


JACQUES LESOURNE
ANDRÉ ORLÉAN
BERNARD WALLISER

Evolutionary Micro- economics



Springer

Evolutionary Microeconomics

Jacques Lesourne · André Orléan
Bernard Walliser

Evolutionary Microeconomics

In cooperation with
Paul Bourguine, Emmanuelle Fauchart,
Jean-François Laslier, Luigi Marengo,
François Moreau, and Gisèle Umbhauer

With 33 Figures and 12 Tables

 Springer

Professor Dr. Jacques Lesourne
Conservatoire National des Arts et Métiers
2, rue Conté
75003 Paris
France
jolesourne@wanadoo.fr

Professor Dr. André Orléan
Paris-Jourdan Sciences Economiques (PSE)
48, boulevard Jourdan
75014 Paris
France
orlean@pse.ens.fr

Professor Dr. Bernard Walliser
Ecole Nationale des Ponts et Chaussées
28, rue des Saints-Pères
75007 Paris
France
walliser@mail.enpc.fr

This book was published under the title
“Leçons de microéconomie évolutionniste”
by J. Lesourne, A. Orléan, B. Walliser (eds.). Odile Jacob, 2002

ISBN-10 3-540-28536-9 Springer Berlin Heidelberg New York
ISBN-13 978-3-540-28536-6 Springer Berlin Heidelberg New York

Cataloging-in-Publication Data
Library of Congress Control Number: 2006922868

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, reuse of illustrations, recitation, broadcasting, reproduction on microfilm or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springeronline.com

© Springer Berlin · Heidelberg 2006
Printed in Germany

The use of general descriptive names, registered names, trademarks, etc. in this publication does not imply, even in the absence of a specific statement, that such names are exempt from the relevant protective laws and regulations and therefore free for general use.

Hardcover-Design: Design & Production, Heidelberg

SPIN 11545491 43/3153-5 4 3 2 1 0 – Printed on acid-free paper

The authors

Paul Bourguine, *research fellow at CREA, Ecole Polytechnique, contributed to chapter 1.*

Emmanuelle Fauchart, *research fellow at Conservatoire National des Arts et Métiers, wrote chapter 6.*

Jean-François Laslier, *research director at CNRS (CECO, Ecole Polytechnique), contributed to chapters 2 and 3.*

Jacques Lesourne, *professor at Conservatoire National des Arts et Métiers, fellow of the Econometric Society, wrote chapter 4 and contributed to introduction and chapter 8.*

Luigi Marengo, *professor at St. Anna School of Advanced Studies, Pisa, wrote chapter 7.*

François Moreau, *professor at Conservatoire National des Arts et Métiers, wrote chapter 9.*

André Orléan, *research director CNRS (PSE, CNRS-EHESS-Ecole Normale Supérieure-ENPC), wrote introduction, chapter 5 and contributed to chapters 2 and 8.*

Gisèle Umbhauer, *maître de conférences at University of Strasbourg, contributed to chapter 3.*

Bernard Walliser, *professor at Ecole Nationale des Ponts et Chaussées and at Ecole des Hautes Etudes en Sciences Sociales, contributed to introduction, chapters 1, 3 and 8.*

Contents

Introduction	1
The standard paradigm	2
Towards an evolutionary paradigm	6
Presentation of the book	7
References	9

Part I: The basic concepts

1 Individual decision	13
1.1 Background and problems	14
1.2 Canonical principles	22
1.3 Some models	31
1.4 Theses and conjectures	39
References	41
2 The elementary market	43
2.1 Background and problems	43
2.2 Canonical principles	48
2.3 Some models	50
2.4 Theses and conjectures	64
References	65
3 Game situations	67
3.1 Background and problems	68
3.2 Canonical principles	76
3.3 Some models	87
3.4 Theses and conjectures	108
References	110

Part II: The markets

4	Market with irreversibilities	115
4.1	Background and problems	116
4.2	Canonical principles	117
4.3	Some models	118
4.4	Theses and conjectures	128
	References	129
5	Mimetic interactions	131
5.1	Background and problems	131
5.2	Canonical principles	140
5.3	Some models	147
5.4	Theses and conjectures	169
	References	171
6	Competition between firms	173
6.1	Background and problems	173
6.2	Canonical principles	180
6.3	Some models	182
6.4	Theses and conjectures	196
	References	198

Part III: The institutions

7	Organization of the firm.....	203
7.1	Background and problems	203
7.2	Canonical principles	205
7.3	Some models	207
7.4	Theses and conjectures	232
	References	233

8 Emergence of institutions 237

8.1 Background and problems 237

8.2 Canonical principles 242

8.3 Some models 247

8.4 Theses and conjectures 256

References 257

9 State and economic system regulation 259

9.1 Background and problems 259

9.2 Canonical principles 263

9.3 Some models 265

9.4 Theses and conjectures 282

References 287

Epilogue..... 291

Introduction

The development of science offers numerous examples of “scientific revolutions” (Kuhn, 1970), which lead to deep changes in the existing “paradigms”. During these revolutions, the prior knowledge, far from being abandoned since it is obsolete, is often reinterpreted in the new paradigm as a limit case of a broader representation. The theory of restricted relativity, proposed by A. Einstein, is an exemplary illustration of such a situation. It elaborates a new conception of energy, mass and time breaking up radically with all was already accepted. Nevertheless, the formula of classical mechanics continues to be valid when considering speeds that are much lower than the light speed. Is microeconomics involved in such a revolution and changing its paradigm? One has to be careful when answering such a question since economics is far from being able to claim the same scientific standards than the hard sciences. Moreover, it is hazardous to speak of a revolution while it is already on the way.

Even if the conception defended in this book is clearly grounded on a refoundation of microeconomics, its point of view is more modest. It rests on four observations: (a) it exists a “standard paradigm” constructed around three key concepts, optimizing rationality, equilibrium and market efficiency, which frames the main classical works in microeconomics; (b) the empirical limits of such a paradigm are obvious since it is unable to explain some major observed economic phenomena; (c) several original models are already available in order to explain at least some of these phenomena; (d) these models express a coherent project, looking as an original paradigm, which integrates standard microeconomics as a limit case. The book aims at designing this new paradigm, which progressively emerges at the crossroads of various modeling streams: evolutionary, cognitivist and institutionalist.

Characterized by its departure from classical economics, the present project has still to be distinguished from another one which inspires today an important part of microeconomics, the “extended standard theory” (Favereau, 1989). The last aims at developing the study of organizations and institutions while staying in the standard paradigm, and is well illustrated by the modern theory of contracts and incentives. The first is interested in institutions too, but is running away from the standard view in

more profound aspects. However, in order to prevent ambiguity, it is necessary to state that it still shares with the classical or extended approach number of principles and problems, for instance the adoption of methodological individualism or a specific interest in price formation.

This introduction is devoted to making precise the four statements which justify the project. The first section is related to assertions (a) and (b) and the second to assertions (c) and (d). A third section presents the structure of the book and its pedagogical aims.

The standard paradigm

The existence of a “standard paradigm” is not unanimously recognized by economists. As spelt out by R. Nelson and S. Winter (1982, p.6), some economists “would strenuously deny there is an orthodox position providing a narrow set of criteria that are conventionally used as a cheap and simple test for whether an expressed point of view on certain economic questions is worthy of respect; or, if there is such an orthodoxy, that it is in any way enforced”. It is right that this notion is mainly put forward by economists willing to differentiate their work from “normal science”. For that reason, it may be endowed with a high critical charge which makes it suspicious to “orthodox” economists. In many cases, it sustains a view which goes beyond a simple objective description of the economists’ achievements in order to induce a new way of dealing with their discipline. This motivation is shared by the authors of this book.

One should nevertheless not under-estimate the difficulties associated with such a goal. Microeconomics is a rapidly developing science which makes use of various concepts and principles in order to cover an always broader field. It is not possible to reduce it to a few notions without making a caricature of it. However, it seems possible to bring out what may be called an “orthodox way” to deal with the usual microeconomic problems. On one hand, it proceeds to a systematic appeal to optimizing rationality and equilibrium as two general categories allowing to think all economic phenomena. On the other hand, it develops a theory of trade order dominated by the assumption of market efficiency. This triptyque will be exploited in order to analyze the standard paradigm.

Optimizing rationality

Adopted by the orthodox approach, optimizing rationality assumes that all agents are endowed with an objective function that they maximize with re-

gard to some constraints. It expresses a specific form of instrumental rationality since it deals with the adequation achieved by an agent between the means at his disposal and the aims he pursues. In order to qualify it, H. Simon (1982) speaks of “substantive rationality” since it is exclusively concerned with the results of the choice process. It is opposed to “procedural rationality” which is mainly interested in the deliberation process leading to some choice. Optimizing rationality is involved in a lot of economic models such as profit maximization by a firm submitted to technological constraints, (discounted) utility maximization by a consumer trading under a (intertemporal) budget constraint, expected utility maximization by a financial investor acting under uncertainty and constrained by his initial wealth.

Optimizing rationality is grounded on several implicit assumptions concerning the agent’s cognition when adapting to market exchanges. First, the agent is always confronted to well defined problems in a transparent environment. Such an assumption is unrealistic since the agent has to search for various information in order to make a more precise view of his environment. He has even to define more accurately what are his own opportunities and preferences since they are not initially given. Second, the agent is endowed with infinite computing capacities. This is really a distinctive feature of the standard approach: the more the situation is complex, the more are the agents endowed with a sophisticated and performant rationality. In the limit, all actual interactions between agents are perfectly simulated by the agents themselves. Such an assumption is again unrealistic since the agents face computation constraints. Hence, optimizing rationality appears at best as a contextual limit case, for instance when the agents are involved in a “small world”.

Equilibrium

In the orthodox view, an equilibrium state is defined as a realizable economic configuration in which no agentw can do better by modifying unilaterally his action. Hence, once an equilibrium state is established, no agent has an incentive to deviate from it. Such a property explains the importance given to that concept: an equilibrium state tends to survive in the absence of changing exogenous factors. In other terms, an equilibrium state is a fixed point of the economic dynamics in a stationary environment. As for optimizing rationality, equilibrium is a general concept which is illustrated in many specific economic models such as Walrasian competitive equilibrium, Cournot oligopolistic equilibrium, monopoly equilibrium or fixed price equilibrium.

Although the study of equilibrium states leads to some fundamental results, the orthodox view stays silent about the way an equilibrium state is reached. The dynamics of what happens out of equilibrium receives little attention. The Walrasian equilibrium is a good illustration of such a lack of understanding. Even if the study of its existence and multiplicity has been fruitfully achieved, and constitutes a powerful achievement of the standard paradigm, it does not exist a satisfying representation of the exchange process leading to it. The Walrasian auctioneer device is, in this respect, very insufficient since it appears as a fictitious entity. In fact, modelling the off equilibrium process is a fundamental requirement, for instance to prove that a competitive economy always stays in a neighborhood of some equilibrium state. Even if it is natural to think that an economy tends to deviate from any non equilibrium position, it does not follow that it converges naturally toward some equilibrium state. The formal study of dynamical systems concludes to the existence of a great variety of attractors even when some fixed point exists somewhere. Hence, by lack of a satisfying analysis, nothing proves that a complete flexibility of prices necessarily leads the economic system to a general equilibrium. Moreover, even for those who stick to the idea that an economy tends to some equilibrium state, the question of the selection between multiple equilibria is still open. This is a common situation in contemporary models. In that case, only a dynamical study is able to select what equilibrium state will prevail as a function of the initial conditions and the history.

When combining optimizing rationality and equilibrium, one obtains an abstract view which seems very far from publicly observed features of a concrete economy. Some orthodox economists were fully aware of that apparent hiatus. It is the case for M. Friedman (1953) in a famous methodological article entitled "The methodology of positive economics". Noticing that the orthodox theory is built on assumptions in obvious contradiction with plain observations, he nevertheless defends them. According to his *as if* argument, what is important is less the adequation of assumptions to observations than the expectations derived from them. Even if the actual behaviors may differ from optimizing rationality, everything goes as if it were valid: the prices and exchanged quantities expected by the model are in conformity with the observations. Such a methodological position is called instrumentalist as opposed to realistic since the assumptions are not chosen for their empirical validity, but are considered as instruments allowing the modeller to infer empirical phenomena. Moreover, M. Friedman and others tried to justify such a position by stressing that non optimizing behaviors may exist, but have a weak impact since the rules of

competition necessarily lead to their removal. According to these theoreticians, modeling an economy as exclusively formed of maximizing agents may be instantaneously wrong, but constitutes nevertheless a good approximation in actual economies. However, if it is the evolution process which produces optimizing behaviors, one has to model it explicitly. Modeling has to think simultaneously the economic phenomena and the conditions of their emergence.

Market efficiency

According to the orthodox approach, the competitive market is the fundamental institutional device allowing an efficient resolution of all coordination problems encountered by mutual exchanges. More profoundly, the competitive equilibrium is endowed with the status of a norm. On one hand, it constitutes the basic reference for evaluating all other equilibrium notions. The notion of “market failure” precisely refers to conditions not satisfied in a competitive market: incomplete information, imperfect competition, sluggish prices. On the other hand, it suggests the way to deal with any new difficulty. The recommendation is to establish or reestablish the institutional conditions for obtaining an equivalent of a competitive equilibrium. For instance, the distribution of “rights to pollute” consists in creating a new market in order to solve an unusual environmental problem.

Such an approach, even if relevant in some instances, conceals great dangers and may lead to important biases. On one hand, the obtention of market efficiency, either allocative or informational, is still an open question and not a dogma. Even when involved with a rigorous proof, as for Paretian efficiency of a competitive market, it rests on many restrictive assumptions on behaviors as well as on goods. On the other hand, the identification of an economy to markets leads to a distorted view of the trade order. It is wrong to consider the market as a natural entity, as a necessary by-product of the rationality of mutual exchanges. The market is a peculiar social construct which needs for coming to maturity a whole set of social conditions. Observing the evolution of capitalism brings to the fore historical phases in which some markets see their role increase or decline. For instance, during the Thirty Glorious years in France, the stock market had a marginal impact. Besides, the competitive forces always coexist with other forms of regulation of same importance, for instance money, hierarchical links, trust, conventions and norms. The prevalence given to the market leads to under-estimate the regulative function of other entities, which act jointly with the market, for instance the firm, the central bank or the law.

Towards an evolutionary paradigm

To the triptyque formed by optimizing rationality, equilibrium and market efficiency, the promoted approach opposes procedural rationality, dynamic processes and plurality of institutions. Hence, it is situated at the junction of several modeling streams which developed with some success these last three categories, namely the cognitivist, evolutionnist and institutionalist approaches. If the term “evolutionary” is chosen to qualify that synthesis, it is not only due to the necessity of a simple denotation, but also to the transversal role played by that notion in the structuration of the set of approaches. As was already stressed, the underlying epistemology of our approach, at odds with Friedmanian individualism, insists on an evolutionary modelling of the processes at work, simultaneously cognitive when individual decision-making is concerned, evolutionist when dynamic interactions are concerned and self-organizational when institutions are introduced. The federating role played by the evolution processes in our analysis explains why the labelling “evolutionist paradigm” is favored¹.

It is obvious that the conceptual achievement of evolutionary economics would have been impossible without the constitution of a set of technical tools allowing for a renewed approach of the economic evolution. For instance, with the mathematical study of non linear dynamic systems, one gets a lot of new concepts and results concerned with stability, bifurcations and various forms of attractors. Likely, with the formal work by physicists on systems of heterogeneous and tightly related entities, one gets richer insights about “self-organization” (Lesourne, 1991) or “emergent phenomena”. Finally, with the development of epistemic logics by philosophers and cognitive scientists, one gets a more accurate view of individual (and collective) beliefs and modes of reasoning.

Despite some external influences, the central theses of the book belong really to economic science, or more generally to social sciences. They induce a conception of economics notably different of the conception which prevails in traditional textbooks. This can be illustrated by three examples

¹ Note however that what the research program called here ‘evolutionary economics’ is not far from what B. Walliser calls elsewhere ‘cognitive economics’ (2000). Conversely, it differs from economic models called ‘evolutionary’ in a strict sense and focalized on a dynamic dimension without replacing it in a cognitive and institutional framework. It differs even more with an approach exclusively grounded on a biological analogy as evolution is concerned.

which depart more and more from the traditional view of an efficient market equilibrium. First, the notion of “path dependency” will be frequently used in order to stress that “history matters”. The state towards which the economic system may converge depends on the internal events that happened along its path and on the external shocks that perturbed its trajectory. Such a notion is not incompatible with an equilibrium analysis, but it restricts its relevance since that analysis has to be completed. Second, the evolutionary dynamics does not necessarily converge towards some optimal state. Contrary to the common vulgate shared by some evolutionnists, evolution does not systematically mimic a global optimization of the system. Not only is the asymptotic state not collectively optimal in some technical sense, but the notion of optimality becomes even problematic. Third, it appears that in many situations, the attractors are not necessarily punctual. The system may stay perpetually in a moving state, and the notion of equilibrium loses its relevance. This is the case when observing limit cycles or chaotic dynamics.

These contributions of evolutionary economics stay compatible with a somewhat mecanist approach of economic evolution. The introduction both of beliefs and institutions is a further step which improves even more the proposed analysis. Especially, it is shown that the beliefs have a proper efficiency since the coordination of individuals depends on how each agent interprets his strategic environment. Likely, the institutional devices influence in several ways the interaction process by coordinating the agents’ beliefs as well as actions. The complex interwaving of these factors leads to an image of economic dynamics which is conceptually better fitted to the economic phenomena and is pragmatically better adapted to the economic problems.

Presentation of the book

The book differs profoundly from preceding books dealing with evolutionary economics too (Witt, 1992; Hodgson, 1996; Schweitzer-Silberberg, 1998; Dopfer, 2001; Foster-Metcalf, 2001; Gandolfi et alli, 2002; Backhaus, 2003; Witt, 2003). These books are litterary presentations of evolutionary economics or proceedings of conferences on the topic. The structure of the present book in three parts manifests a progression from the presentation of basic concepts to the analysis of complex situations. In fact, it follows more or less the structure in traditional textbooks.

The first part deals with the basic concepts concerning individual behavior and mutual interactions. Beginning with individual decision in chapter 1, it presents essentially the notion of procedural rationality. It shows how the bounded rationality of some actor may be compensated by learning along time. Chapter 2 studies a very simple form of interactions on a market where, in conformity with traditional analysis, the behavior of the agents is purely reactive. It is the question of the emergence of a unique price which is essentially studied. The situations of strategic interactions are introduced in chapter 3, devoted to games. It is mainly concerned with evolutionary game theory, and more precisely with learning processes.

The second part introduces more complex market configurations than the first one. In chapter 4, markets with irreversibilities are considered, leading to multiple prices. Chapter 5 considers mimetic interactions and the collective dynamics they involve, leading for instance to the emergence of financial bubbles. Chapter 6 studies various forms of dynamic competition between firms and the results in terms of prices and market organization.

The third part is devoted to institutional devices working as complements to the market. In chapter 7, it is the firm which is analyzed as concerns its internal organization. A more general taxonomy of institutions linked to their emergence conditions is developed in chapter 8. Chapter 9 concludes with the economic role played by the state.

For pedagogical reasons, all chapters are designed along a same structure. A first section recalls the origins and features of the point of view developed by the standard economic approach. It is called "Background and problems". A second section presents the general notions and principles suggested by the evolutionary approach in order to deal with its object. It is called "Canonical principles". A third section studies the consequences of specific assumptions gathered in contrasted models relative to a peculiar situation. It is called "Some models". A last section tries to generalize the partial results obtained in specific contexts and to open unexplored roads. It is called "Theses and conjectures".

This book is mainly oriented toward students having already acquired the basic knowledge of economic theory. Based on an extensive presentation of its concepts and principles, it aims at making them familiar too with the tools and models of evolutionist economics. This is why sufficiently simple and transparent models have been selected in order to analyze easily all their mechanisms. The book, which is not looking for exhaustivity or premature synthesis, stays upstream from more complete and specialized ones.

References

- Backhaus, J. (2003): *Evolutionary economic thought*, Edward Elgar.
- Dopfer, K. (2001): *Evolutionary economics: program and scope*, Kluwer.
- Foster, J. and Metcalfe, S. (eds.) (2001): *Frontiers of evolutionary economics*, Edward Elgar.
- Friedman, M. (1953): The methodology of positive economics, in *Essays in Positive Economics*, Chicago University Press.
- Favereau, O. (1989): Marchés internes, marchés externes, *Revue Economique*, 40(2), 274-328.
- Fisher, F. (1991): La formation des grandeurs économiques: déséquilibre et instabilité, in J. Cartelier (ed.), *La Formation des Grandeurs Economiques*, PUF, Paris, 19-55.
- Gandolfi, A. et al. (2002): *Economics as an evolutionary science*, Transaction publishers.
- Hodgson, G. (1996): *Economics and evolution*, University of Michigan Press.
- Hodgson, G. (ed.) (2002): *A modern reader in institutional and evolutionary economics*, Edward Elgar.
- Kirman, A. and Salmon, M. (ed.) (1995): *Learning and rationality in economics*, Blackwell.
- Kuhn, (1970): *The Structure of Scientific Revolutions*, 2nd ed., University of Chicago Press.
- Lesourne, J. (1991): *Economie de l'ordre et du désordre*, Economica, Paris.
- Nelson, R. and Winter, S. (1982): *An Evolutionary Theory of Economic Change*, Harvard University Press.
- Schweitzer, F. and Silverberg, G. (eds.) (1998): *Evolution and self-organization in economics*, Duncker and Humblot.
- Simon, H. (1982): *Models of bounded rationality*, MIT Press.
- Tisdell, M. (1997): *Bounded rationality and economic evolution*, Edward Elgar.
- Walliser, B. (2000): *L'Economie Cognitive*, Odile Jacob, Paris.
- Witt, U. (ed.) (1992): *What evolutionary economics is about*, Edward Elgar.
- XXX (2002): Symposium on evolutionary economics, *Journal of Economic Perspectives*, 16(2).

1 Individual decision

In economics, it is traditionally assumed that an agent's behavior can be broken down into a series of parallel or sequential actions, chosen as the result of a process of mental deliberation. The agent thus appears as an autonomous decision-maker who chooses, either consciously or implicitly, in a situation that can be isolated from its context, between the various alternatives presented to him. Furthermore, this decision-making process is assumed to be rational, by virtue of two remarkable properties. Firstly, the agent is "consequentialist" in the sense that he chooses his action solely according to its (foreseeable) consequences; secondly, he is "utilitarian" in the sense that he evaluates the effects of his action by weighing up its costs and advantages. Consequently, such an agent is restricted to a minimal psychological framework, insofar as his choices are governed exclusively by three personal choice determiners: his opportunities (delimiting the space of his possible actions), his representations (enabling him to predict the consequences of his action) and his preferences (inducing a judgment on these consequences). These three determiners are further combined in a choice rule which characterizes more precisely the rationality of the decision-maker.

In the classical approach, the decision-maker is animated by very strong rationality relying on three assumptions. First, given his prior beliefs, he is capable of perfectly anticipating the effects of his actions. Second, he judges his actions on the basis of one unique synthetic criterion, utility, which sums up their costs and advantages. Third, he adopts optimising behaviour, in the sense that he seeks the action that maximises his utility (defined directly on the actions beyond their effects) under certain constraints (those limiting the set of his possible actions). These assumptions have been progressively weakened, but only to a limited degree. When dropping the first assumption, the decision-maker only possesses imperfect information about his environment. The more complex modification of the second assumption gives us a decision maker using multiple, but nevertheless commensurable, criteria of choice. The third assumption is generally kept and assumes that the decision-maker makes his choice without having any real difficulty in calculating what his optimum action is.

In the evolutionary approach, the rationality of the decision-maker is much more limited and is situated within a dynamic perspective. His information is reduced and derives not so much from his prior knowledge as from his past observations, which accumulate and enable him to revise his beliefs. His utility is not necessarily pre-defined, but built as a function of his past experience in analogous situations. Above all, the decision maker's deliberative process is constrained by his limited ability to calculate, and this internal constraint must be added to the external constraints. However, this cognitive limitation can be compensated for by the work of time, at least if the decision maker carries out a succession of repetitive choices. In this case he finds himself involved in a learning process which can, over the long term and in some circumstances, converge towards an optimal action, but the medium term trajectory of this learning process is in itself of interest to the modeler.

This chapter explores precisely the passage from the first approach to the second. The first section reviews the principles of classical decision theory, namely the classically proposed rules of choice, both static (§1.1) and dynamic (§1.2), illustrated by a prototypical example (§1.3), the justifications (axiomatic, operational, evolutionary) that have been presented for it (§1.4) and the criticisms (empirical, theoretical, logical) that have been levelled at it (§1.5). The second section defines the principles of evolutionary behavior, setting out different concepts of rationality (§ 2.1) and then successively examining the processes of prediction (§2.2) and selection (§2.3) carried out by the decision-maker, giving rise to the problem of the value of information (§2.4) and to the exploration-exploitation dilemma (§2.5). The third section describes some recently-developed evolutionary models, firstly models of choice with limited rationality (§3.1), then learning processes applied to repeated decision situations, both static (§3.2) and dynamic (§3.3), possibly simplified (§3.4), these processes being illustrated by the earlier prototypical example (§3.5).

1.1 Background and problems

1.1.1 The choice rules in static situations

In classical decision theory, in its static form, the decision-maker finds himself faced with an environment called "nature". The decision-maker takes actions and nature assumes states. The instantaneous conjunction of an action and a state results in consequences that are certain. These are of-

ten expressed in a monetary form. The “normal form” of the decision problem is expressed by a matrix which indicates the consequences resulting from each action-state pair. Here, the choice rules of the decision-maker rely on three ingredients which formalize his choice determiners (opportunities, representations, preferences):

- a predefined set of strategies, whether this involves actions (defined by their sure consequences) or lotteries (defined by their consequences conditional on the states);
- a belief about the occurrence of states, expressed in particular in the form of objective (proportions or frequencies) or subjective (degrees of belief) probabilities;
- a utility function defined on the certain consequences of the actions, which can be ordinal (only the orders are significant) or cardinal (the numerical values are significant).

Nature is assumed to be passive in the sense that it assumes its states mechanically (they are not the result of a decision process) and according to an exogenous rule (the states are insensible to the actions of the decision-maker). Depending on the decision-maker’s uncertainty about this rule and about the state of nature actually produced, situations of uncertainty can be divided into four main categories:

- certainty: the decision-maker knows the state of nature produced (whatever the rule producing it);
- probabilistic uncertainty: the decision-maker knows the probability distribution according to which the state of nature is produced;
- set-theoretic uncertainty: the decision-maker only knows the list of states of nature, without knowing which of these states may be produced;
- radical uncertainty: the decision-maker does not know the list of states of nature.

Of course, there are intermediate situations, for example a second order uncertainty when the decision maker knows that the rule governing the production of states is probabilistic, but only has partial information about this probability distribution.

The whole subsequent history of decision theory can be summed up as a series of attempts to provide the choice rules of the decision-maker in one

or another of the main situations of uncertainty. The earliest and simplest of these rules are:

- the rule of maximisation of utility under certainty (Debreu 1954)
- the rule of maximisation of (objective) expected utility under probabilistic uncertainty (von Neumann-Morgenstern 1944)
- the rule of maximisation of (subjective) expected utility under set-theoretic uncertainty (Savage 1954).

More sophisticated rules have been proposed more recently, generalising the above rules:

- the rule of maximisation of rank-dependent expected utility under probabilistic uncertainty (introducing a function of deformation of probability distribution);
- the rule of maximisation of credibilist expected utility under set-theoretic uncertainty (introducing “non-additive probabilities”).

1.1.2 The choice rules in dynamic situations

In the dynamic form of classical decision theory, the decision-maker and nature intervene sequentially. The conjoined consequences of a succession of actions and states are only defined at the end of the sequence. The “extensive form” of the decision problem is expressed by a “decision tree”. The decision-maker and nature play alternately at successive nodes, and the vertices issued from each non-terminal node represent the options available to the agent whose has the move. Each terminal node expresses the consequences (usually in monetary terms) for the decision-maker of the trajectory leading to this node. Nature is always independent of the decision-maker. Its successive moves may be independent, but may be correlated too. Especially, Nature may first define a state and further supply messages which specify this state. Moreover, the law governing the production of states is assumed to be stationary. Finally, in an extensive form game, a “strategy” of the decision-maker is the prior choice of an action at each node where he may play.

Within this framework, the choice rules defined in statics are extended and a new principle appears: the “backward induction principle”. This postulates that the decision-maker determines his actions by starting from the horizon of the decision tree and progressively working backwards in time along the decision tree. For example, for a (sequential) decision problem

under objective uncertainty, he progressively moves back along the nodes of the tree (from the terminal nodes through to the initial node) by considering, if the node corresponds to a move by nature, the expected utility on all the possible resulting states and, if the node corresponds to one of his moves, the maximum utility on all his possible actions. Expected utility is measured with the probabilities attributed to each state, which are conditional on the information already received in the past about the trajectory considered in the tree.

A slightly more general representation of a decision problem is provided by the “stochastic decision theory” (although it can also be expressed as a decision tree). If the decision-maker and nature always play sequentially, the global system can assume a certain number of finite “configurations” c_h . The system may go through the same configuration several times, thus introducing loops into the history of the process. Because of the influence of Nature, the transition from one configuration h to another configuration k , conditional on an action i , is expressed by a probability of transition p_{hk}^i . Furthermore, the decision-maker chooses his action according to the configuration of the system, and this defines a strategy $s_i = \pi(c_h)$. Lastly, a utility of transition u_{hk}^i is associated with each transition from one configuration to another through a certain action; all the utilities gathered by the decision maker along his trajectory are finally aggregated into a synthetic utility U introducing an appropriate discount factor δ .

For the decision-maker, knowing both the probabilities and utilities of transition, the optimal strategy is that which maximises the discounted sum of expected utilities over an infinite horizon. One can demonstrate that this optimal strategy is deterministic (the chosen action in each configuration is non probabilistic), Markovian (the chosen action is independent of past states) and stationary (the chosen action is independent of time). The optimal strategy $\pi^*(c_h)$ is again obtained through a backward induction procedure. The last has to consider the maximal utility U_h^i that the decision-maker can obtain when starting from the configuration h and taking the action i and the maximal utility U_h that he can obtain when starting from the configuration h . These utilities accord with the Bellman equations, defining a fixed point:

$$\begin{aligned}
 U_h^i &= \sum_k p_{hk}^i (u_{hk}^i + \delta U_k) \\
 U_h &= \max_i U_h^i \\
 \pi^*(c_h) &= \arg \max_i U_h^i
 \end{aligned}
 \tag{1.1}$$

1.1.3 An example of dynamic choice

As an illustration, take the example of Savage’s omelette (Savage, 1954), in which a cook wishes to make an omelette constituted of n eggs. He has at his disposal a batch of eggs, a bowl B and a saucer S. By hypothesis, the egg has a cost a and is good (with a probability $1 - p$) or bad (with a probability p). For making his omelette, the cook can break each egg directly in the bowl or provisionally in the saucer. Breaking an egg provisionally in the saucer has the advantage of not spoiling the whole content already in the bowl, but at some tranfert cost b . When the bowl contains n eggs, the omelette is cooked and sold at price c and the cycle starts again.

The possible configurations of the system are the $(n + 1)$ situations corresponding to the number of eggs in the bowl (from 0 to n). A strategy of the cook consists in deciding, in each configuration, whether to break the next egg in the bowl or the saucer. The problem to be solved by the cook is to determine the strategy to be followed in order to maximize his profit.

In the case of an omelette with only 2 eggs, we can present the process (fig 1.1.) in the following manner (the nodes of the decision maker are represented by squares in which the configuration attained is noted and the nodes of nature are represented by circles):




Fig. 1.1. Omelette decision graph

The optimal strategy is obtained by writing the Bellman conditions on the three possible configurations and by grouping together the common consequences of the alternative actions:

$$\begin{aligned}
U_0 &= a + \delta p U_0 + \delta(1-p)U_1 + \max(0, b) \\
U_1 &= a + \delta(1-p)U_2 + \max(\delta p U_0, \delta p U_1 + b) \\
U_2 &= c + \delta U_0
\end{aligned} \tag{1.2}$$

When the bowl is empty, as b is negative, the egg must always be broken into the bowl, so that the first equation becomes:

$$(1 - \delta p)U_0 = a + \delta(1-p)U_1 \tag{1.3}$$

When the bowl contains one egg, U_1 is obtained by the previous equation and the following equation:

$$U_1 = a + \delta(1-p)(c + \delta U_0) + \max(\delta p U_0, \delta p U_1 + b) \tag{1.4}$$

There are two optimal strategies depending on the values of the parameters; more precisely, the probability of the second egg being bad admits a critical threshold p_c such that:

- if $p < p_c$, always break the egg into the bowl;
- if $p > p_c$, always break the egg into the saucer.

1.1.4 The justifications of the choice rules

Static choice rules receive an axiomatic justification, in the sense that they are the result of a set of axioms defined on the global preferences of the decision-maker concerning strategies (actions or lotteries). On the one hand, these axioms make it possible to give a particular form to the choice rule. Thus, all rules have the form of maximisation because they require the decision-maker to define a complete order of preferences on strategies. On the other hand, their analytic form depends on the additional axioms that are imposed. Thus, in the criterion of maximisation of expected utility, the probabilities are separated from the utilities on the certain consequences by means of either the independence axiom (under probabilistic uncertainty) or the sure thing axiom (under set-theoretic uncertainty). Furthermore, if the decision-maker satisfies the axioms, his beliefs and preferences can, under some conditions, be revealed by his chosen actions. For instance, according to the rule of maximisation of expected utility, the subjective probabilities and the certain utilities of the decision-maker can be reconstructed from the elementary choices he makes between well selected lotteries. Dynamic choice rules have also been the subject of axiomatic

justifications. The backward induction principle, at least in combination with a static criterion of choice, can thus be justified axiomatically. Essentially, this makes it possible to ensure the “dynamic coherence” of the decision-maker, i. e. that a decision taken today for tomorrow will not be called into question tomorrow.

Choice rules, both static and dynamic, have also been given an operational justification, linked to the performances they make it possible to attain (in relation to their cost). In particular, a decision maker could not accept a sequence of choices of which the outcome would necessarily represent a loss for him. The money pump argument is used to justify the fundamental axiom of the transitivity of preferences. An agent with cyclical preferences can be proposed a series of certain choices that can only lead to his ruin. The Dutch book argument demonstrates that the agent’s beliefs should be of a probabilistic form. An agent whose beliefs do not respect the Kolmogorov axioms governing probabilities can be proposed a series of bets that result in an inevitable loss for him. However, the effective range of these arguments is limited, insofar as the agent is hardly likely to find himself actually faced with such artificially constructed sequences of choice.

Finally, some choice rules have received an evolutionary justification, namely that a learning or evolution process can push the decision-maker to follow the rule in question. In general, however, it is not the effective rule of choice of the decision-maker which converges towards a given rule, but his strategy which converges towards the strategy advocated by the rule. This is to say that everything happens “as if” the decision-maker was using such or such a rule asymptotically. Thus Friedman, following on from Alchian, upheld the thesis that in a context of competition between agents, only those who adopt an optimising behavior will survive. This thesis has given rise to formalised studies exploring a situation of repeated interactions, both in the context of learning by agents (see section 3) and in the context of (biological) selection between agents (see chapter 3). This will be taken up more specifically in relation to competition between firms (see chapter 6). All these works conclude that the exclusive survival of the optimising agents only occurs under very specific conditions and in very particular contexts. In fact, in a competitive situation, the relative performance of agents in the presence of others is more important than their performance in absolute terms.

1.1.5 The criticisms of the choice rules

Choice rules have all been the object of empirical criticism, often taking the form of “empirical paradoxes”. In the earliest days of decision theory,

the “Allais paradox”, using the example of a decision maker choosing between cleverly constructed lotteries, challenged the rule of maximisation of (objective) expected utility. Likewise, the “Ellsberg paradox” brought into question the probabilistic reasoning in the rule of maximisation of (objective) expected utility. Following the work of Kahnemann and Tversky (Tversky-Kahnemann 1986), a large movement has developed in favor of an “experimental decision theory”, conducted in the laboratory. This approach counters an old argument which says that the rational decision model is tautological insofar as whatever the actions taken by the decision-maker, there always exist beliefs and preferences that can account for them. In fact, as soon as the rational decision model is sufficiently well specified (precise context, choice rule with a specific analytic form), it becomes refutable. This approach has also led to recognize that, if the laboratory conditions are not far removed from field conditions, the results obtained in the first are representative of the second. Of course, decision-makers acting as laboratory subjects are removed from the social pressures they might otherwise feel, and the stakes are not the same as in a real situation, but it is nevertheless considered that their behavior can reasonably be generalised.

In general, when these experiments result in a choice rule being cast into doubt, the axiom(s) that must be called into question cannot be sorted out solely by logical reasoning (the “Duhem-Quine problem” in epistemology). Nevertheless, in the case of the experiments by Allais, it can be demonstrated that the axiom that has been violated is in fact the independence axiom. When a choice rule is actually called into question, one generally considers weakening the axioms assumed to be at fault. This has the effect of leading to more generalised choice rules, including the former rules as particular cases. Following the experiments of Allais, the consideration of the weakened axiom of “comonotonic independence” has resulted in a generalized choice rule, i. e. maximisation of rank-dependent expected utility. An important issue however, when rules are weakened, is that they also become less refutable in the Popperian sense of the word, as the number of circumstances capable of challenging them is reduced.

Choice rules have also been subjected to much more general theoretical criticism, focusing on the need to take the decision maker’s cognition into account. The decision-maker is, in fact, required to choose a strategy at the conclusion of a process of explicit or implicit deliberation. And yet, except for very simple situations (the “small worlds” of Savage), he has a limited capacity to gather and process information. Consequently, one must add internal cognitive constraints to the external, material constraints imposed on the agent by his environment. These internal constraints are, in particu-

lar, expressed by costs, either the cost of the search for information or the cost of processing the information obtained. As the choice rules become more complex, so these constraints become stronger, for refinements in the usual rules paradoxically call for additional cognitive capacities. This criticism is theoretical if we remain within the traditional context, in which only the agents' actions can be observed. However, mental states (beliefs, preferences) are, to an ever greater extent, considered to be observable through introspection. One could even go further and argue that the process of deliberation is observable, in which case theoretical criticism concerning the limited capacities of the decision maker becomes empirical criticism.

Lastly, choice rules can be the subject of logical criticism, demonstrating that procedural rationality gives rise to an infinite regression on choices (Mongin-Walliser 1989). Taking the above theoretical criticism as a starting point, one can consider that when a decision-maker wishes to optimise, he comes up against an optimisation cost which is often quite high and which may lead him to choose other, less costly choice rules. He is then obliged to carry out a process of meta-optimisation, in which he chooses a rule of choice (from a certain set) by weighing up the loss of utility it represents (compared with optimisation) against the reduced costs of calculation it involves. However, this meta-optimisation itself has a cost, so that the decision-maker finds himself faced with the same problem, but on a higher level. The only way to escape from this infinite regression is by placing oneself at an arbitrary level in the hierarchy of choices, optimising at this level and then descending the lower levels one by one. This provides the price at which it becomes rational not to optimise at these lower levels.

1.2 Canonical principles

1.2.1 Forms of rationality

Depending on the epistemological interpretation given to the choice rule of the decision-maker, Simon (1976) proposed the use of two concepts of rationality:

- “substantive rationality” means, from an instrumentalist perspective, that the choice rule is only evaluated in relation to the validity of its predictions in terms of the actions chosen by the decision maker;
- “procedural rationality” means, from a realist perspective, that the choice rule must be judged by the measure of empirical validity of the decision process actually used by the decision maker.

Friedman (1953) argued for substantive rationality as a means of reinforcing the optimising model. He assumed that “everything happens as if” the decision-maker optimises without having to pass judgment on the process of deliberation by which he arrived at his choice. He cites the example of the billiards player who plays “as if” he was optimising the ricochets of the ball against the sides of the table, without this entailing any conscious or indeed unconscious calculation by the player. The optimising model then functions as a tool, the sole objective of which is to predict the decision-maker’s choice. Procedural rationality was defended by Simon (1976) in order to criticise this same optimising model. Simon argued that it is necessary to explain the concrete deliberation process of the decision maker in terms of the prospection and computation he performs in order to define his choice. He gives the example of the chess player who uses various heuristic search and selection rules described in the science of Artificial Intelligence and which can, in certain cases, take the form of algorithms. However, the optimising model is itself capable of functioning as a realist model, bringing into use a particular heuristics, taking the form of a “gradient algorithm”, for example.

Depending on the cognitive requirements imposed on the decision-maker’s choice rule, we can once more define two alternative forms of rationality:

- strong rationality assumes that the decision-maker is endowed with infinite calculating abilities, enabling him successfully to conclude any process of deliberation he may have to perform;
- bounded rationality assumes that the decision-maker has a limited capacity for gathering and processing information, preventing him from carrying out operations of prospecting and calculation that are too complicated.

Bounded rationality of the decision-maker has strong links with procedural rationality (Laville 2000). On the one hand, bounded rationality only has significance in a context of procedural rationality, for the limited abilities of the decision-maker, often expressed in the form of costs, are only relevant in this one context. On the other hand, bounded rationality of the decision-maker leads us to examine procedural rationality more closely as a means of bypassing the cognitive limits imposed on him. Strong rationality is postulated by the classical choice rules and has the great advantage of appearing in a univocal form. Bounded rationality is taken into account by more recent choice rules, bearing in mind that these rules are developing in multiple directions (Conlisk 1996). A spectrum of models has been devel-

oped in an attempt to mark out the field of possibilities, but at the present time no canonical model of bounded rationality exists that can be substituted for the optimising model. Moreover, models of a certain type can often be rewritten as another type (including in the form of optimisation).

Lastly, choice rules bring into play two types of rationality which intervene complementarily (Walliser 1989):

- cognitive rationality expresses the degree to which the beliefs the agent constructs about himself and his environment are appropriate to the information he possesses;
- instrumental rationality expresses the degree to which the objectives that the agent pursues are appropriate to the means at his disposal, taking his beliefs into consideration.
- both types of rationality can rely on a strong form. Strong cognitive rationality means that the decision-maker is able to form perfect (or rational) expectations; he reasons like a perfect statistician who, on the basis of perfect and complete information, minimises the error of prediction on the variable he is expecting. Strong instrumental rationality signifies that the agent is capable of determining an optimising action on the basis of his opportunities and preferences, given that his beliefs (whatever they may be) are fixed. The two types of rationality can also rely on weaker forms. In fact, since bounded rationality reflects limited cognitive capacities, it is quite naturally of a cognitive type. But the difficulties encountered in expressing it can lead to its formulation as bounded instrumental rationality (see § 3.1). More generally, one may be tempted to reduce cognitive rationality to instrumental rationality by considering that the agent is making the best use of the information available to him. It would be more correct to say that instrumental rationality can be reduced to cognitive rationality, for it influences not only the way the agent considers his environment and his choice determiners, but also his capacity to combine these elements in order to choose his action.

1.2.2 Procedures of prediction

The decision-maker possesses certain structural information, arising from his various acquired knowledge and past experiences, about the decision problem with which he is faced. His first problem involves categorising the situation of the decision. This categorisation operation is performed on

the basis of primitive concepts, namely the actions carried out, the possible states of nature, the resulting consequences and the utilities that are obtained. Firstly, it consists in defining a general framework of choice situations, in the form of either a typology of possible situations or a list of situation prototypes. For example, the decision-maker can distinguish between actions of a material nature and those of an informative nature, between random events of nature with a technical, behavioral or social character or again between material, financial and symbolic consequences. Secondly, this categorisation consists in defining the concrete situation in which he finds himself, either by location within a possible type or by comparison with a prototype. For example, the decision-maker can specify the environmental configurations encountered in the past which he judges to be similar to the one under examination.

The decision-maker can then receive factual information about the actions and states, the consequences and utilities relating to past decisions (his own or those of other agents) in a similar situation. The second problem he faces is that of the internal structuring of the decision situation. This structuring consists in relating the elements of the situation to each other in order to bring out its regularities. On the one hand, this involves defining his beliefs, namely the law governing production of states and the law connecting the consequences with the actions and states. For example, the decision-maker will construct a mental model expressing the causalities he believes there to be between various exogenous factors and the effects of his action. On the other hand, it involves defining his opportunities and preferences, in other words all the strategies available and the relation connecting the consequences with the utility. For example, the decision-maker may, through a process of pre-selection, only consider a subset of different strategies that are assumed a priori to be the most effective. Likewise, he may, through a process of simplification, only consider one family of choice criteria that he finds relevant, either combined in a heuristic evaluation function or remaining separate.

The decision-maker's beliefs, imperfect and incomplete, have been the subject of formalisations within the framework of "epistemic logic". They are expressed either in a syntactic context, where the agent works on the propositions and is endowed with a belief operator indicating which propositions he knows, or in a semantic context, where the agent considers possible worlds and is endowed with a relation of accessibility indicating the worlds between which he is capable of distinguishing. In the semantic context, these beliefs may take an all-or-nothing form, in which the agent contents himself with either knowing or not knowing, or a probabilist form

in which he attributes probabilities to his assertions; we can also envisage mixed hierarchical forms of belief (non-additive probabilities). This epistemic representation enables us to specify which axioms are satisfied by agents' beliefs: logical omniscience (an agent knows all the consequences of what he knows), veridicity (what the agent knows is true), positive introspection (the agent knows what he knows) and negative introspection (the agent knows what he does not know). The failure to satisfy one or other of these axioms is often at the origin of a decision-maker's bounded rationality (this is particularly true for logical omniscience, which endows him with unlimited computational abilities). In the semantic framework, when all the axioms are satisfied, the decision-maker possesses an "information partition" about the possible worlds, indicating the private information available to him; often he also has at his disposal a probability distribution on the worlds, expressing public information about their material aspects.

The decision-maker predicts the effects of his actions on the basis of his beliefs, firstly in terms of objective consequences, and subsequently in terms of criteria of evaluation (or utility). However, he may use very crude models of the functioning of his environment, bypassing certain relations, in order to carry out his predictions. It is in this way that he constructs indexes which group together and summarise his past experience (and possibly that of other decision makers) and make prediction possible. On the one hand, an index may cover states of nature in the form of an index of the past frequency of different states (or a more complex index in which each state is weighted in proportion to its recentness). Assuming that the law governing production of the state is stationary, the decision maker can infer an expectation about the future state from it by translating past frequency into future probability; he can also use very simple, extrapolatory or even adaptive methods of prediction. On the other hand, the index may cover the utilities obtained in the past, in the form of an index of the aggregated utility associated with each action (this index may present average utility, discounted utility or accumulated utility). Assuming here again that the observed performance persists in the future, the decision-maker can infer from it an expectation of the future performance of the action (without passing via the states).

The reappraisal of decision-makers' beliefs has also been the subject of formalisations in epistemic logic. This still involves moving from an initial belief and a message to a final belief. We can distinguish between two contexts of reappraisal: revising when the message supplies additional information about a world considered to be unchanging, and updating when

the message provides an indication of the way in which an evolutive world has changed. In both of these contexts it is possible to obtain change rules expressed in semantics (selection of possible worlds) from change axioms expressed in syntax (properties of belief operators). These rules, particularly interesting when the message contradicts the initial belief, are defined clearly both in the all-or-nothing context and in the probabilist context. In particular, Bayes rule, traditionally used by economists to reappraise probabilities, only proves to be justified in a context of revising and with very demanding axioms (Walliser-Zwirn, 2002).

1.2.3 Procedures of selection

The decision-maker must choose an action according to his opportunities and preferences, and according to his expectations. If a prediction of the future state of Nature is available to him and if his preferences can be reduced to a synthetic utility function, he can perform a maximisation of his action. But he may renounce this optimisation by simply favoring actions with a high utility and ignoring the others, without seeking exclusively the action with the maximal utility. If a prediction of the future state is available to him and if his preferences are expressed in terms of multiple criteria, he must implement one “multicriteria rule” from among a set of such rules. He may also consider an “aspiration threshold” on these partial criteria, and choose an action when its effects exceed this aspiration threshold. If he does not have a prediction of the future state at his disposal, but he does have an index of the utility of actions, he contents himself with strengthening the actions that have performed well in the past and inhibiting those that have performed poorly.

As with his beliefs, the decision-maker can also adjust his opportunities and preferences over the passage of time. On the one hand, he may modify his set of choices by incorporating new strategies. In particular, he may carry out actions related to those he has already tested with success. On the other hand, he may adapt his preferences according to the past utilities he has actually obtained with his actions and which may differ from his expectations. In particular, he may raise or lower his aspiration thresholds according to the ease with which he has attained them in the past. He will, of course, bring his index up to date using his most recent observations of the state produced and the utility obtained. Lastly, he may modify his rule of choice itself, if he feels that he is “locked in” to an action that is performing poorly compared with an external reference (bearing in mind that

he does not know how far he is from a possible optimum) or if he observes that the environment is evolving significantly.

The term “adaptive rationality” is sometimes used to explain the way in which the decision-maker modifies his choice rules with the help of meta-rules, during a learning process which operates on several functional levels (a higher rule acts on lower rules) and several temporal levels (a higher rule changes more slowly than a lower rule). In practice, it is not very easy to distinguish between rules and meta-rules, insofar as the former already incorporate a process of adaptation to the environment, even if its structural characteristics, unlike its parameters, remain fixed. Of course, meta-rules, even more than rules, originate in a rationality that is both procedural and limited, and they are themselves chosen by imitation of the rules adopted by others or by reinforcement in relation to their effectiveness. They also raise the problem of innovation of rules, insofar as the modeler always presupposes the availability of a fixed set of rules, whereas the agent does not consider them all at the same time and is obliged to favour certain ones, even if this means renewing them.

Finally, the decision-maker can implement two types of action, possibly mixed. If the objective of operational actions is to transform a system considered to be unsatisfactory, the aim of informational actions is to gather information in order to feed operational actions. Information gathering can be carried out by means of two extreme paths. The decision-maker can either obtain information in an exogenous and costly manner from specialised entities (autonomous informational action) before acting operationally, or he can obtain it endogenously and free of cost as a by-product of the normal course of the decision process (spontaneously information-bearing operational action). An intermediate situation exists, in which the decision maker arbitrates between two different behaviors in relation to his current action. Exploration behavior consists in defining an action that enables the agent to obtain the largest possible amount of information. Exploitation behavior consists in using the already-existing information as efficiently as possible. In a repeated situation, the arbitration between exploration and exploitation consists in favoring exploration at the beginning (by testing new actions) and exploitation at the end (by using the most effective actions).

1.2.4 The value of information

Consider a decision process which can be broken down into two periods. In the first period, the decision-maker acquires factual information about

the state of nature in the form of a message (from a set of possible messages) considered to be true. If the message is all-or-nothing, it marks out a subset of states for each state, whereas if it is probabilist, it is characterised by its conditional probability for each state, the limit being the situation in which the message specifies the state (certain message). In the second period, the decision-maker modifies his beliefs as a function of the message and chooses an action accordingly. The ex post value of the information is simply the difference in the utility obtained by the agent depending on whether the chosen decision is taken before or after reception of the message. The ex ante value of the information is the expected ex post values for all the possible states (and messages), in other words calculated on average before knowing the message actually received. The decision maker consequently chooses to gather the information if its ex ante value is higher than its cost (inasmuch as he can calculate this value).

The ex post value of the (non certain) information can be positive or negative. The decision-maker may receive an improbable message that incites him to make a bad decision. On the contrary, a fundamental result affirms that the ex ante value of the information is always positive if the decision-maker uses the maximisation of expected utility as his choice rule. This means that the decision-maker, having received a (true) item of information, cannot find himself in a worse situation than the one he was in before receiving it. However, this result is invalidated if the decision-maker uses a choice rule other than the maximisation of expected utility. It is also invalidated if the message takes a form other than that described (non-partitional all-or-nothing message, probabilist message with non-additive probabilities).

This framework can be extended to the case of the endogenous acquisition of information between two actions. In the first period, a decision-maker has the choice between a reversible operational action and an irreversible one. In the second period, this action supplies a message about the state of nature that has occurred. In the third period, the decision maker can take advantage of this message to amend the reversible action, whereas the irreversible action is definitive. If the state is favourable, the utility of the irreversible action is higher than that of the reversible action, even when the latter is amended; if the state is unfavorable, the reversible action is preferable exactly because it can be adapted. It is then possible to demonstrate that the reversible action possesses a certain bonus in comparison with the irreversible action, and this bonus is in fact equal to the value of the information supplied by the message.

1.2.5 The exploration-exploitation dilemma

Consider a decision process repeated over an infinite number of periods. In each period, nature randomly draws a state according to a probability law that remains identical for all the periods. In each period, the decision-maker implements an action which simultaneously supplies him with information about the current state of nature and provides him with an operational utility. He uses this information to make the law of production of states more precise and to improve his future decisions. The exploration-exploitation dilemma is then expressed by trade-off between a short term loss in utility from not taking the best action (opportunity cost of the information) and a long term gain in utility due to the additional information obtained (decisional value of the information). To carry out this arbitration, the decision-maker must ensure that he possesses second order information about the form of the law governing the production of states.

The exploration-exploitation compromise has an optimal solution for particular classes of decision process, more especially for k -armed bandits installed in casinos. An arm i is assumed to provide a gain of g_{ik} in the state k with a probability of p_{ik} ; successive draws of the state are assumed independent and performed according to a law that is invariant over time but unknown to the agent, who nevertheless formulates a hypothesis about its type (normal, Bernoulli, etc.). To simplify matters, we consider a two-armed bandit, each arm i giving a gain of 1 with probability p_i and a gain of 0 with probability $1 - p_i$, and we assume that the decision-maker knows a distribution of p_i . The decision maker must work one arm in each period over an infinite length of time, bearing in mind that his choice rule is the intertemporal maximisation of expected gain with a discount coefficient δ .

It has been demonstrated (Gittins, 1989) that the problem is solvable by backward induction and its solution is given by the Gittins rule, a very elegant solution from the point of view of procedural rationality as it transforms a k -dimensional problem into k one-dimensional problems. The Gittins rule consists in attributing to each arm and for each period a "Gittins index", so that in each period the agent chooses the arm with the highest index and updates the index of this arm according to the result he obtains. This rule leads with positive probability to one sole arm being consistently chosen after a certain time, in other words exploration is abandoned in favor of exploitation. However, as the process is strongly dependent on the path chosen, there is a non-zero probability of choosing the bad arm. This probability decreases as the discount coefficient of the decision-maker increases. If the discount coefficient tends to 1, the agent

will take a very long time to explore before switching to exploitation (the cost of exploration having very little impact on intertemporal utility).

The Gittins index is calculable as a function of the distribution type of the random variable, but its expression generally remains very complicated. This leads to it being approximated asymptotically, for probability distributions of states with finite variance, by indexes relating to the normal law (by virtue of the law of large numbers). For the normal law, a value that does itself approximate the index of the arm i in the simplified example is the following:

$$v_i = m_i + a(\delta)s_i / n_i \quad (1.5)$$

where m_i and s_i are the empirical mean and standard deviation, n_i is the number of tests and $a(\delta)$ is a function asymptotically equivalent to $1/\sqrt{2(1-\delta)}$. This expression is the sum of two terms, the first expressing an “exploitation value” and the second an “exploration value”; the latter decreases rapidly as the number n_i of tests rises (faster than the uncertainty on the mean, which is of the order of $s_i/\sqrt{n_i}$), but increases rapidly with the discount coefficient δ as this tends to 1. For example, with a discount coefficient of 0.98 (corresponding to a process with indefinite horizon, of an average 50 periods), an arm which has been used 20 times and has given a positive result 15 times is equivalent to an arm which has been used 6 times and has given a positive result 1 time.

1.3 Some models

1.3.1 Models of decision under bounded rationality

A first model of choice under bounded rationality is the “satisficing model” proposed by Simon (1982) in opposition to the classic “optimizing” model. The decision-maker judges actions by means of partial criteria u_k , to which are attributed the aspiration thresholds σ_k ; he examines the actions in a predefined order and chooses the first one to attain the aspiration thresholds for all the criteria: s_i such that $u_k(s_i) \geq \sigma_k$. As a particular case, one can consider a unique criterion u (as in the case of optimisation), with its aspiration threshold ε ; the decision-maker chooses the action s_i such that $u(s_i) \geq \sigma$. At first sight, the ε -rationality model of Radner fits this definition, by considering that the decision-maker chooses the first action that approaches to within ε of the optimum: $u(s_i) \geq \max_j u(s_j) - \varepsilon$, but here the aspiration threshold actually depends on the maximum attainable

utility, which is generally unknown to the decision-maker. It can be observed that the satisficing model admits the optimizing model as limiting case when the aspiration thresholds are high enough. However, the satisficing model is directly expressed in terms of bounded instrumental rationality and not bounded cognitive rationality. For this latter to appear, we must examine a process of deliberation by the decision-maker that brings into play cognitive constraints such that he is led to seek a satisfactory action. Such a process, which would have the advantage of endogenising the aspiration thresholds of the decision-maker, has not yet been proposed.

A second model of choice under limited rationality is the “probabilist choice model” (Anderson, de Palma, Thisse, 1992). From a finite set of possible actions, the decision maker chooses the action i with probability p_i such that: $p_i = w_i / \sum_j w_j$, where w_i is a propensity to choose the action i linked to an index of utility u_i of the action i . In the linear model, the parameters w_i are proportional to the index of utility: $w_i = u_i$. In the multinomial logit model, the parameters w_i are written in exponential form: $w_i = e^{\mu u_i}$, with the convenient introduction of a parameter μ . Here again, the logit model converges towards the optimising model when the parameter μ tends to infinity; the decision-maker acts then no more in a stochastic manner, but in a determinist manner (except in the case of indifference between two actions). Conversely, the logit model tends to a purely random choice model when μ tends to zero. The parameter μ thus appears to reflect the limited cognitive capacities of the decision maker, but yet again it operates in a model expressing limited instrumental rationality. However, two cognitive justifications of this model, endogenising the parameter μ , have been put forward. In the first, the decision-maker is endowed with a random utility function, but remains optimising to such an extent that he implements each action with the probability that it is the optimising one. When the law of probability of the utility is chosen correctly (doubly exponential), the logit model is obtained. In the second justification (Mattsson-Weibull, 2002), the decision-maker chooses an action by arbitrating between its utility and a control cost in relation to a reference action. When the control cost is chosen correctly (in the form of entropy), the logit model is again obtained.

Other models directly introduce calculation costs or cognitive constraints sustained by the decision-maker (Binmore, 1988; Rubinstein, 1998). A first example is the “model of choice under costly deliberation”, which presents choices at n successive levels. On the first level, the decision maker chooses, for different procedures of choice, the most effective action. On the second level, the decision maker chooses, according to a

meta-procedure of choice, a procedure of choice by comparing the performance of the action chosen and the cost of implementing the procedure. On subsequent levels, he chooses a procedure of choice for selecting a lower-level procedure of choice. On a finite upper level, he chooses the procedures by optimisation, ignoring the cost, in order to avoid an infinite regression (see §1.5). This assumes that the decision-maker possesses an a priori list of procedures of choice and that he is capable of evaluating their cost of implementation and above all their results without actually having implemented them. A second example is the “finite automaton model”. The decision-maker is assimilated with an automaton whose calculation capacities are such that it only has a finite set of internal states at its disposal. It is therefore incapable of performing calculations which exceed a certain degree of complexity.

1.3.2 Models of learning in static situations

The “fictitious play model” assumes that the decision-maker, during a repeated process of decision, is capable of predicting the future states of nature. Moreover, this model essentially expresses exploitation behavior. The decision-maker observes the past frequency of states of nature, deduces from it a distribution of probabilities on future states and chooses, for each period, the action which maximises his expected utility according to this distribution. Exploration behavior can be introduced through voluntary deviation from the above behavior, and this deviation can take two forms. In the “ ε -greedy fictitious play” model, the decision maker can either use the optimum action with the probability $1 - \varepsilon$, or use another action drawn uniformly at random with the probability ε . In the “disturbed fictitious play” model, the decision-maker uses the logit (and no longer optimising) choice rule with, as index of utility, the expected utility calculated for each action. For the standard fictitious play, one can easily demonstrate that the decision process will converge towards the optimal action (in the sense of maximisation of expected utility) simply by means of the law of large numbers (the frequency of appearance of each state tends to its probability). For the variations proposed, on the contrary, this convergence is not sure because the random component generated by exploration does not disappear asymptotically.

The “CPR model” (Laslier-Topol-Walliser, 2000) is a model of reinforcement (Roth-Erev, 1995) which assumes that the decision-maker only observes the past performance of his actions and no longer observes the states of nature. It considers that the decision-maker adopts, as index of

utility, the cumulative utility obtained for each action and that he chooses his future action with a probability proportional to this index. This model presents good properties as regards the exploration-exploitation dilemma. At the beginning of the process, as the indexes are often initialised uniformly, the decision-maker carries out a systematic exploration of all the actions. At the end of the process, if the index of one action becomes predominant in relation to the others, exploitation becomes very strong, although exploration is never abandoned (every action possesses a residual probability of being chosen). What is more, if one increases (decreases) the parameter μ , one moves the exploration-exploitation compromise towards more exploitation (exploration). For $\mu = 0$, there is pure exploration because all the actions are used with the same probability; for $\mu = \infty$, there is pure exploitation because only the action with the maximum index of utility is used. It can be demonstrated that the learning process thus defined converges towards the optimal action (still in the sense of expected utility) because the good actions are played more and more often, due to a retroactive effect of the cumulative utility, whereas exploration tends to zero.

The “threshold choice model” is a dynamic version of the satisficing model, which, like the previous model, no longer requires observation of the states of nature. On the one hand, the aspiration thresholds are adapted over the passage of time according to the results obtained. If the decision-maker has easily found a satisfactory action in the past, he raises his thresholds, whereas if he has had difficulty in finding a satisfactory action in the past, he lowers his thresholds. For example, if the criterion of utility is unique, the decision maker increments his threshold by a constant bonus ρ if the past action gave a better result and vice versa:

- if $u_i(t) \geq \sigma(t)$, then $\sigma(t+1) = \sigma(t) + \rho$
- if $u_i(t) < \sigma(t)$, then $\sigma(t+1) = \sigma(t) - \rho$

On the other hand, a reference action exists which is generally the past action in relation to which the future action is defined. In this way, the past action will be the first to be examined for the future, the others often being treated globally. For example, if the criterion of utility is unique, the decision-maker can repeat the past action if the utility obtained exceeded the aspiration threshold, otherwise he will choose any of the other actions (with a certain law of probability), thus introducing exploration behaviour. Here again, under certain conditions, the process converges towards the optimum action.

1.3.3 Models of learning in dynamic situations

If, while retaining the hypothesis of a repeated decision problem, we move from a static decision problem to a dynamic one, two types of learning models can be envisaged. Firstly, we can continue to apply the above models of learning while adapting them to a dynamic context. One possibility consists in translating the decision problem, expressed in extensive form, into a normal form by the introduction of strategies of the decision maker and then applying the above methods to the strategies. Thus, the CPR model is applicable to the decision-maker's strategies when their performances can be observed. Another possibility is to keep the decision problem in an extensive form, but to apply the above methods to each node of the decision tree. Hence, the CPR model is applicable by considering that, for each successive occurrence in the decision process, the utility obtained by the decision-maker is attributed simultaneously to all the actions appearing in the trajectory followed in the decision tree. Secondly, we can draw directly on the classical rules of choice proposed for dynamic decision situations. This is all the more necessary as these choice rules, based on the backward induction procedure, require high capacities for the processing of information (Sutton-Barto, 1998).

A model of learning proposed early in Artificial Intelligence is the "Q-learning model" (Watkins, 1989), which applies to a stochastic decision process. A reinforcement model, it does not presuppose a priori knowledge of the characteristics of the decision process (probabilities and utilities of transition), although such knowledge helps to accelerate the process. This model leads to revision of "expected local utilities" U_h^i each time the decision maker uses the action i in the configuration h (which he does for the n_h^i th time) to find himself in the configuration k , obtaining the utility u_{uk}^i . The rule of revision is adapted from the Bellman equation and is written:

$$\Delta U_h^i = a(n_h^i)[\delta U_k + u_{uk}^i - U_h^i] \quad (1.6)$$

where $a(n_h^i)$ is a decreasing averaging function (often $a(n) = 1/n$).

The Q-learning process converges insofar as, when the number of tests rises, the averaging function tends to zero and reduces the correction of utilities further and further. Above all, it converges towards the fixed point of the Bellman equations under conditions that impose few constraints. In fact, Watkins (1989) demonstrated that if the underlying decision process is effectively Markovian, if each action is tested an infinite number of times in each configuration and if the averaging function satisfies

$\sum_n a(n) = \infty$ and $\sum_n a^2(n) < \infty$, then the Q-learning process converges towards the optimal solution.

However, the Q-learning process is only defined perfectly when one specifies, in addition to the rule of revision of local utilities according to the information acquired, the choice rule used by the decision-maker according to the local utilities. Watkins proposed using the multinomial logit rule, in other words choosing the action i in the configuration h according to the probability:

$$p_h^i \propto e^{\mu U_h^i} \quad (1.7)$$

The logit rule has the advantage of causing the decision-maker to go through each configuration and each associated action an infinite number of times and thus to obtain the optimal local utilities. On the other hand, it is only at the end of this convergence that one can deduce the optimal strategy of local utilities; in fact, the action resulting from the logit model does not itself converge towards the optimal action (unless the parameter μ itself evolves over the passage of time and tends to infinity). The Q-learning process thus requires infinite exploration before the execution of instantaneous exploitation. Furthermore, it can be demonstrated that, for certain particular exploration tasks, the time it takes for the Q-learning process with the logit rule to converge is an exponential function of the depth of the decision tree. Of course, we can also use, as the choice rule associated with the Q-learning process, the optimal strategy associated with each step in the revision of local utilities, but then there is no longer any guarantee that the process will converge.

1.3.4 Associated models

Local strategies, which associate an action i with each configuration h , can be generalised in the form of “rules” or “classifiers” (Holland, 1987). In this case, a rule associates an action Y_i (possibly pluridimensional) with a set of configurations X_h following the principle: “if condition X_h , then action Y_i ”. The condition of the rule groups together the configurations between which the decision-maker makes no distinction, either because of an error in perception on his part or because the action involved does not require any distinction to be made. It can be considered as an operation of categorisation performed by the decision-maker and therefore expresses the degree of granularity with which he apprehends his environment in relation to the action. A rule is activated by the decision-maker if one of the

configurations of its condition is actually produced. Of course, several rules may be activated in the same configuration, in which case they find themselves in competition. Moreover, certain rules will be used in a chain to obtain a certain result.

To each rule is attributed a utility or “force” U_h^i which evolves over the passage of time according to an algorithm close to Q-learning, the algorithm of the “chain of bearers”. In each configuration h , the admissible rules make “bids” μU_h^i and one of them is chosen with a probability dependent on its bid:

$$p_h^i \propto e^{\mu U_h^i} \quad (1.8)$$

This rule loses its bid, but receives a reward from two sources:

- from the external environment (if the rule acts on the external environment through the action i by providing a utility u_h^i)

$$\Delta U_h^i = u_h^i - \mu U_h^i \quad (1.9)$$

- from the internal environment (if the rule acts on the internal environment by causing transition to the state k , thus triggering a new rule, of which the action is j and from which it receives the bid):

$$\Delta U_h^i = \mu U_k^j - \mu U_h^i \quad (1.10)$$

Recompenses from the external environment are thus retroceded in a cascade over the whole chain of rules that have contributed to the recompensed action. Over the long term, the utilities end up by converging towards an intrinsic “force” of each classifier.

The above mechanisms, exploited for a given field of rules, can be complemented by mechanisms of exploration. To do so, the conditions and actions are encoded, generally in binary form. The mechanisms of exploration are then based on “genetic algorithms” and perform a partial re-categorisation of the configurations. “Mutation” consists in modifying a character in the coding of the conditions of a rule, while “crossing-over” consists in mixing the codings of the conditions of two different rules. Over the very long term, the new rules replace the worst-performing old rules if they have acquired sufficient force. Such processes often converge more rapidly than in the absence of rule renewal. Above all, they make it possible to adapt to an evolutive environment (see chapter 8).

1.3.5 An example of dynamic choice

Coming back to the example of Savage's omelette (1954), we can observe that the above formalisation is based on a very precise categorisation of the decision problem. So, the state of the egg could be described with more precision, involving the date of production for example, which could influence the cook's choice. The action of the cook could itself be more discriminating, depending on whether or not he candlers the egg before breaking it into the bowl or the saucer. Above all, the consequences could be much more precise, as to the cost of manipulating the instruments or the ecological cost of throwing away an egg, for example. As usual, the problem has been stylised by the modeler in order to construct a "small world" in which it is solvable. In fact, what is important is the way in which the cook himself categorises and interprets the problem, for it is on this basis that he will make his choice.

By applying the Q-learning process to the two-egg omelette, remembering that each of the three possible configurations allows two or one associated actions respectively, the cook revises one of the five local utilities U_h^i at each step. He does not have to know in advance either the probability of the egg being bad or the costs and advantages he incurs, it is sufficient for him to experience the effects of his actions during his successive experiments. However, with the logit rule of decision, the cook will continue to test all the actions in all the configurations until the local utilities become stable and contrasting enough for him to choose his strategy according to the maximising decision rule. The use of a CPR decision rule would make it possible to get learning on local utilities to coincide with actions, while at the same time converging towards the optimum strategy (see chapter 3).

In an example of an omelette with seven eggs, one can also start with four rules:

- R1: from the fourth egg onwards, the cook uses the saucer
- R2: from the fourth egg onwards, the cook uses the bowl
- R3: up to the third egg, the cook uses the saucer
- R4: up to the third egg, the cook uses the bowl.

There are 8 possible configurations (0 to 7 eggs in the bowl) and they can be encoded in binary form (from 000 to 111). Similarly, there are 2 possible actions which can be encoded by 0 (bowl) and 1 (saucer). If we now introduce the "joker" \natural as a dumb symbol (expressing 0 or 1), rule R1, linking several configurations to a unique action, can be written: " $1\natural\natural \rightarrow 1$ ". If the probability of getting a bad egg is very low, the force of rules R2 and

R4 will increase; if the probability is very high, the force of rules R1 and R3 will increase, and for an intermediate probability, the force of rules R1 and R4 will increase. Further, a mutation performed on rule R1 consists, for example, in replacing the first joker by 1, which gives the new rule “from the sixth egg onwards, the cook uses the saucer”.

1.4 Theses and conjectures

Contrary to intentionalist decision theory, which works with few and well-established choice models, evolutionary decision theory proposes a range of models that are still diversified and lacking in unifying principles. The deliberation process highlighted the role played by two principles: the principle of prediction (to predict the effects of possible actions) and the principle of selection (to choose an action to implement). It is preceded by a third principle which appears as more and more important: the principle of categorisation (to apprehend one’s determiners and one’s environment). For each principle, according to procedural rationality, prototypical (possibly parameterised) “rules” must be established to stylise as well as possible the reasoning of the agents while at the same time covering the whole field of likely reasonings. Partial rules associated with each principle must then be grouped together into a small number of global rules, respecting the conditions of coherence between these partial rules.

For a repeated decision problem, evolutionary decision theory has established a hierarchy of interlocking time scales, which intervene hierarchically in the implementation of rules and need to be refined. Over the short term, faced with his (generally random) environment, the decision-maker implements a global rule (possibly probabilist) in order to choose an action according to his observations (possibly noised). Over the long term, the process may converge towards a punctual or cyclical attractor, and this convergence depends on the context and history, because of the random events introduced by the environment and by the decision-maker himself (through his observation and action). Over the very long term, the appearance of new possible actions (resembling random mutations) can in turn shift the position of the previous asymptotic state, in particular in order to adapt to a modification of the environment.

To judge the performance of a learning process, without claiming to attain optimal learning, the modeler (or even the decision-maker) evaluates the global rules by means of multiple, partially antagonistic criteria. On the one hand, the rule is judged by its capacity to lead asymptotically to the

optimum action, in the sense of maximisation of expected utility, at least in a stationary environment (by finding a good compromise between exploration and exploitation). It is also judged by its capacity to cope with an evolutive environment (by always maintaining a sufficient proportion of exploration). On the other hand, the rule is judged by its (generalised) cost of implementation, in terms of the information it requires, the calculations it implies and the convergence time it involves. In particular, one can examine the level of “complexity” of the rule, a concept much studied in Artificial Intelligence and as yet little integrated by economists.

In addition, the modeler can examine the empirical realism of the rules, by testing their relevance in various circumstances. He can test them through field studies, when they are put to use for political or economic choices, especially for financial choices involving high levels of uncertainty. Above all, he can test them in laboratory conditions, by subjecting the decision-maker to more or less abstract choices, which are repeatable and of which certain factors can be controlled. These tests can, under certain hypotheses, enable us to separate the rules used by the decision-maker for prediction (belief dynamics) from those used for selection (sequentiality of choices). They also bring to the fore either the asymptotic character of the process, which is, however, often attained only after a large number of periods, or the transitory character, which is itself of interest.

The rules that have just been presented can be influenced by the social network in which the decision-makers operate, even when remaining in a context where each agent acts in the face of a passive environment that is common to all. Each decision-maker can directly imitate the action chosen by others (assumed to be better informed) or draw inspiration from the best-performing actions of others, this imitation being often carried out in a limited neighbourhood of information. Each decision-maker can also imitate the rules used by others (at least if he can observe or expose them), according to their assumed virtues. This type of imitation still remains little studied. If phenomena of mimicry consequently lead to a correlation between actions and agents, the joint learning process of the decision-makers depends on context and history and may converge towards actions or rules that stay heterogeneous between different decision-makers.

If the above rules are applied to an individual choice process taking place in a random and passive environment, they must be adapted to a context comporting a large number of rational decision-makers. In fact, the number of strategies at their disposal becomes higher and their environment becomes more complex, which implies that they face harder computational constraints. The rules can easily be applied to a parametric context in

which the decision-makers only react to common signals treated as being exogenous (exactly like prices), for example in the study of elementary markets (see chapter 2). They can, more subtly, be applied to a strategic context in which decision-makers choose their actions according to the actions of the others (who do the same), namely in game situations (see chapter 3). The question of whether the selective pressure then imposed on decision-makers in their social environment does actually lead to an optimal action, or even an optimising choice rule, can then be explored once more (see chapter 7).

References

- Anderson, S.P., de Palma, A., and Thisse, J. (1992): *Discrete choice theory of product differentiation*, MIT Press.
- Binmore, K. (1988): Modeling rational players (2), *Economics and Philosophy*, 4, 9-55.
- Conlisk, J. (1996): Why bounded rationality? *Journal of Economic Literature*, 34, 669-700.
- Debreu, G. (1954): Representation of a preference ordering by a numerical function, in R.M. Thrall et al. (eds.), *Decision processes*, John Wiley & Sons, 159-165.
- Friedman, M. (1953): The methodology of positive economics, in *Essays in positive economics*, University of Chicago Press, 3-43.
- Gigerenzer, G. and Selten, R. (eds.) (2001): *Bounded rationality*, MIT Press.
- Gittins, J.C. (1989): *Bandit Processes and Dynamic Allocation Indices*, John Wiley & Sons.
- Holland, J.H. (1975): *Adaptation in natural and artificial systems*, Michigan Press.
- Holland, J.H. (1987): Genetic algorithms and classifier systems: foundations and future directions, Proc. of second ICGA.
- Laslier, J.F., Topol, R., and Walliser, B. (2001): A behavioral learning model in games, *Games and Economic Behavior*, 37, 340-366.
- Laville, F. (1998): Modélisation de la rationalité limitée: de quels outils dispose-t-on?, *Revue Economique*, 49, 335-369.
- Mattsson, L.-G. and Weibull, W. (2002): Probabilistic choice and procedurally bounded rationality, *Games and Economic Behavior*, 41, 61-78.
- von Neumann, J. and Morgenstern, O. (1944): *Theory of games and economic behavior*, Princeton University Press.
- Newell, A. and Simon, H.A. (1972): *Human problem solving*, Prentice Hall.

- Roth, A.E. and Erev, I. (1995): Learning in extensive form games: experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior*, 8, 164-212.
- Rubinstein, A. (1998): *Modeling bounded rationality*, MIT Press.
- Savage, L.J. (1954): *The foundations of statistics*, John Wiley.
- Simon, H.A. (1976): From substantive to procedural rationality, in Latsis S.J. (ed.), *Method and appraisal in economics*, Cambridge University Press, 129-148.
- Simon H.A. (1982): *Models of bounded rationality*, MIT Press.
- Sutton, R.S. and Barto, A.G. (1998): *Reinforcement learning*, MIT Press.
- Tversky, A. and Kahneman, D. (1986): Rational choice and the framing of decision, *Journal of Business*, 59, 251-278.
- Walliser, B. (1989): Instrumental rationality and cognitive rationality, *Theory and Decision*, 27, 7-36.
- Walliser, B. and Zwirn, D. (2002): Can Bayes rule be justified by cognitive rationality principles?, *Theory and Decision*, 53, 95-135.
- Watkins, C. (1989): Learning with delayed reward, PhD thesis, Cambridge University Psychology Department.

2 The elementary market

The aim of economic theory, in the broadest definition of the term, is to analyse the interactions between agents devoted to production, exchange and consumption of goods and services in the widest sense. This analysis must begin by establishing the general principles that govern individual decisions, and this is exactly what the previous chapter was studying. However, this knowledge alone is clearly not enough, because economic interactions are also dependent on the institutional contexts within which they take place. Economic schemes can take one of several different forms, according to the nature of the game rules imposed on the agents. These institutional contexts may vary considerably, as we shall have the opportunity to study throughout this book, but one of them plays a central role in economic analysis, the market, considered as the real or virtual place in which sellers and buyers meet to exchange goods or services. The market, either in its competitive form or in other forms, is the principal subject of microeconomics handbooks, and it is to the market, in its simplest form, that this chapter is devoted.

2.1 Background and problems

The pure perfect competition market is the institution which satisfies most exactly the requirements of the individualist approach. Sellers and buyers are anonymous and equal before the law, even though they may possess different economic weights due to the inequality of initial endowments. Each agent has access to the same information, namely the exogenous quality of the goods and the publicly announced price. The available possibilities of exchange are taken up by the private agents when they are seen to represent a means of increasing their utility. In such a context, the pursuit by each individual of their personal interest results in voluntary transactions. These transactions are only mutually compatible when the total of individual demands is equal to the total of individual supplies. The market is then said to be in equilibrium. The key theoretical question posed by these markets is that of their self-regulation: are the forces of competition powerful enough to result in the market necessarily attaining its equilib-

rium? This belief is shared by many economists. It lies at the heart of standard microeconomic theory.

If we consider that the essential question facing individualist societies can be expressed as follows: “how can a multitude of private decisions, taken independently by each agent on the sole basis of his individual preferences and beliefs, be made into a coherent whole?”, then the pure perfect competition market provides an exemplary answer. Firstly, the market respects the autonomy of agents in the determination of their objectives and preferences. Coordination in trading is always *a posteriori*: it involves no a priori restriction on preferences, no prior subjection that may limit the freedom of the traders by obliging them *ex ante* to respect certain collective objectives that are deemed to be legitimate or desirable. Each individual pursues what he considers to be his own personal interest. Secondly, this scrupulous respect of individual autonomy, espoused by liberal economists as the cornerstone of ethical values in trading relations, does not result in social anarchy. Competition, through the operation of flexible prices, produces a structure of mutually advantageous transactions such that individual intentions end up converging. In this perspective, it is the mechanism of price flexibility which ensures the global coherence of individual actions, what is commonly referred to as the “invisible hand”.

The characterisation of a situation of pure and perfect competition is based on a certain number of classic conditions: exogeneity of prices for the agents; homogeneous and divisible goods; perfect information; transactions without constraint. Thus, E. Malinvaud (1969) wrote: “perfect competition exists when the price of each good is the same for every agent and every transaction, when each agent considers this price to be independent of his own decisions and when he can buy or sell whatever quantity of the good he desires at this price”. It is often assumed that these assumptions require the existence of a large number of atomised agents, each of a sufficiently low weight as to have negligible influence on prices. The formalisation of the pure perfect competition market that is most widely accepted by economists is that provided by the Walrasian market. We shall now briefly describe its main principles.

2.1.1 The Walrasian market

To simplify, we shall consider an exchange economy. This is an economy without production, constituted solely of n consumers. Let us assume that this economy contains m homogeneous and perfectly divisible goods, denoted $k \in \{1, 2, \dots, m\}$. Each consumer i is endowed with an initial vector of resources, denoted $w_i = (w_i^1, \dots, w_i^k, \dots, w_i^m)$. In such an economy, the func-

tion of the market is to distribute the initial global resources between the n consumers, according to their individual preferences. Let $p \in \mathcal{R}_+^m$ be the vector of the m prices, quoted on the m markets of goods: $p = (p_1, \dots, p_k, \dots, p_m)$. It is then possible to determine, for each agent i , the excess demand for good k which we denote $e_i^k(p)$. It is calculated as the difference between his gross demand and his initial resource. From all individual excess demands, it is easy to determine the total excess demand for good k when the price equals p , as follows:

$$e^k(p) = \sum_{i=1}^{i=n} e_i^k(p) \quad (2.1)$$

Equilibrium in the market of good k requires:

$$e^k(p) = 0 \quad (2.2)$$

When this equation is satisfied for all m goods markets, we obtain a general equilibrium of the economy: all the markets are then simultaneously in equilibrium.

The foundation of this presentation is the hypothesis of price exogeneity for the agents. As we have seen above, this is one of the central hypotheses of pure perfect competition². This exogeneity can be said to be “subjective” in the sense that each economic agent assumes that the price is a given signal over which he has no influence. Under these conditions, each trader is simply adapting to the market price he observes. The agents are said to be “price-takers”. This situation is described as “parametric rationality”, as opposed to the strategic rationality of game theory which presents decisions based on the analysis of the behaviour of others, where these others are considered to be influenced by oneself (see chapter 3). As we postulate that the price is unique for all transactions rather than agreed bilaterally during each transaction, it follows that the Walrasian market is a fundamentally centralised structure. Each agent is indifferent to the action of his neighbours; his only concern is the level of prices. This essential characteristic of the Walrasian market immediately raises a question: if all the agents are “price-takers”, how are the prices formed?

² The three other hypotheses, namely perfect information, transactions without constraints and homogeneous, divisible goods, have been assumed implicitly. Firstly, all the agents know perfectly the price p , otherwise they would not be able to calculate their excess functions. Secondly, the agents have calculated their excess demands by assuming that any quantity supplied or demanded will be satisfied (absence of quantitative rationing constraints). Lastly, the goods are homogeneous and divisible.

The importance of this question can be better understood if we observe that the equation (2.2) which determines the equilibrium price in no way clarifies the process by which the market price converges to this value. The only thing this equation expresses is that the equilibrium price, if it exists, must satisfy this condition. In other words, this formalism only allows us to deal with the question of the existence of a market equilibrium, and the more difficult question of the existence of general equilibrium when the m markets are considered simultaneously. To do this, we must find a price $p \in R_+^m$ such that³:

$$e^k(p) = 0 \quad \forall k \tag{2.3}$$

Walrasian theory has successfully answered this question by proposing a set of conditions of existence such that when these conditions are satisfied, we can be sure that at least one general equilibrium exists (Debreu, 1966; Arrow and Debreu, 1954). However, this work neglects another, equally central question, that of the process which leads market prices to converge, or not, to an equilibrium value. In other words, we have demonstrated the possible existence of a price vector capable of rendering individual decisions compatible; we have yet to demonstrate that this price vector is actually attained by the market. This is a very different question to the previous one. Traditionally, it is referred to as the question of *stability*. It involves specifying the dynamic process of interactions between buyers and sellers in a market out of equilibrium and analysing its properties. Given the epistemological reflections made in the introduction to this book, it will be understood that from the perspective of evolutionary theory, the question of stability is of primordial importance. It has been stated that our approach systematically favours an understanding of economic phenomena which locates agents within their context, specifies the cognitive resources at their disposal and emphasises a sequential analysis of their interactions. This is indeed the type of project that must be pursued if we are concerned with the stability of equilibrium. We need to understand how the market behaves when it is not in equilibrium. How do the agents react? How do prices evolve? A first illustration of this approach is provided by Walrasian theory itself, in its attempt to explain the mechanism of price adjustment out of equilibrium. It is called *tatonnement*.

³ Strictly speaking, if we want to take into account goods which may not be desired, we should replace the system of equations by a system of inequations of the type: $e^k(p) \leq 0$.

2.1.2 The Walrasian tatonnement

The central hypothesis put forward by Walrasian theory to explain the source of prices is the famous *auctioneer*: it is assumed that a certain individual dedicates himself to the common cause and takes upon himself the task of setting prices. At each time t , he announces a price vector $p(t) \in R_+^m$. On the basis of this price, all the dealers communicate their excess demands to him, $e_i^k(p(t))$. Using these data, the auctioneer calculates the value of the total excess demand $e^k(p(t))$ for each good k , by adding up all the individual excess demands, as follows:

$$e^k(p(t)) = \sum_{i=1}^{i=n} e_i^k(p(t)) \quad \forall k \in \{1, 2, \dots, m\} \quad (2.4)$$

What happens when not all the markets are in equilibrium? If general equilibrium has not been reached, no transaction is carried out and the auctioneer reviews his prices. He does so by increasing the price in markets where demand exceeds supply, in other words those where excess demand is positive, and reducing the price in markets where supply exceeds demand⁴. This rule for the evolution of prices is what Franck Hahn (1982, p. 745) calls “the law of supply and demand”. It is supposed to simulate what happens spontaneously in a market in disequilibrium with perfect price flexibility. We can then write the dynamic process of evolution followed by prices in the context of Walrasian tatonnement as follows:

$$\frac{dp_k}{dt} = H_k[e^k(p(t))] \quad \forall k \in \{1, 2, \dots, m\} \quad (2.5)$$

with:

$$\left\{ \begin{array}{l} \text{if } x > 0, H_k(x) > 0 \\ \text{if } x = 0, H_k(x) = 0 \\ \text{if } x < 0, H_k(x) < 0 \end{array} \right. \quad (2.6)$$

H_k are functions which do not change sign. Their specific form depends on the “psychology” of the auctioneer, in other words, the strength of his reaction to disequilibrium. Walrasian theory seeks to display properties of convergence that are independent of the specific form of the H_k . This is because, strictly speaking, the auctioneer has no “psychology”, being no

⁴ One good, whose price remains fixed, is chosen as the reference price (numeraire).

more than a metaphor representing a pure mechanism: that of supply and demand coordination.

It should be emphasised that transactions only take place when the process of tatonnement has converged to an equilibrium price vector. As there is no money, the auctioneer is responsible for realising the transactions. To do so, he centralises all the excess supplies, and then distributes them among the buyers. The phase of action (distribution of goods) only begins when the phase of communication has ended (setting the equilibrium price).

This representation of the market is hardly satisfactory from the point of view of evolutionary economists, and it has been the subject of numerous criticisms (Fisher, 1991). Firstly, the structure proposed is extremely centralised, and so bears little resemblance to the generally accepted idea of trade exchanges (except for special cases such as some stock and raw material markets). Secondly, no exchange is allowed during the period of tatonnement before equilibrium is attained, which is hardly realistic. Lastly, in any event, even in the context of these particular hypotheses, the dynamic system (cf. 2.5) only converges to Walrasian equilibrium for very particular specifications of the excess demand functions⁵. This is a devastating result for Walrasian theory, which cannot demonstrate that price flexibility is sufficient to obtain equilibrium!

In the next part of this chapter, we shall attempt to propose an alternative framework for analysis of the market, one which meets the requirements of the evolutionary approach. We shall examine a simple market, also known as an “elementary market”, similar in its abstraction to that studied by Walrasian theory. Markets approaching more closely to the complexity of real markets will be analysed in the last section.

2.2 Canonical principles

The idea underlying the evolutionary approach is to abandon the hypothesis of centralisation, which lies at the heart of Walrasian tatonnement, and to replace it by decentralised processes of information, negotiation and exchange. Here we are poles apart from the theory described above: in our simple market, agents pair up randomly and can contract, if they so wish, at the prices they negotiate, even outside a situation of equilibrium. *A priori*,

⁵ Sufficient conditions do exist concerning excess demand functions, but they are very restrictive. For example, under the condition of gross substitutability of all the goods, the general equilibrium is unique and the tatonnement process converges.

therefore, there is nothing in such a structure to ensure price unicity. Each elementary transaction can take place at its own specific price. Contrary to the previous approach, price unicity, if it appears, must be interpreted as an emergent property, as a pure product of competitive forces, and not as the expression of a *prior* postulate provided by the hypothesis of the auctioneer. Consequently, the evolutionary approach involves the construction of an analysis in which the role and importance of decentralised transactions are fully recognised.

Before going any further in the presentation of this framework of analysis, we need to answer a preliminary question: can we justify the identification of such a structure as a global market, or should we see it rather as a series of small local markets? To put it another way: how far can we go in the decentralisation of transactions? To answer this question, we must examine the specific properties of the processes of bilateral interaction and the search for information. If the economic space is “connected” in the sense that two agents can always meet and local information is available at an affordable cost, then we can say that we are dealing with a unified market. If not, then the whole set of agents must be broken down into a group of disconnected classes, each class constituting a small elementary market.

In addition to the decentralisation of exchanges, the evolutionary approach is clearly differentiated from the Walrasian model by two further hypotheses: adaptive rationality and sequentiality. In accordance with the analyses developed in chapter 1, the economic agents considered in our model do not possess optimising rationality. Their behaviour is essentially adaptive in the sense that their action is not based on the search for a maximum, but on the comparison between what they are proposed and what they have already obtained, or between what they are proposed and what others obtain. In other words, these agents are essentially motivated by the desire to improve their situation, not (at least not directly) by the desire to optimise it. Agents reason on the basis of data provided by their environment, spatial or temporal, using what certain authors have called “situated rationality” (Orléan, 1994). There is no postulate of perfect information shared by all agents. For this reason, the cognitive dimension plays a fundamental role in the evolutionary approach. The way in which agents interpret their environment depends heavily on their past history and on their capacity to analyse it. This conception of economic rationality only assumes its full significance within a sequential context. It is then possible to understand how individual choices are progressively formed and modified. In other words, learning, in the wide sense of the term, is a central characteristic of the evolutionary agent: what he knows and what

he desires is always considered as the product of a specific history. His view of his environment depends on the path he has followed and the experiences he has had. We do not presuppose that all the agents share the same model of their environment or the same conscience of goals to be attained. Consequently, sequentiality imposes itself as an essential characteristic of our models, enabling us to consider the differentiated evolution of individual agents and the different hypotheses they adopt in their representations of the economy.

2.3 Some models

In this section, we shall examine a particular model which contains all the above hypotheses: local transactions, adaptive rationality and sequential dynamics. This model, which involves the labour market, was proposed by G. Laffond and J. Lesourne (1981). In the first sub-section, we present the model from a theoretical perspective. This presentation is necessarily a rather heavy undertaking, as the processes of search, information and negotiation must all be specified. This is the price to be paid for adopting an approach which seeks to define the context of interactions without relying on the hypothesis of perfect information. For the same reason, this modeling strategy necessarily brings into play a large number of random elements. For example, when we do not postulate *a priori* that all the agents know all the characteristics of the economy in which they operate, it is reasonable to suppose that the employees' information concerns random samples of jobs, as is the case when employees acquire their knowledge from reading only certain newspapers, drawn at random from among all the papers available, and only on certain dates. It follows that the meetings between employees and firms are equally random. Despite the complexity of the model, it has been possible analytically to demonstrate a certain number of fundamental results on the attainment of a unique price. These will be presented in this first sub-section. In the second sub-section, we shall propose an approach to the same model using numerical simulations. This will enable us not only to "visualise" the theoretical results obtained beforehand but also to bring out their significance more clearly. For example, although we can demonstrate that convergence time is a surely finite random variable and define its moments, signifying that the system actually does reach equilibrium after a certain time, we have not been able to state clearly its probability law, or even give an order of size in relation to the parameters of the model. Simulation makes it easier to understand what is

happening. We shall show that the dynamic convergence process has two phases, a rapid phase of adaptation which leads the market close to equilibrium, followed by a much slower phase which redistributes the jobs between the different employees according to their characteristics. Lastly, in the third sub-section, we shall demonstrate that the results obtained in the first two sub-sections can be generalised. To do so, we shall draw on the work of J. Lainé (1989).

2.3.1 The theoretical model

The labour market analysed by Laffond and Lesourne is composed of m workers and n jobs offered. Each worker is characterised by his reservation wage, in other words the minimal wage he is prepared to accept for working. We denote of \underline{w}_k , the minimum wage required by the individual k . Each job is characterised by a maximal wage, above which the job is no longer profitable for the firm. We denote and \underline{v}_i the maximal wage the firm will pay for the job i . With the exception of these data, assumed to be exogenous, the workers are identical to each other, as are the jobs. Note that this exogeneity disappears in a long-term context. Technical progress, the general movement of wages and the evolution of industrial relations lead to an endogenous transformation of \underline{w}_k and \underline{v}_i , but we shall not go into that question here.

Walrasian analysis of this market is based on supply and demand functions. To perform this analysis, we consider a price p announced by the auctioneer and we analyse the quantity of workers supplied and demanded. Using the data given above, it is easy to obtain that the supply of labour at price p is equal to the number of individuals prepared to work at this price, so that:

$$S(p) = \{\# \text{ of individuals } k \text{ such that } \underline{w}_k \leq p\} \quad (2.7)$$

As expected, this is an increasing function of price p . If we rank the employees in increasing order of \underline{w}_k and if we assume that each value of \underline{w}_k is unique, we obtain (figure 2.1)

$$\left\{ \begin{array}{ll} S(p) = 0 & \text{if } p \in [0, \underline{w}_1[\\ S(p) = k & \text{if } p \in [\underline{w}_k, \underline{w}_{k+1}[\\ S(p) = m & \text{if } p \in [\underline{w}_m, +\infty[\end{array} \right. \quad (2.8)$$




Fig. 2.1. Labour supply function




Fig. 2.2. Labour demand function

Likewise, the demand function, when the salary announced by the auctioneer is equal to p , can be written (figure 2.2)

$$D(p) = \left\{ \# \text{ of jobs } i \text{ such that } \bar{v}_i \geq p \right\} \quad (2.9)$$

This is a decreasing function. If we rank the jobs in decreasing order of \bar{v}_i and if we assume that each value taken by \bar{v}_i is unique, we obtain:

$$\begin{cases} D(p) = n & \text{if } p \in [0, \bar{v}_n] \\ D(p) = i & \text{if } p \in]\bar{v}_{i+1}, \bar{v}_i] \\ D(p) = 0 & \text{if } p \in]\bar{v}_1, +\infty[\end{cases} \quad (2.10)$$

The Walrasian equilibrium is defined by the equality of supply and demand, which we can write:

$$S(p) = D(p) \quad (2.11)$$

As we are considering a situation in which these curves are not continuous, there is no guarantee that such an equilibrium exists. This is because we have not considered the classic assumption of perfectly divisible goods. Furthermore, if we assume that the \underline{w}_k all have different values and that the same is true for all \bar{v}_i , then we can easily demonstrate that such an equilibrium does always exist, but that it is not necessarily unique. This is due to the particular form of the demand and supply curves deduced from equations (2.7) and (2.9), which are step functions. To avoid these difficulties without theoretical significance, we shall assume from now on that there exists a number K such that:

$$\underline{w}_K = \bar{v}_K = p^* \quad (2.12)$$

Under this condition, we can easily demonstrate that there exists one and only one value such that (2.11) is satisfied, namely p^* . What happens when we consider this market from an evolutionary perspective?

We must start by defining the search process by which the two parties (workers and firms) obtain information about the offered jobs and the workers seeking jobs, and then define the process of negotiation by which individuals modify the prices at which they are prepared to transact. The model distinguishes between periods of prospecting and periods of work. During a period of work, some individuals are unemployed while others work. The wages of those individuals who have jobs are generally differ-

ent from each other, because they are the result of independent bilateral transactions between firms and individuals. For a job i occupied by employee k , we necessarily have a salary p_i^k such that:

$$\underline{w}_k \leq p_i^k \leq \bar{v}_i \quad (2.13)$$

We assume that contracts are only signed for one sole period and that changing jobs or employees has no cost, either for the employees or for the firms.

The period of work is followed by a period of global prospection. A period of global prospection is composed of a succession of m periods of elementary prospections, each of which only concerns one individual. Each individual is active once and once only during a period of global prospection. During a period of elementary prospection, the individual concerned sends letters of enquiry to a random sample of jobs. This process is assumed to be cost-free and such that any job may appear in the sample. This condition, combined with the condition ensuring that any employee can work in any job, entails that we are in the presence of a connected market (see the first part of this chapter).

Using this information, the individual enters into negotiations. Negotiations follow a rule of adaptive rationality. More precisely, they follow a model of dynamic “satisficing” (see chapter 1) in which a decision-maker increases or reduces his aspiration threshold in relation to the ease with which he succeeded in attaining this threshold during the previous period. As a first step, these negotiations involve setting a level of requirement, in other words a minimal acceptable wage. Requirements can only be revised at the start of the period of global prospection, by means of the following procedure: an individual who has not found a job becomes more conciliatory, whereas an individual in employment becomes more demanding. In other words, an unemployed individual reduces his requirements by one unit, unless he has already reached his minimum; an individual in employment will only take another job if it offers him a wage that is higher by one unit. Firms follow a symmetrical rule: the firm whose job is always occupied will reduce its wage offer. More precisely, for a job that has been vacant for two successive periods, the firm increases its proposed wage by one unit if the maximal wage has not been reached; for a post that was only vacant during the previous period, it offers the last wage paid; for an occupied job, it reduces its wage offer by one unit. These are the requirements discovered by the employee in reply to his requests for information. How does he react to the salary requirements of the firm?

He examines the job which offers him the highest salary out of the sample of replies he has received. If this salary falls below his own modified requirements, he remains in the same position as the previous period, either with or without a job. If the wage is acceptable, he is interested in the job, and two situations are possible:

- (i) the job is vacant, in which case the individual takes it up;
- (ii) the job is occupied by an individual who refuses the new wage, in which case our individual takes the job.

If, for an occupied job, at the end of the period of global prospection, the firm has found no candidates prepared to accept a lower wage, and the incumbent worker has found no firms offering a higher wage, then these two agents renew their joint contract at the existing wage, for one period.

Taken together, these hypotheses determine a Markovian stochastic dynamic process, of which we can determine the “absorbing states”. Mathematically, these are states such that once the system attains them, it cannot leave. The states for which all unemployed individuals have reduced their requirements to the minimum, all vacant jobs offer the maximum wage and the wage offered for each occupied job is lower than the requirements of all individuals that could occupy it, are clearly absorbing states. When the process has converged to this position, it remains there. It is then possible to demonstrate that, in every absorbing state, the first K individuals are employed and the first K jobs are occupied, the wages observed being limited to the pairs $(p^* - 1, p^*)$ or $(p^*, p^* + 1)$. The states only differ in the appointment of different individuals to different jobs. The fact that the price can vary by one unit has no economic significance, because of the discrete character of the model. The price can be said to be unique.

Here we can observe a self-organisation of the market, in that the population of firms and employees separates progressively into two groups: employed and unemployed. In addition, the authors have demonstrated that the market converges in probability to an absorbing state within a finite time. Here we find the strict equivalent of the Walrasian equilibrium, the process of which constitutes an “evolutionary” justification (see chapter 3). It is therefore possible to explain the emergence of a unique price simply in terms of competitive behaviour. This model provides us with a formalisation of the market in keeping with the economic intuition that the law of supply and demand is a consequence of every individual’s search for the best exchange opportunities. This does not need to be postulated here, as it was in the model of Walrasian tatonnement. It is a natural product of competitive behaviour. But this model enables us to go even further.

We can demonstrate that some assumptions are essential for obtaining convergence to a Walrasian equilibrium state. Firstly, information is assumed to be “extensive”, so that any job can be discovered and occupied by any worker. This has been called the “connexity hypothesis”. Secondly, it is assumed that, during negotiations, a firm whose job is occupied and which has received no other candidate agrees to keep the incumbent worker at the same wage, just as an individual in a job who cannot find any better wage offer agrees to remain in the same job at the same wage. This is called an “inertia hypothesis”. Convergence to the state of equilibrium would be destroyed by a frenzied attempt by firms to reduce wages, or a frenzied attempt by workers to raise wages. The market would then fluctuate so much that wages would remain constantly dispersed. This result demonstrates that competition is not always a good thing. This is a fundamental point. Price unicity is only obtained when the information and negotiation processes satisfy some constraints.

If, in such a model, convergence to a unique price can be demonstrated mathematically, it can also be illustrated by numerical simulations. These will enable us to deepen our presentation of the model and to understand better the behaviour of convergence time.

2.3.2 Simulation analysis

The general model presented above is specified by assuming that there is an equal number of jobs and of workers $m = n$. We also assume that the value of workers’ minimum wage requirements and firms’ maximum wage requirements are uniformly distributed:

$$\left\{ \begin{array}{l} \underline{w}_k = k \text{ with } k = 1, 2, \dots, n \\ \bar{v}_i = i \text{ with } i = 1, 2, \dots, n \end{array} \right. \quad (2.14)$$

Furthermore, as we have seen, the economic agents modify their requirements during the process in relation to the situations in which they find themselves. Let $w(k, t)$ denote the requirement of worker k at time t and $v(i, t)$ the requirement of firm i at time t . We identify the firm and the job as above. Their requirements must satisfy the following constraints:

$$\left\{ \begin{array}{l} w(k, t) \geq \underline{w}_k \\ v(i, t) \leq \bar{v}_i \end{array} \right. \quad (2.15)$$

At time t , the state of the market is defined by the data $w(k,t)$ and $v(i,t)$ and by the position of each agent: whether the individual k is employed, and if so, at what wage, and likewise for the firms. We therefore define:

$tra(i,t)$: the name of the worker employed by firm i at time t with the convention $tra(i,t) = 0$ if the firm's job is vacant.

$sd(i,t)$: the level of the salary paid by firm i to its employee, with $sd(i,t) = 0$ if i has no employee.

Likewise, we denote $ent(k,t)$ the firm employing worker k at time t and $sr(k,t)$ the wage received by k . Consequently, we have $tra(i,t) = k$ if and only if $ent(k,t) = i$, and in this case: $sd(i,t) = sr(k,t)$.

For the simulations presented, the initial situation is always one in which nobody is yet employed, so that $tra(i,t) = 0$ for all i , and in which all requirements are minimal, so that $w(k,0) = \underline{w}_k$ and $v(i,0) = \bar{v}_i$ for all k and i . We shall now see how the system evolves between t and $t+1$.

It should be noted straight away that, in accordance with the assumptions presented in the first section, if a worker is employed at time t , his requirement can be written: $w(k,t) = sr(k,t) + 1$, meaning that this worker will agree to change jobs for any salary higher than the one he currently receives. Symmetrically, for any firm i , if $tra(i,t) \neq 0$, then $v(i,t) = sd(i,t) - 1$. The requirements, at time t , of unoccupied agents are determined in relation to their values at $t-1$, as we shall see later.

Firstly, we establish contact between a worker and a firm. To do so, we draw at random an individual ic , the "individual contacted", and a firm fc , the "firm contacted". The pair (ic, fc) is drawn in such a way that ic cannot be the same worker employed by fc at time t , but apart from this restriction any pair may be drawn. Having established contact, ic announces that he is prepared to work for fc for any wage equal to or greater than his requirement $w(ic,t)$ and fc announces simultaneously that it is prepared to employ ic at a wage equal to or lower than its requirement $v(fc,t)$. One of two situations may then arise, depending on the relative values of the two parties' requirements.

2.3.2.1 $w(ic,t) > v(fc,t)$: The matching comes to nothing

The structure $(tr(\cdot), sr(\cdot))$ is not modified at time $t+1$, in other words the same workers are employed in the same firms for the same wages. The requirements of the agents other than ic and fc do not change either. The only variables that can change are the requirements of fc and ic . If the individual ic was employed at time t , then he remains employed at time

$t+1$ for the same wage, and his requirement therefore remains constant: $w(ic, t+1) = sr(ic, t+1) + 1 = w(ic, t)$. If, on the other hand, ic was unemployed at time t , then he considers that the matching was fruitless because his requirement was too high. He will therefore reduce it, taking into account his minimum requirement threshold \underline{w}_k . The individual adapts his requirements progressively. We therefore obtain:

if $ent(ic, t) = 0$, then $ent(ic, t+1) = 0$ and $w(ic, t+1) = \sup(w(ic, t) - 1, \underline{w}_{ic})$.

In the same way, for the firm contacted, we obtain:

if $tra(fc, t) = 0$, then $tra(fc, t+1) = 0$ and $v(fc, t+1) = \inf(v(fc, t) + 1, \bar{v}_{fc})$.

2.3.2.2 $w(ic, t) \leq v(fc, t)$: the matching is successful

In this case, fc employs ic and the new wage can logically take any value between the requirement of the employee, denoted $w(ic, t)$, and the requirement of the firm $v(fc, t)$. We shall model the setting of the wage by drawing at random from the set $\{w(ic, t), w(ic, t) + 1, \dots, v(fc, t)\}$, as our reasoning up until now has assumed values in N .

If the firm fc already had an employee at time t , then this employee, denoted $k^* = tra(ic, t)$, is laid off, with no possibility of immediate renegotiation. He finds himself without a job: $ent(k^*, t+1) = 0$. This individual reduces his wage requirement: $w(k^*, t+1) = \sup(w(k^*, t) - 1, \underline{w}_{k^*})$.

Symmetrically, if the worker ic was employed at time t by another firm i^* , he leaves this firm, which raises its wage requirement: $v(i^*, t+1) = \inf(v(i, t) + 1, \bar{v}_{i^*})$. The other variables remain the same.

As we have seen above, the stable states of the process are those which satisfy the two following properties: no matching can be successful; all unoccupied workers and firms have reduced their requirements to their minimum (or maximum) thresholds, namely \underline{w}_k or \bar{v}_i . We can then demonstrate that the stable states correspond to states of economic equilibrium in terms of equality of supply and demand and price unicity. In a stable state, all the wages paid are equal (to within one unit, because of the discrete character of the model). Convergence time is, by definition, the first period in which the economic system enters the stable state.

Another useful concept for studying the system is that of surplus. By definition, the surplus produced by firm i employing worker k is equal to $v_i - \underline{w}_k$. The total surplus at time t is the total of all the surplus produced by the firms whose jobs are occupied at that time. One can demonstrate that, in keeping with economic intuition, when a state is stable, its surplus

is maximal. However, the contrary statement is not true. It is possible to have a situation in which the populations active at time t are those that will be active in the stable state, but where wages have not yet levelled out. In this case, the surplus is at its maximum but the system has not yet converged. Furthermore, for the system to stabilise, it will be necessary for some firms to lay off workers, and the surplus will therefore decrease temporarily, before rising later on.

Figure (2.3) illustrates a typical trajectory for surplus. From the initial situation, where it is nil, it increases sharply during the first hundred periods. Subsequently, it rises more slowly. It reaches its maximum for the first time at about $t = 250$, then decreases until $t = 300$ before reaching its maximum for the second time at about $t = 340$. From this moment on, the surplus no longer changes. However, the system has not yet attained its stable state. Convergence takes place at $t = 400$. What happens between these two periods? During this time, the requirements of unoccupied workers and firms continue to fall until they reach their minimum values.

By carrying out numerous simulations using the same initial situation, we can construct an empirical distribution of the convergence time. This distribution turns out to be highly skewed. Figure (2.4) represents the chart




Fig. 2.3. Evolution of surplus




Fig. 2.4. Chart of convergence time distribution

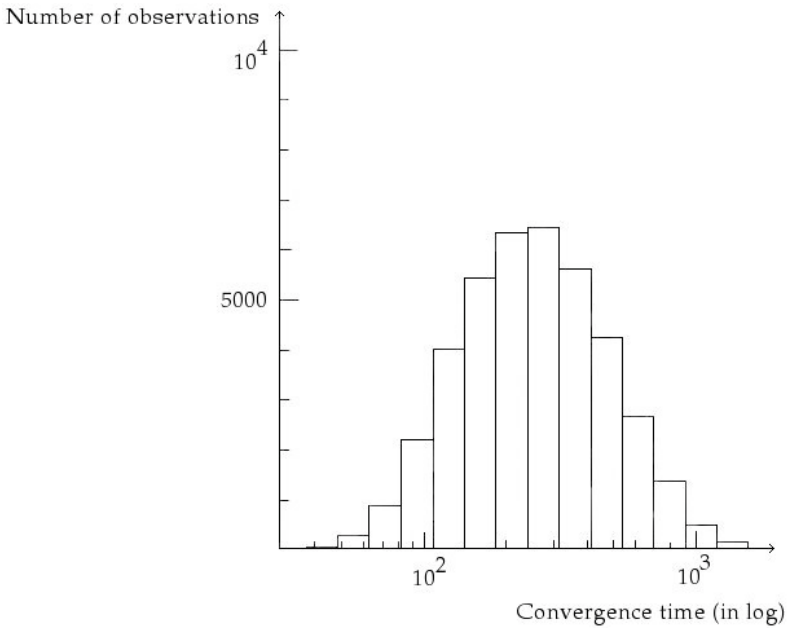


Fig. 2.5. Chart of distribution of convergence time logarithms

of convergence time distribution for the model described above with 10 workers and 10 firms and with over 40000 simulations. To recover symmetry, we can convert to logarithm. Figure (2.5) thus represents the chart of distribution of convergence time logarithms, constructed using the same data as above.

This approach, using simulation, finds an echo in experimental market economics. This latter makes it possible to observe transactions one by one, in a controlled environment. It sheds light on the way in which simple markets equilibrate over time, like the theoretical model described above. In particular, certain famous experiments have explored “double auction” markets, in which agents artificially endowed with private, exogenous reservation values submit supply and demand prices (Gode and Sunder, 1993; Easley and Ledyard, 1993).

2.3.3 A generalisation

We shall now see how the ideas presented above can be generalised to apply to exchange economies comporting m goods. Let $I = \{1, 2, \dots, n\}$ be the set of agents and $M = \{1, 2, \dots, m\}$ the set of goods. As above, we shall assume that the transactions are performed bilaterally. Formally, the meeting between two individuals i and j is represented naturally as a function π of I in I such that:

$$\left\{ \begin{array}{l} \pi(i) = j \\ \pi(j) = i \\ \pi(k) = k \text{ for } k \text{ different from } i \text{ and } j \end{array} \right. \quad (2.16)$$

Traditionally, such an application is called a “transposition”. Consequently, the succession of bilateral meetings is represented by a series $\pi^t_{t \in \mathbb{N}}$ of transpositions on I : $\pi^t(i) = j$ meaning that the pair $\{i, j\}$ is formed at time t . The properties of this series can be variable: endogenous or exogenous, determinist or random. In keeping with the canonical principles defined earlier, the very idea of a unified market in which competition can fully exert its influence requires a certain level of connexity in meetings, to the extent that every agent must be able to compare himself with every other agent. Subsequently, following the work of J. Lainé (1989), we shall examine a regular series of meetings, in other words a series of meetings constituted of an infinite succession of cycles of the same finite length such that every pair of agents is formed at least once during each cycle.

Now that this process has been defined, we must specify the transactions that will occur between agent i and agent j . One of the two agents is chosen, at random, as leader. Let us assume that this is agent i . He is then given the task of proposing the prices that will be applied during the transaction. As in the previous model, this relinquishment of the “price-taker” hypothesis is necessary because we no longer assume the presence of an auctioneer. In responding to the prices proposed by the leader, agent j is assumed to act passively and non-strategically. His sole reaction consists in proposing either (1) to carry out the transactions, within the budgetary set imposed by the leader, which enable him to maximise his utility, if doing so improves the situation he was in before the exchange; or (2) to keep the *status quo* if all the transactions, at the prices imposed, worsen his initial situation. We assume that the leader is perfectly informed about his partner’s behaviour. Consequently, he chooses the price system which maximises his own utility, taking into account the reaction of the other agent. Such an asymmetric procedure of negotiation is called a “Stackelberg procedure”. This is not to be confused with the Edgeworth procedure, a symmetric bilateral procedure in which, in a not completely explicit way, the two agents set prices so that the carrying out of exchanges improves both of their utilities simultaneously.

This negotiation procedure defines a monotonic process, in other words a process which guarantees the non-decrease in the agents’ utility at each period. This procedure is undeniably very basic, as J. Lainé observed (1989, p. 49-50): “If, in a situation of unequal distribution of information, the establishment of a Stackelberg equilibrium appears probable in a game where the players only play once, repetition of the game and the acquisition of information about the strategies of others that this entails can only encourage the passive agent to manipulate the leader’s choices to his own advantage, by misrepresenting his preferences or by lying about the state of his endowment, for example. Insofar as that we neglect any dynamic interaction here between successive negotiations, the evolution of the market is the result of the superposing of totally myopic and amnesic strategies. Moreover, no description is offered of the means of selection of the leader”. The same myopia and amnesia were present in the previous model. It should be noted that, in this type of model, it is the large number and regular nature of the meetings which operates automatically to spread information and to enable adaptation of behaviour. Here, time is a central variable in the conception of individual and collective learning. This is a characteristic of the evolutionary approach: the duration of interactions compensates for the limited rationality of individual agents.

For each series $\pi^t_{t \in N}$ and each initial state ω^0 , there is a corresponding set of paths of allocations $\{x^t\}_t$, as well as of prices $\{p^t\}_t$. The initial state ω^0 specifies how, at time $t=0$, the resources are shared out between the n economic agents. Then, the meetings take place in accordance with $\{\pi^t\}_{t \in N}$. At each time t , following the meeting π^t , a new price vector p^t is proposed and exchanges take place which lead to a new distribution x^t of resources between the economic agents with $x^t \in R_+^{mn}$. No consumption can occur until the process has converged. What interest us here are the asymptotic properties of these paths. If the above assumptions, together with the traditional assumptions about the form of individual preferences⁶ are met, one can demonstrate the following theorem: if the sequence $\{\pi^t\}_{t \in N}$ is regular, all paths $\{p^t, x^t\}$ converge to a price equilibrium $\{p, x\}$. Equilibrium signifies (1) that each individual allocation maximises the utility of all the agents, taking into account budgetary constraints and (2) that the allocation x is realisable, so that $\sum_{i \in I} x_i \leq \sum_{i \in I} w_i^0$ with $x = (x_i)_{i \in I}$ and $w^0 = (w_i^0)_{i \in I}$, where these latter denote the initial allocation of resources for all agents. Naturally, several accessible price equilibria exist *a priori*.

This result provides a partial generalisation of the results obtained earlier for the labour market. It represents an important step towards the satisfactory modeling of a decentralised economy. The interest of this result lies in its capacity to demonstrate that a succession of mutually independent bilateral negotiations, which therefore exhibits *a priori* strongly different prices from one period to another, can lead the market to a situation in which a system of perfectly unified prices prevails, a system which no agent wishes to bring into question. However, production is not taken into account here. We are analysing an economy which is limited to sharing out the initial resources between the consumers. Furthermore, contrary to the labour market presented earlier, this model retains the hypothesis that consumption is not permitted until equilibrium has been attained. The final step must be reached before consumption can be authorised. The relinquishing of this hypothesis can be very costly, as the total stock of goods exchanged will then vary from one period to the next. In the labour market model, working energy is consumed during each period, but it remains constant throughout the process because it is reproduced at an identical level from one period to the next. This can be explained by the particular nature of work as a commodity.

⁶ Strictly convex, weakly monotonic and differentiable preferences.

2.4 Theses and conjectures

These first models have brought to the fore a property of capital importance in the evolutionary approach, namely self-organisation. Self-organisation exists when a process generates properties or qualities at the global level which did not exist at the individual level. For example, the labour market analysed in this chapter draws a clear frontier, within the body of workers and firms, between two rigorously distinct groups – the active and the inactive – where analysis of the individual situations cannot demonstrate any qualitative difference. This emergence of new qualities reveals to the analyst the fact that the economic system has attained a higher form of order. We can interpret price unicity in the same way. It expresses the emergence of an ordered structure under the effect of competition, when certain conditions are satisfied.

A second fundamental characteristic of these models can be found in the difference that exists between, on the one hand, the sophistication of the emergent collective properties and, on the other, the simplicity of individual behaviour. The modeling that has been presented does not assume sophisticated cognitive capacities on the part of economic agents. Most often, they are not even capable of maximisation⁷. They content themselves with comparing situations in order to choose the one which they consider the most profitable. This crude rationality (myopia, amnesia, responsiveness, comparison, inertia) is offset by the evolutionary process itself, which leads to the efficient adaptation of both the agents and the system. In evolutionary models, the duration of interactions and their repetitive character compensate for the limited rationality of individual agents. Intelligence resides in the duration of collective learning. Everything happens as if the extreme sophistication of strategic thinking – constantly expecting the actions of others, once they have expected one's own action, – was rendered unnecessary by the effective movement of exchanges and meetings, which imitates reasoning.

However, placing the emphasis on temporal dynamics obliges us to give their full weight to the accidents which punctuate the process and which lead to modifications in beliefs and behaviour. We use the term “path dependency” to describe a final state which depends on the random events of

⁷ In the last model, we assume that the leader sets the price in such a way as to maximise his utility once he knows the passive reaction of his trading partner. However, such behaviour is not introduced in the labour market model.

the trajectory. From an evolutionary perspective, history matters. For example, in the generalised exchange model, contrary to the labour market model, the final price obtained depends on the history of the meetings. Likewise, in the case of the labour market, the way in which the jobs are allocated between the workers depends on the history of the system.

The main difficulty lies in the generalisation of the model to deal with a group of markets. The generalised exchange model does lead to the decentralised formation of prices, but only when we hypothesise a preliminary phase of sequential price determination during which no concrete transaction is authorised. A combined mechanism of price formation and effective exchanges, like the one presented in the labour market model, remains to be formalised for the case of multiple goods. In addition, the production of goods also remains to be incorporated into the model. So it can be seen that we are still far from possessing a truly satisfactory formalism, beyond the limited progress which has been covered in this chapter.

On the other hand, certain phenomena can be examined, taking the form of “friction”: search costs, transaction costs and irreversible decisions. These phenomena, which can be internalised more easily, will be dealt with in chapter 4.

References

- Arrow, K. and Debreu, G. (1954) “The existence of equilibrium for a competitive economy”, *Economica*, vol. 20, 265-290.
- Debreu, G. (1966): *La théorie de la valeur*, Dunod.
- Easley, D. and Ledyard, J.O. (1993): “Three theories of price formation and exchange in double oral auctions”, in D. Friedman and J. Rust (eds.), *The Double Auction Market*, SFI Studies in the Sciences of Complexity, Proceedings, vol. XIV, Addison Wesley, 63-97.
- Fisher, F. (1991) : “La formation des grandeurs économiques: déséquilibre et instabilité” n: Cartelier (ed.), *La formation des grandeurs économiques*, Paris, PUF, 19-55.
- Gode, D.K. and Sunder, S. (1993): “Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality”, *Journal of Political Economy*, vol. 101, 119-137.
- Laffond, G. and Lesourne, J. (1981): “Markets dynamics and search processes with information costs”, communication at the European Conference of the Econometric Society, Amsterdam.

- Hahn, F. (1982): "Stability", in K. Arrow and M. Intriligator (eds.), *Handbook of mathematical economics*, vol. II, North-Holland Publishing Company, 746-793.
- Lainé, J. (1989): "Processus d'échanges bilatéraux et auto-organisation du marché", *Economie Appliquée*, XLII, n° 3, 41-61.
- Laslier, J.F. (1994): "Note sur la simulation du marché simple auto-organisé", mimeo.
- Lesourne, J. (1991): *Economie de l'ordre et du désordre*, Economica.
- Malinvaud, E. (1969): *Leçons de théorie microéconomique*, Dunod.
- Orléan, A. (1994): "Vers un modèle général de la coordination économique par les conventions" in A. Orléan (ed.), *L'analyse économique des conventions*, Paris, PUF, coll. "Economie", 9-40.

3 Game situations

Since its foundation, “classical game theory” studies the strategic relations established between several actors of a social system and holds privileged links with “standard economic theory”. Like economic theory, game theory is grounded on two main assumptions asserting respectively that each actor follows a strongly rational behavior and that all actors are coordinated through some equilibrium notion. But contrary to economic theory, game theory considers direct relations between actors (such as bargaining relations) rather than relations mediatized by prior institutions (such as market transactions). Moreover, game theory is defined at a higher level of generality than economic theory as concerns the actor’s relations (any type of multilateral actions instead of exchanges of goods), hence the actor’s characteristics (any type of preferences instead material interests).

In a first step, game theory appeared as a convenient tool able to treat the empirical discrepancies observed with the Walrasian general equilibrium model, when some strategic relations between actors, not previously considered, are introduced. Some situations of imperfect competition were formalized in which a reduced number of producers are mutually confronted in relation with passive consumers and fix simultaneously the quantities and prices of exchanged goods. In a second step, game theory allowed to interpret the Walrasian general equilibrium as a limit case of a game equilibrium when the strategic dimension of agents’ interactions disappears. The economic system is seen as a big game between many producers and consumers, endowed with specific characteristics and interconnected by a price system, their number being progressively and proportionally increased.

More recently, “evolutionist game theory” points to the dynamic interactions between actors inserted in a social network, and becomes very close to a parallelly developing “evolutionist economic theory”. More precisely, both theories are grounded on bounded rationality of the actors and study interaction processes repeated over time, an equilibrium state being only eventually obtained as an asymptotic emergent pattern. Both theories consider that each actor has mainly direct relations with other ones in some neighborhood, with some institutions which appear themselves as emergent structures facilitating those relations. By such, evolutionary game theory abandons its strategic dimension, by neglecting the

crossed expectations of actors about their respective actions in favor of purely reactive actions of the actors in response to the observation of their environment.

The only remaining difference of evolutionist game theory with regard to evolutionist economic theory lies, as for the classical versions of both theories, in its higher level of generality. It proposes a general language for representing the adaptive behavior of individual actors and the collective consequences deriving from their interactions, with some economic applications to auction mechanisms or contractual agreements. It proves theoretical results concerning the transitory behavior and essentially the asymptotic behavior of the interaction processes, with some economic applications to emergence of institutions or diffusion of technologies. More globally, it suggests a conceptual frame which is not restricted to economic phenomena, but extends to political and social ones, even if the results obtained are drastically limited by the complexity of the phenomena at hand.

The first part of the chapter examines the methodological foundations of evolutionist game theory with reference to classical game theory. The second part makes explicit the canonical principles already worked out or potentially existing for elaborating some specific model. The third part introduces prototypical classes of models, based on learning or evolution processes, and summarizes the result they provide. The fourth part examines the achievements of evolutionist game theory and its position with regard to the social sciences.

3.1 Background and problems

3.1.1 Principles of classical game theory

Initiated by the founder book of von Neumann and Morgenstern (1944) and pursued by the pioneering work of Nash (1951), classical game theory is grounded on some simple interaction principles between players. The players match in either sequential or simultaneous meetings, define and implement some corresponding actions and get payoffs from the consequences of the combination of all actions. The players choose their actions in a strongly rational way, hence optimize their expected utility under various constraints, with regard to their beliefs about their material environment and their opponents. The players are coordinated in some equilibrium state, i. e. a stationary state (in absence of perturbations from outside

the system) from which no player has an interest to deviate unilaterally. These three principles are examined successively.

In the first place, the game is structurally described by the opportunities of each player, the utility he gets from the conjunction of opportunities and the beliefs he holds about the preceding characteristics. In “normal form” (or “strategic form”), the game is formalized by a game matrix, which combines the possible actions of the players (and eventual states of nature) and expresses the utilities which are obtained for each combination of actions (and states). In “extensive form” (or “developed form”), the game is formalized by a game tree, which details the successive alternate moves of the players (and eventual states of nature) and indicates the utilities which are obtained for each path in the tree. The extensive form can be reduced to the normal form thanks to the notion of strategy, a strategy being defined by the action the player plays at each node where he has the move.

Then, each player is endowed with a strong rationality, both cognitive rationality in order to adapt his representations to his information and instrumental rationality in order to adapt his means to his objectives. Strong cognitive rationality implies that the player forms a perfect expectation of the future play of the game, by considering a correct specification of the game structure and a perfect and complete information on its past play. Strong instrumental rationality involves that the player chooses an action by maximizing his expected utility (taking into account the random environment) under various individual and institutional constraints. Note that the same rationality applies to the actor for individual decision in a risky environment and in game theory, the “strategic uncertainty” about the opponent being in some sense “naturalized” in a “physical uncertainty”.

Finally, an equilibrium notion characterizes the social configurations in which the players’ actions are compatible, and leads to study the existence and multiplicity of the corresponding equilibrium states. In a normal-form game, a Nash equilibrium is a state where each action of a player is a best response to the others (equilibrium) actions, hence results from a fixed point of the best response functions. In an extensive-form game, a sub-game perfect equilibrium is obtained by a backward induction procedure, each player playing at each time his best action, knowing the future best actions of all players. If these definitions explain the stability of some equilibrium state once established, they do not explain how an equilibrium state is achieved, except by introducing a fictitious entity, the Nash regulator, who appears as the precise counterpart in game theory of the Walrasian auctioneer.

3.1.2 Limits of classical game theory

Classical game theory, in its original form, is based on very stringent assumptions as concerns the behavior and coordination of the players, but has tried to weaken them progressively. The players may be endowed with uncertain beliefs, either structural uncertainty on game rules and opponents' (or nature's) characteristics, or factual uncertainty on past opponents' actions (or nature's states) (Harsanyi, 1967). The players are endowed with bounded individual rationality, which constrains their reasoning abilities and leads them to chose suboptimal actions (Rubinstein, 1994). The players' actions are no more coordinated towards an equilibrium state by the Nash regulator, but by crossed expectations on their actions and on their underlying determiners (Aumann, 1974). These limits will be examined successively.

First, classical game theory attempted to internalize the informational limits of the players by introducing sophisticated forms of uncertainty into the usual equilibrium notions. On the one hand, if he is not well aware of the opponents' determiners assumed to be summarized in their types, a player gets endowed with a probability distribution on their possible types, eventually obtained from a common probability distribution. On the other hand, if he is not aware of his opponents' past actions (symbolized by different nodes in the game tree), a player treats globally as an "information set" the nodes of the tree between which he cannot discriminate. More generally, it is usually assumed that any type of uncertainty can be reduced to a probabilistic form and even that this uncertainty is common knowledge among players ("each knows that the other knows that ...").

Second, classical game theory tried to internalize the computational limits of the players by introducing various forms of (cognitive) bounded rationality in the equilibrium notions. On the one hand, a player may be considered as acting as a finite state automaton, which is only able to implement strategies which are not too complex, especially strategies which only involve a bounded memory. On the other hand, a player is considered as a limited reasoning device, constrained by computing costs or as trained to finite internal states, which prevents him for instance to consider crossed beliefs above a given level. More generally, the players may still try to optimize their behavior on simplified decision problems or to treat the real decision problems with methods that are only approached ones.

Third, classical game theory gave "cognitive justifications" to its usual equilibrium functions, by assuming that the players are able, by their sole

reasoning, to substitute for the Nash regulator. Contrary to weaker equilibrium notions, the justification of Nash equilibrium requires not only common knowledge of the game structure, of the players' rationality and of the players' independence of play and a common prior on actions, but even common knowledge of their respective conjectures, i. e. their expectations of others' actions. The subgame perfect equilibrium is easier to justify, since it just needs a common knowledge of the game structure and players' rationality, as well as of players' independence of play. In fact, the players are again endowed with exceptional informational and computational capacities, not only in order to single out an equilibrium notion, but to select an equilibrium state in case of multiplicity.

3.1.3 Principles of evolutionist game theory

Evolutionist game theory, initiated twenty years ago (Maynard Smith, 1982) and still developing (Weibull, 1995; Young, 1998), aims at introducing soft assumptions concerning the adaptation of players along their dynamic interactions. The players have some reduced information on their opponents' characteristics together with a limited spatially local and temporally bounded information on the past play of the game. The players are endowed with bounded rationality, their beliefs reducing at best to extrapolative expectations and their actions contenting with going in the right direction with regard to objectives. However, they are engaged in sequential encounters, in such a way that the repetitive work of time makes up for the role of the Nash regulator or of the crossed expectations. These three features will now be specified.

In the first place, evolutionist game theory considers the information gathered by a player as a by-product of his actions, considered as material rather than informational ones. Information concerns exceptionally some aspects of the game structure, either the utility (for the concerned player or his opponents) which may result from players' conjoint actions or even the opponents' prior beliefs. Information concerns more usually the sequence of actions implemented by the player's opponents, at least in some "information neighborhood", the player being assumed to know his own past actions. Information concerns finally the consequences of conjoint past actions, the utilities that the player gained by these consequences, and even the utilities that the other players got through their own actions.

As well, evolutionist game theory conceives the player's rationality as essentially adaptive, in the sense that the player reacts to the past information by adjusting the parameters of a given choice rule. Cognitive rational-

ity is grounded on a stationarity assumption as concerns player's environment, the observed past actions or utilities being supposed to continue in average in the future. Instrumental rationality is grounded on an improving rather than an optimizing point of view, the choice rules being conceived in order to shift incrementally the action towards increasing utility. Moreover, the choice rules, now made duly explicit, incorporate the possibility of random deviations from their normal course, in order to explore from time to time the consequences of actions not spontaneously implemented.

Finally, evolutionist game theory postulates that the players, sometimes gathered in homogenous populations, meet systematically or randomly in a repeated and generally infinite sequence of repetitions of the same stage game. Each player meets only players situated in an "interaction neighborhood", grounded on a spatial proximity criterion or a qualitative similarity criterion and possibly evolving endogenously over time. The players are engaged in a global dynamic process which, in some cases, converges asymptotically towards an equilibrium state, and gives a "evolutionist justification" to the associated equilibrium notion. Moreover, if the process converges, it leads (at least in probability when including random terms) towards some equilibrium state univocally defined by initial conditions, given context and past history, and gives therefore a constructive answer to the equilibrium multiplicity problem.

3.1.4 Taxonomy of evolutionist game models

In a first family of game models, the players follow an "epistemic learning" where they revise progressively their beliefs according to their observations and where they define consequently their actions. Each player is endowed with structural information about his own utilities stated in the stage game matrix and receives factual information about the opponents' past sequence of moves. He has a weakened cognitive rationality which allows him to revise his beliefs on his opponent through various heuristic rules and to expect (at least in a probabilistic form) his future action. He has a myopic instrumental rationality since, if he always maximizes his utility function given the expectation of other's behavior, his best response is computed only as concerns the present period and not the future ones.

In a second family of game models, the players follow a "behavioral learning" where they reinforce the actions which have shown a good performance in the past and inhibit the actions which have shown a bad performances. Each player possesses no more structural information about others' characteristics and obtains factual information only on the utilities

he got with his own past actions (and eventually the utilities got by the others). He has a cognitive rationality reduced to the conviction that the actions which succeeded in the past, in the sense of an aggregated utility index associated with each action, will still succeed in the future. His instrumental rationality is even more bounded, since he chooses randomly an action, however without neglecting anyone, with a probability which is increasing with its index.

In a third family of game models, the players follow an “evolutionary process”, inspired by evolutionary biology, where they reproduce in proportion to the payoffs they obtain in random and local interactions. Each player is decomposed in several subpopulations of agents using all the same strategy, subpopulation which is characterized by its proportion in the whole player-population. Each agent, hence endowed with a fixed strategy, renounces to any information (except for implementing his strategy if depending on the past) and has a cognitive as well as instrumental rationality reduced to nil. Each agent reproduces increasingly with the utility he obtains through a random interaction, in virtue of a “selection mechanism” which is not made explicit, and may nevertheless change his strategy along an exogenous random law, reflecting a “mutation mechanism”.

These families of models are relevant according to the game context, epistemic learning allowing a fast adaptation to a regular environment while behavioral learning allows a slow adaptation to a turbulent environment. Moreover, the models can be combined in hybrid ones, for instance when players are guided by some type of learning and simultaneously selected along their expectation or choice rules. Finally, the behavioral learning models are isomorphic to the evolutionary ones, the probability for a player to choose a strategy being replaced by the proportion of agents of a given subpopulation to use that strategy. The two first families appear then as the most interesting since they obey precisely to methodological individualism and allow the players with a sufficient rationality (Walliser, 1998).

3.1.5 Limits of evolutionist game theory

Evolutionist game theory, under its primitive form, is grounded on rather restrictive assumptions, even if it tries to compensate for them in its contemporary developments. Information gathering, reduced to a passive form, implies no voluntary process induced by the player in order to get original data and is insensible to the problem of ambiguity of the obtained messages. Behavior rules, always defined exogenously, do not make precise the cognitive processes necessary to implement them, and do not con-

sider their possible dynamic adaptation to context and history. The global system stays stationary in its structure, without a chance for the environment to evolve or innovate and without the consideration of nested time scales leading to corresponding equilibria. These limits are successively examined now.

First, evolutionist game theory considers information as resulting exclusively from the immediate and transparent observation of past spontaneous actions of the players and of their consequences. Players have no direct ways of communicating (except when one introduces initially communication actions), which could allow them to share their plans of action and even exchange their objectives or beliefs. They do not try to reveal from the opponents' observed actions their underlying preferences or representations in order to rebuild at least some part of the missing structural information. They face no difficulty in giving a clear and univoque sense to the information they get, and even succeed in giving a common interpretation to the shared information without need for further concertation.

Second, evolutionist game theory assumes that the players follow expectation and choice rules which, even if of various degrees of sophistication, are endorsed definitely and implemented in a mechanistic way. These behavior rules, even if they precisely associate a final action to the past observations, are not adapted to context since a player may refine his behavior according to its analyzed complexity and instability. These behavior rules, even if their parameters are adjusted to past experience, are not adapted to history since a player may try to change his mind if he observes that he is "locked in" a suboptimal situation. These behavior rules, even if flexible, are not integrated in a whole hierarchy of learning modes where a player shifts from one level to another if his actions are not producing the expected results.

Third, evolutionist game theory assumes that the game is implemented by the players without surprises, since it is astrained to a fixed structure as concerns the players' characteristics and their common environment. At short term, the players adjust progressively their actions to their observations and innovate at best by their purely random deviations from their spontaneous actions without considering original actions. At long term, if the system converges, it is towards a ponctual, cyclical or chaotic attractor, even if the random factors may direct him towards such and such equilibrium state or even lead it to jump from one equilibrium state to another. Globally, the process is not structured along several temporal scales, defined in function of the speed of evolution of environmental variables and the speed of adaptation of players' action variables.

3.1.6 Epistemological status of evolutionist games

Evolutionist game theory, which develops models which are more diversified but less sophisticated than those of classical game theory, differs from it on epistemological grounds too. An evolutionist model is no more studied by exclusive appeal to analytical resolution methods, but is analyzed by simulation methods which aim at marking out the space of consequences. An evolutionist model still attributes to players some mental states, but these are less revealed through players' actions than directly assessed by interrogation of players about their expectations or utilities. An evolutionist model is not considered as normatively recommending prescriptions to the players, but essentially considered as positively providing descriptions of players' behaviors. These questions are successively examined.

First, since the classical models are in few numbers due to the unique expression of the players' strong rationality and the lack of prior players' networks, the consequences in terms of equilibrium states are calculated by analytical way. Conversely, since the evolutionist models are far more various and complex due to the extreme combinatority of players' behaviors and relationships, the dynamic consequences are sometimes necessarily computed by simulation. Such a simulation, achieved by the modeller and unreachable by the players, describes the transitory and asymptotic consequences of any system and allows to study its robustness. It also brings to the fore the multiplicity of models explaining the same phenomenon, hence the necessity to compute all the testable consequences of each model, in order to be able to empirically differentiate them.

Second, for classical models, under the assumption of players' strong rationality, it is usual to consider that their beliefs and preferences may be and have to be revealed from the chosen actions, which are the only observables. At the contrary, for evolutionary models, the mental states shift from an instrumentalist status to a more realistic status, especially as concerns the utility attributed by the players to the consequences of their actions. In learning models, utility is always treated as a mental state, but one assumes that it can be directly expressed by the players, either experienced as concerns its past occurrences or expected as concerns its future occurrences. In evolutionary models, utility is reduced to "fitness", which is unconscious for the players and is ideally observed by the modeller through the reproduction rates assumed to be proportional.

Third, in classical models, the equilibrium strategies are often considered as strategies suggested by the modeller to strongly rational players, and are judged as concerns their mutual compatibility ensuring their stability. Conversely, in evolutionist models, the dynamic strategies are always

considered as resulting from behaviors described by the modeller, without a precise justification from the point of view of the players. According to that shift from a normative to a positive vision, in order to confront the models to the facts, the laboratory experiments are increasing, at the frontier of economics and cognitive psychology. It follows essentially a “projective form”, when testing the consequences of pre-given models, rather than an “inductive form”, when extracting some behavior regularities from free observations.

3.2 Canonical principles

An evolutionist game can be defined by means of five principles stated by the modeller: satisfaction, confrontation, information, evaluation and decision. The satisfaction principle refers to the utility obtained by each player, which is supposed to depend on the strategies played by all the interacting players. The confrontation principle gives insights into the way people interact, namely by stating the meeting assumptions. The information principle is about information gathering: each agent is meant to collect information on the way the game is played, for example by looking for past played actions and/or past associated utilities. The evaluation principle makes precise the way information is processed and interpreted, in order to get condensed aggregated information, on which the future choices may be based on. The decision principle refers to the many potential ways players take their decision, given the aggregated information at their disposal.

3.2.1 The satisfaction principle

The satisfaction principle is the only principle common to both evolutionist and classical game theory. It requires that, in addition to the set of players, the set of strategies available to each of them, one specifies for each player a utility function, which defines his payoff at each issue of the game (which may be repeated). It specifies the set of players, the set of strategies available to each of them as well as his utility function, i. e. his payoff at each issue of the stage game. The opportunities and preferences of the players are supposed to be given (to the modeller) at the beginning of the game and stay unchanged.

The standard illustration, systematically used in the future, is the “technology game”. It represents the coordination problem faced by two firms, 1 and 2, which have to choose among two technologies, A and B. The technology B, in contrast to the technology A, is a state-of-the-art tech-

nology: both firms, if they both choose B, better perform -they both get 4 - than if they both choose A -they both get 2. When one firm uses technology A and the other technology B, the first firm gets b and the second c , these parameters being further specified. Note that this game is symmetric whatever the values assigned to b and c (firms A and B have the same strategy set and achieve the same payoffs in symmetric issues). The corresponding game matrix, a symmetric matrix, is the following:

Table 3.1. Technology game

	A2	B2
A1	(2,2)	(b,c)
B1	(c,b)	(4,4)

In a first variant of the technology game, looking like Rousseau’s stag hunt game, one states $b=1$ and $c=0$. It means that the inferior technology A, better mastered than technology B, can be used alone (with a reduced payoff). Conversely, the superior technology B, which needs to develop, is bad when used alone. In this variant, the technology A is less risky than technology B, in that it yields a payoff which less depends on the choice of the other firm (the payoff is between 1 and 2 for technology A, in contrast to 0 and 4 for technology B). The corresponding matrix is the following (variant 1):

Table 3.2. Technology game, variant 1

	A2	B2
A1	(2,2)	(1,0)
B1	(0,1)	(4,4)

The stage game may be isolated or part of a bigger game, which is potentially much more complex. A stage game is isolated when the payoffs of the confronted players at a given period are independent of the behavior of all other players, which potentially play the same game, at the present (and preceding) periods. Some degree of isolation is necessary for a game to be studied in an evolutionist way; more precisely, in case of lack of the isolation assumption, it is possible to address the game only if the stage game

smoothly changes from one period to the other, due to the evolution of the global game, whose impact is felt only progressively. So the technology game is perhaps not necessarily an isolated game: the prices of both technologies A and B, and therefore the benefits achieved with each of them, may depend on the behavior of other firms having to choose among the technologies A and B (and possibly other technologies), as well as on the behavior of consumers supposed to buy the product produced by means of the different technologies.

3.2.2 The confrontation principle

Players play repeatedly, a finite or infinite number of times, a n -player stage game. The stage game is a non cooperative game in the usual sense, in normal or extensive form. The payoffs of the players are aggregated thanks to some discount rate (see chapter 1).

The n -player stage game is supposed to be played by n players or, more usually in evolutionist games, by n populations of agents, each agent of the population i playing the role of player i . In each period, several n -uplets of individuals are randomly drawn from the n populations (or subsets of these n populations), one agent from each population, and each n -uplet of agents plays the game. The interactions may be more or less numerous at each period. On one side, a single n -uplet is constituted in a random way. On the other side, all possible n -uplets are formed. In the *technology game*, if each firm is represented by an equal number of agents (each agent having a given technology), each agent of one population may meet one firm of the other population or many combinations can be sorted out.

A usual distinction about meetings concerns the “multi-population” or the “mono-population” approach. The multi-population approach is available for any game and corresponds to differentiated players. Agents from a given population meet agents from the other populations. By contrast, the mono-population approach is reserved to symmetric games when players are considered as interchangeable. Agents form a unique population and meet any agents from that population. For a symmetric game, the two approaches are then available while only the first is available for non symmetric games. For example, one may study the technology game with two populations of agents, namely if the two firms are not located in the same place and if the game can only confront firms not located in a same place, but also with only one population of agents if the firms are interchangeable.

A second distinction about meetings makes a difference between “global interaction” and “local interaction”. In a local interaction model, the agents which an agent may meet in a given period belong to subsets of the other population(s). These subsets figure agent’s local “interaction neighborhood”. For example, in some traditional evolutionist games, agents are located on a one or two dimensional space, like a circle or a torus, and the interaction neighborhood spontaneously includes agents who are physically near them. By contrast, in a global interaction model, each agent may meet any agent in the opponent population(s). Global interaction is just a special case of local interaction (the size of the subsets is the cardinality of the populations). For instance, the technology game may be played by firms situated on a circle and a firm meets only the firms at its right and left.

Large and global interactions assist anonymity. By anonymity, one means that the agents who meet do not know each other; the important point is that, given that they unlikely meet again, they are unable to develop strategies that require a long term interaction between them. It automatically follows that the stage game is really the game played and subject to learning over time. If anonymity were not respected, the played and learned game would become the repeated stage game, in which it is possible to build threats and reputation effects, which is not the aim of the evolutionist approach of a game. Of course, this does not mean that evolutionary game theory cannot address repeated games: if the purpose is to learn about a T period repeated game, then one has to suppose that the n agents of a same n -uplet meet exactly T times before one randomly draws new n -uplets of agents. It derives that local interaction models, to avoid the switch from the study of the stage game to the study of the repeated stage game, have to introduce assumptions, like limited memory or limited rationality, in order to prevent repetition effects. It also follows that not all games adapt to the evolutionist framework. For example, if the technology game is played by two oligopolistic firms, then there is no anonymity: both firms, called on to make a technology choice in each period, link future choices to past choices of the usually met opponent, hence construct strategies over a finite number of periods.

In case of large and global interactions (Kandori-Mailath -Rob, 1993; Axelrod, 1984), each strategy is confronted to many, sometimes all possible strategies of the opponents. It follows that each strategy provides a mean payoff, which may be rapidly learned by each agent if the informational context is rich enough. By contrast, in case of few interactions, or if each agent in each population is confronted to only one randomly drawn $(n-1)$ -uplet in the $(n-1)$ opponent populations (Robson and Vega Redondo,

1996), the randomness of the different meetings may have an impact on the evolution of the system. Namely, a strategy may diffuse only because the agents who played it, met (according to the random draw process) agents whose behavior made this strategy, fortunately or unfortunately, successful. For example in the first variant of the technology game, viewed as a mono- population game, the game, starting from a context with an equal number of agents playing A and of agents playing B, can switch to a state where everybody chooses technology A, as well as to a state where everybody chooses technology B. In fact, if every A playing agent meets a B playing agent, A performs better than B (1 is higher than 0), which may lead everybody to choose the A technology (by imitation of the best performing strategy). By contrast, if the A playing agents only meet A playing agents and if the B playing agents luckily only meet B playing agents, then B performs better than A (4 against 2) and everybody may choose the B technology.

3.2.3 The information principle

In classical economics, the players are endowed with much structural information and the structure of the game may even be common knowledge. In evolutionist game theory, the required structural information is much less constraining. A player usually (at least partially) knows his strategy set, but he may ignore his preferences over the possible issues. A player may ignore the characteristics of the other players (their strategy sets, their information, their preferences), the way he is confronted to them, and even that he is involved in a game. For example, in the technology game, each player may partly or completely ignore the costs and the efficiency of each technology given the choice of the other firm, and he may be unaware of the fact that the opponent has also to choose between the same two technologies.

In evolutionary games, factual information grows thanks to repetition. The repetition of the same stage game allows each agent to gather two types of information. On the one hand, he gathers information on the past actions played by his opponents (he is supposed to know his own past actions). On the other hand, he gets information on the utility provided by each of his past played action. He may also collect some information about the payoffs achieved by the opponents for different profiles of played strategies. Information is supposed to be non ambiguous, in the sense that one can categorize it (information on actions, information on utilities) and that it is one to one associated to a given player and to a given action. In-

formation is also more or less reliable; so the information may consist in an interval of values including the true one, sometimes in probability distributions over such an interval. In the technology game for example, repetition of the game may lead each firm to observe the technology chosen by the opponent and may provide, with more or less reliability, information on the payoffs achieved thanks to this choice.

Usually, the information which can be collected by a player over time is limited in space: an agent is namely unable to gather information outside of his “information neighborhood”. This neighborhood is generally included in the interaction neighborhood, but not necessarily. Even inside the information neighborhood, a player may only get information on a random sample of agents. The size of this sample may range from one agent to the cardinality of the neighborhood set. A player may even only learn (from an organization external to the game) aggregated information about the most played past action or statistics about the payoffs assigned to some actions. Finally information gathering is not systematic during time. In the technology game for example, a firm may deliberately choose to observe the opponent’s technology choice only in some periods, more or less spaced out, the interval between two information periods being regular or random.

Up to now, we only talked about passive information, that is to say information which automatically derives from the play of the game. But a player may also be more active in his search of information. On the one hand, he may buy the information collected by external specialized organizations, or get information by communicating with the other players. On the other hand, he may actively construct his information, for example by deliberately testing new strategies, in order to lead the players to a new trajectory of actions and payoffs, which provides him new original information. Such a behavior is frequent when the stage game is not well-known in the early periods of play. If so, exploration of new actions allows to investigate new parts of the game matrix (or of the game tree); for example, it allows to discover personal payoffs assigned to actions and reactions to actions, which have not been played up to now. In the technology game, a firm, even if satisfied with the older technology A, may test the new technology B, in order to discover if it really leads to the expected efficiency; a player may also discover, by switching to the new technology, that this switch leads the opponent to opt for a third technology, not available to himself, a possibility he ignored before the switch.

3.2.4 The evaluation principle

In a first step, a player may mainly focus on the distribution of past played strategies in his information neighborhood. Limited memory and limited ability of the player induces limits on the history of plays observed. An agent may forget any information which dates back more than k periods, and he may only focus on a sample of actions, possibly randomly drawn among the k last plays. On basis of this information, an agent can calculate the frequency with which each strategy has been played, as well as many other statistical properties of the sequence of past observed strategies, which may have an impact on his future play. More ambitiously, an agent can try to discover patterns of behaviors, like cycles of opponents' actions or types of reactions to actions. For example, in the technology game, a firm may be content with observing the technology chosen by its opponent in the three last periods in order to choose the most often selected technology; but a firm may also look for some regularities, like for example a systematic imitation by the opponent agent in each period of her technology choice in preceding periods.

In a similar way, a player may focus on the payoffs associated to different past strategies, played by himself and by the agents in his information neighborhood. Of course, limited memory and limited ability may again limit the sequence of observed past payoffs. The gathered information may be used to construct statistical indicators for each strategy, like the mean payoff or the weighted sum of payoffs provided, the weights (possibly) decreasing with earlier periods. He may also observe some structural patterns about payoffs such as their dispersion. In the technology game, a firm may study the sequence of payoffs assigned to both technologies, in order to adopt the most efficient one; but it may also focus on the dispersion of the obtained payoffs, in order to adopt the less risky technology.

In a second step, after having collected and calculated the above factual information, a player can try to infer some structural information about strategies and preferences and even beliefs. Observing that an opponent plays many different actions allows to deduce that he has a large strategy set, and leads to enrich this set each time a new strategy is observed. Observing that the opponent's actions exhibit some regularities may lead to infer that he follows a dynamic strategy which can be function of the past history of plays. Observing that the payoffs he obtained with a same action are highly dispersed, allows him to deduce that the opponents have large strategy sets in the stage game. Likewise, the collected information on the past actions of an opponent helps to infer information on his preferences -

provided that he plays in a rational way- and therefore to build (part of) the utility matrix of the game. But of course inferring opponents' preferences is often a complex, sometimes impossible task, because it requires strong assumptions (even stronger than in decision theory) about their rationality. Yet, fortunately, this is not always true. In the technology game, a firm may discover that he plays a coordination game, for example by observing both that the other firm always strives to choose in a period the technology he chose himself in the preceding period, and that his own payoffs are higher if both firms choose the same technology than in the reverse case.

Of course, the revelation and deduction of structural characteristics of the game require that the structure of the game is sufficiently stationary. Learning is possible only if the speed of learning (endogenous change) is higher than the speed of evolution of the learned characteristics of the game (exogenous change). If the context of play (stage game, interaction or information neighborhoods) changes too fast, agents have to content with observing this fact and the best they can do is to adopt an adapted way of behavior. For example, they may choose a cautious strategy, which leads to payoffs that do not much depend on the play of the other players and on the evolution of the context. By contrast, if the context of play changes in a slowly way, the players may learn about the way some structural variables change over time. For example, a player may be able to learn the evolution of his interaction neighborhood and the way it changes (endogenously by creation or destruction of links between players, exogenously by addition or deletion of a player), just by observing its evolution in the past, provided the evolution is not too fast and not too anarchical. In the technology game, a firm can easily observe that some other firms appear or disappear.

In a third step, each player may form anticipations on the future actions of his opponents, and these anticipations, as well as the way they are formed, are more or less sophisticated. A player may build anticipations in a simple way: for example, banking on a stationarity assumption, he just supposes that the probability with which the opponents play their different strategies in the future is the frequency with which they played them in the past. A more elaborated anticipation uses the regularities observed in the past behavior of the opponents to infer their potential strategies in the future. Even more sophisticated anticipations can be formed if an agent is partly informed about the beliefs and preferences of his opponents, in which case he may strive to predict their behavior (given an assumed level of rationality of the opponents). In the technology game for example, each firm may base on the frequency of the opponent's past choices of each

technology to extrapolate the opponent's future behavior. But he may also consider that the other firm has a dynamic strategy on which he formed a probability distribution in order to expect his next choice.

Likely, given his past payoffs, each player can define an "aspiration threshold", i. e. a normative indicator of the level of utility to reach (or to avoid). Such a threshold, which is not conditional to each action, is adjusted over time, according to the more or less easiness with which the agent (or agents in his neighborhood) has reached it in the past. In fact, the agent uses a dynamic "satisficing" model, but there is a (unique) threshold defined on the synthetic utility rather than (several) thresholds defined on partial objectives. The threshold may be adjusted at each period or only after a given number of periods (this number may evolve and depend on the size of the difference between the aspiration level and the observed payoffs). The threshold is lowered if the observed payoffs are durably lower than the threshold, it is risen in the reverse case. In the technology game for example, each firm can define aspiration levels it strives to reach or even to go beyond, albeit it has only two actions at disposal to achieve this aim.

3.2.5 The decision principle

A player draws on all the preceding calculated indices to make his choice. In classical game theory, it may be difficult to calculate the player's chosen strategy, for instance when the game has no or many equilibria. By contrast, in evolutionist game theory, an agent always plays, each time he is called on to do so, despite his possible lack of information and rationality. In fact, the way he plays just takes these failings into account, being sometimes very unsophisticated from a strategic point of view. The important point is that the mere fact of playing has two consequences. On the one hand, regardless of the strategic content of the actions, playing draws the game toward a particular direction depending on the played actions. On the other hand, playing diffuses information on the game through the played actions. Diffusion of information takes a strategic dimension: each agent may be conscious that the information conveyed by his actions can be used by other players to their advantage. Acquisition of information is favored in some respects. The exploitation behavior, taking advantage of the already obtained information is often followed and completed by an exploration behavior, which aims at providing new information.

Exploitation behavior may be more or less sophisticated. At one extreme, it may reduce to "inert behavior", which consists in repeating the

same action, or rather in keeping over time the same probability distribution over actions. At the other extreme, “optimizing behavior” consists in only choosing the best action, as regards the value of an index associated to each action. Inbetween, “probabilistic behavior” consists in playing each action with a probability proportional to the value of an index assigned to each action. In fact, the probabilistic behavior has two limit cases: the one is random behavior (which leads to the play of every action with the same probability), the other is optimising behavior (only the best action, as regards the value of the index, is played). Probabilistic and optimising behavior can be adjusted to all kinds of information available, namely information on actions and information on payoffs. For example, an imitation model, in which each action is played with a probability that is function of the frequency of its play in the past, is a probabilistic behavior model, based on an action frequency index. A reinforcement model, in which the probability of play of an action is function of its performance in the past, is also a probabilistic behavior model, based on a past observed payoff index. A best-reply model, in which an agent plays the action that maximizes his payoff given the past profile of opponents’ strategies, is an optimisation model, based on a calculated utility derived from a past action index.

The behaviors of the different players are not necessarily similar. One may suppose that some players behave in a very mechanical way, whereas others behave in a more strategic way and even internalize the behavior of the first ones. The behaviors are also more or less synchronized given the more or less similarity either of the choice rules or of the information of the agents. For instance, imitation behaviors create strong correlation among actions, whereas reinforcement behavior leaves more scope to personal stories. Likely, behaviors may adjust more quickly if the agents choose to diffuse all their information through their actions than if they choose to impede the diffusion of information.

More generally, the way of behavior may evolve over time, according to different setting off mechanisms. In one way, a player may switch from a passive behavior to a more sophisticated behavior when he gets more information and observes that he is locked-in a suboptimal situation. For example, a player may observe that his imitation behavior, very fruitful at the beginning of the game, becomes harmful after a given time, and therefore may turn to a more autonomous way of behavior. In the other way, a player may switch from an optimisation behavior to a more inert behavior, when he evolves from a stable interaction context to a turbulent context, in which the value of the utility index assigned to each action rapidly and strongly changes over time. But the transition, sometimes very rough, from

one way of behavior to another one, requires decisions at two levels of learning (action and way of behavior) and is still not much studied in the evolutionary literature.

Exploration behavior is generally more simple. In fact, exploration appears frequently as a perturbation of the exploitation process and is implemented in specific circumstances. More precisely, a player may introduce randomly some strategy or a new agent-strategy may be randomly introduced in the population of agents. In fact, the selected strategy (or agent-strategy) may be randomly chosen in the initial strategy set or even in an extended strategy set. In the last case, the original strategy is designed in the neighborhood of an already used strategy (mutation) or by mixing between already existing strategies (crossing over). However, a strategy may be selected in a more oriented manner, along a probability distribution on the initial strategy set. If an original strategy is introduced, it is around a strategy having worked well in the past. The introduction of a new strategy may be infinitesimal or more important, on one period or on successive periods.

Exploration behavior is also characterized by the way it is correlated among the players. In particular, a binomial model (in which the mutation probabilities of the agents are independent of time and independent among the agents) can give rise to situations where any number of agents may explore in each period. By contrast, other models ensure exploration by a constant fraction of each population in each period. Finally, a common exploration behavior may follow an event exogenous to the game. In chess for example, the development of the computer capacities allowed the players to simultaneously explore many strategies they had not exploited before.

In fact, the exploitation and exploration behaviors are frequently mixed into a unique behavior. For instance, the probabilistic model assigns strictly positive probabilities to all actions (which ensures the exploration behavior), but contrasted probabilities to these actions (for example higher probabilities to best-reply actions, which expresses the exploitation behavior). Likely, a player may play the same strategy as long as his result is above a utility threshold (which may itself change by mutation) and shift towards another behavior if his utility falls under the threshold.

A trade-off between exploitation and exploration is implicitly achieved by any way of behavior. Moreover, this trade-off may naturally vary with time. For example, a player may intensively explore in the early periods by trying many different strategies, and thereafter assign more and more weight to exploitation behavior, once a whole set of consequences have been observed. This is exactly what happens with the probabilistic model

where the probabilities of the actions become more and more concentrated (but is never zero for any strategy) and may even focus asymptotically on a unique strategy. In the technology game for example, after a procedure of trial and error at the beginning of the game, in order to discover the reaction of the opponent and possibly the nature of the game, the firm may stick to the current best technology (according to a given utility index), and reduce exploration to a random draw of technologies from time to time, in order to check if the usual chosen technology is still the best one.

3.3 Some models

Technically, an “evolutionist game” is defined by a stage game and a dynamic process. The stage game is stated in normal or extensive form, and only defines the utilities associated with the possible combinations of actions. The dynamic process specifies the principles of confrontation, information, evaluation and decision which are at work in the repetition of the stage game. It is thus by specifying the dynamics that the modeler can introduce the actual processes she thinks are pertinent, dealing with best responses, learning, selection, imitation, reinforcement, and so on.

Two kinds of contributions characterize evolutionist game theory. On one hand, the theory tried to find descriptions of the temporal behavior of the system that do not depend too much upon the detail of the considered dynamic process. Such results have been obtained for some classes of “simple” dynamic processes that only incorporate a selection or a reinforcement mechanism, coupled with a mutation or exploration mechanism of low intensity. On the other hand, evolutionist game theory has been able to provide new insights on some specific questions like the diffusion of technologies or the genesis of institutions. Several monographs provide surveys of the technical achievements of evolutionist game theory, for instance Maynard Smith (1982), Weibull (1995), Skyrms (1996), Vega-Redondo (1996), Samuelson (1997), Young (1998) or Fudenberg and Levine (1998).

The models and interpretations of evolutionist game theory fluctuate between two poles: population dynamics and individual learning. In population dynamics, there is a large number of individuals, each of whom being endowed with a strategy he does not control. This strategy remains constant, except in case of mutations, which are random and relatively rare. Within one population, the variations of the proportion of users of a strategy are led by the selection process, which favors reproduction of the best fitted individuals. By contrast, the theory of individual learning makes the hypothesis

that each individual can vary his strategy with time. This variation is stochastic when it incorporates a certain degree of active experimentation. If it is not supposed that the number of individuals is large, one nevertheless supposes that the individuals endlessly apply the same learning rules.

Not without generating some confusion, it is most remarkable that similar formal models appeared in such different fields as theoretical biology and cognitive psychology. Even if its interpretations are completely distinct, the theory of individual learning, including mimetic phenomena, turns out to make use of similar models and results as population dynamics. The fundamental idea -- which needs to be made more explicit -- is to interpret the probabilities that appear in “mixed” strategies not as proportions of individuals using the various pure strategies, but as propensities, for a single individual, to use these pure strategies. Moreover, it is possible to build hybrid models, with populations of learning individuals. Finally, even if the microeconomist is mainly interested by how interactive individuals learn, it cannot be excluded that population dynamics reveals useful to understand some specific social and economic problems.

The exposition of the models of evolutionist game theory will follow this distinction. The first part presents the models of population dynamics. It first presents the notion of an evolutionary stable state, before presenting different evolutionary dynamics and their asymptotic behavior. The second part deals with the models of individual learning, when considering interacting individuals. We distinguish two families of models: learning by reinforcement and exploration. The third part treats of the link between reinforcement and exploration. First, we examine the different interpretations of the formalism, then we present a hybrid model. But before all, some general definitions need to be formally introduced.

3.3.1 Definitions

All concepts will be introduced with the formalism of two-player normal form games. Let X_1 be the set of strategies available to player 1 and, likewise, let X_2 be player 2 strategy set (elements of these sets are “pure” strategies). For $x_1 \in X_1$ and $x_2 \in X_2$, the utilities associated with the strategy profile (x_1, x_2) are denoted $u_1(x_1, x_2)$ for player 1 and $u_2(x_1, x_2)$ for player 2. A game is said to be *symmetric* if both players play the same role in that game, that is to say they have the same set of available strategies ($X_1 = X_2 = X$) and the payoffs are such that $u_2(x_1, x_2) = u_1(x_2, x_1)$. The utility of a player choosing x when the other player chooses y is then simply denoted $u(x, y)$, thus $u(x, y) = u_1(x, y) = u_2(y, x)$.

In the technology game taken as an example, $X_1 = \{A1, B1\}$, $X_2 = \{A2, B2\}$, and the utilities are given by Table 3.1:

$$\begin{aligned} u_1(A1, A2) &= 2, \quad u_1(A1, B2) = b, \\ u_1(B1, A2) &= c, \quad u_1(B1, B2) = 4 \end{aligned} \quad (3.1)$$

and likewise for player 2. The technology game can be seen as a symmetric game with $X = \{A, B\}$.

If p_1 and p_2 are two probability distributions over X_1 and X_2 respectively (such probability distributions are called “mixed strategies”), one also denotes by u_1 the expected utility:

$$\begin{aligned} u_1(p_1, x_2) &= \sum_{x_1 \in X_1} u_1(x_1, x_2) p_1(x_1) \\ u_1(p_1, p_2) &= \sum_{x_1 \in X_1} \sum_{x_2 \in X_2} u_1(x_1, x_2) p_1(x_1) p_2(x_2) \end{aligned} \quad (3.2)$$

and likewise for u_2 . Note that the preceding definition leaves open the interpretation of the probabilities. They can be proportions of users in a population, individual propensity to use pure strategies, or temporal frequencies of the occurrence of pure strategies.

For non-cooperative game theory, the central concept is *Nash equilibrium*. For a two-player game, a Nash equilibrium is a couple of strategies such that the strategy chosen by each player is a best response to the (equilibrium) strategy chosen by the other player. This definition makes sense with pure and mixed strategies; in the case of mixed strategies, a couple (p_1^*, p_2^*) is a Nash equilibrium if for any deviations $p_1 \neq p_1^*$ and $p_2 \neq p_2^*$, one has: $u_1(p_1^*, p_2^*) \geq u_1(p_1, p_2^*)$ and $u_2(p_1^*, p_2^*) \geq u_2(p_1^*, p_2)$. A Nash equilibrium is stable with respect to unilateral deviations of the players. The equilibrium is *strict* if the preceding inequalities are strict; in that case, each player is choosing a *unique* best reply to the other player’s choice. With a finite number of pure strategies, there always exists at least one Nash equilibrium in mixed strategies, but there does not always exist a strict Nash equilibrium. When a strict equilibrium exists, it only involves pure strategies.

In the first variant of the game of technologies, with $b = 1$ and $c = 0$, the reader will easily check that there are three Nash equilibria. Two of them are pure and strict: $(A1, A2)$ and $(B1, B2)$. The first equilibrium provides utility 2 to each player and the second one provides utility 4 to each player. In the third equilibrium, both player choose the mixed strategy $p_* = (3/5, 2/5)$. In that equilibrium, the utility to each player is:

$$(9/25) \times 2 + (6/25) \times 1 + (4/25) \times 4 = 8/5 \quad (3.3)$$

3.3.2 Population dynamics

3.3.2.1 Evolutionary stability

The notion of evolutionary stability was first defined, without reference to an explicit dynamic process, and in the context of symmetric games. One considers a population of individuals that choose pure strategies in X . If p is a mixed strategy, for any pure strategy $x \in X$, the number $p(x)$ can be interpreted as the proportion in the population of individuals that use the strategy x . The expected utility $u(x, p) = \sum_{y \in X} u(x, y)p(y)$ is the utility obtained on average by the use of x , in uniform interactions with a population characterized by p . Let p and q be two mixed strategies, and consider the mix of two sub-populations respectively characterized by p and q , in proportions $(1 - \varepsilon)$ and ε . In the whole population, individuals of the first sub-population get on average the utility $u(p, (1 - \varepsilon)p + \varepsilon q)$ and the other individuals the utility $u(q, (1 - \varepsilon)p + \varepsilon q)$. Population dynamics will make the assumption that, in such a situation, the ratio $(1 - \varepsilon)/\varepsilon$ will be modified in favor of the sub-population that gets, on average, the largest utility (see next section).

By definition, a mixed strategy p is evolutionary stable if, for any other mixed strategy $q \neq p$, there exists an “invasion barrier” $\varepsilon > 0$ such that:

$$\forall \varepsilon \in]0, \varepsilon[\quad u(p, (1 - \varepsilon)p + \varepsilon q) > u(q, (1 - \varepsilon)p + \varepsilon q) \tag{3.4}$$

This means that a population that plays the mixed strategy p cannot be “invaded” by any deviant small sub-population. Indeed, if such a small sub-population were to appear, its size would immediately decrease even further.

The notion of an evolutionary stable strategy can be compared to the notion of a Nash equilibrium. *An evolutionary stable strategy p necessarily forms a symmetric Nash equilibrium (p, p) .* In fact, if (p, p) were not such an equilibrium, there would exist a response q such that $u(q, p) > u(p, p)$; the continuity of u (with respect to its second variable) shows that, for ε small enough, one would also get: $u(q, (1 - \varepsilon)p + \varepsilon q) > u(p, (1 - \varepsilon)p + \varepsilon q)$. Conversely, suppose that p is evolutionary stable, and let $q \neq p$. If $u(q, p) < u(p, p)$, then q does not invade p . If $u(q, p) = u(p, p)$, then the inequality: $u(p, (1 - \varepsilon)p + \varepsilon q) > u(q, (1 - \varepsilon)p + \varepsilon q)$, appearing in the definition of evolutionary stability, simply becomes: $u(p, q) > u(q, q)$. That is to say that, for q to invade p , it must be the case either that q performs strictly better than p against p (first order condition), or that q performs

as well as p against p , but strictly better than p against q (second order condition). Thus, *any strict Nash equilibrium is evolutionary stable*, but a non-strict Nash equilibrium may also be evolutionary stable.

In the technology game, one can compute the invasion barrier that protects strategy A . In a population that contains a fraction ε of B -users, utilities for A and B respectively are:

$$\begin{aligned} u(A, (1-\varepsilon)A + \varepsilon B) &= 2(1-\varepsilon) + 1\varepsilon, \\ u(B, (1-\varepsilon)A + \varepsilon B) &= 0(1-\varepsilon) + 4\varepsilon, \end{aligned} \quad (3.5)$$

therefore mutation B invades population A if:

$$2(1-\varepsilon) + \varepsilon < 4\varepsilon, \quad (3.6)$$

that is:

$$\varepsilon > 2/5 \quad (3.7)$$

The “invasion barrier” is $2/5$. Likewise, the “invasion barrier” that protects B from A is $3/5$.

In the technology game, the mixed equilibrium p_* is not evolutionary-stable; it is indeed invaded by A as well as by B : For the invasion by A ,

$$\begin{aligned} u(p_*, (1-\varepsilon)p_* + \varepsilon A) &= (8/5)(1-\varepsilon) + (6/5)\varepsilon, \\ u(A, (1-\varepsilon)p_* + \varepsilon A) &= (8/5)(1-\varepsilon) + 2\varepsilon, \\ u(A, (1-\varepsilon)p_* + \varepsilon A) &> u(p_*, (1-\varepsilon)p_* + \varepsilon A). \end{aligned} \quad (3.8)$$

And likewise for the invasion by B ,

$$\begin{aligned} u(p_*, (1-\varepsilon)p_* + \varepsilon B) &= (8/5)(1-\varepsilon) + (11/5)\varepsilon, \\ u(B, (1-\varepsilon)p_* + \varepsilon B) &= (8/5)(1-\varepsilon) + 4\varepsilon, \\ u(B, (1-\varepsilon)p_* + \varepsilon B) &> u(p_*, (1-\varepsilon)p_* + \varepsilon B). \end{aligned} \quad (3.9)$$

The following variant of the technology game (variant 2) shows that evolutionary stability is more demanding than Nash equilibrium, even if one restrict attention to pure equilibria. Let $X = \{A, B\}$, and:

$$\begin{aligned} u(A, A) &= 2, u(A, B) = 2, \\ u(B, A) &= 2, u(B, B) = 4. \end{aligned} \quad (3.10)$$

The corresponding matrix is the following:

Table 3.3. Technology game, variant 2

	A	B
A	(2,2)	(2,2)
B	(2,2)	(4,4)

This game correspond to the variant $b = c = 2$ and can again be interpreted as a competition between two technologies. Technology A is a classical communication technology (ordinary cellular phone), which is compatible with the other. Technology B is more powerful, but only matched with itself (UMTS technology). This game has two symmetric Nash equilibria: (A, A) and (B, B) , but the equilibrium (B, B) is evolutionary stable whereas the equilibrium (A, A) is not. In effect, the strategy A defines a Nash equilibrium because $u(A, A) \geq u(B, A)$ (this equilibrium is non-strict). Likewise, B defines a Nash equilibrium because $u(B, B) \geq u(A, B)$ (this equilibrium is strict). Moreover, B is evolutionary stable because $u(B, B) > u(A, B)$ (the first order condition is satisfied). But A is not evolutionary stable because $u(A, A) = u(B, A)$ and $u(A, B) < u(B, B)$ (the second order condition is not satisfied).

The following variant of the technology game shows that evolutionary stability does not require, in the framework of symmetric games, the strict equilibrium condition. Let $X = \{A, B\}$, and $b = 4, c = 3$:

$$\begin{aligned}
 u(A, A) = 2, u(A, B) = 4, \\
 u(B, A) = 3, u(B, B) = 4.
 \end{aligned}
 \tag{3.11}$$

The corresponding matrix is the following:

Table 3.4. Technology game, variant 3

	A	B
A	(2,2)	(4,3)
B	(3,4)	(4,4)

One can again think of technologies A and B as two communication technologies. B is more costly but more efficient, and when A and B are matched, both users benefit from B , but the A -user is better off than the B -user. This game has a unique symmetric Nash equilibrium (B, B) . As the

reader can check, this equilibrium is non-strict and is evolutionary stable; the reason is that, even if A does as well as B against B (notice that $u(B, B) = u(A, B)$), B does strictly better than A against A (notice that $u(B, A) > u(A, A)$).

The definition of evolutionary stability in the case of a non-symmetric game is not so easy, even if a generalisation of the preceding definition seems natural. One now considers that the two players are two populations of individuals. If p_1 is a mixed strategy, for each pure strategy x_1 , $p(x_1)$ is interpreted as the proportion of individuals in population 1 who use x_1 . The expected utilities are average utilities in each population. Consideration of such averages is intuitively justified by imagining a random matching process between the individuals of the two populations.

By definition, a couple of mixed strategies (p_1^*, p_2^*) is a *Two-population Evolutionary Stable* state if, for any mixed strategy $p_1 \neq p_1^*$, for player 1, there exists an invasion barrier $\varepsilon_1 > 0$ such that:

$$\forall \varepsilon \in]0, \varepsilon_1[, u_1((1-\varepsilon)p_1^* + \varepsilon p_1, p_2^*) < u_1(p_1^*, p_2^*), \quad (3.12)$$

and similarly, for any mixed strategy $p_2 \neq p_2^*$ for player 2, there exists $\varepsilon_2 > 0$ such that:

$$\forall \varepsilon \in]0, \varepsilon_2[, u_2(p_1^*, (1-\varepsilon)p_2^* + \varepsilon p_2) < u_2(p_1^*, p_2^*). \quad (3.13)$$

Here again two populations that are at a two-population evolutionary stable state cannot be invaded by small mutant sub-populations.

One then gets, in the two-population case, the following link with the Nash equilibrium concept. A couple of (mixed) strategy is a two-population evolutionary stable state if and only if it is a strict Nash equilibrium. To prove this result, consider a two-population evolutionary stable state (p_1^*, p_2^*) and a deviation p_1 for the first player. There exists $\varepsilon > 0$ such that:

$$u_1((1-\varepsilon)p_1^* + \varepsilon p_1, p_2^*) < u_1(p_1^*, p_2^*). \quad (3.14)$$

By the linearity of expectation, one gets that $u_1(p_1, p_2^*) < u_1(p_1^*, p_2^*)$. The same reasoning for the other player proves that (p_1^*, p_2^*) is a strict equilibrium. Conversely, suppose that (p_1^*, p_2^*) is a strict equilibrium. For any deviation $p_1 \neq p_1^*$ for player 1, $u_1(p_1, p_2^*) < u_1(p_1^*, p_2^*)$. By linearity, for any $\varepsilon \in]0, 1[$, $u_1((1-\varepsilon)p_1^* + \varepsilon p_1, p_2^*) < u_1(p_1^*, p_2^*)$; therefore, the evolutionary stability criteria is satisfied for player 1. Such is also the case for player 2, hence the result.

But these elementary considerations may sometimes hide some subtleties as one can check by considering again the example 3, treated this time as a two-population game (example 3bis). Consider the two-player game defined by the bi-matrix (whose payoffs happen to be symmetric):

Table 3.5. Technology game, variant 3bis

	A2	B2
A1	(2,2)	(4,3)
B1	(3,4)	(4,4)

The game is the same as before, once strategies $A1$ and $A2$ are identified with A , and $B1$ and $B2$ with B . It was seen that, according to the first definition of evolutionary stability, that is for a one-population model, (B, B) is an evolutionary stable state. But the equilibrium $(B1, B2)$ of the two-player game is not strict, thus $(B1, B2)$ is not evolutionary stable according to the second definition (two-population evolutionary stability). The reason here is clear. In the two-population model, if all individuals in one population play $B1$, the individuals in the other population can play $B2$ as well as $A2$. The fact that the payoff for A is lower against A and larger against B is only destabilizing in the one-population model.

3.3.2.2 Replicator dynamics

An evolutionary dynamics partially specifies the intuitions that were evoked by the preceding definitions of evolutionary stability. The use of a strategy tends to increase if this strategy performed relatively well in the past. The canonical example here is the “replicator dynamics”, which appears as a central concept in many selection processes. There exist several variants of that dynamics, defined at the aggregate level, leaving unspecified the matching process between individuals. The common principle is that the proportion of individuals who use a given strategy changes over time; the speed of this change being proportional to the difference between the average utility obtained by the users of this strategy and the average utility in the whole population, at the current date. Replicator dynamics is a model of how the best performing strategies within a population are selected. Whenever a strategy is not present at all in the population, no mutation can make it appear in the future.

Consider first a single population, and thus a symmetric game. Let p^t be the vector that describes, at date t , the proportion, among the population, of users of the different pure strategies. The expected utility to the pure strategy x is $u(x, p^t)$, interpreted as the average payoff to the users of strategy x when they are randomly matched with any individual. The average utility in the population is $u(p^t, p^t)$. One supposes that, in continuous time, the speed of variation $\dot{p}^t(x)$ is exactly proportional to $u(x, p^t) - u(p^t, p^t)$. One then obtains the one-population replication equation:

$$\dot{p}^t(x) = \nu p^t(x) [u(x, p^t) - u(p^t, p^t)] \quad (3.15)$$

with ν a speed coefficient, which can depend upon time, but is the same for the different strategies (one usually takes $\nu=1$). This equation has a dominant interpretation in terms of biological selection, but can also be interpreted in terms of social imitation, as is proven in the third sub-section.

In the two-population case, several generalisations of the replication equation make sense. One can for instance propose the “standard two-population replicator dynamics”:

$$\begin{aligned} \dot{(p_1^t(x_1) = p_1^t(x_1)[u_1(x_1, p_2^t) - u_1(p_1^t, p_2^t)]} \\ \dot{(p_2^t(x_2) = p_2^t(x_2)[u_2(p_1^t, x_2) - u_2(p_1^t, p_2^t)]} \end{aligned} \quad (3.16)$$

Figure (3.1) shows the trajectories of the standard two-population replicator dynamics for the game of technologies (variant 1). The horizontal variable x is the frequency of the strategy $A1$ for the first player, and the vertical variable y is the frequency of $A2$ for the second player. The differential equations (3.16) writes here:

$$\begin{aligned} \dot{x} &= x(1-x)(5y-3) \\ \dot{y} &= y(1-y)(5x-3) \end{aligned} \quad (3.17)$$

One can see that the two pure equilibria (corresponding to the corners $x = y = 0$ and $x = y = 1$ of the picture) are stable whereas the mixed strategy equilibrium (corresponding to the point $x = y = 3/5$ inside the picture) is not.




Fig. 3.1. The technology game (variant 1): Standard two-population replicator dynamics.

The evolutionary dynamics share the following property. If a strategic configuration is a Nash equilibrium, then the state of the system corresponding to this configuration is a rest-point of the selection process (the system stays at that point if it is initially there). Indeed, at a Nash equilibrium, all pure strategies which are used with some positive probability provide exactly the same utility, equal to the average utility. As a consequence, the reproduction dynamics based on the utility difference with that average do not modify the frequency of such pure strategies. Without mutations, unused pure strategies remain so. A Nash equilibrium is thus a rest-point. If the evolutionary process moreover embodies a mutation process, this property cannot be guaranteed in general. Apart from Nash equilibria, other configurations can be rest points.

Nevertheless, if all Nash equilibria are rest points, it is not the case in general that the evolutionary system asymptotically converges towards such a state. A state is said to be asymptotically stable if the state of the system converges towards it as soon as the system is initially close enough to that state. The conditions for a Nash equilibrium to be *asymptotically stable* depend on the detail of the dynamics, but they are most often very tight, at least as tight as the conditions for evolutionary stability. Mixed equilibria often fail asymptotic stability. A typical result is the following. *A couple of strategies is asymptotically stable for the standard bi-population replicator dynamics if and only if it is a two-population evolutionary stable equilibrium, that is a strict Nash equilibrium.*

Another familiar concept in game theory is the concept of dominated strategy. The elimination of strictly dominated strategies is a feature common to most models of evolution by selection. In particular, *in the standard replicator dynamics, any strictly dominated strategy is eliminated.* This result can be very easily proved in the two-population case. Let x_1 and y_1 be two strategies for player 1, which are played with some positive probability at the initial date t_0 , and such that x_1 strictly dominates y_1 . Denote:

$$\delta = \min_{x_2 \in X_2} [u_1(x_1, x_2) - u_1(y_1, x_2)] \quad (3.18)$$

then $\delta > 0$ and:

$$\frac{d}{dt} \log \frac{p'_1(x_1)}{p'_1(y_1)} = u_1(x_1, p'_2) - u_1(y_1, p'_2) \geq \delta \quad (3.19)$$

It follows that $p'_1(y_1)$ tends exponentially towards 0, meaning that the dominated strategy y_1 is quickly eliminated.

The question of the elimination of weakly dominated strategies is much more delicate. This question is tackled using an equilibrium concept which is more restrictive than Nash equilibrium: the (subgame) perfect equilibrium. By definition, a perfect equilibrium of an extensive-form, two-player finite game is a pair of strategies (one for each player) that defines a Nash equilibrium in the considered game and all its sub-games. Such an equilibrium can be obtained by the backward induction process, starting from the terminal nodes, each player choosing its best action given what is to happen further in time.

Consider the classical deterrence game proposed by Selten and called chain store paradox (example 4). There are two players: the challenger C

and the defender D . The challenger moves first by deciding to attack or not. In case of an attack, the defender answers by either yielding or counter-attacking. An illustration for this game is a market on which a monopoly D operates. Another firm C can challenge the monopoly by entering the market. If the challenger enters the market, the monopoly can either yield by sharing the market, or set up a price war. The normal form of the game is in Table 3.6.

Table 3.6. Dissuasion game

	D yields	D counter-attacks
C does not attack	(0,2)	(0,2)
C attacks	(1,1)	(-1, -1)

The game has two Nash equilibria: “ C does not attack and D counter-attacks in case of an attack”, and “ C attacks and D yields to an attack”. Call “dissuasion” the first equilibrium and “splitting” the second. In the normal form, one can see that the strategy “counter-attack” for D is weakly dominated by the strategy “yield”, which corresponds to the fact that counter-attacking is either trivially equivalent to yielding (when D does not have to play), either strictly worse (if C attacks). The two Nash equilibria define two courses of play in the game. According to the dissuasion equilibrium, C does not attack and, according to the splitting equilibrium, C attacks while D yields. But the course of the game for dissuasion is not reasonable in the sense that the counter-attack threat is not credible. In effect, in the subgame induced by C ’s attack, there is only one player, D , left to move, and this player does not maximize its utility. In fact, only the splitting equilibrium is subgame perfect, dissuasion is not.

The standard argument in favor of the elimination of those equilibria which fail subgame perfectness is typical of the way of reasoning of classical game theory, because it rests on the players’ rationality: C attacks because she knows that D will prefer to yield when C will have attacked. In the evolutionary framework, one conceives of a player as a population. In front of a fixed population of challengers whose larger part is aggressive, selection pressure tends to eliminate the counter-attack strategy since, in such an environment, yielding is the best strategy. But unless the counter-attack strategy is not totally eliminated, there remains a selective pressure on the challengers in favor of non attacking. The question of the

speed of adjustment then becomes crucial. For some reasonable dynamics (see Samuelson, 1997), it is in fact possible to obtain at the limit of the selection process a situation such that: (i) the challengers do not attack; (ii) certain defenders nevertheless continue to carry in their genes the counter-attack strategy. If, for one reason or another, the situation were to change, the counter-attack behavior could again be observed.

3.3.3 Individual learning

3.3.3.1 Reinforcement learning

In *reinforcement learning*, an individual has a propensity to use strategy x which increases with the utility obtained when using x in the past. This first type of learning is purely behavioral and does not require that the individual be conscious of playing a game repeatedly. Some of these models are indeed pure models of individual decision-making, applied to the interactive case. Reinforcement learning supposes that each player builds an index of the past performance of each strategy and chooses to favor her best-performing strategies. The basic model here is the *Cumulative Proportional Reinforcement rule* (“CPR” rule). Introduced by Bush and Mosteller (1955), this rule is used under different variants to fit experimental data by Roth and Erev (1995), and it is studied from the theoretical point of view by Laslier, Topol et Walliser (2001). By definition, the CPR rule consists in choosing strategy x with a probability that is proportional to the past cumulated utility that x provided in the past to the individual. One can see that this rule realizes a certain compromise between exploration of the different strategies and exploitation of the best ones since, on one side, any pure strategy is chosen again at each date with some positive probability, even if this strategy was not very successful in the past (exploration), and on the other side, successful strategies are reinforced and used with increasing frequencies (exploitation).

Let $x_1 \in X_1$ be a strategy for player 1. Denote by $p_1^t(x_1)$ the probability that 1 chooses x_1 at date t . One has:

$$p_1^t(x_1) = \frac{CU_1^t(x_1)}{\sum_{y_1 \in X_1} CU_1^t(y_1)} \quad (3.20)$$

where $CU_1^t(x_1)$ represents the total (cumulated) utility provided by x_1 up to date $t-1$. In a two-player game, the evolution of this quantity is thus as follows: if, at date t , x_1 is played by 1 and x_2 is played by 2, then 1

earns $u_1(x_1, x_2)$ and thus $CU_1^{t+1}(x_1) = CU_1^t(x_1) + u_1(x_1, x_2)$; if 1 chooses some other strategy than x_1 , then $CU_1^{t+1}(x_1) = CU_1^t(x_1)$. Finally one can write:

$$CU_1^{t+1}(x_1) = CU_1^t(x_1) + u_1(x_1, x_2)\varepsilon(t, x_1, x_2) \tag{3.21}$$

with $\varepsilon(t, x_1, x_2)$ being 1 with probability $p_1^t(x_1)p_2^t(x_2)$ and 0 with the complement probability.

The preceding equations define a *discrete time stochastic process*. Take as state variable, the past frequency $f^t(x_1, x_2)$ of the couples of strategies (x_1, x_2) up to date t . Going from t to $t + 1$, this frequency goes, in the case where (x_1, x_2) is not played, from $f^t(x_1, x_2)$ to $f^t(x_1, x_2)\frac{t}{t+1}$; and it goes to $\frac{1}{t+1} + f^t(x_1, x_2)\frac{t}{t+1}$ in the opposite case, which happens with probability $p_1^t(x_1)p_2^t(x_2)$. The expected value of $f^{t+1}(x_1, x_2)$ given f^t can thus be written as a function of f^t and t , which justifies the choice of f^t as state variable. The discrete time deterministic process associated with the considered stochastic process is then obtained by identifying in an iterative way f^{t+1} with its expected value given f^t . After some manipulations, one gets:

$$f^{t+1}(x_1, x_2) - f^t(x_1, x_2) = \frac{1}{t+1} [p_1^t(x_1)p_2^t(x_2) - f^t(x_1, x_2)] \tag{3.22}$$

with $p_1^t(x_1)$ and $p_2^t(x_2)$ which can be written in (3.20) as functions of f^t since, for instance,

$$CU_1^t(x_1) = t \sum_{y_2 \in X_2} u_1(x_1, y_2) f^t(x_1, y_2) \tag{3.23}$$

The results obtained about reinforcement learning models of the CPR type are even more modest than those obtained for the replicator dynamics, due to the stochastic nature of the process. Nevertheless, one can show that the *rest points of the deterministic dynamics are Nash equilibria of the game restricted to those pure strategies that are actually used with some probability*. Deeper, one can show that the stochastic process converges with positive probability towards any strict Nash equilibrium. For certain types of mixed strategy Nash equilibria, one can show that the process converges with zero probability towards them. Finally, strictly dominated strategies are eliminated in the deterministic cumulated reinforcement process.

The Prisoner's dilemma (see table 3.7) is the prototype of a game in which both players not using dominated strategies lead to an inefficient outcome. It is represented by the following bi-matrix:

Table 3.7. Prisoner's dilemma

	A2	B2
A1	(2,2)	(5,1)
B1	(1,5)	(4,4)

The strategies $A1$ for player 1, and $A2$ for player 2 are “defection”, $B1$ and $B2$ being “cooperation”. For each player, defection strictly dominates cooperation, but the situation in which both players cooperate is strictly better for both players to the situation in which both defect. As a possible economic interpretation, consider two firms that can engage in advertising (strategy A) or not (strategy B). Whatever the other does, it is better for a firm to advertise, either to overtake the opponent (if the opponent does not advertise) or keep up with her (if she advertises). Nevertheless, if both advertise, they hinder each other, and they would be better off if neither were to advertise. This game belongs to the family of “technological games” we already used as examples (for $b=5$ and $c=1$). In such a game, reinforcement individual learning, just like replicator dynamics, does not make cooperation emerge.

3.3.3.2 Extrapolation learning

In *extrapolation* learning, an individual has a propensity to use strategy x which increases with the efficiency of this action in front of the past actions of the other players. This type of learning requires cognitive capacities which are higher than those required for reinforcement learning, because the individual now is supposed to estimate anticipated utilities. These capacities remain smaller than what is supposed in the *rational learning* models (Kalai and Lehrer, 1993), which are not studied here. Learning through extrapolation requires that each player observes the past moves of her opponent, modify accordingly her anticipation of her opponent's future strategy, and chooses a best responds to that anticipated strategy. The prototypical model is Fictitious Play (FP). It was proposed by Robinson (1951) as an algorithm for finding the equilibria of a game, and

it has been the object of many theoretical and empirical studies. According to this model, each individual simply chooses a best response to the past statistical frequency of her opponent's strategies.

About this model, asymptotic theoretical results are not clearcut. In fact one can be interested both in the convergence of players' strategy choices, and in the convergence of the players' beliefs. It was early recognized that the process generated by FP can lead to cycling in some games. Some games have the property that the generated sequence of statistical distributions converges to a (mixed strategy) Nash equilibrium, meaning that the players' beliefs are in the long run close to an equilibrium, even if the chosen strategies are not. But this does not hold for all games (Monderer and Shapley, 1966, 1996). The difficulties here are due to the deterministic character of the process, which forbids any deviation. One is thus led to consider perturbed versions of Fictitious Play. A simple version consists in handling individual choice through a logit choice rule, rather than an exact maximization. A more elaborate version will be presented in the next section.

3.3.4 Hybrid models

3.3.4.1 *The interpretation of replicator dynamics*

The original interpretation of the replicator dynamics equation is linked with *biological reproduction*. Let $n^t(x)$ be the number of individuals that use strategy x at date t and let thus $n^t = \sum_{x \in X} n^t(x)$ be the total number of individuals in the population. Suppose that, at date t , one individual i , using the strategy x , is chosen at random in the first population to reproduce (asexual reproduction). This individual is randomly matched with another individual of the other population. Denote by y the strategy used by this other individual. One supposes that the number of offsprings of individual i , in the next generation, is proportional to the utility $u(x, y)$ of her strategy in this matching. Utility is here directly the measure of reproductive "fitness". Under these hypotheses, the probability of a match of type (x, y) is $p^t(x)p^t(y)$ and the number $n^{t+1}(x)$ of x -users at date $t+1$ is equal to $n^t(x)$ with the probability $1 - p^t(x)$, and to $n^t(x) + \lambda u(x, y)$ with the probability $p^t(x)p^t(y)$, where λ is some proportionality constant. On average, one thus has:

$$n^{t+1}(x) - n^t(x) = \lambda p^t(x) \sum_{y \in X} p^t(y) u(x, y) = \lambda p^t(x) u(x, p^t). \quad (3.24)$$

The evolution of the total size of the population is:

$$n^{t+1} - n^t = \lambda \sum_{x \in X} p^t(x) u(x, p^t) = \lambda u(p^t, p^t) \quad (3.25)$$

In terms of the proportion of x -users, one gets:

$$p^{t+1}(x) = \frac{n^{t+1}(x)}{n^{t+1}} = \frac{n^t(x) + \lambda p^t(x) u(x, p^t)}{n^t + \lambda u(p^t, p^t)} \quad (3.26)$$

If the population is large:

$$\begin{aligned} p^{t+1}(x) &\approx \frac{n^t(x)}{n^t} \left(1 + \frac{\lambda}{n^t(x)} p^t(x) u(x, p^t) - \frac{\lambda}{n^t} u(p^t, p^t) \right) \\ &= p^t(x) + p^t(x) \frac{\lambda}{n^t} [u(x, p^t) - u(p^t, p^t)] \end{aligned} \quad (3.27)$$

One can find here, going to continuous time and for $v = \frac{\lambda}{n^t}$, the replicator equation, according to which the variation in the proportion of x -users is proportional to the difference between the utility for x and the average utility.

A second interpretation of the replicator is in terms of *social mimetism*. Let again $n^t(x)$ be the number of individuals which use, at date t , the strategy x , in a population whose size is now constant $n^t = n$ and large.

One still denotes by $p^t(x) = \frac{n^t(x)}{n^t}$ the proportion of x -users. One supposes that the individuals meet frequently so that, during period t , using strategy x provides the average level of utility $u(x, p^t)$. Suppose now that at date t one individual, say i , is chosen at random and is given the possibility to change her strategy. This individual gets to know the level of utility obtained by another, randomly chosen individual, and the strategy of this individual. Call x the strategy of i at the beginning of period t and call y the strategy of the other individual. The individual i considers switching from x to y . Suppose that the probability that i switch to y is:

$$\alpha + \beta [u(y, p^t) - u(x, p^t)], \quad (3.28)$$

where α and β are two positive constants. One can see then that the number of users of the strategy x is increased by one unit with the probability:

$$\sum_{y \in X} p^t(x) p^t(y) (\alpha + \beta [u(x, p^t) - u(y, p^t)]) \quad (3.29)$$

and is decreased by one unit with the probability:

$$\sum_{y \in X} p^t(x) p^t(y) (\alpha + \beta [u(y, p^t) - u(x, p^t)]) \quad (3.30)$$

On average, the evolution of the population size is:

$$\begin{aligned} n^{t+1}(x) &= n^t(x) + \sum_{y \in X} p^t(x) p^t(y) 2\beta [u(x, p^t) - u(y, p^t)] \\ &= n^t(x) + p^t(x) 2\beta [u(x, p^t) - u(p^t, p^t)] \end{aligned} \quad (3.31)$$

The size of the population being constant, this equation can be written in terms of the variable p :

$$p^{t+1}(x) = p^t(x) + p^t(x) \frac{2\beta}{n} [u(x, p^t) - u(p^t, p^t)] \quad (3.32)$$

Once again, going to continuous time and for $v = \frac{2\beta}{n}$, one can recognize the replicator dynamics equation (notice that the constant α disappears).

Finally, starting with the continuous-time deterministic dynamic process of the CPR model, one can again show the link with the replicator dynamics. The variation of $f^t(x_1, x_2)$ is inversely proportional to the time elapsed since the beginning of the process and one can write:

$$\dot{f}^t(x_1, x_2) = \frac{1}{t} [p_1^t(x_1) + p_2^t(x_2) - f(x_1, x_2)] \quad (3.33)$$

From such a formula, one can show that, for large t , one has approximately:

$$\begin{aligned} p_1^t(x_1) &\approx \frac{1}{CU_1^t} p_1^t(x_1) [u_1(x_1, p_2^t) - u_1(p_1^t, p_2^t)] \\ p_2^t(x_1) &\approx \frac{1}{CU_2^t} p_2^t(x_2) [u_2(p_1^t, x_2) - u_2(p_1^t, p_2^t)] \end{aligned} \quad (3.34)$$

which allows for a comparison between reinforcement learning and replicator dynamics. Looking at equations (3.16) and (3.34) one notices that the CPR rule defines a slowing down version of replicator dynamics; in effect, the cumulated utilities CU_i^t which appear in (3.34) are approximately proportional to the length of play.

3.3.4.2 A mixed learning model

Consider coordination games in which the individuals (all) prefer to coordinate but in which different conventions for coordination are possible. Such situations are very common: efficiency is improved when technological standards are in place, relations between landlords and exploitants are made easier by the existence of conventional crop sharing rules. When the evolution of such a system is caused by a selection process coupled with a mutation process, one has to compare the time scales needed for these two processes to act. In the *long* run, when the effect of mutations can be neglected, selection pressure leads to a stable state, even if different stable states can be attained. In the *very long* run, when the effect of mutations cannot a priori be neglected, they make possible jumps from one stable state to the other, and the behavior of the system depends upon the detail of the mutation process. For instance, it may be the case that in the long run any one of two possible conventions emerges, whereas in the very long run, one of them is selected much more frequently than the other.

In order to analyse this phenomenon, Young (1993a) introduced a stochastic perturbation in the *Fictitious Play* learning model, together with the consideration of populations of individuals. The following hypothesis are respective applications of the three principles of information, evaluation and decision, to hold at each date t :

1. Each individual has a *memory of depth* m , that is to say that she has no access to events that occurred before period $t - m$.
2. Each individual randomly draws a *sample* of size s of strategies chosen by the other individuals, among the strategies chosen at dates $t - m, \dots, t - 1$.
3. With probability $1 - \varepsilon$, the individual chooses a *best response* to the static distribution she observed by sampling, and with probability ε , the individual chooses at random her strategy (the number ε is an error rate).

The three parameters m, s, ε define a dynamic process called *adaptive process with memory* m , *sample size* s and *error rate* ε . “Fictitious play” corresponds to an adaptive process with unbounded memory size, exhaustive sampling and zero error rate.

Because the transitions do not depend on time, and because the agents all have a bounded memory, one can take, as the state space, the set of all sequences of length m . Such a sequence is simply called an history. Given

positive error rate, such a system explores all possible states. But if the error rate is very small, mutations are rare, and some states are visited very rarely. By definition, a state is *stochastically stable* if the probability that this state be regularly visited remains strictly positive when the error rate tends to zero. Notice that some states are very simple: for instance m time repetition of the same strategy profile is a state. By definition, a *stochastically stable equilibrium* is a strategy profile whose m -repetition is a stochastically stable state. The notion is a refinement of the notion of a Nash equilibrium.

One can see this notion at work in *two-player, two-strategy symmetric coordination games*. The bi-matrix for such a game is the following:

Table 3.8. Coordination game

	A2	B2
A1	(a,a)	(b,c)
B1	(c,b)	(d,d)

with the strict inequalities:

$$\begin{aligned} a &> c \\ d &> b \end{aligned} \tag{3.35}$$

which guarantee that the game has two strict equilibria.

The two equilibria correspond to two possible conventions for coordinating strategies. The variants of the “technological game” presented above are symmetric coordination games with $a = 2$ and $d = 4$. The equilibrium $E1 = (A1, A2)$ provides payoff d to each player and the equilibrium $E2 = (B1, B2)$ provides payoff d to each player. If a is larger than d , the equilibrium $E1$ is efficient and Pareto-dominates the equilibrium $E2$. Example 1 corresponds to the parameters:

$$a = 2, b = 1, c = 0, d = 4. \tag{3.36}$$

Another criterion for comparing equilibria can be defined by taking into account potential deviations from predicted play. For player 1, in the equilibrium $E1$, if the opponent deviates with probability e , the expected payoff for strategy $A1$ is: $(1 - e)a + eb$ and the expected payoff for strategy $B1$ is: $(1 - e)c + ed$. The optimal choice is $A1$ as long as the first expression is larger than the second one, that is as long as the opponent is making mistake with a probability lower than the *tolerance threshold* e_1 :

$$e_1 = \frac{a-c}{a-c+d-b} \quad (3.37)$$

For the equilibrium $E2$, tolerance for player 1 is:

$$e_2 = \frac{d-b}{a-c+d-b} \quad (3.38)$$

The equilibrium $E1$ is said risk-dominant if the tolerance is larger (for both players), in this case:

$$a-c > d-b \quad (3.39)$$

Notice that the loss $(a-c)$ is how much the player loses when she *her-self* deviates from equilibrium play. When the opponents deviates, the loss is $(a-b)$, a quantity which does not appear in the definition. Clearly, the risk-dominant equilibrium can be Pareto-efficient or Pareto dominated.

In the simple coordination games that are mentioned above, one can show that the stochastically stable equilibria are risk-dominant. For instance, in the Example 1, the stochastically stable equilibrium is $E2 = (B1, B2)$, which is risk-dominant (and Pareto-optimal). In the variant depicted in Table 3.9, which corresponds to: $a = 2, b = 3, c = 0$, and $d = 4$, the equilibrium $(A1, A2)$ is risk-dominant and Pareto-dominated.

Table 3.9. Technology game, variant 4

	A2	B2
A1	(2,2)	(3,0)
B1	(0,3)	(4,4)

Start now from the Example 2, in which $a=b=c=2$ and $d=4$, and change c . The two situations $E1$ and $E2$ remain Nash equilibria for all $c \leq a=2$. The Pareto-dominated equilibrium $E1$ (old technology) is risk-dominant and thus evolutionary stable if $2-c > 2$, that is $c > 0$; and the equilibrium $E2$ (new technology) is risk-dominant and evolutionary stable if $c > 0$. In more complex coordination games, the two notions diverge. But in bargaining games, stochastically stable equilibria can sometimes be shown to correspond to equitable agreements (for instance 50/50 split), which provides an evolutionary justification for equity norms (Young 1993b).

Intuitively, in the very long run, starting from an equilibrium state $E1$ and under the influence of random mutations, the system will some day jump to another equilibrium. But if the system has more than two equilibria, some equilibria are more likely than others to succede to $E1$. The fact that a particular equilibrium $E2$ has replaced $E1$ will be more or less surprising for an outside observer, depending on the “evolutionary distance” from $E1$ to $E2$. This distance measures the number of independent mutations that are needed for destabilizing $E1$ in the direction of $E2$. For instance, in a coordination game, an equilibrium whose tolerance is small requires, for being destabilized, a smaller number of mistakes than an equilibrium whose tolerance is large. In the very long run, the system visits the different equilibria, but spends more time in risk-dominant equilibria. Historical breaks that lead from one convention to another cannot be predicted as to their moments of happening, but their sequence can partially be understood.

3.4 Theses and conjectures

The models considered by evolutionist game theory appear very sensible to the details of modelization and may only be studied for specific classes of (repeated) games and (dynamic) processes. The spectrum of usual games goes from classical 2×2 normal form games (prisoners dilemma, battle of sexes, stag-hunt) to more complex extensive form games (ultimatum, chain-store, centipede). The spectrum of processes goes from the most simple ones (fictitious play, CPR, replicator) to processes combining ideas from various origin (reinforcement effect, aspiration levels, imitation mechanism). One constructs then some prototypes of well explored “game-processes”, from which it is necessary to interpolate or extrapolate in order to obtain general results, formal or more qualitative ones.

From a formal point of view, the most robust result concerns the elimination of strictly dominated strategies in any process which is utility-improving, a condition which is usually satisfied. The reason is that, if a process reinforces the pure strategies providing a utility greater than its average value, it implies progressively the extinction of the dominated strategy as much as the dominating strategy is effectively represented. Another rather robust result asserts that a Nash equilibrium is always a rest state (a state where a player stays if it is already there) for any purely deterministic (without randomness) process. The reason is that if the system is in a Nash equilibrium state, there is no incentive to induce a player to deviate from that state, since his utility would, at least locally, decrease.

An asymptotic result which is less robust asserts that the strict pure strategy Nash equilibrium states are attractors of many usual processes, contrary to the mixed strategy Nash equilibrium states. The reason is that a player's deviation from a pure strict Nash equilibrium in any direction leads to a strict loss of utility, which induces that player to come back to the initial situation. A complementary result states that, for a process with permanent randomness, the system stays for a long time in a Nash equilibrium, then may shift roughly to another Nash equilibrium. The reason is that, if some random factors sufficiently dispersed effectively lead the system to stay in a neighborhood of an equilibrium, more polarized random factors may drive the system out of its attraction basin and reach another equilibrium.

From a more qualitative point of view, it appears that the existence of random factors modifies seriously the transitory and asymptotic properties of the system with regard to a deterministic system. The random terms act in a differentiated way depending on their origin, either the occurrence of players' interactions, the statistical conditions of players' observation, the definition of players' expectations or the players' exploration behaviors. In the same way, it appears that a limited interaction neighborhood for players modifies deeply the system properties with regard to a homogenous interaction network. This neighbourhood acts in a very different way depending on its structure, either immediate neighbours on a line or a circle, or more distant neighbours on a multidimensional structure.

In a less precise way, the path of the system depends as much on the interaction networks than on the behaviour rules, at least if the last are reasonably improving. In fact, the transitory phenomena of diffusion of the players' strategies are strongly influenced by the degree of connexity of their relations, the behaviour's reactivity depending essentially on the speed of evolution. In other respects, the asymptotic states of the system are not homogenous, but manifest outstanding spatial and qualitative configurations, reinterpreted by the modeller as emergent structures. It is possible to obtain a segmentation of players' strategies in geographic or qualitative areas, or formation of privileged relations between players forming endogenous information networks (Durieu-Solal, 2001).

The asymptotic properties of evolutionist systems, well studied for normal form games, have to be extended to extensive form games as well. Specifically, it is interesting to know if a subgame perfect equilibrium state may be obtained as the limit of a learning process, several processes having already been studied. (Laslier-Walliser, 2004). For instance, a player may attribute a utility index at each branch of the game tree where he has

to move, according to the past payoff obtained when that branch was used, and apply the probabilistic model in order to explore all branches. The prototype here is the chess game, a zero-sum game with perfect information, for which one knows that it exists some couples of equilibrium strategies, however practically not computable due to the high number of possible strategies.

Besides, if the asymptotic properties of the usual processes are now well explored, the transitory properties are less studied even if as insightful. In fact, the convergence speed of some process is few studied by the game theorists contrary to what is done in Artificial Intelligence where many results were proved. Moreover, it is interesting to see if some typical patterns do not appear at the beginning or middle step of the transitory phase, and especially more or less rough discontinuities. Finally, multiple simulations allow to make precise the influence at middle term of the various random factors which are integrated in the process.

The economic applications, badly explored till now, have to be more systematic, in order to extend the scope initially limited to the adoption of technologies or the duopolistic competition (Ania-Tröger-Wambach, 2002). The general results obtained on specific classes of games can be directly transferred to specific economic models expressed in form of games belonging to that class. Conversely, some learning assumptions proposed in economics and unknown in game theory can be transposed in order to be studied more systematically. Such a point of view was adopted in the case of the duopoly model (Kirman, 1995), where the learning is an epistemic one of a Bayesian nature (adjustment by mean squares).

References

- Ania, A.B., Tröger, T., and Wambach, A. (2002): An evolutionary analysis of insurance markets with adverse selection, *Games and Economic Behavior*, 40, 153-184.
- Aumann, R. and Brandenburger, A. (1995): Epistemic conditions for Nash equilibrium, *Econometrica*, 63(5), 1161-80.
- Axelrod, R. (1984): *The Evolution of Cooperation*, Basic Books.
- Bush, R. and Mosteller, F. (1955): *Stochastic Models for Learning*, Wiley.
- Durieu, J. and Solal, P. (2002): Adaptive play with spatial sampling, *Games and Economic Behavior*, 43 (2), 189-195.
- Fudenberg, D. and Levine, D. (1998): *The Theory of Learning in Games*, MIT Press.

- Harsanyi, J.C. (1967): Game with incomplete information played by “Bayesian” players, *Management Science*, 159-82, 320-34, 486-502.
- Harsanyi, J.C. and Selten, R. (1988): *A general theory of equilibrium selection in games*, MIT Press.
- Kalai, E. and Lehrer, E. (1993): Rational learning leads to Nash equilibrium, *Econometrica*, 61, 1019-1045.
- Kandori, M., Mailath, G., and Rob, R. (1993): Learning, mutation and long run equilibria in games, *Econometrica*, 61, 27-56.
- Kirman, A. (1995): Learning in oligopoly: theory, simulation and experimental evidence, in A. Kirman and M. Salmon (eds.), *Learning and rationality in economics*, Blackwell.
- Laslier, J.-F., Topol, R., and Walliser, B. (2001): A behavioral learning process in games, *Games and Economic Behavior*, 37, 340-366.
- Laslier, J.-F. and Walliser, B. (2004): A reinforcement learning process in extensive form games, *International Journal of Game Theory*, 33, 219-227.
- Maynard Smith, J. (1982): *Evolution and the Theory of Games*, Cambridge University Press.
- Monderer, D. and Shapley, L. (1966): Fictitious Play property for games with identical interests, *Journal of Economic Theory*, 68, 258-265.
- Monderer, D. and Shapley, L. (1996): Potential games, *Games and Economic Behavior*, 14, 124-143.
- Nash, J. (1951): Non cooperative games, *Annals of Mathematics*, 54, 289-93.
- Nöldeke, G. and Samuelson, L. (1993): An evolutionary analysis of backward and forward induction, *Games and Economic Behavior*, 5, 425-454.
- von Neumann, J. and Morgenstern, O. (1944): *Theory of Games and Economic Behaviour*, Princeton University Press.
- Robinson, J. (1951): An iterative method of solving a game, *Annals of Mathematics*, 54, 296-301.
- Robson, A. and Vega-Redondo, F. (1996): Efficient equilibrium selection in evolutionary games with random matching, *Journal of Economic Theory*, 70, 65-92.
- Roth A. and Erev, I. (1995): Learning in extensive form games: experimental data and simple dynamic models in the intermediate term, *Games and Economic Behavior*, 8, 164-212.
- Rubinstein, A. (1994): *Models of bounded rationality*, MIT Press.
- Samuelson, L. (1997): *Evolutionary Games and Equilibrium Selection*, MIT Press.
- Skyrms, B. (1996): *Evolution of the Social Contract*, Cambridge University Press.
- Vega-Redondo, F. (1996): *Evolution, Games and Economic Behavior*, Oxford University Press.

- Vega-Redondo, F. (1997): The evolution of Walrasian behavior, *Econometrica*, 65, 375-84.
- Walliser, B. (1998): A spectrum of equilibration processes in game theory, *Journal of Evolutionary Economics*, 8, 67-87.
- Weibull, J. (1995): *Evolutionary Game Theory*, MIT Press.
- Young, H.P. (1993a): The evolution of conventions, *Econometrica*, 61:57-84.
- Young, H.P. (1993b): An evolutionary model of bargaining, *Journal of Economic Theory*, 59, 145-168.
- Young, H.P. (1998) *Individual Strategy and Social Structure, an Evolutionary Theory of Institutions*, Princeton University Press.

4 Market with irreversibilities

In chapter 2, it was demonstrated that, under certain hypotheses, the unique price that becomes established in the market is constructed progressively by agents entering into contact with each other and dividing simultaneously into two groups, those who buy or sell and those who no longer operate.

Four hypotheses were made in order to reach this result:

- The good being exchanged undergoes no modification over time.
- There are no information costs or adaptation costs of any kind for any of the agents.
- Any agent can enter into contact with any of the other agents at any time and together with his memory these contacts constitute his only source of information.
- Each agent is only concerned with obtaining the best price – the highest if he is a seller, the lowest if he is a buyer. The minimum price acceptable to the seller and the maximum price acceptable to the buyer are invariant.

These hypotheses entail a concept of time with no direction. Whatever the state of the market at a given time, its final state will be the same. In other words, there are no irreversibilities, i. e. transformations of the good, constraints or costs implicit in the change of the market from one state to another.

Many different forms of irreversibilities are encountered in economics. In this chapter we shall consider three of them:

- Costs that we can consider as being connected with *frictions*. We shall examine two aspects of these costs successively: *information costs* sustained by an agent looking for a job or for a worker and *transition costs* (costs of laying-off or training workers for a firm, costs of adaptation to a new job for a worker),
- *Irreversibilities of investment*, which we shall examine in the context of one particular case: that of workers who invest by accepting the expense involved in moving locality and in changing from one labour market to another,

- *Irreversibilities of knowledge*, as the development of the market enables agents to acquire particular skills and knowledge.

Introducing irreversibilities into the traditional market equilibrium model is no easy task, as they only become apparent in a dynamic of evolution through a succession of states. On the other hand, they can be taken into account quite naturally in models derived from the one given in chapter 2.

The presence of irreversibilities generates a much wider range of results. The stable state characterised by a unique price independent of the history of the market is often replaced by multiple stable states associated with dispersed prices. The attainment of any one of these states is dependent on the random events affecting the progression of the system.

4.1 Background and problems

The impact of irreversibilities on the functioning of markets has of course preoccupied economists from the beginning. We could take as an example the many works devoted to the imperfection of information. For a long time, analysis was hindered by the domination of static models, which sometimes even resulted in inaccurate conjectures. However, certain works published over the last few decades have contributed to progress on this question. Here we will limit ourselves to citing the studies of Stigler on the economics of information (1961) and the analyses of J. Stiglitz (1967) and of S. Salop and J. Stiglitz (1982) on equilibrium in markets with imperfect information. These works have started to bring to the fore phenomena studied for some time by physicists working on the dynamics of irreversible systems, but they have yet to establish the general framework.

In fact, irreversibility implies “*a breaking of symmetry between before and after*” (Prigogine and Stengers, 1988). Hence the possibility of the occurrence of *events*. “By definition, an event cannot be deduced from a determinist law. It implies, in one way or another, that what has occurred could have not occurred, therefore reflecting a realm of possibles that no knowledge can reduce.” Nevertheless, an event is only of interest if it is significant, in other words likely to transform future evolution and thus generate new coherences.

These aspects, often hidden in the literature on the subject, will be brought out clearly in the models in this chapter.

4.2 Canonical principles

We retain the conceptual framework used in chapter 2:

- time is broken down into a series of periods,
- different categories of agents coexist in the market,
- the agents are constantly carrying out activities of exploration and exchange,
- during each period, the agents seek, discover and receive information and either modify or do not modify their exchange requirements accordingly,
- during choices, meetings or discoveries, random phenomena occur and the evolution of the market is thus described by a stochastic process.

On the contrary, the simplifying hypotheses that ensured the absence of irreversibilities have been removed. Now, if agents benefit from the utility obtained from exchanges, they also sustain information or adaptation costs, lay out expenditure in investment or see their performances or the conditions of their choices modified.

Unlike the simple market, which converges towards a stable state with a unique price, independent of its history (and therefore predictable), models with irreversibilities generate much more varied configurations.

- They may converge to a stable state dependent on the history of the market or endlessly generate fluctuations.
- From a given initial state, the market may end up in any one of a multiplicity of stable states, in each of which the jobs occupied and the workers employed are different.
- In each stable state, we may observe a unique price or price dispersion.
- The evolution of the markets may even, in certain models, lead to more original phenomena, such as the emergence of new markets.

4.3 Some models

4.3.1 Model with information costs

We take the labour market model described in chapter 2 as our starting point.

When information costs occur, individuals must assess, before they buy a unit of the good, whether or not it is preferable to sustain the information cost. They are confronted with an exploration-exploitation dilemma and they must consequently estimate the prospects of the price that they could obtain if they started seeking. They can no longer content themselves simply with reacting to the propositions they receive. *Anticipation is necessary even within the framework of procedurally rational behaviour.*

Individuals are interested in discovering the jobs available for two different reasons: actually to become candidate a job or to learn about the situation of the market in order to decide whether or not they should continue their search.

To take these two aspects into account, we need to modify the labour market model of chapter 2. We shall now briefly describe the changes made to the model.

Each period t is now broken down into two sub-periods.

At the start of the first sub-period, each individual knows (as a consequence of the second sub-period of the previous period), a subset of jobs and, if he is employed, his current job. In addition, each firm has announced the salary that it will offer for the period $(t + 1)$.

Individuals then come into the market one after another, once each and in random order. Each individual visits all the jobs he knows and applies for the best job he finds if it is attractive, in other words if it offers a salary above the level that he has previously fixed for himself. If the job is already occupied, the incumbent individual has a pre-emptive right to the salary offered on condition that he has not yet come into the market. The job is given either to the first applicant or to the incumbent and the contract is binding for the period $(t + 1)$. There is one exception to this rule: if the firm has not yet found an applicant when the incumbent comes into the market, it will offer him *the same salary* for the period $(t + 1)$ as it paid him for the period t .

During the second sub-period, there are four categories of agents with different motivations operating in the market.

- Individuals who have found a job for $(t + 1)$ content themselves with searching for information and registering for the subset of jobs which they will consider during the first sub-period of $(t + 1)$ in order to find a job for $(t + 2)$.

- The individuals who have not yet found a job for $(t + 1)$ add the desire to find a job for $(t + 1)$ to the motivation of acquiring information.
- The firms in which the jobs are occupied for $(t + 1)$ seek only to register individuals as potential applicants during the first sub-period $(t + 1)$.
- The firms in which the jobs are still unoccupied for $(t + 1)$ also wish to find recruits for this period.

Firms in the first category, therefore, only have one salary offer to make (for the period $(t + 2)$), whereas those in the second category must propose a salary for $(t + 1)$ and a salary for $(t + 2)$.

During this second sub-period, the individuals come into the market one after another, once each and in random order. The individual present in the market draws a random sample of firms, learns the salary offer(s) made and decides, depending on his situation:

- whether to be candidate for a job for $(t + 1)$ and if so, which one,
- whether to register with a firm to be taken into consideration during the first sub-period of $(t + 1)$.

All individuals are accepted for registration and the first applicant for a job available for the period $(t + 1)$ obtains the job at the salary offered.

In such a model, there may be several types of information costs. We shall limit ourselves to the case in which there is only *one fixed search cost* which an individual must pay to obtain the right to draw a sample of firms.

We must now describe the behaviour of individuals. This behavior covers three activities: *observation of the market*, the *act of applying for a job* and the *adaptation of anticipations*.

As regards *observation*, one possible way of describing individual behaviour consists in assuming that at the start of the second sub-period of the period t , the individual makes five different estimations:

- the first is the value he attaches to immediate observation of the market,
- the second is the value he attaches to observation of the market during the period $(t + 1)$ if he has not observed the market during the period t ,
- the third is the value he attaches to observation of the market during the period $(t + 1)$ if he has observed the market during the period t ,

- the fourth is the salary he thinks he can obtain for the period $(t + 1)$ if he seeks a job for this period during the second sub-period of t ,
- the fifth is the salary he hopes to obtain for the period $(t + 2)$ if he registers during the second sub-period of t .

The individual then compares the present value (this assumes the introduction of an individual discount coefficient *over the shortest possible significant horizon, i. e. two periods*), of his expected income as a result of the search with the sum obtained by adding the search costs to his income in the absence of any search. Several possibilities can then be examined, but there is little interest in going into the details in this book.

Job application behaviour is obvious for the second sub-period: any job which offers the individual more than the minimum salary is preferable to unemployment. It is more subtle in the first sub-period, as the individual must define the level x above which he will accept offers rather than take his chances in the second sub-period.

There remains for us to consider the *adaptation of anticipations*. To do so, we shall introduce the following hypotheses:

- for each individual, there exist upper and lower limits to the value of observation of the market; this value grows when the individual does not search and reaches its maximum within a finite time; it does not grow when the individual searches.
- If, during T consecutive observations, an individual has not found during the second sub-period a job available for the following period offering more than y , he does not expect, during the second sub-period of t , to obtain more than y for the period $(t + 1)$. Naturally, an analogous hypothesis is made for the period $(t + 2)$.

These hypotheses are essential. They signify that the agents *learn* and *adapt* their anticipations accordingly.

It should be noted, before we go any further, that the introduction of information costs and therefore of anticipations requires a subtler description of the market and generates more complex individual behaviour, even if the rationality is procedural and the horizon of anticipation is short-term.

As in chapter 2, we denote by K the number of workers who will be employed and the number of jobs that will be occupied in the market in the absence of information costs.

Under the above hypotheses, the market converges in probability towards a stable state, but numerous stable states exist in addition to the traditional

equilibrium. These states are not efficient (the equilibrium is not represented simply by the employment of the first K workers by the first K firms) nor do they have a unique price (several salaries coexist in a stable state).

As one might expect, the individuals separate into two categories:

- *passive individuals*, for whom the maximum value given to observation of the market is lower than the search costs,
- *active individuals*, for whom this is not the case.

In a stable state, active individuals in employment continue searching, from time to time, to verify that they cannot find a better offer in the market. Passive individuals, having searched in a transient way during the process, put a definite stop to their search after a certain date, because they consider that the estimated extra gain they could hope to obtain does not compensate for the loss resulting from the search costs.

In a stable state, the situation of these two categories of individuals is radically different. Active individuals in employment receive the highest salary that can be observed in the market. This salary is at least equal to the equilibrium salary of chapter 2, and it may be considerably higher. Passive individuals in employment receive any one of the salaries present in the market.

Employed active individuals work in one of the first K firms. Passive individuals, on the other hand, may work in a firm that is not among the first K firms, in a job, therefore, that would not be occupied in the traditional equilibrium.

It can be demonstrated that as the number of active individuals increases, the maximum salary decreases. The appearance of new stable states originates in the existence of passive individuals among those who would be employed in the traditional equilibrium.

This model illustrates two important consequences of the presence of irreversibilities in an evolutionary context.

1. *The final state of the system cannot be predicted. It depends on the history* which randomly orientates the system towards one state or another, so that the traditional equilibrium now appears as just one out several possible outcomes.
2. Even if all the agents have an interest in obtaining a higher salary, the diversity in psychological attitudes towards searching have a profound influence on the process: a market with few active agents will only accomplish a very partial sorting of workers and jobs and will maintain a *high level of salary dispersion*.

In this way, the introduction of irreversibility in the form of information costs generates very different dynamics to the dynamic of “reversibility” presented in chapter 2. As the future is no longer written in advance, a whole range of stable states may appear and themselves constitute the starting point of different future trajectories. History and its random nature may lead the market towards a promising stable state or trap it within states involving the definitive abandonment of certain expectations.

Figure 4.1 uses a simple case to give an example of a stable state that is not efficient and does not have a unique price.

On the left (4.3.a) the market of chapter 2, in which, in stable state, 6 individuals are employed with salary p , the seventh is unemployed. Likewise, the first six jobs are occupied, the seventh not. This stable state is concentrated (only one salary on the market) and efficient (the first six individuals are employed and the first six jobs occupied). On the right (4.3.b) the market with information costs of this chapter. One possible stable state is the following:

- individual 1, assumed passive, can occupy job 7,
- individuals 2, 3, 4, 5 and 6 can occupy jobs 1, 2, 3, 4, 5 with salary v_5 , greater than p
- individual 7, active or passive, can be unemployed.

This stable state does not have a unique price as salaries v_5 and v_7 exist. Neither is it efficient, because job 7 is occupied whereas job 5 is not.

4.3.2 The existence of friction costs

We shall now modify the chapter 2 model to introduce adaptation costs with the character of friction costs. We introduce the two following transformations.

- With each job i ($1 \leq i \leq n$) we associate a cost $c_i \geq 1$ such that, if the job is occupied, the firm will not offer a new candidate the same salary reduced by one unit, but the same salary reduced by c_i , for when the firm replaces the incumbent worker it is assumed to sustain a cost $(c_i - 1)$ representing the expenses involved in laying-off the incumbent and in training the new recruit, the psychological cost of the change, etc.
- With each individual k ($1 \leq k \leq m$) we associate a cost $d_k \geq 0$ such that, if the individual is already employed, he will only accept a new




Fig. 4.1. Example of a market with 7 individuals and 7 jobs.

job if his salary increases by $d_k + 1$, because if he changes job, the individual k may have to pay a sum d_k representing the monetary and psychological cost of his re-adaptation.

It is immediately obvious that this representation of friction costs is rather cursory: it is not simply over the next period that the firm hopes to make savings of more than c_i but over the whole length of occupancy of the job by the new recruit; likewise it is not simply over the next period that the individual hopes to obtain an increase in income of more than d_k but over the whole duration of his new job. Let us, nevertheless, boldly assume that the agents only take into account their immediate advantages, thus exempting them from making anticipations which, strictly speaking, are indispensable.

The introduction of friction costs widens the set of stable states and produces, in addition to the traditional equilibrium:

- stable states that are efficient but may not have a unique price;
- stable states in which the first K individuals are employed but some of the first K jobs are unoccupied;

- stable states in which the first K jobs are occupied but some of the first K individuals are unemployed.

The dispersion of salaries in a stable state, measured by the difference between the highest and lowest salaries paid in the market, can be given an upper limit. This limit brings into play the sums $(c_i + d_k)$ that can be associated with the pairs (i, k) of each job and the individual occupying this job in a given state of the market. The analysis thus confirms that the higher the friction costs, the wider the range of possible dispersion of salaries.

In general, therefore, friction costs prevent the construction of a unique price.

4.3.3 The presence of geographically dispersed markets

In a very simplified form, this new model will illustrate phenomena of considerable importance in geographical economics and which traditional microeconomics has always had difficulties in dealing with.

The jobs are spread geographically over markets located in L different localities. Individuals must live in the locality in which they work. However, they can move from one locality to another to gain access to new jobs.

An individual living in one locality and thinking of moving to another directly compares the salaries offered to him with a *mobility cost* associated with the move from the first locality to the second. This cost includes not only the monetary and psychological cost of moving but also the values he attaches to his future work and his salary prospects in the two cases. This cost therefore implicitly includes the estimates resulting from the individual's anticipations. This manner of formalising the way in which the individual takes the future into account is rudimentary and much less satisfactory than the method used in the model with explicit information costs. It nevertheless provides some interesting results. Naturally, mobility costs can be considered as the investment required to move from one locality to another.

Here, a state of an economy is characterised by a geographical distribution of individuals between the localities and by the allocation of individuals to different jobs or to unemployment.

We shall now introduce, in a simplified form, the concept of *collective utility*.

By definition, collective utility is the sum of the utilities of all the agents, both individuals and firms; the utility of an individual is equal to the sum of a constant (his utility when unemployed) and the difference, if he is

employed, between the salary he is paid and his minimum salary \underline{w}_k ; similarly, the utility of a firm is the sum of a constant (the utility of the job if it is unoccupied) and the difference between the maximum salary \bar{v}_i and the salary paid by the firm.

These definitions are equivalent to using the monetary values paid or received to determine utilities.

We are thus led to consider two concepts of optimum state for the whole set of markets. These concepts bring into play the differences in collective utility between two states, by considering that all the mobility costs incurred during the change from one state to the other represent a loss sustained by the individuals.

1. Firstly, a state is said to be a *conditional optimum* if, *starting from this state*, no possible changes exist that would result in an increase in collective utility.
2. Secondly, a state is said to be a *global optimum in relation to an initial state*, if, taking mobility costs into account, the increase in collective utility associated with such a change is equal to or greater than that associated with any other change from the same initial state.

Two concepts of stability can be associated with these two concepts of optimality.

1. *Conditional stability*: a state is said to be conditionally stable if, for every individual, his salary in this state is higher than that offered to him for any other job, once the mobility costs – if there are any – have been deducted.
2. *Global stability*: a state is said to be globally stable in relation to an initial state if it is not rejected by any individual or any coalition of an individual and a job, taking into account the states that these coalitions can attain when starting from the initial state.

It is possible to demonstrate a double equivalence: firstly between the set of conditional optimums and the set of conditionally stable states and secondly between the set of global optimums and the set of globally stable states.

Now let us return to the chapter 2 model, introducing a distribution of jobs to the different localities (this distribution is unchanging) and an initial distribution of individuals who can move from one locality to another, paying mobility costs to do so. Starting from this initial state, the dynamic processes described in chapter 2 ensure convergence in probability of all

the markets towards a *conditionally stable* state, but there is no guarantee that this state will be globally stable in relation to the initial state.

This result can be understood immediately by means of a simple example. Let us assume that the global optimum implies that the individual k_1 , initially unemployed and living in $l(k_1)$, occupies the job i_1 located in $L(i_1)$. It is possible that the firm i_1 makes an offer $s(i_1)$ for this period that is too low to interest the individual k_1 :

$$s(i_1) - \underline{w}(k_1) - d[l(k_1), L(i_1)] \leq 0 \quad (4.1)$$

where $d[l(k_1), L(i_1)]$ denotes the mobility cost of moving from $l(k_1)$ to $L(i_1)$. The individual can then accept an offer from another firm and move to $L(i_2)$. But once he is in $L(i_2)$, the offer of i_1 may never become interesting, for:

$$\bar{v}(i_1) - \bar{v}(i_2) - d[L(i_2), L(i_1)] \leq 0 \quad (4.2)$$

The individual k_1 can thus get trapped in a conditionally stable state.

This would not have occurred if the individual k_1 and the firm i_1 had formed a coalition to refuse any move that did not provide them with at least:

$$\bar{v}(i_1) - \underline{w}(k_1) - d[l(k_1), L(i_1)] \quad (4.3)$$

This result illustrates a central aspect of the geographical evolution of an economy: the spatial distribution of individuals is not predetermined, but a result of the history of adaptation of the markets, and from a given initial state, mobility costs may shut the economy into a state very different from the global optimum. The combination of the irreversibility generated by mobility costs and the absence of coordination between agents results in a distribution of economic activities that is only conditionally optimal.

4.3.4 The acquisition of skills

We shall now illustrate the irreversibility that results from improvements in workers' knowledge: the particular case of the transformation of the good exchanged during endogenous evolution of the market.

We again take the labour market model of chapter 2 as our starting point, this time adding certain possibilities.

To begin with, all the workers have the same professional ability level (ability 1 of "unskilled workers"); however, some of them may randomly acquire a higher level of ability when they are in employment (ability 2 of "specialised workers").

To begin with, all firms only offer jobs of ability 1, corresponding to the use of a basic technology (technology 1), but they can also, later on, envisage a more complicated technology (technology 2) – either known from the start or randomly discovered – requiring the employment of an individual of ability 2.

An individual of ability 2 can occupy (and accept) jobs of either ability 1 or ability 2.

Unlike the model presented in chapter 2, and simply to facilitate the presentation, it is now the firms that enter the market in random order and draw random samples of individuals. For each period, a firm defines the salaries it offers for a job in technology 1 or technology 2. Individuals either apply for the jobs or not, and the firm then chooses to keep the positions vacant if there are no applicants, or the only possible organisation if there are no applicants for the other organisation, or the most profitable of the two organisations if it has the choice.

An individual k of ability 1 employed by the firm i has the probability $p_{ki} \geq 0$ of acquiring the ability 2 (this probability depends on the employer, for being employed by a firm which only gives its personnel the strict minimum in terms of training is not equivalent to being employed by a firm with a positive training policy – such as IBM, for example). All individuals are assumed initially to have an ability level of 1.

In this model, irreversibility is constituted by the jump in the ability level of certain individuals. This jump occurs randomly as the market evolves.

We introduce two definitions that appear quite natural:

- *promotable* individual is an individual such that there is a date t and a succession of non-zero probability market states e_0, e_1, \dots, e_t such that this individual has acquired ability 2 on date t ,
- a *promoted* individual is a promotable individual such that there exists a date on which this individual has acquired ability 2 following the actual functioning of the market.

Under the hypotheses of chapter 2, the market thus described converges in probability towards a stable state, but there are two categories of stable state:

- stable states that one can qualify as *finished* because all promotable individuals who would have been employed if they had had ability 2 from the start have indeed been promoted and employed,
- *unfinished* stable states that do not have this property.

When the market ends up in an unfinished stable state, it does not succeed in realising all the potential abilities latent within all the individuals.

In this model, *the market can either realise or sterilise potentialities*. Individuals who have acquired ability 2 are sorted out from those who remain at ability 1 and, simultaneously, a second market for individuals of ability 2 is created alongside the initial market. The characteristics of this new market vary depending on the history.

By generalising the lessons drawn from this model, we can obtain an important result: for an economy to be trapped in a stable state that sterilises a proportion of potential resources may have more far-reaching consequences than the simple respective size of the two labour markets. There is a strong chance that, at some time in the future, innovations will appear, the implementation of which will require a sufficient volume of skilled labour. If the economy does not have this required volume, it will be obliged to renounce the adoption of these innovations. It is well-known, for example, that a country possessing only a handful of engineers or scientists cannot nurture the hope of developing – or even simply using – certain new technologies. Even so, the difference that grows between two initially identical economies is *quantitative* during a first stage – through variations in the respective numbers of workers of different categories – and then becomes *qualitative* during a second stage, depending on whether or not each economy can gain access to certain classes of innovation. It is therefore possible that economies originally sharing the same prospects for growth move ever further apart in terms of their respective levels of development.

4.4 Theses and conjectures

Two observations can be drawn from the above panorama:

1. The evolutionary paradigm proposes a framework within which irreversibilities can be represented “naturally”. These irreversibilities constitute an essential element of social and economic phenomena, whether in the field of investments, training, adaptation costs, etc., and in turn make it logical to introduce agents’ anticipations – or more generally their projects – even if it is possible to construct models in which the agents are myopic. We thus obtain the triptych of necessity, chance and will, without which no satisfactory description of economic evolution is possible.
2. In the presence of irreversibilities, the stable states towards which markets converge may be multiple and the stable state actually

reached depends on both the initial conditions and the history. This multiplicity of futures can contribute to the birth of institutions, as we shall see in chapter 8. A brief glimpse of this possibility is shown in the model of the evolution of abilities. In models of this type, the openings allowed by the innovation or the traps resulting from certain choices appear naturally as consequences of the hypotheses made, without having to be artificially added to the theoretical analysis.

References

- Lesourne, J. (1991): *L'Economie de l'ordre et du désordre*, Economica.
- Prigogine, I. and Stengers, I. (1988): *Entre le temps et l'éternité*, Fayard.
- Salop, S. and Stiglitz, J. (1982): The Theory of Sales: a simple model of equilibrium price dispersion with identical agents, *American Economic Review*.
- Stigler, J. (1961): *The economics of Information*, Journal of Political Economy, 69.
- Stiglitz, J. (1979): Equilibrium in Product Market with Imperfect Information, *American Economic Review*.

5 Mimetic interactions

Population dynamics is one of the most studied subjects in evolutionary economics. We have already seen many examples in the preceding chapters. It concerns analysis of the ways in which a large number of interacting individuals evolve. The particularity that distinguishes this type of analysis from the standard Walrasian approach is the fact that it integrates direct interactions between the agents. The classical theory of Walrasian competition, on the contrary, assumes that individuals have no direct relation with each other; they are only connected indirectly through the intermediary of the prices set by the auctioneer, as chapter 2 has already explained. Furthermore, this theory postulates exogenous preferences such that the end objectives of each agent are in no way influenced by what the rest of the group decides. Thus, each individual is isolated. This is the classic conception of individual independence in relation with the group. The fact, for example, that a certain number of individuals decide to buy the good k has no direct effect on the behaviour of a given individual A . It in no way increases his desire to purchase the same good k . The only effect is indirect, through the price variations it causes: the initial purchases provoke a price rise which, in general, will lead the individual A to reduce his purchases of good k . A vast range of possible interactions breaking with the Walrasian framework is presented in chapter 3, together with the formal tools enabling them to be analysed. Among these interactions, one plays a particularly important role: imitation. Chapter 3 provides a first analysis of it. We propose to go further in the present chapter, by studying its properties in more detail. Focusing particular attention on imitation can be justified by the tremendously important role it plays in all the social sciences.

5.1 Background and problems

In economics, the role of imitation was only recognised relatively recently, although some heterodox works in financial economics highlighted very early on the primordial importance of imitative behavior for anyone seeking to understand the dynamics of speculation (Aglietta and Orléan, 1980; Kindleberger, 1978). The central obstacle to the acceptance of imitation

derived from the too-exclusive identification of the rational individual with the Walrasian individual, – an individual isolated from the others, perfectly independent and who has no need to observe other people to decide what he wants. From this point of view, strong links with the group – such as the imitation and collective manias which are its strongest form of expression, – appear to belong to the domain of sheer irrationality. Imitative man has been equated with the man of the masses, perceived as being the absolute opposite of *homo oeconomicus*: blinded by collective passion, he abandons his own interests to follow unquestioningly the group movement. In economic analysis, the propensity to imitate was thus considered an archaic remainder, an anachronism that no longer had its place in the individualistic world of interest, markets and rational calculation. It was seen as a phenomenon of which the study was best left to sociologists, social psychologists or even anthropologists, – but certainly not to economists.

However, as our understanding of economic dynamics has deepened, the oversimplified nature of this view has become more apparent. For example, in the field of finance, it would be inconceivable to continue as if herd behavior or contagious processes play no role, when they are constantly appearing in the works of analysts and historians studying stock markets or currency markets (Kindleberger, 1978). Furthermore, advances in the microeconomic analysis of interactions have, little by little, demonstrated that imitation is not incompatible with rationality, – far from it – although it may lead, globally, to sub-optimal configurations. Three rational motives for imitation have been demonstrated (Orléan, 2001). We shall now examine them in one after another.

5.1.1 Informational imitation

When two individuals are confronted with the same problem, individual *A* may decide rationally to imitate individual *B* because he believes that *B* possesses information or knowledge that he himself does not have, so that he interprets *B*'s action as being more efficient. Take, for example, the situation where *A* and *B* are in a room with two doors. Only one of these doors leads outside, and *A* does not know which it is. We then assume that a fire breaks out. If *A* sees *B* get up and run towards a certain door, it is rational for *A* to follow him. Either *B* knows no more than *A*, and following him is simply a way for *A* to make a random choice between the two doors, or *B* does know something more, in which case it is advantageous for *A* to follow him. This is an example of informational imitation: one person copies another because of the information the latter is assumed to

possess. This is perfectly rational behaviour. It has been analysed by S. Bikhchandani, D. Hirshleifer and I. Welch (1992). Note that in this type of imitation, the intrinsic utility of individuals is not affected by the behavior of others. We do not assume the presence of positive externalities. In our example, individual A does not obtain any extra utility from the fact that he chooses the same door as B. We could even consider that the opposite is true, as two individuals may hinder each other while trying to get through the same door, which would oblige us to take negative externalities into account.

5.1.2 Normative imitation

When a group adheres collectively to a norm, we can see that transgressors are liable to sanction (Elster, 1989; Orléan, 1997). This is the case when, for example, not adopting a certain form of behavior will tarnish an individual's reputation in the eyes of the other members of the group. Under these conditions, it can be rational to imitate the collective behavior in order to avoid the various possible sanctions. The existence of these sanctions drives individuals towards conformism. This is what we shall call "normative imitation". The norm is not necessarily exogenous. It can emerge unintentionally, when individuals are afraid to stray from the majority opinion, without knowing exactly what this majority opinion is. The illustration we shall propose for exploring normative imitation will be of this type. More generally, normative imitation is also used to describe an individual making a choice simply because the same choice has been widely selected within the group and has thus become "legitimate and normal" in the eyes of the individual.

5.1.3 Preferential imitation⁸

Here we place ourselves in situations where the preferences of each individual are the combined result of two components, one of which is intrinsic and the other collective (Schelling, 1979). Consequently, the utility

⁸ In Orléan (2001), which serves as a guiding thread through this paper, the term „self-referential imitation” was preferred. But the reality is the same. The adjective „preferential” emphasises the source of the phenomenon, namely preferences. The adjective „self-referential” emphasises the fact that when choosing his behaviour, each agent takes as his main reference the endogenous behavior of the group, equally constituted of imitators.

$U^i(\alpha)$ that the individual i obtains when he chooses the option $[\alpha]$ is written in the form:

$$U^i(\alpha) = U^i(u_\alpha^i, n_\alpha^i) \quad (5.1)$$

where u_α^i is the component that determines the intrinsic utility of the option $[\alpha]$ for the individual i and n_α^i designates the number of individuals who have chosen the option $[\alpha]$ within what we shall call the “reference group” of the individual i , in other words the subgroup of individuals on whom i depends for his economic activity. This may, for example, involve “neighbors” with whom he has regular exchanges. We are in a situation of “preferential imitation” when the equation (5.1) obtained is such that when n_α^i increases, the utility $U^i(\alpha)$ grows. In other words, the higher the number of individuals in his reference group who have chosen the option $[\alpha]$, the more positively individual i evaluates the utility of this option⁹. One particular case is that in which the reference group of each individual is the whole set of individuals: $n_\alpha^i = n_\alpha$. In the case of preferential imitation, we take into account the positive externalities that drive the individual to tend to choose the option already chosen by those with whom he has the most frequent contact. The most classic example comes from the study of technological innovations and what Brian Arthur (1988) calls “network externalities” and “increasing returns of adoption”: the higher the number of individuals choosing a technical innovation, the more desirable this innovation becomes in the eyes of the remaining individuals.

Note that there may be a certain ambiguity between normative and preferential imitation, insofar as that in both cases the source of the imitation lies in the extra individual utility generated by conforming to majority behavior. Consequently, the justification for differentiating between these two forms of imitation lies in the distinction between the sources of utility: in one case, the utility considered is the traditional utility of economists, as in the analyses of Brian Arthur; in the other case we introduce a specific utility, connected to the recognition by the group of the legitimacy of an action or opinion. The formal model may be identical in these two situations, but the interpretation is different. In a financial market, for example

⁹ It may be observed that here we assume implicitly that any two agents in the reference group have an identical influence on i . Subsequently, we shall be able to abandon this hypothesis and assume that individuals have varying degrees of influence, depending, for example, on their „proximity“ to the individual in question.

(Orléan, 2001), we cannot equate the proposition: (a) the fund manager A has bought the security X because he thinks that the other fund managers will buy it, thus causing the price to rise (preferential imitation) with the proposition: (b) the fund manager X has bought the security X because to do otherwise would have tarnished his reputation, made him look incompetent and lost him some of his customers (normative imitation).

The integration of imitation into models of interaction has decisive consequences. Hans Föllmer (1974) was one of the first researchers to grasp the full significance of this. Analysing works that introduced random elements into the preferences of individuals, he observed that as long as the random events were assumed to be independent, this introduction did not profoundly modify the classical results: “randomness alone, without interaction, does not seriously affect the existence of price equilibria” (p. 51). But, he continued, we obtain a completely different picture when the random events disturbing the preferences of an individual are dependent on the preferences of his reference group, in this case his “neighbours”. He wrote: “It is through waves of imitation that interacting preferences become an important source of uncertainty” (p. 52). And then what happens? “the microeconomic characteristics may no longer determine the macroeconomic phase” (p. 52). This is the essential point, and it is important to understand just how much is at stake here. The central idea is that when we introduce imitative behaviour, the global result to which the dynamics of interaction leads us can no longer be interpreted solely in terms of the intrinsic characteristics of the individuals, – what are usually called the “fundamentals”. We must also take into account the dynamics itself. This now plays an essential role that can no longer be passed over in silence. This break with the fundamentalist interpretation which has traditionally enjoyed exclusive dominance in economics is certainly the most interesting characteristic of imitation models, as Föllmer so well observed. We shall now present a first illustration, drawn from the work of Mark Granovetter (1978), enabling us to define our concepts more precisely.

In the article in question, this author studies a group composed of N individuals faced with a binary decision. This is a situation in which the individual’s choice is limited to two alternatives: E or not- E , denoted NE . Granovetter uses the following decision as an example: “to riot or not to riot”. Considering the choice to be a rioter, Granovetter observes that the utility attached to this option is an increasing function of the total number of rioters because “the cost to an individual of joining a riot declines as riot size increases, since the probability of being apprehended is smaller the larger the number involved” (p. 1422). In other words, the higher the

number n_E of rioters in the group, the more the costs fall, thus increasing expected profits. We are in a situation where the reference group for all individuals is the whole group. Under these conditions, we have: $U^i(E) = U^i(u_E^i, n_E)$, a strictly increasing function of n_E where E is the choice “to be a rioter” and where n_E is equal to the total number of rioters in the group. The specific attitude of the individual i towards the choice “to be a rioter” is taken into account through the variable u_E^i . We therefore find ourselves in a situation of preferential imitation. Granovetter says nothing explicitly about the utility of being a non-rioter. For the sake of simplicity, we shall assume that the utility associated with this option has a value independent of the choices of the group¹⁰.

We therefore obtain:

$$U^i(NE) = u_{NE}^i \tag{5.2}$$

In addition, Granovetter proposes the hypothesis that individuals are rational, inasmuch as they choose the option that provides them with the greatest utility. Under these hypotheses, it follows quite naturally that the




Fig. 5.1. Individual utility curve

¹⁰ We could have also assumed the presence of preferential imitation for this second option.

agents display threshold behavior (figure 5.1): for every individual i , there is a number n_i^* of rioters such that if $n_E < n_i^*$, the individual chooses not to riot, whereas if $n_E \geq n_i^*$, he decides to participate in the riot. This is a direct consequence of preferential imitation. We define individual i 's threshold as the proportion of individuals in the group who must have chosen to riot before the individual i will himself choose to riot, i. e. $s_i = n_i^* / N$. "The threshold is simply that point where the perceived benefits to an individual of doing the thing in question (here, joining the riot) exceed the perceived costs" (p. 1422). The work of Granovetter emphasises the dispersion of individual thresholds, expressing wide diversity in individual attitudes towards the riot. He distinguishes between "radical" individuals with a low threshold and "conservatives" who have a high threshold. The instigators, who have a threshold of 0%, riot even when nobody has rioted before. At the other extreme, the "non-violent" individuals, who have a threshold of 100%, will never riot.

In this model, the fundamental data are provided by the set of thresholds $\{s_i\}_{i=1, \dots, N}$. When we know these numbers, we know perfectly the characteristics of all the N members of the group. We can then plot the curve of accumulated distribution: $F(x) = \{\# \text{ of individuals with a threshold below } x\}$. A cruder characterisation of the overall attitude of the group can be obtained by simply considering the average threshold of the individuals constituting the group:

$$\bar{s} = \frac{1}{N} \sum_{i=1}^{i=N} s_i \quad (5.3)$$

By definition, a state of equilibrium is a situation in which no individual is incited to alter his behaviour. It can be characterised by the fact that the percentage x^* of rioters is such that $F(x^*) = x^*$. On a graph, the equilibria are obtained at the intersection of the distribution curve and the first bisector. These equilibria can be multiple. If we assume a normal distribution of thresholds with an average m ¹¹, we obtain figure 5.2, in which three equilibria appear: x_1^* and x_3^* are stable; $x_2^* = m$ is unstable.

To illustrate his analysis, Granovetter first considers a group of 100 individuals¹², with the following distribution of thresholds: "there is one

¹¹ To do so, we consider that individuals with a negative threshold are „instigators” and those with a threshold above 1 are „non-violents”.

¹² As this hypothesis enables us to identify the threshold and number, it makes it easier for us to describe the dynamics of imitation.




Fig. 5.2. Cumulative distribution function of thresholds and equilibria

individual with threshold 0, one with threshold 1, one with threshold 2, and so on up to the last individual with threshold 99” (p. 1424). What happens when we have an initial situation where there are no rioters, i. e. $n_E(t=0)=0$? The individual who has a threshold of 0 will adopt the behaviour of a rioter. He is an instigator. Now, the individual whose threshold is 1 will follow him. Then the behaviour of these two individuals will provoke the reaction of the individual whose threshold is 2, and so on until all 100 individuals have become rioters. Consequently, the final value obtained is $n_E=100$. This is the equilibrium value.

From this moment on, the choice of all the members of the group will remain constant. As a second step, Granovetter very slightly modifies the distribution of thresholds as follows: the individual whose threshold was 1 now has a threshold of 2, while all the other thresholds remain unchanged. If we compare these two configurations in terms of the fundamental data, notably the indicator \bar{s} of the average propensity to riot, we can see that they are almost identical. And yet the behaviour of the group is radically altered. Starting from a situation where $n_E=0$, the individual whose threshold is 0 still decides to riot, but now the riot goes no further, because

all the other members of the group have thresholds strictly superior to 1. The equilibrium value we now observe is $n_E = 1$. Note that in this second configuration, the value $n_E = 100$ is still an equilibrium, in other words a value which, once attained, remains constant.

Here we have a good example of the inadequacy of the fundamentalist diagnosis when imitative dynamics is taken into account. To make this clear, Granovetter considers a hypothetical newspaper reporting on the two situations, namely $n_E = 100$ and $n_E = 1$: “Newspaper reports of the two events would surely be written as, in the first case, “A crowd of radicals engaged in riotous behavior”; in the second, “A demented troublemaker broke a window while a group of solid citizens looked on”. We know, however, that the two crowds are almost identical in composition” (p. 1425). In other words, the newspaper interprets the difference between the two scenes as being the result of differences in the intrinsic propensities of the two groups to riot: a “crowd of radicals” compared with “a group of solid citizens”. However, we know that this diagnosis is totally false. The two groups are composed of the same individuals, except for one unique element. This error is not exclusive to journalists. It is very common and “natural”. Economists and sociologists both show a tendency, when they observe differences between two populations, spontaneously to reduce those differences to the fundamental data characterising the populations. The example proposed by Granovetter demonstrates that in situations of complex interactions – such as those in which imitation plays a role – *dynamics counts*. Depending on the initial situation and the random shocks encountered by the economic system, the final equilibrium may display considerable modifications. Drawing an analogy with financial phenomena, we could use the term “bubble” to designate this disconnection between the fundamental characteristics and the final equilibrium.

To put it another way, using a concept introduced in chapter 2, this first example presents imitation as an economic force that produces self-organisation. A central characteristic of self-organisation is that it gives rise to new qualities, – qualities which are absent from the analysis when only the elements of the system are considered, one by one. For this reason, the traditional fundamentalist description is necessarily inadequate when confronted with this type of phenomenon. A speculative bubble can be analysed as a form of self-organisation of financial markets. We shall now move on to the systematic analysis of imitation.

5.2 Canonical principles

In recent years, the case most studied in economic literature has been that of binary choices. We can, in fact, write the general model for this case. This is the canonical model that we shall be describing in this section. In section 5.3, we shall examine some illustrations of this general model, focusing in succession on informational, normative and preferential imitation.

5.2.1 A general model

Let us consider a population of N agents. Each agent is denoted i , with $i \in \{1, 2, \dots, N\}$. We assume that each agent is faced with a binary choice: he either chooses option [1] or option [2]. We denote $n_1(t)$ the number of individuals who, at time t , have chosen the option [1] and $n_2(t)$, the number who have chosen the option [2]. Thus, we have:

$$n_1(t) + n_2(t) = N \quad (5.4)$$

To discover the distribution of choices within this population, therefore, we only need to know one of these numbers. The other can then easily be deduced. By convention, we shall choose, as the state variable describing the system, the variable $n_1(t)$, which we rename $n(t)$ or the variable

$$f(t) = n(t)/N. \quad \text{The variable denoted } x(t) = \frac{n_1(t) - n_2(t)}{N} = 2f(t) - 1,$$

which measures the proportional deviation between those who have chosen option [1] and those who have chosen [2], may also be chosen as state variable. $x(t)$ varies between -1 and $+1$.

We are exploring situations of imitative interaction, expressed by the fact that each individual's choice between [1] and [2] depends "positively" on the behaviour of the other agents. There are several different ways of integrating this imitative effect. If we assume that the individuals are "spatialised", we must introduce neighbourhoods, in which case each individual is only influenced by his neighbours. We shall not introduce this complexity within the context of this chapter, except to touch on it briefly in the section "Conclusions and conjectures". We shall assume, more simply, that every individual is influenced to the same extent by each of the other members of the group, as in the example proposed by Granovetter. In the most general form of this hypothesis, we write that an individual changes his choice from [1] to [2] with the probability p_{12} and that he changes from [2] to [1] with the probability p_{21} . These probabilities are respec-

tively decreasing and increasing functions of n and therefore of f . We have, therefore:

$$\begin{cases} p_{12}(n): \text{probability of transition from [1] to [2]} \\ p_{21}(n): \text{probability of transition from [2] to [1]} \end{cases} \quad (5.5)$$

We must underline the fact that the truly economic aspect lies in the determination of the probabilities $p_{12}(n)$ and $p_{21}(n)$. For the rest, it is simply a question of mathematically solving the model thus constructed¹³. To do this, however, we must first specify precisely how the sequence of individual choices unfolds. As we shall see, the economic model described by the equations (5.5) can result in the definition of two different dynamics, which we shall denote D_1 and D_2 , depending on the way in which this sequence of individual choices is defined.

5.2.2 The D_1 dynamic

The first dynamic, denoted D_1 , assumes that at each instant t , an individual is drawn at random from the group and chooses his option in accordance with the probabilities (5.5). We can see immediately that we are dealing with a first-order Markovian stochastic process. It is stochastic because if, at instant t , the state variable equals $n(t) = n$, then at $t + 1$, $n(t + 1)$ may equal $n + 1$, n , or $n - 1$. It is a first-order Markovian process because the transition probabilities at instant t only depend on the value of $n(t)$. For example, to change from n to $n + 1$, the individual chosen at random must have previously chosen [2], an event of which the probability is $\frac{N - n}{N} = 1 - f$, and this individual must now choose [1], i. e. an event of which the probability is $p_{21}(n)$. We calculate the probability of changing from n to $n - 1$ in the same way. We thus obtain:

$$\begin{cases} \text{Prob}(n \rightarrow n + 1) = (1 - \frac{n}{N}) p_{21}(n) = w_+(n) \\ \text{Prob}(n \rightarrow n - 1) = \frac{n}{N} p_{12}(n) = w_-(n) \end{cases} \quad (5.6)$$

¹³ Note that the solution proposed is independent of the form of the probabilities of transition. It can therefore be used even in non-imitative cases, i.e. in situations where the probability of choosing an option is not an increasing function of the numbers of individuals who have already chosen the same option.

It follows that the correct way to describe the system at instant t is not in terms of its state variable $n(t)$ but of its law of probability, which we denote $p(n;t)$. Following what has just been said, we can then write the law of evolution of this probability. To do so, we write that between t and $t+1$, the “weight of probability” in n is increased by the flow coming from $n-1$ and $n+1$, but is decreased by the flow leaving n for $n-1$ and $n+1$. The calculation of these inputs and outputs immediately leads to the master equation:

$$\frac{dp(n;t)}{dt} = w_+(n-1)p(n-1;t) + w_-(n+1)p(n+1;t) - w_+(n)p(n;t) - w_-(n)p(n;t) \quad (5.7)$$

5.2.2.1 Solution of the model

Our aim is to analyze the asymptotic behavior of this process when t tends to infinity. In other words, we want to determine the stationary distribution(s) of the process (5.6). When there is one and only one stationary distribution, the process is said to be ergodic. Following the study of Markovian processes, we can say that the process (5.6) is ergodic when all the transition probabilities $p_{21}(n)$ and $p_{12}(n)$ are strictly positive for all values of n , except $p_{12}(N)$ and $p_{12}(0)$ which have no economic significance. This is the situation we are going to study in more detail in the present subsection. We shall show that when the positivity of transition probabilities is verified, we can calculate explicitly the asymptotic distribution of the process (5.6) (Weidlich and Haag, 1983).

To do so, let us first observe that the stationary probability, denoted $p_s(n)$, for n varying between 0 and N , verifies the following equation:

$$\frac{dp_s(n;t)}{dt} = 0 \quad (5.8)$$

If we take into account the limit conditions, namely

$$w_-(0) = w_+(N) = 0 = w_-(N+1) = w_+(-1) \quad (5.9)$$

and then use equation (5.7), we find that:

$$p_s(n) = p_s(0) \prod_{k=1}^{k=n} \frac{w_+(k-1)}{w_-(k)} \text{ for } 1 \leq n \leq N \quad (5.10)$$

This calculation is possible because the positivity of the transition probabilities entails the positivity of $w_-(k)$, for all $k > 0$. If we take into account the fact that the sum of probabilities $\sum p_s(n)$ is equal to 1, we can

easily deduce the value of $p_s(0)$, then the value of $p_s(n)$. We can then calculate the stationary probability exactly. When the process is not ergodic, these calculations are no longer valid. The stationary probability is no longer unique.

The analysis of this stochastic process can be approached in another way. This involves re-writing the equations, using the variable f instead of the state variable n to which it is equivalent, and then presenting a continuous approximation on the former variable. In fact, for large N , it is possible to consider that f behaves as a real variable belonging to the segment $[0,1]$. We can then rewrite the master equation (5.7) in the following form:

$$\frac{\partial P(f;t)}{\partial t} = -\frac{\partial}{\partial f} [K(f)P(f;t)] + \frac{\varepsilon}{2} \frac{\partial^2}{\partial f^2} [Q(f)P(f;t)] \tag{5.11}$$

with:

$$\begin{cases} W_+(f) = (1-f) \cdot p_{21}(f) \\ W_-(f) = f \cdot p_{12}(f) \end{cases} \tag{5.12}$$

$$\begin{cases} K(f) = W_+(f) - W_-(f) \\ Q(f) = W_+(f) + W_-(f) \end{cases} \text{ and } \varepsilon = \frac{1}{N} \tag{5.13}$$

Equation (5.11) is a classic diffusion equation, known as the Fokker-Planck equation. $K(f)$ represents the drift coefficient and $Q(f)$ represents the fluctuation coefficient. The significance of this equation can be grasped intuitively. When $K(f)$ is positive, this means that the probability $W_+(f)$ is greater than the probability $W_-(f)$. f therefore tends to increase. $Q(f)$ measures the overall noise generated by the system even if the pluses and the minuses cancel out. This intuition can be confirmed by demonstrating that the stochastic process described by the Fokker-Planck equation (5.11) is equivalent to the following stochastic differential equation:

$$df(t) = K[f(t)]dt + \sqrt{\varepsilon Q[f(t)]}dW(t) \tag{5.14}$$

5.2.2.2 Properties of the stationary probability

In this sub-section we shall continue to assume that the process is ergodic, so that we can refer without ambiguity to the stationary distribution and continue to examine the continuous approximation given by equation (5.11). $P_{st}(f)$ denotes the stationary probability. In this context, what most interest us are the properties of the stationary probability, – and par-

ticularly its extrema. If asymptotic variance is low, we know that f will be situated somewhere in the neighborhood of a maximum of this probability. Bearing in mind that the stationary probability satisfies the following property:

$$\frac{\partial P(f, t)}{\partial t} = 0 \quad (5.15)$$

following equation (5.11) and the limit conditions, we have:

$$K(f)P_{st}(f) = \frac{\varepsilon}{2} \frac{d}{df} [Q(f)P_{st}(f)] = \frac{\varepsilon}{2} P_{st}(f) \frac{d}{df} Q(f) + \frac{\varepsilon}{2} Q(f) \frac{d}{df} P_{st}(f) \quad (5.16)$$

As the extrema f_m of the stationary probability are such that:

$$\frac{dP_{st}}{df}(f_m) = 0 \quad (5.17)$$

with:

$$\begin{cases} \frac{d^2 P_{st}}{df^2}(f_m) < 0 & \text{for maximum } f_m \\ \frac{d^2 P_{st}}{df^2}(f_m) > 0 & \text{for minimum } f_m \end{cases} \quad (5.18)$$

it therefore follows, when the terms $\varepsilon Q'(f_m)$ and $\varepsilon Q''(f_m)$ can be ignored¹⁴ in relation to the drift terms (this is true for $\varepsilon \ll 1$), that the extrema are such that:

$$K(f_m) = 0 \quad (5.19)$$

with:

$$\begin{cases} \frac{dK}{df}(f_m) < 0 & \text{for maximum } f_m \\ \frac{dK}{df}(f_m) > 0 & \text{for minimum } f_m \end{cases} \quad (5.20)$$

¹⁴ This is not always the case. We can cite two models in the literature in which this approximation cannot be made (Kirman, 1993) and (Orléan, 1990).

Conditions (5.19) and (5.20) are interesting because they enable us to determine the overall form of the stationary probability without having to perform the rigorous calculation (5.10) of the stationary distribution, which can be complicated.

The intuition behind this result is that if we ignore the fluctuations, then the deterministic movement associated with the stochastic process under consideration is of the form:

$$\frac{df}{dt} = K[f(t)] \quad (5.21)$$

This result can be deduced immediately from equation (5.14). The fixed points in this process are the extrema of the stationary probability. These are maxima if the fixed points are stable; they are minima if the fixed points are unstable.

5.2.2.3 An important special case: the \mathcal{I} hypothesis

One special case is particularly important: that in which the probability of choosing option [1] or option [2] is independent of the previous choice made. In the text that follows, we shall refer to this as the \mathcal{I} hypothesis. When the \mathcal{I} hypothesis is adopted, we have:

$$p_{21}(n) = p_1(n) \quad \text{and} \quad p_{12}(n) = p_2(n) = 1 - p_1(n) \quad (5.22)$$

The only thing that counts is the probability, at a given moment, of choosing [1] or [2]. By writing $P(f) = p_1(n)$, we can deduce:

$$K(f) = (1 - f)P(f) - f[1 - P(f)] = -f + P(f) \quad (5.23)$$

Consequently, the extrema are such that:

$$P(f_m) = f_m \quad (5.24)$$

We obtain the equilibrium states of Granovetter's model, for if we were to adapt that model to the present context, $F(x)$ would be the probability of an individual participating in a riot when a proportion x of individuals have already chosen to riot.

5.2.3 The D_2 dynamic

Using the same economic model (5.5), we can describe a second dynamic context in which, at the instant $t + 1$, all the individuals make their choices

simultaneously, instead of the one individual drawn at random that we had in the D_1 dynamic. We shall call this the D_2 dynamic. Applying the law of large numbers, it follows that f_{t+1} – the proportion of individuals that choose option [1] at $t+1$ – is equal to the probability of choosing this option when, at $t+1$, all the agents observe the same value f_t when making their choice. The law of large numbers applies because each agent chooses independently in relation to the random event in question. The fact that every agent uses f_t when making his choice has no incidence on this result. By using transition probabilities, it follows that:

$$f_{t+1} - f_t = (1 - f_t)P_{21}(f_t) - f_t P_{12}(f_t) = K(f_t) \quad (5.25)$$

This modelisation is used by A. Corcos, J.P. Eckmann and A. Malapinas (1998). We can immediately verify that the approximation in continuous time of this discrete process is written:

$$\frac{df}{dt} = K[f(t)] \quad (5.26)$$

So we obtain once again the deterministic dynamics associated with the D_1 dynamic. It is not difficult to see that the stable fixed points of the dynamics (5.25) are the maxima (5.19) of the stationary probability of the D_1 dynamic.

Furthermore, if we add the hypothesis referred to above as the \mathcal{I} hypothesis, according to which the choice at any given instant is independent of the previous choice, we can write:

$$P(f) = P_{21}(f) = 1 - P_{12}(f) \quad (5.27)$$

We then obtain:

$$f_{t+1} = P(f_t) \quad (5.28)$$

To sum up, one can always choose between one of two equivalent dynamic specifications for the economic model (5.5): either the individuals choose one by one, which we have called the D_1 dynamic, or they all choose at the same time, which we have called the D_2 dynamic. In the former case, we obtain convergence towards a stationary probability; the latter converges towards stationary states. If there is ergodicity, the formal equivalence between the two is expressed by the fact that the maxima of the stationary probability are fixed points of the deterministic process. The observed values of f are very close in both contexts.

5.3 Some models

We shall now use the general model described above, together with its two dynamic specifications D_1 and D_2 , to study imitative interactions. This part of the chapter is divided into three sub-sections, each dealing with one of the three forms of imitation. From an economic point of view, the essential point is to write the transition probabilities (5.5), then to use our general results concerning the solution of the model to analyse the results and propose an economic interpretation of them.

5.3.1 Informational imitation

To analyse this form of imitation, we shall consider a population constituted of N individuals, denoted i with $i \in \{1, 2, \dots, N\}$. Let θ denote the exogenous state of nature, or the state of the world. This is equal to either $\{H\}$ or $\{L\}$ and is assumed to be fixed. The task of each individual is to discover the value of this state. Thus, each individual has two options: option [1] consists in declaring that the state is $\{H\}$ and option [2] that it is $\{L\}$. The context is therefore one of binary choice, in accordance with the general model presented in the previous section.

The difficulty of the task confronting the individuals derives from the fact that they cannot observe the state of nature directly. To discover the state, they have access to a random noisy signal σ , equal to either $\{+\}$ or $\{-\}$, the value of which is connected to the state of the world through the following conditional probabilities¹⁵:

$$\begin{cases} P(\sigma = + | \theta = H) = P(\sigma = - | \theta = L) = p > 0.5 \\ P(\sigma = - | \theta = H) = P(\sigma = + | \theta = L) = 1 - p < 0.5 \end{cases} \quad (5.29)$$

In other words, the closer p gets to 1, the more precise the signal, and the higher the probability of discovering the correct state of the world. At the limit, for $p = 1$, observation of the signal is sufficient to determine the state of the world without any ambiguity. We assume that each agent i observes an independent draw of σ , denoted σ^i . In addition, each individual knows that the prior probability of the two states is $1/2$.

¹⁵ This analytical framework was proposed by Bikhchandani, Hirshleifer and Welch (1992). However, their analysis considered a succession of agents entering the market sequentially, whereas in our work all the agents are present simultaneously.

5.3.1.1 Choices in the absence of imitation

To start with, we consider the situation in which each agent makes his choice solely on the basis of his private information σ^i . In this case, to determine his choice, the agent calculates the probability conditional on the messages: $P(H/\sigma^i)$ and $P(L/\sigma^i)$:

$$P(H | \sigma = +) = P(H | +) = \frac{P(+ | H)P(H)}{P(+)} \quad (5.30)$$

with

$$P(+)=P(+ | H)P(H)+P(+ | L)P(L) \quad (5.31)$$

Bearing in mind the fact that the states of the world are *a priori* equiprobable, we obtain:

$$P(H | +) = p \quad \text{and} \quad P(H | -) = 1 - p \quad (5.32)$$

On the basis of this calculation, the agent chooses the state with the greater probability, i. e. a probability higher than $1/2$. This result is justified by the fact that we assume a symmetric utility function, in other words the utility obtained from choosing H when the state of the world actually turns out to be H , which we denote $U(H | H)$, is equal to the utility $U(L | L)$ obtained from choosing L when the state of the world actually is L . We therefore have $U(H | H) = U(L | L) = a$. Likewise, we assume that $U(L | H) = U(H | L) = b < a$. As $p > 1/2$, it follows that the individual who observes $\{+\}$ will choose H and the one who observes $\{-\}$ will choose L , as one would intuitively expect. Because the signal is noisy, however, this rule can lead to an incorrect choice. This is the case when the state of the world is $\{H\}$ and, through bad luck, the individual observes the signal $\{-\}$, which leads him to choose L . This only occurs, however, with the probability $1 - p$. The more precise the signal, the lower the probability of error.

Using this decision rule, we can establish the transition probabilities corresponding to the equations (5.5) of the general model. We have:

$$\begin{cases} p_{21} = \text{the probability of observing } \{+\} = p \\ p_{12} = \text{the probability of observing } \{-\} = 1 - p \end{cases} \quad (5.33)$$

Note that these probabilities do not depend on n , because the individuals are independent and not imitative. Let us now consider the D_1 dynamic, in which all agents behave in the same way. What happens? To find out, we assume that the state is $\{H\}$. The options are either to choose H , option

[1], or to choose L , option [2]. Let n denote the number of individuals who have chosen H . We can then use equation (5.10) to calculate the stationary probability:

$$p_s(n) = p_s(0) \prod_{k=1}^{k=n} \frac{(N-k+1)p}{k(1-p)} = p_s(0) \binom{N}{n} \left(\frac{p}{1-p} \right)^n \quad (5.34)$$

As we want $\sum p_s(n) = 1$, we have:

$$p_s(0) + p_s(0) \sum_{n=1}^{n=N} \binom{N}{n} \left(\frac{p}{1-p} \right)^n = p_s(0) \left(1 + \frac{p}{1-p} \right)^N = 1 \quad (5.35)$$

from which follows:

$$p_s(0) = (1-p)^N \quad (5.36)$$

and consequently:

$$p_s(n) = \binom{N}{n} (p)^n (1-p)^{N-n} \text{ for all } n \quad (5.37)$$

Here we recognise the binomial law $B(p, N)$ of the parameters p and N . This was to be expected, because when $\{\theta = H\}$, the agents independently choose H with probability p or L with probability $1-p$, which is the very definition of the binomial law. It can easily be demonstrated that the sequential process in which the agents choose one after another converges towards the same law. The binomial distribution has a unique mode for $n_m = pN$, or again $f_m = p$, in accordance with equation (5.24), in which the probability of choosing option [1] is a constant function, written $P(f) = p$. In a configuration such as this, a proportion $(1-p)$ of individuals in the group make the wrong choice. The same result is obtained when the state of the world is $\{L\}$. So, whatever the state of the world θ , we have a proportion $1-p$ of individuals in the population who, on average, make the wrong choice. Now we shall see whether the introduction of imitation leads to any improvement in this result.

5.3.1.2 The introduction of imitation

Intuitively, it seems reasonable to believe that the above result can be improved when the agents are no longer limited exclusively to their private information. Essentially, we shall examine situations in which agents can observe the actions of others, but have no access to their private information.

This hypothesis corresponds to a market configuration. In a market, agents do not know the information possessed by the others, but they can observe them buying or selling the product in question. More specifically, we only assume that agents observe the past price, which reveals part of the information possessed by private agents. This is the specific hypothesis that we shall adopt. In our context, it is the variable f , the collective choice, which plays the role of price. We assume that this value is announced publicly at each instant. It is obvious that the aggregate information given by the variable f is decisive and can help the private agents to improve their assessment of θ . To demonstrate this, we shall consider the situation presented above, in which individual choices are independent.

It appears that an *external* observer can discover the correct state of the world solely on the basis of his observation of the collective opinion f . This can be understood simply by observing that when N is large, the state chosen by the majority of the group, namely H when $\{f > 1/2\}$ or L when $\{f < 1/2\}$, has a very high probability of indicating the real state of the world. The probability that f is lower than $1/2$ when f is a random binomial variable of the parameters p and N , with p greater than $1/2$, rapidly tends to 0 when N tends to infinity. As an example, for $N=50$ and $p=0.7$, this probability is equal to 0.0013, very much lower than the probability of error obtained solely from observation of σ , which is equal to 30% in this case. Under these conditions, the pure imitative rule that consists in copying the majority choice of the group has a very high probability of leading to the correct result. It therefore seems reasonable to believe that the observation of f by agents can increase their ability to discover the real state of the world. The imitation introduced here is informational: the individual imitates because he thinks that the information of the group is of a higher quality than his own.

Expressed slightly differently, this analysis tells us that collective opinion, in the case of independent choices, is “informationally efficient”, inasmuch as, for large enough N , it reveals the underlying state of the world. Imitation of the majority choice then stands out as an efficient option. This is true for the external observer, but the members of the group cannot benefit from it because they do not use this information. There is still a proportion $1-p$ who make the wrong choice. The issue now is to analyze whether imitation of the majority choice can also improve the individual performances of the members of the group. In other words, what happens when the agents are aware of the informational efficiency of the collective opinion? Does f remain informationally efficient when the agents themselves seek to exploit it to refine their assessments?

Presented in this way, the problem takes on a form that will be recognised by researchers in the field of finance. It is an expression of the famous paradox of Grossman and Stiglitz (1980), who examined one of the most fundamental properties of financial market theory: the informational efficiency of prices, which states that “the price reveals perfectly all the information available in the market” (Fama, 1970, p. 383). Grossman and Stiglitz made the following observation: if investors know that the price is informationally efficient, then all incentive to seek information will disappear, because simple observation of the price will enable an investor to do just as well as those who seek information, but without sustaining any information costs. However, if every investor reasons in this way, then informational efficiency will disappear because no one will seek information, and consequently the price will no longer contain any information. The authors conclude that informational efficiency is impossible. Expressed in our terms, this result affirms that if every agent ignores his private information and simply imitates the majority choice – which is the rational thing to do if the majority choice is informationally efficient – then this majority choice will cease to be efficient. It could be anything. Note that Grossman’s and Stiglitz’s result is strictly negative. It expresses impossibility, without commenting on what might actually happen in such a situation. Our model will enable us to examine the same question in a dynamic context.

For this purpose, we introduce a new fact into the previous model: agents can now observe the value of f . Under these conditions, the information set of the individual i now has two elements: his private information and the collective opinion. It is written $\{\sigma^i, f\}$. The individual is now faced with one of two possible situations. In the first, the two elements of information are compatible, in that they both indicate the same choice. This is the case when $\{\sigma = +\}$ and $\{f > 1/2\}$ and also when $\{\sigma = -\}$ and $\{f < 1/2\}$. In this situation, the individual has no problem in making his choice. In the second situation, the two elements of information are contradictory. Take, for example, the situation in which $\{\sigma = +\}$ and $\{f < 1/2\}$. What will the individual do in such a situation? That depends on the way he assesses the relative precision of the two elements of information. If he believes in the theory of efficiency, he will ignore his individual information and follow the majority choice, in this case L . If, on the contrary, he has little confidence in the group, he will follow his private signal and choose H . We shall adopt a probabilist strategy to take this ambiguity into account: the individual imitates the majority choice with a probability of μ ; he follows his private information with a probability of

$1 - \mu$. μ is an exogenous, fixed parameter which measures the strength of the propensity to imitate. It varies between 0 and 1. We denote $q_\mu(\sigma, f)$ the probability that an individual will choose H when he is characterised by the imitation parameter μ and he has observed σ and f . We can then write:

$$\left\{ \begin{array}{ll} \text{if } f < 0.5 & q_\mu(+, f) = 1 - \mu \\ \text{if } f > 0.5 & q_\mu(-, f) = 1 \end{array} \right. \text{ and } \left\{ \begin{array}{ll} \text{if } f < 0.5 & q_\mu(-, f) = 0 \\ \text{if } f > 0.5 & q_\mu(+, f) = \mu \end{array} \right. \quad (5.38)$$

Such is the decision rule of an individual in the group. If we assume that all the individuals in the group are characterised by the same parameter value μ and if we adopt a D_1 dynamic, we can then determine the stationary distribution. The interest of this presentation lies in the comparative static analysis of how this distribution varies with μ . In this way we can directly evaluate the influence of informational imitation on the collective opinion. Does more imitation lead to more efficiency, or the opposite?

5.3.1.3 Solution of the model

Let us consider a D_1 dynamic. At the instant $t=0$, we assume that $f(t=0)$ takes a value, any value, denoted f_0 ¹⁶. At each instant $t > 0$, an individual is drawn at random. He observes the value f_{t-1} announced publicly at the end of the last period and acquires information about the state of nature for the period t through observation of σ_t^i , which is drawn at random. He then chooses between H and L , in accordance with the probabilities (5.38). Here, we postulate that his decision is independent of his previous choices. It depends exclusively on the two elements of information $\sigma_t^i = \sigma$ and $f_{t-1} = f$. We are therefore in the context of what we have called the \mathcal{I} hypothesis. Solution of the problem is thus reduced to calculation of the transition probability $p_{21}(f_{t-1}) = P(f)$, in other words the probability of choosing H , which we shall denote P_H . This depends on the exogenous state of the world θ . By taking certain liberties with strict mathematical expression, we can write this as:

$$P(f) = P_H(f, \theta) = \alpha(\theta) \cdot q_\mu(+, f) + [1 - \alpha(\theta)] q_\mu(-, f) \quad (5.39)$$

where $\alpha(\theta)$ denotes the probability of drawing ($\sigma = +$) when the state of the world is θ . We therefore have $\alpha(H) = p$ and $\alpha(L) = 1 - p$. It is easy to

¹⁶ We could even assume that at the instant $t=0$, the value of f is drawn from any probability distribution whatever: $P(f, 0)$.

see that this transition probability lies strictly between 0 and 1 when μ is strictly lower than 1. If this is so, it is also true for the transition probability $p_{12}(f) = 1 - P(f)$. Therefore if $\mu < 1$, we are sure that the process is ergodic. Equation (5.10) gives us the exact form of the stationary distribution.

We can also place the problem within the context of continuous approximation (5.11). This makes it easier to discuss the form of the stationary distribution, subsequently returning to equation (5.10) to verify the correctness of our approximation. To analyse the properties of the stationary probability, we shall examine the case in which ($\theta = H$); the case $\theta = L$ can then be immediately deduced by symmetry. We obtain:

$$P(f) = P_H(f, H) = p \cdot q_\mu(+, f) + (1 - p) \cdot q_\mu(-, f) \quad (5.40)$$

The extrema are given by equation (5.24). We can then use the conditions (5.20) to determine whether these are minima or maxima. The simple form (5.38) of the functions $q_\mu(\sigma, f)$ provides us with a simple solution of these equations. We obtain the following results.

When the value of the parameter μ remains lower or equal to a certain value $\mu^* = \frac{2p-1}{2p}$, the equation $P(f) = f$ only has one solution, which we shall denote f_i . We can then demonstrate that the stationary distribution is unimodal (figure 5.3), with mode f_i .

More precisely, when we assume that ($\theta = H$), we obtain:

$$f_i(\mu, H) = p + (1 - p)\mu \quad \text{when } 0 \leq \mu \leq \mu^* \quad (5.41)$$

Note that when $\mu = 0$, we find ourselves back with the previous case, without imitation. It appears that imitation is efficient inasmuch as it improves the collective result, because the proportion f_i of agents finding the correct value of θ becomes higher than p . This result is easy to interpret. Without imitation, only the $p\%$ of individuals who had, by pure chance, received the signal $\{+\}$ made the correct choice, while those who observed the signal $\{-\}$ made the wrong choice. Now, because of imitation, this is no longer the case. A proportion μ of the $(1 - p)$ economic agents who are unlucky enough to observe the signal $\{-\}$ nevertheless make the right choice, in this case H , because they copy the majority choice. Under these conditions, the higher the value of μ , the larger the proportion of agents making the correct choice H . However, for this to be the case, the majority choice itself must be correct. This is not necessarily so when μ is too large, as we shall now see.

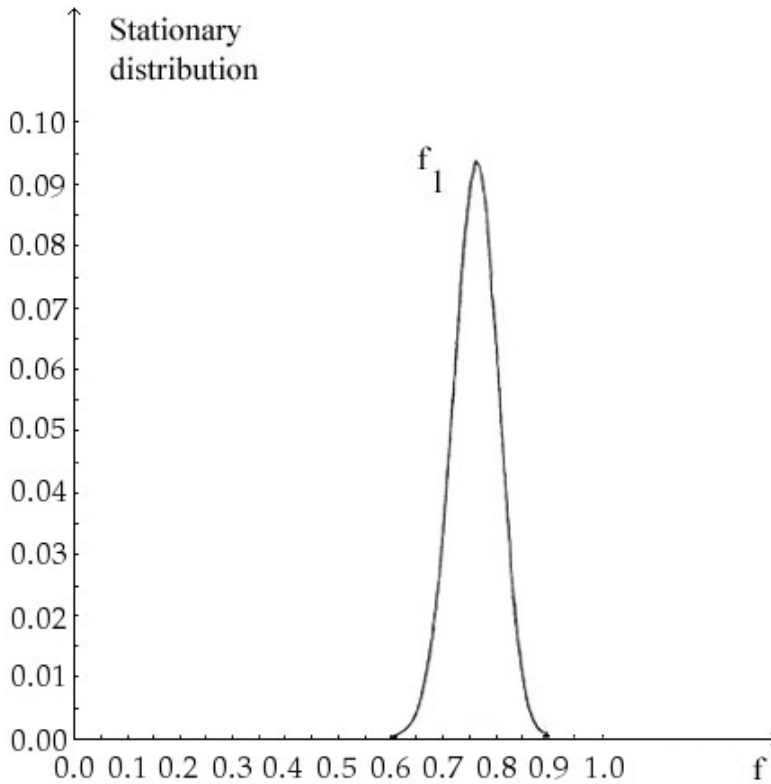


Fig. 5.3. Stationary distribution with $\mu = 0.2$ and $p = 0.7$ when $\theta = \{H\}$

When μ is greater than μ^* but still lower than 1, the form of the stationary distribution changes: it becomes bimodal (figure 5.4).

The equation $P(f) = f$ now has two solutions¹⁷, which we shall denote f_M and f_I , such that:

$$f_I = p + (1 - p)\mu \text{ and } f_M = p(1 - \mu) \tag{5.42}$$

This situation expresses a profound change in the economic properties. The propensity to imitate has now become so strong that it dominates private information. In other words, if the majority opinion is in favour of option [L], although many individuals observe the information $\{+\}$ because

¹⁷ This result may surprise readers acquainted with dynamic systems, who would expect to see three solutions. The reason is the discontinuity of the function $P(f)$ at $1/2$. Otherwise, there would indeed be three solutions.




Fig. 5.4. Stationary distribution with $\mu=0.8$ and $p=0.7$ when $\theta = \{H\}$

($\theta = H$), their propensity to imitate is so strong that they will, on average, continue to prefer the existing majority. Under these conditions, the proportion f of choices in favour of H remains¹⁸ lower than $1/2$, despite a large flow of information supporting option $[H]$. To sum up, if μ is too high, the information-spreading process is dominated by the process of self-validation of the existing majority, whatever that may be. For this reason, we can observe two modes, one above $1/2$ which we have called f_I and the other below $1/2$, which we have called f_M . Under these conditions, imitation is no longer efficient. It starts generating bubbles, in other

¹⁸ Of course, we may observe a transition from one mode to the other if, by chance, a large number of agents observe $\{+\}$ signals and do not imitate, but this event is so rare that its probability is very low. Practically, for large enough N and for reasonable lengths of time, it is never observed.

words situations in which the majority opinion is disconnected from the objective state of nature: the state of nature may be [H] and agents may receive numerous information signals $\{+\}$, majority opinion will nevertheless decide that option [L] is the right choice.

5.3.1.4 The ambivalence of informational imitation

For a deeper understanding of these results, we must shift the focus of our attention away from the stationary distribution and onto the dynamic that follows the variable f over the passage of time when we introduce exogenous changes into the state of nature θ . Figure 5.5 illustrates a situation in which $f_0 = 0.5$, μ is lower than μ^* and the state θ changes in value exogenously every 1000 periods. We can observe that the majority opinion always converges towards the correct choice. It changes from $f_i(\mu, H)$ to $f_i(\mu, L)$. For example, when $\mu = 0.2$ and $p = 0.7$, f moves from the neighborhood of $f_i(0.2, H) = 0.76$ to that of $f_i(0.2, L) = 0.24$. In both cases, the percentage of error is equal to 24%: thus below the 30% observed when the individuals do not imitate. We can therefore say that this imitative system is efficient. It faithfully follows the variations in fundamental data, in this case θ . It should be noted that in this configuration, the value of f_0 has no impact on the equilibrium values.




Fig. 5.5. Evolution of f when $\mu = 0.2$ and $p = 0.7$

When μ is greater than μ^* , while still not exceeding 1, the operating logic of the economic system is radically transformed, as figure 5.6 illustrates. We have chosen the same parameter values as those used in figure 5.5. It can be seen that the majority opinion is disconnected from the value of θ . During the first “cycle” $[0,1000]$, we assume that $\theta = H$. Here, the majority opinion converges to the neighborhood of $f_I(0.8, H) = 0.94$. Now, only 6% of individuals make the wrong choice, so it looks as if the system performs much better than before. In fact, this is not at all the case. This result is obtained through sheer chance. The majority opinion could just as easily have converged to the neighborhood of $f_M(0.8, H) = 0.14$. This becomes much apparent during the next “cycle” $[1001,2000]$, when θ is equal to $[L]$. The majority opinion remains fixed on H . It converges to the neighborhood of $f_M(0.8, H) = 0.14$, which corresponds to 86% error. This happens in spite of the large flow of $\{-\}$ signals received by the agents. As each agent individually attaches great importance to the majority opinion he observes, namely f , this private information is not expressed in f . This corresponds to what Bickchandani, Hirshleider and Welch call an “informational cascade”. The system is locked. The price no longer reveals private information. As figure 5.6 demonstrates, the majority opinion no longer moves, no matter what changes θ undergoes. The system has lost its efficiency and become rigid. It can be observed that on average the error




Fig. 5.6. Evolution of f when $\mu = 0.8$ and $p = 0.7$

is greater than $(1-p)$. This is a general result for $\mu > \mu^*$. It would have been better for the agents to rely only on their private information.

Nevertheless, it must be noted that the ergodicity of the process does indeed mean, in the case of bimodal distribution, that the value f ends up by moving from the neighborhood of one mode to the neighborhood of the other. The “position of f at infinity” is independent of the initial value f_0 . However, the transition probability is very weak. It decreases in e^{-N} . Over reasonable lengths of time, it cannot be observed. Everything happens as if there were two distinct equilibria. Under these conditions, to all practical purposes, the equilibrium attained does depend on the initial value f_0 . This becomes rigorously true when $\mu = 1$. Under this hypothesis, individuals become totally imitative. They no longer take their private information into account, but content themselves with following the majority opinion. In this situation, it is clear that the system converges towards the majority opinion indicated by the value f_0 . The process is no longer ergodic. It has two stationary distributions, the two Dirac distributions in $(f=0)$ and $(f=1)$, i. e. δ_0 and δ_1 .

To sum up, this analysis shows us that imitation can be either positive or negative depending on whether the group’s overall propensity to imitate is weak or strong, what we shall call the “ambivalence of informational imitation”. This is an intuitive result: “It is rational for me to imitate others as long as they are better informed than I. Imitating individuals who are themselves imitators, on the other hand, is inefficient.” The best collective performance is obtained for $\mu = \mu^*$. This does not correspond to efficiency, however, because numerous agents continue to make the wrong choice.

However, we must now take this line of reasoning further by abandoning the unrealistic hypothesis of an exogenous μ . We must integrate the concept of imitation chosen as a rational reaction to the behavior of the group. What happens now? If the individual A knows that μ of the group is lower than μ^* , then he knows that the majority opinion has a very high probability of being correct. Under these conditions, his best course is to choose a μ equal to 1. He becomes totally imitative. If, on the other hand, μ is greater than μ^* , he knows that by following the majority opinion he has a probability lower than p of being correct. He will therefore choose a μ equal to 0, a rule which gives the correct choice with the probability p . He totally abandons imitation and only uses his private information. In other words, the individual A chooses to be imitative when the others are not, and abandons imitation if the others are imitative. However, it is impossible to do the opposite to what the others are doing when all agents are acting simultaneously and in the same way. There is no equilibrium but complex dynamics depending on how agents react to disequilibrium. A modelisation can be found in Orléan (1998).

5.3.2 Normative imitation

Weidlich and Haag (1983) present an approach aimed at modeling the way in which members of a society divide up between two fundamental political choices, of the type “right” versus “left”. The authors adopt the hypothesis that each individual’s opinion is the result of a deliberation involving not only his individual preferences, but also his “desire to adapt to the prevailing opinion, desire which increases in strength in proportion to the predominance of this opinion” (p. 41). They distinguish between “liberal” societies, in which individual choice is free, and “totalitarian” societies which exert strong pressure on deviants. This is a clear example of what we have called “normative imitation”. More precisely, if we denote the two political opinions under consideration [1] and [2], the authors hypothesise that the dynamic of individual choices is described by the following transition probabilities:

$$\begin{cases} p_{21}(f) = g(f) \exp\left[h + k\left(f - \frac{1}{2}\right)\right] \\ p_{12}(f) = g(f) \exp\left\{-\left[h + k\left(f - \frac{1}{2}\right)\right]\right\} \end{cases} \quad (5.43)$$

in which, as throughout this chapter, the variable f represents the proportion of individuals having chosen opinion [1]. Once again, therefore, we obtain our general model (5.5) of binary choice.

The parameter h describes the objective weight of the two opinions within the population under consideration. $h > 0$ signifies that opinion [1] is intrinsically preferred; $h < 0$ signifies that opinion [2] is preferred. The parameter k is always positive. It measures the extent of normative imitation, i. e. the strength of the propensity to follow the “majority choice”: if $f > 1/2$, the higher the value of k , the higher the probability of choosing opinion [1]. The function $g(f) > 0$ determines the frequency of changes. This is a variable of reactivity. Weidlich and Haag assume it to be constant. It can be seen that the qualitative properties of the stationary distribution remain valid with $g(f) > 0$. The exponential form of these probabilities is drawn directly from physics, from what is called the “Ising model”, but, as we have seen in chapter 1, this type of specification can be introduced quite naturally using the logit function.

5.3.2.1 Properties of the stationary distribution

As these probabilities are always strictly positive, we can be sure of the ergodicity of the stochastic process associated with the D_1 dynamic. We can also calculate the stationary probability, by means of equation (5.10). Let us assume $N = 2V$ and denote l such that $n = V + 1$. We therefore have l which varies between $-V$ and $+V$. Once the calculations have been carried out, we obtain, for all values of l :

$$p_s(l) = p_s(0) \frac{(V!)^2}{(N!)} \binom{N}{V+l} \exp\left[2hl + \frac{k}{V}l^2\right] \tag{5.44}$$

By using the fact that the sum of the probabilities is equal to 1, we can determine $p_s(0)$, and then the values of $p_s(l)$. Note that if h is 0, then this distribution is symmetrical: $p_s(l) = p_s(-l)$.

It is easier to analyse the form of the stationary distribution by adopting the context of continuous approximation (5.11). Using the probabilities given by the equations (5.43), we can demonstrate that the functions $K(f)$ of drift and $Q(f)$ of fluctuation are written:

$$\begin{cases} K(f) = g(f) \left\{ \sinh\left[h + k\left(f - \frac{1}{2}\right)\right] - (2f - 1) \cosh\left[h + k\left(f - \frac{1}{2}\right)\right] \right\} \\ Q(f) = g(f) \left\{ \cosh\left[h + k\left(f - \frac{1}{2}\right)\right] - (2f - 1) \sinh\left[h + k\left(f - \frac{1}{2}\right)\right] \right\} \end{cases} \tag{5.45}$$

We can then calculate the extrema of the stationary probability using (5.19), i. e. $K(f_m) = 0$. After (5.45), simplifying matters by taking as our variable either $x = 2f - 1$ or $x = l/V$, where x varies between -1 and $+1$, we obtain:

$$x_m = \tanh(h + k'x_m) \quad \text{with} \quad k' = \frac{k}{2} \tag{5.46}$$

This is a transcendental equation.

To start with, let us take the case $h = 0$. As figure 5.7 illustrates, if the gradient of the tangent at the origin, namely k' , is lower than the gradient of the first bisector, then this equation only has one solution, $x_m = 0$, which is a maximum. If, on the other hand, the gradient of the tangent at the origin is greater than 1, then there are three solutions: $x_m = 0$ and $\pm v$. 0 becomes a minimum, and $\pm v$ become two maxima.

The results are of the same type when h is different from 0. For $k' < 1$, only one solution exists. This solution is a maximum of the stationary distri-




Fig. 5.7. Fixed points

bution. For $k' > 1$ and for $|h| > h_0$, there is still only one solution. Finally, for $k' > 1$ and for $|h| < h_0$, there are three solutions, denoted x_0 , x_+ and x_- . To find the value of h_0 , note that this critical case corresponds to the fact that the line $y = -\frac{h}{k'} + \frac{1}{k'}x$ is tangent to the curve $\tanh x$. We can then demonstrate that the value of h is determined by the following equation:

$$\cosh^2 \left[h_0 - \sqrt{k'(k'-1)} \right] = k' \tag{5.47}$$

For an economic analysis of these mathematical results, we shall concentrate on the case $h = 0$, as the configuration $h \neq 0$ has no qualitative difference.

5.3.2.2 Interpretation of the results

In their analysis, Weidlich and Haag focus essentially on the fact that a high k corresponds to a “totalitarian” society. In other words, they believe that the imitation involved in individual decision-making is of the normative type¹⁹:

¹⁹ Of course, Weidlich and Haag do not use these terms, which are specific to us.

if the individual is influenced by the opinion of others, it is because he is afraid of the sanctions incurred by holding an opinion different to that of the group. As we have observed, this imitation is very close to preferential imitation. The informational dimension, on the other hand, is only of secondary importance to these authors. They write: “Clearly, this effect is independent of the truth of the opinions [1] and [2]” (p. 47). Alongside this imitative effect, formalised by the parameter k , their model also takes into account the intrinsic opinions of the population, through the parameter h . The situation corresponding to $h=0$ describes a configuration in which, from the perspective of the intrinsic preferences of the agents, the two opinions [1] and [2] are of equal importance. If we limit ourselves to this “fundamentalist” description, without the effects of imitation, we would expect the population to be equally divided between the two opposing political choices. What do we observe when k is different from 0?

For small k' , lower than 1, the equation (5.44) gives us a symmetric, unimodal stationary distribution centered on ($f=1/2$), as illustrated in figure 5.8. This distribution fits the fundamentalist evaluation perfectly: f lies in the neighbourhood of 0.5. In other words, the normative influence of the group is too weak to overcome the strength of individual opinions. However, under the effect of random events specific to the model, the collective opinion f can stray temporarily from the value 1/2. This effect is all the more powerful when the variance of the stationary distribution is high. When N increases, the dispersion of f around 1/2 decreases in proportion to \sqrt{N} .




Fig. 5.8. Stationary distribution $P_{st}(x)$ for $k'=0.5$ and $h=0$. (Ref : Weidlich and Haag, 1983, p.46.)

When N is constant but k' increases, while remaining lower than 1, the distribution remains unimodal, but its variance increases (figure 5.9). This is because interactions between agents become more intense. In other words, the population is still, on average, divided in two, following the intrinsic preferences of individuals, but collective opinion can be subjected to strong fluctuations, arising from normative influences and pushing it temporarily away from its equilibrium point. It must be emphasised that this variability does not result from the naturally random character of the fundamental characteristics²⁰, but derives from imitative interaction within the group. It is reasonable to believe that this type of variability plays an important role in many markets. Imitation functions as an amplifier of natural, exogenous fluctuations.

When k' is greater than one, we can see a qualitative transformation in the stationary distribution (figure 5.10). In mathematical terms, when we examine the deterministic process associated with the stochastic process (equation 5.21), we can see a “bifurcation”²¹: there are changes in the number




Fig. 5.9. Stationary distribution $P_{st}(x)$ for $k'=1.0$ and $h=0$. (Ref : Weidlich and Haag, 1983, p.46.)

²⁰ This naturally random character appears, for instance, in the situation where $k=0$. We then have: $p_{21} = Cste$, generating fluctuations around the average value of $1/2$.

of fixed points (from 1 to 3), and in their stability ($1/2$ becomes an unstable fixed point)²¹. As for the distribution itself, $f = 1/2$ becomes a local minimum, and two symmetric maxima appear, which we denote f_+ and f_- , such that: $0 < f_- < \frac{1}{2} < f_+ < 1$. This transformation is surprising (figure 5.10). It tells us that, even when there is no objective reason to prefer one opinion to the other, the effects of interaction are powerful enough to render the configuration ($f = 1/2$) extremely improbable. The system now has two favoured positions, corresponding to two symmetric modes, f_- and f_+ . These are “bubbles”, the first of which is negative and the second positive, as the majority chooses opinion [2] and opinion [1] respectively, even though the intrinsic preferences of the agents are perfectly indifferent. Once again, we are dealing with a phenomenon of self-organisation, inasmuch as one of the two opinions acquires the status of a norm accepted by the group, even though the two opinions are equally present in the population. This is the kind of situation that Mark Granovetter was looking for in his article on “Threshold models of collective behavior” (1978).




Fig. 5.10. Stationary distribution $P_{st}(x)$ for $k'=1.3$ and $h=0$. (Ref: Weidlich and Haag, 1983, p.47.)

²¹ This is referred to as a supercritical pitchfork bifurcation. The same bifurcation can be found in Orléan (1990).

The foundations of these two configurations reside in the processes of emergence and self-validation of shared beliefs. When the parameter k is high, the private agents rely heavily on the average opinion to determine their choice. This generates a phenomenon of self-validation of the majority opinion, powerful enough to ensure that the emergent opinion is lastingly disconnected from intrinsic preferences and stabilises around the two symmetric values. As in the previous model, when N is high, the probability of transition is low. In “finite time”, the system remains in the neighborhood of one of the two opinions. The choice between f_- and f_+ is then strongly dependent on the initial situation.

5.3.3 Preferential imitation

In this sub-section, we introduce the fact that the utility associated with each different option changes positively with the number of individuals who have chosen it, what we have called “preferential imitation”. More specifically, based on the work of Gérard Weisbuch (1995, 1996) we consider the consumption of two goods, denoted [1] and [2], such that the utility obtained can be written:

$$\begin{cases} U_1 = u_1 + Jn_1 \\ U_2 = u_2 + Jn_2 \end{cases} \quad (5.48)$$

In other words, consumption of good [1] has two components: an intrinsic component, denoted u_1 , and a component that is dependent on the choices of others, denoted Jn_1 , where n_1 is the number of individuals in the group who have purchased good [1]. We assume that J is strictly positive. Put another way, the higher the number of other individuals who choose this option, the more utility it provides. J is the coefficient measuring the strength of preferential imitation. We assume that consumption of good [2] follows the same logic. Consequently, we once again obtain our general model (5.1).

How can we describe the individual choices? Following Weisbuch (1995), we have chosen the “logit” model of probabilist choice described in chapters 1 and 3, namely:

$$\begin{cases} P_1(f) = \frac{\exp \beta U_1(f)}{\exp \beta U_1(f) + \exp \beta U_2(f)} \\ P_2(f) = \frac{\exp \beta U_2(f)}{\exp \beta U_1(f) + \exp \beta U_2(f)} \end{cases} \quad (5.49)$$

How do we interpret these equations? For $\beta = 0$, the equations (5.49) tell us that individuals choose between options [1] and [2] with a probability equal to $1/2$, whatever the values of the utilities $U_1(f)$ and $U_2(f)$. The more β increases, the more the choice becomes concentrated on the option with the greater utility. At the limit, for $\beta = +\infty$, the choice ceases to be probabilist and centres uniquely on the product with the greater utility. If we interpret this property in terms of the arbitration between exploration and exploitation, then a low β expresses behavior concentrating mainly on exploration, in which the observed value of utilities only plays a secondary role. A high β , on the other hand, means that individuals' choices are largely dictated by the utilities obtained. As these utilities are increasing functions of f , it follows that as β rises, the level of imitation also rises, to the detriment of random exploration of the products. Consequently, the overall level of imitation in the model depends on two effects: firstly, the influence of the other individuals on the utilities, as measured by the parameter J ; secondly, the influence of this utility on the actual choices, as measured by the parameter β . It is not surprising, therefore, that in what follows, the relevant parameter for understanding the results of the present model is βJ , which synthesises these two effects. To see this, we must first obtain the transition probabilities. After a series of manipulations²², we obtain:

$$P_{21}(f) = g(f) \exp \left\{ \beta \left[\frac{u_1 - u_2}{2} + JN \left(f - \frac{1}{2} \right) \right] \right\} \tag{5.50}$$

Once again, therefore, we obtain the previous model, well-known under the name of "Ising model". If we adopt the D_1 dynamic, we know that the process converges towards a unique stationary distribution, the extrema of which confirm the equation (5.45), so that:

$$x_m = \tanh \left[\beta \frac{u_1 - u_2}{2} + \beta \frac{JN}{2} x_m \right] \tag{5.51}$$

where the variable $x = 2f - 1 \in [-1, +1]$.

²² It should be noted that: $\frac{e^A}{e^A + e^B} = \frac{e^{A-B}}{e^{A-B} + 1} = \frac{e^{(A-B)/2}}{e^{(A-B)/2} + e^{-(A-B)/2}}$ with:

$$\frac{A-B}{2} = \beta \frac{u_1 - u_2 + JN(2f - 1)}{2}$$

Whether or not there exists a multiplicity of solutions depends on a threshold value of β , namely:

$$\beta_c = \frac{2}{JN} \quad (5.52)$$

We can observe two dynamic regimes separated by an abrupt transition, depending on the values of β . This is often referred to as a “phase transition” – a term originating from statistical physics.

If $\beta < \beta_c$, the equation (5.51) only has one solution. In the symmetric case where $u_1 = u_2$, we find $x_m = 0$: the two products share the market fifty-fifty. In the asymmetric case, there is a preference for one of the two products. This preference is proportional to the difference between the intrinsic preferences, i. e. $(u_1 - u_2)$, with a small β factor. This regime can be qualified as chaotic, insofar as the preferences are not very marked and the two products share the market on a roughly equal basis.

For the values $\beta > \beta_c$, the equation (5.51) has three solutions. The preferences are very marked: one of the products largely dominates the other in terms of market share, even in the symmetric case. Which of the two products becomes dominant is essentially determined by the initial conditions and not by the intrinsic utilities.

Remember that the existence of three solutions in the asymmetric case imposes an additional constraint of the type: $|u_1 - u_2| < h_0$. In this case, if we write:

$$r = \frac{u_1 - u_2}{u_1 + u_2} \quad (5.53)$$

the three solutions continue to exist as long as:

$$\left(\frac{3r}{2}\right)^2 < \left(1 - \frac{\beta}{\beta_c}\right)^3 \quad (5.54)$$

This condition is obtained by development of the hyperbolic tangent in the neighborhood of $r = 0$ and of the critical point $\beta = \beta_c$.

We shall now examine the way f varies as the intrinsic utilities change. To do so we shall switch from the context of the D_1 dynamic to that of the D_2 dynamic. We have seen that these two contexts are equivalent as regards their economic conclusions, but as we wish to bring out what is classically referred to as the phenomenon of “hysteresis”, it is more ap-

appropriate to place ourselves within the natural context for the presentation of this phenomenon, i. e. the theory of deterministic dynamic systems. We know that within the D_2 dynamic, in continuous time, the dynamic of f is governed by the following deterministic dynamic system of the first order:

$$\frac{df}{dt} = K[f(t)] \quad (5.55)$$

Graphically, to represent the fixed points of this process given by equation (5.51), it is useful to rewrite this equation in the following form:

$$\tanh(X) = \frac{2}{\beta JN} X - \frac{u_1 - u_2}{JN} \quad (5.56)$$

It follows that the fixed points are obtained at the intersection of the hyperbolic tangent and a straight line which varies with the parameters $(u_1 - u_2)$ while remaining parallel to itself. The higher the value of $(u_1 - u_2)$, the lower the intersection with the X -axis, as illustrated in figure 5.11. From now on, we suppose that $\beta > \beta_c$.

Let us assume that we start with a u_1 greatly superior to u_2 (the line denoted [0]). Individuals prefer product 1 (fixed point A). If the intrinsic quality of 2, namely u_2 , changes positively, then individuals will increase the share of product 2, while remaining favorable to product 1, even if u_2 becomes superior to u_1 . In figure 11, the straight lines move from the line [0] to the line [1], then to the lines [2] and [3], so that the equilibrium points move successively from A to B, then to C and finally to D. In D, though u_2 is greater than u_1 , x remains positive, i. e. [1] is preferred to [2]. However, when u_2 exceeds u_1 to such an extent that we arrive at the situation where $|u_1 - u_2|$ is greater than h_0 , there is no longer an equilibrium clearly favoring product 2 (line [4]). We then obtain the equilibrium point E. If, from this situation, u_1 starts to increase, the equilibrium points move from E to F, to G and then to H. Although the parameter $(u_1 - u_2)$ returns to its previous values, the equilibrium has changed. It no longer favours product [1] but product [2]. We call this a phenomenon of hysteresis. The system possesses a memory. When a parameter moves from h_0 to h_1 and then back to h_0 , the system does not return to its initial situation. This is because the underlying dynamic system has several possible equilibria and we pass from a “regime” favoring one product to a “regime” favoring the other.




Fig. 5.11. Hysteresis

5.4 Theses and conjectures

The detailed study of imitative interactions provides a perfect illustration of the evolutionary economists' watchword: "dynamics counts". In all the models studied, we have seen that the conventional fundamentalist approach that seeks to analyze economic phenomena simply through elementary characteristics, considered one by one, is not relevant, because we are within configurations where a multiplicity of equilibria is the rule rather than the exception. Self-organisation, bifurcation or path dependence are constantly present. "Phase transitions" have demonstrated, as in Granovetter's threshold model, that small variations in the control parameters, and particularly in those measuring the level of imitation, can generate radical qualitative changes.

In addition, it is important to note that these models call into question the traditional frontiers between different disciplines, because they systematically highlight the role of the interactions and influences connecting an individual to his neighbors and to the group as a whole. Essentially, this transdisciplinarity is expressed by the concept of network, which seeks to describe precisely the way each individual is linked to the others – to what we have called, in this chapter, the reference group. From this perspective,

the introductory reference to Mark Granovetter is of programmatic value. The explicit aim of Granovetter, in proposing the concept of network²³, is to stimulate a revival of socio-economics through the removal of the existing barriers between sociology and economics, barriers which have hindered the growth of both disciplines. Books written by pure economists, such as the one by Cohendet et al. (1998) or the more recent work by Durlauf and Young (2001) bear clear witness to the fact that this desire for opening also exists on the economic side. The evolutionary approach has an important role to play in this interdisciplinary revival. The transdisciplinary dimension of these models is also evident in the encounters they have provoked, as we can find physicists working in quantitative sociology (Weidlich and Haag, 1983), social psychology (Galam and Moscovici, 1991) or economics (Weisbuch, 1995 and 1996) just as we find economists (Durlauf and Young, 2001; Orléan, 1998) exploring sociological themes.

The form of transition probabilities was the first subject in the development of this general model, with its reciprocal influences. While a certain number of economists retained the “logit” form drawn from the “Ising model” (Lux, 1995), others have sought to deduce these probabilities from more solid microeconomic foundations, for example (Orléan, 1990, 1998; Kirman, 1993). It has also provided the opportunity for certain researchers to investigate the coexistence of imitative trends and opposing trends towards diversity (Gaio et al., 2001). One illustration of this is given by the “minority game” (Arthur, 1994), in which the individual seeks to adopt the minority opinion. It is used as an abstract plan for the examination of financial dynamics.

The most striking growth, however, has focused on analysis of the “spatialisation of interactions”. This concept arises naturally from the introduction of the idea that the reference group of an agent does not comprise the whole population but only his “neighbors” – whatever the notion of neighborhood may be taken to mean. Many works have concentrated on this question (Kindermann and Snell, 1980; Durlauf, 1993; Berninghaus and Schwalbe, 1996). It is worth noting that most often the qualitative results obtained are those presented here, for example equation (5.46) and the type of phase transition that it formalises. It is to be hoped, however, that these approaches will now extend beyond pure formalisation and be strengthened by empirical studies, which are at present all too scarce.

²³ After others, including Harrison White.

References

- Aglietta, M. and Orléan, A. (1980): *La violence de la monnaie*, Presses Universitaires de France.
- Arthur, B. (1988): "Competing Technologies: An Overview" in G. Dosi et al. (eds.), *Technical Change and Economic Theory*, Pinter Publishers, 590-607.
- Arthur, B. (1994): "Inductive Reasoning and Bounded Rationality", *American Economic Review*, vol. 84, n°2, May, 406-411.
- Berninghaus, S. and Schwalbe, U. (1996): "Conventions, Local Interaction, and Automata Networks", *Journal of Evolutionary Economics*, vol. 6, n°3, 297-312.
- Bikhchandani, S., Hirshleifer, D., and Welch, I. (1992): "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades", *Journal of Political Economy*, vol. 100, n°5, 992-1026.
- Cohendet, P., Llerena, P., Stahn, H., and Umbhauer, G. (eds.) (1998): *The Economics of Networks*, Berlin, Heidelberg and New York, Springer-Verlag.
- Corcos, A., Eckmann, J.P., and Malaspinas, A. (1998): "Self-Referred Decision Rules and Chaos", mimeo, 12 pages.
- Durlauf, S. (1993): "Nonergodic Economic Growth", *Review of Economic Studies*, vol. 60, 349-66.
- Durlauf, S. and Young, P. (2001): *Social Dynamics*, The MIT press.
- Elster, J. (1989): *The Cement of Society*, Cambridge University Press.
- Fama, E. (1970): "Efficient Capital Markets: A Review of Theory and Empirical Work", *Journal of Finance*, vol. 25, 383-417.
- Föllmer, H. (1974): "Random Economies with Many Interacting Agents", *Journal of Mathematical Economics*, vol. 1, n°1, March, 51-62.
- Gaio, L., Kaniovski, Y., and Zaninotto, E. (2001): "On Bubbling Dynamics Generated by a Stochastic Model of Herd Behavior", mimeo, 16 pages.
- Galam, S. and Moscovici, S. (1991): "Towards a Theory of Collective Phenomena: Consensus and Attitude Changes in Groups", *European Journal of Social Psychology*, vol. 21, 49-74.
- Granovetter, M. (1978): "Threshold Models of Collective Behavior", *American Journal of Sociology*, vol. 83, n°6, May, 1420-1443.
- Grossman, S.J. and Stiglitz, J.E. (1980): "On the Impossibility of Informationally Efficient Markets", *The American Economic Review*, vol. 70, n°3, June, 393-408.
- Jones, S. (1984): *The Economics of Conformism*, Oxford, Blackwell.

- Kindermann, R. and Snell, L. (1980): "On the Relation between Markov Random Fields and Social Networks", *Journal of Mathematical Sociology*, vol. 7, 1-13.
- Kindleberger, C.P. (1978): *Manias, Panics and Crashes*, London and Basingstoke, Macmillan Press.
- Kirman, A. (1993): "Ants, Rationality and Recruitment", *Quarterly Journal of Economics*, vol. 108, February, 137-56.
- Lux, T. (1995) "Herd Behaviour, Bubbles and Crashes", *The Economic Journal*, vol. 105, July, 881-896.
- Orléan, A. (1990): "Le rôle des influences interpersonnelles dans la détermination des cours boursiers", *Revue Economique*, vol. 41, n°5, September, 839-868.
- Orléan, A. (1997): "Jeux évolutionnistes et normes sociales", *Economie Appliquée*, vol. L, n°3, 177-198.
- Orléan, A. (1998): "The Evolution of Imitation", in P. Cohendet, P. Llerena, H. Stahn and G. Umbhauer (eds.), *The Economics of Networks*, Berlin, Heidelberg and New York, Springer-Verlag, chapter 15, 325-339.
- Orléan, A. (2001): "Psychologie des marchés. Comprendre les foules spéculatives" in J. Graverau and J. Trauman (eds.), *Crises financières*, Paris, Economica, 105-128.
- Schelling, T. (1979): *Micromotives and Macrobehaviour*, Norton.
- Weidlich, W. and Haag, G. (1983): *Concepts and Models of a Quantitative Sociology*, Springer-Verlag.
- Weisbuch, G., Kirman, A., and Herreiner, D. (1995): "*Market Organisation*", Santa Fe Institute Working Paper, n°95-11-102.
- Weisbuch, G., Chenevez, O., Nadal, J.P., and Kirman, A. (1996): "A Formal Approach to Market Organization: Choice Functions, Mean Field Approximation and Maximum Entropy Principle", mimeo.

6 Competition between firms

This chapter completes the description of the evolutionary conception of markets by introducing what competition is about in this framework. Chapter 4 has introduced irreversibilities and has given insights on how they may affect the functioning of markets and the aggregate outcomes by acting on the adjustment of individual actions. Chapter 5 has introduced the possibility that agents act mimetically thus affecting the way individual actions aggregate into market outcomes. This chapter introduces competition as a process that forces agents – here firms – to act in certain ways or make them vanish if they do not. The three chapters of this second part of the book thus provide insights to the reader on markets as institutions that make individual actions interdependent and that therefore affect agents' actions, the way they adjust, and the way they aggregate.

6.1 Background and problems

6.1.1 Reference models

In order to define what is “competition” in the standard approach to markets, two emblematic models of competition are briefly recalled: the Walrasian model of “perfect competition” and the Cournotian model of “imperfect competition”. While numerous developments have taken place in the domain of imperfect competition in the last thirty years (see Tirole, 1995), they do not change the conception of competition brought about in the standard approach to economics and even reinforce it by enlarging the variety of situations that are described in that framework.

The Walrasian model describes a market situation in which several “perfect competition” conditions are met. The exchanged goods are homogeneous, divisible and non stockable. The buyers (consumers) and sellers (firms) know the available goods and act optimally with regard to their information. In the short run, the agents are small and have no market power. In the long run, there is free entry and exit from the market. From these primary conditions, the theory infers that the transactions are realized under a price system satisfying secondary conditions. For each good, there

exists a unique price for all agents and transactions. Such a price is perfectly known by all agents. The agents act optimally by considering the price as exogenous (they are “price takers”). Each agent may proceed to any transaction at that price. From these last conditions, the theory infers that each agent expresses a supply or demand by equalizing the price with his marginal cost or utility. In the short run, the price of each good equalizes the total demand with the total supply. In the long run, the price converges towards the minimum of the mean cost, which implies that the firms make no profit.

The Cournot model describes a market situation in which some “perfect competition” conditions are relaxed. The firms are no longer small and have some market power, contrary to the buyers. They know not only the goods, but the characteristics of their rivals. Hence, they become “price makers” since they impose the prices of the goods they produce. They act optimally by expecting the behavior of their rival. Finally, each firm defines a reaction function which tells what quantity to supply for each possible quantity supplied by its rival. The price results from the compatibility of the reaction functions.

Those two emblematic models are similar in one essential dimension: their reliance on a notion of equilibrium. The concrete process leading to such an equilibrium is not precisely stated. However, it is possible to consider that it results from “rational expectations”. The agents have a perfect factual and structural information on their interaction context and form optimal expectations with regard to that information. In the Walrasian model, each firm perfectly expects the prices and reacts to them in a parametric way. In the Cournot model, each firm simulates the opponent’s behavior and fixes his own supply accordingly.

In fact, the standard approach to competition is a theory of coordination in the allocation of resources. As emphasized by J. Vickers (1993) “competition is an equilibrium state of a market dependent on those fundamental forces of demand and cost structure that determine the number of viable survivors. That is to say the state of competition is equated with the structure of the market which is measured by the number and relative size of the surviving firms. From this follows the familiar taxonomy with perfect competition at one end of the spectrum and monopoly at the other, each defined in terms of a relation between market share and the consequential ability to increase price above marginal costs. Between these extremes lies less clear cut territory either in relation to the idea of monopolistic competition or in relation to concepts of oligopolistic interaction in which expectations of rivals behaviours have overwhelming significance”.

Interestingly enough, a parallel can be stated between the theory of competition respectively developed in biology and in economics. Biologists distinguish between “exploitation competition”, in which two species have equal access to a common resource and find a way to use the resource without fighting over, and “interference competition”, in which both species develop strategies to prevent the other one to access the resource and thus to fight over it. The economic model of perfect competition is a model of exploitation competition while the imperfect competition models is a step towards interference competition. However, the imperfect competition models still fundamentally conceive firms’ actions as optimal reactions in contexts where all the other things are given. Even when these models are dynamic and consider successive periods, firms do optimize through backward induction.

6.1.2 Main critics

In front of this conception of competition, various economists have long opposed a conception of competition as a dynamical process rather than an equilibrium state. The evolutionary conception of competition presented in this chapter draws heavily on this heritage. Hayek (1968) for instance denounces the “absurdity of the usual procedure of starting the analysis with a situation in which all facts are supposed to be known. [...] It leaves no room whatever for the activity called competition, which is presumed to have already done its task.” In other words, the conception of competition which is behind standard models of markets has nothing to do with the idea of a struggle with unpredictable outcomes between rival firms but has rather to do with studying the value that competitive variables may take once the competition is supposed to have played its role. Even if the Cournot model and other models of imperfect competition are about strategic interdependence among firms, they are not about any struggle since the problem of firms in those models is to coordinate their actions, not to fight. More precisely, three main strands of critics to the standard conception of competition can be asserted.

First, the standard approach assumes that firms have a complete information on all the profitable opportunities, available technologies, and even other’s preferences. For Hayek (1968), the spirit of competition is precisely to be a *discovery procedure*: “competition is important primarily as a discovery procedure where entrepreneurs search for unexploited opportunities that can also be taken advantage of by others. [...] Market theory often prevents access to a true understanding of competition by proceeding

from the assumption of a “given” quantity of scarce goods. Which goods are scarce, however, or which things are goods or how scarce or valuable they are is precisely one of the conditions that competition should discover: in each case it is the preliminary outcomes of the market process that inform individuals where it is worthwhile to search”. In other words, the essence of competition is to make rivals search for what consumers prefer and what their willingness to pay is. The idea of a discovery procedure is consistent with the idea of a trial and error process.

Note that standard economics has developed models with imperfect information that lead to equilibria with a unique price different from the competitive one, a multiplicity of prices or no equilibrium at all (Salop and Stiglitz, 1977 and 1982; Reinganum, 1979; Diamond, 1971). But the type of uncertainty is perfectly known by the agent who optimizes in consequence. There is also a literature on auctions and related situations (contests, tournaments) that studies how the competitive setting affects the prices (or information) that decentralized agents reveal and the price that is eventually elicited (McAfee and McMillan, 1999). This literature thus provides insights on prices in the absence of an auctioneer but it is rather deductive, showing the influence of given competitive settings on individual behaviors, rather than representing in an inductive way how decentralized behaviors can give rise to different competitive settings.

Second, the standard approach assumes that there is a given set of opportunities already existing or which is predictable. For Schumpeter (1942, p.84), the essence of competition is precisely to promote *novelty* in the market: what matters is “the competition for the new commodity, the new technology, the new source of supply, the new type of organization – competition which demands a decisive cost or quality advantage and which strikes not at the margins of the profits and the outputs of the existing firms but at their foundations and their very lives”. The purpose of innovation is precisely the capturing of custom and of monopoly rents away from competitors in a context where resources are limited (both resources to produce goods and services and resources to purchase those goods). It is the introduction of a new good or technology or service which will, for some time at least, give a competitive advantage to some agent. In the Schumpeterian tradition, the introduction of change is a process of “*creative destruction*” in which firms split into winners and losers, and winners at a point in time may become losers when novelty is later on introduced by rivals.

Standard models of course deal with innovation, but they consider innovation in two ways. On one hand, it is a pure random variable, in which

case it is a predictable outcome that can be anticipated *ex ante* and that induces an *ex ante* optimal reaction. On the other hand, it is a variable that can be optimized as all the other competitive variables. In both cases, the purpose of innovation is reduced to inducing an instantaneous and passive reallocation of resources for all firms.

Third, the standard approach to competition considers the firm as an entity able to perform complex calculus in order to maximize its profit. For the managerialist and behavioralist economists, such a description of firm behavior is unrealistic. On one hand, Baumol (1959), Marris (1964) and Williamson (1964) propose that profits are only one possible objective of the firms and that in particular firms managed by professional managers have multidimensional objective functions. On the other hand, Tinter (1941), Simon (1955, 1959), Cyert and March (1963) and Winter (1964) draw on their empirical and theoretical studies to show that firms behave according to simple decision rules rather than to optimizing principles: they tend to “satisfice” rather than to optimize. For Nelson and Winter (1982), the standard reliance on optimization behavior leads to “disregard essential features of change: the prevalence of Knightian uncertainty, the diversities of viewpoint, the difficulties of the decision process itself, the importance of highly sequential ‘groping’ and of diffuse alertness for acquiring relevant information, the value of problem solving heuristics, the likely scale and scope of actions recognized *ex post* as mistaken”.

Here again, standard models try to introduce limited information and bounded rationality of the agents. They consider for instance the “discrete choice model” where the agents choose an action with a probability which is proportional to their utility. But the prices still result from an equilibrium, *i. e.* a state where no boundedly rational player has an interest to deviate.

6.1.3 Evolution and competition

Besides, a long trend examined economics in a strict evolutionary way, that is with a biological analogy. Alchian (1950) was the first to suggest an approach of competition that “embodies the principles of biological evolution and natural selection by interpreting the economic system as an adaptive mechanism which chooses among exploratory actions generated by the adaptive pursuit of ‘success’ or ‘profits’”. He assumes that behaviors are adaptive, imitative and trial and error rather than maximizing. He introduces a reflection on what is profit maximization in a world of uncertainty and what could be the chance of a maximizer to survive in an evolving world.

Friedman (1953) responded to this whole reflection by suggesting that “the profit maximization hypothesis [...] is an appropriate summing up of the condition of survival” meaning that only those firms that indeed maximize their profits will survive to natural selection. Friedman seeks to show that Alchian’s argument can be reconciled with the standard approach that assumes that only maximisers will exist by conceding that competitive forces can be seen as a temporal process of selection which retains maximising firms and makes them prosper while eliminating non-maximising firms. However, as suggested by Alchian himself, if there is radical uncertainty, profit maximization cannot be a guide to action at all and therefore there is no way it can summarize the conditions for survival.

In fact, Friedman’s argument raises a certain number of questions. To begin with, even admitting that some firms are maximising, the argument is only valid under a certain number of conditions. Firstly, the maximising and non-maximising firms must produce a homogeneous good under the same cost conditions. Otherwise, one can easily imagine situations in which non-maximising firms which possess a cost advantage or sell a product of a higher quality will eliminate maximising firms. Secondly, maximising firms must systematically reinvest their profits in additional production capacity. If, on the contrary, they distribute most of their profits in the form of dividends to shareholders, they may find themselves at a disadvantage faced with non-maximising firms which implement an aggressive investment policy. Thirdly, the flow of non-maximising firms entering the market must not be significantly higher than the flow of maximising firms.

Friedman’s argument also raises more fundamental problems for any economist whose objective is the formal representation of the temporal process of selection that gives consistency to competitive forces. Firstly, Friedman assessed, as a means of supporting his argument, that “the immediate determinant of behavior (can be) anything – routine, chance, or no matter what else. As soon as the determinant leads to behaviour compatible with rational profit maximisation, the firm prospers and acquires the resources needed to grow; if, on the contrary, it does not lead to such behaviour, the firm tends to lose resources and can only survive if it manages to obtain resources from outside”. From the moment that the aim is to describe a temporal process of selection, Friedman’s argument cannot guarantee that a firm finding itself by chance maximising at a certain period will still be maximising at a later period, nor that a firm having a routine that enables it to become maximising at a certain period will have survived until that period. Secondly, Friedman’s argument assumes that there is always, in each period, a sufficient number of firms adopting actions that

lead to profit maximisation. Consequently, this assumes that if there are few or no such firms during a certain period, then firms possessing this characteristic will enter the market. Economically, this hypothesis does not appear very obvious.

Many formal models have sought to test whether firms that maximise profits rather than another objective function have the highest chance of long term survival. Winter (1964, 1971) and Nelson and Winter (1982) provide a coherent framework for reconciling the intuition of natural selection operating on effective actions and the necessary continuity in firm behaviours to make natural selection precisely operating and selecting the most performant actions while eliminating the others. Their hypothesis is that firms act according to behavioral rules that they maintain as long as they yield good results and seek to change when they do not. Chiappori (1984) formalizes the survival of firms in a Markovian context where the processes of entry, growth, decline and exit are specified on the basis of transition probabilities from one state to the other. The disappearance of sub-optimizing firms only occurs under very restrictive conditions requiring either that the probability of degeneracy is nil and no sub-optimizing firm can enter or that an infinite number of optimizing firms enter at each period. Dutta and Radner (1999) have shown that when random events (to which probabilities can be attributed) are introduced, firms seeking to maximise other objective functions have a greater chance of surviving at equilibrium than firms with profit maximisation as their objective.

More generally, for an economist interested in the explicit way in which competition maintains some behaviors or firms and eliminates others, it is necessary to take into account various different phenomena, such as incomplete information, costly choices, irreversible investment decisions -all elements included in standard models but as affecting the allocation of resources in equilibrium rather than as affecting the ability to compete with others- as well as the differential abilities of firms to get and combine resources²⁴ and the fact that conditions of efficiency and survival are modified endogenously. A more elaborate theory of the firm in interaction is required in order to extract the heuristic value of the representation of the process of selection, i. e. the action of competitive forces. This involves explaining, in a non-trivial way, which behaviors are capable of surviving in one or another competitive environment. In fact, the evolutionary conception of competition set forth in this chapter builds on those critics of the

²⁴ It is well known in cooking concourses that no one does the same thing with exactly the same ingredients. And it does not even include expectations.

standard approach to competition and on the willingness to provide models of competition as a process. But they rest on a broader conception of evolution which departs from the biological analogy in order to introduce true human learning principles.

6.2 Canonical principles

In this chapter, the firm is still considered as a unique centre of decision (see chapter 7 in this book for entering the black box of the firm). The crucial departure from the standard approach is to consider that the firm has a limited information and a bounded rationality, but is able to learn through time. The bounded character of firm's rationality does not derive from its intentions, that is clearly to make acceptable profits, but from the fact that it is endowed with constrained capacities of gathering and treating information. It has an incomplete and imperfect information about its opponents and the consumers' demand and is no more optimizing its price or supply. However, it makes expectations which are confronted to realizations and it takes actions which are regularly tested. The fixing of prices is thus decentralized rather than centralized as in the perfect competition model. It is adjusted through a dynamic process from which emerge various structures (prices, firm concentration, permanent market links).

The evolutionary models of competition that articulate with the developments in this book are characterized by the following "building blocks":

Principle 1: Firms are searching which actions will be successful.

Since there is no *deus ex machina* to provide a complete information, the firms have to supply for it by revising their expectations of consumer preferences and willingness to pay as well as their expectations of rivals' actions. Unable to implement an optimal *ex ante* calculus, they must decide on the actions to adopt by trying different actions and observing the market feedback. Thus, the firms get information that has a local character and a momentary value. The differential success of alternative actions becomes a central ingredient of an evolutionary theory of competition.

Principle 2: Firms compete for implementation of novelty.

An essential feature of competitive behaviors is to attract custom by offering them goods and services with distinctive characteristics. In this perspective, competition drives to the creation of new items, such as new technologies and new products. Rather than considering that firms behave competitively in a well known context (in the sense of optimizing in order

to get the maximum custom under constraints), the evolutionary conception of competition assumes precisely that competition is the inducement for the creation of new items. What is emphasized is the intensity (or lack of intensity) in the willingness and possibility of firms to find original means of capturing a large part of the custom at the disadvantage of their rivals.

Principle 3: Firms must build capabilities to compete.

Firms must invest in resources for competing, whether human capabilities, technical capacities or else research and development. But those investments both are motivated by competition and affect by feedback the terms of competition. For instance, when German chemical firms started to build R&D capabilities in the 19th century, it changed the nature of the competition in this industry at the international level. Hence, investments in resources for competing are not any more the result of an optimization calculus, nor are they constraints to the optimization calculus. In an evolutionary conception to competition, they are part of what can make a firm successful rather than vanish.

Implication: Aggregate structures emerge from the decentralized actions undertaken by the firms

It is from the active competition preceedingly sketched that will emerge prices, networks and even complete market structures. Standard theory argues that costs are a very strong determinant of market structure, with decreasing returns inducing many firms to coexist. However, if novelty is the fuel to competition, a breakthrough product offered by only one firm may be supplied scarcely in regards of demand and the firm should capture a big rent. If rivals are not able to provide an equivalent product, even decreasing returns won't do anything to the serving of the market. Indeed, this intuition is in the standard theory since imperfect competition results from firms introducing differentiation of their products and reducing the elasticity of the demand that addresses to them, therefore weakening the impact of the cost structure on the degree of competition. But the standard theory postulates the degree of competition and whether costs play an important role in this degree. An evolutionary theory of competition, on the contrary, shall derive whether costs are important as an outcome of firms behaviors, in particular as an outcome of the type of means that firms are using to attract custom. Thus, there is no discontinuity in the evolutionary theory of competition between "perfect" and "imperfect" competition, but rather a continuum of possible situations deriving from the effective behaviors and actions explored and implemented by the firms.

6.3 Some models

Three polar models of an evolutionary theory of competition between firms are presented. Each of these models focuses on one significant phenomenon and therefore tends considerably to simplify the other aspects. They allow to compare the predictions of the standard models with the predictions of the evolutionary models regarding the relationship between competition, on the one side, and market structures, on the other side. They thus make it possible to assess the validity of standard microeconomic results when competition is explicitly represented as a process rather than an equilibrium state.

The first model (Lesourne and Caron Salmona, 1985) describes the process of decentralized tatonnement about which behavior to adopt, with no reference to any equilibrium. On the contrary, it is firms' behaviors and procedures to test and revise those behaviors that leads the price structure to emerge. However, it keeps as much as possible the other aspects of the perfect competition model such as an homogeneous good and the free entry and exit.

The second model (Nelson and Winter, 1982) describes a competition process driven by strategies over novelty. In this model, there is no firm tatonnement per se as behaviors are captured by stable routines. But this model allows to draw implications about possible trajectories and outcomes of competition, especially as concerns the formation of market structures.

The third model (Coursac et al, 1998) describes a competition process in which firms do simultaneously build (selling) capacities in order to compete and compete effectively on the market. Standard models already deal with firms investment in capabilities (such as R&D for instance), but they deal with such decisions as with any other decision, by stating it *ex ante* without further revision.

6.3.1 A model of decentralised tatonnement

The first model (Caron-Salmona and Lesourne, 1985) examines the conditions under which a decentralised market can lead to the emergence of a unique price equal to the marginal cost. The model describes the functioning of a competitive market that has characteristics that makes it comparable with the "perfect market" of the standard approach: numerous buyers and sellers are in interaction; the exchanged good is divisible, homogeneous and non-stockable; none of the agents has the power to impose the price of

the good; buyers seek to obtain the good at a reasonable price; sellers seek to make a profit or at least to survive. With this description as its basis, the model formalises the sequential process of adjustment of the exchange prices and the quantities of the good supplied and bought. It enables to demonstrate how traditional microeconomic results are affected by original hypotheses about behaviour, essentially the fact that sellers have to search for acceptable prices and quantities rather than take the price for given and calculate accordingly the optimal quantity to supply.

In the market under consideration, there are K buyers, indexed $k = 1, \dots, K$, and J sellers, indexed $j = 1, \dots, J$. Each buyer k has an individual demand function decreasing as a function of the price p , i. e. $D_k(p)$ such that $D_k(p) = 0$ if $p \geq p_k$. Each seller j is a retailer supplying a quantity $q_j(t)$ at a unit price $p_j(t)$. The good is acquired by the retailers at a constant unit cost c that is identical for all of them. Thus, each unit sold at the price $p_j(t)$ provides the retailer j with a positive profit $\prod_j(t)$ if $p_j(t) > c$.

In each period t , the buyers k enter the market one after another in a random order. Each buyer k draws a random sample of retailers, to which he adds the retailer from whom he bought during the previous period, if this retailer is unique. This is the "favored" retailer of buyer k during the period t . The buyer k then examines the price offered by each retailer in the sample and chooses the one offering the lowest price, i. e. the price $p_k = \min(p_j)$, if $p_j < p_k$ and for the p_j included in the sample. If more than one retailer offers this price, the buyer will choose his favored retailer if he is part of this subset, otherwise he will choose one at random. The buyer k then buys the quantity desired at the price p_k , i. e. $D_k(p_k)$. If the chosen retailer cannot supply the whole quantity desired, the buyer k goes to the other retailers, choosing them in order of increasing price as long as $p_j < p_k$. Of course, if the quantity $D_k(p_k)$ can be obtained from several different retailers who all sell at the price p_k , the buyer k obtains $D_k(p_k)$ at the price p_k . If, however, the other retailers whom the buyer goes to see sell at a price higher than p_k , the buyer must review his demand downwards, on the basis of the average price he would pay if the remainder of his demand was satisfied at this stage.

At the end of each period t , each retailer observes his sales and his unsold goods, if there are any. The hypothesis is adopted that the retailer also remembers the results he obtained for the period $t-1$, to ensure the intertemporal coherence of his strategy. According to his results for the periods t and $t-1$, he must choose the quantity and the price he will offer for

the period $t+1$. However, as the demand that is addressed to each retailer is a random variable $D_j(t)$, and he does not know the distribution of probability underlying the random sampling of the buyers, he cannot calculate the associated expected profit. He therefore makes use of a simple rule to choose the price and quantity of the good that he will offer for the next period.

From one period to the next, the price is adjusted by a maximum of one unit upwards or downwards. The retailer does not know the distribution of prices in the market and adopts a process of tatonnement around the price he offered for the previous period. Two different cases can be distinguished:

- if $\prod_j(t) > \prod_j(t+1)$, the retailer adjusts the price to bring supply and demand closer together:
 - if $q_j(t) > D_j(t)$ then $p_j(t+1) = p_j(t) - 1$
 - if $q_j(t) = D_j(t)$ then $p_j(t+1) = p_j(t) - 1$ or $p_j(t)$ or $p_j(t) + 1$
 - if $q_j(t) < D_j(t)$ then $p_j(t+1) = p_j(t) + 1$
- if $\prod_j(t) < \prod_j(t-1)$, the retailer does not persevere with the price policy he used for the period t :

$$p_j(t+1) = \begin{cases} p_j(t) + 1 & \text{if } p_j(t-1) - p_j(t) > 0 \\ p_j(t) - 1 & \text{if } p_j(t-1) - p_j(t) < 0 \end{cases} \quad (6.1)$$

From one period to the next, the quantity supplied is adjusted by assuming that the retailer believes that demand is a decreasing function of price and that he does his best to improve his profit. Returning to the two cases just presented:

- if $\prod_j(t) > \prod_j(t-1)$, the retailer adapts the quantity he supplies to a level consistent with the price change that he has already decided upon. For example, if he has decided to lower his price because he had unsold goods, he will now supply a lower quantity than he supplied for the previous period, but a higher quantity than he actually sold, $D_j(t)$. Thus
 - If $p_j(t+1) < p_j(t)$ then $q_j(t+1) \geq D_j(t)$
 - If $p_j(t+1) = p_j(t)$ then $q_j(t+1) = D_j(t)$
 - If $p_j(t+1) > p_j(t)$ then $q_j(t+1) \leq D_j(t)$.

It can be demonstrated that there always exists a choice $[p_j(t+1); q_j(t+1)]$ such that in principle (because the demand is a random variable), the expected profit increases: $q_j(t+1)[p_j(t+1) - c] \geq \prod_j(t)$.

- If $\prod_j(t) < \prod_j(t-1)$, the retailer chooses $q_j(t+1)$ in such a way as to expect a higher profit than in the period before last: $q_j(t+1)[p_j(t+1) - c] \geq \prod_j(t-1)$.

Taken together, these mechanisms determine a stochastic dynamic process of adjustment of prices and quantities exchanged. A stable state is defined by two characteristics: no retailer thinks he can further increase his profit by modifying the price and quantity he offers; no buyer thinks he can find a retailer offering a better price than his favoured retailer.

This model produces a series of results:

1. The process does not systematically converge towards a stable state and the prices may fluctuate endlessly. However, the process does converge in probability towards a stable state within a finite time under one particular condition. This condition holds that there exists at least one prudent retailer, in other words a retailer who never proposes more than once a price p_j such that $D_j(t)[p_j(t) - c] < \prod_j(t-1)$, in other words a price such that if the retailer satisfied the demand at this price, his profit would be reduced.
2. In a stable state, the selling price is unique. All the retailers end up by selling at the same price and the buyers purchase from their favoured retailers.
3. The value of the unique price that rules in the stable state cannot be predicted a priori. It depends on the initial conditions and the history of the process. Therefore, this price may very well be higher than the marginal cost c , and provide a rent to all the retailers.
4. To find the standard result, according to which competition makes the price converge towards the marginal cost c , there must exist, in addition to the prudent retailer, a competitive retailer, in other words a retailer who considers that, in every period, reducing the price is a more profitable strategy than increasing it because this generates a strong increase in demand.
5. The monopoly price p^* rules in the stable state if the three following conditions are satisfied:

- the buyers all have the same demand function $D(p)$ and the profit generated per unit sold is a concave function, maximal for $p^* (< p)$;
 - a buyer who has a favored retailer j in the period $t-1$ buys from this retailer in the period t if $p_j < p$, otherwise he buys nothing; a buyer who has no favored retailer chooses one at random;
 - all the retailers have the same behavior, that of learning by reinforcement: all the prices which have, in the past, resulted in a fall in profit are definitively rejected (as in the case of the prudent retailer) and it is assumed potentially profitable to try again any price that has generated a rise in profit in the past.
6. If the consumers fall into two categories, the “mobiles” who examine the prices before choosing and the “conservatives” who keep their favored retailer as long as he is satisfactory, a stable state can be characterised by a unique price or price dispersion. A unique price is only obtained in the stable state if all the active retailers have mobile buyers. In a stable state with price dispersion, the mobile buyers necessarily purchase at the lowest price observed on the market.
7. This model demonstrates that, for given cost and demand characteristics, the buyers’ and sellers’ rules of behavior strongly condition the properties of the competitive processes. Depending on these rules of behavior, we can observe the convergence or the absence of convergence towards a stable state, convergence towards the minimum price of perfect competition or towards the monopoly price, convergence towards a unique price of equilibrium or towards price dispersion. Thus, rules of behavior create evolution in the functioning of markets. If the retailers have price strategies that are not very coherent over time, prices may fluctuate without end. The appearance of a retailer with prudent behavior can disturb this process and lead to the emergence of a unique price. If a retailer with competitive behavior then appears, this unique price moves towards the minimum price. Lastly, if all the retailers adopt learning behavior and the buyers are passive, a unique monopoly price may be the outcome. The model thus makes it possible to explain phenomena that traditional analysis does not take into account. For example, a change in the behavior of a retailer or his entry into the market with competitive behavior can transform the state of the market even if the fundamentals of the activity remain unchanged.

6.3.2 A model of Schumpeterian competition

The second model (Nelson and Winter, 1982, chapters 12 and 13) studies how firm behaviors generate market structures. Firms are in competition for an homogenous good and seek to get the maximum profit through minimizing their production cost. At each period, they are “price takers”, relatively to a market price that results from the equalization of supply and demand. This setting is interesting in the sense that it keeps many characteristics of a situation of perfect competition. What is suggested here however is that even when firms compete to sell an homogeneous good for which the elasticity of the demand addressed to each firm tends to be infinite, they may indeed displace the competition in the technological realm, making bets on the ability of technological innovation or rather imitation for gaining market shares. Whereas in the standard framework of competition, firm behavior was dictated by cost fundamentals and the associated univocal market structure, in this evolutionary framework of competition, firm behaviors induce which market structure will eventually emerge as well as the sustainability of that structure.

More precisely, in the market for a homogenous good, there are n firms without entry. At each period t , firm i is characterised by a given technology, its stock of capital $K_i(t)$ and its productivity by unit of capital $A_i(t)$. The quantity $q_i(t)$ produced by each firm i in period t is given by the simple production function with constant returns and complementary inputs: $q_i(t) = K_i(t)A_i(t)$. The cost c of a unit of capital is similar for all firms and all technologies. The unique price $p(t)$ results from the equalization of the total supply of all firms $S(t)$ and the total demand $D(t)$, without specifying the process leading to it.

From one period to the next, firm i can change its situation in the market through different actions:

- the firm can modify its productivity $A_i(t+1)$. Two strategies are available: either imitating the best available technique (activity m) and/or searching for a better technique, hence innovating (activity n). Both strategies require that the firm invests in R&D and both strategies have an uncertain outcome, either a failure or a success. This depends on (a) the money invested in research: r_{im} per unit of capital for imitative R&D and r_{in} per unit of capital for innovative R&D; (b) the efficiency of the research activity, a_m for imitative R&D and a_n for innovative R&D. The probability of success is equal to $\min[a_m r_{im} K_i(t), 1]$ for imitative R&D and $\min[a_n r_{in} K_i(t), 1]$ for in-

novative R&D. When imitation succeeds, firm i gets the best available productivity of period t , hence $A(t)$. When innovation succeeds, the productivity that comes out of the process is a random variable $A_i(t)$. If technical progress is rather cumulative, firm i will draw a productivity level slightly higher than $A_i(t)$ with a high probability, whereas if technical progress is not cumulative, it will face (like all other firms) a productivity level independent of its current productivity $A_i(t)$. Thus, on a whole, the firm gets the best productivity it could have obtained:

$$A_i(t+1) = \max[A_i(t), A(t), A_i(t)] \quad (6.2)$$

- the firm can invest in production capacities for increasing $K_i(t)$. Between two periods, the stock of capital varies according to investment $I_i(t)$ and depreciation factor δ , such that $\Delta K_i(t) = I_i(t) - \delta K_i(t)$. The investment depends on three factors: the market share of firm i equal to $\frac{q_i(t)}{S(t)}$, the efficiency per unit produced equal to $\frac{p(t)A_i(t)}{c}$, the profit per unit of capital equal to $\prod_i(t) = p(t)A_i(t) - c - r_{im} - r_{in}$. Hence, $I_i(t) = F\left(\frac{p(t)A_i(t)}{c}, \frac{q_i(t)}{S(t)}, \prod_i(t)\right)$, where F is increasing in its first argument and decreasing in the two others. Depending on the precise form of this function, firms have a rather “restrictive” or rather “aggressive” policy for increasing their production capacities. Through their investment policy, firms are able to indirectly affect the market price: aggressive policies contribute to reduce the market price while restrictive policies contribute to maintain the price. In this respect, firms with a larger market share have more power on the market price than firms with a smaller market share.

The model describes a stochastic dynamical process which is studied by simulation. The following qualitative results are obtained.

1. *What structure emerges from this competition?* Consider first that imitation and innovation behavior are absent as well as investment. If all firms have the same productivity level at the outset, the structure of the model induces that all firms share the market equally. If a subset of firms are more efficient and would produce enough to supply the market, then the less efficient firms would be driven out

of the market. Consider now that the firms can change their technology and invest. The structure of the model induces that whatever the initial structure –rather concentrated at the outset or rather non concentrated at the outset – the market structure at the end of the process is concentrated (hence a small number of firms have a significant share of the market). Obviously, the less concentrated the structure at the outset, the more significant the process of concentration for describing the market dynamics. The central phenomenon in generating this tendency to concentration is differential firm growth. Of course the level of concentration after some time depends on the extent to which the most successful firms do grow in capacity. When the most successful firms invest in capacities, they induce a rise in total supply that drives the market price down. Therefore, the less efficient firms exit from the market because their productivity is too low to enable this price to cover their production cost.

2. *What factors affect the process of concentration of the market structure?* The factors that affect the process of concentration of the market structure are those that affect the differential growth of firms. What induces concentration to arise basically is that some firms grow faster than others. So the investment policy is a central determinant of the concentration of the market structure. If successful firms were not investing, they would just make more money, but the market structure would not be affected. What modifies market structure over time is that firms seek to grow and take the market from the less efficient ones. In this model, investment depends upon efficiency and profit as well as current market share. Better performance thus translates into an increase in production capacity. The larger this increase as a function of current performance, the larger the increment in market share realized by the best performers and the quicker the convergence to a concentrated structure. Depending upon the ease of imitation – whether it is easy to reach $\bar{A}(t)$ –, the nature of technological progress – whether the random variable representing the innovative draw is firm dependent or firm independent –, and the respective probabilities to imitate and innovate, firms that make an increment keep a more or less durable advantage over the others and are thus more or less able to force the others out of the market.
3. *Are some behaviors more likely to survive than others?* The structure of the model induces innovators to do bad as compared to imitators. In other words, it considers the case where innovation is not

beneficial and where therefore it is not worth supporting innovative R&D investment. So, it provides only a partial analysis of behavior differential performance. It nonetheless shows that when there is a diversity of behaviors, some may be unfit for survival because they do not allow firms supporting those behaviors to have positive differential growth.

In synthesis, the tendency towards concentration is all the stronger when technical progress is fast, the probability of innovation is low, imitation is difficult and investment policies are aggressive, because those factors induce durable differential firm growth to occur: those firms which have good draws have durable advance over their competitors which they deepen by investing in capacities and thus increasing rapidly their market share. Hence, differences between firms in market share more or less average away. They average away if firms can imitate the most efficient firms easily, if the most efficient firms do not exploit their advantage by increasing their capacities, if technical progress is slow and induces only small increments. On the contrary, they tend to be magnified if the less efficient firms can not imitate the most efficient ones, if the most efficient ones have aggressive investment policies, if technical progress is cumulative and makes the most efficient ones still more likely to increase their efficiency. Hence, fundamentals at some date are not sufficient to predict market structure and characteristics at further dates, even when competition is simple and mimics this described by standard economics, due to the time-consuming and unequal construction of durable competitive advantages.

6.3.3 A model of market interdependence

The third model (Coursac et al., 1998) studies the dynamics of a market for a differentiated product distributed in shops on a discrete spatial structure. The consumers have to support a transportation cost for acquiring the good due to both their own localization and the firms' localization. The firms have localization strategies since they can buy and sell shops or set up shops in new locations. Moreover, the firms have industrial strategies since they can purchase other firms. Once again, the fundamentals of the product market are not sufficient to predict the dynamics of the market structure. The type of competition between firms, submitted to local markets but able to create new locations, plays a crucial role. The evolution in the structure of shops' localization induce changes in the local environment of each firm that modifies its profit opportunities, hence the structure of shop's localization. The essential sequentiality of the decisions (as in a stock market for

instance) makes profit opportunities to be modified endogeneously. Firms therefore are not able to calculate ex ante an optimal number of shops nor their optimal location because opportunities for opening or closing shops are modified endogeneously and cannot be expected.

However, due to its complexity, the model can only be studied by simulation. After a general description of the market structure, the strategies of the firms concerning the localization of shops and the purchase of firms will be successively studied.

6.3.3.1 Description of the market

Each firm j ($j \in \{1, 2, \dots, J\}$) supplies a product defined by two physical characteristics, denoted (ξ_{1j}, ξ_{2j}) . The price is obtained by a combination of these characteristics: $p_j = \alpha \xi_{1j} + \beta \xi_{2j}$. The two coefficients α and β are identical for all firms. The values of the characteristics are drawn at random from a uniform distribution $[0, N]$. A firm that is active on the market possesses n_j shops ($n_j > 0$). Each shop k_j ($k_j \in \{1, \dots, n_j\}$) is situated on a spatial structure in the form of a flat grid and its location is defined by two coordinates (a_{k_j}, b_{k_j}) . One of these shops is defined as the head office of the firm, and its address is denoted (a_j, b_j) .

Each consumer i ($i \in \{1, 2, \dots, I\}$) is situated on a point (a_i, b_i) in the two-dimensional space, with a_i and b_i obtained by drawing at random from a uniform distribution $[0, N+1]$. Each consumer has a preference peak for certain values (ξ_{1i}, ξ_{2i}) of the product characteristics, also drawn at random from a uniform distribution $[0, N]$. Each consumer buys one unit of product in each period. Each consumer chooses to buy this unit in the shop k_j which maximises his utility $U_i(k_j)$. This utility increases in proportion to the closeness of the quality offered by the shop to his preference peak and also in proportion to the geographical closeness of the shop to his own location. The utility is assumed to be quadratic:

$$U_i(k_j) = -\mu_i \left[(a_i - a_{k_j})^2 + (b_i - b_{k_j})^2 \right] - (1 - \mu_i) \left[(\xi_{1i} - \xi_{1j})^2 + (\xi_{2i} - \xi_{2j})^2 \right] \quad (6.3)$$

The demand received by each shop k_j is denoted $d(k_j)$. The profit of each shop k_j , i. e. $\prod k_j$, is written:

$$\prod k_j = d(k_j) \cdot (p_j - C(n_j)) - D \quad (6.4)$$

where $C(n_j)$ is the unit production cost of a firm possessing n_j shops and D is a fixed cost. The total profit of the firm j is the sum total of the profits of its shops, i. e.:

$$\prod(j) = \sum_{k_j=1}^{k_j=n_j} \prod k_j \quad (6.5)$$

The unit production cost depends solely on the number of shops owned by the firm and decreases as this number rises

$$C(n_j) = C \cdot \frac{\gamma + e^{-n_j}}{\zeta + e^{-n_j}} \quad (6.6)$$

where C , γ and ζ are constants and $\zeta > \gamma$. For a shop to be sold, the price p_j must be higher than the production cost.

6.3.3.2 Strategies about shops

Each firm possesses equity capital EC_j , which

- increases thanks to the income from bonds, i. e. $d \cdot EC_j$, where d is the rate on the bond market, and thanks to the profit made in each period;
- decreases because of the payment of corporate income tax (on profit and income from bonds) at rate c .

Thus, equity capital changes from one period to the next according to the equation:

$$EC_j(t) = EC_j(t-1) + (1-c) \cdot [d \cdot EC_j(t-1) + \prod_j(t-1)] \quad (6.7)$$

The procedure followed by each firm j active on the market during period t is as follows. Any shop which has made a loss in the last two periods, i. e. $\prod_{k_j}(t-1) < 0$ and $\prod_{k_j}(t) < 0$, is closed. The cost of closing CD is identical for all shops. A firm which has no shops left at the end of this closing process is no longer active in the market but survives as a legal entity as long as it possesses positive equity capital, namely $EC_j > 0$. If the firm j has not closed any shop during period t , it examines the different strategies for growth: the installation of a new shop in a free location or the purchase of a shop from a competitor.

The opportunity to set up a new shop is appraised by the firm after evaluation of the demand which the shop could receive. This evaluation is made using a random sample of competitors' consumers, by comparing the past utility of each of these consumers with the utility he would obtain by

going to the new $(n_j + 1)^{th}$ shop. The estimated demand, calculated for each free point on the grid, is denoted $d(n_j + 1)$. The extra profit that this $(n_j + 1)^{th}$ shop would procure for the firm j is written:

$$\prod(n_j + 1) = d(n_j + 1)[p_j - C(n_j + 1)] - D + \sum_{k_j=1}^{n_j} d(k_j)[C(n_j) - C(n_j + 1)] \quad (6.8)$$

where the second term indicates the reduction in production costs generated in all the shops of j by the increase in the number of shops.

However, for the installation of a new shop to be possible, two conditions must be satisfied. Firstly, the installation cost CI must be lower than the equity capital EC_j . Secondly, the discounted expected profit of the shop

$\prod \frac{n_j + 1}{1 - \lambda}$, where λ is the discount rate, must be greater than CI .

To evaluate a purchase opportunity, the firm j searches for a shop which would provide it with sufficient expected profit and which the competitor is actually willing to sell. The latter condition is satisfied if j offers its competitor l a price p higher than the discounted expected profit of this shop k_l , in other words $p > \frac{\prod(k_l)}{1 - \lambda}$, $\prod(k_l)$ being the last profit made in this shop by l . For j to find the operation acceptable, the same two conditions as required for the installation of a new shop must be satisfied, but with CI replaced by p .

At the conclusion of these calculations, the firm j :

- does nothing, if its equity capital is insufficient for the installation of a new shop or the purchase of an existing one;
- chooses the more advantageous operation, either the best installation or the best purchase, if it has the necessary equity capital.

In each period, the firms appear successively in the market, starting with the active firms.

6.3.3.3 Main results

The results obtained suggest that the structure of a market for a product can change endogenously without any change taking place in the nature of the competing products, because of the firms' strategies concerning assets. Two essential factors influence the structure of the market: the cost of installation CI and the fixed cost D .

When the cost of installation is high, the geographical space rapidly becomes blocked with shops. The shops have a long lifespan and belong to firms which also have a long lifespan in the market. In this case, the majority of movements involve firms trying to establish themselves in the market. They have little success in this endeavour, and consequently enter and leave the market frequently. As the strategy of growth is costly, the largest firm never reaches a very great size. Once the space is blocked, the dominant strategy concerning assets is the purchase of shops from competitors.

When the cost of installation is low, the space is less blocked, as spaces are regularly freed by the closure of shops. If purchasing becomes the dominant strategy after a certain time, it does not totally exclude the installation of new shops. Movements are more likely to be due to changes in the ownership of shops than to firms entering and leaving the market. Firms explore the space more easily and may well install shops in locations that would be more profitable for other firms, and these other firms will consequently seek to buy these shops. The maximum size of firms is higher because a successful firm can expand more easily, as the strategy of growth is less costly.

When the fixed cost is high, the profitability of shops is low and equity capital is similarly low. Installations and purchases are therefore rare. Movements involve firms rather than shops.

When the fixed cost is low, the market becomes blocked fairly rapidly and purchasing shops becomes the dominant strategy. Movements involve shops rather than firms, as the latter are easily profitable.

To sum up, a rise in the cost of installation will reduce the number of shops, have little influence on the number of firms, increase the lifespan of shops, reduce the maximum size of firms, reduce the rotation of shops and increase the rotation of firms. A rise in the fixed cost will reduce the number of shops and the number of firms, reduce the rotation of shops and increase the rotation of firms.

More generally, variation in the number of firms and in the number of shops is all the higher when the cost of installation and the fixed cost are low. Moreover, this variation usually takes a cyclical form. For example, very active purchasing phases follow on from calm phases. This cyclical aspect derives from the endogenous and autonomous dynamic of the assets market, in which the driving force is provided by the relation between assets operations and changes in the equity capital of firms. When a firm sells a shop, it increases its equity capital, thus widening its scope for potential action. A cycle of purchases then tends to develop. Phases of com-

petition and of monopolisation of the market can also be observed to succeed one another.

6.3.3.4 The possibility of merger between two firms

The model can be extended by introducing a new strategy concerning assets: the merger of two firms. This strategy is compared with installation and purchase by following the same procedure as before, in other words by comparing the discounted expected profits of the different options.

In this model, merger is the result of an unfriendly takeover bid. When a firm j evaluates the opportunity of merging with a firm l , it calculates the additional profit \prod_{jl} expected from this merger by adding the demand effect and the cost effect.

$$\prod_{jl} = \frac{1}{1-\lambda} \left[\prod l + \sum_{k_j=1}^{n_j} d(k_j) \cdot (C(n_j)) - C(n_j + n_l) + \sum_{k_l=1}^{n_l} d(k_l) \cdot (C(n_l) - C(n_l + n_j)) \right] \quad (6.9)$$

For this takeover bid to be viable, the firm j must be able to pay the firm l the value of its discounted profit, i. e. $\frac{\prod(l)}{1-\lambda}$, as well as sustaining a merger cost, denoted $CMER$. The equity capital of the firm j must be such that $EC_j \geq \frac{\prod(l)}{1-\lambda} + CMER$. The firm j examines the potentialities of the different firms with which it has the resources to merge, and considers as candidate for this merger the firm which maximises expected additional profit after payment of the purchase and the merger cost. If the firm decides to merge, it can then choose whether to supply the two products or uniquely its own.

As regards the influence of merger costs, the following results are obtained:

- the possibility of merger reduces the number of shops;
- the higher the merger cost, the stronger the tendency towards monopolisation;
- the possibility of merger leads to a reduction in the diversity of the supply of products, as merged firms tend to supply only some of the products they are capable of producing;
- mergers reduce the rotation of both shops and firms; the merger cost does however have a negative influence on the movements of firms;

- the blocking-up of the market, made possible by a low merger cost, leads to rapid and almost exclusive domination by the strategy of purchase.

6.4 Theses and conjectures

This chapter has depicted competition as a dynamic process that select among the effective actions of firms rather than a static equilibrium state associated with an efficient allocation of resources among firms. Hence, this chapter proposes that an evolutionary approach to competition abandons the concept of equilibrium and its related idea that firms are in competition if they take into account in their optimization calculus the fact that there are other firms in the market trying to capture custom as well under conditions known perfectly by the optimizing firm. In contrast it offers a view of competition as the process that forces firms to search for actions and strategies that allow to take away or maintain custom from other firms.

The three models exposed in this chapter shed partial light on the qualitative results expected from an evolutionary approach. They bring the following insights:

First, *a decentralized tatonnement can generate multiple market structures and outcomes for given "initial conditions"*. The crucial element in determining which outcomes will be generated from an interaction process is the diversity of the effective behaviors present in the market. For instance, the first model suggests that the structure of prices differs substantially whether there is a competitive retailer or not, in particular in regards of the margins that the firms extract. The structures that emerge shall be characterized by their degree of "order", i. e. the degree with which the actions that drive to success have been sorted out by the process of competition. Contrary to equilibrium, order is about sorting out through a dynamic and ex post process actions that will reveal successful and will prosper at the disadvantage of unsuccessful actions. The market "order" that shall emerge from the exploration by the firms of the possible behaviors can be more or less stable. It is stable if it is insensible to the introduction of a new behavior or a new technology; it is unstable if it is on the contrary sensible to such an introduction.

Second, *the introduction of novelty as a fuel to competition changes drastically the prediction relative to the relation between degree of competition and market structures*. Different strategies over the introduction of novelty (innovation versus imitation) generate different market structures

and outcomes. For instance, the second model suggests that concentrated structures can be the result of competition rather than the result of the absence of competition as suggested by the standard framework. The introduction of novelty is about making rival products obsolete or introducing a more efficient technology that gives a cost advantage over rivals. If the novel idea is unique and successful and if no imitator or innovator is able to compete with it, the original innovator shall serve the market and possibly earn a monopolistic rent. If imitation occurs or alternative innovating products are introduced, the market structure shall become less concentrated over time.

Third, the *outcome of competition depends also on the ability of firms to build capacities and to access to resources in other markets*. The ability of firms to be competitive in their final markets shall in part depend on their differential ability to compete for resources (qualified labor, raw material, ...) in upstream markets. This induces to take precisely into account the interdependence between markets when modeling competition among firms.

More generally, some conjectures about the role of competition can be suggested. On one hand, when it is not disturbed by other phenomena, competition plays a “positive” role in several ways: it provokes the exploration of a wide set of possibilities; it obliges firms to adopt the best possible adjustment of their actions; it eliminates the least efficient behaviors and firms. On the other hand, competition presents at least one danger: it may eliminate creative behavior, which is a vector of medium and long term economic progress, if other behaviors are more rewarding. Lastly, competition is not sufficient to guarantee the presence of behavior enabling maximum economic efficiency to be attained. Firms may seek to maintain their market share but none of them may be interested in capturing custom away from its competitors.

There is still a lot of future work needed to complete an evolutionary conception of competition. Among the types of model that are lacking; one can mention (a) models which reconcile a modeling of the decentralized tatonnement of firms in search of successful behaviors (as done in the first model) and a modeling of the introduction of change by firms (as done in the second model; (b) models which reconcile the building of capabilities by firms (as done in the third model) and its relations with the nature of competition and the emergence of aggregated structures (prices, markets, ...). A crucial stake of those models shall be to identify patterns of competition that would have some genericity and would allow predictions. Predictions cannot be made on the precise behaviors that firms will seek to implement and test since they are by definition unpredictable and various,

but predictions should certainly be made on the patterns of competition likely to be encountered in markets.

References

- Alchian, A. (1950): "Uncertainty, evolution, and economic theory", *Journal of Political Economy*, 58.
- Baumol, W. (1959): *Business behaviour, Value and Growth*, Mc Millan, New York.
- Caron-Salmona, H. and Lesourne, J. (1985): "Dynamics of a retail market with search processes", *European Economic Review*, 31.
- Chiappori, P.A. (1984): "Sélection naturelle et rationalité absolue des entreprises", *Revue Economique*, 1.
- Coursac, I., Février, P., Malandain, J. et Roy, V. (1998): "Etude d'une dynamique endogène de localisation et croissance des firmes", Report under the direction of J. Lainé and G. Laffond, ENSAE.
- Cyert, R.M. and March, J.G. (1963): *A Behavioral Theory of the Firm*, Prentice-Hall, Englewood Cliffs.
- Diamond, P.A. (1971): "A model of price adjustment", *Journal of Economic Theory*.
- Dutta, P.K. and Radner, R. (1999): "Profit maximization and the market selection hypothesis", *The Review of Economic Studies*, 66.
- Fisher, F.M. (1970): "Quasi competitive price adjustment by individual firms: a preliminary paper", *Quarterly Journal of Economics*.
- Friedman, M. (1953): *Essays in Positive Economics*, Chicago Press.
- Hayek, F. von (1968): "Competition as a discovery procedure", Series *Kieler Vorträge*, N°56, University of Kiel.
- Nelson, R.R. and Winter, S. (1982): *An Evolutionary Theory of Economic Change*, The Belknap Press of Harvard University Press.
- Reinganum, J.F. (1979): "A simple model of equilibrium price dispersion", *Journal of Political Economy*.
- Salmona, H. (1985): "Equilibre sur le marché d'un bien en information imparfaite: une analyse de la littérature", *Economie Appliquée*, 3/4.
- Salop, S. and Stiglitz, J. (1977): "Bargains and ripoffs: a model of monopolistically competitive price dispersion", *Review of Economic Studies*.
- Salop, S. and Stiglitz, J. (1982): "The theory of sales: a simple model of equilibrium price dispersion with identical agents", *American Economic Review*.
- Schumpeter, J. (1942): *Capitalism, Socialism and Democracy*, New York, Harper.

-
- Simon, H. (1955): "A behavioural model of rational choice", *Quarterly Journal of Economics*, 69, 99-118.
- Simon, H. (1959): "Theories of Decision Making in Economics", *American Economic Review*, 49, 255-256.
- Stiglitz, J. (1989): "Imperfect information in the product market", in Schmalensee R. et Willig R.D. (eds.), *Handbook of Industrial Organization*, vol.I, Elsevier Science Publishers.
- Tinter, (1941): "The Theory of Choice Under Subjective Risk and Uncertainty", *Econometrica*, vol.9, pp.298-304.
- Tirole, J. (1995): *Théorie de l'organisation industrielle*, Economica.
- Vickers, J. (1993): "Concepts of competition", *Oxford Economic Papers*.
- Williamson, O.E. (1964): *The Economics of Discretionary Behaviour: Managerial Objectives in a Theory of the Firm*, Prentice Hall, N.J.
- Winter, S. (1964): "Economic natural selection and the theory of the firm", *Yale Economic Essays*, 4.
- Winter, S. (1971): "Satisficing, selection and the innovating remnant", *Quarterly Journal of Economics*, 85.

7 Organization of the firm

This chapter examines one of the most promising and lively areas of research within evolutionary theory, the one concerned with the study of organizations and business firms in particular. In their book Nelson and Winter devoted an important part of their analysis to business firms, also referring to the Schumpeterian recognition of their central role in determining the pace and direction of technological change.

This chapter briefly presents the basics of the evolutionary perspective on firms and organizations and some examples of a growing family of models.

7.1 Background and problems

Standard neoclassical theory identifies firms with production functions. Production possibilities are perfectly known and the production function is the efficiency frontier which separates known from unknown techniques. In this framework technological change is a sort of magic by which “certain production possibilities are suddenly extracted from the Unknown and added to the Known. There is ordinarily no analysis as to why particular possibilities rather than others should be thus discovered” (Winter, 1968). The reduction to production functions prevents any serious analysis of the role of firms as fundamental loci where productive technical knowledge is generated, stored, socialized, transmitted and, especially, modified and, therefore, as major sources of technological changes (Schumpeter, 1955). Because it puts a particular emphasis on technological, organizational and, in general, economic change, evolutionary theory assigns a central position to theorizing about firms and their internal mechanisms which generate such changes (Nelson-Winter, 1982).

Also neoclassical theory has recently questioned and overcome the most radical standard assumptions and developed more sophisticated theories which open the firm’s black box. Agency theory has studied owner-manager-worker contractual relations acknowledging that firms are made of a multiplicity of agents with heterogeneous objectives and asymmetric information and that their actions must be guided by appropriate incentives

in order to make them comply with the organizational objectives. Transaction costs economics on the other hand has moved from the recognition of bounded rationality and contract incompleteness and analyzed firms as coordination devices alternative to and complementary with the market.

Evolutionary economics acknowledges the importance of information asymmetries in explaining organizations and their structure, and certainly shares with transaction costs economics the emphasis on bounded rationality, but stresses the cognitive aspect of organizations which is mostly neglected by those theories. Firms are there first of all to do things. They produce, store, transmit, socialize and modify the knowledge which is necessary to produce (useful) things (Winter, 1982). Such knowledge cannot normally be detained by a single individual, nor can it be normally put together by simple market coordination, but must be “organized”. Market mechanisms do not in fact convey enough information and knowledge to manage complex distributed knowledge, moreover an important part of the relevant knowledge is tacit and requires direct interaction in order to be apprehended, transmitted and shared.

Evolutionary theory is thus trying to develop theories and models based upon knowledge, cognition and learning, the latter seen not simply as information processing but as structural modification of cognitive representations, action repertoires and organizational architectures, in other words as genuine technological and organizational innovation.

Evolutionary theory in general, and the part concerned with organizations in particular, strongly believes in bounded rationality, usually in a rather strong version thereof. In this and other respects evolutionary theory has strong connections with the behavioural theory of the firm (March-Simon, 1958, Cyert-March, 1963). Agents and organizations face problems they only partly understand and have to build representations of such problems and problem-solving procedures which are inevitably imperfect and continuously subject to revision through learning. An important part of problem-solving is the activity of subproblem decomposition (Simon, 1981): large, complex and new problems have to be decomposed into smaller, more manageable and familiar ones. In organizations problem decomposition maps into division of labour and modularization of artifacts (Baldwin-Clark, 2000) and is usually a precondition for the formation of new markets. Most of the time in fact it is the processes of division of labour, standardization and modularization within organization which creates new technological interfaces whose coordination can then possibly be transferred to the market.

7.2 Canonical principles

In standard neoclassical theory agents behave according to the principles of utility maximization and coordination among them can only be achieved by putting into place the appropriate monetary rewards. This is normally done by market forces, but whenever the latter cannot operate optimally, alternative (second-best) contractual arrangements are devised within organizations. In evolutionary theory instead intra-organizational coordination is mainly carried out by routines (Nelson-Winter, 1982). According to Cohen et al.: “A routine is an executable capability for repeated performance in some context that has been learned by an organization in response to selective pressure” (Cohen et al., 1996).

The first key term, capability, is defined as the capacity to generate an action pattern that has been stored in localized or distributed form within the organization. As Nelson and Winter suggested, this capacity entails the ability to know what action pattern to perform and when to perform it. This implies that individual agency within organizations is not only guided by the logic of consequence of utility maximization but also, and mainly, by the logic of appropriateness (March, 1994), according to which members of the organization act according to rules they believe are appropriate in their role and given a situation they recognize as familiar.

This also implies that organizational members must be able to receive and interpret inputs generated, within a specific context, both by humans and non-human devices. Inputs may include, for example, formal orders from a superior, informal suggestions, descriptions of the situation, the activation of another action pattern by other members, a particular date on the calendar or a particular time on the clock (Nelson-Winter, 1982). While the ability to receive inputs could be independent of the specific organizational context, the ability to recognize and interpret messages and inputs – as well as to generate appropriate messages and input – and to recall the appropriate pattern of action, is part of the knowledge and problem-solving repertoire that the single member learn within the organization. This view of routines as stored and shared problem-solving skills is commonly, and fruitfully, adopted in both theoretical and empirical research on the origin, development and change of organizational capabilities.

At the same time, growing effort has gone also into the formal representation of processes of search, recombination, reinforcement of sequence of elementary operations yielding particular problem-solving procedures. However, routines emerge and are implemented in organizations com-

posed of a plurality of individuals who might have conflicting interests. Certainly, a firm can be considered as a hierarchy of routines. This hierarchy, however, also entails a mechanism of exercise of authority and governance of the admissible behaviours by which individual members can pursue their interests. This is indeed acknowledged by Nelson who suggest that routines can be seen also as “truces” amongst potentially conflicting interests, but this complementary nature of routines has been so far relatively neglected in that literature which explicitly builds upon evolutionary ideas.

Finally, as concerns methodology, evolutionary modelling about organizations has received a great impetus from the development of agent-based simulation techniques, which are at the heart of all the models presented in this chapter. Computer simulations are very useful in at least two respects: they allow both for the construction of “behaviourally rich” models and for the study of explicit mechanisms of interaction among (heterogeneous) agents.

The former means that once we abandon the (expected) utility maximization assumption, we can introduce a variety of more plausible and richer behavioral assumptions. These can be normally expressed algorithmically and their properties and implications studied by means of computer simulations. Typical examples are adaptive algorithms of various kinds for decision making and learning, such as genetic algorithms, neural networks, classifiers systems and many others.

The latter involves that, by means of computer simulations, we can implement and study rich interaction environments in which a multiplicity of heterogeneous agents can explicitly interact through a variety of organizational rules. This possibility is very important for evolutionary models in general, because heterogeneity is a fundamental property of evolution, which is meaningless in a representative agent reduced form, and because out-of-equilibrium dynamics is usually more important than asymptotic properties. But this is even more important in models of organizations, because organizational behaviours are themselves the result of the interaction among heterogeneous agents. The rules of interaction define the organizational structure similarly to the definition given by Herbert Simon: “[...] [the structure of an organization] designates for each person in the organization what decisions that person makes, and the influences to which he is subject in making each of these decisions” (Simon, 1976).

7.3 Some models

We will examine four models: the former concentrate on information processing features, the latter on structural evolution.

7.3.1 An information processing model

The properties of organizations have long been investigated outside an evolutionary perspective: among the many examples, the most notable are Aoki (1986), Mount-Reiter (2002), Radner (1992), Sah-Stiglitz (1986)²⁵. Generally speaking, such models consider a set of information processing units with some form of bound in their computational capabilities and study optimal organizational structures which maximize aggregate performance under the assumption of a commonly shared objective function. For instance, consider individual agents who must evaluate projects, but have a limited ability to do so. They may incur into two types of errors: approve a bad project or reject a good project and the aggregate error is analyzed for different organizational arrangements. In particular, Sah compare basic hierarchical and decentralized structures and show that the former reject more projects (including good ones) than the latter, while the latter accept more projects (including bad ones) than the former²⁶.

John Miller develops a model in the same vein but within an adaptive/evolutionary framework (Miller, 2001). Organizations are tree-like structures in which each node represents a basic information processing capability (e. g. a processor which given two integers delivers their sum) and edges are communication channels in which information can flow between the two connected nodes. All organizations possess a root node, which has to deliver the final result. Then the organization is built recursively: some “child” node(s) are probabilistically created which are connected to the root and with some probability the root node also has an edge connecting it to the input data queue. The same process is repeated recursively for each newly created child node. Nodes which do not have any “child” are terminal nodes and are automatically connected to the input data queue.

²⁵ Radner (2000) provides a general discussion of information processing and decision making in teams.

²⁶ Interesting extensions and refinements have been recently suggested by Christensen-Knudsen (2002) and Ioannides (2003).

Organizations face a set of fully decomposable²⁷ problems and their performance is given by the speed at which they are able to deliver a correct solution.

Figure 7.1 presents a simple example of an organization, and its functioning can be described as follows. Nodes are activated either in a random or ordered sequence, in the latter case activation begins with the highest numbered node. When activated, a node which is attached to the data queue picks a number – if any – from it and adds it to its memory, when this operation has been finished the memory content is transmitted to the parent node. A node which is not attached to the queue waits for input from its child node(s). As an example, suppose that the organization in figure 7.1 has to sum up five integer numbers $\{a, b, c, d, e\}$ contained in the data queue and activation is ordered. At the first stage nodes 1, 2 and 3 accumulate, respectively, c , b and a ; then, at the second stage, node 3 picks d and adds it to a , node 2 picks e and adds it to b , node 1 waits because the data queue is now empty. At the next stage node 3 communicates $a + d$ to its parent, node 1, which adds it to the content of its memory, which now becomes $a + c + d$. Finally, node 2 transmits the content of its memory, i. e. $b + e$ to node 1, whose memory content becomes $a + b + c + d + e$. The task has been correctly completed in three time steps.

Miller studies the generic properties of randomly generated organizations and then introduces an adaptive process for the analysis of organizational evolution. In the first set of simulations a population of organizations are randomly generated with a number of nodes which can range from 1 to 50 and their performance is measured by the time steps they require to successfully complete one or more addition problems under either random or ordered node firing. The first interesting result is that even such randomly generated organizations exhibit “order for free”, that is they are capable of solving problems with a good performance. In particular, the following results are noteworthy. First, there are economies of scale with respect to the number of problems: when organizations solve multiple problems, the average time required for each of them decreases – especially for organizations with a high number of nodes – because new problems can be processed by lower nodes while older problems have moved to higher level nodes. Second, ordered firing is more efficient than random firing, but the difference between the two vanishes as the number of problems increases.

²⁷ See section 7.3.3 below for models in which organizations face non-decomposable or nearly-decomposable problems.




Fig. 7.1. Miller's model

Third, concerning the relation between number of nodes and performance, it can be shown that small organizations (with four or fewer nodes) tend to have high variance in their performance, regardless the firing procedure. With a number of nodes above eight the variance becomes instead rather low and generic organizations of the same size tend to display very similar performance. Under ordered firing there is a downward trend in the relationship between the number of time steps and the number of nodes, while under random firing this relation is U-shaped and the computation time reaches its minimum when the number of nodes is eight.

In a further set of exercises, Miller lets the population of randomly generated organization evolve according to an algorithmic procedure which resembles John Holland's genetic algorithms (Holland, 1975) and John Koza's genetic programming (Koza, 1993). Each organization is assigned a "fitness" measure, which is the reciprocal of the time required to solve the problem(s). Fitter organizations survive and produce copies of them-

selves, while bad performing ones tend to be suppressed. Moreover good performing organizations also produce offsprings which are modifications of themselves via two genetic operators. The first genetic operator, crossover, consists of taking two organizations (chosen with probability proportional to their fitness) and creating two offsprings which are copies of the parents except that two randomly chosen sub-trees are swapped. The second genetic operator is mutation which can consist of randomly activating or deleting a connection between one node and the data queue, or adding a terminal node, or deleting an existing node. Then the fitness values for the new population are measured and the process is repeated.

Miller presents some simulations of this model of evolving organizations, whose main results can be summarized as follows. First, evolution does indeed improve considerably performance. After fifty generations the best evolved organizations display an improvement of performance ranging from 28 percent (multiple problems and random firing case) to 47 percent (single problem and ordered firing case) relatively to the randomly generated organizations.

Second, this improvement can be explained partly by the selection mechanism, which in each generation makes copies of the best performing organizations and eliminates the worst ones, and partly by the operation of the two genetic operators, crossover and mutation, which recombine and modify “at the margin” the structural characteristics of the best performing organizations.

Third, in some cases evolution tends to produce, starting from the high heterogeneity of a random first generation, rather homogeneous structural characteristics: this is the case of single problem and random firing, where evolved organizations possess on average only three nodes (standard deviation 1.1); of multiple problems and ordered firing, where on the contrary evolved organizations tend to have a number of nodes which approaches the maximum of 50 (48 on average with 1.9 of standard deviation); of multiple problems with random firing, with an average number of nodes of 12 and standard deviation 2.4. On the one hand in fact additional nodes can improve performance of the organization because they increase computational power, but on the other hand they increase the number of levels and thus information transmission costs. The latter of course are particularly important in the random firing cases, where in fact evolution pushes towards organizational structures with few nodes.

All in all these results show a trade-off between the higher processing power of larger organizations and the higher coordination costs they imply and diminishing returns are quickly reached to the addition of new nodes.

7.3.2 An organizational learning model

The model previously presented belongs to the family of models of information processing organizations, in which the information processing capabilities of individual agents are given. Marengo (Marengo, 1992 and Marengo, 1996) present a model which focuses instead upon the modification of such information processing capabilities, i. e. a process of structural learning. Individual agents are imperfect adaptive learners, as they adjust their information processing capabilities through local trial-and-error. Of course this adaptive learning is driven by information coming from the environment and/or from other members of the organization. The model shows that the architecture of such information flows plays a crucial role in determining the learning patterns and the performance characteristics of the organization.

Consider the following coordination problem: an organization has to respond to an exogenous environment by implementing some collective action. Suppose for instance that a firm can produce a certain number of product types, which are demanded by an exogenous market, and that the production process is divided into several operations, each being carried out by a different shop. The problem is therefore to detect correctly which product type is being demanded (state of the world) and to coordinate the actions of the shops so that the correct production process is implemented.

More specifically, suppose that there exist eight possible product types, called respectively “1”, “2”, ..., “8”. The firm’s production possibilities set is represented by sequences of eight operations and each operation can be of two types (P and Q). Such sequences have all the same length and map into a product type, which is conventionally designated by the number of operations of type “ P ” which are used in its production (the product of type “ q ” is produced when q operations are of type “ P ”). The problem of the firm is therefore to forecast the product type which will be demanded by the market and to implement the correct production process by coordinating the operations of the two shops. The payoff is the following: if the firm produces the correct product type it receives a payoff of 5 units; if it does not produce the correct output it receives a negative payoff, given by the difference between the actual product type and the required one (for example, if the market demands type “7” but the firm produces type “5”, it will receive the payoff -2).

Suppose now that the firm is formed of three units, the management and two shops and is in relation with some environment. The shops proceed to the production operations (four operations each). The exchanges




Fig. 7.2. Structure of firm

of information are given in the figure 7.2. The management observes the environmental message by channel 1(the actual state of the world), interprets it and sends a message to the two shops by channels 2.

Each of the two shops observes three kinds of signals and develops an interpretative model for them. The environmental signal (last observed state of the world) follows channel 1B, the message sent by the management (and based on its own interpretation of the environment) follows channel 2, and the signal sent by the other shop (i. e. its last action) follows channel 3. The latter two messages are coordinating devices, respectively a centralized and a decentralized one, which allow the shops to coordinate their action, whereas the former allows the two shops to form their own independent (from the management’s) model of the world.

The manager and the shops face a standard learning problem (see chapters 1 and 3). Let $S = \{s_1, s_2, \dots, s_N\}$ be the set of the N possible states of nature and $A = \{a_1, a_2, \dots, a_K\}$ the set of the K possible actions the decision-maker can undertake. The payoff to the agent is given by a function: $\pi : A \times S \rightarrow R$, where the agent’s payoff to action a_i when the state of the world s_h occurs will be indicated by π_{ih} . The action the agent chooses depends obviously on the level of his/her knowledge about the state of nature. The agent’s knowledge (or information processing capabilities) can

be represented by a collection of subsets $P(s_i) \subseteq S$ where $P(s_i)$ is the set of states of nature which the agent considers as possible (or cannot tell apart) when the real state is s_i .

The (reinforcement) learning model which is considered is based on classifiers systems (Holland, 1986), but with some substantial differences and simplifications. The basic component is a condition-action rule R_j , the condition and the action being represented by strings of symbols:

$$R_j : c_1, c_2, \dots, c_N \Rightarrow a_1 a_2 \dots a_k \quad (7.1)$$

with $c_i, a_h \in \{0,1\}$.

The execution of a certain action is conditional upon the agent's perception that the present state of nature falls in the given condition. The condition part c_1, c_2, \dots, c_N is a "category", that is a subset of the states of nature (or an 'event'). It is activated when the actual detected state of the world falls in such a subset. Practically, it is satisfied whenever the actual state of nature corresponds to a position where a "1" appears: if s_i is the last observed state of the world, we have $c_i = 1$. The action part $a_1 a_2 \dots a_k$ corresponds to a given action. Practically, it contains a unique symbol a_h equal to 1 (all others being equal to 0) which corresponds precisely to the chosen action. The decision maker can be therefore represented by a set of such condition-action rules: $R = \{R_1, R_2, \dots, R_q\}$.

In addition, each rule R_j is assigned (at date t) a "strength" $St(R_j, t)$ and a "specificity" $Sp(R_j)$. Strength basically measures the past usefulness of the rule, that is the rule's cumulated payoff. Specificity measures the strictness of the condition: in our case the highest specificity (or lowest generality) value is given to a rule whose condition has only one symbol "1" and therefore is satisfied when and only when that particular state of nature occurs, whereas the lowest specificity (or the highest generality) is given to a rule whose condition is entirely formed by "1" and is therefore always satisfied by the occurrence of any state of nature.

At the beginning of the process, the decision maker is supposed to be completely ignorant about the characteristics of the environment he or she is going to face: all the rules initially generated have the highest generality. The action parts are instead randomly generated. The decision maker is also assumed to have limited computational capabilities, therefore the number of rules stored in the system at each moment is kept constant and relatively small in comparison to the complexity of the problem which is being tackled.

This set of rules is processed in the following steps throughout the simulation process:

1. Condition matching: a message is received from the environment which informs the system about the last state of nature. Such a message is compared to the condition of all the rules and the rules which are matched, 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 designate the one which is allowed to execute its action. 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): $Bid(R_j, t) = (k_1 + k_2 Sp(R_j)) St(R_j, t)$ where k_1 and k_2 are constant coefficients. The winning rule is chosen randomly, with probabilities proportional to such bids.
3. Action and strength updating: the winning rule executes the action indicated by its action part and 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 l -th rule is the winner of the competition, we have: $St(R_l, t + 1) = St(R_l, t) + \pi(t) - Bid(R_l, t)$
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 already existing and most successful rules, introduce new ones which could improve the performance of the system. Thus, new rules are constantly injected into the system and scope for new opportunities is always made available.

Genetic operators generate new rules which explore other possibilities in the vicinity of the currently most successful ones, in order to discover the elements which determine their success and exploit them. Search is not completely random but influenced by the system’s past history. New rules take the place of the currently weakest ones, so that the total number of rules is kept constant. Two genetic operators have been used for the condition and one for the action part. The latter is a simple type of local search and is simply a mutation in the vicinity: the action prescribed by the newly generated rule is chosen (randomly) in the close proximity of the one pre-

scribed by the parent rule. The former deserve more attention since they operate in opposite directions:

- specification: a new condition is created which increases the specificity of the parent one. Wherever the parent condition presents a 1, this is mutated into a 0 with a given small probability;
- generalization: the new condition decreases the specificity of the parent one. Wherever the latter presents a 0, this is mutated into a 1 with a given small probability.

Different degrees of specification and generalizations can be simulated. They rely on different probabilistic combinations of the two genetic operators and on the variation of the coefficient k_2 in the bid equation: the higher this coefficient, the more highly specific rules will be likely to prevail over general ones. A specificity coefficient summarizes the overall inclination of the system toward the search for specific rules.

Coming back to the firm's structure, the weights with which the three types of messages (from environment, management or other shop) enter the shops' decision processes define the balance between differentiation and commonality of knowledge. Such weights are represented by the specificity coefficients which express the agent's search for a precise model which interprets the corresponding type of message. A high specificity coefficient for messages 1B coming from the environment implies that the shops aim at building a detailed and autonomous model of the world. A low coefficient implies instead that shops do not pay much attention to the environment, but rely more on the world's interpretation given by the management. A high specificity coefficient for messages 2 coming from the management implies that shops attribute great importance to the correct interpretation of the coordinating messages which are sent by the management. A low coefficient implies instead that shops are not seeking centralized coordination on the organizational collective knowledge, the management having a reduced role. Finally, a high specificity coefficient for the messages 3 coming from the other shop implies that shops are attaching high importance to mutual, decentralized coordination. A low coefficient means that the organization is without any form of decentralized coordination, i. e. no inter-shop communication.

Marengo (Marengo, 1992 and Marengo, 1996) present a set of simulations in various environments, whose main results can be summarized as follows. In stationary environments, i. e. when the state of nature does not change, agents can achieve coordination without building any model of the

environment and resorting only to trial-and-error with selection. If instead they try to learn, i. e. to build such a model and constantly improving it, they need also to learn a model for the interpretation of coordinating messages: messages 1 and/or 1B are not sufficient, and messages 2 or 3 are also needed. If the environment undergoes cyclical and predictable changes, high specificity coefficients on the shops' conditions which classify environmental messages 1B are needed in order to exploit the environmental regularity. Shops need to have a direct access to environmental information in order to develop the necessary decentralized learning. Finally, if the environment undergoes frequent and unpredictable changes, the organization has to develop stable routines which give a "satisficing" average result in most conditions. In this case, decentralized learning is detrimental, because the stability of such routine is continuously jeopardized by individual efforts to grasp the unpredictable environment; shops should rely on the management's message.

All in all, in order to exploit a regularly changing environment, a high amount of knowledge about the environment itself is required: the model must distinguish between the states of nature and connect them diachronically. It is not surprising therefore that the most appropriate organization in such circumstances is the one which, by partly decentralizing the acquisition of knowledge about the environment, can achieve higher levels of sophistication in its model of the world, provided the coordination mechanisms – which are here centralized – are powerful enough to enable the organization to solve conflicts of representations. On the other hand, this very decentralization of the acquisition of knowledge can be a source of loss when it is more profitable for the organization to cling to a robust and stable set of routines. This situation requires strong coordination in order to make the entire organization implement coherently such a set of robust routines. Autonomous and decentralized experimentation can only disrupt such a coherence.

7.3.3 A structural evolutionary model

In the last few years, a new family of evolutionary models of organizations has developed, inspired by biologist Stuart Kauffman's so-called "NK model" (Kauffman, 1993). The "NK model" is a model of selection and adaptation in complex environments where the evolving entities are not "simple" and uni-dimensional, but are complex structures, made up of many components which interact among each other non linearly. Kauffman shows that in this case the outcome of evolution is the result of the in-

terplay between selection forces and the mechanisms that regulate the inner working of the entities under selection, i. e. the network of interactions among components. If the latter are “strong” and “diffused” (in a sense we will specify shortly), then the power of selection is highly limited and evolutionary dynamics is mostly determined by the structural characteristics of the entities, displaying strong path-dependency and sub-optimality.

The interest of this kind of models for the economic theory of organization resides precisely in showing that selection is not necessarily conducive to optimality, irrespective of the presence of informational asymmetries and transaction costs. In particular, the likelihood of reaching optimality through selection is especially low when the entities under selection are complex ones, made of many interacting components. In the following, we will first provide a short and basic presentation of the NK model (more details can be found in Kauffman, 1993, a discussion of its economic implications and limits with some extensions can be found in Frenken et al., 1999).

A NK system is described by a string (or configuration) of elements, denoted i and in number N , which refer to subsystems, components or attributes of the system. Each element i can be in one and only one possible states out of a finite set A_i (taken as $\{0,1\}$). The possibility space $A_1 \times A_2 \times \dots \times A_N$ contains all the possible strings of the system. Moreover, a fitness function F assigns a real value (taken without loss of generality in the interval $[0,1]$) to each possible string. It measures its relative performance in the environment and it depends obviously on the states of each component: $F : (A_1 \times A_2 \times \dots \times A_N) \rightarrow [0,1]$. Moreover, when a fitness value f_i can be ascribed to each element i , the overall fitness is generally considered to be the average of them:
$$F = \sum_1^N \frac{f_i}{N}.$$

The distribution of fitness values of all possible strings is called the *fitness landscape* of a system. Its appearance depends on the system inner structure captured by its K -value, i. e. the number of “epistatic” relations among elements. An epistatic relation implies that the contribution of one element to the overall fitness of the system is dependent both upon its own state and upon the state of K other elements. Two limit cases of complexity can be distinguished. Minimum complexity is obtained when elements are independent: $K = 0$, the fitness of each element depending only on its own state. The maximum complexity is obtained when elements are completely intertwined: $K = N - 1$, the fitness of each element depending simultaneously on the state of all elements. In applications, the fitness of each ele-

ment is drawn randomly from a uniform distribution between 0 and 1. If $K = 0$, the draw is unique for each element; if $K = N - 1$, it depends on the whole string.

Suppose now that a string can be modified by means of a simple “random one-bit mutation” into a neighbouring string which differs by only one bit. The new string can have a higher, lower or identical fitness value. Now we can introduce a simple selection process (or, equivalently, an adaptive learning process) by which fitness improving mutations are retained while fitness reducing ones will be discarded. In this way, we define a path (or random walk) in a fitness landscape which, starting from an initial string, “climbs” towards higher fitness values until no further fitness-improving mutations are possible. Where will such a walk end up to?

The local optima are obviously evolutionary attractors: once an agent occupies a local optimum, it cannot escape this optimum by a one-bit mutation. An important property of evolutionary attractors is the number of starting points that can end up in the attractor state (including the optimum itself). This number characterizes the size of the basin of attraction of the local optimum.

Kauffman (Kauffman, 1993) provides a general analysis of the properties of NK landscapes for all possible values of the N and K parameters. Among the many properties he finds, the following are particularly noteworthy:

1. The number of local optima increases exponentially as the K -parameter increases from its minimum value $K = 0$ (where there is a unique local optimum which is also global) to its maximum value $K = N - 1$.
2. The locations of local optima in the possibility space are correlated (they have many states of elements in common) for low K -values, but correlation vanishes as K increases towards its maximum.
3. The overall fitness values of local optima of systems with a low K value are higher on average.
4. The basins of attraction of local optima of systems with a low K are larger on average, and more generally, the higher their fitness, the larger their basin of attraction. As a consequence, an adaptive walk has higher probability of locating a local optimum with a high valued peak.

One of the earliest and most interesting applications of the $N - K$ model to organizational evolution has been presented by Daniel Levinthal (Levin-

thal, 1997), who makes the assumption that an organization can be represented as a string of (binary) features (policies, routines, standard operating procedures, etc.) linked together by complex interdependencies²⁸. A population of randomly generated organizations evolves on a fitness landscape, the evolution is driven by variation, selection and information passing processes.

Variation, i. e. the generation of variety, is provided by two mechanisms:

1. *local search*: existing organizations mutate one feature, while keeping all the other constant. If the new string has higher fitness, then the mutation is retained, otherwise it is discarded.
2. *radical changes* (“long jumps”): at times, existing organizations can mutate many (possibly all) features. If the new string has higher fitness, then the mutation is retained, otherwise it is discarded.

Selection is given by a simple birth and death process: some organizations die with a probability inversely proportional to their relative fitness and are replaced by newly born ones. Some of these new organizations are randomly generated, owing possibly no resemblance to the existing ones, while other are replica of existing successful organizations.

Information passing among generations is warranted by two mechanisms:

1. *retention*: successful existing organizations have a high probability of surviving. Their features tend therefore to survive.
2. *replication*: some of the newly born organizations which replace bad performing ones, which are selected out, are copies of the most successful existing organizations. The features of the latter tend therefore to spread in the population.

This very basic evolutionary model applied in complex environments provides some interesting results obtained by simulation (Levinthal, 1997).

Suppose first that the initial system is formed by a large population of randomly generated organizations. It evolves according to the mechanisms

²⁸ By interdependencies we mean something more general than the complementarities analyzed by super-modular models of production and organization (Milgrom-Roberts, 1990): interdependencies can in fact be non-monotonic, meaning that changing the i^{th} feature from, say, 0 to 1 may sometimes produce an increase of the overall fitness and sometimes a decrease, depending upon the values of the remaining features.

of selection and information passing just mentioned, but variation can be only local, i. e. only one bit at a time can be mutated for every organization. Local adaptation and selection will reduce the heterogeneity of the population: bad performers will be selected out and replaced by copies of good performers. In the meantime, good performers will climb with local mutations the fitness peaks they are located on. However, the final outcome of evolution will vary considerably depending upon the value of the K parameter. With $K = 0$, local adaptation will quickly take all organizations to the only (global) optimum of the landscape: thus selection and adaptation will completely wipe out the initial heterogeneity of the population and cause a fast convergence to a unique, optimal, organizational form. For higher values of K , this is less and less the case: the landscape will display an increasing number of distinct local optima on which subsets of organizations will converge according to their initial location. Selection and adaptation will reduce the initial variety but will never make it disappear.

This result, rather obvious in this framework, must not be overlooked, as it provides a simple and intuitive explanation of the persistence of heterogeneity among firms, a piece of evidence widely reported by the literature²⁹, but at odds with neoclassical theory, according to which deviations from the only best practice should be only a transient property inevitably due to fade out as market selection forces operate. Note also that as K increases, not only does the number of local optima increase, but also the size of the basin of attraction of each of them will shrink. It is possible therefore that none of the organizations is located in the basin of attraction of the global optimum and therefore no organization will ever find the globally optimum configuration³⁰.

In complex environments, diversity of forms is not only a persistent property, but can even emerge out of homogeneity. Assume that the initial condition is now represented by a population of identical organizations. Because of random local search, they start mutating different attributes and moving in different directions in the landscape. Again, if $K = 0$, the landscape is single-peaked and such diversity will quickly decrease until all organizations will converge to the unique global optimum. If K is larger,

²⁹ An example is provided, e.g. by Jensen and McGuckin (1997).

³⁰ The belief shared by many economists that selection brings optimality is, generally speaking, unjustified, as strongly argued by, among the others, Herbert Simon: „...in a relative sense, the fitter survive, but there is no reason to suppose that they are fittest in any absolute sense” (Simon, 1983).

such initial random mutations will take organizations in the basins of attraction of different local optima. Selection and adaptation will only partially reduce such diversity.

Suppose now that organizations can perform more radical changes (“long jumps”). With large K , heterogeneity tends – though very slowly – to disappear because organizations located on sub-optimal peaks can always perform – though with low probability – a radical mutation which allows them to jump on a higher fitness peak, until they reach the highest one of the global optimum. However, if N is large enough, such a possibility may have a very low probability and not make any real impact on the medium term evolution of the population.

Consider finally the case of environmental changes, which can be modeled by re-drawing the fitness contributions of some features after the population has evolved and stabilized over the local optima. Suppose first that such a change concerns only one feature. With minimal complexity ($K = 0$), the global optimum will either remain where it was or move to a point which is at most one mutation away. Thus, if the population has already evolved and is located on the global optimum, it can easily and quickly adapt and move to the new global optimum. Simulations show that all incumbent organizations survive to such an environmental change. If instead the complexity of the landscape is high ($K \gg 0$), even the modification of the fitness contribution of just one attribute can cause a large alteration of its shape. In high dimensional landscapes with a large N , local optima can move far away. This implies that a population which has settled on the local optima of the initial landscape will find it much more difficult to adapt to the change. Mortality of incumbents will rapidly rise as K increases.

If the environment changes more radically, i. e. the fitness contributions of many (possibly all) the attributes are re-drawn, we get a different picture. As we have already argued, in a “simple” landscape with $K = 0$, all organizations quickly converge to the same configuration, which corresponds to the unique global optimum and diversity dies out. If a dramatic environmental shock happens for which the global optimum moves far away from its initial position, the entire population will find itself in a low fitness area of the landscape and incumbent organizations will be outperformed by newly created ones with random configuration. If, on the contrary, K is high, the population remains distributed over a large number of local optima and there is a high likelihood that a subset of the population will find itself in or close to a high fitness portion of the landscape after the environmental shock has occurred. Preserving diversity helps the population adapt to dramatic environmental changes.

7.3.4 A model of problem-solving by division of labour

One of the main limitations of the previous model is that it reduces organizational search processes either to simple one-bit mutations or to totally random “long jumps”. Actual organizations do not search randomly nor are they confined to strictly local search. As already reminded, they exist first of all to do things and put into place various arrangements to look for better ways to do things, according to an at least partly conscious and explicit cognitive representation of the problem faced. In this respect, firms’ activities are very similar to problem-solving and the way labour is divided and organized can be analyzed as a form of problem-decomposition. Marengo and Dosi (2005) present a model for this purpose which we briefly describe.

Let us assume that solving a given problem requires the coordination of N atomic components (“elements”, “actions” or “pieces of knowledge”), each component taking alternative states. For the sake of simplicity, the number of states is reduced to two, labelled 0 and 1. Hence, the set of components is: $C = \{c_1, c_2, \dots, c_N\}$ with $c_i \in \{0, 1\}$. A configuration, that is a possible solution to the problem, is a string: $x^i = c_1^i c_2^i \dots c_N^i$. The set of configurations is: $X = \{x^1, x^2, \dots, x^{2^N}\}$. An ordering is defined over the set of possible configurations: $x^i \geq x^j$ (or $x^i > x^j$) whenever x^i is weakly (or strictly) preferred to x^j . In order to avoid technical complications, we assume, for the time being, that there exists only one configuration which is strictly preferred over all the other configurations (i. e. a unique global optimum). This simplifying assumption will be dropped below. A problem is defined by the pair (X, \geq) .

As the size of the set of configurations grows exponentially in the number of components, the state space of the problem becomes too large to be extensively searched by agents with bounded computational capabilities. One way of reducing its size is to decompose³¹ it into sub-spaces. Let $I = \{1, 2, \dots, N\}$ be the set of indexes and let a block³² $d_i \subseteq I$ be a non-

³¹ A decomposition can be considered as a particular case of search heuristics: search heuristics are, in fact, ways of reducing the number of configurations to be considered in a search process.

³² Blocks in our model can be considered as a formalization of the notion of modules used by the flourishing literature on modularity in technologies and organizations (Baldwin-Clark, 2000) and decomposition schemes are a formalization of the notion of system architecture which defines the set of modules in which a technological system or an organization are decomposed

empty subset of it, with cardinality $|d_i|$. We define a decomposition scheme (or simply decomposition) of the problem (X, \geq) a set of blocks: $D = (d_1, d_2, \dots, d_k)$ such that $\bigcup_{i=1}^k d_i = I$. Note that a decomposition does not necessarily have to be a partition.

Given a configuration x^i and a block d^j , we call block-configuration $x^i(d_j)$ the substring of length $|d_j|$ containing the components of configuration x^i belonging to block d_j : $x^i(d_j) = x_{j_1}^i x_{j_2}^i \dots x_{j_{|d_j|}}^i$ for all $j_h \in d_j$. We also use the notation $x^i(d_{-j})$ to indicate the substring of length $N - |d_j|$ containing the components of configuration x^i not belonging to block d_j . Two block-configurations can be united into a larger block-configuration by means of the \vee operator so defined:

$$x(d_j) \vee y(d_h) = z(d_j \cup d_h) \quad (7.2)$$

where $z_{\{v\}} = \begin{cases} x_v & \text{if } v \in d_j \\ y_v & \text{otherwise} \end{cases}$

We can therefore write $x^i = x^i(d_j) \vee x^i(d_{-j})$ for any i . Moreover, we define the size of a decomposition scheme as the size of its largest defining block: $k = 0$.

Coordination among blocks in a decomposition scheme may either take place through market-like mechanisms or via other organizational arrangements (e. g. hierarchies). Dynamically, when a new configuration appears, it is tested against the existing one according to its relative performance. The two configurations are compared in terms of their ranks and the superior one is selected, while the other one is discarded³³. More precisely, let us assume that the current configuration x^i and take block d_h with its current block-configuration $x^i(d_h)$. Let us now consider a new configuration $x^j(d_h)$ for the same block, if:

$$x^j(d_h) \vee x^i(d_{-h}) \geq x^i(d_h) \vee x^i(d_{-h}) \quad (7.3)$$

then $x^j(d_h)$ is selected and the new configuration $x^j(d_h) \vee x^i(d_{-h})$ is kept in place of x^i , otherwise $x^j(d_h)$ is discarded and x^i is kept.

In terms of a given division of labor structure (the decomposition scheme) within firms, individual workers or organizational sub-units specialize in various segments of the production process (a single block). Dy-

³³ As a first approximation, we assume that this sorting and selection mechanism is errorless and operates at no cost and without any friction.

namically, different inter-organizational decompositions entail different degrees of decentralization of the search process. The finer the inter-organizational decompositions, the smaller the portion of the search space which is being explored by local variational mechanisms and tested by market selection. Thus, there is inevitably a trade-off: on one hand, finer decompositions and more decentralization make search and adaptation faster (if the decomposition is the finest, search time is linear in N); on the other hand, they explore smaller and smaller portions of the search space, thus decreasing the likelihood that optimal (or even good) solutions are ever generated and tested.

A decomposition scheme is a sort of template which determines how new configurations are generated and can be tested afterwards by the selection mechanism. In large search spaces in which only a very small subset of all possible configurations can be generated and undergo testing, the procedure employed to generate such new configurations plays a key role in defining the set of attainable final configurations. We will assume that boundedly rational agents can only search locally in directions which are given by the decomposition scheme: new configurations are generated and tested in the neighborhood of the given one, where neighbors are new configurations obtained by changing only some (possibly all) components within a given block.

Given a decomposition scheme $D = \{d_1, d_2, \dots, d_k\}$, we say that a configuration x^i is a preferred neighbor or simply a neighbor of configuration x^j with respect to a block $d_h \in D$ if the following three conditions hold:

1. $x^i \geq x^j$,
2. $x_v^i = x_v^j \forall v \notin d_h$,
3. $x^i \neq x^j$.

Conditions 2 and 3 require that the two configurations differ only by components which belong to block d_h . According to the definition, a neighbor can be reached from a given configuration through the operation of a single decentralized coordination mechanism.

We call $H(x, d_h)$ the set of neighbors of a configuration x for block d_h . The set of best neighbors $B(x, d_h) \subseteq H(x, d_h)$ of a configuration x for block d_h is the set of the most preferred configurations in the set of neighbors:

$$B(x, d_h) = \{y \in H(x, d_h) \text{ such that } y \geq z, \forall z \in H(x, d_h)\} \quad (7.4)$$

By extension from single blocks to entire decomposition schemes, we can give the following definition of the set of neighbors for a decomposition scheme:

$$H(x, D) = \bigcup_{h=1}^k H(x, d_h) \tag{7.5}$$

A configuration is a local optimum for the decomposition scheme D if there does not exist a configuration y such that $y \in H(x, D)$ and $y > x$. A search path or, for short, a path $P(x^i, D)$ from a configuration x^i and for a decomposition D is a sequence, starting from x^i , of neighbors:

$$P(x^i, D) = x^i, x^{i+1}, x^{i+2}, \dots \text{ with } x^{i+m+1} \in H(x^{i+m}, D) \tag{7.6}$$

A configuration x^j is reachable from another configuration x^i and for decomposition D if there exists a path $P(x^i, D)$ such that $x^j \in P(x^i, D)$. Suppose configuration x^j is a local optimum for decomposition D ; we call the basin of attraction of x^j for decomposition D the set of all configurations from which x^j is reachable:

$$\psi(x^j, D) = \{y, \text{ such that } \exists P(y, D) \text{ with } x^j \in P(y, D)\} \tag{7.7}$$

Now let x^0 be the global optimum and let $Z \subseteq X$ with $x^0 \in Z$. We say that the problem (X, \geq) is locally decomposable in Z by scheme D if $Z \subseteq \psi(x^0, D)$. If $Z = X$, we say that the problem is globally decomposable by scheme D ³⁴.

We can soften the perfect decomposability requirement into one of near-decomposability. We no longer require the problem to be decomposed into completely separated sub-problems, i. e. sub-problems which fully contain all interdependencies, but we content with finding sub-problems which contain the most “relevant” interdependencies, while less relevant ones can persist across sub-problems. In this way, optimizing each sub-problem independently will not necessarily lead to the global optimum, but to a “good” solution³⁵. In other words, we construct near-decompositions which give a precise measure of the trade-off between decentralization and opti-

³⁴ A special case of decomposability, which is generalized here, is presented in Page (1996), and is called dominance. In our terminology, a block configuration $x^j(d_h)$ is dominant when $x^j(d_h) \vee x^i(d_{-h}) \geq x^i$ for every configuration $x^i \in X$.

³⁵ This procedure allows us to also deal with the case of multiple global optima and thus we can now also drop the assumption of a unique global optimum.

mality. Higher degrees of decentralization, while generally displaying a higher adaptation speed, are likely to be obtained at the expense of the asymptotic optimality of the solutions which can be reached.

Let us rearrange all the configurations in X by descending rank $X = \{x^0, x^1, \dots, x^{2^N-1}\}$ where $x^i \geq x^{i+1}$, and let $X_\mu = \{x^0, x^1, \dots, x^{\mu-1}\}$ with $0 \leq \mu \leq 2^N - 1$ be the ordered set of the best μ configurations. We say that X_μ is reachable from a configuration $y \notin X_\mu$ and for decomposition D if there exists a configuration $x^i \in X_\mu$ such that $x^i \in P(y, D)$. We call the basin of attraction $\psi(X_\mu, D)$ of X_μ for decomposition D the set of all configurations from which X_μ is reachable. If $\psi(X_\mu, D) = X$, we say that D is a μ -decomposition for the problem. The μ -decompositions of minimum size can be found algorithmically with a straightforward generalization of the above algorithm which computes minimum size decomposition schemes for optimal decompositions.

It is easy to show (Marengo-Dosi, 2005) that as μ increases, we can generally find finer near-decompositions. As Figure 7.3 illustrates, for randomly generated problems³⁶, if second-best solutions are accepted, we can have considerable reductions in the decomposition schemes size and in the expected search time. This shows that the organizational structure sets a balance in the trade-off between adaptation speed and optimality. It is easy to argue that in complex problem environments, characterized by strong and diffused interdependencies, such a trade-off will tend to produce organizational structures which are more decomposed and decentralized than what would be optimal given the interdependencies of the problem space. This property is shown in Figures 7.3 and 7.4, which present the typical search paths on a non-decomposable problem of two search processes driven, respectively, by decompositions:

$$\begin{aligned}
 D1 &= \{1, 2, \dots, 12\} \\
 D12 &= \{\{1\}, \{2\}, \dots, \{12\}\}
 \end{aligned}
 \tag{7.8}$$

³⁶ In this figure and the following (with the exception of Figure 7.1) we indicate on the vertical axis the rank of configurations re-parametrized between 0 (worst) and 1 (best).




Fig. 7.3. Near decomposition




Fig. 7.4. Fitness evolution 1

Figure 7.4 shows the first 180 iterations in which the more decentralized structure (D12) quickly climbs the problem space and outperforms the search based on a coarser decomposition. If there were a tight selection environment, a more than optimally decentralized organizational structure would quickly displace structure D1, which reflects the “true” decomposition of the underlying problem space.

However, the search process based on the finest decomposition quickly reaches a local optimum from where no further improvement can occur, while the process based on the coarser decomposition keeps searching and climbing slowly. Figure 7.5 shows iterations between 3000 and 3800, where the finest decomposition is still locked-into the local optimum it reached after very few iterations, while the coarsest one slowly reaches the global optimum (normalized to 1). Strong selective pressure therefore tends to favor organizational structures whose degree of decentralization is higher than what would be optimal from a mere problem-solving perspective.

This result is even stronger in problems that we could define as “modular”, i. e. characterized by blocks with strong interdependencies within blocks and much weaker interdependencies between blocks. In these




Fig. 7.5. Fitness Evolution 2

problems, higher levels of decompositions can be achieved at lower costs in terms of sub-optimality.

Another important property concerns micro (“idiosyncratic”) path-dependencies of organizational forms and their long-term persistence. If finer-than-optimal decompositions tend to emerge and to spread because of their “transient” evolutionary advantages, then one will generally observe also long-term diversity in the population of organizations in terms of (i) the decomposition they are based upon; (ii) the problem solutions they implement; and (iii) the local peaks they settle into³⁷. This is easily shown by a simulation in which we generate 100 organizations characterized by a randomly generated decomposition and a random initial string and let them search a randomly generated indecomposable problem. The 10 worst performing organizations are selected out every 10 iterations and replaced by 10 new organizations where 5 are randomly generated and 5 have the same decomposition scheme of the best performing ones but are placed on a randomly chosen configuration.

Figure 7.6 plots the number of diverse organizational forms at every iteration. Initially, diversity does indeed sharply decrease because of selective pressure but then it stabilizes on numbers consistently and persistently higher than 1. A very similar trend describes the number of different surviving configurations, which reflects the fact that the population of organizations settles onto several local peaks of similar value.

In another set of simulations where, at given intervals, the current problem has been changed – at regular intervals – with another one having exactly the same structure in terms of decomposability, but with different – randomly generated- orderings of configurations. This can be taken as a metaphorical proxy for environment volatility. For instance, consumers might have changing preferences over a given set of characteristics, or producers face changing relative input prices. Interestingly enough, it turns out that, even with totally decomposable problems, as the change ranks becomes more frequent, the population is entirely invaded by organizations characterized by coarser and coarser decompositions and at the limit, by organizations which do not decompose at all. This robustly suggests that growing volatility has stronger consequences than does growing interdependence. The reason why this happens is shown in Figures 7.7 and 7.8 which present, respectively, the expected improvements and the probability of improvement for searches based on the finest (D12) and coarsest

³⁷ On this latter point a similar result is obtained by Levinthal, 1997.




Fig. 7.6. Evolution of organization forms

(D1) decomposition schemes in a fully decomposable problem³⁸. It is shown that, when starting from low rank configurations, a search based upon coarser decomposition has a higher probability of finding a better configuration and, when such a better configuration is found, its expected rank is higher for coarser decompositions. This is due to the fact that finer decompositions search only locally; and this, on average, cannot produce large improvements in fully-decomposable problems. When the problem space is highly volatile – though always fully decomposable – sooner or later every organization will fall into an area of very “bad” configurations from which coarser decompositions have a higher chance of promptly re-covering.

³⁸ Figures 7.7 and 7.8 refer to the fully decomposable search space given by the binary numbers between 0 and $2^N - 1$. But the same qualitative results are obtained for any kind of fully decomposable search space.




Fig. 7.7. Expected improvements




Fig. 7.8. Probability of improvement

7.4 Theses and conjectures

Theories of firms and organizations have always played a central role in evolutionary economics, because firms in turn play a central role in the processes of technological and economic change. Evolutionary theorizing has focused upon firms as loci of generation of productive knowledge and has been producing some – mainly agent-based – models in which the organizational structure is analyzed for its knowledge generation and distribution properties. Such properties have been generally overlooked by neo-classical theory, which has on the contrary concentrated almost entirely on the firm as a system of contractual monetary incentives designed in such a way as to align individual actions with the organization's goals.

Evolutionary theory takes an opposite point of view, emphasizing the cognitive aspect of organizations and leaving in the background issues of incentives, governance and power. This hypothesis seems often as extreme as the neoclassical one. According to the latter business firms are places in which strictly self-centered selfish individuals spend their life devising ways to cheat on and profit from each-other. The role of the organization is therefore to implement complex contractual arrangements which limit the damages of self-interested behavior.

A large part of evolutionary theory on the contrary, as shown in the models briefly outlined in this chapter, implicitly assumes away any incentive issue: organizations are teams in which agents basically cooperate but have heterogeneous knowledge. Thus evolutionary theory does not have much to say at time being on such extremely important topic as incentives, motivation, governance, power and, especially, on how these interact with the primary activity of knowledge generation and sharing.

A major task for future evolutionary theorizing and modeling will be integrating such issues and providing its own original view. Nelson and Winter's notion of routine as a truce among conflicting goals and claims in the organizations can represent a useful starting point³⁹. However the choice between the two extreme of self-interested rational behavior on one hand and routine behavior on the other can also be seriously reconsidered in the light of the recent research – mainly experimental – which has been heavily questioning that monetary incentives be the only motive of eco-

³⁹ For some preliminary speculations on such integration see Coriat-Dosi, 1998; for a tentative model which integrates conflict within a model similar to the one presented in the previous paragraph see Dosi-Levinthal-Marengo, 2003.

conomic agency. Experimental evidence shows a variety of motives across different subjects and different contexts and clearly shows that organizations can indeed shape individual motives, as they shape individual and collective knowledge, and not simply adapt to them (Fehr-Fischbacher, 2002).

References

- Aoki, M. (1986): Horizontal vs. Vertical Information Structure of the Firm, *American Economic Review*, 76, 971-983.
- Baldwin, C.Y. and Clark, K.B. (2000): *Design Rules: The power of modularity*, MIT Press.
- Cohen, M. and Burkhart, R. (1996): Routines and other recurring action patterns of organizations: Contemporary research issues, *Industrial and Corporate Change*, 5, 653-698.
- Christensen, M. and Knudsen, Th. (2002): *The Architecture of Economic Organization: Towards a general framework*, University of southern Denmark.
- Coriat, B. and Dosi, G. (1998): Learning how to govern and learning how to solve problems: on the co-evolution of competences, conflicts and organizational routines, in: A. Chandler, P.P. Hagström, and O. Sölvell (eds.), *The dynamic firm*, Oxford University Press.
- Cyert, R.M. and March, J.G. (1963): *A behavioral theory of the firm*, Prentice Hall.
- Dosi, G., Levinthal, D., and Marengo, L. (2003): Bridging contested terrain: linking incentive-based and learning perspectives on organizational evolution, *Industrial and Corporate change*, 12, 413-436.
- Fehr, E. and Fischbacher, U. (2002): Why social preferences matter: the impact of non-selfish motives on competition, cooperation and incentives, *Economic Journal*, 112, C1-C33.
- Frenken, K., Marengo, L., and Valente, M. (1999): *Interdependencies, nearly-decomposability and adaptation*, Kluwer Academic.
- Holland, J.H. (1975): *Adaptation in natural and artificial systems*, University of Michigan Press.
- Holland, J.H., Holyoak, K.J., Nisbett, R.E., and Thagard, P.R. (1986): *Induction: processes of inference, learning and discovery*, MIT Press.
- Ioannides, Y.M. (2003): Complexity and organizational architecture, Dept. of economics, mimeo, Tufts University.
- Jensen, J.B. and McGuckin, R.H. (1997): Firm performance and evolution: Empirical regularities in the US microdata, *Industrial and corporate change*, vol. 6, pp 25-47.
- Kauffman, S.A. (1993): *The origins of order*, Oxford University Press.

- Kosa, J.R. (1993): *Genetic Programming*, MIT Press.
- Levinthal, D. (1997): Adaptation on rugged landscapes, *Management Science*, 43, 934-950.
- March, J.G. and Simon H.A. (1958): *Organizations*, Wiley.
- March, J.G. (1994): *A primer on decision making*, Free Press.
- Marengo, L. (1992): Coordination and organizational learning in the firm, *Journal of evolutionary economics*, 2, 313-326.
- Marengo, L. (1996): Structure, competence and learning in an adaptive model of the firm, in: G. Dosi and F. Malerba (eds.), *Organization and Strategy in the evolution of the enterprise*, MacMillan, 124-154.
- Marengo, L., Dosi, G., Legrenzi, P., and Pasquali, C. (2000): The structure of problem-solving knowledge and the structure of organizations, *Industrial and corporate change*, 9, 757-788.
- Marengo, L. and Dosi, G. (2005): Decentralization and market mechanisms in collective problem-solving, *Journal of economic behavior and organization*, 58, 303-326.
- Milgrom, P. and Roberts, J. (1990): The economics of modern manufacturing, *American Economic Review*, 80, 511-528.
- Milgrom, P. and Roberts, J. (1992): *Economics, organization and management*, Prentice Hall.
- Miller, J.H. (2001): *Evolving information processing organizations, Dynamics of organizations: computational modeling and organization theories*, MIT Press, 207-327.
- Mount, K.R. and Reiter, S. (2002): *Computation and complexity in economic behavior and organization*, Cambridge University Press.
- Nelson, R.R. and Winter, S.G. (1982): *An evolutionary theory of economic change*, Harvard University Press.
- Page, S.E. (1996): Two measures of difficulty, *Economic Theory*, 8, 321-346.
- Radner, R. (1992): Hierarchy: the economics of managing, *Journal of Economic Literature*, 30, 1382-1415.
- Radner, R. (2000): Costly and bounded rationality in individual and team decision-making, *Industrial and corporate change*, 9, 623-658.
- Sah, R.K. and Stiglitz, J.E. (1986): The architecture of economic systems: hierarchies vs. polyarchies, *American Economic Review*, 76, 716-727.
- Schumpeter, J.A. (1955): *The theory of economic development*, Harvard University Press.
- Simon, H.A. (1976): *Administrative behavior*, The Free Press, 3rd edition.
- Simon, H.A. (1981): *The science of the artificial*, MIT Press, 2nd edition.
- Simon, H.A. (1983): *Reason in human affairs*, Stanford University Press.

- Williamson, O.E. (1975): *Markets and hierarchies: analysis and antitrust implications*, Free Press.
- Winter, S.G. (1968): Towards a neo-schumpeterian theory of the firm, mimeo, RAND Corporation,
- Winter, S.G. (1982): An essay on the theory of production, in: H. Hymans (ed.), *Economics and the world around it*, University of Michigan Press, 55-93.

8 Emergence of institutions

Concerning the two main entities forming an economic system, one may say that “at the beginning are the agents, at the beginning are the goods”. These entities can be described in physical terms, respectively as structured collections of individuals and as coherent bundles of material (or immaterial) items. But they are rather defined functionally, as centres of decision in the first case, as objects of exchanges in the second case. Even if they may come in or disappear for economic reasons, they are at first sight considered as primitive entities which exist spontaneously.

A third type of entity is frequently introduced, namely institutions, even this hotchpotch concept gathers indistinctly what is not reducible to agents and goods. They are seldom described by their very nature since they exist in a physical as well as psychical form, or most often in a hybrid one. They are essentially defined by their role of mediation between agents, especially in order to help them to face exogenous or endogenous uncertainty. They may be considered as already there, but their emergence is frequently sketched as the result of a combination of agents’ actions.

The present chapter analyzes the functional roles and the evolution modes encountered for economic and non economic institutions. In fact, it emphasises a specific type of institution, called a “convention”, which acts as a behavior rule able to coordinate the agents (part 1). It examines how different types of conventions may emerge and evolve in order to resolve some game failures (part 2). It considers more precise models of emergence, by a learning or an evolutionary process, of conventions concerned with money, wages or resources (part 3). It suggests some directions for a further study of more complex institutions (part 4).

8.1 Background and problems

8.1.1 Economic theory and institutions

Economic theory introduces various institutions in order to deal with production, exchange and consumption of goods. The central institution is the “market” which confronts the offers and demands of some good and

renders them compatible through a price. Substitutive institutions follow the same aim of regulating the exchanges, but along more or less different procedures: asymmetric auction mechanisms, planning procedures, bilateral negotiations. Complementary institutions act as support of the market by framing the exchanges: exchange rights fix legal conditions for exchange, money fluidifies the implementation of exchanges, trust makes possible the bilateral exchange between goods and money when non synchronized. Further institutions frame the global transaction system: technical conventions codify the goods and technologies, moral norms determine the exchangeable goods, property rights fix the owner of the goods.

Several taxonomies of institutions have been proposed in the literature, however introducing fuzzy categories. Structurally, an “organic institution” appears as a human organization while an “epistemic institution” is a mental representation endorsed by all agents. For instance, the state, a central bank, a firm or a syndicate are of the first kind while money, trust, a shared belief or an allocation norm are of the second kind. A “constitutive institution” makes possible a new social activity while a “regulative institution” acts as a mediator in an already existing activity. For instance, a new type of money, an artificial financial market or an original technical language are of the first kind while patents, business contracts or traffic rules are of the second kind. In the following, only epistemic and regulative institutions, which act directly on agents’ beliefs and preferences, will be considered. However, firms are studied in chapter 7 and syndicates are treated through a model in section 3.

In economic theory, different schools of thought stressed the basic role of institutions around and outside the market. The “institutionnalists” (Veblen, 1899; Commons, 1934; Mitchell, 1949) remarked that institutions are habits or customs which influence the agents’ behaviors and facilitate their expectation. The “cognitivists” (Hayek, 1973) considered that institutions (like markets) produce signals (like prices) which are a good summary for the information needed by agents. The “transaction costs” theory (Williamson, 1975) explained the choice of various coordination modes – market, hierarchy – by the reduction of the transaction costs involved. The “property rights” theory (Coase, 1937, Alchian-Demsetz, 1973) stressed the importance of well-defined property rights when evaluating the efficiency of an economic system. The “regulationists” (Boyer, 2004) explained the heterogeneity of institutions by their adaptation to different places and periods.

8.1.2 Conventions as institutions

Among all types of institutions, the literature privileges the “conventions” which can be expressed by “behavior rules” of the agents. A behavior rule is stated in the form “if context C, then action A” and associates a given action to some “signals” associated to the context. A signal is an exogenous characteristic attached to the agent himself or to his material environment. Hence, a convention appears as some regularity of agent’s behavior which relates exclusively on observable variables. For the modeler, it can be interpreted positively as some behavior which is spontaneously followed by the agents in given circumstances. But it may also be interpreted normatively as a behavior that is imposed to the agents under a social pressure. In the last case, a convention has to be interpreted by the agent himself in order to decide if it applies or not to a given situation.

Conventions are generally studied in the framework of non cooperative game theory, which has many advantages. First, game theory works at a level of generality which is higher than economic theory since players and actions are undifferentiated. Second, game theory respects methodological individualism since an equilibrium state is only generated by players’ actions. Third, game theory considers that players’ interactions are not mediated by previous institutions, at least explicit ones. Such a view was already defended in the seminal book “Theory of games and economic behavior” (von Neumann-Morgenstern, 1944) where the equilibrium states are interpreted as “standards of behavior”. It was developed in a whole trend where an institution is seen as a component of an equilibrium state chosen by the players (Schotter, 1981; Kreps, 1990).

The “evolutionist” view of conventions is based on the idea of “spontaneous order” issued from interactions of boundedly rational players. First, it means that the order is not a voluntary design, in the sense that it is not intended by a regulating authority. Second, it means that the order is not even conscious, in the sense that it is not expected or recognized by the players. Third, it means that the order needs not to be enforced by some legal authority in order to hold. Spontaneous order was initially stressed by Hayek (1973) who considers the market as a coordinating entity which results from some evolutionary process. It was essentially developed by Sugden (1989) who considers that “individuals living together in a state of anarchy tend to evolve conventions or codes of conduct that reduce the extent of interpersonal conflict”.

8.1.3 Types of conventions

In classical game theory, three main “game failures” are recorded, analogous to the “market failures” in economic theory. Each failure is associated with a class of games and is illustrated by some 2x2 game concerned with car driving (see chapter 3). First, the “coselection problem” stems from the existence of several equilibrium states among which one has to be sorted out. It appears in games like the crossroads game (where one equilibrium is better for one player and conversely), the technology game (where one equilibrium Pareto-dominated is better for both players), or the intermediary driving side game (where the equilibria are equivalent for the players). Second, the “cooperation problem” stems from the existence of an equilibrium state which is not Pareto-optimal. It appears in games like the car lights game. Third, the “codetermination problem” stems from the absence of any equilibrium state. It appears in games like the Indian file game.

The basic games always consider two drivers who, respectively, may pass or not at some crossroads, drive right or left, buy an electric or petrol car, use full or dipped lights, speed up or slow down when driving in line, with the following payoffs:

1/2	stop	pass
stop	(2,2)	(2,3)
pass	(3,2)	(0,0)

Crossroads game

1/2	Left	right
left	(1,1)	(0,0)
right	(0,0)	(1,1)

Driving side game

1/2	electric	petrol
electric	(3,3)	(0,1)
petrol	(1,0)	(1,1)

Technology game

1/2	full	dipped
full	(1,1)	(3,0)
dipped	(0,3)	(2,2)

Car lights game

1/2	speed	slow
speed up	(1,3)	(3,1)
slow down	(2,0)	(0,2)

Indian file game

Fig. 8.1. Basic games

A specific type of convention may be associated with each type of game failure. As concerns the coordination problem, a “convention of coordination” makes precise how a player is directed towards a specific equilibrium state among several ones. For instance, in the crossroads game, a “priority convention” expresses that priority is given to players coming from the right. In the technology game, a “standardisation convention” asserts a technology to choose. In the driving side game, an “orientation convention” specifies what side of the road is privileged. In the hawk-dove game (see

chapter 3), similar to the crossroads game, a “property convention” states what player will get the resource. As concerns the “cooperation problem”, a “reciprocity convention” specifies how a player may be directed towards a Pareto-optimal state rather than an equilibrium state. For instance, in the repeated car lights game, a “threatening convention” suggests that the players play as their opponent in the last period.

8.1.4 Explanation of conventions

In a first stage, game theory assimilates a convention to an equilibrium state, in the framework of classical game theory. Outside the fact that an institution has a coordination aim like an equilibrium, an institution has many features of an equilibrium (Aoki, 2001). It is endogenous, self-enforcing, robust to exogenous changes and often multiple. Theoretically, for any institution, one may consider that a game failure is detected in some game and an institution is chosen in an auxiliary game in order to compensate for it. Practically, for a convention, a game failure is defined in some game and the convention is defined as an equilibrium state in the same game. However, classical game theory raises an “implementation problem” since it is just defined as a fixed point by the modeler. No constructive process leading to it is exhibited, except if considering a fictitious entity, the Nash regulator.

In a second stage, game theory assimilates the genesis of an institution to a process leading to an equilibrium state. Two types of processes were proposed in order to remedy the “implementation problem” (Binmore, 1992). The “eductive process” considers that the players are able to get on an equilibrium state by their sole reasoning. The players are perfectly rational and simulate their opponent’s behavior under the assumption of common knowledge of the game structure and of the players’ rationality. Such an approach was applied to a convention by Lewis (1969), who restricts to pure coordination games. The “evolutionary process” considers that the players come to an equilibrium state by convergence of a dynamic process. Such a process can be a learning process, either belief-based learning or reinforcement learning (see chapter 3). Such an approach is precisely applied in the “evolutionist view” of conventions.

When adopting such a view, a third stage is however needed in order to separate the institution from the underlying equilibrium and give it a normative foundation. The “naturalization problem” is solved in two steps while the “normatization problem” is solved in two further steps. The “recognition step” concerns the observation by the players that some regu-

larity appeared, regularity that anybody can observe. The “polarization step” concerns the treatment of the observed regularity as an institution, in the sense that the players react now directly to it rather than to their opponents. The “legitimation step” transforms the institution, even if it is already self-enforcing, into a moral obligation sustained by reciprocal expectations of players. The “legalization step” imposes to the institution a lawlike force sustained by incentives and sanctions. These four steps can be easily described, but are not readily formalized in the game theory framework.

8.2 Canonical principles

8.2.1 Evolutionary game theory

The genesis of a convention is studied in the standard framework of evolutionary game theory (see chapter 3). In a large population, players interact infinitely many times, but only through their actions without direct communication. They interact on some network, each player meeting more or less randomly the opponents situated in some neighborhood. Each player follows a given strategy, and is moreover characterized by some exogenous “labels” related to its position or role in the game. He gets a payoff at each period depending on all players’ actions (through a payoff matrix) and this payoff is observed by him instantaneously. Moreover, each player observes precisely or on average the others’ actions as well as the others’ labels, but the interactions are other way anonymous. For instance, if the label takes value A or B, each player interacts only with players with opposed labels and adapts his strategy to the label observed.

According to the precise specification of his behavior, each player follows various learning processes. In a belief-based learning process, each player implements an action which is a best response to the observed frequency of others’ past actions. In a reinforcement learning process, each player implements an action according to the performance obtained in the past with that action. The last process is formally similar to an evolutionary process where each player reproduces according to the past performance obtained in the past with his fixed strategy. In any case, the learning processes of the players define conjointly a dynamical stochastic process. Such a process can be summarized in a phase diagram showing in what directions the system evolves from a given position. It can be noticed that nothing implies that the mean utility increases in such a process, evolution and utility being disjointed.

The process may adopt a chaotic behavior, cycle indefinitely or converge towards some asymptotic state. The last is moreover submitted to

stability conditions which are more or less strong (asymptotic stability, dynamic stability). When players interact in a systematic way, the asymptotic state may be an equilibrium state of the basic game which is repeated, such as a Nash equilibrium. Especially, the process may converge towards an “evolutionary stable equilibrium”, which is just a refinement of a Nash equilibrium. In 2x2 games, two situations are considered. When the players are not labelled, an equilibrium state is a value p where each player plays the first action with probability p . When the players are labelled, an equilibrium state is a couple (p, q) where the first player (labelled A) takes the first action with probability p and the second (labelled B) takes the first action with probability q . When the players are situated on a network, equilibria are more complex since a given state is achieved in some areas and another in other areas.

8.2.2 Emergence of a convention

The process leading to the emergence of a convention will be illustrated by the crossroads game (Sugden, 1986, 1989). Consider first that the drivers are “myopic” and aware of no label able to differentiate them. The dynamic process converges towards the unique stable state $(2/3, 2/3)$ where each driver stops two times on three. The consequences are rather bad since collision happen one time on nine and lead to an average utility of 2. In order that a priority convention arises, one driver at least has to give sense to a distinction between two positions or roles (A or B) of the drivers. The label may be associated with a big car against a small car, a car coming from the right against a car coming from the left, a fast car against a slow car or a big road against a small road. Detecting or even creating such an asymmetry between the drivers imposes in fact some conceptual jump for the drivers.

Consider now two subpopulations of players, the “smart” drivers who recognize some label and the “myopic” drivers who do not. The smart drivers follow some convention “pass if A, stop if B” while the myopic drivers pass with probability r . A stable equilibrium state exists for some specific value r^* of r . The equilibrium value r^* increases with the proportion x of smart drivers and becomes equal to one over some threshold $x^* = 2/3$. It can easily be shown that, at the equilibrium state, the myopic players have always a smaller utility than the smart ones. The myopic drivers may then become conscious that the smart ones use a convention allowing them to do better. Moreover, the smart drivers have themselves an interest in the myopic ones following the convention. Hence, when a direct com-

munication is possible, they make them aware of the convention and try to persuade them to follow it.

For a given label, two competing conventions are often available associating an action with one occurrence of it. The choice of one specific form of label may be favoured by some structural asymmetry, either informational or preferential. The first case is illustrated by the fact that a driver on a main road is less conscious of a crossroads than a driver on a secondary road. When a driver on the main road becomes aware of the crossroads, he observes that the driver on the secondary one stops more often than passing, hence tends to pass. Such a process is self-enforcing when more drivers become conscious and leads to a priority for the main road. The second case is illustrated by the assumption that if two drivers speed when coming to a crossroads, a big car is less hurt than a small car. The driver with the big car has a little incentive to pass and the self-enforcing process leads again to a priority for big cars. In both cases, the vital stage of the emergence of a convention is the initial one.

Finally, several independent labels may be considered in some situation and lead different groups to follow competing conventions. In general, the biggest group obtains the highest payoff, inducing the others to change their convention in order to better perform. Here again, it is the label which is the first to become conscious which has the highest chance to be selected. In fact, the labels are suggested either by asymmetries associated with the specific situation or by analogies with similar situations. Hence, the asymmetries or analogies which strike most people most quickly have more chance to be implemented. Especially, some conventions are more fertile than others since they can be applied to many different situations. For instance, priority to the north-south road over the east-west road is harder to apply than the priority to the right.

8.2.3 Change of convention

The process leading to the change of some convention will be illustrated by the technology game (Boyer-Orléan, 1992). If the drivers follow some learning or evolutionary process, they converge to a situation where all use the same type of cars, either petrol or electric ones. Considered as a convention, the final state depends on the initial conditions, the dynamic rules and random factors. Without the last, the final state is defined by some threshold p^* characterizing the proportion of players using initially petrol cars or electric cars. If all drivers use petrol cars, a change of convention is clearly desirable since the convention is Pareto-inefficient. But it is often hard to obtain since the stability conditions of a convention work against

its modification. In fact, such a change of convention stems from exogenous factors which can be categorized into three classes, formally examined in an evolutionary game framework.

First, the change of convention can be attributed to a population change driven by external factors. Such a change modifies the endogenous performance of each possible convention and may induce a shift. This is the case when a group following a given convention invades the group following the existing convention. For instance, the proportion of electric cars may be increased suddenly by some fashion effect inspired by outside. If it exceeds the threshold p^* , all petrol cars will be progressively converted into electric ones. This is even more the case when a population following some convention more or less collapses. For instance, many petrol cars can be destroyed during a war and allow a new social experimentation. A new initial situation is created where electric cars have a new opportunity to emerge and supplant the petrol cars.

Second, the change of convention can be attributed to the existence of an exogenous network between players. The relevant performance of a convention becomes then a local one rather than a global one. This is the case when an interaction network is designed where the players meet only opponents in some neighborhood. For instance, if a driver is surrounded in a large neighborhood by essentially petrol cars, he has no incentive to shift to a petrol car. But if the neighborhood is smaller, he may face enough electric cars and shift, and the better technology locally spreads out and may even progressively dominate the overall network. This is also the case when an information network is designed where the players are informed only about neighboring opponents. For instance, if due to random samples, he observes that his neighbors use mainly electric cars, he will adopt it and the technology will again develop locally and even globally.

Third, the change of convention can be attributed to a change in the structural properties of the conventions. The intrinsic performance of existing and new conventions is then directly evolving through time. This is the case when some new opportunities appear, especially opportunities in-between existing ones. For instance, mixed cars may appear which work with petrol as well as batteries. The petrol cars may then be converted into mixed cars, then the mixed cars into electric cars. That is even more the case when the intrinsic utility of each convention changes due to technological improvement. For instance, the electric cars are progressively produced at lower cost for exogenous reasons or since they are more produced. The threshold p^* between electric cars and petrol cars becomes lower and favours the acquisition of electric cars, especially when the population is itself changing.

8.2.4 Enforcement of a convention

A convention may progressively appear as the unconscious output of some learning or evolutionary process. However, its emergence is facilitated by the fact that it is recognized during the process by some of the players. Moreover, its emergence is enforced by the fact that the convention is recognized afterwards by all players. Since it mimics an equilibrium state, the convention is naturally self-enforcing: each player has an interest to follow the convention if the others do so. However, the conditions of an equilibrium state may not be precisely put together, inducing some agents to deviate from the convention. Some players may not have a complete view of the situation and act locally in an opportunist way. Some players may not really understand the coordinating role played by the convention and act against it.

Hence, the convention may be further enforced and become a norm, sustained by various social pressures. A “legitimated convention” (Orléan, 2004) is more precisely sustained by the fact that the players are induced to conform to it. Such a conformity relies on the dissatisfaction of not following a norm considered as fair. It also relies on the pressure exerted by the disapproval from others to not conform to the norm. It can be defined in purely psychological terms or be translated into material or legal incentives and sanctions. The conformity effect can be treated as a social interest which guides the players, a player following a convention because he thinks he has socially to do so. The social utility can be integrated as a new argument in the player’s utility function besides its intrinsic personal utility. In order to reinforce a convention, the social utility has just to work in the same direction than the individual utility.

More generally, the social interest may be relevant in situations where a convention is not self-sustaining. This is especially the case when there exists a Pareto-optimal state which is not an equilibrium and is taken as a convention of cooperation. The social interest acts then against the personal interest and is even predominant in order to impose the convention. For instance, in the car lights game, using dipped lights may become a convention even if it not an equilibrium state in the static game (but it may be in the repeated one, as the notion of convention suggests). The drivers are incited to use their dipped lights since their is is a heavy social disapproval of using full lights. The same is true for the “participation convention” in the financement of a public good (Sugden, 1986), for the “effort convention” in the interactions between employers and employees (Leibenstein, 1982) or for the “equity convention” in the fixing of wages (Akerlof, 1980).

8.3 Some models

Three models are presented. The first deals with the emergence of money, which has a material support, but acts as an epistemic institution and even a convention. The second deals with the emergence of a distributive norm which is directly an epistemic institution. The third deals with the emergence of a syndicate, which is essentially an organic institution

8.3.1 Emergence of money

Sethi (1999) provides a model for the emergence of money. It is not aiming at representing the historical process by which money or even some occurrence of money was created, but to suggest an evolutionary mechanism able to get such a result.

The model considers three populations of agents and three goods. An agent of type i consumes only good i and produces only good $(i + 1)$ modulo 3 (cf. figure 8.1). Goods are indivisible, costly to store, and at each period, an agent can store only one good. When an agent succeeds in exchanging the stored good against his consumption good, he gets some utility and produces immediately his production good. In such a structure, the bilateral exchanges cannot stem from “the double coincidence of needs” since the agent i wants to consume the production of agent $(i + 2)$, but this agent does not want to consume the production of agent i .

In order to simultaneously achieve their consumption needs, the agents have to exchange goods according to figure 8.2.

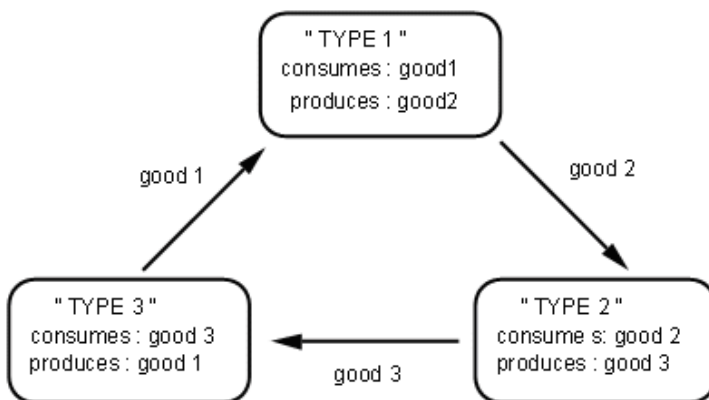


Fig. 8.2. Exchange of goods

Agents i have two exchange strategies at disposal: only accept his consumption good i (and storing then only production good $(i + 2)$) or accept good $(i + 2)$, i. e. consider it as a mean of exchange or “money” (and storing either good $(i + 1)$ or good $(i + 2)$). These strategies are respectively followed by agents $i\alpha$ and $i\beta$.

In each period, denote:

- $s = (s_1, s_2, s_3)$ the vector where s_i is the proportion of agents i accepting i as money
- $p = (p_1, p_2, p_3)$ the vector where p_i is the proportion of agents i storing their production good $(i + 1)$.

Hence, among agents i , the proportion of agents $i\alpha$ is $(1-s_i)$, the proportion of agents $i\beta$ storing good $(i + 2)$ is $(1 - p_i)$ and the proportion of agents $i\beta$ storing good $(i + 1)$ is equal to $s_i - (1 - p_i)$ or $(s_i + p_i - 1)$.

At each period, two agents meet randomly and exchange if and only if they take a benefit, inducing a modification of the stock distribution p . Assuming that the populations of agents are initially equally distributed, the meeting probabilities follow as well as the probabilities of stocks.

In the short run, the vector s of strategies of the agents is considered as fixed. The distribution of stocks p evolves with the exchanges, according to a probabilistic law which has the following deterministic approximation:

$$dp_i / dt = (1 - p_i)p_i + 1 - (s_i + p_i - 1)p_i + 1 \tag{8.1}$$

In fact, the proportion of agents i storing their production good $(i + 1)$ varies after two types of meetings, profitable for one agent at least:

- an agent $i\beta$ sells good $(i + 2)$ against good i , then transforms it immediately in good $(i + 1)$. This happens when meeting an agent $(i + 2)$ storing his production good i (hence the term p_{i+2}), while agent i stored good $(i + 2)$ (hence the term $1 - p_i$),
- an agent $i\beta$ sells good $(i + 1)$ against good $(i + 2)$. This happens when meeting an agent $(i + 1)$ storing good $(i + 2)$ (hence the term p_{i+1}), while agent i stored good $(i + 1)$ (hence the term $s_i + p_{i-1}$).

Note that an agent $i\alpha$ who obtains the good i consumes it immediately in order to produce good $(i + 1)$, leaving the stock unchanged. Likely, note that an agent never buys his own production good.

The dynamics of stocks has the following asymptotic states:

- if $s \neq (1,1,1)$, there is a unique stable equilibrium state,

- if $s = (1,1,1)$, there are two equilibrium states, only one being stable

In any case, it exists a *temporary equilibrium* called *monomorphic* if all agents of a same population use a same strategy ($s_i=0$ ou 1) and *polymorphic* otherwise.

In the long run, one assumes that the dynamics of stocks is fast with regard to the dynamics of strategies. Denote $p_1(s), p_2(s), p_3(s)$ the distribution of stocks for the three goods at the temporary equilibrium when the vector s is fixed.

When the storing cost of good i is c_i , one can compute the expected revenues $\Pi_{i\alpha}(s)$ and $\Pi_{i\beta}(s)$ of the six categories of agents at the temporary equilibrium. Consider for instance agents $i\alpha$ storing always good $(i+1)$ with cost c_{i+1} . In order for such an agent to get one unity of his consumption good i , he needs to meet either an agent $(i+2)$ who accepts good $(i+1)$, what happens with probability $(p_{i+2}(s) + s_{i+2} - 1)/3$, or an agent $(i+1)$ storing good i , what happens with probability $(1 - p_{i+1}(s))/3$. By adding both probabilities and normalizing the consumption utility, one obtains:

$$\Pi_{i\alpha}(s) = (p_{i+2}(s) + s_{i+2} - p_{i+1}(s))/3 - c_{i+1} \quad (8.2)$$

Likely, the expected revenue for an agent $i\beta$ is:

$$\Pi_{i\beta}(s) = a_i(s)(p_{i+2}(s) + s_{i+2} - p_{i+1}(s))/3 - a_i(s)c_{i+1} - (1 - a_i(s))c_{i+2} \quad (8.3)$$

with $a_i(s) = p_{i+2}(s)/(p_{i+1}(s) + p_{i+2}(s))$.

The average revenue of population i is:

$$\Pi_i = (1 - s_i)\Pi_{i\alpha} + s_i\Pi_{i\beta} \quad (8.4)$$

The situation is that of a game with three players i and two strategies (α ou β) for each player. Along the principles described in chapter 3, the agents are assumed to modify their strategy according to the difference between their observed revenue with their strategy and the average revenue of their population. More precisely, the dynamics of strategies is a “replicator dynamics” given by three differential equations:

$$\dot{s}_i/s_i = \Pi_{i\beta} - \Pi_i \quad (8.5)$$

According to the storing costs, two *environnements* are possible:

- environnement A with $c_1 < c_2 < c_3$
- environnement B with $c_1 < c_3 < c_2$

For *environnement A*, the following results hold:

- if $3(c_3 - c_2) > 0,5$, it exists a stable asymptotic state for $s = (0,1,0)$
- if $\sqrt{2} - 1 < 3(c_3 - c_2) < 0,5$, it exists a stable asymptotic state for $s = (x, 1, 0)$ with $x \in [0,1]$
- if $3(c_3 - c_2) < \sqrt{2} - 1$, it exists a stable asymptotic state for $s = (1,1,0)$

These results can be interpreted by considering exchanges (i,j) where good i is sold against good j :

- if $(c_3 - c_2)$ is high, agents 1 proceed to exchanges $(2,1)$, agents 2 to exchanges $(1,2)$ and $(3,1)$, agents 3 to exchanges $(1,3)$. Good 1 is considered as money since it is part of all exchanges:




Fig. 8.3. First structure of exchanges

- if $(c_3 - c_2)$ is low, agents 1 proceed to exchanges $(2,1)$, $(2,3)$ and $(3,1)$, agents 2 to exchanges $(3,2)$, $(3,1)$ and $(1,2)$, agents 3 to exchanges $(1,3)$. Either good 1 or 3 can be accepted conventionally as money.




Fig. 8.4. Second structure of exchanges

In an intermediate situation, only a fraction of agents 1 accept the good 3. Only the good 1 can be considered as money.

For *environment B*, the following results hold:

- it exists a stable asymptotic state if $s = (1, 1, 0)$.
- if $3(c_3 - c_1) < 1 - 0,5\sqrt{2}$ and $3(c_2 - c_1) < \sqrt{2} - 1$, it exists a second stable asymptotic state for $s = (1, 0, 1)$.

The first stable state is the same than for environment A. Goods 1 and 3 can act as money. The second stable state is of same type as the preceding one, but with a permutation of goods. Goods 2 and 3 can be taken for money.

In environment B, a bifurcation leads from one to two stable states according to the value of costs.

Sethi makes a distinction between two categories of equilibria, the fundamental equilibrium and the speculative equilibrium. In the first, it is the good with lowest storage cost which is chosen as money (i.e. good 1). Agents exchange either to obtain their consumption good or to diminish their storing cost. In the second, an agent at least accepts some good which is neither his consumption good, nor a good with a lower storage cost. His article shows that it is not necessary to have drastic information and rationality conditions in order to explain the emergence of money. It is enough to assume that it exists some prior degree of specialization in production and consumption, which is progressively increased as money is used as a mean of exchange.

8.3.2 Emergence of a distributive norm

Young (1993b, 1996) presents a model of a sharecropping rule relative to some land. A sharecropping rule specifies the fraction of the harvest attributed to the landowner and to the laborer. Such a model is supposed to be compatible with stylized facts, especially observed in Indian villages. In a given village, the rule is the same for all (local conformity effect). In different villages, the rules may notably differ (global diversity effect). At some times, a sharp change of rule happens in some village (punctuated equilibrium effect).

The model considers as the one-shot bargaining game the Nash demand game, where the landlord and the tenant make demands about the part of the crop they want. If the demands are compatible, each player obtains what he asked for. If the demands are incompatible, they get nothing. This

game is discretized, each agent having only three possible demands: high H (3/4), medium M (1/2) and low L (1/4). Assume moreover that the players are risk-neutral, hence that the utility is proportional to the share of the crop he gets.

The payoff matrix is then the following (the landlord being the first player):

Table 8.1. Share cropping game

	H	M	L
H	(0,0)	(0,0)	(75,25)
M	(0,0)	(50,50)	(50,25)
L	(25,75)	(25,50)	(25,25)

The game has three pure Nash equilibria: (1/4,3/4), (1/2,1/2) and (3/4, 1/4). Hence, three conventions can be associated to them.

The model considers an evolutionary process called “adaptive play” (Young, 1993a), which is just a generalization of the fictitious play process (see chapter 3). It considers that the players are only partially informed, express boundedly rational responses and are also making random deviations (interpreted as involuntary mistakes or voluntary experiments).

More precisely, it is based on four assumptions applied at each period:

- there is a population of landlords and of laborers and one landlord is randomly matched to a laborer, each population being large and finite
- each agent observes only the m preceding bargains (where the demands were compatible or not), forgetting then all past miscoordinations
- each agent considers only a sample of past demands of size s and computes the frequency in the past of the other’s possible demands
- each agent forms a random expectation of the other’s future demand and chooses a best reply with probability $1-e$ and a random demand with probability e .

The model considers the stochastic process induced by the agents. In the middle term, the system may stay for a long time in some specific state, then shift rapidly to an other state. However, at long term, it converges as-

ymptotically towards a “stochastically stable” equilibrium. Such a state is defined in order to be robust to small, random and persistent shocks.

If the amount of information is sufficiently large and sufficiently incomplete (s/m is under some threshold) and if the experiment process is sufficiently low (e under some threshold), the selected convention is $(1/2, 1/2)$, hence an equal share of the crop. No consideration of fairness grounded on prior moral norms is necessary. No idea of exogenous “focal point” grounded on cultural features is needed. A specific share is just selected because it is stable in the long run.

8.3.3 Emergence of a trade union

Consider the elementary market, described in chapter 2, where workers and firms providing jobs are following a learning process. If information and search costs are nil, the process converges asymptotically towards a stable equilibrium with a unique price p . On such a market, it is possible to describe the emergence of a trade union or syndicate, which allows his members to get a better remuneration than p .

The assumed dynamics of the market is the following:

- it exists, during the market life, a worker who plays the role of a germ and thinks that by forming a coalition, some workers would do better. In each period t , a contract is proposed by the germ to each worker. According to this contract, the worker is committed to work only for the firms which are signing the contract:
- each worker makes his decision to accept or refuse the contract according to his past experience and present information. The set M_t of workers having accepted the contract in the period forms (if not empty) the present syndicate.
- considering the number of its committed members, the syndicate announces to each firm a wage requirement $\delta_t \geq p$. It is assumed that in first period, $\delta_1 > p$. Moreover, the syndicate is assumed to disappear if it is constrained to bring back its price requirement under the threshold p .
- each firm accepts or refuses the proposition of the syndicate according to his past experience and present information. She knows that, when rejecting the proposition of the syndicate, all members will refuse to work for her. In that way, a set (eventually empty) N_t of contracting firms is formed. The contracting workers and firms form an *organized market*:

- the workers and the firms who have not contracted (if existing) get on a *free market* and compare their demands. A unique wage s_t rapidly forms on this market and can be observed by anybody.
- at the end of the period, the workers and the firms, according to their last experience, adapt their proposition to join the syndicate or to contract with it. As for the syndicate, it considers the last required wage and its consequences in order to fix the required wage for the next period.

The behavior of workers, firms and syndicate can be made more precise, with regard to both markets.

In each period t , the worker k is characterized by his *degree of attraction* $a_t(k)$ towards the syndicate and adapts his action in consequence:

- if $a_t(k) \leq 0$, the worker does not join the syndicate
- if $a_t(k) > 0$ and
 - $a_{t-1}(k) \leq 0$, the worker joins the syndicate with a given probability
 - $a_{t-1}(k) > 0$, the worker joins the syndicate with probability 1.

In other terms, the degree of attraction has to be positive in the last two periods in order for the worker to join surely the syndicate.

The probability Π_t that the member of a syndicate gets a job on the organized market equals the rate between the number of jobs offered by the firm having contracted with the syndicate and the number of members of the syndicate. With probability Π_t , the worker gets his syndical wage and with probability $1 - \Pi_t$, he gets his minimal wage \underline{w}_k (assumed to be equal to its unemployment allocation). Hence, his expected average wage at period t is: $h_t(k) = \Pi_t \delta_t + (1 - \Pi_t) \underline{w}_k$

The worker compares this expected revenue to the wage s_t on the free market. He increases his degree of attraction for the syndicate if $h_t(k) > s_t$ and reduces it if not.

In each period t , the firm i is characterized by a *negotiation threshold* $b_t(i)$ towards the syndicate and adapts her action in consequence:

- if $b_t(i) < \delta_t$, the firms refuses the proposition of the syndicate
- if $b_t(i) \geq \delta_t$, the firm agrees with the offered contract

Of course, the negotiation threshold $b_t(i)$ cannot exceed the present maximal wage. The firms adapts her threshold for the next period according to his present experience:

- if the firm refused the offer the syndicate, she increases the threshold $b_t(i)$ if observing that the wage on the free market is higher than on the organized one: $s_t > \delta_t$, maintains the threshold if $p \leq s_t \leq \delta_t$ and reduces the threshold –hence hardens his position- when $s_t < p$.
- if the firm accepted the offer of the syndicate, she does not modify her threshold.

In each period t , the syndicate looks at increasing the required wage, but has to deal with the risk of seeing the number N_t of contracting firms decreasing:

- if $N_t > 0$, the syndicate increases his requirement by one unit if $s_t > \delta_t$, and maintains or increases its requirement with some probability if $s_t \leq \delta_t$
- if $N_t = 0$, the syndicate maintains his offer if $s_t > \delta_t$, and reduces his offer if $s_t \leq \delta_t$.

From the preceding assumptions, it is possible to infer the following results. From any initial configuration, it exists a date T such that, if the syndicate is not already dead, it lives for ever. In that case, the labour market has three characteristics:

- the required syndical wage is bounded by two fixed limits $\underline{\delta}$ and $\bar{\delta}$ depending on the market past evolution
- the number of contracting firms is bounded by two values \underline{n} et \bar{n} depending on the market past evolution
- the set of workers is split in three groups the composition of which depending on the past market evolution: a (non empty) group of workers joining the syndicate in each period, a (possibly empty) group of workers belonging sporadically to the syndicate, a (possibly empty) group of workers never joining the syndicate, either because their requirements are too low or because their requirements are too high.

Hence, the market dynamics leads, when the syndicate does not die, to a syndicate which is formed of a “core” of convinced workers and a “corona” of unsatisfied workers who always oscillate between the organized and free markets. Such a dynamics is history-dependent since the syndicate may die or reach a variable size according to random shocks.

8.4 Theses and conjectures

As shown by this chapter, the emergence of conventions can be conveniently treated within the assumptions developed in the preceding chapters. Players are boundedly rational, actions are sequential, information is progressively acquired, random factors are considered. More precisely, the usual learning and evolutionary processes can be applied to the emergence of all types of conventions, especially in the framework of evolutionary game theory (Walliser, 2005). This is the case for conventions of coordination such as driving rules or conventions of reciprocity such as resource splitting. It is even the case for more fundamental economic conventions such as money or wage rules.

Other institutions than conventions can readily be analyzed along similar schemes, especially using reinforcement mechanisms. First, it is the case for social beliefs (expectations, theories, opinions) which act as cognitive institutions. It is possible to consider that each agent has at each period some degree of adherence to the belief. This degree is progressively modified by the confirmation he obtains for the belief. Second, it is the case for coalitions of players (syndicates, lobbies) which are organizational institutions. It is again possible to consider that each agent has some degree of adherence to a group. This degree evolves with the performance each agent achieves thanks to its group.

More generally, the problem is not restricted to the emergence of a given institution in some given environment. First, an institution appears, lasts and dies in some cycle where the decline is generally faster than the outbreak. This is due to the fact that the reasons for which an institution holds can be different from the reasons why it emerges. An institution is a device which may keep its structure unchanged while changing its functions through time. Second, there exists a whole system of institutions which are sustaining one another in a coherent way. Hence, the appearance or disappearance of some specific institution has some driving consequences on all related institutions.

As economic institutions are concerned, the genesis of an institution depends on the action of the state, but also on the influences of the agents themselves. Any economic agent tries to achieve his objectives by acting simultaneously at two structural and temporal levels. At short term, he considers the many institutions as fixed and tries to do his best in a more or less strategic environment. For instance, a firm wants to reduce its costs, but also to increase his market shares with regards to its competitors. At long term, he intends to modify the economic rules of the game in order to

construct a protected environment. For instance, a firm tries to modify the competition or distribution rules in order to become a local monopoly.

References

- Akerlof, G. (1980): A theory of social custom, *Quarterly Journal of Economics*, 94, 749-75.
- Alchian, A. and Demsetz, H. (1973): The property rights paradigm, *Journal of Economic History*.
- Aoki, M. (2001): *Toward a comparative institutional analysis*, MIT Press
- Binmore, K. (1992): *Fun and games*, Heath.
- Boyer, R. (2004): *La théorie de la régulation*, Repères, La Découverte.
- Boyer, R., and Orléan, A. (1992): How do conventions evolve?, *Journal of Evolutionary Economics*, 2, 165-177.
- Coase, R.H. (1937): The nature of the firm“, *Economica*.
- Commons, J.R. (1934): *Institutional Economics*, University of Wisconsin Press.
- Demsetz, H. (1967): Toward a theory of property rights, *American Economic Review*.
- von Hayek, F. (1973): *Law, legislation and liberty*, University of Chicago Press.
- Kreps, D. (1990): Corporate culture and economic theory, in J.E. Alt - K.A. Shepsle (eds.), *Perspectives on positive political economics*, Cambridge University Press, 90-143.
- Leibenstein, H. (1982): The prisoner's dilemma in the invisible hand: an analysis of intra-firm productivity, *American Economic Review*, 2, 92-97.
- Lewis, D. (1969): *Convention, a philosophical study*, Harvard University Press.
- Mitchell, W.C. (1935): Commons on institutional economics, *American Economic Review*.
- von Neuman, J. and Morgenstern, O. (1944): *Theory of games and economic behavior*, Princeton University Press.
- Orléan, A. (2004): L'économie des conventions: définitions et résultats, in A. Orléan (ed.), *Analyse Economique des conventions*, Presses Universitaires de France, 9-48.
- Schotter, A. (1981): *The economic theory of social institutions*, Cambridge University Press.
- Sethi, R. (1999): Evolutionary stability and media of exchange, *Journal of Economic Behaviour and Organization*.
- Sugden, R. (1986): *The Evolution of Rights, Cooperation and Welfare*, Basil Blackwell.

- Sugden, R. (1989): Spontaneous order, *Journal of Economic Perspectives*, 3, 85-97.
- Ullman-Margalit, E. (1977): *The emergence of norms*, Oxford University Press.
- Veblen, T. (1899): *The theory of the leisure class*, Mc Millan.
- Walliser, B. (2005): Game theory and emergence of institutions, mimeo.
- Williamson, O. (1975): *Markets and hierarchies*, Free Press.
- Young, H.P. (1993a): The evolution of conventions, *Econometrica*, 61(1), 57-84.
- Young, H.P. (1993b): An Evolutionary model of bargaining, *Journal of Economic Theory*, 59, 145-68.
- Young, H.P. (1996): The economics of convention, *Journal of Economic Perspectives*, 10(2), 105-122.

9 State and economic system regulation

From an evolutionary point of view, what role can or should the State play in the economic system? This chapter demonstrates that the State's role extends far beyond the simple task of correcting failures in the functioning of the market – the role to which public intervention is limited in standard microeconomic theory. However, the conditions for the success of this intervention are complex and delicate, and the more realistic framework of the economic system proposed by the evolutionary approach, while providing wider scope for intervention, restricts the degree to which the State can control the evolution of this system. When the neoclassical paradigm – and with it the criterion of Pareto-optimality – are called into question, it becomes much more complicated to define the function that the State should fulfil. The objectives of public intervention no longer appear so clear-cut. The formulation of a set of prescriptive criteria to guide this intervention therefore becomes indispensable.

In this chapter, the State's role is limited to interventions that are more of a microeconomic order (environmental policy, technology policy, management of the labour market, etc.). Policies of a more macroeconomic nature (budgetary and monetary policies, for example) have been deliberately excluded from the analysis. The evolutionary approach undeniably sheds new light on the objectives and conditions for the success of the latter, but this book is devoted to the presentation of an evolutionary microeconomic theory.

9.1 Background and problems

In neoclassical microeconomic theory, the legitimacy of public intervention in the economic system derives from several different sources. Generally, for economists, and more particularly for neo-institutionalists (see chapter 8), it is the State's duty to define the institutions and, on a wider scale, the framework – through the introduction and enforcement of laws – in which individuals and firms carry out their economic activities. The State therefore establishes the rules of the game and ensures their application. Here, its objective is to maximise an intertemporal collective surplus,

possibly under the constraint of equity satisfied by means of tax and redistributive tools. However, this chapter is not concerned with such a wide view of public intervention⁴⁰, but with a more specific justification, namely the correction of failures in the market. Such failures correspond to situations in which the free action of market forces does not spontaneously lead to a Pareto optimum, i. e. a situation in which the position of one individual cannot be improved without deterioration in the situation of another. The possibility of decentralised functioning of the economic system is highlighted by the two theorems of welfare economics, which stipulate, firstly, that competitive equilibrium always leads to a Pareto optimum situation and secondly, that a given optimum can always be attained through the competitive market provided that the appropriate modifications have been made to the initial agents' endowments. Consequently, State intervention can only be justified in situations where the market leads to inefficient or inequitable states. Traditionally, *three situations leading to market failures have been identified: (I) the existence of a natural monopoly, (II) the existence of public goods and external effects and (III) the imperfection of information.* We should emphasise from the start that an evolutionary approach in no way calls this point of view into question.

A situation of *natural monopoly* refers to the existence of growing returns in the industry concerned. Considering the size of the market, potential economies of scale justify the existence of a monopoly rather than the presence of many different firms. In this case, the public authorities are encouraged to regulate entry into this market, but also to control the activity of the monopoly so that it does not take undue advantage of its dominant position. Since the 1990s, neoclassical theory has given increasing importance to the role of competition policy, far beyond the simple regulation of natural monopolies, faced with firms using a whole battery of anti-competitive weapons (collusion, barriers to entry, predation), as the new industrial economics has illustrated only too well. In this field, as we shall see later, the evolutionary approach makes it possible to bring the effectiveness of classic antitrust measures into proportion.

The existence of *public goods* also justifies public intervention. A public good possesses two characteristics: non-rivalry or shared consumption (consumption by one individual does not reduce the consumption possibili-

⁴⁰ On the question of the evolutionary analysis of the formation and evolution of laws, and more generally of the political process, see Slembeck (2003) and Pelikan (2003).

ties of other individuals) and non-exclusion (once the good or service has been produced, it is impossible to reserve its consumption exclusively to those who pay for it). Any good or service presenting these two characteristics is produced in sub-optimal quantities by the market. Likewise, the existence of *positive externalities* (positive effects from which certain agents benefit without payment of any consideration) during production or consumption of a good results in a sub-optimal level of production in a market system, whereas *negative externalities* (negative effects sustained by certain agents who receive no compensation in return) generate over-production. In these cases, the State may carry out an internalisation of these external effects (through taxes or subsidies, for example).

Finally, the market can also fail in a situation of *imperfect information*, notably when this is distributed asymmetrically. It is well known (Akerlof, 1970) that situations of adverse selection may then exist (due to an asymmetry of information about the characteristics of an agent). For example, the incapacity of suppliers to signal the quality of the goods supplied to consumers leads the latter only to accept a price corresponding to an “average” quality, and suppliers of “high quality” prefer to withdraw from the market. In this case, the State may decide to impose quality controls.

Literature on market failures has been enriched since the middle of the 1980s (Greenwald and Stiglitz, 1986; Stiglitz, 1989; Arnott et al., 1994). Thus, state intervention can become legitimate when markets are incomplete, in other words when a good or service that is the object of solvent demand is not supplied, a situation which, notably, prevents agents from protecting themselves against uncertainty because of the absence of contingent markets. Likewise, situations of moral hazard may occur (due to an asymmetry of information about the actions of an agent), making it impossible, for example, to be certain that a contractor subsequently respects the conditions stipulated beforehand in a contract. Williamson (1975, 1985) highlighted the problems that the existence of such opportunist behaviour can cause in the functioning of a market with transaction costs. To minimise these costs, the functioning of the market and the behaviour of agents can be supervised by institutions (private or public) endowed with incentive and/or coercive powers.

In these situations of market failures, what makes the State the best-placed body for regulating the market? The answer lies most notably in the fact that the government possesses many more means of applying incentives to agents than any private firm. The State enjoys four main privileges: the power to impose and collect taxes, the power to prohibit certain behaviours on the part of individuals or firms, the power to punish viola-

tions of the law and the possibility of reducing transaction costs. For some, even this particularly restricted domain of competence is questionable. The principle of Coase thus establishes that, in the presence of externalities, for example, private parties can carry out transfers of rights which internalise these externalities in a way that is mutually beneficial to all parties. However, as Coase (1992) himself emphasised, from the moment that transaction costs are non-zero, the State can recover an undeniable advantage over the market when it comes to correcting certain failures in the market.

However, to legitimise public intervention it must be demonstrated that the shortcomings of the market are of a higher order than those of the government. The State is also, for example, subject to the constraint of imperfect information. In the relationship between the regulator and a natural monopoly, there is clearly an asymmetry of information in favour of the latter, which has access to more information than the government about the structure of costs. Generally, the various sources of government shortcomings identified in literature on the subject are, in addition to the imperfect nature of the information available to the public policy-maker, the fact that the individuals who benefit from State intervention do not coincide with those who finance it, excessive bureaucracy, the activities of lobby groups and the myopia of deciders who are essentially preoccupied with their own re-election. These are so many reasons that may lead governments to implement inappropriate policies.

This brief overview of the neoclassical arguments in favour of public intervention would remain incomplete if we limited ourselves solely to the literature on market failures. From a more macroeconomic point of view, two categories of developments grant a wider role to public intervention than the simple correction of these failures: strategic trading policy⁴¹ and the theory of endogenous growth⁴². Although the first current finally concludes that it is difficult to define clear instructions for the public policy-maker, and the scope for public action is restricted in the

⁴¹ Following the work of Brander and Spencer and of Krugman (1988), works on strategic trading policy proliferated during the 1980s. Under the hypothesis of an international duopoly, these works show that the imposition of import taxes or export subsidies by the government can lead to a Pareto improvement.

⁴² According to the theory of endogenous growth (Romer, 1986), the State can play a role ensuring the existence of an institutional environment favourable to training and innovation (of which the collective returns are greater than the private returns). Public policy is presented in this theory as a fundamental determinant of the growth rate of an economy.

second current⁴³, the works on strategic trading policy and the theory of endogenous growth have two interesting characteristics in common. (i) They introduce the problem of efficiency in a dynamic context. They explicitly consider time as an essential phenomenon and relativise the a-historicity of neoclassical theory. (ii) They attribute a more active role to the State. Public intervention is viewed as a factor that stimulates national growth or the competitiveness of national firms. These almost heterodox views on the field of public intervention in a neoclassical context are in fact at the heart of the evolutionary approach.

9.2 Canonical principles

From an evolutionary point of view, the redefinition of the role of the State is founded on a paradigm, the main characteristics of which are as follows. Above all, evolutionary theory shifts the emphasis away from the idea of equilibrium and towards the idea of a process that may or may not tend towards stable states. The economic universe is bounded neither by time nor by the opportunities that are presented to agents. So, while pursuing their own objectives, both individuals and the State are continually modifying their plans in the light of new information. In this process of trial and error, transitory states assume as much importance as the asymptotic states to which they may eventually lead.

Uncertainty, in Knight's sense of the word, is an important dimension in evolutionary analysis. We do not systematically consider a risky or uncertain world in which agents know the possible states of the world and attribute objective or subjective probabilities to them. The State, just like the other agents, is faced with radical uncertainty insofar as that most of the future states of the world cannot even be conceived. The economic agent is constantly having to adapt to changes in his environment, changes which, most of the time, he has not anticipated. Thus, evolutionary theory is characterised not only by the existence of unpredictable changes and by the behavioural adjustments necessary to adapt to or exploit these changes, but also by the fact that these adjustments and adaptations themselves – whether private or public – do not generally lead to predictable results. For better or for worse, economic life is an adventure (Nelson and Winter,

⁴³ The public policy-maker's field of intervention is largely concentrated on policies of training for individuals and the encouragement and orientation of firms' research and development activities.

1982). So, even if the government possesses an incomparably greater capacity to model the economic system than that of private agents, fundamentally it remains subject to the same informational and behavioural constraints.

The concept of irreversibility is also present in evolutionary theory. Past choices exercise a constraint on all future choices; the economic process resembles a trajectory with branches in which it is often impossible to retrace one's steps or to switch from one branch to another. In particular, once the system is directed towards a given attractor, its progression becomes ineluctable and it is then too late to redirect the system towards another attractor. Thus, as we shall demonstrate in this chapter, public interventions in the economic sphere are subject to phenomena of irreversibility and their effects depend on the specific path followed.

Bounded rationality is another essential component of the evolutionary approach. Faced with a fluctuating and uncertain environment and possessing only limited reasoning capacities, agents can do no more than implement adaptive procedures. In this respect, the situation of the State is no different, despite its greater possibilities for the gathering and processing of information. On the other hand, when the situation makes it impossible to refer to a Pareto optimum, the State must define the criteria used for the evaluation of its policies. The adoption of universal objectives, valid for every problem, appears to be beyond the reach of the State. It must content itself with aims that are more specific. Likewise, if it can no longer content itself with the simple correction of market failures, its field of intervention must be specified on a case-by-case basis.

If the evolutionary public policy-maker cannot adopt optimising behaviour, can he nevertheless be considered benevolent, in other words concerned with intertemporal social welfare and the respect of individual preferences? The neoclassical position casts doubt on this hypothesis: the real objective of public authorities can at times diverge significantly from the maximisation of an intertemporal collective surplus (due to bureaucracy, lobbying, etc.). Nevertheless, for the rest of this chapter we shall retain the strong hypothesis of the benevolence of the public authorities. We shall therefore assume (i) that the State's objective is intertemporal social welfare (although we shall also explore the question of the measurement of this welfare); and (ii) that the State is free of any internal conflict between its various components (ministries, public agencies, local government, etc.). There is in fact no need to consider the existence of conflicts of interest within the public sphere or of diverging objectives in the search for social welfare to enrich the analysis of public intervention in the economic system.

We conclude this section on the canonical principles governing public intervention in an evolutionary context by observing that the evolutionary State does not limit itself to correcting market failures: it can also facilitate the market process (accelerating convergence towards a unique equilibrium), guide the market process (directing the economic system onto what is judged to be a satisfactory path) and play a creative role by favouring the emergence of conditions enabling agents to reach situations which market forces alone would not have rendered possible. As for the objective of public intervention, this is considerably more multiform in the evolutionary context than in the neoclassical approach. Subsequently, however, the State is always assumed to adopt a “Pareto philosophy” in its interventions, even if it is not capable of fully realising this philosophy.

9.3 Some models

In this section, three tasks of the State will be illustrated by models (Moreau, 2004): facilitating the functioning of market forces; guiding the economic system onto a satisfactory path (firstly under the hypothesis of State omniscience, then under the hypothesis of a State with imperfect information); and enabling private agents to satisfy needs not covered by the market, by acting as a creator of novelties.

9.3.1 The State as a catalyst of market forces

To illustrate this first role of the State, very few of the traditional hypotheses need be removed. We simply consider that agents, now endowed with imperfect information and bounded rationality, find it difficult to generate an equilibrium, even if it is unique. The State consequently has a role to play as catalyst in the functioning of the market.

9.3.1.1 An illustrative model

Let us take the labour market model presented in chapter 2. In this market, economic agents are assumed to display very unsophisticated adaptive behaviour. Each employee k has a minimum wage $w(k)$, below which he refuses to work. Each firm i , which employs one person at the most, offers a maximum wage $\bar{v}(i)$, above which it prefers not to employ anyone. The functions $\underline{w}(\cdot)$ and $\bar{v}(\cdot)$ thus represent the habitual supply and demand functions of this labour market. However, there is no Walrasian auctioneer

to coordinate the different requirements of these firms and workers. In each period, a pair (worker, firm) is drawn at random and the result of this meeting can be positive (an employment contract is signed) or negative. In both cases, the result modifies the requirements of the workers and the firms. For example, a worker already in a job will only accept another position if the salary is higher. If, on the other hand, an unemployed worker demands a salary so high that the firm refuses to employ him, he will reduce his demands. Likewise, a firm that does not succeed in recruiting a worker for a vacant position will raise its salary offer. But under no circumstances can the salary demanded by a worker k fall below $\underline{w}(k)$ or the salary offered by a firm i exceed $\bar{v}(i)$. After a process of exploration, the duration of which is random, the system constituted by these interacting workers and firms converges towards a stable state that is characterised by the fact that no further contract can be signed between a worker and a firm. All unemployed individuals have reduced the salaries they demand to their minimum threshold levels and every firm with vacant positions has raised its salary offer up to the maximum threshold. This stable state is also characterised by the fact that all the salaries are equal (price uniqueness) and supply equals demand. The total surplus, which can be defined as the sum total of the differences for each job between the maximum salary that the firm was prepared to offer and the minimum salary that the worker was prepared to accept, is now maximal. Apart from a few exceptions, this was not the case during the transitional periods. There is therefore clearly a loss of surplus during periods of exploration of the system. Shortening the duration of this period of exploration can therefore undeniably be considered efficient.

9.3.1.2 A wider vision of the State's role of coordination

The only difference between this model and the neoclassical market is that convergence towards equilibrium is not immediate, because there is no Walrasian auctioneer. Convergence is the fruit of a process of random meetings between workers and firms and the resulting, crude adjustments in their behaviour. What role can the State play in such a configuration? If we assume an omniscient State that can gather all the information from the economic agents without cost, the State can directly impose the equilibrium price. In this case, it is fulfilling the function of the auctioneer. More realistically, the State can content itself with favouring coordination between workers and firms, in other words facilitating the operation of market forces. Here we find Hayek's idea that the main problem faced by the

economic system is that of coordination between individuals and the only function the State should fulfil is that of furthering this coordination. Within the framework of the model presented here, this consists in enabling every firm and individual to meet as many different potential partners as possible in each period. This increases the probability of seeing each position rapidly occupied and accelerates the process of adjusting their requirements for both individuals and firms. *The role of the State may therefore be to create institutions capable of enabling workers and firms to multiply multilateral contacts in each period.* This is, of course, the role played by public employment agencies intended to improve the circulation of information between workers and employers.

Wakeley (1998) developed a similar argument to justify certain measures of technology policy. He considered the case of competition between several firms in a given technological environment, where this environment is unknown to the agents beforehand. Some of the technological choices made by firms are nevertheless clearly superior to others and lead to a maximum surplus. Wakeley demonstrated that, even if market forces do systematically lead to the most appropriate market structure for the best exploitation of the technological environment, the process of convergence could take a long time. This waste of time can result in significant losses in the welfare of the economic agents. *In this case, State intervention can consist in encouraging firms (by means of subsidies, for example) to explore the technological environment to the maximum with the aim of accelerating the discovery of good technological options.*

In conclusion, when we really take into account the dynamic aspect of economic processes and the generally adaptive behavior of economic agents, a new role appears for the State, even if the system is destined to converge towards a unique equilibrium. This role is that of market catalyst. *The State then contributes to the speeding-up of a process of convergence that would eventually lead to a situation of equilibrium in any case, but which generates losses of surplus along the way, during the periods of transition.*

9.3.2 The State as an omniscient guide of the market process

Public intervention gains considerably in depth when we consider the case in which several punctual attractors exist. The public policy-maker knows that the economic system possesses several punctual attractors. However, we maintain a strong hypothesis: the omniscient State seeks to maximise a function of social welfare. Thus, among these different attractors, the State

can identify without ambiguity the one that will maximise the welfare of the society. This provisional hypothesis of State omniscience is highly unrealistic, but it has the virtue of making it easier to compare the lessons to be drawn from the neoclassical and evolutionary approaches concerning the economic role of the State.

This problem quite naturally evokes the works of Arthur (1988, 1989) and David (1985) on the processes of technological competition. In these studies, two goods or two technologies are generally in competition and the probability for each of them of being chosen in the future depends on its past rate of adoption, due to the existence of positive externalities (network externalities, pecuniary externalities, etc.). These models display the distinctive characteristics of evolutionary systems: irreversibility, path dependence, the crucial nature of random events. However, models of technological competition often assume questionable hypotheses. In the models of Arthur, for example, technologies appear simultaneously in a virgin market. It would be more realistic to adopt the hypothesis of a new technology entering into competition with an older, already established technology. We shall therefore illustrate our argument with the help of the model developed by Laffond et al. (2000), which is more convincing in this respect. The model proposed by Malerba et al. (2001), analysing competition policy in an evolutionary context, will also be briefly presented.

9.3.2.1 *A first illustrative model*

We assume there are two products, imperfectly substitutable, produced using two different technologies by two firms indexed k ($k = 1, 2$). Thanks to the use of a clean technology, firm 1 is able to propose an environmentally friendly product. As for firm 2, it sells a product manufactured using a pollutant technology. The technology used by firm 1 is still emergent and entails a higher unit production cost than the technology used by firm 2. Therefore, at the start of the competitive process, which is spread over T periods, the unit cost of product 1 is higher. However, if the rhythm of growth in the market share of product 1 is fast enough in the future, its unit cost will fall significantly (economies of scale and of learning). The unit cost of product 2, on the contrary, as it issues from a mature technology, can only decrease slowly and at a constant rhythm when its market share increases.

Let $c_1(t)$ denote the unit cost of product 1 compared to that of product 2 for each period t and let $P_1(t)$ denote the market share of firm 1 in the same period. For each date t , there is a critical market share part for firm 1, $R_1(t)$, such that:

- if $P_1(t) \geq R_1(t)$ then $c_1(t+1) \leq c_1(t)$
- if $P_1(t) < R_1(t)$ then $c_1(t+1) > c_1(t)$

In other words, when enough consumers switch to product 1, its unit cost will decrease faster than that of the competing product. In no circumstances, however, can the relative unit cost of product 1 fall below a minimum threshold c_1^* , assumed to be lower than 1, which is considered to be a technological limit. When the number of consumers switching to product 1 is insufficient, the unit cost of product 1 in relation to product 2 rises, hindering its capacity to increase market share in the following period.

Let us now examine the characteristics of demand in this market. During the period t , buyers come into the market and each of them buys one unit of one of the two products. In relation to the product k , a buyer i has a utility of the form $U_k^i = \omega_k^i / c_k$ where ω_k^i represents a parameter of preference that is independent for each buyer and distributed according to a uniform law on $[0,1]$. The utility obtained by an individual i from consumption of a product k is therefore proportional to a factor of preference that is specific to him, and inversely proportional to the unit production cost (and therefore the price⁴⁴) of the product. For each period, a number N of buyers is drawn at random; every buyer has the same probability of being drawn and each one chooses to consume the product that procures him more utility. The market share of product 1 on date t is therefore a random variable, because in each period the agents who come into the market can have a higher or lower intrinsic preference for this product (depending, for example, on their ecological sensibilities).

The dynamics of the process of competition between the two firms 1 and 2 depends as much on the initial conditions (in other words the initial relative unit cost and the size of the phenomena of economies of scale and learning) as it does on the history of the process itself (through random events on demand). Generally, the process of competition can be represented with the help of figure 9.1. Here, we assume that the number of buyers in each period is high enough for the random character of demand to be disregarded. On any date t , if the unit production cost of product 1

⁴⁴ To simplify matters, it is assumed that the utility of consumers is a function of the unit cost of the products and not their selling price. In fact, this simply means considering that firms determine their prices by increasing their unit production cost by a proportional margin that is constant and identical for all competitors.

compared with product 2 is greater than a , then this relative cost can only increase from one period to the next (we are in the zone where $P_1(t) < R_1(t)$) and firm 1 will see its market share tend towards zero. When $c_1(t) < a$, on the contrary, $P_1(t) > R_1(t)$, the relative cost of product 1 will therefore decrease towards c_1^* and its market share will tend towards one.




Fig. 9.1. The competitive process

This process contains two punctual attractors, each corresponding to a monopoly situation for one of the two products. Let us assume that a monopoly for product 1 is socially preferable to a monopoly for product 2. If, at the beginning of the competitive process, $c_1(0) < a$, then the functioning of this market alone, characterised by heterogeneous preferences and by positive pecuniary externalities (the fact that an individual buys a product in the period t enables the other consumers to buy it at a lower price in the period $t + 1$), is enough to lead the process towards the socially preferable attractor. In the opposite case, State intervention is required to prevent the system from becoming locked in an inferior attractor. Intervention by the public authorities can naturally take the form of a modification in the relative cost of the two products. For example, by imposing an appropriate rate of tax on product 2, a public policy-maker can enable firm 1 to establish its product on the market despite the initially unfavourable conditions. Note that if $c_1(0) \gg a$, the rate of tax would be so high as to represent a pure and simple prohibition of product 2. Figure 9.1 illustrates the impact of

such a tax, which results in an increase in the critical market share from a to b . On any date t , the imposition of a tax rate τ on product 2 makes it possible, if $a < c_1(t) < b$, to ensure the convergence of the system towards the socially preferable attractor.

9.3.2.2 *The strengths and weaknesses of an omniscient guiding State*

In traditional microeconomic theory, the case of negative externalities connected with pollution is often restricted to the question of the sub-optimality of the market equilibrium and the overproduction of the pollutant product. State intervention is consequently aimed at modifying the market mechanisms in order to reduce pollution to an optimum level. Analysis of public intervention is considerably enriched by the introduction of the more realistic hypothesis that the same services can be supplied to consumers with the help of goods using more or less pollutant processes, and that the economic system therefore possesses several punctual attractors. Firstly, as we have just emphasised, failures in the market are not ineluctable. On the other hand, when public intervention proves to be necessary to guide the system towards the desired attractor, two of its characteristics must be underscored: (i) this intervention may be only temporary; (ii) it may not always be effective.

Temporary public policy. The often-irreversible character of public interventions can generally be imputed to political motives. For example, taxes that are imposed for a specific purpose, or to raise money for a particular end, often outlive the original problem because of budgetary constraints. Yet the State can act simply as a momentary pilot of the system. A temporary intervention (stimulus) can suffice to guide the system towards the desired attractor, with accumulative endogenous phenomena (economies of scale and learning) then taking over from the exogenous intervention. Let us illustrate this proposition. Figure 9.1 demonstrates that if taxation remains necessary on any date t such that $a < c_1(t) < b$, it is no longer required after the date t' (with $t' > t$) where $c_1(t') < a$. From the period t' onwards, the relative production costs of the two competing products, in the absence of any public intervention, provide firm 1 with a large enough market share to generate dynamic increasing returns. The tax can then be abolished.

The narrow window dilemma. This major dilemma in public policy, highlighted by David (1987), results from the fact that the period during which the public policy-maker can act successfully on the system is often brief.

In such a dynamic system, any delay in the introduction of the tax can prove to be fatal to public policy. So, returning to our example in figure 9.1, if, on the date t , the policy-maker receives the information that $c_1(t) < b$, he can then decide to introduce a tax at the rate of τ on the price of product 2. But if we assume that the tax will only become effective from the date t' and that $c_1(t') > b$, then this tax will not succeed in stopping the erosion in the market share of product 1 and will not prevent its future disappearance from the market. Of course, the rate of tax can be increased, but the same phenomenon will reoccur. In other words, the later the intervention, the more prohibitive its cost will be. The government must therefore pay particular attention to the behaviour of the system in the proximity of transitory states. In our example, these states correspond to the periods during which it is still possible to redirect the course of the system (i. e. to ensure the long-term survival of product 1) by a modification in market mechanisms and not by the simple prohibition of one of the competing products (by prohibitive taxation of product 2).

The role of random events and of the heterogeneity of consumer preferences. The above analysis assumes that the number of consumers coming into the market in each period is high enough to neutralise the heterogeneity of their preferences. Removing this hypothesis renders public intervention more complicated. So, in the example of figure 9.1, for a given rate of tax on product 2, the dynamics of the system will be orientated in one direction or the other by the choices of the first cohorts of consumers to come into the market, depending on whether these consumers are sensitive to ecological issues. A low rate of tax can suffice to direct the system towards the desired attractor (a monopoly for firm 1) if the first consumers to come into the market are ecologically minded. On the contrary, if the first consumers have no ecological sensibilities, even a higher rate of tax will not dissuade them from buying product 2. In other words, one same rate of tax may just as well ensure the survival of product 1 as it may fail to prevent the accumulative increase in the market share of product 2, because of the unequal nature of the distribution of ecological sensibilities in consumers and the fact that these consumers enter the market in a random order, spread out over time.

9.3.2.3 A second illustrative model

The model of Malerba et al. (2001) provides another example of the problem posed for public authorities concerning the timing of their intervention. This model studies the computing sector, reproducing its principal

empirical characteristics: (i) very significant increasing returns both on the supply side (successful R&D turns into economic success and consequently into an increased R&D effort leading to yet further successes) and on the demand side (consumer loyalty); (ii) the sequential entry of new firms following the innovators. Here, the aim of the public authorities is to prevent over-concentration of the market and their weapon is anti-trust policy. This stipulates that any firm whose market share exceeds 75% can be split in two, either immediately or later. As in the previous model, the date on which the State intervenes proves to be crucial in determining its capacity to limit concentration in the market. If the anti-trust authorities act too soon (as soon as the market share of the dominant firm reaches the predefined threshold), the potential for increasing returns is still high and few firms have yet entered the sector: a monopoly will quickly re-emerge from the dominant firm that has been split in two. If, on the contrary, the anti-trust authorities intervene too late, the act of splitting the dominant firm will simply give rise to a duopoly, because the potential for increasing returns will be low (after a certain number of years, the technological opportunities will have been almost entirely used up), thus preventing the trajectories of the two duopolistic firms from diverging and also preventing firms in the competitive fringe from catching up. If public intervention occurs just at the right time, neither too soon nor too late, firms in the competitive fringe will have already entered the market, and when the monopolistic firm is split both the size and the efficiency of the different firms becomes homogeneous. The emergence of a monopoly will then take much longer and on average, the market will remain less concentrated than in the previous case. This model therefore illustrates the difficulties and disappointments that the public authorities can encounter when they intervene in systems characterised by increasing returns, and notably the obstacles that arise when the State attempts to counter the specific dynamics of such a system. Here again, the window of opportunity during which public intervention can hope to succeed is narrow.

In conclusion, when there is a multitude of punctual attractors, each with their different merits, irreversibility takes on its full significance and State intervention finds a new justification. A similar observation appears in neoclassical theory on the subject of the existence of multiple equilibria. The State is one of various institutions liable to select a particular equilibrium. However, even under the hypothesis of a benevolent public policy-maker who knows the respective merits of the various punctual attractors, public intervention proves to be more problematical than standard theory claims. The dynamic nature of economic processes, uncertainty,

path dependence and the heterogeneity of agents all greatly complicate the State's task. In particular, the public policy-maker encounters difficulties in controlling private behaviour simply by means of the incentives provided by the price system and also in discovering the moment at which his intervention on a system possessing its own specific dynamics would be effective.

9.3.3 The State as a pilot with imperfect information about the economic system

The State is now endowed with bounded rationality and imperfect information. It does not know the respective merits of the different punctual attractors. Therefore, the State cannot identify beforehand the products that present a danger to the environment. It discovers them through a learning process consisting in having research carried out into the characteristics of the different products. Of course, this learning process develops in parallel with the process of competition between firms.

9.3.3.1 *An illustrative model*

By extending the model with two technologies presented in the previous section, we assume that the danger Y^k presented by a product k can be located on a one-dimensional scale $[0, \max Y^k]$. However, Y^k remains unknown to the public policy-maker, because of the imperfection of information. On each date t , research is carried out upon request in order to discover the danger. The result of this research is a progression of independent random messages X_t^k with a mean of Y^k and standard deviation of σ_k . This hypothesis is not contradictory to the fact that environmental risks can be discovered accidentally, as demonstrated by the revelation of the dangers of DDT or the hole in the ozone layer. Although new research is carried out in each period, the risk assessment used by the public authorities generally displays a certain degree of inertia. Thus, in the period t , the estimation used of the level of risk Y_t^k presented by a product k takes the following form:

$$Y_t^k = (1 - \beta)Y_{t-1}^k + \beta X_t^k \quad (9.1)$$

where β represents the parameter of the degree of inertia displayed by the public policy-maker in the adjustment of risk assessments. When β tends towards zero, the weight of the first discoveries will preponderate, as all future research results can do no more than marginally correct the estimated level of risk. On the contrary, if β tends towards one, the risk esti-

mation is adjusted almost instantly to the – possibly widely divergent – research results.

When treating these environmental risks, the public policy-maker finds himself confronted with an exploration/exploitation dilemma that he cannot resolve in the “optimal” way (see chapter 1). The nature of the risks makes it impossible to prolong exploration indefinitely, because irreversible damage could be caused. In practice, therefore, the only policy that can be adopted is one of satisficing. For example, two levels of risk threshold are defined, Y_{\min} and Y_{\max} . If, on any date t , $Y_{\min} \leq Y_t^k < Y_{\max}$, then the product k is taxed at the rate τ . If, on the contrary, $Y_t^k \geq Y_{\max}$, then the product k is prohibited. Finally, if $Y_t^k < Y_{\min}$, the product k is authorised and left untaxed.

Let us now add two extra hypotheses to the model. Firstly, the decisions of the public policy-maker are assumed to be irreversible, whatever the evolution in the estimations of risks generated by research results. Secondly, we introduce a new category of product, so that three products are now in competition: one produced using a pollutant process (product 3), one produced using an initially pollutant process to which an end-of-pipe system of pollution control has been added (product 2), and one produced using what is called a clean process, where the environmental aspect has been taken into account from the start (product 1).

The complexity of this model, arising notably from the combination of the two random processes of competition and risk research, leads us to use simulations to analyse it. Tables 9.1 and 9.2 give the parameter values chosen for the simulations that are presented afterwards (for $\max Y^k = 10$).

Table 9.1 Parameter values chosen for simulations⁴⁵

Number of periods	$T = 100$
Number of new consumers entering the market in each period	$N = 500$
Initial level of risk assessment	$Y_0^1 = Y_0^2 = Y_0^3 = 0$
Initial unit cost of product 1	$c_1(0) = 1.5$
Minimum unit cost of product 1	$c_1^* = 1.1$
Initial unit cost of product 2	$c_2 = 1.3$
Initial unit cost of product 3	$c_3 = 1$

⁴⁵ To simplify matters, we assume that the unit cost of products 2 and 3 is constant; so that $c_1(t)$ here refers to an absolute cost and no longer a relative cost.

Table 9.2. Distributions of probabilities for the random draw of each risk level for the three products during public research


X_t^k	0	1	2	3	4	5	6	7	8	9	10	Y^k
$k = 1$	0.05	0.20	0.50	0.10	0.05	0.04	0.02	0.01	0.01	0.01	0.01	2.56
$k = 2$	0.01	0.02	0.05	0.12	0.50	0.15	0.10	0.02	0.01	0.01	0.01	4.24
$k = 3$	0.01	0.01	0.02	0.02	0.02	0.04	0.13	0.23	0.40	0.08	0.04	7.10

9.3.3.2 *The intervention traps facing an imperfectly informed public policy-maker*

The objective pursued by the State is simple: to ensure the survival of the most environmentally friendly product. This aim does not appear to be beyond the reach of the public policy-maker, as the learning process should lead him to tax and/or prohibit products 2 and 3, thus ensuring the survival of product 1. Yet, frequently, this scenario does not arise and product 1 disappears from the market, due to its initially higher unit cost and to the vagaries of research. The simulations that one can carry out on this simple model are rich in lessons to be drawn about key points, and notably traps, in public intervention in an evolutionary framework, where the State is endowed with bounded rationality and has to obtain information about the system it wants to regulate.

The exploration/exploitation dilemma. When the State is no longer assumed to have perfect information about the economic system, it is confronted with a crucial dilemma. On the one hand, for any intervention to be efficient, a sufficient mass of information must be collected beforehand, simply in order to identify clearly the most environmentally friendly product. On the other hand, any delay in the adoption of measures in support of this product is the source of a double loss of welfare. It prolongs the large-scale consumption of pollutant products and it can lead to the disappearance of the clean product. The exploration/exploitation dilemma confronting the public policy-maker can thus be expressed in the following manner. Concentrating too early on exploitation and neglecting all exploration may result in the wrong choices being made, based on a too-partial view of reality. On the contrary, prolonging the exploration phase too far can result in the disappearance of the best product from the set of opportunities. The State must determine the best moment to move from the phase in which

exploration is favoured (where the public authorities try to identify the real characteristics of the different products) to the phase in which the emphasis is on exploitation (where the State modifies the functioning of the market process in order to favour one of the competing products).



$$\theta = 0.1; \tau = 1.4; \beta = 0.3$$

◆ = product 1 disappears; Δ = product 1 survives

Fig. 9.2. Date of intervention and survival of the clean product

Figure 9.2 illustrates the difficulties encountered by the public policy-maker in ensuring the survival of product 1 when he delays his intervention too long. This graph displays the results from a hundred simulations. The differences in the results derive uniquely from the random nature of the research into the risks of the different products. Each simulation is located by means of two coordinates: the date on which product 2 was taxed (x -axis) and the date on which product 3 was prohibited (y -axis). Two results are possible: either product 1 survives at the end of this competitive process (white triangles), or product 1 disappears during the process (black diamonds). It can be observed that if product 3 is not prohibited before the date $t = 30$ and product 2 is not taxed before $t = 20$, the market share of product 1 will be insufficient to generate the decrease in its unit cost required to guarantee its survival. The causes of such lateness in the application of measures may be due to government failing, but it can also result simply from the random nature of the research process.

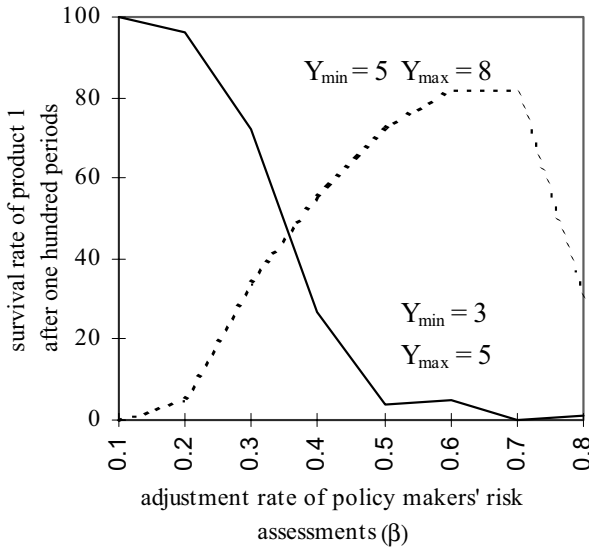


Fig. 9.3. The role of the interaction between the intervention thresholds and the adjustment rate of risk assessments in the survival of the clean product

Although public intervention should not be delayed too long, if it is to have a chance of significantly influencing the course of the economic system, any premature intervention presents the disadvantage of being based on a limited mass of information. David (1987) refers to this situation as the “blind giant dilemma”. Figure 9.3 illustrates this phenomenon. The graph presents the survival rate of product 1 after 100 periods (y -axis) and the adjustment rate β of policy-makers’ risk assessments (x -axis). The curve described by the unbroken line corresponds to low intervention thresholds Y_{min} and Y_{max} , and therefore to an early (on average) intervention. The curve described by the dotted line, on the contrary, corresponds to high thresholds and therefore a later (on average) intervention. Figure 9.3 demonstrates that if the public authorities try to intervene too early (unbroken curve, high β), product 1’s chances of survival are low. Two winning strategies emerge from these simulations: (i) setting low intervention thresholds but adjusting risk assessments slowly (unbroken curve, low β); (ii) defining higher intervention thresholds but giving more importance to the most recent research results (dotted curve, high β). The second of these strategies appears to be the only realistic one. It is hard to believe that the public authorities could deliberately ignore the most recent research results without deleterious effects on their standing in public opinion.

However, the implementation of this strategy by a benevolent government is not always very easy, given the extent to which public policy is subject to both external and internal constraints.

Public policy under external constraints. One of these constraints is classic. It derives from lobbying by firms aiming to delay public intervention by persuading the policy-makers to adopt the highest possible intervention thresholds and the slowest possible adjustment rate of risk assessments. As figure 9.3 demonstrates, such a configuration is unfavourable to the survival of product 1 and therefore to the success of public intervention. But the tendency to adjust risk assessments slowly may also derive from the difficulties the public authorities encounter in *gathering information*. In the model, we assume that the public authorities have direct access to the results of risk research and then decide on the speed of their integration into the mass of information already acquired. Often, however, this information is possessed by firms, which can therefore adopt strategies of holding information back from the State.

Public policy under internal constraints. The fact that the prohibition and taxation of products are considered to be irreversible in the model is consistent with empirical observations (examples of taxes imposed for a specific purpose and abolished once the purpose has been achieved are very rare), but this irreversibility can be detrimental to the survival of the most environmentally friendly product. When intervention thresholds are low or the adjustment rate of risk assessments is high, product 1 may be the only one of the three products to be affected by a tax or prohibition measure arising uniquely from an accident in the risk research. In the next period, new research results would probably demonstrate the overhasty nature of the State intervention. However, public policy-makers are generally reluctant suddenly to cast doubt on their past knowledge and policies. They are then incapable of exploiting the information that exploration has enabled them to gather. They cannot transform the identification of a problem into an appropriate response. Public policies generally evolve in reaction to information gathered about the economic environment. They may also reflect less objective changes such as a shift in the relative power of different interest groups.

To sum up, public intervention in the economic system is often a process than unfolds in parallel to the search for information enabling the details of this intervention to be clearly defined. When the State is no longer considered an omniscient agent, new stakes appear for public policy, but

new constraints also appear that limit its capacity to achieve its objectives. The State's primordial problem now becomes resolution of the exploration/exploitation dilemma. During the process of accumulation of information that enables the projected intervention to be more effectively defined, the autonomous dynamics of the system may evolve in such a way as to make later intervention more costly. This is all the more true when, for private agents, extension of the period of exploration increases the risk that certain choices (the purchase of durable goods, geographical location of a firm, technological choices, etc.) turn out to be in contradiction to the policy finally implemented. However, the capacities of the public policy-maker to resolve this dilemma are weakened by two factors that tend to prolong the period of exploration. One is the powerful inertia of the organisational routines implemented by the public authorities; the other is the difficulty that may be encountered in gathering the relevant information, which is often scattered among the private agents.

9.3.4 The creative State

Up to now, we have only envisaged one aspect of public intervention: the control and piloting of the course of the economic system. In other words, in an evolutionary context, only the policies influencing the process of trajectory selection have been considered. However, a relevant framework for the analysis of public policies must also consider those policies whose influence bears on the creation of variety (Metcalfe, 1995). Increasing the variety of the set of future economic choices appears to be a sensible policy in a situation of radical uncertainty. One of the key functions of public policy is then to minimise the risk of technological or behavioural lock-in, by maintaining a certain level of diversity in the characteristics of the agents in different markets, in order to increase the probability and profitability of experimental and exploratory behaviour. Thus, an evolutionary policy-maker no longer focuses solely on efficiency but also on creativity. As Nelson and Winter (1982) emphasise, when one no longer reasons, as in neoclassical theory, in terms of a given set of opportunities, the role of the State is no longer simply to choose within this set but to enlarge it. Over and above the State's traditional roles, described above, as a controller or pilot of the market process, the creative role of the State is now brought to the fore. As the State has the power and capability to structure the framework within which markets and other social institutions operate, its role is in fact incomparably more important than that played by the

other economic entities (Whalen, 1992). We can envisage several modes of public intervention:

- the State can carry out its own research and development activities in two forms: (i) create public firms whose objective is to develop environmentally friendly products, or, more likely, (ii) integrate this objective into the missions of existing public firms;
- the State can, by means of taxes, subsidies or the prohibition of certain technological choices, encourage private agents to extend their sets of opportunities independently. One of three scenarios may then arise. Firstly, the agents may react to the public intervention not by innovating but by choosing another option in the set of existing opportunities. State intervention is a failure. Secondly, the agents may extend their sets of opportunities, but choose an option that does not satisfy the objective of improving social welfare. State intervention fails again. Thirdly and finally, the private agents may extend their sets of opportunities and choose an option that does satisfy the objective of improving social welfare. State intervention is a success.

To illustrate these three scenarios, we shall take the example of the taxation of pollutant products and somewhat modify the hypotheses of the model presented above. Let us now assume that product 1 can only appear following the imposition of a tax on one of the other two products. This taxation then plays the role of a signal announcing the opening of a “launch window” for a clean product. However, firms may not necessarily be ready to lay out the necessary investments at this moment in time. In the case of the first scenario, firms react defensively to the State intervention. They do not innovate. Product 1 does not enter the market and consumers are therefore still restricted to a choice between products 2 and 3. In the second scenario, the firms innovate. Now, however, they simply modify the characteristics of their products, without altering their pollutant nature, in order to circumvent the law. For example, they could bring onto the market a product 1 that is not genuinely environmentally friendly but a modified version of product 2 or 3 designed simply to pass the tests of dangerousness carried out by a public institution. In the third scenario, if the firms react by innovating and developing a truly clean product (product 1), the public intervention can be considered a success for it will have favoured the creation of a genuinely new product.

To sum up, public intervention can be seen as a means of inciting innovation among risk-averse agents little inclined to explore certain promising

but hazardous paths. This is the case when the State carries out or funds basic research, the prospects and profitability of which are too uncertain to draw investment from private agents. Likewise, as a shareholder in large industrial groups the State is sometimes less risk-averse than private shareholders. The creative State can also incite private agents to produce innovations compatible with public objectives. The public authorities and the agents targeted by the intervention may then establish a form of cooperation, in order to prevent either a reduction in the incentive to innovate or attempts at innovation with the simple aim of circumventing the legislation.

The State can also be at the origin of the creation of original institutional forms enabling economic agents to benefit from more advantageous trajectories than those they would have followed without public intervention. Let us return to the example of the emergence of a trade union presented in chapter 8. The State is quite capable of playing the role of “germ carrier” that will enable workers organising a coalition to improve their situation. Nevertheless, once again, public intervention can never be guaranteed success. On the one hand, the union may well disappear over the course of the history of the labour market in question. On the other hand, the creation of a union may be the source of inefficiency if the balance of power between employers and unionised employees swings too far towards the latter.

9.4 Theses and conjectures

The simultaneous consideration of an economic system possessing several punctual attractors of which the prior exhaustive identification is impossible, of a universe where radical (unpredictable) novelties occur and of public and private agents all characterised by bounded rationality, opens the way to a fresh analysis of public intervention from a prescriptive point of view. From an evolutionary perspective, the State must fulfil several new roles, from the simplest, the State as market catalyst, through to the “richest”, the State as creator of the conditions enabling the economic system to reach states socially preferable to those that market forces would attain on their own, via the “pilot” State with imperfect information about the merits of the different punctual attractors, whose role is to prevent the system becoming locked-in to less socially satisfactory attractors.

This chapter has presented the obstacles that can hinder such interventions. Thus, the dynamics of information gathering can be out of rhythm with the windows of opportunity for efficient State intervention on the economic system. Consequently, the intervention may turn out to be pre-

mature or, on the contrary, too late. The causes of such mistiming are connected with the random nature of discoveries and the occurrence of radical novelties. This discordance also derives from a strategic manipulation of information by certain agents, or more simply from the absence of an appropriate system to incite them to reveal their private information. Furthermore, public intervention can also fail because of an over-evaluation by the policy-maker of his power over the economic system, either because he has under-estimated the power of the system's own independent dynamics (under the hypothesis of increasing returns, for example), or because he has under-estimated the complexity of the interaction between private behaviour and public intervention. The agents are heterogeneous in their preferences and therefore in their decisions, and their adaptive behaviour in the face of a government-driven modification in their environment is not always fully predictable.

Is it possible to define rules of conduct that help to reduce the danger of public intervention falling into these traps? Trying to draw up such rules proves to be quite a perilous endeavour. In orthodox microeconomic theory, the role of the State is limited and the probability that public intervention will end in failure is quite high. But the criterion by which this intervention is judged is, in itself, simple. State intervention is desirable if it can enable the economic system to be guided into a socially preferable situation whose attainment could not be guaranteed solely by the functioning of the market. Adopting an evolutionary view of the economic system leads us to extend the motives of public intervention well beyond the simple correction of market failures. However, drawing up prescriptive criteria as powerful as the criterion of optimality proves to be a difficult and complicated task. Even if we assume that the State remains focused on social welfare (a strong hypothesis), this welfare can no longer be associated with the concept of Pareto optimality so dear to neoclassical theory. According to neoclassical theory, the market fails when it does not lead to an optimum equilibrium; according to evolutionary theory it fails when it does not allow a desired trajectory to be followed.

Various works have indeed explored the question of rules of conduct that could structure the State's action. For Nelson and Winter (1982), if we must abandon reference to equilibrium and optimum, it becomes more appropriate to reason in terms of management and adaptation to changes. We must then content ourselves with the much more modest objectives of identifying the problems and searching for possible improvements. For Gerybadze (1992), we need to define a method for resolving multiple problems that can help to determine, at the start and then at certain key moments,

whether certain minimum conditions for the success of public intervention are satisfied, whether the key agents will behave in the predicted way, and whether it is possible to avoid an unacceptable risk of the intervention failing. In other words, evolutionary models can contribute to the development of well-argued recommendations identifying the domains in which the probability of a useful intervention is high (Lipsev and Fraser, 1997).

The prescriptive proposals set out below are drawn both from the models presented in this chapter and from the above recommendations.

Take advantage of the flexibility offered by the market. The strength of the market remains the efficiency and flexibility of its decentralised mechanism for the gathering, storage and diffusion of information. Public policy can favour the dissemination of this information. In the context of the labour market model, the State can accelerate convergence towards the equilibrium by improving the diffusion of information about labour supply and demand. Public policy can also contribute to the resolution of problems of critical mass (when a sufficient number of agents make a given choice, the system moves onto the desired trajectory on its own) or of coordination between agents (of the “prisoner’s dilemma” type). However, the State should never overlook the independent behaviour of private agents when taking public policy decisions. The individual strategies of private agents must be rendered, if not convergent, then at least compatible with the objectives of the policy-makers. Perhaps the general loss of confidence in the explicit and directive interventionist approach can be explained by the reluctance of policy-makers to acknowledge the necessarily limited nature of their means of action and the difficulties they may encounter in anticipating and thus influencing private behaviour. Policy-makers must therefore adopt more adaptive behaviour and seek to guide the market process in the desired direction rather than impose predetermined solutions. In the model of competing technologies, when the State is faced, at the beginning of the process, with a lack of information about the technologies to favour, it is preferable not to impose technological choices on firms, but to orientate them, through moderate taxation of pollutant technologies, towards what appear to be the desirable solutions at that moment in time.

Act incrementally. In a world ruled by radical uncertainty, where the future evolutions of the environment are generally uncontrollable, incrementalism can be a non-loss strategy. Proceeding by incremental changes enables the public authorities to guard against major errors in several ways (Linblom, 1959). Firstly, the results of the previous steps in a policy generally

provide the policy-maker with information about the probable consequences of the following steps. Secondly, a policy-maker should not implement policy measures aimed at immediately attaining his objectives if these measures require anticipations that are beyond his capacities or those of the best-informed agents. His choices are only successive steps in the resolution of the problem under consideration. Thirdly, by proceeding incrementally a policy-maker can test the relevance of his past predictions. Lastly, he can correct a past error easily and quickly – or at least more easily and quickly than he could if his decisions were more widely spaced in time. In other words, the objective is not to define an optimal policy but a policy that is reasonable over the short term. The model of environmental policy illustrates the validity of incrementally adjusting decisions to new information.

Do not neglect exploration. Nevertheless, the adoption of incrementalism should not lead the policy-maker to abandon phases of exploration. Renouncing the possibility of radical changes in strategy can in fact prevent the system from shifting away from the initial trajectory. The question then arises of how to resolve the exploration/exploitation dilemma. In the environmental policy model, this question focuses on the duration of the phase of research into environmental risks. In a world open to the occurrence of radical novelties, it is hard to define a robust criterion by which to resolve this dilemma. However, it does appear reasonable for the State to intensify exploration when approaching branches in the trajectory and to favour exploitation once the system is orientated onto one of the alternative branches (Saviotti, 1995).

Establish institutions to ensure that advantage can be taken of market flexibility, incrementalism and exploration. If it is advisable to explore in order to be able to modify the trajectory of the system when this proves to be necessary, still the appropriate institutional structure must be in place to learn from the exploration and, above all, make it possible to react to what has been learnt. The State must create a structure capable of questioning itself, capable of modifying its organisational routines. To change the trajectory of an economic system it is also, and above all, necessary that the possibility still exists of modifying its path. Consequently, diversity and flexibility also become sub-objectives of public intervention. Maintaining the diversity of potential trajectories reduces the risk of the system becoming locked-in on a trajectory that could prove to be unsustainable over the

longer term⁴⁶. Of course, maintaining the diversity of future economic choices requires the encouragement of innovative behaviour among private agents. How can we judge and manage diversity? Studies in this domain – notably those on biodiversity (Weitzman, 1992; Merick and Weitzman, 1998) – propose few precise rules of conduct. Generally, their recommendations are limited to stressing the necessity of favouring the survival of options that are as far apart from each other as possible. To sum up, the State should establish institutions that can minimise the risks of path dependence on public policy, ensuring the flexibility of economic choices and thus the diversity of possible trajectories, so maintaining the possibility of changing the trajectory if this should prove to be necessary.

An evolutionary approach to public intervention in the economic system provides a more faithful picture, not only of the size of the potential role of the State and the real actions that it can perform, but also of the obstacles in its way. Neoclassical theory concludes that, apart from obstacles connected with the imperfection of information, government failure can only be imputed to the policy-makers (due to rent seeking or myopia, for example). The evolutionary approach, on the contrary, affirms that even a benevolent State can encounter numerous obstacles. This affirmation is in accordance with the observation that policy-makers have moved from a certain level of euphoria in the 1960s to convinced scepticism in the 1990s as regards their ability to correct market failures (Wegner, 2003). However, compared with the well-oiled mechanism of public intervention with its unique objective and clear prescriptions as described by the neoclassical approach, some people may find the interest of an evolutionary analysis of the State's role less convincing. In fact, an increase in the complexity of the analysis and an impoverishment of policy-making prescriptions appears to be the price we must pay in order to gain realism in our hypotheses about information and the behaviour of economic agents.

Finally, the modelling of the State itself now remains to be enriched by relaxing some of the hypotheses assumed in this chapter. The main extensions that can already be envisaged are the following:

1. *Removing the hypothesis of the monolithic State.* The public sphere is in fact composed of a multitude of institutions, each one possessing its specific, possibly contradictory, objectives in addition the their shared objective.

⁴⁶ This non-sustainability can be perceived from a social (unemployment, exclusion, etc.) or from an environmental point of view (climate change, etc.).

2. *Analysing the formation of the demand expressed by economic agents to political bodies.* This demand is expressed in different ways (elections, lobby groups, strikes, etc.) and to different bodies (political parties, parliaments, governments, etc.).
3. *Representing the role of the State as guarantor of a social contract.* Through specific policies of wealth redistribution or training, the State creates the conditions of social cohesion favourable to the growth and competitiveness of a country.

References

- Akerlof, G.A. (1970): "The Market for 'Lemons': Qualitative Uncertainty and the Market Mechanism", *Quarterly Journal of Economics*, 84: 488-500.
- Arnott, R., Greenwald, B., and Stiglitz, J.E. (1994): "Information and Economic Efficiency", *Information Economics and Policy*, 6(1): 77-88.
- Arthur, B. (1988): "Competing Technologies: an Overview", in G. Dosi *et al.* (eds.), *Technological Change and Economic Theory*, Pinter, London.
- Arthur, B. (1989): "Competing Technologies, Increasing Returns and Lock-in by Historical Events", *Economic Journal*, 99: 116-131.
- Coase, R.H. (1992): "The Institutional Structure of Production", *American Economic Review*, 82(4): 713-719.
- David, P. (1985): "Clio and the Economics of QWERTY", *American Economic Review*, 75(2): 332-337.
- David, P. (1987): "Some New Standards for the Economics of Standardization in the Information Age", in P. Dasgupta and P. Stoneman (eds.), *Economic Policy and Technological Performance*, Cambridge University Press, Cambridge.
- Gerybadze, A. (1992): "The Implementation of Industrial Policy in an Evolutionary Perspective", in U. Witt (ed.), *Explaining Process and Change: Approaches to Evolutionary Economics*, The University of Michigan Press.
- Greenwald, B., and Stiglitz, J.E. (1986): "Externalities in Economies with Imperfect Information and Incomplete Markets", *Quarterly Journal of Economics*, 101(2): 229-264.
- Krugman, P.R. (ed.) (1988): *Strategic Trade Policy and the New International Economics*, MIT Press, Cambridge.
- Laffond G., Lesourne, J., and Moreau, F. (2000): "Interaction between Public Policies and Technological Competition under Environmental Risks", in: U. Cantner, H. Hanusch, and S. Klepper (eds.), *Economic Evolution, Learning and Complexity*, Physica Verlag, Heidelberg.

- Lindblom, C.E. (1959): "The Science of 'Muddling Through'", *Public Administration Review*, 19(1): 79-88.
- Lipsey, R.G., and Fraser, S. (1997): "Technology Policies in Neo-Classical and Structuralist-Evolutionary Models", *STI Review*, 22: 31-73.
- Malerba F., Nelson, R., Orsenigo, L., and Winter, S. (2001): "Competition and Industrial Policies in a 'History Friendly' Model of the Evolution of the Computer Industry", *International Journal of Industrial Organization*, 19(5): 635-664.
- Merick, A. and Weitzman, M.L. (1998): "Conflicts and Choices in Biodiversity Preservation", *Journal of Economic Perspectives*, 12(3): 21-34.
- Metcalfe, S. (1995): "Technology Systems and Technology Policy in an Evolutionary Framework", *Cambridge Journal of Economics*, 19(1): 25-46.
- Moreau, F. (2004): "The Role of the State in Evolutionary Economics", *Cambridge Journal of Economics*, 28(6): 847-874.
- Nelson, R. and Winter, S. (1982): *An Evolutionary Theory of Economic Change*, Belknap Press, Cambridge.
- Pelikan, P. (2003): "Why Economic Policies Need Comprehensive Evolutionary Analysis", in P. Pelikan and G. Wegner (eds.), *The Evolutionary Analysis of Economic Policy*, Edward Elgar, Cheltenham.
- Romer, P.M. (1986): "Increasing Returns and Long-run Growth", *Journal of Political Economy*, 94(5): 1002-1037.
- Saviotti, P.P. (1995): "Renouveau des politiques industrielles: le point de vue des théories évolutionnistes", *Revue d'Economie Industrielle*, 71: 213-226.
- Slembeck, T. (2003): "Ideologies, Beliefs, and Economic Advice – a Cognitive-Evolutionary View on Economic Policy-Making", in: P. Pelikan and G. Wegner (eds.), *The Evolutionary Analysis of Economic Policy*, Edward Elgar, Cheltenham.
- Stiglitz, J.E. (1989): "On the Economic Role of the State", in A. Heertje (ed.), *The Economic Role of the State*, Basil Blackwell, Oxford, 9-85.
- Wakeley, T.M. (1998): "Schumpeterian Process Competition, Welfare and Laissez-Faire: an Experiment in Artificial Economic Evolution", *Journal of Evolutionary Economics*, 8(1): 45-66.
- Wegner, G. (2003): "Evolutionary Markets and the Design of Institutional Policy", in P. Pelikan and G. Wegner (eds.), *The Evolutionary Analysis of Economic Policy*, Edward Elgar, Cheltenham.
- Weitzman, M.L. (1992): "On Diversity", *Quarterly Journal of Economics*, 107(2): 363-405.

- Whalen, C.J. (1992): "Schools of Thought and Theories of the State: Reflections of an Institutional Economist", in W. M. Drugger and W. T. Waller (eds.), *The Stratified State: Radical Institutional Theories of Participation and Duality*, M.E. Sharpe, Armonk.
- Williamson, O.E. (1975): *Markets and Hierarchies: Analysis and Anti-trust Implications*, The Free Press, New York.
- Williamson, O.E. (1985): *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, The Free Press, New York.

Epilogue

Having read through the nine chapters of this book, what conclusions can we make about the current state and future prospects of evolutionary microeconomics?

Naturally, the answer to this question depends on the point of view adopted. Three different approaches are possible:

- a theoretical approach focused on the models and their behaviour;
- an empirical approach which examines the consistency between theoretical analyses and experimental results and observations;
- a normative approach which focuses on economic policy.

We shall explore each of these three approaches in turn.

On the theoretical level, we believe that several positive observations can be drawn from the developments in the book. However, certain weaknesses and dangers also appear.

The first point is that it is possible to construct a coherent microeconomic discourse without resorting to tools – in their traditional form – such as utility and production functions and without focusing the argument on the properties of equilibrium and the search for optima. But this discourse naturally keeps the fundamental concepts of preferences, constraints and substitutions that lie at the heart of microeconomics. As for equilibrium, it is no longer considered as an emerging asymptotic state.

Secondly, by placing itself within the framework of stochastic dynamic processes, evolutionary microeconomics avoids some of the difficulties of the traditional approach:

- it tolerates complex environments for each agent;
- it brings into play simultaneously the three elements so dear to prospectivists: *necessity*, *chance* and *will*. *Will* is incorporated on the one hand into the anticipations and beliefs of agents, tending to relate to what is “true” (cognitive rationality), and on the other hand into their actions, relating more to what is “good” (instrumental rationality). *Chance* is introduced in various forms, not only in the re-

sults of an action (effects of a search) or on the occasion of interactions between agents (modes of matching), but also during the agents' processes of reasoning (observations, calculations) or choice of behaviour (discrete choice models). *Necessity*, finally, appears above all in the physical or technical environment of the agents or sometimes in their institutional environment (if the agents are socially conditioned).

These three elements are certainly present in classical microeconomics, but they sometimes have difficulty in coexisting, either because of the static framework used or because of constraints imposed on the dynamics.

Next, we should emphasise the fact that evolutionary microeconomics is not "anti-economics", in the sense that this term has been used by certain critics of neoclassical economics. The traditional results can always be found among the possible states if the dynamic process converges towards stable states, and they are the only ones to appear if we introduce the appropriate simplifying hypotheses.

On the other hand, evolutionary microeconomics brings out quite naturally, on the basis of certain canonical principles, a multitude of situations which traditional theory can only present by modifying its usual hypotheses. Thus, stable states may be multiple, of different "qualities" (some of them are "metastable") or not exist at all if the system fluctuates endlessly or generates unpredictable evolutions.

Lastly, evolutionary microeconomics seeks to explore the birth, future and disappearance of institutions, whereas these latter are usually introduced exogenously into classical models. And yet these institutions are used by the agents to modify profoundly the course of economic processes.

It thus appears that the evolutionary approach is the product of a paradigm particularly well-adapted to the analysis of economic phenomena.

Naturally, as a counterpart to these positive observations, there exist certain disadvantages to this approach:

- All evolutionary models require the assumption of numerous hypotheses, whether in terms of rules of procedural rationality, characteristics of random factors or the construction of anticipations. In the present state of research, the results obtained are therefore often particular and it is not always easy to specify the "classes of hypothesis" that are actually equivalent in terms of their results. However, this apparent complexity does have one merit: it often makes it possible to demonstrate that a change in one hypothesis (by the intro-

duction of a germ, or an agent with different behaviour) profoundly modifies the course of the process.

- The abundance of possible models, although it enables better understanding of phenomena, also presents the risk that arbitrary models may be constructed on the basis of hypotheses chosen to fit the circumstances. Whence the danger that evolutionary microeconomics may become a heterogeneous assortment of the good seed and the tares, lacking – at least at the present time – strict criteria of acceptability.
- The fragility of propositions based on evolutionary models is often further weakened by the mode of demonstration of their properties. It is, of course, sometimes possible to demonstrate rigorously the properties of a stochastic process when time increases indefinitely, and priority should be given to this approach. Evidently, in certain cases, the demonstration by simulation that at least one progression possessing a given characteristic exists is sufficient to prove the existence of a history exhibiting this characteristic, but the complexity of the process and the high number of parameters that it brings into play often require the use of a high number of simulations to observe either that they all share common properties for a sufficiently distant time or that the differences in their properties are generated by changes in the value of certain parameters. Strictly speaking, we should refer to conjectures rather than theorems in such cases. This intrinsic insufficiency of simulation techniques in no way detracts from the major advantage that they represent: the possibility of quickly analysing the “type” of properties appearing in stochastic dynamic models defined by a high number of hypotheses and bringing into play a multitude of variables.

We must make one last remark before leaving the theoretical approach, and we have kept it until the end of this section to underline its importance: nowhere in this book has evolutionary microeconomics been presented as an analogy or cloning of the theories of evolution found in the life sciences. The fact that certain parallels exist in the intellectual approaches in no way signifies that evolutionary microeconomics represents some sort of transposition of models developed in other fields. Evolutionary microeconomics is entirely immersed in the science of economics and seeks to answer questions raised within the context of this science. Debate

on this point is now obsolete, even if the criticism was justified for the early economic literature of an evolutionary nature.

On the empirical level, several points of view must be examined in succession.

Firstly, it can be observed that evolutionary models often bring to the fore quite naturally phenomena that have already been observed and recognised by economists, but which remain on the fringes of traditional theory because they do not lend themselves easily to integration within the usual corpus. This book makes it possible to draw up a first, far from exhaustive list of these phenomena: the dispersion of prices that generates information costs in a market (chapter 4), the appearance of bubbles in financial markets (chapter 5), the differences in the structure of one same industry in different countries (chapter 7), the diversity of the institutional forms regulating the labour market in different countries (chapter 8), the interaction, in one branch, of the market for goods and the financial market (chapter 8), the divergences between the microeconomic policies of different governments pursuing the same objectives (chapter 9), the consequences of past historical choices in terms of the location of activities or congestion costs (chapter 4), the emergence of successive forms of money (chapter 8), etc.

This observation should not prevent us from recognising the differences between the models presented in terms of the “quality” of their representation of reality. Some of them constitute acceptable caricatures; others simply prove that we can conceive of processes leading to the emergence of the phenomena that we are trying to explain.

To test the models proposed by evolutionary microeconomic theory more strictly, four directions are open for exploration.

1. The first is the use of experimental results drawn from cognitive science: this path enables the evaluation of microeconomic models in terms of the rules chosen governing agents’ behaviour.
2. The second is the use of experiments: the use of experimental economics, born around thirty years ago, to test evolutionary microeconomic theories should provide more fertile results than it has for traditional theory. The highlighting of “bubbles” is an excellent example in this respect.
3. The third direction lies in the use of long term historical analysis. By comparing the development of the electricity industry in different

countries, for example, we can identify the decisions and more generally the phenomena that constitute the branching points in evolution. We could also cite the well-known example of the QWERTY typewriter keyboard. The huge advances made in economic history over the last half century are of a nature likely to reinforce the future growth of evolutionary microeconomics.

4. The fourth path is the use of the econometric techniques customarily used when precise series of observations are available. Here we can cite the work of Kirman on the Marseilles fish market. Securities markets also represent a field that is beginning to be explored.

Naturally, a lot remains to be done to anchor evolutionary microeconomics firmly within experience and observation, but the prospects are promising.

We must, however, acknowledge that evolutionary microeconomics is open to a classic epistemological criticism: as it makes use of more generalised models, it succeeds in explaining more phenomena, but it is less refutable in the Popperian sense. We can only hope to escape from this dilemma by generalising the models to begin with and then specifying them in the determined directions.

On the normative level, the contribution of evolutionary microeconomics lies between two extremes.

- When, as in the simple market model, reference to a Pareto optimum still has significance, evolutionary microeconomics can give meaning not only to the differences between the possible stable states and the traditional stable state, but also to losses due to the length of time – of random value – taken to achieve stability. This opens the way to a comparison of the institutional forms of market as a function of the goods concerned (labour, capital, raw materials, etc.).
- When the evolution generates random trajectories over an infinite period and the concept of a collectively optimum strategy loses its meaning, evolutionary microeconomics explores more limited decisions in the terms in which they are generally tackled by governments (e.g. competition policy, environmental policy, national and regional development, etc.). The propositions made by evolutionary microeconomic theory may become more modest, but they do have the advantage of dealing with more concrete questions.

It can be observed that we have not touched on the more general case of the government itself being composed of agents involved in evolutionary interaction. No doubt such formulations will be explored over the medium term.

Finally, we do not believe it unreasonable to end this book with the expression of a double conviction:

- Evolutionary microeconomics raises the hope of significant progress in the science of economics in the 21st century: – the hope of a more all-encompassing theory, including classic microeconomics as a particular case; the hope of stronger coherence between economic theory, observations and experiments; the hope of greater understanding between economists and decision-makers.
- Evolutionary microeconomics thus opens up an immense field of exploration, a field in which we have only just started to clear the way, but which offers countless possibilities of discovery to those who venture into it.