relations among autonomous agents. In *Decentralized AI—3*, E. Werner & Y. Demazeau (eds). Amsterdam: Elsevier.
- Cesta, A. & M. Miceli 1993. In search of help: strategic social knowledge and plans. Proceedings of the 12th international workshop on distributed AI, Hidden Valley, Philadelphia.
- Champion, A. 1989. *Counterurbanization*. London: Edward Arnold.
- Chattoe, E. 1993. An evolutionary simulation of price and quantity setting with methodological implications for economics. Unpublished MS, Lady Margaret Hall, Oxford University.
- Chauvin, Y. 1988. *Symbol acquisition in humans and neural (PDP) networks*. Unpublished doctoral dissertation, University of California, San Diego.
- Christaller, W. 1933. *Die Zentralen Orte in Süddeutschland*. Jena: Fischer.
- Cigler, A.J. & B.A. Loomis 1986. *Interest group politics* (2nd edn). Washington, DC: Congressional Quarterly Press.
- Clancey, W. 1989. Ontological commitments and cognitive models. *Proceedings of The eleventh annual conference of the cognitive science society*. Ann Arbor, Michigan.
- Clark, E. 1983. Convention and contrast in acquiring the lexicon. In *Cognitive development and the development of word meaning*, T.B. Seiler & W. Wannenmacher (eds), 67–89. Berlin: Springer.
- 1987. The principle of contrast: a constraint on language acquisition. In *Mechanisms of language acquisition*, B. MacWhinney (ed), 1–33. Hillsdale, New Jersey: Lawrence Erlbaum.
- Clark, K. & S. Gregory 1987. A relational language for parallel programming. In *Concurrent Prolog: collected papers*. E. Shapiro (ed.), 9–26. Cambridge, Mass.: MIT Press.
- Coase, R. 1937. The nature of the firm. *Economica* **4**, 386–405.
- Cohen, M.N. 1985. Prehistoric hunter-gatherers: the meaning of social complexity. In *Prehistoric hunter-gatherers: the emergence of cultural complexity*, T. Douglas-Price & J.A. Brown (eds), 99–119. New York: Academic Press.
- Coleman, J.S. 1993. The rational reconstruction of society. *American Sociological Review* **58**, 1–15.
- Collins, R. 1990. Stratification, emotional energy and the transient emotions. In *Research agendas in the sociology of emotions*, T.D. Kemper (ed.) 27–57. State University of New York Press.
- Collins, R.J. & D.R. Jefferson 1991. Ant farm: towards simulated evolution. In *Artificial life II*, C. Langton (ed), 579–601. Reading, Mass.: Addison-Wesley.
- Conte, R., & C. Castelfranchi 1993. Norms as mental objects: from normative beliefs to normative goals. Proceedings of MAAMAW '93. Neuchâtel: Switzerland.
- 1994. Mind is not enough: precognitive bases of social interaction. In *Simulating societies*, N. Gilbert & J. Doran (eds), 267–86. London: UCL Press.
- 1995. *Cognitive and social action*. London: UCL Press.

- Conte, R., M.Miceli, C.Castelfranchi 1991. Limits and levels of cooperation: disentangling various types of prosocial interaction. In *Decentralized AI—2*, Y.Demazeau & J.P.Muller (eds). Amsterdam: Elsevier.
- Conte, R. & J.Sichman. DEPNET: how to benefit from social dependence. *Journal of Mathematical Sociology*. Special issue on sociological algorithms, in press.
- Corbara, B. 1991. L'organisation sociale et sa genèse chez la fourmi *Ectatomma ruidum* Roger (Formicidae, Ponerinae). Doctoral thesis, University of Paris XIII. Grenoble: Microfiche ANRT/91/PA13/1007.
- Corbara, B., D.Fresneau, J-P.Lachaud, Y Leclerc 1991. Evolution de la division du travail dans les jeunes sociétés de *Neoponera apicalis* et d'*Ectatomma ruidum* (Formicidae, Ponerinae). *Actes Collectif Insectes Sociaux* **7**, 189–94.
- Couclelis, H.1985. Cellular worlds, a framework for modelling micro-macro dynamics. *Environment and Planning A* **17**, 585–96.
- 1986. Artificial intelligence in geography: conjectures on the shape of things to come. *The Professional Geographer* **38**, 1–11.
- 1988. Of mice and men: what rodent populations can teach us about complex spatial dynamics. *Environment and Planning A* **20**, 99–109.
- Courrège, P. 1965. Un modèle mathématique des structures élémentaires de parenté. *L'Homme* **5**, 248–91.
- Cusack, T.R. & R.J.Stoll 1990. *Exploring realpolitik: probing international relations theory with computer simulation*. Boulder, Colorado: Lynne Rienner.
- Davidson, R.H. & W.J.Oleszek 1985. *Congress and its members*. Washington, DC: Congressional Quarterly Press.
- Davis, R. & R.G.Smith 1983. Negotiation as a metaphor for distributed problem-solving. *Artificial Intelligence* **20**, 63–100.
- De Vries, J. 1984. *European urbanization, 1500–1800*. London: Methuen.
- Denaro, D. & D.Parisi 1994. *Imitation and cultural transmission in populations of neural networks*. Submitted to 1994 Conference of the Cognitive Science Society.
- Doran, J.E. 1992. A computational investigation of three models of specialization, exchange and social complexity. In *Representations in archaeology*, J-C.Gardin & C.S. Peebles (eds), 315–28. Bloomington, Indiana: Indiana University Press.
- Doran, J. & N.Gilbert 1994. Social simulation: an introduction. In *Simulating societies: the computer simulation of social phenomena*, N.Gilbert & J.Doran (eds), 1–18. London: UCL Press.
- Doran, J., H.Carvaja, Y.J.Choo, Y.Li 1991. The MCS.multi agent testbed: developments and experiments. In *Proceedings of the international working conference on cooperating knowledge based systems (CKBS)*, S.M.Deen (ed.), 240–54. London: Springer.
- Doran, J., M.Palmer, N.Gilbert, P.Mellars 1994. The EOS project: modelling Upper Palaeolithic social change. In *Simulating societies*, N.Gilbert & J.Doran (eds), 195–222. London: UCL Press.
- Dosi, G. & M.Egidi 1991. Substantive and procedural uncertainty: an exploration of economic behaviours in complex and changing environments. *Journal of Evolutionary Economics*, **1**, 145–68.
- Dosi, G., C.Freeman, R.R.Nelson, G.Silverberg, L.Soete (eds) 1988. *Technical change and economic theory*. London: Pinter.
- Dosi, G., L.Marengo, A.Bassanini, M.Valente (1993). *Microbehaviors and dynamical systems: economic routines as emergent properties of adaptive learning*. Berkeley, California: University of California Press, CCC working paper no. 93–13.
- Doyle, J. 1992. Rationality and its roles in reasoning. *Computational Intelligence* **8**, 376–409.
- Drogoul, A., B.Corbara, D.Fresneau 1992a. Applying ethomodelling to social organization in ants. In *Biology and evolution of social insects*, J.Billen (ed.), 375–83. Leuven University Press.

- Drogoul, A., J.Ferber, B.Corbara, D.Fresneau 1992b. A behavioural simulation model for the study of emergent social structures. In *Toward a practice of autonomous systems*, F.J.Varela & P.Bourgine (eds), 161–70. Cambridge, Mass.: MIT Press.
- Duby, G. 1984. *Histoire de la France urbaine*. Paris: le Seuil, 4 vols.
- Dunbar, R.I.M. 1993. Co-evolution of cortex size, group size and language in humans. *Behavioral and Brain Sciences* 16, 681–94.
- Durfee, E. 1988. *Coordination of distributed problem solvers*. Dordrecht: Kluwer.
- Durfee, E.H., V.R.Lesser, D.D.Corkill 1987. Coherent cooperation among distributed problem solvers. *IEEE Transactions on Computers* 36, 1275–291.
- 1989. Cooperative distributed problem solving. In *The handbook of artificial intelligence*, vol. IV, A.Barr, P.R.Cohen, E.A.Feigenbaum (eds). Reading, Mass.: Addison-Wesley.
- Durkheim, E. 1895. *The rules of sociological method*. Readings from Emile Durkheim. Chichester: Ellis Horwood.
- Egidi, M. 1992. Organizational learning and the division of labour. In *Economics, bounded rationality and the cognitive revolution*, M.Egidi & R.Marris (eds). Aldershot: Edward Elgar.
- Eibl-Eibesfeldt, I. 1967/1978. *Grundri der vergleichenden Verhaltensforschung: Ethologie*. Munich: Piper.
- Elster, J. 1982. Marxism, functionalism and game-theory: the case for methodological individualism. *Theory and Society* 11, 453–81.
- Emerson, E.A. 1990. Temporal and modal logic. In *Handbook of theoretical computer science*, J.van Leeuwen (ed.), 996–1072. Amsterdam: Elsevier.
- Emmeche, C. 1994. Is life a multiverse phenomena?. In *Artificial life III*, C.G.Langton (ed.), 553–68. Menlo Park, California: Addison Wesley.
- Epstein, J.M. 1985. *The calculus of conventional war: dynamic analysis without Lanchester theory*. Washington, DC: Brookings.
- Ferber, J. 1989. *Des objets aux agents*. Doctoral thesis, University of Paris VI.
- 1991. *Conception et programmation par objets*. Paris: Hermès.
- Ferber, J. & P.Carle 1991. Actors and agents as reflective concurrent objects: a MERING IV Perspective. *IEEE Transactions on Systems, Man and Cybernetics*.
- Ferber, J. & A.Drogoul 1992. Using reactive multi-agent systems in simulation and problem-solving. In *Distributed artificial intelligence: theory and praxis*, N.M.Avouris & L. Gasser (eds), 53–80. ECSC-EEC-EAEC, Brussels and Luxembourg.
- Fik, T.J. 1988. Hierarchical interaction: the modelling of a competing central place system. *Annals of Regional Science* 22, 48–69.
- Fik, T.J. & G.F.Mulligan 1990. Spatial flows and competing central places: towards a general theory of hierarchical interaction. *Environment and Planning A* 22, 527–49.
- Findler, N.V. 1987. Air traffic control: a challenge for artificial intelligence. *AI Expert Magazine* 2, 59–66.
1989. Distributed knowledge-based systems in manufacturing. Keynote speech, *Proceedings of the international conference on artificial intelligence in engineering*, Cambridge, England. Also reprinted in *Artificial intelligence in manufacturing*, G.Rzevski (ed.), 17–33. New York: Springer.
- 1990. *Contributions to a computer-based theory of strategies*. New York: Springer.
- 1991. *An artificial intelligence technique for fact retrieval: an application in medical knowledge processing*. Cambridge, Mass.: MIT Press.
- 1992. Distributed artificial intelligence: distributed planning and problem-solving, and multi-agent systems. Invited tutorial, *Proceedings of the first Singapore international conference on intelligent systems*, Singapore, 1992.
- Findler, N.V. & J.Gao 1987. Dynamic hierarchical control for distributed problem-solving. *Data and Knowledge Engineering* 2, 285–301.
- Findler, N.V. & Q.Ge 1989. Perceiving and planning before acting—an approach to enhance global network coherence. *International Journal of Intelligent Systems* 4, 459–70.

- Findler, N.V. & Q.Ge 1994. Distributed goal-oriented dynamic plan revision. *International Journal of Intelligent Systems* **9**, 183–210.
- Findler, N.V. & R.Lo 1986. An examination of distributed planning in the world of air traffic control. *Journal of Parallel and Distributed Processing* **3**, 411–31.
- 1991. A distributed artificial intelligence approach to air traffic control *IEE Proceedings, Part D, Control theory and applications* **138**, 515–24.
- 1993. Distributed air traffic control—Part I: Theoretical studies; and Part II: Explorations in a testbed. *Journal of Transportation Engineering* **119**, 681–94.
- Findler, N.V. & U.K.Sengupta 1992. Multi-agent planning and collaboration in dynamic resource allocation. *Proceedings of the first international conference on AI planning systems*, 305–6, University of Maryland, College Park.
- Findler, N.V. & J.Stapp 1992. A distributed approach to optimized control of street traffic signals. *Journal of Transportation Engineering* **118**, 99–110.
- Finger, M., M.Fisher, R.Owens 1992. METATEM at work: modelling reactive systems using executable temporal logic. Submitted to the *Sixth international conference on industrial and engineering applications of artificial intelligence and expert systems (IEA/AIE '93)*. Switzerland: Gordon and Breach.
- Fisher, M. 1993. Concurrent METATEM—a language for modelling reactive systems. In *Parallel architectures and languages, Europe (PARLE), Munich, Germany*, 185–96. Berlin: Springer.
- Fisher, M. & H.Barringer 1991. Concurrent METATEM Processes—a language for distributed AI. In *Proceedings of the European simulation multiconference*, Copenhagen, Denmark.
- Fisher, M. & R.Owens 1992. From the past to the future: executing temporal logic programs. In *Proceedings of logic programming and automated reasoning (LPAR)*, St Petersburg, Russia (published in *Lecture Notes in Computer Science* **624**).
- Fisher, M. & M.Wooldridge 1993. Specifying and verifying distributed intelligent systems. In *6th Portuguese conference on artificial intelligence (EPIA '93)*, 13–28. Berlin: Springer.
- Fletcher, R. 1986. Settlement in archaeology: world-wide comparison. *World Archaeology* **18**, 59–83.
- Floreano, D. 1993. Emergence of nest-based foraging strategies in ecosystems of neural networks. In *From animals to animats 2*, J.A.Meyer, H.L.Roitblat, S.W.Wilson (eds). Cambridge, Mass.: MIT Press.
- Fodor, J. 1976. *The language of thought*. Sussex: Harvester Press.
- Fontana, W. 1991. Functional self-organization in complex systems. In *1990 lectures in complex systems, SFI studies in the sciences of complexity, vol. III*, L.Nadel & D.Stein (eds), 40–26. Redwood City, California: Addison-Wesley
- Forrest, S. 1990. Emergent computation. *Physica D* **42**, 1–11.
- Forrest, S. & J.H.Miller 1990. Emergent behaviour in classifier systems. In *Emergent computation*, S.Forrest (ed.), Cambridge, Mass.: MIT Press.
- Frankhauser, P. 1990. *Beschreibung des Evolution urbaner Systeme mit der Mastergleichung*. Ph.D. thesis, University of Stuttgart.
- Fresneau, D., B.Corbara, J.P.Lachard 1989. Organisation sociale et structuration spatiale autour de Couvain chez *Pachycondyla apicalis* (Formicidae, Ponerinae). *Actes Collectif Insectes Sociaux* **5**, 83–92.
- Freyd, J. 1983. Shareability: the social psychology of epistemology. *Cognitive Science* **7**, 191–210.
- Gamble, C. 1991. The social context for European Palaeolithic art. *Proceedings of the Prehistoric Society* **57**, 3–5.
- Gasser, L., C.Braganza, N.Hermann 1987. MACE: a flexible testbed for distributed AI research. In *Distributed artificial intelligence.*, M.Huhns (ed), 119–52. San Mateo, California: Morgan Kaufmann.
- Geertz, C. 1973. Thick description: toward an interpretive theory of culture. In *The interpretation of cultures*, 3–30. New York: Basic Books.

- Georgeff, M.P. & F.F.Ingrand 1989. Decision-making in an embedded reasoning system. In *Proceedings of the eleventh international joint conference on artificial intelligence (IJCAI)*, 972–8. Detroit, USA.
- Georgeff, M.P. & A.L.Lansky 1987. Reactive reasoning and planning. In *Proceedings of the American association for artificial intelligence (AAAI)*.
- Gervet, J. 1992. Sociétés animales et société humaine: de l'évolution à l'histoire. In *Darwinisme et société: de l'évolution à l'histoire*, 501–18. Paris: Presses Universitaires de France.
- Giarratano, J. & G.Riley 1989. *Expert systems—principles and programming*. Boston, Mass.: PWS-Kent.
- Giddens, A. 1984. *The constitution of society*. Cambridge: PolityPress.
- Gilbert, M. 1987. Modelling collective belief. *Synthese* **73**, 185–204.
- 1989. *On social facts*. London: Routledge.
- Gilbert, N. & J.Doran (eds) 1994. *Simulating societies: the computer simulation of social phenomena*, London: UCL Press.
- Glace, N.S. & B.A.Huberman 1993. The outbreak of co-operation. *Journal of Mathematical Sociology* **17**, 281–302.
- Glowczewski, B. 1991. *Du rêve à la loi chez les Aborigènes. Mythes, rites et organisation sociale en Australie*. Paris: Presses Universitaires de France, coll. Ethnologies.
- Goldberg, A. & D.Robson 1983. *Smalltalk 80: the language and its implementation*. Reading, Mass.: Addison-Wesley.
- Goldberg, D.E. 1989. *Genetic algorithms in search, optimization, and machine learning*. Reading, Mass.: Addison-Wesley.
- Goody, J. 1977. *The domestication of the savage mind*. Cambridge: Cambridge University Press.
- Guérin-Pace, F. 1993. *Deux siècles de croissance urbaine*. Paris: Economica.
- Haag, G., M.Münz, D.Pumain, T.Saint-Julien, L.Sanders 1991. Interurban migration and the dynamics of a system of cities. *Environment and Planning* **24**, 181–98.
- Hägerstrand T. 1952. The propagation of innovation waves. *Lund studies in geography*. Ser.B.Human geography 4.
- Hamilton, W.D. 1964. The evolution of social behavior. *Journal of Theoretical Biology* **7**, 1–52.
- Hammond, M. 1990. Affective maximisation: a new macro—theory in the sociology of emotions. In *Research agendas in the sociology of emotions*, T.D.Kemper (ed.) 58–81. State University of New York Press.
- Hardin, G. 1968. The tragedy of the commons. *Science*, **162**, 1243–48.
- Harnad, S. 1994. Artificial life: synthetic vs virtual. In *Artificial life III*, C.G.Langton (ed), 239–52. Menlo Park, California: Addison Wesley.
- Hazlehurst, B. 1991. The cockpit as multiple activity system. Unpublished manuscript, University of California, San Diego.
- 1994. *Fishing for cognition: an ethnography of fishing practice in a community on the west coast of Sweden*. Unpublished doctoral dissertation, University of California, San Diego.
- Hebb, D.O. 1949. *The organization of behavior*. New York: John Wiley.
- Hepper, R.G. (ed.) 1991. *Kin recognition*. Cambridge: Cambridge University Press.
- Hewitt, C. 1977. Viewing control structures as patterns of passing messages. *Artificial Intelligence* **8**, 323–64.
- Hinton, G. & S.Becker 1989. An unsupervised learning procedure that discovers surfaces in random-dot stereograms. Unpublished manuscript, University of Toronto, Canada.
- Hirschman, A.O. 1970. *Exit, voice and loyalty: responses to decline in firms, organizations, and states*. Cambridge, Mass.: Harvard University Press.
- Hogeweg, P. & B.Hesper 1985. Socioinformatic processes: MIRROR modelling methodology. *Journal of Theoretical Biology* **113**, 311–30.
- Holland, J.H. 1992. Complex adaptive systems. *Daedalus* **121**, 17–30.
- 1975. *Adaptation in natural and artificial systems*. Ann Arbor, Michigan: University of Michigan Press.

- 1986. Escaping brittleness: the possibilities of general purpose learning algorithms applied to parallel rule-based systems. In *Machine learning II*, R.S.Michalski, J. G.Carbonell, T.M.Mitchell (eds). Los Altos, California: Morgan Kaufmann.
- Holland, J.H., K.J.Holyoak, R.E.Nisbett, P.R.Thagard 1986. *Induction: processes of inference, learning and discovery*. Cambridge, Mass.: MIT Press.
- Huberman, B.A. & N.S.Glance 1993. Diversity and collective action. In *Interdisciplinary approaches to complex nonlinear phenomena*, H.Haken & A.Mikhailov (eds). New York: Springer.
- Huff, J.O. 1976. A hierarchical migration model of population redistribution within a central place hierarchy. *Geographical Analysis* **8**, 231–54.
- Hutchins, E. 1980. *Culture and inference*. Cambridge, Mass.: Harvard University Press.
- 1990. The technology of team navigation. In *Intellectual teamwork: social and technical bases of cooperative work*, J.Galegher, R.Kraut, C.Egido (eds), 191–219. Hillsdale, New Jersey: Lawrence Erlbaum Associates.
- 1991. Social organization of distributed cognition. In *Perspectives on socially shared cognition*, L.Resnick, J.Levine, S.Teasley (eds), 283–387. Washington, DC: The American Psychological Association.
- 1993. How a cockpit remembers its speed. Unpublished manuscript, University of California, San Diego.
- in press. *Cognition in the wild*. Cambridge, Mass.: MIT Press.
- Hutchins, E. & B.Hazlehurst 1991. Learning in the cultural process. In *Artificial life II*, C.Langton, C.Taylor, D.Farmer, S.Rasmussen (eds), 689–706. Redwood City, California: Addison-Wesley.
- Hutchins, E. & T.Klausen in press. Distributed cognition in an airline cockpit. In *Cognition and communication at work*, Y.Engeström & D.Middleton (eds). New York: Cambridge University Press.
- Huth, P. 1988. *Extended deterrence and the prevention of war*. New Haven, Connecticut: Yale University Press.
- Jen, E. (ed.) 1990. *1989 lectures in complex systems*. Santa Fe Institute Studies in the Sciences of Complexity. Redwood City, California: Addison-Wesley.
- Johnson, A.W. & T.Earle 1987. *The evolution of human societies: from foraging group to agrarian state*. Stanford, California: Stanford University Press.
- Johnson-Laird, P.N. 1983. *Mental models*. Cambridge, Mass.: Harvard University Press.
- Kahneman, D. & A.Tversky 1984. Prospect theory: an analysis of decisions under risk. *Econometrica* **47**, 263–91.
- Kahneman, D., P.Slovic, A.Tversky 1982. *Judgement under uncertainty: heuristics and biases*. New York: Cambridge University Press.
- Kauffman, S. 1988. The evolution of economic webs. In *The economy as an evolving complex system*, P.W.Anderson & K.J.Arrow (eds), 125–45. Reading, Mass.: Addison-Wesley.
- Kaufman, L. & P.J.Rousseuw 1990. *Finding groups in data: an introduction to cluster analysis*. New York: John Wiley.
- Kennedy, P. 1987. *The rise and fall of the great powers*. New York: Random House.
- Kinny, D., M.Ljungberg, A.Rao, E.Sonenberg, G.Tidhar, E.Werner 1992. Planned team activity. *Proceedings of MAAMAW '92 conference.*, July 29–31, S.Martino al Cimino (Viterbo), Italy.
- Koza, John R. 1992. *Genetic programming*. Cambridge, Mass.: MIT Press.
- Kreps, D., P.Milgrom, J.Roberts, R.Wilson 1982. Rational cooperation in the repeated prisoner's dilemma. *Journal of Economic Theory* **27**, 245–52.
- Kugler, R.L. 1990. NATO's future rôle in Europe—toward a more political alliance in a stable "11/2 bloc" system. *Technical report R-3923-FF*, The Rand Corporation.
- Lachaud, J-P. & D.Fresneau 1987. Social regulation in ponerine ants. In *From individual to collective behavior in social insects*, J.M.Pasteels & J-L.Deneubourg (eds), *Experientia Supplementum* **54**, 197–217.
- Lakoff, G. 1987. *Women, fire, and dangerous things*. Chicago: University of Chicago Press.

- Lanchester, F.W. 1916. *Aircraft in warfare: the dawn of the first arm*. London: Constable. Key chapters reprinted as: Mathematics in warfare. *The world of mathematics*, J. Newman (ed.) 1956. New York: Simon and Schuster.
- Langton, G.C. (ed.) 1994. *Artificial life III*. Menlo Park, California: Addison Wesley
- Levinson, S.. in press. Interactional biases in human thinking. In *The social origins of intelligence*, E.Goody (ed.). New York: Cambridge University Press.
- Lewis, D. 1969. *Convention*. Cambridge, Mass.: Harvard University Press.
- Lomborg, B. in press. Game-theory vs multiple agents: the iterated prisoner's dilemma. In *Decentralised AI—4*, C.Castelfranchi & E.Werner (eds). Berlin: Springer.
- Lumsden, C.J. & E.O.Wilson 1985. The relation between biological and cultural evolution. *Journal of Social and Biological Structures* **8**, 343–59.
- March, J.G. & H.A.Simon 1958. *Organizations*. New York: John Wiley.
- Marengo, L. 1992. Coordination and organizational learning in the firm. *Journal of Evolutionary Economics* **2**, 313–26.
- Maruichi, T., M.Ichikawa, M.Tokoro 1990. Modelling autonomous agents and their groups. In *Decentralized AI—proceedings of the first European workshop on modelling autonomous agents and multi-agent worlds (MAAMAW)*, Y.Demazeau & J.P.Muller (eds), 215–34. Amsterdam: Elsevier/North-Holland.
- McCain, R.A. 1979. Reflections on the cultivation of taste. *Journal of Cultural Economics* **3**, 30–52.
- 1981. Cultivation of taste, catastrophe theory, and the demand for works of art. *American Economic Review* **71**, 332–4.
- 1982. Optimal subsidies to the arts in a shortsighted world. *Journal of Cultural Economics* **6**, 15–32.
- 1986. Game theory and the cultivation of taste. *Journal of Cultural Economics* **10**, 1–15.
- 1991a. Groping: toward a behavioral metatheory of choice. In *Handbook of behavioral economics*, R.Franz, H.Singh, J.Gerber (eds). Greenwich, Connecticut: JAI Press, 495–508.
- 1991b. Deontology, consequentialism, and rationality. *Review of Social Economy* **49**, 168–95.
- 1991c. A linguistic theory of rationality. *Social Science Information* **30**, 233–55.
- 1992a. *A framework for cognitive economics*. New York: Praeger Publishers.
- 1992b. Cultivation of taste and bounded rationality: overview. Third ICARE conference on the economics of the arts, Venice, Italy (December).
- McCain, R.A. Cultivation of taste and bounded rationality: some computer simulations. *Journal of Cultural Economics*, in press.
- McClelland, J. 1988. *Explorations in parallel distributed processing*. Cambridge, Mass.: MIT Press.
- McClelland, J., D.Rumelhart, the PDP Group 1986. *Parallel distributed processing: vol. 2, explorations in the microstructure of cognition*. Cambridge, Mass.: MIT Press.
- Mellars, P.A. 1985. The ecological basis of social complexity in the Upper Paleolithic of southwestern France. In *Prehistoric hunter-gatherers: the emergence of cultural complexity*, T. D.Price & J.A.Brown (eds), 271–97. New York: Academic Press.
- Menczer, F. & D.Parisi 1992. Recombination and unsupervised learning: effects of cross—over in the genetic optimization of neural networks. *Network* **3**, 424–42.
- Mercier, P. 1968. Anthropologie sociale et culturelle. *Ethnologie Générale*, 887–1036. Paris: Ed. Gallimard.
- Minsky, M. 1985. *Society of the mind*. New York: Simon and Schuster.
- Mokken, R.J. & F.N.Stokman 1985. Legislative analysis: methodology for the analysis of groups and coalitions. In *Coalition formation*, H.A.M.Wilke (ed.). Amsterdam: Elsevier.
- Moriconi, F.. 1993. *L'urbanisation du monde, 1950–1990*. Paris: Economica.
- Morrill, R.L. 1962. Simulation of central place pattern over time. In *Proceedings of the IGU symposium in urban geography*, K.Norborg (ed.), 109–120. Lund: Glerup.

- Moses, Y. & M.Tenneholtz 1992. On computational aspects of artificial social systems. Proceedings of the 11th DAI workshop, Glen Arbor (February).
- Moszkowski, B. 1986. *Executing temporal logic programs*. Cambridge: Cambridge University Press.
- Murdock, P.F. 1949. *Social structure*. New York: Free Press.
- Nelson, R.R. & S.G.Winter 1982. *An evolutionary theory of economic change*. Cambridge, Mass.: Harvard University Press.
- Nowak, A. & B.Latané 1993. Simulating the emergence of social order from individual behaviour. In *Simulating societies: the computer simulation of social phenomena*, N.Gilbert & J. Doran (eds). London: UCL Press.
- Nye, J.S., Jr 1990. *Bound to lead, the changing nature of American power*. New York: Basic Books.
- O'Neill, J. (ed.) 1973. *Modes of individualism and collectivism*. London: Heinemann.
- Oleszek, W.J. 1989. *Congressional procedures and the policy process*. Washington, DC: Congressional Quarterly Press.
- Parisi, D. in press. Science as history. *Information en Sciences Sociales*.
- Parisi, D., F.Cecconi, S.Nolfi 1990. Econets: neural networks that learn in an environment. *Network* **1**, 149–68.
- Parisi, D., S.Nolfi, S.Cecconi 1991. Learning, behavior, and evolution. In *Toward a practice of autonomous systems*, F.Varela & P.Bourgine (eds). Cambridge, Mass.: MIT Press.
- Parisi, D., U.Piazzalunga, F.Cecconi, M.Gerace in press. Spatial distribution and sociality. In *Artificial sociality*, C.Castelfranchi & E.Werner (eds). New York: Springer.
- Parsons, T. 1952. *The Social System*. New York: Free Press, Glencoe.
- Pingle, Mark 1992. Imitation versus rationality. Conference of Society for the Advancement of Socio-Economics, Irvine, California.
- Pirenne, H. 1925. *Les villes médiévales*. Reprinted 1971. Paris: Presses Universitaires de France.
- Pnueli, A. 1977. The temporal logic of programs. In *Proceedings of the eighteenth symposium on the foundations of computer science*.
1986. Specification and development of reactive systems. In *Information processing '86*. Amsterdam: Elsevier/North-Holland.
- Powell, R. 1990. *Nuclear deterrence theory: the search for credibility*. Cambridge: Cambridge University Press.
- Pred, A. 1977. *City system in advanced economics*. London: Hutchinson.
- Pumain, D. 1991. Settlement size dynamics in history. 7th European colloquium on theoretical and quantitative geography, Stockholm.
- Radner, R.R. 1992. Hierarchy: the economics of managing. *Journal of Economic Literature* **30**, 1382–415.
- Rao, A.S. & M.P.Georgeff 1991. Modelling agents within a BDI-architecture. In *International conference on principles of knowledge representation and reasoning (KR)*, R.Fikes & E. Sandewall (eds), 473–84. Cambridge, Mass.
- Rawlins, G.J.E. (ed.) 1991. *Foundations of genetic algorithms*. San Mateo, California: Morgan Kaufman.
- Reddaway, W.B. 1936. Irrationality in consumers' demand. *Economic Journal* **46** (September), 419–23.
- Reddy, M. 1979. The conduit metaphor: a case of frame conflict in our language about language. In *Metaphor and thought*, A.Ortony (ed.), 284–324. New York: Cambridge University Press.
- Rumelhart, D., G.Hinton, R.Williams 1986. Learning internal representations by error propagation. In *Parallel distributed processing: vol. 1, foundations*, D.Rumelhart, J. McClelland, the PDP Group (eds), 318–62. Cambridge, Mass.: MIT Press.
- Rumelhart, D., J.McClelland, The PDP Group 1986. *Parallel distributed processing: Volume 1, foundations*. Cambridge, MA: MIT Press.
- Sanders, L. 1992. *Systèmes de villes et synergétique*. Paris: Economica.

- Schelling, T.C. 1960. *The strategy of conflict*. Oxford: Oxford University Press.
- Schumpeter, J. 1934. *Theory of economic development*. Cambridge, Mass.: Harvard University Press.
- Shoham, Y. 1989. Time for action: on the relation between time, knowledge and action. In *Proceedings of the eleventh international joint conference on artificial intelligence (IJCAI)*, Detroit, USA, 954–9.
- 1990. Agent oriented programming. Technical Report STAN-CS-1335–90, Department of Computer Science, Stanford University, California, USA.
- Shoham, Y. & M.Tenneholtz 1992a. On the synthesis of useful social laws for artificial agent societies (preliminary report). *Proceedings of the AAAI Conference*, 276–81.
- 1992b. Emergent conventions in multi agent systems: Initial experimental results and observations. *Proceedings of the 3rd International Conference on KR&R*. Cambridge, Mass., 225–32.
- Shore, B. 1991. Twice-born, once conceived: meaning construction and cultural cognition. *American Anthropologist* **93**, 9–27.
- Simon, H.A. 1962. The architecture of complexity. *Proceedings of the American Philosophical Society* **106**, 467–82.
- 1981. *The sciences of the artificial*. Cambridge, Mass.: MIT Press.
- Slobin, D. 1985. Crosslinguistic evidence for the language-making capacity. In *The crosslinguistic study of language acquisition*, D.Slobin (ed.), 1157–260. Hillsdale, New Jersey: Lawrence Erlbaum.
- Sobel, J. 1992. How to count to one thousand. *Economic Journal* **102**, 1–8.
- Sperber, D. 1985. Anthropology and psychology: towards an epidemiology of representations. *Man* **20**, 73–89.
- 1987. Les sciences cognitives, les sciences sociales et le matérialisme, in *Introduction aux sciences cognitives*, 39–446. Paris: Ed. Gallimard.
- Stein, D.L. (ed.). 1989. *Lectures in the sciences of complexity*. Redwood City, California: Addison-Wesley.
- Stein, D. & L.Nadel (eds) 1991. *1990 lectures in complex systems*. Santa Fe Institute Studies in the Sciences of Complexity. Redwood City, California: Addison-Wesley.
- Stigler, G. and G.Becker 1977. De gustibus non est disputandum. *American Economic Review* **67**, 67–90.
- Terna, P. 1993. Imitazione nelle reti neurali. *Sistemi Intelligenti* **5**, 169–73.
- The raulaz, G., G.Goss, J.Gervet, J-L.Deneubourg 1990. Task differentiation in *Polistes* wasp colonies: a model for self-organizing groups of robots. In *Simulation of adaptive behaviour: from animals to animats*, J-A.Meyer & S.W.Wilson (eds), 346–55. MIT Press/Bradford Books.
- Thom, R. 1972. *Stabilité structurelle et morphogenèse: Essai d'une théorie générale des modèles*. New York: Benjamin.
- Thomas, B., Y.Shoham, A.Schwartz, S.Kraus 1991. Preliminary thoughts on an agent description language. *International Journal of Intelligent Systems* **6**, 497–508.
- Tilly, C. 1985. War making and state making as organized crime. In *Bringing the state back in*, P.B.Evans, D.H.Rueschemeyer, T.A.Skocpol (eds). Cambridge: Cambridge University Press.
- 1990. *Coercion, capital, and European states, AD 990–1990*. Oxford: Basil Blackwell.
- Tobler, W.R. 1979. Cellular geography. In *Philosophy of geography*, S.Gale & G.Olsson (eds), 279–386. Dordrecht: Reidel.
- Trivers, R.L. 1985. *Social evolution*. Menlo Park, California: Benjamin Cummings.
- Tuomela, R. 1991. We will do it: an analysis of group-intentions. *Philosophy and Phenomenological Research* **51**, 249–77.
- 1992. Group beliefs. *Synthese* **91**, 285–318.
- Vines, G. 1993. Planet of the free apes? *New Scientist* **138**, 5 June 1993, 39–42.
- von Hayek, F.A. 1948. Economics and knowledge. In *Individualism and economic order*. Chicago: University of Chicago Press.

- Waal, F.de 1982. *Chimpanzee politics*. London: Jonathan Cape.
- Waldrop, M.M. 1992. *Complexity: the emerging science at the edge of chaos*. New York: Simon & Schuster.
- Watkins, W.J.N. 1955. Methodological individualism: a reply *Philosophy of science* **22**.
- Werner, E. 1990. Cooperating agents: a unified theory of communication and social structure In *Distributed artificial intelligence: vol. II*, L.Gasser & M.N.Huhns (eds). San Mateo, California: Morgan Kaufmann.
- Werner, G.M & M.G.Dyer 1991. Evolution of communication in artificial organisms. In *Artificial life II*, C.G.Langton, C.Taylor, J.D.Farmer, S.Rasmussen (eds). Reading, Mass.: Addison-Wesley.
- White, R.W. 1977. Dynamic central place theory: results of a simulation approach. *Geographical Analysis* **3**, 226–46.
- 1978. The simulation of central place dynamics: two-sector systems and the rank-size distribution. *Geographical Analysis* **10**(2), 201–8.
- 1991. A cellular automata approach to the evolution of urban land use patterns. 7th European colloquium on theoretical and quantitative geography, Stockholm.
- Williamson, O. 1975. *Markets and hierarchies: analysis and antitrust implications*. New York: Free Press.
- Wilson, E.O. 1985. The sociogenesis of insect societies. *Science* **228**, 1489–95.
- Zeigler, B.P. 1990. *Object-oriented simulation with hierarchical, modular models: intelligent agents and endomorphic systems*. New York: Academic Press.

Index

abstraction 123
action 147, 155, 215
adaptive cognition 10
adaptive learning 51
agency 8
aggression 253
alliance 21, 110, 113, 212
alliance, strategic 225
alternative futures 121
altruism 210, 239
ant 192
anthropomorphism 95
archaeology 2
artificial ethology 205
artificial intelligence 88
artificial life 2
artificial society 2, 6, 12
attachment behaviour 239
autoassociator 162
awareness 10
axiomatization 61

back-propagation of error 164
bias, sub-cognitive 10
bucket-brigade mechanism 56

cellular automata 88, 147
central place theory 87
chaos 123
chromosome 128
classifier systems 51
clobbering 106
co-evolution 41, 148
co-operation 6, 10, 210, 266, 268
co-ordination 7, 40, 49, 110
coalition formation 235
coercion 21
cognition 104, 121, 157
 adaptive 10
cognitive agent 157, 266
cognitive anthropology 158
cognitive capabilities 49
cognitive capacity 137

- cognitive processes 128, 158
- cognitive science 126, 158
- collaboration 7
- commitment 24, 29, 235, 281
- communication network 94
- community 182, 213
- competition 44, 219, 265, 268
- competitive selection 148
- complex adaptive system 19, 150
- complexity 124
- concurrency 270
- condition-action rule 52, 147
- conduit theory of meaning 181
- connectionist networks 162
- consumer demand 126
- cultural diversity 84
- cultural evolution 84
- cultural inheritance 60

- decision rule 23
- decision-making 103
- dialect 178
- differential equations 87, 90
- discretionary behaviour 42
- distributed artificial intelligence 89, 103, 144, 212, 268
- distributed cognition 160
- distributed decision-making 41
- division of labour 40, 190, 204
- dominance relationship 204

- economic organization 40
- economic rationality 121
- economic theory 126
- economics
 - experimental 126
 - neo-classical 42
- emergence 2, 8, 41, 89, 103, 144, 147, 212, 240, 266
- emergent behaviour 212
- emergent organization 19
- emergent phenomenon 137
- emergent political actors 35
- emergent properties 19, 144
- emotion 121
- emulation 129
- evolution 68, 84, 87, 93, 127, 212, 229, 239
 - cultural 84
- exogamy 60
- experiment 83, 126
- experimental economics 126
- explanation, functionalist 9
- exploitation 48

exploratory simulation 4

feature extraction 164
feed-down 149
filibuster 226
formal models 157
functionalist explanation 9

game theory 3, 5, 20
gene 244
genetic algorithm 2, 51, 66, 127
genetic coding 84
geography 86
goal structure 216
groping model 131

heuristic rules 126
hierarchy 113, 204
human societies 144
hunter-gatherers 104, 251

imitation 250
imperial overstretch 32
individualistic approach 11
institution 151
inter-agent communication 124
interaction 155

kin relations 239
kin selection theory 239
kinship 59, 104
kinship rule 60
kinship structure 60

language 155, 157, 177
 computer 268
language acquisition 179
leader-follower relationship 110
learning 46, 148, 164, 177, 250
learning-by-consuming 134
levels of emergence 149
lexicon 157

management 126
market 41, 94
meaning 184
mental representation 104, 266
mentalism 282
meta-norm 213, 227

methodological holism 82, 145, 152
methodological individualism 5, 82, 146, 152
methodology 4, 121
micro-macro link 11, 88, 145
migration 95
model, social 105
multi-agent modelling 88, 190
multiple agent plan 105
mutation 53, 84, 128, 148
mutual aid 265

natural selection 84
neo-classical economics 42
neural net 2, 162, 240
norm 145, 212, 226, 252
norm generation 228
normative reasoning 266

object-orientation 192
optimization 3

permutation group 60
plan, multiple agent 105
planning 42, 105, 108
politics 2, 19
prehistory 104
prestige 105
Prisoner's Dilemma 6, 22
problem-solving 43
production 94
public hidden units 187

random walk 127
rational choice 20
rationality 126
reactive system 269
reflexivity 144
representation 159, 183
 mental 104, 266
returns to scale 45
rule of interaction 64
rule
 condition-action 52
 heuristic 126
 kinship 60
 utilitarian 259

selection 148
settlement system 86, 91
shared symbols 183

- simulation
 - exploratory 4
 - multi-agent 190
 - stochastic 25
- Smalltalk 95
- social behaviour 82, 239
- social cognition 157
- social complexity 103
- social conflict 265
- social dynamics 104
- social group 246
- social impact 148
- social interaction 183
- social learning 3
- social model 105, 153
- social representation 145
- social strategy 3
- social structure 147, 204, 238
- sociality 4
- sociogenesis 190, 195, 205
- sociology 144, 238
- sociotomies 209
- space 239
- spatial distribution 86
- spatial relationships 59
- stochastic simulation 25
- stock market 213
- strategic alliance 225
- structural functionalism 5, 11
- structuration 146, 152, 154
- structure 147, 155
- structuring properties 147
- sub-cognitive bias 10
- symbol 157
- symbolic representation 183
- system, settlement 91

- target of simulation 2
- teleological conservatism 131
- teleologically conservative algorithm 132
- temporal logic 269
- theory building 122
- theory
 - central place 87
 - conduit 181
 - economic 126
 - kin selection 239
 - structuration 154
- tragedy of the commons 19
- tribute model 21

urbanization 86
utilitarian rule 259
utility 126, 252
utility function 133

validation 196
validity 121