

Routines, Hierarchies of Problems, Procedural Behaviour: Some Evidence From Experiments¹

Massimo Egidi

Department of Economics
University of Trento
Via Inama, 1
38100 TRENTO- Italy
Tel: +39-461-882223
Fax: +39-461-882222
megidi@risc1.gelso.unitn.it

WP 1995-03

ABSTRACT

A laboratory experiment was performed as replication of the original one created by M. Cohen and P. Bacdayan at Michigan University. It consists in a two-persons card game played by a large number of pairs, whose actions are stored in a computer's memory. In order to achieve the final goal each player must discover his sub-goals, and must coordinate his action with the partner's one. The game therefore involves the division of knowledge and cooperation among players, and gives rise to the emergence of organizational routines. It is suggested that the organizational routines, i.e. the sequences of patterned actions which lead to the realization of the final goal, cannot be fully memorized because of their variety and number. It is shown that players do not possess all the knowledge needed by an hypothetical supervisor to play the best strategy: they generally explore only a limited part of the space of the potential rules, and therefore learn and memorize a simple, bounded set of "personal" meta-rules. These meta-rules, also called "production rules" in standard Cognitive Science's language, are of the form <If "Condition" then "Action">. Each "Condition" can concern either the game configurations or the partner's action. In the former case the identification of an appropriate "Action" depends on the sub-goals exploration. In the latter it depends on the recognition (or discovery) of interaction rules; in this eventuality the production rule embodies a dynamic - and possibly cooperative - reaction to the partner's action. Organizational procedures (routines) therefore emerge as the outcome of a distributed process generated by "personal" production rules. These routines, as in von Hayek's view, "can be understood as if it were made according to a single plan, although nobody has planned it." (Hayek, 1980, p. 54). Empirical evidence is provided to support the above statements.

¹Paper prepared for the International Economic Association Conference on Rationality in Economics, (ICER, Torino, dec. 1993), forthcoming in Arrow et al. (editors), *The Rational Foundations of Economic Behaviour*, London, MacMillan.

1. Introduction: procedural rationality , organizational routines and problem solving.

The idea that the learning activity plays a central role in human decision making derives from the pioneering work of Simon, March and Newell. In 1956, in a path-breaking article which constituted a first crucial step in analyzing rationality within organizations, Cyert, Simon and Trow carried out an empirical analysis of managerial decisions which revealed an evident "dualism" of behaviour: on the one hand, a coherent choice among alternatives; on the other, a search for the knowledge required to make the choice. For several months, an observer monitored the decisions made by the executives in a particular firm and recorded a number of features entirely at odds with the prescriptions of the decision theory then current. The principal finding was that, when decisions are made in conditions of high uncertainty - conditions, that is, poorly structured in terms of knowledge, beliefs, information - their outcomes are not easily assessed. This activates a *search* process intended to frame all the elements involved in the decision.

A dichotomy between types of decision consequently arises, in relation to the different level to which the problems involved are cognitively structured :

"Decisions in organizations vary widely with respect to the extent to which the decision-making process is programmed. At one extreme we have repetitive, well defined problems (e.g., quality control or production lot-size problems) involving tangible considerations , to which the economic models that call for finding the best among a set of pre-established alternatives can be applied rather literally. In contrast to these highly programmed and usually rather detailed decisions are problems of non-repetitive sort , often involving basic long-range questions about the whole strategy of the firm or some part of it, arising initially in a highly unstructured form and requiring a great deal of the kinds of search processes listed above." (Cyert, Simon and Trow, 1956, p.238)

What are the features of decisions taken in a highly programmed decision context? Many field observations of human behaviour in organizations show that in well-structured conditions, where subjects must implement sequences of choices vis-à-vis alternatives well known to them, their behaviour becomes *routinized*. The sequence of choices confronted by individuals performing an organizational task constitutes a repetitive *procedure* which is memorized, becomes familiar to those executing it, and presents well-defined alternatives codified according to the variants arising from changing external circumstances. Most human activity within economic organizations takes the form of this procedural and routinized behaviour.

If individuals are able to memorize repeated sequences of decisions deriving from their interactions with others, and then execute at least parts of them "automatically", the role of routines becomes clear: they enable individuals to save on "rational computation" and radically reduce the complexity of individual decisions. This narrowly restricts the area in which substantive rationality needs to be exercised, and therefore reduces decision "errors".

It must be emphasized that this role of routinized behaviour within the decision process - in order to reduce the need for rational computation - is only a working hypothesis, and must be analyzed carefully in order to understand its consequences and to find empirical evidence for it. An effort to explore more precisely how routinized behaviour reduces the need for "rational computation" has been made by Nelson and Winter, on the basis of the methodological principles enunciated by M. Polanyi in *Personal Knowledge* (1958). They note that some behavioural sequences consist of actions which are often partially inarticulate, i.e. they are not expressed linguistically, and need not to be transmitted in the form of messages. This feature leads Nelson and Winter to the problem of how *tacit knowledge* is

formed, transferred and stored in memory. This is an interesting starting point for exploring how cognitive skills, which arise through experience and cooperation, are stored in memory and by consequence become building blocks for subjects who have to solve problems. Along this line of research, Cohen and Bacdayan (1991) suggest that routines are stored as *procedural* memory; following Squire's (1987) distinction between procedural and declarative memory they claim that " *procedural* memory appears to be the form that stores the components of individual skilled actions - for both motor and cognitive skills. It is distinguished from *declarative* memory, which provides the storage of facts, propositions, and events." (1991, p. 5). They use a laboratory experiment to analyze the emergence of procedural behaviour by two subjects involved in a game which requires coordination and cooperation, and its "sedimentation" in memory. I will return later to the specific issues raised by their article, because the present paper is based on a replication of their experiment. The general point at issue here is how the acquisition and memorization of cognitive skills takes place, and how its transfer is possible, i.e. how skills can be re-used. Unfortunately, as Singley and Anderson show (1989), the range of transfer of procedurally encoded skills seems to be very restricted; as the two authors demonstrate, the progress in this domain have been very limited and further progress is required before we can adequately understand the phenomenon. It is clear that any future success along this direction will progressively shed light on the unsolved aspects of the process of skills acquisition and the creation of new routines.

My purpose here is to explore some of the features of human problem solving, skills creation and the emergence of routines by conducting careful analysis of the results yielded by replication of Cohen and Bacdayan's experiment.

To clarify the question, let us frame the problem of skills creation in relation to decision making processes. When a decision is highly unstructured, we are in the situation that Cyert and Simon called non-programmed decisions, where the predominant role is played by the search for the knowledge required to solve a problem. In this situation, not only must subjects gather information, they must also be able to select the information and knowledge that is effectively relevant to their purposes and to assimilate it into the system of knowledge that they already possess. To do so, they must have a "level of competence" adequate to the situation of choice; they must, that is, implement skills of learning and problem solving. The core of the decision-making process is therefore the activity of search and learning that furnishes actors with the information and knowledge they require to achieve their goals.

The conditions for standard choice theory to be applied are entirely lacking, because the preferences orderings are highly incomplete, decisions are intertemporally inconsistent and choices are largely ineffective effective in relation to the goals to be pursued. The most important part of the process is driven by the ability of the subjects to formulate and solve problems. The problem is how to model these kind of situation.

One way of tackling the problem is based on the idea that the solution to a task can be found by recursively decomposing it into a set of simpler interrelated tasks. This idea, now widely used in the problem solving theoretical domain, originates from March and Simon's book "Organizations" (1958). To better understand its origin, let us distinguish between two different aspects of routines: on the one hand, there is the aspect that I have emphasized above, i.e. their relationship to the cognitive structures and decision processes of individuals; on the other, routines can be regarded as elementary units which form the basic structures internal to organizations.

Within the organization, we can consider as *routine* any procedure which provides for the execution of a specific task; it is therefore a procedure which solves a set of problems internal to the organization. A procedure can be described as a set of instructions determining the actions to be taken when dealing with a particular circumstance.

It seems natural, therefore, to model a procedure as a *program*, in the specific sense given to the term by computation theory, as a list of instructions in an artificial language. This enables us to represent procedures formally and to model *procedural rationality* (March and Simon, 1958, Ch. 6).

Of course, if we try to describe the functioning of an organization as governed by a hierarchy of procedures, we *cannot* attribute to individuals the ability to memorize, routinize and execute them with precision: if we assume that individuals behave like automata endowed with unlimited memory, we fall into an error similar to that committed by the proponents of "Olympian" rationality when they attribute to economic agents unlimited computational skills.

In reality, individuals do not usually possess precise and detailed knowledge of organizational procedures; they have "incomplete" knowledge, and they are able to complete it by *recreating its missing components*.

The last observation sums up the central problem of the present paper: how can one to explain the ability of individuals to create procedures or recreate missing parts of it while they are coordinating their skills and intellectual efforts to achieve a given goal.²

As I have said, an important approach used in the cognitive sciences adopts the idea that individuals solve problems by decomposing them into a set of interrelated sub-problems. The methodological status of this approach is twofold: on the one hand, the approach originates from empirical observation of *real human behaviours*; on the other, it is the outcome of careful analysis of the *nature of problems* (mainly in games and other formalized contexts) and therefore is normative in character. As usual, the normative and descriptive aspects of this approach cannot be separated: the models based on the decomposition of problems - like Laird, Newell and Rosebloom's (1987) *Soar* - do not pretend to be accurate replications of the deep psychological mechanisms behind human behaviour; they merely seek to replicate human behaviours or at least successfully to compete with them. By consequence, an implicit "as if" hypothesis is presupposed by this approach. It is my purpose to try to reduce the domain of validity of the "as if" hypothesis and to verify - by means of laboratory experiments - to what extent the problem decomposition approach can be considered as providing a good explanation of human learning process, at least in playing games.

I will examine how individuals learn to identify the sub-goals, to link them to each other, and to build "production" rules and procedures for solving elementary sub-goals. On this basis I will explore the emergence of behavioural rules - which involve cooperation - and their relationship with routinized behaviours.

One of the most interesting results which emerge from the experiments is this: if a problem to be solved has a variety of different aspects, as happens in card games when there is a large number of different initial configurations, the learning process dominates the actors' activity. To achieve their goals, players make sequences of moves which depend on the configurations of the game; these sequences are *organizational routines*, which cannot be memorized by players because of their variety and number. Players do *not* keep all knowledge and information they need to play stored in memory: they create and memorize a set of simple "meta rules"³ which allow them to generate the organizational routines . These rules are

²These two abilities are quite different the one from the other: the first one concerns the design and planning ability, while the second one concerns adaptation and learning. The combination of the two activities suggest to describe the organizational evolution as a "punctuated equilibria" process, with discontinuities: the planner , or a hierarchy of planners, suddenly modify the organizational structure, and the employees adapt the "informal" organization to the new plans; while the idea of organizational learning as the result of the interaction (coordination) among local forces driving the local solution of problems, leads to a darwinian vision of the organization, chacterized by a continuous process of internal evolution.

³ The experiment design does not allow to verify if players are aware of the meta-rules they use, and consequently if these rules can be considered as *tacit* ones.

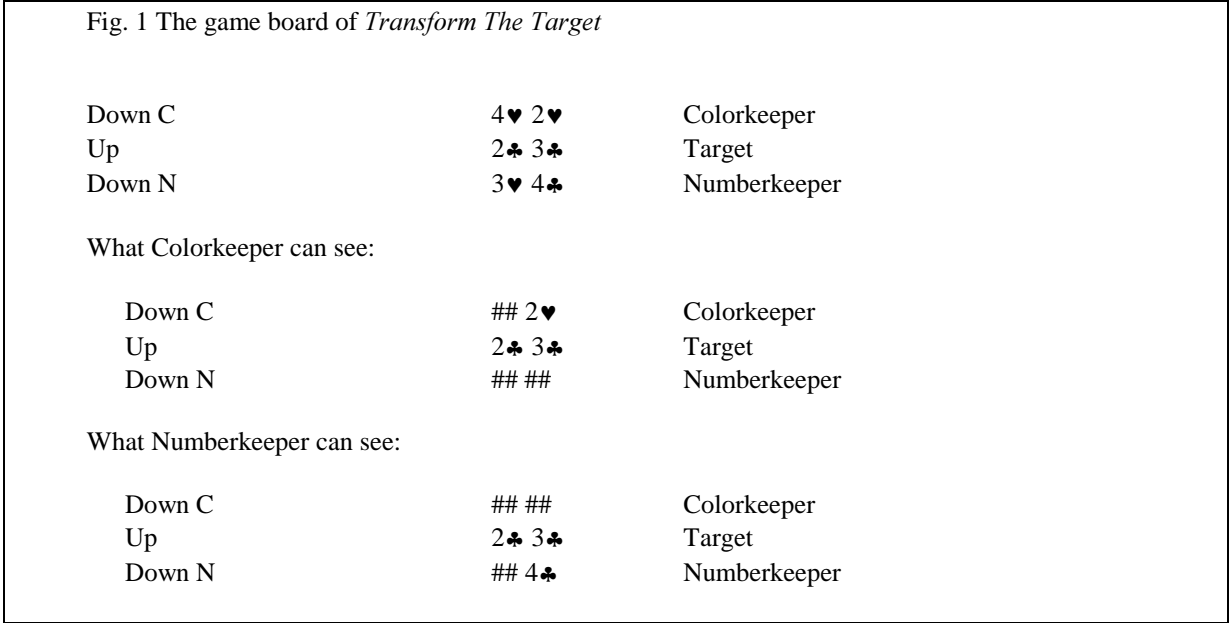
elementary "production rules" (in the standard sense of cognitive sciences), which are the result of sub-goals identification and allow players to recreate the missing knowledge at any particular moment.

2. Transform The Target.

The game "Transform The Target", described here, was first devised by Michael Cohen and Paul Bacdayan for the experimental study of the development of behavioural routines in a cooperative context. It is a card game in which two players must cooperate in order to achieve a pre-established goal. A fixed payoff is awarded to each pair of players when the tournament is finished. This payoff is the greater the more rapidly the players complete their games, and it is awarded to the pair, not to the individual, so that both players have an incentive to cooperate.

The game is played on the computer rather than with cards. The players move the cards they see on the computer screen using the mouse. The two players' moves and all the movements of the mouse are recorded, thus providing a detailed information base on the games. A tournament of 40 games involving 32 pairs of players was organized and recorded, replicating the original tournament by Cohen and Bacdayan. I give below the rules of the game as described by the two authors.

The game is played by two players using six cards: the 2,3 and 4 of a red suit, and the 2,3 and 4 of a black suit. The players are called respectively Numberkeeper and Colorkeeper. When the cards are dealt, each player can see only : the card in his hand, the card on the target and the card in "Up" position. The board is illustrated in Fig. 1. Colorkeeper cannot see the card that Numberkeeper has in his hand and vice versa.



In each hand that is played, the ultimate object is to put the red two into area marked "Target". A move in the game is an exchange of the card in a player's hand with one of the cards on the board (or a "Pass", making it the other player's turn).

Each player is subject to a restriction on moves: Colorkeeper may make exchanges with the target area only if the colour in the target is preserved. Numberkeeper may exchange with the target only if the action preserves the number in the target area. Exchanges with board

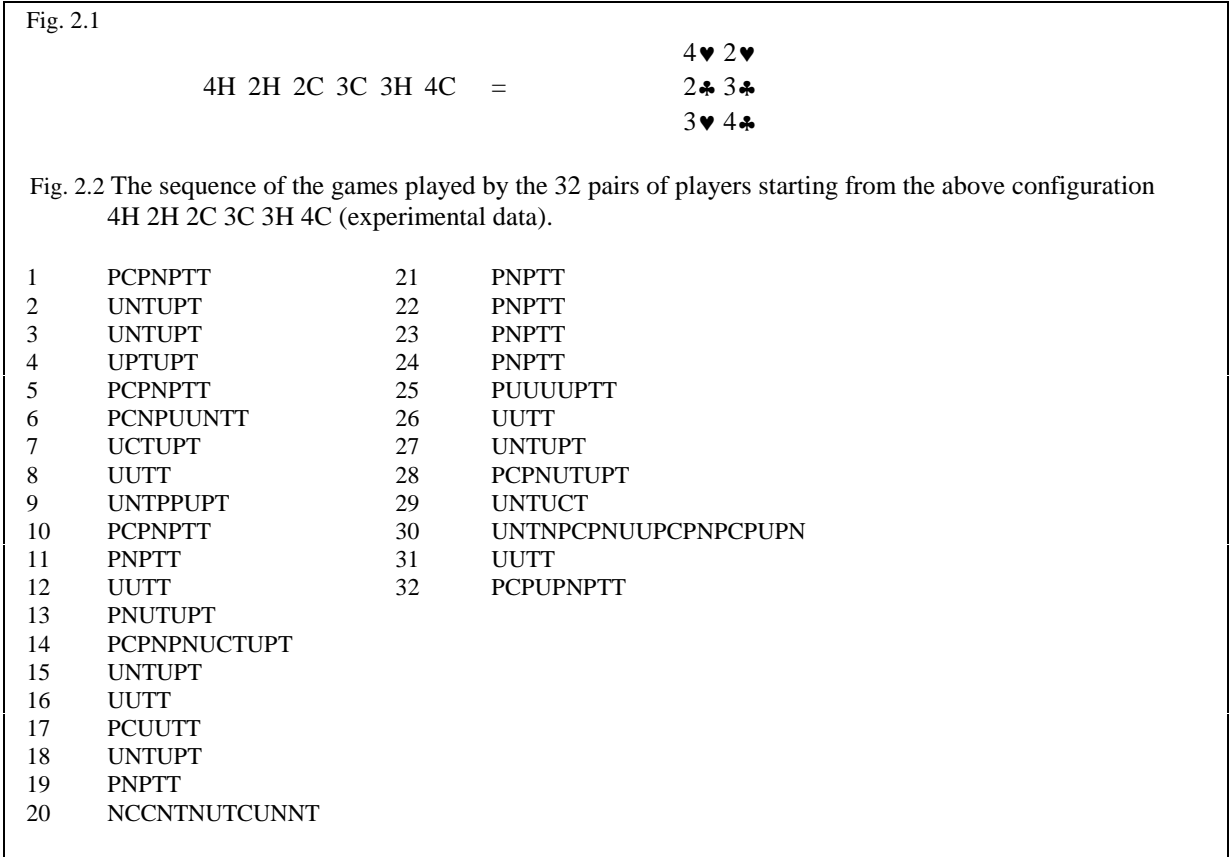
areas other than the target are not restricted. (Cohen and Bacdayan, 1991). Colorkeeper moves first.

Summarizing, during the game a player can make one of the following moves:

- U - exchange his card with the card Up
- C - exchange his card with the face-down card on the left of Ck's card
- N - exchange his card with the face-down card on the left of Nk's card
- T - exchange his card with Target
- P - pass

A state of the game is defined by the distribution of the cards on the board. Players do not have full and direct information about the current state of the game, because there are covered cards. Their knowledge of the situation is incomplete, and they must conjecture what state of the game is going on. The more the game proceeds, the wider is the information collected by players (if they use information in an intelligent way).

Each of the 32 pairs of players played the tournament using the same distribution of cards as all the others. Figure 2 shows the sequence of the games played by the 32 pairs with an initial configuration of 4H 2H 2C 3H 4C (nH here denotes card n of Hearts, and mC denotes card m of Clubs). This sequence represents the initial board of the game according to the following convention: the cards placed on the board are treated as the elements of a matrix with 3 rows and two columns B(i,j), (i=1..2, j=1..3). Fig. 2.1 gives an example.



For a given initial distribution of cards, each player has a wide set of possible strategies at his disposal. To find a strategy (and if it exists the best one) a player can follow the standard von Neumann-Morhgenstern procedure: he can build the game tree by applying in order all the rules to the possible initial states.

However, whereas in the case of (perfect) information games like chess where the initial state is known to both players, so that a tree of strategies can be constructed right from the unique initial state, in this game the initial state is unknown to the two players. At the starting configuration of a hand, each of the two players ignores the 3 covered cards : each player must therefore deduce the 6 possible configurations that may occur (counting the disposition of the three unknown cards in the three covered positions, we have 6 possible configurations).

At the beginning of the game neither of the players knows which of the six possible configurations is the 'real' one. This they can only discover by acquiring further information as they play their hand. For every initial board each player should therefore prepare a strategy. This manner of proceeding obviously would require a memory and a computing capacity which far exceeds normal *non-specialized* human capacities. These are the typical conditions of computational complexity which highlight the relevance of the assumption of bounded rationality .

Moreover, the experiment shows that individuals do not proceed in the manner suggested by the classic model of Olympian rationality. The players are unable to use all the information available, and in many cases they are not able to memorize all of it.

This depends on the constraints imposed by two types of memory - i.e. the ability to process a little information instantaneously, and the ability to process a lot of information but only through storage in the long-term memory - which make it impossible to adopt a plan which incorporates all 6 possible configurations (Cohen and Bacdayan 1991). The experimental data (see below) confirm the fact that the players do not use all the available information. This is because they cannot memorize all of it simultaneously, and because they do not conduct a thorough search of all the rules space.

If each player is unable to memorize all the configurations and to discover the best strategy available, which kind of search process can he activate to find a (good) strategy? The classic analysis of the process of strategic search - Nilsson (1980), Pearl (1984), Newell (1990) - is based on the idea that to each state of the tree can be assigned an evaluator which enables the player to choose among the different local strategies which appear (locally) optimal. The problem with this approach is that the evaluators are supposed to be "given" exogenously; there is no method by which the evaluators can be generated endogenously.

A different approach is based on the idea that I discussed in the Introduction, namely that the solution to a task can be found by decomposing it into a set of simpler interrelated tasks.

As we shall prove experimentally, players do not normally explore all the possible strategies, but move on the space of strategies and sub-strategies, solving local problems and using them as building blocks with which to solve further problems.

We shall see that the search in the space of sub-problems is a highly uncertain and conjectural process. In fact when a problem has been decomposed into a set of sub-problems, generally not all of the sub-problems will be immediately solvable. The main goal - at this stage of the exploration - is to understand the relations among sub-problems. During their attempts to understand the connections among problems, players do not focus their attention immediately on the existence of procedures which allow the sub-problems to be solved. The problem space is explored in the sense that new sub-problems and new connections among the sub-problems are discovered; the aim being to reduce recursively complex problems to

simpler ones and to understand their reciprocal relationships. This exploration, however, does not require that the solutions of the 'simpler' problems should be known.

3 The space of the sub-goals


Let us see how the space of the sub-problems for the game described above (henceforth *Transform The Target*) can be constructed. As I have said, the problem is to put 2♥ in the Target. At the beginning of the card game any other card may be in the Target.

Reasoning 'backwards' and using the rules of the game, one finds that 2♥ can be put into the Target area only if (Fig. 3):

1. 3♥ or 4♥ is in the Target and the player with 2♥ in his hand is the Colorkeeper;
2. 2♣ is in the target and the player with 2♥ in his hand is the Numberkeeper.




The problem can therefore be solved if at the beginning of the game one of the three cards 4♥, 3♥, 2♣ is in the Target area. Let us instead suppose that one of the remaining cards, i.e. 3♣, or 4♣ is in the Target. In this case, the game can be solved if we are able to reduce it to the situation above, i.e. to put one of the cards 4♥, 3♥, 2♣ in the Target. Applying the rules we see that this is effectively possible, and by reasoning backwards as before we obtain the following relations:



We have therefore decomposed the problem into its sub-goals. Of course, this is only one of the possible decomposition's that could be performed, and takes account only of the situation in the Target. As we shall see from experimental examination of the sequences, there are good grounds for considering it of particular importance.

If we combine these relations among sub-goals we obtain the graph given in Fig. 5.

Fig. 5 The graph of sub-goals.



Colorkeeper's legal moves _ _ _
 Numberkeeper's legal moves -----

The nodes of the graph represent the cards that can be in the Target position of the game board. The nodes adjacent to a given node indicate the cards that can be placed by Nk or by Ck on the Target to replace the card currently on it. For example, if 4♣ is in the Target area, the rules allow its exchange with 4♥, with 3♣ or with 2♣. If we concentrate only what happens on the Target during the game, we find that the sequence of cards follows the connections in figure 1.

Of course, the rules state that certain moves can only be made by Nk and others only by Ck. The graph is arranged so that all the horizontal lines represent permissible moves by the Numberkeeper *but not* by the Colorkeeper. Conversely, all the vertical and oblique lines represent moves that Ck is permitted to make but Nk is not.

It is possible to reconstruct the sequences of cards in the Target during a game. This can be done by following the paths in the graph which begin with the card that was on the Target at the beginning and finish with the card on the Target at the end. For example, if the initial card in the Target was 4♣ and the final card is 2♥, then we have the following paths on the graph:

Fig. 6 Paths and Hamiltonians in the sub-goals space

- 1 - 4♣ 4♥ 2♥
- 2 - 4♣ 2♣ 2♥
- 3 - 4♣ 3♣ 3♥ 2♥
- 4 - 4♣ 3♣ 2♣ 2♥
- 5 - 4♣ 3♣ 3♥ 4♥ 2♥
- 6 - 4♣ 3♣ 3♥ 4♥ 4♣ 2♣ 2♥
- 7 - 4♣

The graph therefore represents the space of the sub-goals that the players may - fully or partially - figure out in order to solve the game. Different paths represent different decompositions of the final goal into partial goals.

One notes immediately that the first 5 paths pass only once through a given node. Every longer path passes more than once through at least one node. In case 6 above, for instance, the path passes twice through node 4♣.

The situation that arises when 4♣ is on the Target is of particular interest for its complexity, and I shall therefore use it as an example. In Fig. 6 the paths in the graph which begin with 4♣ are ordered according to the (increasing) number of nodes that are passed through. Paths 1 and 2 are equivalent because they both comprise three nodes, while all the others contain a larger number of nodes. It is usually possible to find the minimal path (or

paths) connecting two nodes in the graph, and in our case it is indeed a trivial undertaking: they are the two paths $4\clubsuit 2\clubsuit 2\heartsuit$ and $4\clubsuit 4\heartsuit 2\heartsuit$, which are therefore optimal.

The final goal of putting $2\heartsuit$ into the Target can be achieved with one move if the card on the Target at the beginning of the game is $3\heartsuit$, $2\clubsuit$ or $4\heartsuit$. These are the cards adjacent to node $2\heartsuit$ in the graph. In all other cases intermediate goals must be accomplished before the final goal is achieved.

4 Conjectural strategies

The graph of the sub-goals introduced above can be used to formulate the following working hypothesis: the players do not learn and memorize a game strategy which enables them to behave optimally. Instead, they learn and memorize the graph of sub-goals, or a part of it, and then use it to devise their own game strategy. Since a strategy must enable a player to decide the move to be made at every stage of the game, the problem is whether there exists a concise way of representing the strategy, or whether the player must remember the entire tree of moves of which it is composed. Only if there exists a set of rules which are sufficiently short and simple to be memorized by the player can we assume that he is able to learn it and use it.

In *Transform The Target* this compact set of simple rules exists, and it derives from the conjectural decomposition of the problems space. The players generate a fraction of strategy at each stage of the game on the basis of whichever sub-goal is important at the time. I would stress that this is a *conjectural* search (Egidi 1991) by pointing out that the existence of a given path in the graph indicates that the problem *admits to* a solution, but that it does not *constructively* provide a solution. The players realize that the problem can be solved, i.e. that there is a procedure which if adopted by them will lead to the desired outcome, but they do not know what this procedure is.

For example, suppose that the card initially on the Target is $4\clubsuit$. In this case one solution is the sequence $4\clubsuit 2\clubsuit 2\heartsuit$, and it is reached as follows: Ck puts $2\clubsuit$ on the Target and Nk puts $2\heartsuit$ on the Target. $2\clubsuit$ and $2\heartsuit$ are the 'key-cards' in this sequence. If the players have them in their hands, they can immediately execute the sequence, otherwise they will have to search for the key-cards and then lay them on the Target. Because the two players can always find the two hidden cards, a solution certainly exists.

Accordingly, a solution to the problem exists, but we have not yet specified it in detail. We know that a sequence of key-cards can be put on the Target in order to achieve the goal. But we do not know exactly what sequence of moves Ck and Nk must perform. This manner of proceeding is entirely different from strategy identification in traditional game theory. In the latter case, in fact, to identify a strategy it is necessary to construct the game tree, and therefore to construct the sequences that lead to the desired outcome, but this is not necessary in the scheme analysed here.

A simple analogy shows that this is an issue with a general bearing on learning and discovery processes. Consider the discovery of theorems in a formalized theory. On the one hand, we have theorems of pure *existence*, which show that the problem is capable of solution. On the other, there are *constructive* theorems which indicate how to find the solution; i.e. ones derived from an algorithm (or procedure) which leads to the solution.

The experimental results discussed below show that this distinction is reflected in the behaviour of the two players, since for the majority of players it is much easier to explore the

space of the sub-problems and find the correct paths than it is to identify the procedures required to accomplish these paths.

5 Coordination and signals

Once both players have identified a path in the graph, i.e. a sequence of key-cards that must be put on the Target, the problem is that they cannot communicate to each other which path they want to follow. Each player can understand which path the other intends to pursue only by observing the moves that his partner makes. If there is only one optimal path, as happens for example when there is a hearts card on the Target (see Fig. 5), the interaction between the two players is very simple, and coordination can straightforwardly take place because each player has only one task to perform. If there are two optimal paths, as happens when the 3 or 4 of clubs is on the Target, the players may find coordination extremely difficult, given that the information that each of them obtains from watching the first moves of the other is normally not enough to deduce which path has been chosen.

In this case, the goals that the two players set themselves may be incompatible, and the game will require an adjustment of goals after a certain number of moves. This therefore raises a first kind of coordination problem: the coordination of goals.

In the case in which the two players' goals are compatible - that is, both of them have chosen to follow the same path in the graph - in order to implement their strategy they must acquire the key-card for that path. To do this, they must conduct a search that can proceed in many different ways, both because of the distribution of the cards and because of the different forms that cooperation between the two players can take. Here a second form of coordination takes place : the coordination of procedures.

In order to examine the features of the two kinds of coordination, we must first clarify how a player decides the move to make . As an example, assume that at the beginning of a hand the card on the target is 4♣. If players are able to signal to each other that they want to follow the path 4♣ 2♣ 2♥, their actions will be the following:

Ck looks for 2♣ and, if he has 2♥ in his hand, he reveals it to Nk. When he finds 2♣, he keeps it in his hand and puts it in the target.

If Nk has 2♣ in his hand, he reveals it to Ck. He then looks for 2♥ and, when he finds it, he keeps it in his hand, and waits until Ck puts 2♣ in the target . He finally puts 2♥ in the target.

Of course the problem is how the players can signal to each other the sub-goal they wish to pursue. Before discussing this crucial point, let us generalize the example above. If we extend the analysis to the other minimal paths in the graph, we obtain the following table

Fig. 7

Conjectural Card on Target Strategy at the beginning	NumberKeeper's sub-strategy	ColorKeeper's sub-strategy	
1 3♥	Reveal 2♥	Seek 2♥	
2 4♥	Reveal 2♥	Seek 2♥	
3 2♣	Seek 2♥	Reveal 2♥	
4 3♣	Seek 3♥	Reveal 3♥	Goto 1
5 3♣	Reveal 2♣	Seek 2♣	Goto 3
6 4♣	Seek 4♥	Reveal 4♥	Goto 2
7 4♣	Reveal 2♣	Seek 2♣	Goto 3

Seek X means: look for card X and put it into the Target area.

Reveal X means: put card X on Up so that your partner can use it.

We have thus identified the strategies that the players must adopt on the basis of minimal paths for each of the possible initial situations of the Target. Fig. 7 requires some comment. The original task (putting the 2♥ into the Target) has been divided into independent tasks each of which must be completed by a different player.

To understand better how players may coordinate themselves, it is convenient to classify the hands into two orders of difficulty, in relation to the different "distances" between the card on the target at the beginning of the hand and the final position (2♥ on the Target). This distance is easily measured in terms of the number of sub-goals which compose the problem to be solved: looking at the graph of sub-goals (Fig. 5) and counting the number of branches which connect the 2♥ with the card in the Target at the beginning, gives us the number of the sub-goals.

In cases 1, 2, 3 in Fig. 7 there is only one branch connecting 2♥ with the card in the Target. I classify this kind of hand as being at a "low" level of difficulty; in cases 4,5,6,7 (Fig. 7) there are two branches connecting 2♥ with the card in the Target, and I will consider hands of this kind as being at a "high" level of difficulty.

There are good reasons for this classification. In the case of configurations with a low level of difficulty, in fact, in order to achieve their goal the players must perform actions which require a very elementary kind of goal coordination, one which does *not* require signalling. One of the players - the action leader - must search for the key-card and place it in the Target. The role of the other player is simply to reveal the key-card if he has it in his hand. Consequently, if the players have a clear idea of their elementary goal, they can accomplish it without any further information. In the opposite case of hands with a high level of difficulty, the players must choose between two different paths (at least), and therefore each player must understand the sub-task chosen by his partner in order to move consistently with him.

To coordinate their tasks, the players must use the information arising from the board regarding the status of the cards and eventually the moves which have been made. Let us see how the board status and the moves can be considered as a signalling system which enables the players to coordinate their goals.

Note that in hands with a high level of difficulty (cases 4, 5, 6, 7 in Fig. 7) the sequences of actions to perform have been divided into two stages according to the two sub-goals pursued. For example, suppose that the card initially in the Target is 4♥ and that the players

adopt strategy n. 7. This strategy consists of a first phase in which Nk must search for 4♥. When Nk has placed 4♥ in the Target, Ck must search for 2♥ (the second phase of the strategy) in order to put it in the Target.

In each of these two phases there is a "Leader" player who must *search* for a key-card, and a "follower" who must *reveal* to him that card if he has it. The follower may even help the leader in his search, and if this happens the problem arises of coordinating the search for the key-card. Involved here is *procedural coordination*, which only appears when it is clear to both players which key-card they must search for. As we shall see below, the helping behaviour of the follower is usually inefficient: the more effective course of action, in fact, is one in which the follower only *reveals* the key-card if he has it in his hand, instead of trying to help the leader in his search.

To return to our example, if during the first phase the follower does not have the leader's key-card (4♥) in his hand, instead of waiting until the leader finds 4♥ and puts it in the Target he can immediately start searching for his own key-card, the 2 of Hearts. In this way the second phase is anticipated and performed in parallel with the first, since the two players search *simultaneously* for their key-card: the Colorkeeper for 2♥ and the Numberkeeper for 4♥.

This obviously raises a serious problem of *sub-goal coordination*: how, in fact, can the two players know the card that the other is searching for? Since they cannot signal their intentions, they can only act according to the information yielded by the board, that is, on the moves that have been made. The question is this: is the information rich enough to allow players to understand unambiguously each other's intention and therefore to coordinate their action?

I will discuss this crucial point in paragraph 8, while in the following paragraph I will go deeper into the distinction between procedural coordination and sub-task coordination.

6 Conjectural strategies and procedural coordination.

Let us return to the analysis of the nature of the strategies listed in Fig. 7. Each strategy can be realized by executing a set of rules, like "Search for 2♥" or "Reveal 4♣".

I shall call 'conjectural strategy' the set of rules which, if followed by the player, lead to completion of the task.

Note that these are not strategies in the traditional sense of the term until the *generic* terms 'Seek' and 'Reveal' are transformed into precise *procedures*. At this stage, we can be sure that the game can be successfully played because we know that the players can always successfully conduct their search for the key-cards, even if we do not know the procedures which allow the effective sequence of moves to be made. This is why I have used the term 'conjectural': a conjecture must be made in order to assume that the procedures which solve the elementary problems (e.g. 'Seek' 2♥) exist and are viable.

To take the final step and transform conjectural strategies into a true procedure (i.e. with a set of rules which state unequivocally the moves that the players must make at every moment of the game), the procedures corresponding to 'Seek' and 'Reveal' must be fully specified. This can be done in many different ways as regards the 'Seek' procedure, depending on how the two players decide to coordinate their actions during the search.

It is obvious that every conjectural strategy identifies a single path. But the relation does not hold in reverse: to each path, in fact, there usually corresponds several strategies.

Let us assume that, in order to implement a strategy, the two players decide to acquire the key-cards of that particular path. To do this, they must conduct a search which can take different forms - both because of the arrangement of the cards and because of the different ways in which cooperation between the two players can take place.

The problem is therefore to discover the relations between each path and the possible solutions it leads to; that is, the strategies which, if adopted by the two players, enable them to follow the desired path.

The difference between the path and the strategy that realizes it can also be viewed in procedural terms. In fact, the existence of a specific path in the graph shows that the problem admits a solution, but it does not constructively provide one. The players realize that the problem can be solved, i.e. that there is a procedure which if adopted by them will lead to the desired outcome, but they do not know the procedure(s).

On the basis of this observation we can distinguish between substantive coordination and procedural coordination. Substantive coordination is coordination which derives from choosing consistent paths. Procedural coordination derives from jointly devising or discovering a procedure with which plans can be implemented .

Obviously the substantive coordination problem takes very different forms depending on the card in Target. In the first three cases in the above table 7, each player has only one possible strategy (optimal or satisficing) available. Hence the substantive coordination can take place as long as each of the two players is aware of the (only) strategy available to the other. The most interesting cases are those in which 3♣ or 4♣ is in the Target area, because the two players have two entirely different strategies available to them. This raises the problem of understanding which of the two strategies the other player intends to use, and therefore of coordinating their actions .

The problem of procedural coordination arises instead in all cases, whatever the card in the Target area may be . It admits to different solutions depending on how the two players intend to interact. In fact, what do "Reveal" and "Seek" mean?

'Reveal' may assume slightly different meanings. It is given a 'minimalist' interpretation when a player has a card in his hand which the other is looking for, and which he must reveal by putting in position U (if he does not have the card in his hand he 'passes', and the other player must perform the search). A stronger interpretation is when a player realizes what card his partner is looking for and seeks it himself in order to reveal the card to him when he has found it. This stronger version entails a procedural coordination problem in the search. The two players may hinder each other if they are both simultaneously seeking the same card. The same holds for the meaning of "Seek" : players can adopt a procedure (UUPT) which involves joint information or a procedure (PCPN) which allows one player only to gather information.⁴

Note that if the players move into the graph of goals while they search for a path connecting the goal (2♥) with the card in the Target (4♣ for example) , if they discover the two minimal paths 4♣4♥2♥ and 4♣2♣2♥, they find the two optimal solutions not as a consequence of an exhaustive search in the graph, but because these two solutions are "easier" to discover. The discovery of the minimal path is therefore not the consequence of choice optimally performed, as the traditional rational choice model suggests: as we will see below, players do not try to make an exhaustive search in the space of the possible procedures before choosing: more "practically" they identify the minimum path simply by making the minimum effort during their exploration. In this kind of game, in fact, the shortest path between the goal and the card in the target is also the easiest to discover.

⁴ See Cohen and Bacdayan (1991).

7 The emergence of meta-routines.

The main point I want to clarify at this stage of the analysis is whether players, after an initial period of learning, behave in a routinized way. We need therefore to define what the routinized behaviour is expected to be in this specific context, and then experimentally check if it emerges from players' actions.

A first possible routinized behaviour would consist in a fixed choice of one only among the different possible paths connecting two nodes in the graph of sub-goals. If players discover only a part of the graph of subgoals, they can be tempted to use only this subset . More specifically, we must check if players discover one only of the two alternative conjectural strategies which can be chosen in order to put 2♥ on the target when the hand is of a high level of difficulty (3♣ or 4♣ on target).

If we observe only the sequence of cards which appear on the target during a hand, we see sequences 4♣4♥2♥ and 4♣2♣2♥, (the two minimal paths) , or more complex ones , like 4♣3♣3♥2♥ . To be brief, hereafter I shall call respectively 442 and 422 the two minimal paths 4♣4♥2♥ and 4♣2♣2♥. Analogously , I shall use 332 and 322 instead of 3♣3♥2♥ and 3♣2♣2♥. (Fig. 8)

Fig. 8

4♣4♥2♥=>442 4♣2♣2♥=>422 3♣3♥2♥=>332 3♣2♣2♥=>322
4♣3♣3♥2♥=>4332 4♣4♥3♥2♥=>4432

Suppose that the card initially on the Target is 4♣. In this case the two minimal paths are 422 and 442, and the key-cards are 4♥, 2♣, 2♥. Note that, whatever path is followed by the players to reach the final goal, only Colorkeeper can make use of 2♣ and only Numberkeeper can use 4♥, while 2♥ can be utilized by both the players. The same holds if 3♣ is initially on the Target. Therefore it is convenient to call 4♥ or 3♥ "the key-card of Numberkeeper" , to call 2♣ "the key-card of Colorkeeper" , and finally to call 2♥ the "double key-card".⁵

Assume that during the first hands the players progressively learn to identify some of the different sub-goals and the elementary relations which link them together. In other words, assume that they discover a part (at least) of the sub-goals space represented in Fig. 5. Suppose for example that they become progressively aware of one of the paths connecting 4♣ with 2♥ , the 442 path. We can say that they behave in a routinized way if, after discovery of path 442, they decide to stick to this solution for every hand starting with 4 clubs on the target. (the same holds for hands starting with 3 of clubs and path 332).

If the great majority of players behaved as if they were routinized , we might observe a large number of couples following either the 442 or the 422 path, and a small number of players choosing the two optimal paths indifferently.

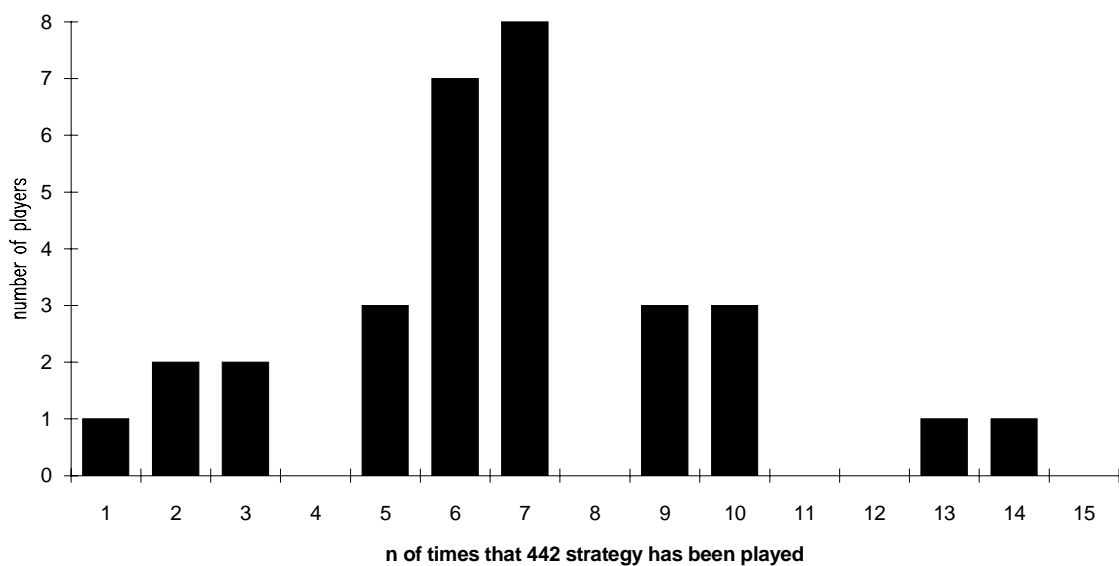
Fig. 9 sets out the empirical results. On the horizontal axis we have the times that the path 442 has been chosen (the total number of hands where either 442 or 422 strategy can be

⁵ I am indebted to Ricardo Pereira for this useful distinction.

applied is 15). On the vertical axis we have the number of couples who have played a given mixed path. It is clear that the players who followed one path only are a minority (they are represented on the two extreme sides of the statistical distribution) while the great majority adopt the 442 path about 50% of times, and 422 the remaining 50%.

The conclusion is striking: most of the players do *not* react in a routinized way to the situations that arise when the 3 or 4 of clubs is on the Target, in the sense that they do not use the same set of rules, to reach the goal whatever the initial configuration of card is on the board: they do not use only one of the two alternative conjectural strategies that can be activated; they use both, depending on the further information arising from the board. There is a normal distribution around a central type which is characterized by a choice for about 50% of times of path 442 and for 50% of path 422. This is a very flexible behaviour in term of goals achievement. To adopt this kind of behaviour, individuals must understand how the space of the sub-problems is structured. At the two extremes of the statistical distribution lie extremely routinized behaviours, in the sense of the rigid choice of one path only (either 442 or 422). Players who adopt this behaviour are subject to a *high* number of procedural errors , a greater number of errors compared with "flexible" players .The same occurs for 332.

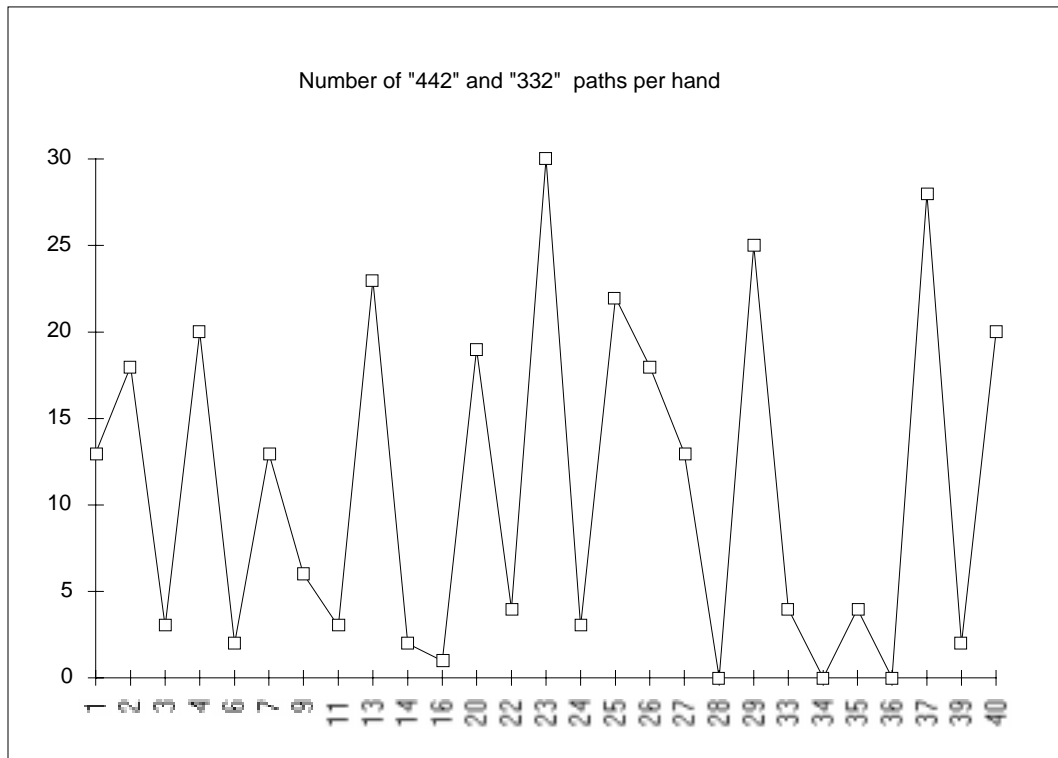
Fig. 9 - The 442 strategy's distribution



This phenomenon can be explained by the difficulty which arises from the attempt to apply the same set of rules (442,332 or 422,322) to every kind of cognitive and informative situation: in fact, not all of the initial card distributions give rise to the same amount of information for the players, and they therefore represent very different cognitive situations . This cognitive difficulty is clearly illustrated by the great difference in the number of players who choose 442,332 (or the complementary 422,322) strategy across hands. This number changes radically, as Fig. 10 and 11 show.

Fig. 10		Path		Path	
	Hand N	442 and 332	422 and 322		Tot
1	1	13	13		26
2	2	18	13		31
3	3	3	29		32
4	4	20	11		31
5	6	2	29		31
6	7	13	17		30
7	9	6	26		32
8	11	3	29		32
9	13	23	9		32
10	14	2	30		32
11	16	1	31		32
12	20	19	13		32
13	22	4	28		32
14	23	30	2		32
15	24	3	29		32
16	25	22	10		32
17	26	18	12		30
18	27	13	19		32
19	28	0	31		31
20	29	25	6		31
21	33	4	28		32
22	34	0	32		32
23	35	4	28		32
24	36	0	32		32
25	37	28	4		32
26	39	2	30		32
27	40	20	12		32

Fig. 11



In order to explain the variety of behaviours exhibited across hands, we must understand how the information arising from the game can determine the moves of the players, and more generally the choice of a path in the sub-goals space.

For simplicity I define as *static information* that which derives from the board status, and as *dynamic information* that which derives from the partner's (present or possibly past) moves.

Looking at the starting configurations of the 40 hands, we can select the ones in which one only key-card appears. In this way we have "pure" configurations, which show us the reactions of the players more clearly. The 40 initial configurations of the tournament are reported on the Appendix.

These configurations can be classified as follows:

1. Colorkeeper has his key-card (2♣) in his hand : hands 3, 6, 28.
2. Colorkeeper's key-card (2♣) is in the Up position : hands 14, 33, 34.
3. Colorkeeper has the double key-card (2♥) in his hand : hands 4, 23, 29.
4. The double key-card (2♥) is in the Up position : hands 1, 26.
5. Colorkeeper has the Numberkeeper's key-card (4♥ or 3♥) in his hand: hand 35.

It is easy to verify (Fig. 10, Fig. 11) that within each of the four groups in cases 1, 2, 3, 5, the players behave in a very regular manner: the majority of the players choose the 422 path in cases 1, 2, 5 and the majority choose the 442 in case 3. Accordingly, in Fig. 11 to the cases 1,2,5 correspond a minimum, and to case 3 a maximum.

Take case 1 as an example. Here 95% of the players follow the 422 path (the number of players who choose the complementary path 442 corresponds to a minimum point in the graph of Fig. 11). Note that in this case corresponding to the 422 path is always the same first move of Ck (to put his key-card on the Target). In this "pure" case, the first move is closely

correlated with the path on the graph; this suggests that we should see whether this property is more general, i.e. if in the pure cases there is an univocal correspondence between the status of the board and the move chosen by the players. Fig. 12 below gives the experimental result of the 5 "pure" cases.

Fig. 12

- 1 If Ck has his key card (2♣) in hand then his goal is to put it the Target ; the path that players follow is 422 (95%).
- 2 If the key-card of Ck (2♣) is in the Up position then his move is to pick it up and the path that players follow is 422 (94%).
- 3 If Ck has the double key card (2♥) in his hand then his goal is to pass until Nk has put his key card (4♥) on the target . The path that the players follow is 442 (80%).
- 4 If the double key-card (2♥) is in the Up position then Ck's goal is either to pick it up (50% and the path followed is 442) or to search (50% and the path followed is 422) for his key card.
- 5 If Ck has the partner' key card (4♥) in his hand, then his goal is to search for his key-card 97% (and the path followed is 422).

On the basis of Fig. 12 we may try to associate either a move or a set of alternative moves for Ck with each of the 5 pure configurations (and obviously to do the same for Nk). Fig. 13 below sets out the mapping configurations-moves based on the empirical data.

Fig. 13 - Mapping between the space of pure configurations and the space of moves (experimental data).

- 1 If Ck has his key card (2♣) in his hand then his move is T (put it in the Target).
- 2 If the key-card of Ck (2♣) is in the Up position then his move is U (Ck picks it up).
- 3 If Ck has the double key card (2♥) in his hand then his move is P.
- 4 If the double key-card (2♥) is in the Up position then the Ck's goal is either U (he picks it up) or C,N (search).
- 5 If Ck has the partner's key card (4♥) in hand, then his move is C,N (search).

Note that some of the above rules, like rule n 3 and 5, are very "naive". In particular, contrary to expectations, the players seem unable to recognize configurations like 5 in Fig. 13, where the obvious move to make is U (Ck reveals the card he has in hand to his partner) . The same cognitive difficulty arises with hands of low level (i.e. starting with 3, 4 of hearts or 2 of clubs on the Target) when one of the players must reveal to the other the key card he has in hand. As we will see below, a learning process arises, and very slowly the players became aware of the need to reveal the key-card of their partner .

If we look the remaining hands, where none of the starting configurations is "pure", it is evident that more than one of the rules in Fig. 13 can be applied: a possible conflict among rules emerges, and therefore players must explore the space of the combinations among rules and decide what move to make for every "mixed" combination.

In fact consider as an example a configuration in which Ck has the double key-card in his hand *and* his key card is in Up position (hands 2,7,27) : two rules (n 2 and 3) which prescribe two different moves (U and P respectively) are simultaneously activated. Fig. 14 below shows all the mixed configurations, and the corresponding hands.

Fig. 14 - The Space of mixed configurations

Ck has in Hand	There is in Up	Hands	Choice of path 442 (Fig. 11)
1 - His Key-card	The Double Key card	24,39	Min.
2 - "	The partner's key card	36	Min.
3 - The Double Key card	His Key card	2,7,27	about 50%.
4 - "	The partner's Key card	4,13,29	Max.
5 - The partner's Key card	The Double Key-card	20,25,37	Max (trend).
6 - "	His Key card	9,22	Max.

In order to solve the conflicts among the different prescriptions arising from the simultaneous application of two rules, and in order to allow players univocally to decide a move for every board configuration, the players may mentally explore the mixed configurations space. To cover exhaustively any possible configuration of key-card, their mental exploration would require great cognitive and memory effort.

As we will see, there is evidence that generally the players do *not* explore all the rules combinations *in advance*. On the contrary they realize the existence of conflicts and ambiguities only during the game, when they directly experiment with the mixed configurations. This behaviour is evidenced - as is clear from Fig. 14 - by the fact that the reactions of the players to the same configuration change over time in relation to the progress of the learning process.

In cases 4 and 5 of Fig. 14, in particular, the players modify their reaction to the configurations very slowly: they have noticeable cognitive difficulty in focusing their attention on the role of the partner key-card and therefore in realizing that it is convenient to "Reveal" to the partner his key card.

Summing up, there is a clear evidence that players map between the key configurations of the game and the moves in such a way as to decide precisely the move to make for every key-configuration. Up to this point, however, the mapping between configurations and moves is not perfectly univocal, because players modify their reactions to the key configurations (pure and mixed) very slowly. The question is whether this process converges or not on a set of rules which define univocally the move to make for every key-configuration. In order to answer this question we need a model of a boundedly rational player; an artificial player which can display all the possible "rational" solutions. This is the task of the next section.

8 An Artificial Player

How can we build "rational" strategies, and assume them as the cornerstone for evaluating the human strategies which emerge from the empirical data? For answer to this point we must build a model of procedurally rational action, generate rational strategies and compare them with the experimental data. As will become clear, not only is rationality involved in the construction of the model, but also expectations and rational expectations.

To build an artificial player, we simply use the set of rules of *Transform The Target* game and assume that the artificial player has been able to explore the space of the rules in such a way as to create the graph of goals. Therefore the artificial player is based only on the division of goals into sub-goals that we have already discussed in detail. His behaviour will be described by a set of "production rules", i.e. rules of the form "If <condition> then <action>", which are typical of problem solving and machine learning.

As before, I will discuss only the configurations which are "difficult" because they admit more than one solution on the graph of subgoals. If a player is supposed to be able to explore the sub-goals graph, he must be able to choose a path in the graph. We assume he is able to choose the minimal path, reminding the reader of the previous *caveat* concerning the difference between the choice of a minimal path and the traditional optimal choice of a strategy. When there are two minimal paths the problem is which of the two possible paths to choose. Is the information available from the cards disposition on the board sufficient to decide a "rational" move univocally?

Remember that *static rules* are those arising from application of the decomposition of goals to the information deriving only from the board status. The static rules impose a set of restrictions on the possible actions of the players.

A model based on these rules describes players as reacting only to information provided by the board: players try to apply the set of static rules on the basis of information available from the cards on the board, and if a rule matches the information, they will move accordingly. Fig.15 and 16 below shows the set of static rules which arise from the choice between the two possible minimal paths. They have been built simply by using the graph of sub-goals in the extended form (Fig. 7).

Fig. 15 - Static rules for Nk ($X\clubsuit$ in Target, $X = 3,4$).

- S-Rule1 - If I have my key-card ($X\heartsuit$) in hand then my goal is either to put it in the Target (with probability $p = AA[1]$ (442)) or to search for $2\heartsuit$ (422) with probability $= 1-AA[1]$.
- S-Rule2- If my key-card $X\heartsuit$ is on the Up position then my goal is either to pick up $X\heartsuit$ (with probability $p = AA[2]$ (442)) or to search for $2\heartsuit$ (422) with probability $= 1-AA[2]$.
- S-Rule3- If the double key-card $2\heartsuit$ is in my hand then my goal is either to wait until my partner puts his key card $2\clubsuit$ in the Target (with probability $p=AA[3]$ (422)) or to search for $X\heartsuit$ (442) with probability $= 1-AA[3]$.
- S-Rule4- If the double key-card $2\heartsuit$ is in the Up position then my goal is either to take it (with probability $p = AA[4]$ (422)) or to search for my key-card $X\heartsuit$ (442) with probability $p = 1-AA[4]$.

Fig. 16 - Static rules for Ck ($X\clubsuit$ in Target, $X=3,4$).

- S-Rule5- If I have my key-card $2\clubsuit$ in my hand then my goal is either to put it in the Target (with probability $p=AA[5]$ (422)) or to search for the double key-card $2\heartsuit$ (442) with probability= $1-AA[5]$.
- S-Rule 6- If my key-card $2\clubsuit$ is in the Up position then my goal is either to pick it up (with probability $p=AA[6]$ (422)) or to search for the double key-card $2\heartsuit$ (442) with probability= $1-AA[6]$.
- S-Rule7-If the double key-card $2\heartsuit$ is in my hand then my goal is either to Pass until my partner puts his key-card $X\heartsuit$ in the Target (with probability $p=AA[7]$ (422)) or to search for my key-card $2\clubsuit$ (442) with probability= $1-AA[7]$.
- S-Rule8- If the double key-card $2\heartsuit$ is in the Up position then my goal is either to pick it up and Pass until my partner has put $X\heartsuit$ on the Target, (with probability $p=AA[8]$ (422)) or to search for my key-card $2\clubsuit$ (442) with probability= $1-AA[8]$.

Note that the Nk's rules are perfectly equivalent to the Ck's ones.⁶

Of course these rules do *not* make a correspondence between pure configuration and moves: to any pure configuration there correspond a couple of possible moves, depending on the path that has been chosen on the sub-goals graph.

The above S-rules, if applied, do *not* uniquely determine a move; they simply reduce the number of possible alternative moves. If we want to build an artificial player which adopts only static rules, we must then further restrict the set of possible actions which stem from the application of a given rule. The restrictions must be applied in such a way that only one move is generated by a rule. One way of doing this is to set the probabilities AA[i] to 0 or 1 values only. Let us therefore assign to probabilities AA[i] only the values 0 or 1.

Three different situations can arise:

- 1 - one rule only matches the information provided by the board.
- 2 - more than one rule can be applied : a possible conflict among rules emerges.
- 3 - no rule matches the information on the board; information is not sufficient to decide the move to make.

In the first case a rule can be applied, and consequently a goal is univocally established.

The second case has a set of consequences which must be discussed in depth. I therefore anticipate discussion of the second case, which is simpler.

In this case, the set of static rules does not allow players to decide a move for a given card configuration. We cannot hope to solve this situation by adding a set of dynamic rules , for the simple reason that dynamic rules , being based on the partner's moves, can be applied only after the first move of any hand. Consequently, if none of the key cards is "visible" at the start of a hand, we cannot use either static or dynamic rules.

This situation is temporary : if we focus our attention , as usual, on the most "difficult" hands, which start with 3 or 4 clubs on Target, we realize that there are three key-cards, so that at least one of the two players must have a key-card in his control (either one of them has a key-card in his hand, or a key card is in position U and therefore both players can see it). Therefore only one of the players at most suffers from a temporary lack of information: but even if the situation is "temporary", the player does not have a rule for deciding his goal. To solve this point, I simply assume that players, in the absence of any information, will search for a key-card. By consequence I add a new rule (search) which covers any card situation which previously did not match the static information. This assumption is largely confirmed by the experimental data.

Now let us turn to the second case, when more than one rule can be simultaneously applied to the board. Here a conflict arises. The conflict must be resolved by providing a specific new rule which decides the priority among the conflicting rules. But how is it possible to select a new rule which allow players to move in a coordinated way? Note that if players behaved on the basis of the static rules only, they would not be able to coordinate their goals. The only way to solve the conflict is therefore to use additional information from the game: this means that , since the board configurations are insufficient for a coordinated set of actions, (rational) players must use the information arising from past configurations and their partner's moves.

⁶ For simplicity I have deliberately omitted the rule regarding the move that Ck must use when he has the Nk's key-card in hand (and viceversa) : it is obviously the move U (Reveal).

To clarify this point, consider the vector AA of probabilities distribution under the assumption that AA[i] can only be 0 or 1, depending on the goal that the players want to pursue.

AA[1]=1 means: If Nk has 4♥ in his hand then his goal is to put 4♥ in the Target.

His move is T and the expected path on graph is (442)

AA[2]=1 means: If 4♥ is in the Up position then Nk's goal is to pick up 4♥. His

move is therefore U and the expected path on the graph is (442)

...and so on.

Having attributed to AA[i] a value 0 or 1 for all the arguments from 1 to 8, we have fully defined a one-to-one mapping between rules and moves. This implicitly defines the features of coordination between the two players.

Distribution AA = (1100-0011) represents a 442 path and distribution AA = (0011-1100) represents a 422 expected path. Any different distribution implies a *non-coordinated* sequence of choices. In fact, for example, for the distribution (1100-1100) Nk's goal is 4♥, and Ck's goal is 2♣, which are incompatible.

As an example, Fig. 17 below shows the different behaviours which arise from different 0/1 probability distributions. Note that if the two players follow a couple of goals which are inconsistent, as in the (1100-1100) distribution, they may simultaneously be self-consistent: in fact, in our context, the self-consistency of a player consists in the ability to pursue a given goal in a coherent way. In the (1111-0000) distribution, the two players are inconsistent with themselves. Therefore two possible kinds of behaviour arise: the player's consistency with himself and with his partner. It is important to note that is possible to identify the player's consistency by analyzing the game's sequences, as shown in the table below.

Fig. 17

Hand n 11 2♣ 3♥ 4♥ 4♣ 3♣ 2♥

Nk's Goal	Ck's Goal	Play	Paths	Nk's rules				Ck's rules			
		Consistent									
4♥	2♥	CUUTT	442	1	2	3	4	5	6	7	8
		NUUTT	442	1	1	0	0	0	0	1	1
		Consistent									
2♥	2♣	CPTT	422	0	0	1	1	1	1	0	0
		NPCPTT	422	0	0	1	1	1	1	0	0
		Inconsistent									
2♣	4♥	NUCTUPT	442	1	1	0	0	1	1	0	0
		CUTUPT	422	1	1	0	0	1	1	0	0
2♥	2♥	NPCPPP..		1	1	0	0	1	1	0	0
		CPNPCP..		1	1	0	0	1	1	0	0

I have shown before that the great majority of couples do not follow the same fixed path for all different hands. We therefore can identify only a very small number of couples as characterized by a given, fixed distribution of 0-1 probabilities.

This means that the couples which follow - for example - the 442 path in hand n. 11, and are characterized by the sequence of static rules 1100-0011, will not follow the same sequence for all remaining hands. Players (slowly) learn to coordinate their sub-goals, i.e. to choose either AA = (1100-0011) or AA = (0011-1100) across hands, in relation to the positions of the key-cards on the board. In consequence, the path to follow on the goals graph is not fixed in advance but depends on the initial distribution of cards.

The problem now is to establish whether information from the board is sufficient to allow either player to understand his partner's intention, and consequently to choose his moves in a coordinated and consistent way. In other words, we must see whether there are signals sufficient for both players to set the distribution of probabilities either to AA = (1100-0011) or to AA = (0011-1100) in a univocal way, avoiding the inconsistent distributions.

Note that when a configuration on the board is mixed, i.e. when more than one key-card appears simultaneously on the board, more than one rule matches the configuration. Here we are in the situation summarized by Fig. 14; readers can verify that these situations do not restrict the choice in such a way as to permit players to make univocally one move. By consequence in both cases, pure and mixed configurations, players need additional information in order to coordinate their goals.

The players decide the most convenient path in relation to the available information: since the information provided by the cards distribution on board is not sufficient to coordinate their goals, each player must also use information about his partner's past moves.

As consequence, we must take the dynamic rules into account. These rules allow players to coordinate their goals as they play the hand. Fig. 18 lists a set of dynamic rules which simply map the key-configuration *plus* the past move *onto* the set of moves .

Fig. 18 - Dynamic rules concerning Nk's decisions.

- D-Rule 1 If Ck's key-card (2♣) was in Up position *and* Ck took it in the last move then Nk searches for the double key card (2♥).
- D-Rule 2 If Ck's key card (2♣) was in the Up position and Ck did *not* take it (either Ck looks for the double key card or keeps it in his hand) then Nk searches for his key card (4♥).
- D-Rule 3 If double key card (2♥) was in the Up position and Ck did take it in the last move, then Nk searches for his key card (4♥).
- D-Rule 4 If the double key card (2♥) is in Up position and Ck did put it the last move, or did not take it the last move, (Ck probably is looking for his key-card or keeps it in his hand) then Nk takes it.

Fig. 18 sets out the Nk's dynamic rules. Likewise, we can build the dynamic rules concerning Ck's decisions, by simply exchanging Nk with Ck and Nk's key-card with Ck's key-card in the figure.

It is very important to note that by memorizing the past moves and recursively using the dynamic rules, the artificial player can use all the available information, and therefore can be fully rational; moreover, if the artificial player attributes to his partner the ability to use all available information, he would interpret the partner's moves as rational: in this case, the artificial player would have *rational expectations* .

Even if rational expectation behaviours can be checked in a limited number of hands, the evidence shows that this kind of behaviour is very rare. As is clear from the data that I set

out in the next section, the great majority of players do not use all available information, and they do not attribute this kind of ability to their partner .

9 Some evidence on learning

The initial card distributions in the first 5 hands have been replicated , in the same order, in the 5 hands after the 25th of the tournament (26 to 30). We consequently have a clear idea of the effects of the learning process, which are given in the Fig. 13 below. The first left column lists the sequences of moves played by the artificial player using static and dynamic rules. The artificial player uses all available information but does *not* assume rational expectations: this means that he decides the move to make by looking at the configurations, plus the move made by his partner, without pretending to interpret the partner' move as rational.

Fig. 19 The Learning Effect after 25 Hands

Number of Couples		32	
Artificial Hands	GOALS	Experimental Data	
		Hand 1	Hand 26
4H 3C 2H 4C 2C 3H			
UCPTT	442	1	7
NUTT	422	2	1
CUNPTT	422	3	6
UNPCPTT	442	4	8
		10	22
	%	31,25	68,75
		Hand 2	Hand 27
3H 2H 2C 4C 4H 3C			
PCPNPTT	442	2	2
UUTT	422	7	15
PNPTT	442	7	9
		16	26
	%	50	81,25
		Hand 3	Hand 28
3H 2C 4C 3C 2H 4H			
TNPT	322	2	12
TCPNPT	322	1	6
NCPNPNTT	332	0	0
CNNCCNPTT	332	0	0
NNPCPTT	332	1	0
CCNNPTT	332	0	0
CNNPCPTT	322	0	0
		4	18
	%	12,5	56,25
		Hand 4	Hand 29
2C 2H 4H 3C 3H 4C			
PNPTT	332	7	12
PCPNPTT	332	4	4
NNCPTT	322	0	0
NNCCTCPT	322	0	0
CNTCPT	322	0	0
NCCNNPTT	322	0	0
PNPCPCPTT	332	0	0
CCTT	322	0	0
		11	16
	%	34,375	50
		Hand 5	Hand 30
3C 3H 4C 4H 2C 2H			
NUUPT	42	3	5
CUUPT	42	6	17
		9	22
	%	28,125	68,75

The Table requires very little comment: in fact the improvement in the players' ability is very clear.

Comparing the game conduct of the same couple faced with the same distribution (hands 1 and 26, 2 and 27, and so on), we have further confirmation that the couples who choose a goal path for a given cards distribution at hand X ($X=1,..5$) do not maintain the same path after 25 hands, i.e. they do not routinize a specific path on the goals graph.

An issue that must be raised at this point, even though I cannot discuss it in depth, is this: how do the static and dynamic rules of the artificial player emerge? A tentative answer is the following: at the beginning of the tournament players create a set of rules by associating a move to every pure key-configuration. They do this in a very naïf way, without fully exploring the space of the subgoals. Then slowly they are forced by the emergence of mixed configurations to compare the rules of thumb and to modify them. Static and dynamic rules are thus gradually generated by comparing the rules of the game against the sub-goals. The evidence from the experiment confirms this description too, although more experimental data and analysis is required to give a careful description of this learning process.

A second aspect of behaviour which concerns procedure formation must be emphasized. If for a large number of hands the card on the Target is the same, the paths that the player must consider to solve the problem do not change for a period of time. This persistence allows players to routinize, i.e. to react to the configurations by "automatically" deciding the move to make; they follow the sets of static and dynamic rules in very precise and quick way. If the sequence of hands which starts with the same card on the target is suddenly interrupted by a simpler distribution (for example, when 2 Clubs occurs in the Target) the persistence of routinized behaviour prevents the players from reacting in a correct way, and is therefore the cause of numerous errors. This aspect is clearly evidenced by the reduced number of errors and deviations from satisficing behaviour recorded when hands are homogeneous in terms of the Target, and by the high increase of errors when suddenly a different, even if very easy, initial distribution appears on the board.

10 Final Remarks

This search relies on the ability to learn. I have attempted to use the experimental data which emerge from Cohen's game replication to verify some of the current assumptions about routines emergence and procedural rationality.

Much has been written over the last twenty years about the experimental study and analytical modelling of learning processes. I can only remember three main lines of enquiry: the first began with the pioneering research by March and Simon (1958) and Newell and Simon (1976), and analyses the learning process as a search in problem space; the second, which originates from McCulloch and Pitts (1943), D.O. Hebb (1949) and von Hayek's (1952)⁷ ideas, develops models based on neural networks; and the third originates with the invention of genetic algorithms by John Holland and subsequent work on classifier systems.

These approaches have been very successful in machine learning area, but have rarely been compared with the empirical results on human learning by means of laboratory studies. Even though I have focused on the approach based on the problem decomposition, it becomes clear, after analysis of the experimental data, that there are connections with the approaches based on "production rules", which I wish now to discuss briefly.

⁷ De Vries (1993) writes: "Hayek's *The Sensory Order* is an intriguing book. Hayek was an outsider. Therefore he did not know D.O.Hebb's *Organization of Behaviour* (1949) until the final version of his own book was practically finished. Hebb's vision was in so many respects similar to his own that Hayek doubted for a while whether he should publish it. However, Hebb was more concerned with physiological details and less with general principles, which were Hayek's main interest."

To summarize some of the previous findings: the consequences of the assumptions of bounded and procedural rationality, if taken seriously, give rise to a model of human decision making which is quite different from the traditional picture of decision as unbounded rational choice: cognitive and memory limitations engender a process of search which is highly asymmetric and path dependent. Players move (conjecturally) in the space of sub-problems as they try to connect the local goals to each other in order to achieve their specific goals. Each of them focuses his attention and forms his beliefs differently. Thus the coordination of beliefs becomes a crucial factor in decision making, which requires cognition and learning.

Each strategy can be realized by executing a set of *production rules*, of the form "if Condition then Action", where action is, for example, "Search for 2♥" or "Reveal 4♣". I have called 'conjectural strategy' the set of rules which, if followed by the player, lead to completion of the task.

Note that these are not strategies in the traditional sense of the term until the *generic* terms 'Seek' and 'Reveal' are transformed into precise *procedures*.

Players discover very slowly a set of production rules, which allow to generate routines. They start from generating a set of naive rules of thumb which maps key-configurations of the game onto moves. The simultaneous matching of this rules, that happens when on the board a mixed configuration appears, give rise to a process of learning (adaptation) which allow players to substitute conflicting rules and provide new rules for the ambiguous signals. The convergence speedness to a set of stable rules depends on the "story", i.e. on the order of the sequence of boards (with different cognitive difficulties).

What is the relationship between these findings on learning at micro-behaviour level and the emergence of organizational routines?

I wish to emphasize that on a micro-behaviour level, the cognitive "atoms" are the elementary production rules (if Condition then Action) which by adaptation to the goals and sub-goals give rise to sequences of action procedurally rational. These sequences are the organizational routines, which are *not* memorized by players. Players do not need to keep all knowledge and information they need to play stored in memory: they only have to remember the cognitive "atoms" which allow to generate the organizational routines. This means that they are able to explore and "recreate" missing knowledge, as I have suggested at the beginning of the present paper.

We therefore must recognize that the notion of organizational routine must be considered as a synonymous with "not completely specified procedure". This assumption is confirmed by field studies of the behaviour of individuals in organizations which have evidenced the open and incomplete nature of routines. Incompleteness gives flexibility to the realization of routines and facilitates their change; a flexibility made possible precisely by the fact that agents are able to complete procedures by means of their ability to learn and to solve problems.

References

- Alchian A. A. (1951) , "Uncertainty, Evolution and Economic Theory", Journal of Political economy, vol. LVII, pp. 211-221.
- Allais M. (1953) "Le comportement de l'homme rationel devant le risque: Critique del postulats et axiomes de l'Ecole Américaine", Econometrica 21, pp. 503-46.
- Arrow K. J. (1971) "Economic Welfare and the Allocation of Resources for Invention", in Lamberton F. (editor), "Economics of Information and Knowledge", London: Penguin Books.
- Arrow K. J. (1974) The Limits of Organizations, New York: Norton.
- Arrow K. J. (1978), "The Future and the Present in Economic life", Economic Inquiry , n. 16, pp. 157-170.
- Axelrod R. (1984) The Evolution of Cooperation, New York: Basic Books.
- Cyert R. M., Simon H.A. e Trow D.B. (1956) "Observation of a business decision", Journal of Business n. 29, pp. 237-248.
- Cyert R. M. e March J. G. (1963) A Behavioural Theory of the Firm, Englewood Cliffs, N.J. : Prentice-Hall.
- Coase R. (1937) "The Nature of the Firm" in Economica, n. 4, pp. 386-405.
- Cohen M. D., March J. G. e Olsen J.P. (1972) "A Garbage Can Model of Organizational Choice", Administrative Science Quarterly, n. 17.
- Cohen M. D. (1991) "Individual learning and organizational routine: Emerging connections". Organization Science, 2 (1) pp. 135-139.
- Cohen M. D. e Bacdayan P. (1991) "Organizational Routines Are Stored as Procedural Memory: Evidence Form a Laboratory Study", University of Michigan: mimeo.
- Cutland N. J. (1988) Computability - An Introduction to recursive function theory, Cambridge University Press: Cambridge
- De Vries R.P. (1994) "The Place of Hayek's Theory of Mind and Perception in the History of Philosophy and Psychology" in Birner J. and Van Zijp (editors) Hayek, Co-ordination and Evolution, Routledge: London.
- Dosi G. ed Egidì M. (1991) "Substantive e Procedural Uncertainty. An exploration of Economic Behaviours in Complex and Changing Environments" in Journal of Evolutionary Economics , n.1, pp. 145-168
- Egidì M. (1992) "Organizational Learning and the Division of Labour", in Egidì M, Marris R. (editors) Economics, Bounded Rationality and the Cognitive Revolution, Aldershot: Edward Elgar Pub. Lim.
- Friedman M. (1953) Essays in positive Economics. Chicago: Chicago University Press.
- Goldberg D. E. (1989) Genetic Algorithms in Search, Optimization, and Machine Learning Reading (Mass), Addison-Wesley
- von Hayek F. A. (1945) "The Use of Knowledge in Society", in American Economic Review, 35, n.4, pp. 519-530.
- von Hayek F. A. (1952) The Sensory Order , London: Routledge & Kegan Paul.
- von Hayek F. A. (1980) "Economics and Knowledge" in Individualism and Economic Order (Reprint of 1948 first printing) Chicago: The University of Chicago Press.
- Hebb D. O. (1949) The Organization of Behaviour, London: Wiley.
- Hey J. D. (1991) Experiments in Economics, Oxford: Basil Blackwell.
- 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., Thagard P.R., (1988) Induction - Processes of Inference, Learning, and Discovery , Cambridge (Mass) : MIT Press
- Holyoak K. J. (1990) "Problem Solving", in Osherson N. , Smith E.E. (editors) Thinking, vol.3 ,Cambridge: The MIT Press.
- Kahneman D., Tversky A. (1986) "Rational choice and the Framing of Decisions", in Hogart R. M. , Reder M. W. Rational choice - The Contrast between Economics and Psychology , Chicago, The University of Chicago Press.
- Laird J.E (1984) Soar User's Manual: Version 4.0, Palo Alto: Intelligent Systems Laboratory, Xerox Palo Alto Research Center.
- Laird J.E., Newell A., Rosebloom P.S. (1987) "Soar: An architecture for general intelligence", Artificial Intelligence 33, pp.1-64.
- Machina M. J. (1982) ""Expected Utility" Analysis without the independence axiom." Econometrica , 50 pp. 277-323.
- March J. G. (1988) Decisions and Organizations, Oxford: Basil Blackwell.
- March J. G. e Simon H. A. (1958) Organizations , New York: John Wiley.
- McGuire C.B., e Radner R., Decision and Organization, Amsterdam: North-Holland
- McLelland J. L., Rumelhart D. E., (1988) Parallel Distributed Processing - Explorations in the microstructure of Cognition , Cambridge (Mass): MIT Press
- Muth J.F.(1961), "Rational Expectations and the Theory of Price Movements", in Econometrica, vol. XXIX, N. 3, pp. 315-335.
- Nelson R. R. and Winter S. (1974) "Neoclassical vs. Evolutionary Theories of Economic Growth: Critique and Prospectus", The Economic Journal n. 4 (336) pp. 886-905.
- Nelson R. R. and Winter S. (1982), An Evolutionary Theory of Economic Change, Cambridge (Mass) :The Belknap Press of Harward University Press .
- Newell A. (1990) Unified Theories of Cognition, Cambridge (Mass): Harward University Press.
- Newell A. and Simon H.A. (1972) Human Problem Solving, Englewood Cliffs, N.J. : Prentice Hall.
- Nilsson, N. (1980) Principles of Artificial Intelligence , Palo Alto (Calif.) :Tioga.
- Pearl J. (1984) Heuristics, Reading, Massachusetts: Addison- Wesley
- Polanyi M. (1958) Personal Knowledge: Towards a Post-Critical Philosophy, London: Routledge and Kegan
- Simon H. A. (1957), Models of Man, New York: Wiley.
- Simon H. A. (1963) "Problem solving machines" in International Science and Technology, 3(36), pp. 48-62
- Simon H. A.(1969), The sciences of the Artificial (Carl Taylor Compton Lectures), Cambridge: M.I.T. Press
- Simon H. A. (1972a), "From Substantive to Procedural Rationality" in McGuire B., Radner R.
- Simon H. A. (1972b), "Theories of Bounded Rationality" in McGuire B. ,Radner R.
- Simon H. A. (1979), "Rational Decision Making in Business Organization", in American Economic Review, Vol LXIX, pp. 493-513
- Simon H.A and Newell A (1972), Human Problem Solving, Englewood Cliffs, N.J: Prentice-Hall.
- Singley M. K. and Anderson J.R. The transfer of Cognitive Skill, Cambridge, Massachusetts, Harward University Press.
- Squire L. R., (1987) Memory and Brain, New York: Oxford University Press.
- Williamson O. (1975) Markets and Hierarchies:Analysis and Antitrust Implications, New York: Free Press.

von Wiese L. (1929) "Die Konkurrenz, vorwiegend in soziologisch-systematischer Betrachtung" in Verhandlungen des 6. Deutschen Soziologentages.

Winter S. G. (1982), "An Essay on the Theory of Production", S. H. Hymans (editor), Economics and the World Around It, Ann Arbor: Michigan University Press.

Appendix

The forty initial configurations of the tournament

1	4H 3C 2H 4C 2C 3H	21	3H 2C 4C 4H 3C 2H
2	3H 2H 2C 4C 4H 3C	22	2H 4H 2C 4C 3C 3H
3	3H 2C 4C 3C 2H 4H	23	3C 2H 3H 4C 2C 4H
4	2C 2H 4H 3C 3H 4C	24	3H 2C 2H 4C 3C 4H
5	3C 3H 4C 4H 2C 2H	25	2C 4H 2H 4C 3C 3H
6	2H 2C 3C 4C 4H 3H	26	4H 3C 2H 4C 2C 3H
7	4H 2H 2C 3C 3H 4C	27	3H 2H 2C 4C 4H 3C
8	2H 3C 3H 4H 2C 4C	28	3H 2C 4C 3C 2H 4H
9	3H 4H 2C 4C 2H 3C	29	2C 2H 4H 3C 3H 4C
10	4C 3C 4H 3H 2C 2H	30	3C 3H 4C 4H 2C 2H
11	2C 3H 4H 4C 3C 2H	31	2C 3H 2H 4H 4C 3C
12	4H 3H 2H 2C 4C 3C	32	4H 3H 3C 2C 2H 4C
13	3H 2H 4H 4C 2C 3C	33	3H 4H 2C 3C 4C 2H
14	4C 3H 2C 3C 4H 2H	34	4H 4C 2C 3C 3H 2H
15	2H 4C 2C 3H 3C 4H	35	2C 3H 4C 3C 4H 2H
16	2C 4H 4C 3C 3H 2H	36	4C 2C 3H 3C 2H 4H
17	4C 3C 2C 4H 3H 2H	37	2C 4H 2H 4C 3C 3H
18	4C 3C 2H 2C 4H 3H	38	3C 2C 4C 4H 2H 3H
19	3C 4C 3H 4H 2C 2H	39	4C 2C 2H 3C 4H 3H
20	2C 4H 2H 4C 3H 3C	40	3C 3H 4H 4C 2C 2H