



International Institute for
Applied Systems Analysis
www.iiasa.ac.at

Norms as Emergent Properties of Adaptive Learning: The Case of Economic Routines

Dosi, G., Marengo, L., Bassanini, A. & Valente, M.

IIASA Working Paper

WP-94-073

August 1994



Dosi G, Marengo L, Bassanini A, & Valente M (1994). Norms as Emergent Properties of Adaptive Learning: The Case of Economic Routines. IIASA Working Paper. IIASA, Laxenburg, Austria: WP-94-073 Copyright © 1994 by the author(s). <http://pure.iiasa.ac.at/id/eprint/4138/>

Working Papers on work of the International Institute for Applied Systems Analysis receive only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute, its National Member Organizations, or other organizations supporting the work. All rights reserved. Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage. All copies must bear this notice and the full citation on the first page. For other purposes, to republish, to post on servers or to redistribute to lists, permission must be sought by contacting repository@iiasa.ac.at

Working Paper

Norms As Emergent Properties of Adaptive Learning: The Case of Economic Routines

Giovanni Dosi
Luigi Marengo
Andrea Bassanini
Marco Valente

WP-94-73
August 1994



International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 71521 □ Telex: 079 137 iiasa a □ Telefax: +43 2236 71313

Norms As Emergent Properties of Adaptive Learning: The Case of Economic Routines

Giovanni Dosi
Luigi Marengo
Andrea Bassanini
Marco Valente

WP-94-73
August 1994

Working Papers are interim reports on work of the International Institute for Applied Systems Analysis and have received only limited review. Views or opinions expressed herein do not necessarily represent those of the Institute or of its National Member Organizations.



International Institute for Applied Systems Analysis □ A-2361 Laxenburg □ Austria
Telephone: +43 2236 71521 □ Telex: 079 137 iiasa a □ Telefax: +43 2236 71313

**NORMS AS EMERGENT PROPERTIES OF ADAPTIVE
LEARNING:
The case of economic routines.**

by

Giovanni DOSI^{*}, Luigi MARENGO^{},
Andrea BASSANINI^{***} and Marco VALENTE^{***}**

^{*} Department of Economics, University "La Sapienza", Rome, Italy, and IIASA, Laxenburg, Austria.

^{**} Department of Economics, University of Trento, Trento, Italy, and IIASA, Laxenburg, Austria.

^{***} Faculty of Statistics, University "La Sapienza", Rome, Italy.

Support to the research at different stages has been provided by the International Institute of Applied Systems Analysis (IIASA), Laxenburg, Austria, the Italian Ministry of University and Research (Murst 40%), the Italian Research Council (CNR, Progetto Strategico "Cambiamento Tecnologico e Sviluppo Economico") and the Center for Research in Management, University of California, Berkeley.

SYSTEMS ANALYSIS OF TECHNOLOGICAL AND ECONOMIC DYNAMICS

This new research project at IIASA is concerned with modeling technological change, and the broader economic developments that are associated with technological change, both as cause and effect. The central purpose is to develop stronger theory and better modeling techniques. The basic philosophy is that such theoretical and modeling work is most fruitful when attention is paid to the known empirical details of the phenomena the work aims to address.

Over the last decade considerable progress has been made on various techniques of dynamic economic modeling. Some of this work has employed ordinary differential and difference equations, and some of it stochastic equations. Several models have been developed in which an economic analogue of “natural selection” winnows out a population whose members have different attributes and different degrees of fitness. A number of efforts have taken advantage of the growing power of simulation techniques. Others have employed more traditional mathematics. As a result of this theoretical work, the toolkit for modeling technological and economic dynamics is significantly richer than it was a decade ago.

During the same period, there have been major advances in the empirical understanding. There are now many more detailed technological histories available. Much more is known about the similarities and differences of technical advance in different fields and industries and there is some understanding of the key variables that lie behind those differences. A number of studies have provided rich information about how industry structure co-evolves with technology. In addition to empirical work at the technology or sector level, the last decade has also seen a great deal of empirical research on productivity growth and measured technical advance at the level of whole economies. A considerable body of empirical research now exists on the facts that seem associated with different rates of productivity growth across the range of nations.

As a result of this recent empirical work, the questions that successful theory and useful modeling techniques ought to address now are much more clearly defined. The theoretical work described above often has been undertaken in appreciation of certain stylized facts that needed to be explained, like the apparent phenomenon of dynamic increasing returns, or in other cases, understanding that in many industries the distribution of firm sizes is approximately log normal. However, the connection between the theoretical work and the empirical phenomena has so far not been very close. The philosophy of this project is that the chances of developing powerful new theory and useful new analytical techniques can be greatly enhanced by performing the work in an environment where scholars who understand the empirical phenomena provide questions and challenges for the theorists and their work.

The research will focus upon the following three areas:

1. Technological and Industrial Dynamics
2. Innovation, Competition and Macrodynamics
3. Learning Processes and Organisational Competence.

Abstract

Strategic interaction among autonomous decision-makers is usually modelled in economics in game-theoretic terms or within the framework of General Equilibrium. Game-theoretic and General Equilibrium models deal almost exclusively with the existence of equilibria and do not analyse the processes which might lead to them. Even when existence proofs can be given, two questions are still open. The first concerns the possibility of multiple equilibria, which game theory has shown to be the case even in very simple models and which makes the outcome of interaction unpredictable. The second relates to the computability and complexity of the decision procedures which agents should adopt and questions the possibility of reaching an equilibrium by means of an algorithmically implementable strategy. Some theorems have recently proved that in many economically relevant problems equilibria are not computable.

A different approach to the problem of strategic interaction is a "constructivist" one. Such a perspective, instead of being based upon an axiomatic view of human behaviour grounded on the principle of optimisation, focuses on algorithmically implementable "satisficing" decision procedures. Once the axiomatic approach has been abandoned, decision procedures cannot be deduced from rationality assumptions, but must be the evolving outcome of a process of learning and adaptation to the particular environment in which the decision must be made. This paper considers one of the most recently proposed adaptive learning models: Genetic Programming and applies it to one the mostly studied and still controversial economic interaction environment, that of oligopolistic markets.

Genetic Programming evolves decision procedures, represented by elements in the space of functions, balancing the exploitation of knowledge previously obtained with the search of more productive procedures.

The results obtained are consistent with the evidence from the observation of the behaviour of real economic agents.

1 - Introduction.

As Kenneth Arrow - himself one of the major contributors to rational decision theory - puts it, a system of literally maximizing norm-free agents "... would be the end of organized society as we know it" [Arrow (1987), p. 233]. And indeed one only rarely observes behaviours and decision processes which closely resemble the canonical view from decision theory as formalized by von Neumann, Morgenstern, Savage and Arrow.

What are then the characteristics of norm-guided behaviours? And where do norms come from? Can they be assumed to derive from some higher-level rational choice? Or can one show different kinds of processes accounting for their emergence?

In this work we shall discuss these issues and present an evolutionary view of the emergence of norm-guided behaviours (i.e. routines) in economics.

We shall call rules all the procedures linking actions and some representation of the environment. In turn, representations are likely to involve relations between environmental states and variables and require the fulfilment of certain conditions (IF-THEN rules). It is a familiar definition in Artificial Intelligence and cognitive psychology (see Newell and Simon (1972) and Holland *et al.* (1986)). Of course representations may encompass both environmental states and internal states of the actor; and the action part may equally be a behaviour in the environment or an internal state, such as a cognitive act ¹.

Further, we shall call norms that subset of rules which pertain to socially interactive behaviours and, in addition, have the following characteristics:

- 1) they are context-dependent (in ways that we shall specify below), and
- 2) given the context, they are, to varying degrees, event independent, in the sense that, within the boundaries of a recognized context, they yield patterns of behaviour which are not contingent on particular states of the world.

This definition of norms is extremely broad in scope and encompasses also behavioural routines, social conventions and morally constrained behaviours ². Thus our definition includes the norm of not robbing banks, but excludes robbing or not robbing banks according to such criteria as expected utility maximization; it includes the "rules of the games" in game theoretical set-ups, but excludes the highly contingent behaviours which rational players are supposed by that theory to engage in thereafter.

Our argument is divided into two parts. First, we ask what is the link between norms, so defined, and the "rational" decision model familiar in the economic literature. In particular we shall address the question whether, whenever one observes those types of norm-guided

¹ Clearly, this very general definition of rules includes as particular cases also the procedures for decision and action postulated by "rational" theories.

² These finer categorization are quite familiar in political sciences: see for example the discussion in Koford and Miller (1991). On the contrary, the broader notion of norms adopted here includes both moral constraints and positive behavioural prescriptions (i.e. both "morality" and "ethicity" in the sense of Hegel).

behaviours, they can be referred back to some kind of higher-level rational act of choice among alternative patterns of action. We shall claim that this is not generally the case. The empirical evidence, even in simple contexts, of systematic departures of judgements and actions from the predictions of the rationality model is now overwhelming³. Here however we are not going to discuss such evidence, rather we shall pursue a complementary line of enquiry and show that, with respect to an extremely broad set of problems, a 'rational' choice procedure cannot even be theoretically constructed, let alone adopted by empirical agents. Drawing from computation theory and from the results of Lewis (1985a) and (1985b), it can be shown that many choice set-ups involve algorithmically unsolvable problems: in other words, there is not and there cannot be a universal rational procedure of choice. An optimization procedure cannot be devised even in principle: this is the negative part of the argument.

But what do people do, then? We shall suggest precisely that agents employ problem-solving rules and interactive norms, which: 1) cannot be derived from any general optimization criterion and, 2) are "robust", in the sense that they apply to entire classes of events and problems (Dosi and Egidi (1991)).

The second part of this work considers the origin and nature of these rules. The cases we shall consider regard the emergence of corporate routines applied to the most familiar control variables in economics, i.e. prices and quantities. However, there appear to be no a priori reason to restrict the applicability of the argument to economic behaviours. In fact, a similar analytical approach could be applied to several other forms of patterned behaviour in social interactions.

Concerning the origin of behavioural norms, we develop a model broadly in the perspective outlined by Holland (1975) and Holland *et al.* (1986): various forms of inductive procedures generate, via adaptive learning and discovery, representations or "mental models" and, together, patterns of behaviour: "the study of induction, then, is the study of how knowledge is modified through its use" (Holland *et al.* (1986), p.5). In our model, artificial computer-simulated agents progressively develop behavioural rules by building cognitive structures and patterns of action, on the grounds of initially randomly generated and progressively improved symbolic building blocks and no knowledge of the environment in which they are going to operate. The implementation technique is a modified version of Genetic Programming (c.f. Koza (1992) and (1993)), in which agents (firms) are modelled by sets of symbolically represented decision procedures which undergo structural modifications in order to improve adaptation to the environment. Learning takes place in an evolutionary fashion, and is driven by a selection dynamics whereby markets reward or penalize agents according to their revealed performances.

3 Cf., for instance, Kahneman, Slovic and Tversky (1982), Kahneman and Tversky (1979), Herrnstein and Prelec (1991).

A major point in the analysis which follows is that representations of the world in which agents operate and behavioural patterns co-evolve through the interaction with the environment and the inductive exploratory efforts of agents to make sense of it (actually, in our model, they cannot be explicitly distinguished)⁴. Indeed, we show that, despite the complexity of the search space (technically, the space of λ -functions), relatively coherent behavioural procedures emerge. Of course, none of us would claim that empirical agents do learn and adapt in a way which is anything like Genetic Programming, or, for that matter, any other artificially implementable formalism (but, similarly, we trust that no supporter of more rationalist views of behaviour would claim that human beings choose their course of action by using fixed-point theorems, Bellman equations, etc.). We do however conjecture that there might be a sort of "weak isomorphism" between artificial procedures of induction and the ways actual agents adapt to their environment.

The final question that we address concerns the nature of the behavioural patterns that emerge through our process of learning and market selection. In particular, in the economic settings that we consider, are these patterns algorithmic approximation to the purported rational behaviours which the theory simply assumes? Or, do they have the features of relatively invariant and context-specific norms (or routines) as defined earlier? It turns out that, in general, the latter appears to be the case: surviving agents display routines, like mark-up pricing or simple imitative behaviour (of the type "follow-the-leader") in all environments that we experimented, except the simplest and most stationary ones. Only in the latter do we see the emergence of behaviours not far from what supposedly rational agents would do (and, even then, cooperative behaviours are more likely to come out than what simple Nash equilibria would predict⁵). The context dependence of emerging routines can be given a rather rigorous meaning: the degrees of complexity of the environment and of the problem-solving tasks can be mapped into the characteristics of the emerging routines. Interestingly enough, it appears that the higher the complexity, the simpler behavioural norms tend to be and the more potentially relevant information tends to be neglected. In that sense, social norms seem to be the typical and most robust form of evolutionary adaptation to uncertainty and change.

In section 2 we shall show that, in general, it is theoretically impossible to assume that the rationality of behaviours could be founded in some kind of general algorithmic ability of the agents to get the right representation of the environment and choose the right course of action. Section 3 presents a model of inductive learning where representations and actions co-evolve. Finally, in section 4 we present some results showing the evolutionary emergence of

4 On the evolution of representations, see also Margolis (1987). In economics, such a co-evolutionary perspective is held by a growing minority of practitioners. More on it can be found in Nelson and Winter (1982), Dosi *et al.* (1988), March (1988), Marengo (1992), Dosi and Marengo (1994), Arthur (1992).

5 This is of course in line with the findings of Axelrod (1984) and Miller (1988).

behavioural routines, such as mark-up pricing. In the appendix, we provide a more detailed treatment of some of the propositions of section 2.

2 - Rational vs. norm-guided behaviour.

Let us start from the familiar view of rational behaviour grounded on some sort of linear sequence leading from 1) representations to 2) judgment, 3) choice and, finally, 4) action. Clearly, that ideal sequence can apply to pure problem-solving (for example proving a theorem, discovering a new chemical compound with certain characteristics, etc.), as well as to interactive situations (how to deal with competitors, what to do if someone tries to mug you, etc.).

At least two assumptions are crucial to this 'rationalist' view, namely, first, that the linearity of the sequence strictly holds (for example one must rule out circumstances in which people act and then adapt their preferences and representations to what they have already done) and, second, that at each step of the process the agents are able to build the appropriate algorithm in order to tackle the task at hand. Regarding the first issue, the literature in sociology and social psychology is rich of empirical counterexamples and alternative theories⁶. Indeed, in the next section of this work, we shall present a model whereby representations and actions co-evolve.

The second issue is even more at the heart of the 'constructivist' idea of rationality so widespread in economics, claiming that agents are at the very least procedurally rational⁷. In turn this implies that they could algorithmically solve every problem they had to face, if they were provided with the necessary information about the environment and the degrees of rationality of their possible opponents or partners. Conversely, the very notion of rational behaviour would turn out to be rather ambiguous if one could show that, even in principle, the appropriate algorithms cannot be constructed.

It happens in fact that computability theory provides quite a few impossibility theorems, i.e. theorems showing examples of algorithmically unsolvable problems. Many of them bear direct implications also for the micro assumptions of economic theory and, particularly, for the possibility of 'naturally' assuming the algorithmic solubility of social and strategic interaction problems⁸. We can distinguish between two kinds of impossibility

⁶ See for discussions, among others, Elster(1986), Luhmann (1979), and with respect to economics, also Dosi and Metcalfe (1991)

⁷ The central reference on the distinction between 'substantive' and 'procedural' rationality is of course Herbert Simon: see especially Simon (1976), (1981), (1986).

⁸ See Lewis (1985a) and Rustem and Velupillai (1990). Note that, loosely speaking, algorithmic solvability means that one is able to define a recursive procedure that will get you, say, to a Nash equilibrium. This turns out to be a question quite independent from proving a theorem which shows the existence of such an equilibrium.

results. First, it is possible to show the existence of classes of problems which are not solvable by means of a general recursive procedure (c.f. Lewis (1985a) and (1985b)). This implies that economic agents who look for efficient procedures for the solution of specific problems cannot draw on general rules for the construction of algorithms, because such general rules do not and cannot exist (c.f., also, Dosi and Egidi (1991)). Broadly speaking, we can say that nobody may be endowed with the meta-algorithm for the generation of every necessary algorithm.

Second, it is possible to prove the existence of single problems whose optimal solution cannot be implemented by means of specific algorithms. Hence one faces truly algorithmically unsolvable problems: economic agents cannot have readily available algorithms designing optimal strategies to tackle such problems. Therefore, unless they have been told what the optimal solutions are by an omniscient entity, they have actually to find other criteria and procedures to solve them in a 'satisfactory' way. In fact, they need novel criteria to define what a satisfactory solution is and inductively discover new procedures to accomplish their tasks (see e.g. Dosi and Egidi (1991)).

Let us briefly examine these two kinds of impossibility results: Lewis (1985a) and (1985b) proves a general result about the uncomputability of rational choice functions (on computable functions see also Cutland (1980) and Cohen (1987)).

Let $P(X)$ be the set of all subsets of a space of alternatives X where an asymmetric and transitive preference relation has been identified, we can roughly define a rational choice function as a set function $C:P(X) \rightarrow P(X)$ such that, for every $A \in P(X)$, $C(A)$ is the set of acceptable alternatives⁹. Lewis considers some compact, convex subset of $\mathbb{R}^n \setminus \{0\}$ as the space X of alternatives. Among these alternatives he takes into account only the set of recursive real numbers in the sense of Kleene and Post, i.e. the set of real numbers which can be codified as natural numbers by means of a particular Gödel numbering (for more details see Lewis (1985a)). Moreover, one operates directly on the codified values (which are called R-indices). Given a preference relation defined only on the space of R-indices and numerically representable by a computable function and given some non-triviality conditions, Lewis does not only show that the related rational choice function is uncomputable but also that so is its restriction over the sole decidable subsets¹⁰. Even more important than the proposition on undecidable sets (since in this case it may seem that the uncomputability of the function necessarily derives from the undecidability of the subsets), the result concerning only its restriction to the decidable subsets of \mathbb{R}^n is quite powerful. It means in fact that the functions are uncomputable even if their domains are computable.

⁹ Given a preference relation $>$ on a set of objects X and a nonempty set A belonging to X , the set of acceptable alternatives is defined as:

$c(A, >) = \{x \in A: \text{there is no } y \in A \text{ such that } y > x\}.$

¹⁰ Broadly speaking, we call a set decidable if there exist an algorithm which is always able to completely identify its elements, i.e. if the membership function which characterizes the set is computable.

Obviously this result does not imply that the optimal solution cannot be algorithmically determined for every $A \in P(X)$. Lewis' theorems actually prove only that no automatic procedure can generate uniformly optimal solutions over the whole family of optimization problems identified by the set of all recursive subsets of R -indices of elements of X . This would be true even if there existed some specific solution algorithm for every single problem of this family (see Lewis (1985a), p. 67). This result shows actually that there exist small enough classes (i.e. not so broad to be meaningless from a decision-theoretic point of view) of well-structured choice problems whose solution cannot be obtained by means of a general recursive procedure.

In economic theory, environmental or social interactions are usually represented by using subsets of R^n as spaces of alternative strategies. Thus, Lewis' results can be naturally extended to give proof of the generic uncomputability of the class of general economic equilibria and consequently of the class of Nash equilibria for games (see Lewis (1987)). Indeed every definition of General Equilibrium requires that agents be able to solve optimally some decision problems. Although Lewis' result is only a general impossibility one, which does not imply the uncomputability of single equilibria, its significance should not be underestimated.

Concerning game theory, it is possible to find even stronger results about the computability of Nash equilibria for specific games. Rabin's theorem (see Rabin (1957) and Lewis (1985a)) shows that there is at least one infinite stage, two-person, zero-sum game with perfect information, whose optimal strategies are uncomputable. This is a particular Gale-Stewart game, which can be described as follows: let $g: N \rightarrow N$ be a predefined total function; player A moves first and chooses an integer $i \in N$; then player B, knowing A's choice, chooses $j \in N$; finally, A, who knows both i and j , chooses $k \in N$. If $g(k) = i + j$, A wins the game, otherwise B does. Gale-Stewart games admit always the existence of one winning strategy: if $N \text{range}(g)$ is infinite, whatever i has been chosen by A, B has at least one reply which let him win the game, otherwise A does. Consequently, it is easy to show that every Gale-Stewart game has infinite Nash equilibria, with at least a subgame perfect one among them. Now, let g be computable and let $\text{range}(g)$ be a recursively enumerable set S such that $N \setminus S$ is infinite and does not have any infinite recursively enumerable subset¹¹. In such a case Rabin's theorem states that this game has no computable winning strategies. Moreover it is important to notice that the existence of simple sets has been proved (see Cutland (1980)), so that Rabin's theorem provides a strong result about the existence of games wherein there exist Nash equilibria which are not algorithmically realizable.

Another example from the second group of results mentioned in the text can be found in the properties of Post systems. A Post system (Post (1943)) is a formal logical system

¹¹ A recursively enumerable set can be defined as a set whose partial characteristic function is computable or, equivalently, as the range of some computable function (see e.g. Cutland (1980) or Cohen (1987)).

defined by a set of transformation rules, which operate on symbolic strings, and by a set of initial strings (for more details see Cutland (1980)). The problem of establishing whether a string can be generated by the initial set of a fixed Post system is called a 'word problem'. Through computability theory it is possible to show that there is an infinite number of Post systems whose word problems are unsolvable even by specific algorithms (see e.g. Thakhtenbrot (1963)). This result has a wide significance in the economic domain. Consider for example production theory: it is possible to show that there is no guarantee that optimal productive processes can be algorithmically identified, even under exogenous technical progress. Therefore it is impossible to assume that economic agents make always use of optimal processes without giving a context-specific proof.

It is worth emphasising that these impossibility results entail quite disruptive implications not only for the 'constructivist' concept of rationality, but also for the so-called *as-if* hypothesis (see Friedman (1953) and the discussion in Winter (1986)). In order to assume that agents behave as if they were rational maximizers, one needs to represent a thoroughly autonomous selection process which converges to an optimal strategy equilibrium, i.e. one must be able to formalize something like an automatic procedure which ends up with the elimination of every non-optimizing agent (or behaviour).

However, the first group of results mentioned above, implies that, for some classes of problems, we are not allowed to assume the existence of a general and algorithmically implementable selection mechanism leading in finite time to the exclusive survival of optimal behaviours. In addition, the second group of results provides examples where one can definitely rule out the existence of every such a selection mechanism.

Moreover, the minimal prerequisite one needs for a selection-based *as-if* hypothesis on behavioural rationality is the existence of some agents who use the optimal strategy in the first place (c.f. Winter (1971)). But, if the set of optimal strategies is undecidable, how can we be sure of having endowed some agent with one optimal strategy? An approximate easy answer could be that if we consider a sufficiently large population of differentiated agents, we can safely suppose that some of them play optimal strategies and will be eventually selected. But how big should our population be, given that we cannot have any idea about the size of the set of possible strategies?

Finally there is also a problem of complexity which arises in connection with rational behaviour (both under a "constructivist" view and under the *as-if* hypothesis). Broadly speaking, we can roughly define the complexity of a problem as the speed of the best computation processes we could theoretically use to solve it (c.f., e.g., Cutland (1980)). But then the speed of environmental change becomes a crucial issue: as Winter (1986) and Arthur (1992) pointed out, the *as-if* view is primarily connected with a situation without change. In fact, even when the only kind of change we allow is an exogenous one, a necessary, albeit by no means sufficient condition for the hypothesis to hold is that the speed of convergence be

higher than the pace of change. However, it is easy to find many examples of games whose optimal strategies, while existing and being computable, require too much time to be effectively pursued even by a modern computer ¹².

Even more so, all these results on uncomputability apply to non-stationary environments, wherein the 'fundamentals' of the economy are allowed to change and, in particular, various types of innovation always appear. Hence, in all such circumstances, which plausibly are the general case with respect to problem-solving and social interactions, agents cannot be assumed to 'naturally' possess the appropriate rational algorithm for the true representation of their environment (whatever that means) and for the computation of the correct action procedures (note that, of course, these impossibility theorems establish only the upper bound of computability for empirical agents).

A fundamental consequence of these negative results is that one is then required to explicitly analyze the **processes of formation of representations and behavioural rules**. This is what we shall do in the next section, by considering the emergence of rules of cognition/action in some familiar economic examples of decision and interaction.

3 - Genetic Programming as a model of procedural learning.

Genetic Programming (cf. Koza (1992) and (1993)) is a computational model which simulates learning and adaptation through a search in the space of representations/procedures. Similarly to John Holland's Genetic Algorithms (c.f. Holland (1975)), Genetic Programming (GP henceforth) pursues learning and adaptation by processing in an evolutionary fashion a population of structures which are represented by fixed length binary strings in the case of Genetic Algorithms and by symbolic functions in the case of GP.

In GP the learning system (an artificial learning agent) is endowed with a set of basic "primitive" operations (such as the four arithmetic operations, Boolean operators, if-then operators) and combine them in order to build complex procedures (functions) which map environmental variables into actions. Each artificial agent is represented by a set of such procedures and learns to adapt to the environment through an evolutionary process which involves both fitness-driven selection among existing procedures and generation of new ones through mutation and genetic recombination (cross-over) of the old ones.

General features of this model are the following:

1 - *Representations and rule behaviour*: a common feature to many computational models of learning, including the one presented here, is that of modeling the learning process not just as acquisition of information and probability updating, but as modification of representations and models of the world. But contrary to other similar models (such as genetic algorithms and

¹² Think for instance to the game of Chess or to the Rubik cube.

classifiers systems), genetic programming (GP henceforth) models learning and adaptation as an explicit search in the space of procedures, i.e. functions in their symbolic representation, which link perceived environmental states to actions ¹³.

2 - *Adaptive selection*: each artificial agent stores in its memory a set of alternative procedures of representation/action and selects at each moment of time a preferred one according to its fitness, i.e. the payoff cumulated by each procedure in the past.

3 - *Generation of new rules*: learning does not involve only adaptive selection of the most effective decision rules among the existing ones, but also generation of new ones. Learning and adaptation require a calibration of the complicated trade-off between exploitation and refinement of the available knowledge and exploration of new possibilities. GP uses genetic recombination to create new sequences of functions: sub-procedures of the existing most successful ones are re-combined with the cross-over operator in order to generate new and possibly more effective combinations.

In GP symbolic functions are represented by trees, whose nodes contain either operators or variables. Operators have connections (as many as the number of operands they need) to other operators and/or variables, if they are variables they do not have, of course, any further connection and constitute therefore the leaves of the tree.

Thus, every node can be chosen in a set of basic function (e.g. the arithmetic, Boolean, relation, if-then operators) plus some variables and constants:

$$BF = \{+, -, *, /, \dots, \text{OR}, \text{AND}, \text{NOT}, >, <, =, \dots, v_1, v_2, v_3, \dots, c_1, c_2, c_3, \dots\}$$

But basic functions can be freely defined depending on the kind of problem which is being faced (see Koza (1993) for a wide range of examples of applications in different problem domains).

The execution cycle of a GP system proceeds along the following steps:

- 0) an initial set of function/trees is randomly generated. Each tree is created by randomly selecting a basic function; if the latter needs parameters, other basic functions are randomly selected for each connection. The operation continues until variables (which can be considered as zero-parameter functions) close every branch of the tree.
- 1) once a population of trees is so created, the relative strength of each function is determined by calculating its own fitness in the given environment.
- 2) a new generation of functions/trees is generated. Two mechanisms serve this purpose: selection and genetic operators. Selection consists in preserving the fittest rules and discarding the less fit ones. Genetic operators instead generate new rules by modifying and recombining the fittest among the existing ones. The generation of new (possibly better) functions/trees in GP is similar to the genetic operators proposed by Holland for the

¹³ A more general formal tool in the same spirit and which we intend to apply in the near future is presented in Fontana (1992) and Fontana and Buss (1994), applied in the domain of biology to self-reproducing systems.

Genetic Algorithms and is mainly based on the cross-over operators ¹⁴. Cross-over operates by selecting randomly two nodes in the parents' trees and swapping the sub-trees which have such nodes as roots.

Consider for example the two parents functions:

$$P_1 := X + (Y * Z) - Z \quad \text{and} \quad P_2 := Z / (Y * X) - A$$

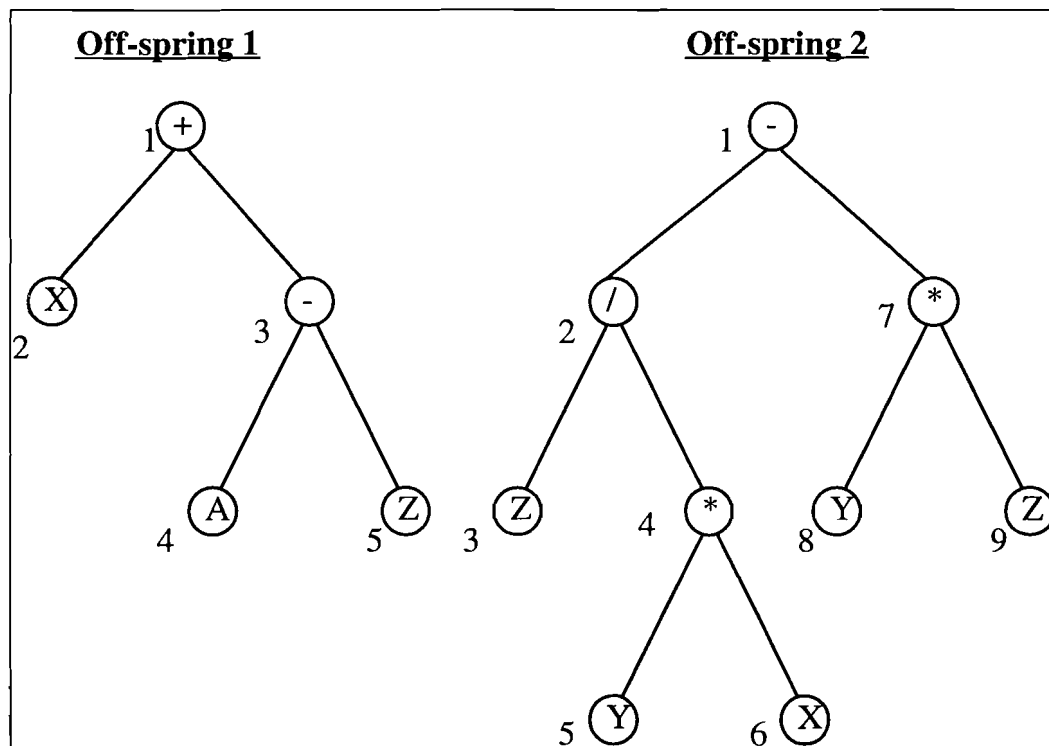
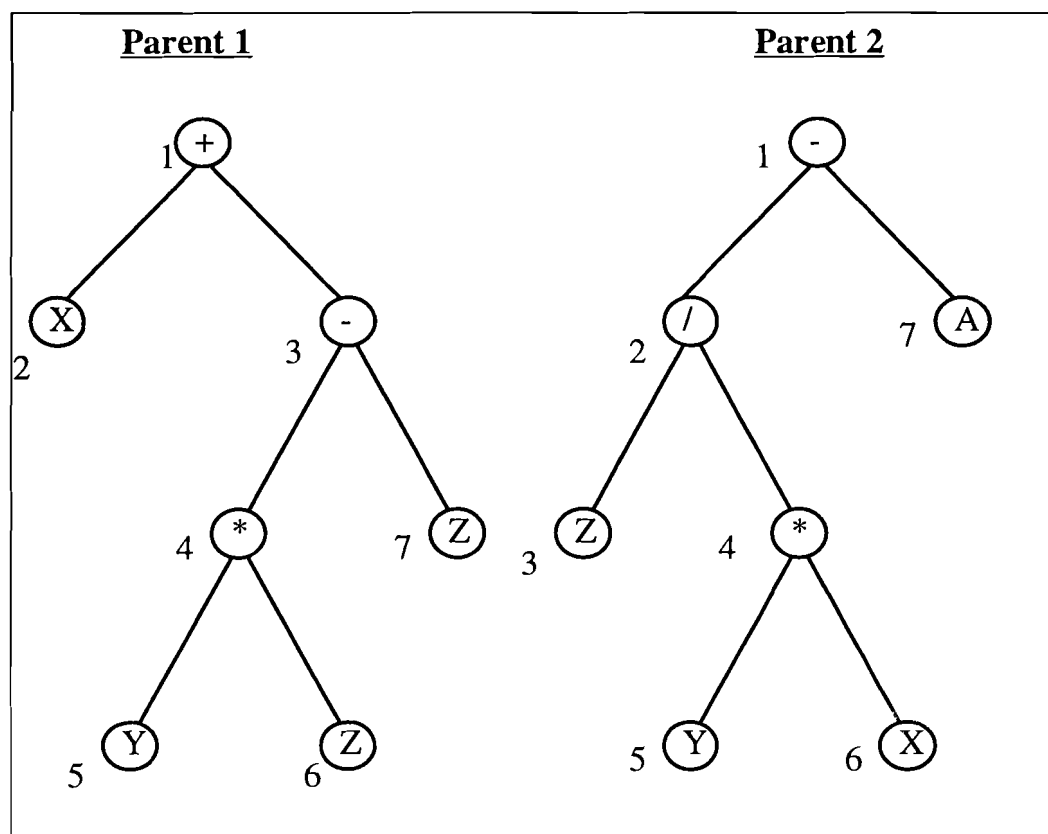
which are depicted below in their tree representation. Suppose that node 4 in the first function and node 7 in the second one are randomly selected: cross-over will generate two new 'off-spring' trees which correspond to the functions:

$$OS_1 := X + (A - Z) \quad \text{and} \quad OS_2 := Z / (Y * X) - (Y * Z)$$

Such off-spring substitute the weakest existing rules, so that the number of rules which are stored at every moment in time is kept constant.

3) go back to 1).

¹⁴ For a discussion of the power of cross-over as a device for boosting adaptation, see Holland (1975) and Goldberg (1989).



4 - Learning pricing procedures in oligopolistic markets.

Let us consider one of the most typical problems of economic interaction, namely, an oligopolistic market. A small group of firms face a downward-sloped and unknown demand curve, and have to set simultaneously their prices at discrete time intervals. To do so they can observe both the past values taken by the relevant market variables (quantity and prices) and the current value of such firm-specific variables as costs. However, they do not know either the parameters of the demand function or the prices competitors are about to set. Once all prices have been simultaneously set, the corresponding aggregate demand can be determined and individual market shares are updated according to relative prices.

This interactive set-up and the substantive uncertainty about both the exogenous environment (i.e. the demand function) and the competitors' behaviour require agents to perform a joint search in the space of representations and in the space of decision functions: GP seems therefore a natural way of modeling it.

Let us examine more precisely the structure of the market we analyse in our simulations. There exist an exogenous linear demand function:

$$p = a - bq \quad a, b > 0 \quad (1)$$

and n firms which compete in this market by choosing a price p_i . Firms are supposed to start-up all with the same market share s_i :

$$s_i(0) = 1/n \quad i = 1, \dots, n$$

Price decisions are taken independently (no communication is possible between firms) and simultaneously at regular time intervals ($t = 1, 2, \dots$). Each firm is supposed to incur into a constant unitary cost c_i for each unit of production. Once all decisions have been taken, the aggregate market price can be computed as the average of individual prices:

$$p(t) = \sum s_i(t) p_i(t) \quad (2)$$

and the corresponding demanded quantity is thus determined. Such a quantity is divided up into individual shares which evolve according to a sort of replicator dynamics equation in discrete time:

$$\Delta s_i(t) = \eta [p(t)/p_i(t) - 1] s_i(t-1) \quad (3)$$

where η is the reciprocal of the degree of inertia of the market¹⁵.

¹⁵ Note that this replication-type dynamics is consistent with the assumptions of homogenous-good output whenever one allows for imperfect information or search cost or inertial behaviour by consumers. The latter are not explicitly modelled here but they implicitly underlay the system-level mechanisms of formation of industry demand and their distribution across firms as defined by eqs. (1) to (3). The possible metaphor of these mechanisms is the following sequence: i) each firm sets its price; ii) a "public statistical office" collects all of them and announces the "price index" of the period (as from eq. (2)); iii) on the grounds of that index consumers decide the quantity they want to buy; iv) as a function of the difference between the announced average price and the price charged by their previous-period suppliers, consumers decide whether to stick to them or go to a lower-price one. Clearly, the stochastic reformulation of eq. form would be more adequate to describe the mechanism, but, for our purposes, the main property that we want to capture - namely, inertial adjustment of the market to price differential - is retained also by the simpler deterministic dynamics. Were

Finally, individual profits are given by:

$$\Pi_i(t) = [p_i(t) - c_i(t)] s_i(t) q(t) - F_i \quad (4)$$

where F_i are fixed costs, independent of the scale of production, but small enough to allow the firms to break-even for an ε excess of prices over variable costs, were they to pursue Bertrand-type competition.

We model these firms as artificial agents, each represented by an autonomous GP system, which, at each time step t , must select one pricing rule among those which it currently stores. Each artificial agent can observe at each moment of time t the following past (i.e. the values taken at time $t-1$) variables:

- average industry price $p(t-1)$,
- aggregate demanded quantity $q(t-1)$,
- individual prices of each agent $p_i(t-1)$, for $i = 1, 2, \dots, n$
- own unitary cost $c_i(t-1)$
- own market share $s_i(t-1)$


moreover it can observe its current unitary cost $c_i(t)$.

Each agent is then endowed with a few basic "elementary" operations, i.e. the four arithmetic operations, if-then operators, Boolean operators and equality/inequality operators, in addition a few integers are given as constant to each GP system.

Each agent's decision rules are randomly generated at the outset, and a preferred one is chosen for action in a random way, with probabilities proportional the payoffs cumulated by each rule in the previous iterations. Periodically, new rules are generated through cross-over and replace the weaker ones.

agents to behave as in conventional Bertrand models, eq. (3) would still converge, in the limit, to canonic Bertrand equilibria.

It must be also pointed out that our model is not concerned with the population dynamics of the industry but primarily with the evolution of pricing rules. Therefore we artificially set a minimum market share (1%) under




In order to test the learning capabilities of the model we started with the simplest model with a single agent in a monopolistic market. As shown in Figure 1, in this case with constant costs and stable demand, price converges rapidly to the optimal one. Figure 2 presents the behaviour ¹⁶ of our artificial monopolist in more complex situations in which both costs and the parameters of the demand curve randomly shift. It can be noticed that, in spite of the complexity of the task, our artificial monopolist "learns" a pricing rule which behaves approximately like the optimal one ¹⁷.

which firms cannot shrink. According to the past performance record, firms may die, in which case they are replaced by a new agent which stochastically recombines some of the behavioural rules of the incumbents.

¹⁶ For an easier interpretation, we plot in these figures only the last 100 iterations of the best emerging rule.

¹⁷ In this case and in the following ones the pricing rules which are actually learnt by our artificial agents are usually long and difficult to semantically interpret, but they behave "as if" nearly optimal pricing rules.


Figure 2 - Monopoly: Random Costs and Demand



Let us now consider an oligopolistic market. We explore two different environmental and learning scenarios. In the first one we suppose that the demand function is fixed and equal to:

$$p = 10000 - 10q$$

Moreover, unitary costs, identical for every agent, are a random variable uniformly distributed on a finite support. Finally, on the representation/action side, our artificial agents are allowed to experiment each set of rules, on average, for 100 iterations.



In figure 3 we report the results of a simulation which concerns an oligopolistic market with 9 firms. The average price is plotted against costs, while in figure 4 we report some price series for individual firms. It appears that many firms, as for example firm 4 in the particular simulation that we show, follow a pricing strategy which strictly follows cost variations.

**Figure 4a - Oligopoly: Inertial Learning
Costs and Prices of Agent '4'**

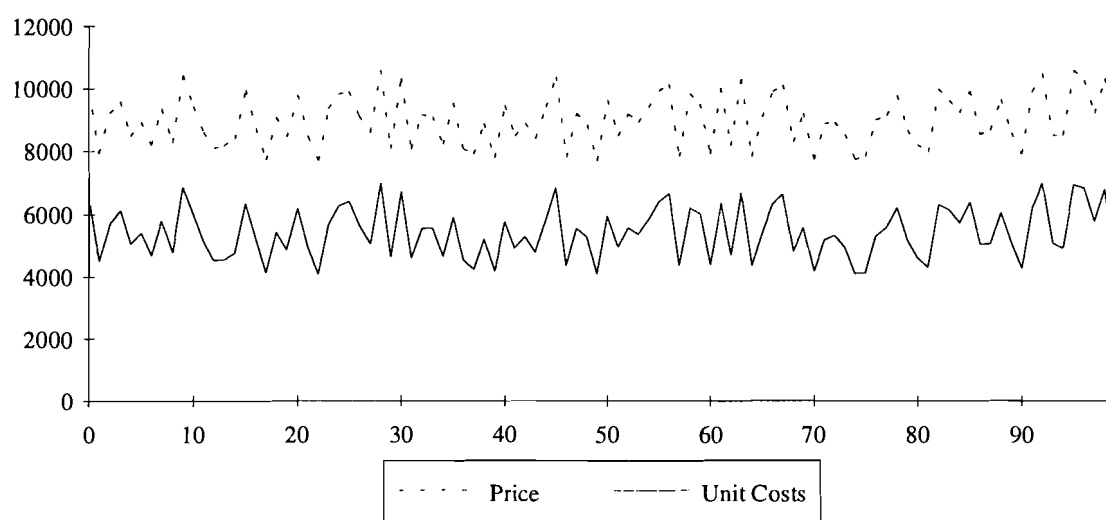




Figure 4b - Prices of Agent '0' and Agent '2'



Although emerging rules are usually quite complex ¹⁸, they behave "as if" they were simple mark-up rules. Another typical behaviour that we observe is a follow-the-leader type of pricing rule - in two out of nine firms (e.g. firm 0). Following the agent with lowest mark-up level is an extremely simple pricing rule (as other agents' prices in $t-1$ are directly observable) which allows anyway a positive average rate of profit. Moreover, with a higher number of

¹⁸ The complexity of the rules is at least partly due to the fact that our agents have to produce constants (such as mark-up coefficients) that they do not possess in their set of primitive operations and have therefore to be obtained by means of such operations on variables as to yield constant values (e.g. $(X+X)/X = 2$).

firms, the complexity of the coordination task increases and this, in turn, favours the emergence of simple imitative behaviour¹⁹.



Under the second scenario, the intercept of the demand function randomly fluctuates, drawing from uniform distribution on the support [8000 - 12000]. In addition, the individual unitary costs are given by the ratio between two variables: a component which is common to the entire industry and is represented by a random variable uniformly distributed over the interval [0,8000] and an individual productivity component, different for each firm, which is a random walk with a drift. Finally, in this scenario, we allow agents to change stochastically their sequences of rule at each period, i.e. to switch among the procedures of representation/action which they store. In this way, one forces behavioural variability (and, of course, this decreases predictability of each and every competitor). This extreme learning set up prevents any rule from settling down and from proving its value in the long term, while facing rather stable behaviours of the competitors. Despite all this, the main conclusions reached under the former scenario hold: mark-up type policies still turn out to be the most


¹⁹ Econometric estimates of the form:

$$\ln p_t = \alpha + \beta_1 \ln p_{t-1} + \dots + \gamma_0 \ln c_t + \gamma_1 \ln c_{t-1} + \dots$$

for the industry as a whole, always yield R^2 above 0.90 with significant coefficients for current costs and the first lag on prices only, and always insignificant lagged costs. Conversely, for the majority of the firms, no lagged variable significantly adds to the explanation: firms appear to follow a stationary rule of the simplest mark-up type, $p_i^t = m_i(c^t)$. However, for some firms (the "imitators") current prices seem to be set as a log-linear combination between costs and lagged average prices of the industry, or the lagged price of one of the competitors (as in the example presented in figure 4).

frequent and most efficient response to environmental uncertainty²⁰. Figure 5 and 6 illustrate costs and price dynamics for the industry

In other exercises, not shown here, we consider similar artificial agents whose control variables are quantities rather than prices. Again, as in the example presented above, a monopolist facing a stationary environment does discover the optimal quantity rule. However, under strategic interactions the agents do not appear to converge to the underlying Cournot-Nash equilibrium, but, rather, cooperative behaviours emerge. In particular, in the duopoly case, the decision rule has "Tit-for-tat" features (cf. Axelrod (1984)) and displays a pattern of the type "do at time t what your opponent did at time $t-1$ ".



It has been already mentioned that a straightforward "semantic" interpretation of the procedures which emerge is often impossible. However, their inspection - in the simplest cases - together with the examination of the behavioural patterns that they entail, allows an assessment of their nature. Some remarkable patterns appear. First, procedures which "look like" optimization rules emerge only in rather simple and stationary environments. Second, as the complexity of the representation/decision problem increases, rules evolve toward simpler

²⁰ As may be expected, estimates of the form presented in footnote 15 yield somewhat lower R^2 as compared to the previous case - both for the industry aggregate and for the individual firms -, but still most often in the range between 0.6 and 0.8.

Also the other properties of individual pricing procedures stand, and in particular simple stationary rules characterize the most successful players, as assessed in terms of cumulated profits or average market shares. Finally, in analogy to the previous learning scenario, the adjustment dynamics in aggregate prices - where the first lag on prices themselves turns out to be significant - appear to be due primarily to an aggregation effect over most often stationary rules (for a general theoretical point on this issue, cf. Lippi (1988)).

ones, involving the neglect of notionally useful information and very little contingent behaviour. More precisely, the procedure which the evolutionary dynamics appear to select either neglect the strategic nature of the interactive set-up - thus transforming the decision problem into a game "against nature" - or develop very simple imitative behaviours. In all these circumstances the resulting collective outcomes of the interaction significantly depart from the equilibria prescribed by a theory of behaviours grounded on standard rationality assumptions (this applies both to the Cournot-Nash and to the Bertrand set-ups, corresponding to quantity-based and price-based decision rules).

5 - Conclusions.

In this work we have begun to explore the properties of the procedures of representation/decision which emerge in an evolutionary fashion via adaptive learning and stochastic exploration in a space of elementary functions. Following a negative argument on the general impossibility of endowing agents with some genetic and natural optimization algorithms, we presented some preliminary exercises on the co-evolution of cognition and action rules. The results highlight the evolutionary robustness of procedures which - except for the simplest environments - have the characteristics of norms or routines, as defined earlier. Of course one can easily object that real agents indeed base their understanding of the world on a pre-existing cognitive structure much more sophisticated than the elementary functions we have assumed here, and that therefore our result might not bear any implication for the understanding of the actual evolution of norms. On the other hand, the problem solving tasks that empirical agents (and, even more so, real organizations) face are several orders of magnitude more complex than those depicted in this work. There is no claim of realism in the model we have presented, however we suggest that some basic features of the evolution of the rules for cognition and action presented here might well hold in all those circumstances where a "representation gap" exists between the ability that agents pre-possess in interpreting their environment and the "true" structure of the latter. This is obviously a field of analysis where stylized modelling exercises on evolutionary learning can only complement more inductive inquiries from e.g. social psychology and organizational sciences.

REFERENCES.

- Arrow, K. (1987), Oral History: An Interview, in G.R. Feiwel (ed.), *Arrow and the Ascent of Modern Economic Theory*, London, MacMillan.
- Arthur, W.B. (1992), *On Learning and Adaptation in the Economy*, Santa Fe NM, Santa Fe Institute, working paper 92-07-038.
- Axelrod, R. (1984), *The Evolution of Cooperation*, New York, Basic Books.
- Cohen, D. (1987), *Computability and Logic*, Chichester, Ellis Horwood.
- Cutland, N.J. (1980), *Computability: An Introduction to Recursive Function Theory*, Cambridge, Cambridge University Press.
- Dosi, G. and M. Egidi (1991), Substantive and Procedural Uncertainty. An Exploration of Economic Behaviours in Complex and Changing Environments, *Journal of Evolutionary Economics*, vol. 1, pp. 145-68.
- Dosi, G., Ch. Freeman, R. Nelson, G. Silverberg and L. Soete (eds.) (1988), *Technical Change and Economic Theory*, London, Francis Pinter.
- Dosi, G. and L. Marengo (1994), Some Elements of an Evolutionary Theory of Organizational Competences, in R.W. England (ed.), *Evolutionary Concepts in Contemporary Economics*, Ann Arbor, University of Michigan Press, pp. 157-78.
- Dosi, G. and J.S. Metcalfe (1991), On some notions of Irreversibility in Economics, in P.P. Saviotti and J.S. Metcalfe (eds.), *Evolutionary Theories of Economic and Technological Change*, Chur, Harwood Academic.
- Elster, J. (1986), *The Multiple Self*, Cambridge, Cambridge University Press.
- Fontana, W. (1992), Algorithmic Chemistry, in C. Langton, J.D. Farmer and S. Rasmussen (eds.), *Artificial Life*, Redwood City, Ca, Addison Wesley.
- Fontana, W. and L. W. Buss (1994), What would be conserved if "the tape were played twice"?, *Proceedings of the National Academy of Sciences USA*, vol. 91, pp. 757-61.
- Friedman, M. (1953), *Essays in Positive Economics*, Chicago, University of Chicago Press.
- Goldberg, D.E. (1989), *Genetic Algorithms in Search, Optimization and Learning*, Reading Mass., Addison Wesley.
- Herrnstein, R.J. and D. Prelec (1991), Melioration: a Theory of Distributed Choice, *Journal of Economic Perspectives*, vol.5, pp.137-56.
- Hirschman, A. (1977), *The Passions and the Interests*, Princeton NJ, Princeton University Press.
- Hodgson, G. (1988), *Economics and Institutions*, London, Polity Press.
- Hogart, R.M. and M.W. Reder (eds.) (1986), *Rational Choice*, Chicago, Chicago University Press.

- Holland, J.H. (1975), *Adaptation in Natural and Artificial Systems*, Ann Arbor, 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*, Cambridge Mass., MIT Press.
- Kahneman, D., Slovic, P. and A. Tversky (eds.) (1982), *Judgment under Uncertainty: Heuristics and Biases*, Cambridge University Press.
- Kahneman, D. and A. Tversky (1979), Prospect Theory: An Analysis of Decision under Risk, *Econometrica*, vol. 47, pp. 263-91.
- Koford, K.J. and J.B. Miller (eds.) (1991), *Social Norms and Economic Institutions*, Ann Arbor, University of Michigan Press.
- Koza, J.R. (1992), The Genetic Programming Paradigm: Genetically Breeding Populations of Computer Programs to Solve Problems, in B. Soucek (ed.), *Dynamic, Genetic and Chaotic Programming*, New York, John Wiley.
- Koza, J.R. (1993), *Genetic Programming*, Cambridge, Mass., MIT Press.
- Lewis, A. (1985a), On Effectively Computable Realization of Choice Functions, *Mathematical Social Sciences*, vol. 10, pp. 43-80.
- Lewis, A. (1985b), The minimum degree of recursively representable choice functions, *Mathematical Social Sciences*, vol. 10, pp. 179-88.
- Lewis, A. (1986), *Structure and Complexity. The Use of Recursion Theory in the Foundations of Neoclassical Mathematical Economics and the Theory of Games*, Ithaca, Cornell University, dept. of Mathematics, mimeo.
- Lewis, A. (1987), *On Turing Degrees of Walrasian Models and a General Impossibility Result in the Theory of Decision-Making*, Technical report n. 512, Institute for Mathematical Studies in the Social Sciences, Stanford University.
- Lippi, M. (1988), On the Dynamics of Aggregate Macro Equations: from Simple Micro Behaviours to Complex Macro Relationships, in Dosi *et al.* (1988).
- Luhmann, N.(1979), *Trust and Power*, Chicester, Wiley.
- March, J.G. (1988), *Decisions and Organizations*, Oxford, Basil Blackwell.
- Marengo, L. (1992), Coordination and Organizational Learning in the Firm, *Journal of Evolutionary Economics*, vol. 2, pp. 313-26.
- Margolis, H. (1987), *Patterns, Thinking and Cognition: A theory of judgement*, Chicago, Chicago University Press.
- Miller, J.H. (1988), *The Evolution of Automata in the Repeated Prisoner's Dilemma*, Santa Fe Institute, working paper.
- Nelson, R.R. and S.G. Winter (1982), *An Evolutionary Theory of Economic Change*, Cambridge, Mass., Harvard University Press.
- Newell, A. and H. Simon (1972), *Human Problem Solving*, Englewood Cliffs NJ, Prentice-Hall.

- Post, E. (1943), Formal Reductions of the General Combinatorial Decision Problem, *American Journal of Mathematics*, vol. 65, pp. 197-215.
- Rabin, M.O., (1957), Effective Computability of Winning Strategies: Contributions to the Theory of Games III, *Annals of Mathematics Studies*, vol.39 pp. 147-57.
- Rustem, B. and K. Velupillai (1990), Rationality, Computability and Complexity, *Journal of Economics Dynamics and Control*, vol.14, pp.419-32.
- Samuel, A.L. (1959), Some Studies in Machine Learning using the Game of Checkers, *IBM Journal of Research and Development*.
- Simon, H.A. (1976), From Substantive to Procedural Rationality, in S.J. Latsis, (ed.), *Method and Appraisal in Economics*, Cambridge University Press, pp. 129-48.
- Simon, H.A. (1981), *The Sciences of the Artificial*, Cambridge Mass., MIT Press.
- Simon, H.A. (1986), Rationality in Psychology and Economics, *Journal of Business*, vol.59, supplement.
- Thrakhtenbrot, D.A. (1963), *Algorithms and Automatic Computing Machines*, Boston, Mass., D.C. Heath and Co.
- Winter, S.G. (1971), Satisficing, Selection and Innovating Resonant, *Quarterly Journal of Economics*, vol.85, pp. 237-61.
- Winter, S.G. (1986), Adaptive Behaviour and Economic Rationality: Comments on Arrow and Lucas, *Journal of Business*, vol.59, supplement.