

Substantive and procedural uncertainty

An exploration of economic behaviours in changing environments

G. Dosi¹ and M. Egidi²

¹ University of Rome “La Sapienza”, Rome, Italy and Visiting Fellow, SPRU, University of Sussex, Brighton, UK

² University of Trento, Italy

Abstract. Different sources of uncertainty are analysed and a representation of decision-making in principle consistent with behavioural evidence is proposed. The endogenous emergence of “innovations”, in the forms of unexpected events and novel behaviours is also examined.

Key words: Innovation – Uncertainty – Problem-solving – Rationality

1. Introduction

In a very general sense, uncertainty in human behaviours stems from incompleteness of the knowledge necessary to forecast future events, undertake any one course of action and control its results. In order to analyse how economic agents behave under uncertainty, one obviously needs to understand how they can reduce their lack of knowledge. In turn, a central aspect of decision making under uncertainty concerns learning processes, involving also the ability of the economic agents of reducing their uncertainty by framing their choices.

Uncertainty may have two origins, i.e. 1) the lack of all the information which would be necessary to make decisions with certain outcomes, and 2) limitations on the computational and cognitive capabilities of the agents to pursue unambiguously their objectives, given the available information.

The former source of uncertainty comes from the *incompleteness of the information set*, and the latter from the inability of the agents to recognise and interpret the relevant information, even when available. In other words, from their *knowledge incompleteness* rather than information incompleteness.

Notably the traditional axiomatic approach to choice under uncertainty implicitly ignores the second of uncertainty: on the contrary, in this work we focus on this aspect of decision.

Of course, different environments and different decision problems entail also different forms of uncertainty of the two types. As a consequence, according to

the circumstances, the agents may be expected to show different kinds of behaviours and different decisional procedures leading to these behaviours.

The traditional point of departure of the axiomatic theory of choice under uncertainty assumes that agents can exhaustively represent all possible events. As known, the theory defines the choice problem in terms of (1) a set of states of the world (or events), (2) actions, and (3) consequences of actions (conditional on the occurrence of each event). The decision problem is then formulated as the choice of the course of action which maximises some sort of goal (or utility) of the decision-maker given his beliefs on the probability distributions of the events¹.

Of course, the beliefs may not be correct and the model – in the case of repeated choices – involves the updating of the subjective probabilities, generally through Bayesian learning.

With respect to such as axiomatic decision theory, there are two sides to the work that follows, a critical one and a propositive one.

On the critical side, we shall argue that in the usual treatment of decision under uncertainty as a straightforward maximisation problem, seemingly innocent assumptions, such as the stationarity of the “world” and the competence of the agents to process the information that the world delivers, constrain its interpretative power to particular classes of decision problems.

On the propositive side, we shall suggest an interpretation of decision processes which holds also (indeed, especially) when environmental conditions or limitations on the “rational” capabilities of the agents (which we shall define) make the usual decision-theoretic models an inadequate representation of economic behaviours under uncertainty.

In Section 2 we shall analyse different sources of uncertainty which affect decisions and behaviours. Some of them can be treated within a maximising decision-theoretic framework, but some cannot. In particular, those cases which cannot be treated relate to environmental complexity or non-stationarity, on the one hand, and to the nature of the problem-solving which particular classes of decisions require, on the other.

In analogy with H. Simon’s distinction between “substantive” and “procedural” rationality we shall introduce the notions of *substantive and procedural uncertainty*. The former is related to some lack of information about environmental events, while the latter concerns the competence gap in problem-solving. We shall argue that most cases of decisions involving procedural uncertainty cannot be dealt with on the grounds of the traditional decision theory. These cases include, we shall see, most things that have to do *lato sensu* with “innovation”.

As a further point, we must acknowledge that theories of “rational” behaviour in economics are sometimes given a different epistemological status, namely as part of a positive theory of economic environments, used in order to define the equilibrium conditions of some variable (prices, quantities, etc.), without, however, any assumption on the actual behaviour of individual agents. Thus, we must also discuss under what circumstances, “as ... if” assumptions on microeconomic behaviours are legitimate devices to analyse the state or dynamics of the economic system.

¹ The classical references are Arrow (1951, 1951 a, 1970, 1971), von Neumann and Morgenstern (1953), Savage (1954); for an overview of subsequent developments, Diamond and Rothschild (1978).

To what extent and under which circumstances can the market surrogate the (uncertainty-bound) fallibility of individual agents? We shall address this question in Section 3, concluding that the most circumstances actual (possibly, non-optimizing) behaviours do matter in terms of aggregate performance of the system.

Then, what do people actually do? How do they make their decisions, and which decisions do they make? In Section 4, we shall suggest a theoretical interpretation of behaviours in these circumstances. In particular, we shall analyse the relationship between decisions and problem-solving in environmental conditions where “uncertainty” stems indeed from the limitations intrinsic to the computational and recursive features of a “rational” decision process. We shall develop upon the notion of procedural rationality and make use of results of computation theory – especially those derived since the '30s by Turing, Markov and Godel in relation to the problem of formalisation of natural reasoning.

Our argument applies, first of all, to decisions and behaviours in environments where procedural and substantive uncertainty jointly appear.

In the framework that we propose, decisions are represented as specific problem-solving procedures in formal languages, which (1) are developed on the grounds of competences that logically pre-exist the acquisition of information from the environment; (2) involve particular rules for the search and selection of information; and, (3) can evolve, generating new languages and new behaviours.

Our interpretation seems quite consistent with the evidence on judgement and behaviour under risk and uncertainty from experimental psychology. Certainly, some of our predictions overlap with (and are distinguishable from) those of the “rational” decision-theoretic models. The simpler the environment or the decision problem, the more redundant our interpretation is (and the more parsimonious is the “rational choice” model). Conversely, the more complex the environment or the decision problem, the less adequate is the “rational” approach: with sufficient problem-solving complexity, the latter does not apply, either because it would be self-contradictory or because it would run against a finite computational limit of empirical agents.

In Section 5 we shall suggest some examples of empirical phenomena and theoretical fields to which our interpretation can be fruitfully applied. The theory of production, the theory of the firm and innovation belong to this category.

2. Nature of the uncertainty, information and choice procedures: substantive and procedural uncertainty

A very common metaphor of uncertainty, at least since the time of Bernoulli, is expressed in terms of flipping a coin, throwing dice, etc. The general applicability of the metaphor to other uncertain situations has been questioned, among others, by Knight, Keynes and Schackle. However, in somewhat more refined formulations, it underpins the general axiomatic model of choice under uncertainty. Flipping coins, gambling, weather forecasts or – in the economic domain – things like life insurance, etc. have some characteristics in common.

First, “events” are states of nature which can hardly be affected by the decision-maker (or, for that matter, other decision-makers). Second, the list of events is known (and preferably finite). Third, the process of attribution of consequences to actions and events is procedurally trivial (... if tomorrow it will rain my harvest will get wet ...; ... if I gamble on the red, and black comes out, I lose ...). Thus,

preferences, states of nature, actions, consequences are easily separable and procedurally can be easily mapped into each other. Uncertainty simply comes from lack of information about the occurrence of future events. Using Heiner's terminology we have here an information gap, but no competence gap of the agents (Heiner 1983, 1987). Call this case, *substantive uncertainty and procedural certainty*. An axiomatic theory of choice easily applies.

Consider now the case of a strategic interaction, say an n-person game. Suppose also that we (as analysts of that environment) know that all the usual assumptions "objectively" hold: the rules of the game are given, there is a pay-off matrix which depends on the states of nature and on the actions of the other players. Theoretically, the uncertainty facing each agent strongly resembles that of the former case, and in fact there is a strong theoretical resemblance between the ways one generally treats "games against nature" and "strategic games".

Let us define *weak substantive uncertainty* (analogous to "risk") as all those circumstances where uncertainty simply derives from lack of information about the occurrence of a particular event within a known list of events, in principle representable as a random drawing by "nature", with a certain known (or at least knowable) probability distribution.

Consider next the case whereby the "events" are not in any proper sense "states of nature" but are partly endogenous to the decision process of the agents, so that events are not independent from actions. Here the metaphor of uncertainty in terms of random natural variables (the weather, etc.) is inappropriate even as a theoretical approximation. We believe this to be a quite general case in economic affairs: externalities, increasing returns, most forms of technical change, path-dependent processes imply, to different degrees, endogeneity of events and *bi-directional interactions between actions, events and outcomes*².

In particular consider the case of non-stationary environments, meaning environments that change, either endogenously or through exogenous shocks, their structural characteristics (technology, tastes, commodities, etc.). Innovation is obviously a phenomenon which implies non-stationarity. In turn, non-stationarity implies that each agent may expect an unknown number of unknown events to occur in the future. Whether the agents interact in strategic fashion or perfect competition prevails, endogeneity of events and non-stationarity introduce a strategic dimension of a higher order in the decision process³.

We define *strong substantive uncertainty* as all those cases involving unknown events or the impossibility, even in principle, of defining the probability distributions of the events themselves.

Clearly, this latter notion of uncertainty closely resembles the earlier suggestions of Keynes, Knight and Schackle, and from a standard informational point, involves near ignorance.

Note also that the boundaries between "weak" and "strong" uncertainty are not so clear-cut. Einhorn and Hogarth fruitfully define a domain of ambiguity defined as "... an intermediate state between ignorance (no distributions of ex-

² An appealing discussion of some cases is in Shelling (1978). Arthur (1985) and Silverberg, Dosi, and Orsenigo (1988) formalise cases of environmental dynamics where this occurs; Silverberg (1988) discusses the general subject.

³ Note, incidentally, that even under these conditions, experimental evidence shows significant departures from "rationality": cf. Tversky and Kahnemann (1974, 1988), Hogarth and Reder (1988); on some results of experimental economics interpretable – in our view – in this perspective, see Plott (1982), V. Smith (1982), Hey (1982).

pected probabilities of events are ruled out) and risk (all distributions but one are ruled out). Thus, ambiguity results from the uncertainty associated with specifying which of a set of distributions is appropriate in a given situation” (Einhorn and Hogarth 1988, p. 45).

So far, we have discussed different categories of uncertainty from a purely informational point of view, which we have called *substantive* forms of uncertainty. Indeed, this has also been the perspective in which its economic analysis has been mostly undertaken, generally limited to “weak” uncertainty, or risk.

Relatedly, a quite diffused theoretical presumption in the economic analysis of information is that agents make the best possible use of the available information, which is generally taken to mean both that they utilize maximisation procedures and that these procedures are somewhat “naturally” associated with their normal cognitive competences. That is, there is no competence gap in their *information processing*. We shall call all the conditions in which this applies as *procedural certainty*.

Of course, one can easily identify a lot of empirical circumstances characterized by weak substantive uncertainty *cum* procedural certainty. From gambling to deciding whether to subscribe an insurance, one can list several choices under uncertainty, whose procedures are nonetheless relatively trivial and can be theoretically assumed to be generally known⁴.

However, even weak substantive uncertainty may well involve competence gaps, notwithstanding the availability of information which would theoretically allow a “rational” solution⁵.

Experimental psychology has provided robust evidence on such departures from “rational” choice procedures essentially related to (1) the ways the problem is posed (whether it is “transparent” or “opaque”); (2) outcomes that are delayed and not easily attributable to particular actions; (3) variability in the environment which degrades the reliability of feedbacks, especially when outcomes of low probability are involved; (4) lack of information about the outcomes of other possible actions (Tversky and Kahnemann 1988, pp. 90–91).

Similar properties of choice under uncertainty can be analysed at a theoretical level. Consider, for example, the case of a strategic interaction whereby we (the analysts) can represent the decision problem within the axiomatic theory of choice. Suppose also that we know that some Nash-equilibria exist (because we have proved some theorem on their existence). Can we expect that, in general, the agents will actually behave according to the choice-theoretic model? A “rational” agent who knows the rules of the game and the list of state of nature would, of course, exhaustively explore all the decision tree, form his beliefs on the probability distributions of the states of nature and choose the strategy accordingly (the

⁴ Note also that this case of procedural certainty is the only one where the analyst can, so to speak, interpret behaviours “working backward” (Arrow 1983, especially p. 23) from revealed actions to the agents beliefs and priors (on probability distributions, degrees of risk-aversion, etc.) and postulate a (conscious or “automatic”) maximising choice, given the information structure and some generic goal (e.g. “make as much money as possible”). Finally, note that only with procedural certainty, maximisation and “revealed” consistency of plans may be considered to be equivalent.

⁵ As an illustration, let us just recall the dramatic increase in computational complexity required to re-establish coherent plans after the re-opening of a spot market in an Arrow-Debreu-Hahn economy in inter-temporal equilibrium (Radner 1968). A fortiori the argument applies to our case.

von Neumann algorithm for a two-person game is the procedural representation of this process). However, the “rational” procedure may quickly run against the computational constraints of empirical agents.

Interestingly, computability may not only be a problem for empirical agents, but also for the theoretician: for example, Nash-equilibria may exist but may not be recursively realisable⁶.

However, the complexity of the decision task (for example, many-person games, non-linearities in the relationships between events, actions and consequences, etc.) may easily determine a competence gap, in addition to the information gap associated with substantive uncertainty, and in addition to computational limitations. Heiner (1988) defines such a competence gap in terms of the difficulties of the agents to map information into the “true” events (that is, as a positive probability of “wrongly” interpreting an environmental message).

More generally, we propose that in most decision problems, “choice” is nothing but the terminal act of a problem-solving activity, preceded by the formulation of the problem itself, the identification of the relevant information, the application of pre-existing competences or the development of new ones to the problem solution and, finally, the identification of alternative courses of action.

In the analysis of this process, we suggest, the focus of the decision task is not on choice as such, but, rather, on problem-solving procedures. Building on Tversky and Kahnemann (1984) and (1988), we call “problem framing” all those activities leading from the identification of the problem to the selection of possible actions. Relatedly, we call cognitive competences the abilities of the agents to “frame” a decision problem and derive a course of action.

With the exception of the simplest decision tasks, nothing allows us to assume that the agents are naturally endowed with cognitive competences sufficient to the identification of solutions of a particular problem, let alone the “optimal” ones. Neither, can one assume that such competences are uniformly distributed amongst the different agents.

We shall define all those circumstances whereby the solution of choice problems is constrained by the computational and cognitive capabilities of the agents as characterized by *procedural uncertainty*.

With both substantive and procedural uncertainty, the agents will certainly try to find “rational” procedures, in the sense of Simon (1957) and (1959), but “rational” may not mean more than robust and computationally efficient. Typically, one is likely to find relatively stable “rules”, and, indeed, Heiner has shown that whenever any competence gap exists, routinized behaviours are more efficient than optimizing procedures (Heiner 1983, 1988).

Are these “rules”, whatever they are, approximations to an optimising behaviour? *Prima facie* the answer is that this will depend on the nature of the environment and the decision problem. Winter (1988) quotes Schumpeter arguing that this is likely to occur only when “things have had time to hammer logic into men”. Certainly, necessary conditions are environmental stationarity, persistence of near-equilibrium conditions and existence of a computable optimal procedure. In turn, these conditions are a sub-set of those which arise under weak substantive uncertainty (risk).

⁶ For a thorough formal analysis of the choice functions which are effectively computable from a theoretical point of view, see A. Lewis (1985, 1986).

More so, strong substantive uncertainty will always be associated with procedural uncertainty. In non-stationary environments, agents are always bound to try to understand – via procedures that almost by definition cannot be derived from the information delivered by the markets – future behaviours of other agents or future events that had never occurred in the past⁷. Those same capabilities which allow some agents to generate unforeseeable changes (e.g., developing a new product, opening a new market, etc.) are endogenous sources of uncertainty for the other agents.

Note that non-stationary environments may not be generally expected to satisfy the completeness-of-events description, and moreover can give rise to *computational failures* of the agents.

Under these circumstances, we suggest, the usual axiomatic theory of choice is neither an approximation to empirical behaviours nor a legitimate theoretical “stylisation”. Not only will agents not behave as literal maximisers, but it will be impossible to construct a theoretical model based on maximising agents, since the theory should construct a model where events are conditional on unknown (unpredictable) behaviours, and, of course, behaviours are also conditional on expectations on events. An infinite regress. Indeed, the very notion of “optimality” becomes an ambiguous theoretical notion.

In these cases, substantive and procedural sources of uncertainty are obviously related. Very little of the “mechanical” nature of the decision process is left.

A typical maximisation problem under uncertainty in standard decision theory is a well structured problem with a well defined solution and a known algorithm leading to it. On the contrary, in the situations we are considering here, one often finds ill-structured problems, which – irrespective of the actual information available and the actual computational complexity – continuously require also the generation of “representations” (models) of the world and models for the analysis and solutions of problems. Expanding upon Simon (1957), we suggest that this is the general case underlying choice/action in most empirical economic environments. Strong procedural uncertainty implies that the main activities involved in the process leading to action are neither choice (in the sense of some kind of comparison among well-defined alternatives) nor updating of probability distributions of known random variables (via some sort of Bayesian procedure), but model-building and problem-solving.

Under these circumstances, a positive theory of decision and behaviour entails the analysis of problem-solving procedures and of the forms of knowledge underlying it. Before turning to these issues, however, let us consider to what extent, in such decision processes, markets can surrogate the agent fallibility.

3. Rational choice as an analytical device, or, can the market surrogate individual fallibility under uncertainty?

The empirical evidence on economic behaviours, it is often conceded, is messy and subject to various sorts of disturbances. Especially under uncertainty, people may

⁷ Possibly the extreme attempt of applying a “rational” choice theoretic model, while recognising the requirement of some sort of “special skills” in information processing and decision-making, is Stiglitz (1986), where maximising agents “learn to learn”. However, the “orthodox agent” faces the same problems of procedural uncertainty discussed here also at the second order information-processing, and so on: an infinite regress is unavoidable.

well make mistakes on their expected probability distributions and in their actions; indeed, they may not literally maximise (in the sense that they do not consciously go through the decision procedures outlined, say, in Arrow (1951 a)) but follow different kinds of – more or less automatic – “rules”. However, it is sometimes argued, the standard representation of agents as literal maximisers (and, perhaps, learners in the sense of Bayes’ theorem) still retains a useful theoretical content since it underpins the analysis of the equilibrium states which characterise the system (the obvious implication being that these states are attained irrespectively of the procedures through which empirical agents make their decisions). In the economic literature, one finds three major arguments in favour of this “as ... if” assumption, namely:

1. Behaviours may neither come close to literal maximisation in terms of decision procedures, nor even in terms of choices and actions that are undertaken, but markets perform as selection environments so that in the asymptotic state of the system only “maximisers” survive, whether they know to be maximisers or not (see in particular Alchian 1950 and Friedman 1953). This is clearly the strongest version of the “as ... if” hypothesis.
2. The actual attempts of empirical agents to make the best use of the available information are such that the maximising choice-theoretic models can be considered as approximations to empirical behaviours. However, with uncertainty, information is imperfect and possibly asymmetric, agents hold different beliefs and “priors”, and their knowledge of the future is limited. They can interpret correctly the signals that they receive (and try to do so), but they do not receive all the signals necessary to make those optimal choices that would have been taken under perfect information. The markets surrogate – to greater or lesser extents – the fallibility of individuals in so far as they allow the exchange of commodities (and claims on them) contingent on particular events. This is of course the theoretical approach of contemporary General Equilibrium Analysis (as in Arrow and Hahn 1971).
3. Empirical behaviours may well show “rules” with varying degrees of automaticity: thus, the actors may not be procedurally maximisers. However, their behaviour may still be fruitfully represented in the standard choice-theoretic way, in so far as the observed “rules” are those which actually lead to optimal actions (see Machlup 1946 and Friedman 1953) (Of course, point (3) may well overlap with point (2), and also coincide with point (1), whenever the world is in that particular “selection equilibrium”).

The “as ... if” assumptions have been critically discussed, in a perspective which we broadly share, in Nelson and Winter (1982) and Winter (1986 and 1988). Here, we shall simply recall some theoretical properties and results which will be useful for our subsequent argument. Of course, the strongest version of the “as ... if” argument is sub (1). In order to substantiate that claim one would ideally need a proof of global stability of the equilibrium model that one is talking about. Lacking that (as one does), one would need at least proofs of local stability of a sufficiently small number of equilibria under rather general conditions and some rather robust ideas on the nature of the adjustment process leading there. In general, none of these has been provided, neither with respect to the level of analysis which is proper to this conjecture, i.e., General Equilibrium, nor even in selection environments where agents do behave rationally and strategically (for a discussion, see Silverberg 1988). At the very least, one should show that the end

point(s) of an unspecified market process (the “as ... if” equilibria) are recursively realisable, which, note, is a much less demanding proof than global stability. Loosely speaking, it simply requires the existence of a finite logical process leading from individual preferences, etc. to the aggregate outcomes: however, exactly the opposite has been proved (Lewis 1986).

Take this argument the other way round: since the strongest version of the “as ... if” hypothesis does not hold, some characterisation of actual behaviours is required, because particular types of microeconomic decision processes and actions do affect the states which the system attains: in fact, a widening theoretical literature shows that equilibria are often path-dependent, behaviour-dependent and institution-dependent.

Consider now the heuristic status and the theoretical domain of applicability of the “weaker” “as ... if” assumption, sub point (2) and, to some extent, (3), above. That approach entails “uncertainty”, meaning the lack of part of the relevant information about the occurrence of future events; complete procedural rationality (the agents solve without particular difficulties the computational problems involved in the relevant choice/action task) and, finally, markets which surrogate fallibility by trading “conditional commodities”. Of course, the obvious boundary on its positive applicability is the existence of the complete set of future markets, which we all know are in fact very rare. However, this is not our point. Consider, indeed, the case when all future markets exist, contingent on the complete set of known events (which are random variables) and known commodities. Even in this case, the “surrogation” of individual fallibility via the market holds under the strict condition that no unpredictable change is allowed to occur in the future, and this must be also believed by the agents. (In this respect change may be a new commodity, a new technology, etc.) To illustrate it, suppose that, on the contrary, new commodities (or technologies) may appear. Then, we have two possibilities.

First, one may add from the start to the commodity space and the event space the “unknown” dimensions corresponding to the exogenous future “shocks”. In other words, it must be possible to (and indeed the agents must – consciously or “automatically”) represent a terminal economy, T , whose commodity space N' strictly includes the initial space N , containing $n'-n$ “present” commodities (here, meaning, simply, commodities that exist also at the present). However, this implies a “complete future” which exhausts all humanly conceivable possibilities. In turn, this takes away a lot of the “exogeneity” and randomness of future changes. In a sense, as Winter (1988) puts it quoting the Ecclesiastes, one must assume (both the theoretician and the agent) that “there is no new thing under the sun”.

Second, and alternatively, one may postulate that the economy is sequential, so that, in each period, new events, new technologies, new commodities are generated, through, say, a random draw from a quite wide and largely unknown set. Thus, starting from the initial economy, one would generate sequences of economies characterised by an increasing number of “dimensions” (in the commodity-space, the event-space, etc.), without a unique terminal economy. Thus, one would have to represent the future (the “terminal economy”) in a space with infinite dimensions⁸. Though this might not be a problem from a mathematical

⁸ A “sequential” approach to economic analysis has recently produced interesting insights into economic dynamics characterised by unpredictable technological change (see Amendola and Gaffard 1988).

point of view, it is certainly one from the point of view of the “rationality” of individual agents, since at each step of an infinite sequence they must be able to recognise correctly “the new”: for example, they must be able to recognise a “new event”, even if, of course, they could still be “uncertain” on its probability distribution.

Under both these theoretical alternatives, uncertainty about the possibility of something new (new commodities, technologies, etc.) seems hardly reconcilable with the “maximising rationality plus market surrogation” hypothesis. In the former case, one must drastically reduce the domain of “newness” to an identifiable set of possible events. Whether they are finite or not, the set must be complete and “realisable”. In the latter case, the constraint is on the computational and interpretative capabilities of the agents: a sequential economy, which the agents know it is sequential, impose on each “rational” agent an increasing (quickly, an infinite) computational burden and/or require some “hidden” or “creative” capability of recognising “newness”, which certainly the decision-theoretical model does not specify.

This is the fundamental point. Irrespective of whether or not contingency markets empirically exist, the domain of theoretical applicability of an axiomatic theory of “rational choice” is that whereby uncertainty regards incomplete information about the occurrence of a *known list* of events⁹.

In other words, axiomatic representations of microeconomic behaviours under uncertainty as straightforward “rational choice” problems are theoretically legitimate (irrespective of the existence of the relevant contingency markets), only under suitable (and quite strong) restrictions on the nature of uncertainty itself and on the nature of the choice-problem. Thus, one should, first, assess the sources and nature of uncertainty in different choice settings. Second, “[a]n analysis of behaviour under uncertainty must answer [the following] questions”: (i) how are uncertain prospects formulated by the unit of decision? (ii) How does the actual occurrence of contemplated events affect the estimation of prospects? ... (iii) How is the choice between uncertain prospects arrived at?” (Hahn 1985, p. 211). We add, (iv) how does the possibility of un contemplated events affect behaviours? Putting it another way, “How is knowledge acquired? ... How is knowledge acted upon?” [Ibid.]. These are precisely the issues we shall discuss in the following.

4. Procedural uncertainty: representation, interpretation, rules and problem-solving

In order to highlight the fundamental problem-solving nature of the decision process with procedural uncertainty, let us consider first those simplest cases where uncertainty emerges only as a result of the complexity of the decision task, despite substantive certainty (that is, despite complete information on “events”).

⁹ Another way of interpreting the foregoing argument is with reference to Shackle’s view of uncertainty (Shackle 1969) in terms of ignorance about “residual” events and (related) “surprise”. Katzner (1986) shows that a behavioural theory of “surprise” (with unknown events) cannot be reduced (i.e., is not formally equivalent to) an axiomatic theory of choice with subjective probabilities about a known list of events. As one shall see, that interpretation of behaviours under uncertainty is easily consistent with our argument and, indeed, possibly also a fruitful conjecture on the empirical regularities in signal-detection and response under several uncertainty conditions.

Take strategic games with complete and perfect information (chess, nim, etc.) and also “puzzles” (e.g., Rubik cube).

In these strategic games and puzzles the procedure of search for a strategy (a winning one in the former case, an efficient one in the latter) may be represented by formal rules. We have “states” or “configurations” (the position of the various pieces on the chess-board, the position of the different colours on each side of the Rubik cube, etc.) and we have “rules” – perfectly known to the players – for the transformation of one state into another one. Each player must solve the problem of choosing a strategy made of a sequence of transformations leading from the initial state to the desired one (Nilsson 1980). Reaching a (the) winning configuration is generally a complex procedure, given the high number of combinatorial possibilities (“states”) that may be generated and must be examined (this applies to both strategic games and puzzles like the Rubik cube).

In two-persons games a winning strategy for either player exists, and, as already mentioned earlier, may be expressed through the well-known von Neumann algorithm. However, such an algorithm requires the generation of all the game-tree and creates problems of computational complexity. The same holds for puzzles, like the Rubik cube. An algorithm for the generation of all final states starting from any initial (scrambled) one exists. In fact, as we shall see later, there are more than one.

Suppose however that an agent faces a *new* problem: say, he is given for the first time a Rubik cube, taught the rules and told the solution concept.

He can follow three procedures, namely:

1. Explore extensively the game-tree according to a general search algorithm (Nilsson 1971) (assuming that this does not exceed his computational capabilities). Remarkably, this is the procedure which in general would correspond to that implied by standard maximising choice-theoretic models. Call this as the *orthodox player*. As well known, an agent behaving accordingly with the procedure may easily fail to reach the task, because of the high computational effort and memory storage generally required.
2. Use locally the algorithm orienting the search through *ad hoc* criteria. He will examine sub-trees and attribute empirical static evaluators to the intermediate positions (states) that are reached. This is the way many people seem to play, for example, chess. It is also the way “artificial players” do it (i.e., chess-playing computer programmes). Finally, it is the way that, in the economic literature, the “bounded rationality” of empirical agents is often interpreted, as an approximation – imperfect for quantitative limitations on computational capabilities – to the complete rationality of the choice-theoretic model (see, for example, Baumol and Quandt 1964). Call this the *satisficing player*.
3. Try to find a new problem-specific algorithm. He will temporarily give up the solution of the problem and try to solve sub-problems, moving in the sub-problem space, i.e., attempt to create new representations of the problem (Newell & Simon 1972). The problem is posed in a very incomplete way (a rather limited number of positions is explored), and the actual sequence of configurations may be redundant (an analyst with the unlimited computational capability apt to pursue procedure (1) could find a shorter sequence), but the level of abstraction is higher than that of both the orthodox and the satisficing players: the problems that he tries to solve do not hold only for a specific state but for an entire class which share the same property. Call this the *innovative player*.

As we will see in the following, neither of the three procedures can be considered “the best”, independently from the nature and conditions of the game (e.g., the complexity of the game-tree, its rules, the number of times that the game is played, etc.). Moreover, it is in principle impossible to establish *ex-ante* (even for an external analyst, more so for the players) whether a procedural choice is better than another one, e.g., in terms of time and costs of search. The knowledge of the set of events which must be generated or explored is necessarily incomplete and so also is the knowledge of the elements which will lead to the solution (types of mental operations, transformation sequences, time, etc.).

Formally, procedures (1) and (2) do not present particular problems of interpretation. However, procedure (3) does involve somewhat subtle issues related to knowledge and capabilities that the agent must pre-possess and that cannot be derived from what he was told on the rules of the game, the solution concepts, etc. In order to analyse them, it might be worthwhile to describe briefly how the “solution of problems” can be represented in formal terms.

Let us assume that the problems are representable in an appropriate formal language. First, each subject (agent) will pre-possess a set of elementary symbols, such as the words in language through which the problem can be described. In the case of the Rubik cube, the elementary symbols may represent the little coloured squares of which the cube is made. A sequence of these symbols, a string, will represent a configuration (a state) of the cube. Second, there are transformation rules, through which each phrase of the language, each string, can be modified to generate other strings. In our example, the rules specify how to rotate the mobile parts of the cube. Solving the problem (reaching the solution) means finding a sequence of transformations which get to a final given state. Thus the agents must have the capability of exploring the tree generated from the initial configuration by applying the transformation rules. The initial string may be formally called an “axiom”, in the sense of Computation Theory, since it is the only initial “postulation” from which the strings (theorems) that belong to the system are generated (see for example Cutland 1980).

Let us summarise the main characteristics of the problem-solving framework that we have defined for strategic games and puzzles. We may establish: (1) an alphabet $S = (A, B, C, D \dots)$ whose strings depict states or configurations of the problem (states of the game); (2) some rules of transformation which act upon the strings of alphabet S , each of which is symbolised by a letter of a second alphabet $P (= a, b, c, d \dots)$.

One usually indicates with S^* the set of all the strings of finite length that can be constructed with letters of alphabet S , and with P^* the set of all the strings of finite length that can be constructed with letters of alphabet P .

Note also that to pass from a position $s_1 \in S^*$ of the problem to another state $s_2 \in S^*$ it is necessary to write a sequence of P symbols, which constitutes a *procedure* describing the solution of the game (we will term such procedures as *routines of programmes*).

It is easy to recognise that in general not all the sequences of P^* constitute a programme, because there will be sequences which one cannot apply to certain configurations of the problem. The set of admissible sequences or programmes, i.e. those sequences that can be applied to a state of the game to produce a new position, constitute a *formal language*.

Any strategic situation which can be described in the normal terminology of games theory can be translated into the terms of a formal language. Such a

language which characterises it is defined by the set of all the programmes which can be put into practise in the context of the given strategic situation.

In fact, what we have described is a special case of a *Post axiomatic system*. All puzzles (like the Rubik cube) can be represented as Post systems, i.e., can be interpreted in a problem-solving specific framework¹⁰. Moreover, we can represent the problem of finding a strategy in a game (with two or more players) by means of a Post System.

Having posed a problem, if a solution exists, will there be a programme (procedure) that represents that solution? How can one reach such a procedure?

Recall the Rubik cube. The initial position is given by chance and the transformation rules are elementary rotations around the vertical and horizontal axes. The solution is a sequence apt to reach a final configuration. Consider the solution procedures of the “orthodox”, “satisficing” and “innovative” players. The “orthodox” one will orderly examine all the possible sequences, i.e., will write all the possible procedures which can be generated from the initial position. Obviously, he needs sufficient memory in order to compare the new sequences with the previous ones and avoid the generation of an infinite tree. The “satisficer” will basically follow the same procedure but explore only a limited number of branches of the tree and perhaps “stop” and “go back” according to ad hoc evaluators. The “innovative” player will try to reach a more synthetic representation of the problem and a more powerful language (which implies more general transformation rules). When the search is successful, this reduces the procedural uncertainty by reducing the computational complexity of the problem.

Let us analyze how this task can be performed, by way of an illustration. Consider an hyper-simplified Rubik cube made of two juxtaposed squares (back and front). Each square is composed of four elements (four little squares) of four different colours. Denote the colours as A, B, C, D. Each element has the same colour in the front and in the back. Adopt the convention of calling the position of an element, clockwise from the top-left element as 1, 2, 3, 4. So, for example, the string ABCD represents a configuration where colour A is in position 1, B in position 2, etc. The rules of transformation are rules of 180 degrees rotation around the two axes with four possibilities, “up” (u) (rotation of the top half of the square around the vertical axis); “down” (d) (rotation of the bottom half around the vertical axis); “left” (l) (rotation of the left half of the square around

¹⁰ Let $\Sigma = \{a_1, a_2, \dots, a_n\}$ be a finite set of symbols, called an *alphabet*. A *string* from Σ is, as we previously stated, any sequence $a_{i_1}, a_{i_2}, \dots, a_{i_m}$ of symbols from Σ . For any alphabet Σ we denote as Σ^* the set of all the strings from Σ . Let S, S_1, \dots, S_n be strings from Σ^* . Post called *Production* a general rule of transformation between strings, which takes the form

$$(a) \ g_1 S_1, g_2 S_2, \dots, g_{m-1} S_{m-1}, g_m S_m \rightarrow h_1 S_{i_1}, h_2 S_{i_2}, \dots, h_n S_{i_n}$$

where

- (i) $g_1, g_2, \dots, g_m, h_1, h_2, \dots, h_n$ are given (fixed) strings
- (ii) the subscript i_1, i_2, \dots in stand for strings all derived from 1, 2, ... m and which need not to be distinct.

A Post (canonical) system consists of

- (1) a finite alphabet Σ
- (2) a finite subset of Σ^* , (the subset A of the *axioms*)
- (3) a finite set of *Productions* of the form (a) above, whose fixed strings are in Σ^* . It can be shown that a system of Post is equivalent to a Turing machine (Cutland 1980, 3rd chapter; M. Minsky 1967).

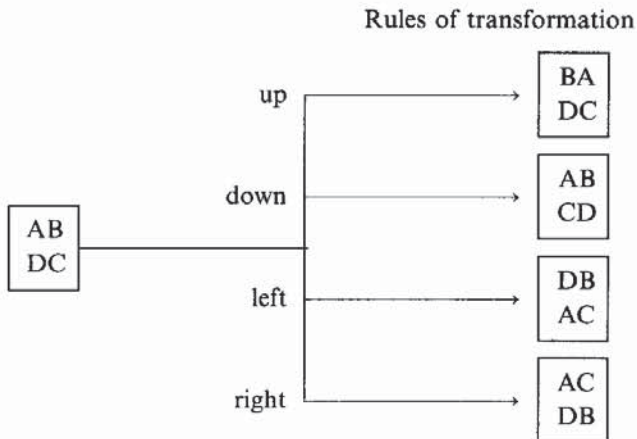


Fig. 1

the horizontal axis); “right” (r) (rotation of the right half of the square around the horizontal axis). Of course, with each rotation, half of the back face comes to the front, and vice versa. The game is easily represented in a system of Post. Each string with the four letters (symbols) A, B, C, D is a configuration of the game. For example, reading clockwise, the string ABCD is depicted in Fig. 1.

One can easily see that each transformation exchanging adjacent letters is admissible (e.g., ABCD into DCBA). One can also consider as adjacent the first and the last letter of each string, by representing each string on a ring divided in four sections. Thus we have:

- (1) an axiom, say, an initial position ABCD
- (2) the rules of transformation. The are
 - a) exchanges between adjacent letters:
 - ABCD \rightarrow BACD (by applying the up rule)
 - ABCD \rightarrow ACBD (left)
 - ABCD \rightarrow ABDC (down)
 - b) exchanges between the leftmost and the rightmost letter
 - ABCD \rightarrow DBCA (right)

- (3) An objective: a certain arbitrary final configuration.

A procedure (a programme) establishes a sequence of transformations (for example **rrllul**). Each procedure can be obtained by combining the four elementary procedures, **u, d, l, r**.

Consider in detail the strategy of the “innovative player”. Suppose he faces with the following:

Problem 0: <reach the configuration ABCD starting from any one scrambled configuration>

This problem can be immediately divided into the following sub-problems:

- (1) move A from the initial position in the string to position 1 (top left)
- (2) examine the element in position 3 (bottom right). If it is not C, then move C (which can only be in positions 2 or 4) to position 3
- (3) if B and D are not in the required places 2 and 4, exchange them.

In order to solve the sub-problems 1 he must solve the following:

- (a) Move a colour from any one position to any other position. This involves in turn sub-problems, i.e., (a1) move a colour to an adjacent position, and (a2) move a colour to the diagonal opposite position on the square.

Programmes with length one (**u**, **d**, **l**, **r**) solve the first sub-problem. Obviously, one solves the second by applying the first transformation twice (hence, procedures of length two, e.g. **ud**, **ul**, ...). Note that repetitive procedures (**uu**, **dd**, ...) do not yield any configuration change and the procedures **ud**, **lr**, **rl**, **du** invert the square to its symmetric opposite. Call all the procedures which solve problem (a) *adjacent change routines*.

The problem (a) is now reduced to elementary sub-problems which can be solved with *adjacent change routines*. Problem (2) is trivially solved by an adjacent change routine.

Sub-problem (3) can be solved if the player is able to:

- (b) Exchange two colours along a diagonal *leaving the rest unchanged*. In turn this involves two sub-problems, i.e., first, (b1) exchange along the north-west/south-east diagonal, and (b2) exchange along the north-east/south-west diagonal.

One can easily see that the procedures **uru** and **dld** solve the first sub-problem and the procedures **ulu** and **drd** solve the second one. Call all the procedures which solve problem (b), *diagonal routines*.

Finally, the player has reduced the original problem to a set of sub-problems which can easily be solved by means of simple routines. It is easy to recognize that any problem like Problem 0 can be solved by a combination of the routines (procedures) which solve problems (a) and (b).

Let us verify this “theorem”, supposing for simplicity that we want to transform any initial scrambled configuration into the final configuration ABCD.

Starting from the scrambled configuration, move A into position 1 (top left). In order to do so, use the adjacent change routines. Then examine the element in position 3 (bottom right). If it is not C, then move C (which can only be in positions 2 or 4) to position 3 using an adjacent routine (**d** or **r**). Finally, if B and D are not at the desired place, use the diagonal routines **ulu** or **drd**. Thus, by simply using the two (“adjacent” and “diagonal”) routines, it is possible to obtain the solution, ABCD, or whatever other desired final configuration.

The entire class of problems of reaching a final configuration starting from a scrambled one is then solved by *recursive decomposition* of the original problem in sub-problems.

However, one must notice that this is not the only way to solve them. We leave to the reader the demonstration that the same result (e.g., configuration ABCD) can be obtained by decomposing the original problem into the following sub-problems:

- (1) Move A to position 1 with adjacent routines;
- (2) Move B to position 2 with adjacent routines, leaving A in position 1.
- (3) If C and D are not in the required place, exchange them using the adjacent routine **d**.

Note that in this case the routines are not the same as the previous “theorem”: no use is made of “diagonal routines” but only a sequence of “adjacent routines”.

It is important to notice that the analysis of the problem has led to a new, more synthetic, representation of the same problem. The *solution of the sub-problems led in fact to the construction of routines, that are the building-blocks of a new, more abstract, language*. In addition to the elementary routines **u**, **d**, **l**, **r** – defined earlier – we can define:

q (= **ulu** = **drd** = **rur** = **ldl**)
(exchange of colours on the right diagonal)

w (= **uru** = **dld** = ...)
(exchange of colours on the left diagonal)

v (= **uu** = **dd** = ...)
(null operation)

s (= **ud** = **lr** = ...)
(specular inversion of the entire square)

The new routines define new “elementary transformations” whose combinations easily allow *any* problem of the class of the previous Problem 0 to be solved and also the general Problem 00: <transform any initial scrambled configuration into any final configuration>.

Note that starting from the original system which operated on strings like ABCD, etc., we have defined a new set of transformation rules by means of an “intelligent”, *inductive* procedure based on problem decomposition. This new language is able to represent the original problem in a more tractable fashion, from the point of view of its complexity.

As we have seen, there are at least two different ways to decompose the original problem, and consequently players can generate at least two different languages. How can a player find a “good” decomposition? Is there a general algorithm which allows us to decompose a problem? Are there many different algorithms for the decomposition such that we can choose the best one?

To answer these questions, let us recall that what we have described is a special example of a Post axiomatic system, i.e., a system which generates “theorems” (new strings of characters) from “axioms” by the application of “transformation rules”. Post used it to describe the working of a *formal logical system*. Puzzles are special examples of Post systems: then, the properties of Post systems can be applied to them and generally to any formalized problem-solving framework.

In general, one can formalize the notion of *problem solving* by means of Post system: solving a problem is simply to reach a final configuration starting from a given one. If it is possible to transform the latter into the former by means of repeated applications of the transformation rules – as in the Rubik simplified puzzle – then one obtains the solution of the problem: one is able to write a “programme”.

However, a fundamental implication of the equivalence between a Post system and a Turing machine (see for example (Cutland 1980; Traktenbrot 1963) is that to both applies the so-called *unsolvibility of the word problem*, i.e.: Given two strings from the same alphabet, and a finite set of transformation rules, there is *no* general algorithm that can say if the two strings *are* derivable one from the other or *not*.

Here a “general algorithm” means an algorithm valid for any possible Post system. As in the previous example, it is possible to solve the *word problem* for a specific Post system, or a set of them; but not for the entire class of Post systems.

It follows that it is impossible to generate a programme which solves a given formalized problem by using a general problem-solver algorithm. In other words, the players who look for efficient algorithms for the solution of specific games *cannot draw on general rules for construction of algorithms, because they do not and cannot exist.*

This is the intrinsic element of procedural uncertainty. In a general problem-solving framework one does not know and cannot know *ex-ante* which or how many algorithms can be generated, or if they can be generated at all. Relatedly, this search for new algorithms involves “true” uncertainty and cannot be simply interpreted as “risk” with a cost and a probability of success.

As a consequence of all this, the main characteristics of the problem solving search, when successful, are the following:

- (1) It produces routines which apply to entire classes of problems (they are higher-level rules of transformation).
- (2) There may be more than one routine (more than one sequence of elementary operations) which solve the problem. Generally, it is impossible to establish which one is “better” independently from the initial configuration of the problem.
- (3) Even if the routines are “derived” from the original system, there is no general algorithm of derivation. That is, the procedure of derivation cannot be automatized.

Let us discuss some consequences of these results in terms of “rational” decision making. We have argued earlier that the theoretical boundary of applicability of a rational choice-theoretical model as a positive theory of behaviour is set by the recursive computability of choice functions. It has been shown (Lewis 1985, 1986) that this condition may not be expected to hold in general also in stationary environments (General Equilibrium models, certain Nash-type models of strategic interaction).

Further restrictions emerge if we formalise the decision problem in a general problem-solving framework: since the problem of generating a solution for *any* class of problems is unsolvable one must also rule out the possibility of automatically deriving uniform decision rules from a general principle of “rationality”. Many classes of problems are of course solvable, but any new class (not included in a previously solved one) requires a solution which can be found only by means of an *inductive search process*.

What do people generally do, then? They use routines which imply “higher level” representations of the problem, de-composition into familiar sub-problems, powerful “rules of transformation” that lead from the information about “the problem” to its “solution” (i.e., to the decision) and that do not only apply to single problems but to entire classes of them.

As we saw, in a general problem-solving framework even when “routines” are “derived” from the original information about the problem, they entail more “information” than that contained in its original formulation. Their “derivation” involves some inductive ability – which the agent must possess – of generating new knowledge which is not contained in the original information.

In line with the description of discovery as a result of a search in the problem-space (Simon 1981), we suggest that this behaviour, based on higher-level representation of problems and routines, can be generally attributed also to the problem-solving/decision activity of economic agents.

Of course, the agents most often use these routines automatically. Indeed, their efficiency and their robustness stems partly from the very fact that they can be used automatically. However, repetition automates an original “creative” act. If we are allowed a metaphor, “mechanical stupidity” hints at an original product of intelligence¹¹. “Routinised” decision processes – in the sense discussed here – reduce procedural uncertainty. Further, the more general the routines are, the more capable they are to deal with variations in the nature of the problem to be solved (i.e., the class of problems that they can efficiently tackle is wide). With sufficient “generality” or “abstraction” they can tackle ill-structured problems which typically arise with strong substantive uncertainty (defined earlier). Conversely, the search for new routines is, as argued, characterised by an intrinsic procedural uncertainty.

Whether the choice/behaviour which one can expect on the grounds of our interpretation correspond to that which could be predicted on the grounds of the “rational” choice theoretic-model (within the limits of its applicability) depends on the nature of the decision-problem. Of course, the nearer one gets to substantive certainty, computational simplicity and tautology, the more likely it is that the two coincide. In the other cases, there is no a priori reason to expect that they generally will. The foregoing interpretation of the relationship between various forms of uncertainty and decision-making, we believe, has important implications in terms of economic analysis. We shall briefly consider some examples in the following section.

5. Routines, production and innovation

One of the direct applications of this analysis of behaviours under substantive and procedural uncertainty is to the theory of production¹². If one recognises the essential problem-solving nature of production activities (certainly infinitely more complex than a Rubik cube, even leaving aside substantive uncertainty), one can easily represent them in the formal language suggested earlier – strings, rules of transformation, etc. The production routines (the technology) codify the procedure and the knowledge involved in the solution of particular classes of production problems. These routines are – to different degrees – specific to these classes and to the people and the organisations (typically, firms) who have developed them. Of course, the transferability of these routines from one organisation to another has to do also with their degrees of tacitness, the nature of the knowledge that their original generation and their implementation involve, etc. The tree of the possible techniques which can nationally be generated on the grounds of a certain set of initial information, skills, chemical/physical principles is plausibly infinite. So, our earlier considerations on the impossibility of automatic derivation of the actual routines strictly applies. Some routines are “stored” in individual

¹¹ Vercelli (1986) quotes Husserl, saying that “... tradition is forgetting the origins...”.

¹² More on this point in Egidi (1986). This analysis is broadly complementary to that of Winter (1982) and Nelson and Winter (1982). What follows is also strictly consistent with the analysis of technology in terms of “technological paradigms” of Dosi (1982 and 1984).

people, many others in organizations, which often reproduce and improve them through their use (Nelson and Winter 1982).

In the framework proposed here, the representation of technical change is straightforward: it is the generation of new routines. These new routines may relate to a new production process – different, hopefully more efficient, rules of transformation of certain inputs and information into a given output – or to the conception of a new product, or a new organizational set-up. Again, the automatic underivability of new routines highlights the meaning of technological uncertainty. There is not, and there cannot be, a general rule which leads innovative search. Indeed, there can only be specific search heuristics, strongly characterised by procedural uncertainty.

The model allows also a theoretically easy distinction between what one of us calls “normal” technical change (i.e., technical change along a certain “technological trajectory”) as distinguished from “extra-ordinary” (paradigm-) changes (Dosi 1982).

In the language introduced earlier, “normal” innovations are generated, although not automatically, by “derivation”, analogy, refinement from a basic set of “axioms” and transformation rules and an underlying specific knowledge, i.e., a given paradigm. A paradigm change occurs when it involves also a change of “axioms” and transformation rules. (Alexander the Great’s solution to the problem of the Gordian Knot is one of the simplest examples of a radically new system of Post leading to a new routine that in turn solves a problem for which no efficient routine could be derived on the grounds of another set of axioms and transformation rules.)

The development, improvement and modification over time of specific routines on the grounds of a certain body of knowledge accounts also for the cumulative and local nature of technical change. That is, one empirically finds that firms and industries do not explore the entire “production possibility set”, whatever that means, but typically improve and develop upon the areas of their existing competence, yielding to quite defined “trajectories” of technical progress¹³.

Both “radical” and “normal” innovations yield routines – of varying degrees of generality – which reduce the procedural uncertainty of the innovators themselves, but, other things being equal, increase the substantive uncertainty of the other agents, in that they continuously introduce novelty (“new events”) in the environment.

In this interpretation, the flexibility of a particular technology (a particular body of knowledge) means, at one level, the opportunity that it offers of generating a wide set of efficient routines (of course, this, strictly speaking, can be assessed only *ex post*). At another level, the flexibility of a particular routine relates to its generality and robustness (i.e., its “abstraction”), and, thus, the width of the set of problems to which it can be successfully applied (this obviously includes the set of products that can be made and the different environmental conditions that it can meet)¹⁴. That is, the flexibility of a technological routine relates to its robustness in dealing with the substantive uncertainty stemming

¹³ For a survey of the evidence, Dosi (1988). On cumulative and local technical progress see also Atkinson and Stiglitz (1969) and David (1975). On the related nature of search processes, Nelson (1981).

¹⁴ An analysis of the relationship between production routines, flexibility and economies of scope is in Colombo and Mariotti (1986).

from the high dimensionality of the event-space, and the incompleteness of the ex-ante knowledge that agents have of them.

One can establish a natural link between the characterisation of production activities in terms of specific knowledge and routines, on the one hand, and the nature of business firms, on the other. Certainly, an important part of the explanation of the existence and forms of internal organisation of modern corporations resides in the problems of performance control and transaction costs (Williamson 1975, 1985). However, equally important, firms – or whatever other form of productive organisation – must exist in so far as the complexity of routines (and related specific knowledge) required by a particular problem-solving activity is beyond the capabilities of an individual agent (for more on the relationship between knowledge specificities and the boundaries of firms see Dosi, Teece and Winter 1991).

Note also the dual characteristics of problem-solving. On the one hand, it implies identification of sub-problems, their codification and, as a consequence, specialisation and division of labour. On the other hand, it involves strict coordination among the different tasks. Those reasons which account for the impossibility of automatic derivability of the routines make also very problematic and inefficient the complete separation of the various pieces of knowledge leading to the solution of a particular task (it would be as if, in the example of the Rubik cube discussed earlier, one would study the movements of colour A, another those of colour B and then trade the results). Indeed, organizations embody, to varying degrees, both specialisation and coordination within company-specific competences.

In our interpretation, it generally happens that, for sufficiently complex problem-solving tasks, in non-stationary environments, the specific knowledge which is required cannot be defined independently from the organisation in which it is developed and applied. Obviously, there are pieces of people-embodied knowledge which are clearly identifiable and traded on the markets for labour or for *specific services*. If a firm needs to solve a problem of basic semiconductor properties of a certain material, of course, it is more likely to hire a physicist than a carpenter. However, several problem-solving routines are embodied in the way the organisation works rather than in precisely identifiable competences of individual people, even if clearly it makes use of these competences. The indeterminacy of labour contracts (on what is actually “delivered”) has partly to do with this phenomenon. So has the interpretation of “de-skilling” which has often been observed in the history of labour-processes: what were previously specific people-embodied competences – and thus traded as such, on the labour market – become attributes of organisational routines, well beyond the knowledge-inputs of the individual performers of certain tasks.

Hence, as argued at greater length in Dosi, Teece and Winter (1991), the general conjecture is that the scope of the firm is defined by the domain of applicability of its specific knowledge and problem-solving routines¹⁵. Innovation, of course, changes (1) the routines, (2) the forms of specific knowledge required for their development and/or implementation and (3) the degree to which a certain problem-solving task can be subdivided into relatively autonomous and general “sub-problems”. Therefore, it changes also the scope of the firm. Seeing it from the mirror opposite angle, it changes the scope for “Smithian” intra- and inter-firm division of labour.

¹⁵ Teece’s analysis of the co-specialised assets which underpin a firm’s competitiveness is clearly consistent with this interpretation (see Teece, 1982 and 1986).

6. Conclusions

Possibly one of the most heard common sense arguments in favour of a “rational” choice-theoretic model of behaviour is that it is simply the theoretical representation of the attempt of empirical agents to make the best use of information that they can, and that, conversely, a routine representation would imply some sort of “mechanicity” or “stupidity” on the part of the agents. We showed that under most circumstances involving various sorts of substantive and procedural uncertainty, precisely the opposite holds true. With sufficiently complex problem-solving tasks as, we argued, most human activities within and without the economic sphere are, “intelligent” actors will look for problem-solving routines which are “abstract” and robust, in the sense that they will not only apply to single problems but classes of them. In fact, the development of these routines involves forms of specific, inductive and synthetic knowledge which is not implied in the “information” about the problem itself. In the most general sense, innovation is the process of discovery of new routines. The nature and, indeed, the *necessity* of such routines can be founded in a quite general theory of problem solving, as pioneered by Simon and, earlier, by the theories of computation and formalization of natural languages. Indeed, such results are highly complementary to the findings on the evolutionary emergence of norms in complex biological and social environments (cfr. Holland 1975; Holland et al. 1987).

In this interpretation, *innovation is an endogenous mechanism of generation of uncertainty in the environment*, since each innovation is in act the appearance of a new, unexpected event. At the same time it is a *procedure through which the innovator*, when successful, solves more efficiently, with greater generality, a particular set of problems, and, hence, *through the resulting routines, reduces complexity and procedural uncertainty*. Given *strong uncertainty*, defined earlier, the main activity of each agent will not reasonably be a Bayesian updating of probability distributions upon a set of events that he considers as entirely “external” and natural. On the contrary, he will try to find, in ways which are agent-specific, new and, so to speak, “deterministic” procedures of control upon, and generation of, new events. The entire process inevitably produces endogenous non-stationarity of the environments. More general routines are developed and, to different degrees, diffused and imitated by other agents and firms. In turn, the general opportunity of development of new routines yields environmental complexity and strong substantive uncertainty¹⁶. The permanent gap between the notionally efficiency of the currently existing routines and those which could be developed on the grounds of all the existing information and knowledge defines the (uncertain) domain of “Schumpeterian” innovative behaviours.

Putting it another way, there are possibly in economics two extreme mirror opposite views. The first is that, the R. Lucas’ words, “there are no \$ 100 bills on the sidewalk”, or a bit more extremely, as a member of a committee encharged to decide whether to invite a Rational Expectation economist as a professor, once

¹⁶ If we are correct in this representation of innovation and, more generally, non-stationarity, one might wonder what is the interpretative status of a growing stream of literature which represents the innovative process (e.g., R&D investments and innovative outputs) in terms of the usual maximisation procedure under uncertainty. As argued in Dosi and Orsenigo (1988), they often yield important “negative” or “a fortiori” results (...“even if the agents could be literal maximisers, the equilibria would be multiple, would not generally be Pareto-optima...”, etc.). However, we believe they offer only limited insights in the processes leading to innovation.

ironically said, “if it were a good idea he would already be here”. The world always runs on the edge of its possibilities and everything worth happening has already happened. At the other extreme, there is the view that “intelligent” processes of discovery, and the intimately associated processes of routinisation and automatic application of routines, always create notional unexploited opportunities.

We stand by this latter view: Yes, there always are \$ 100 bills on the sidewalk, even if it is impossible to axiomatise the process leading to their discovery. The changing focus that we have proposed here from a purely informational notion of uncertainty to another – procedural – one, associated with the problem-solving competences of the agents, allows precisely the analysis of economic behaviours that often appear “patterned” and routinized, but also present recognisable changes and discoveries.

Acknowledgements. Comments on an earlier draft by the participants at the International Workshop on “Programmable Automation and New Work Modes”, Paris, 2–4 April 1987 and at the International Political Science Workshop, Dubrovnik, October 1988 are gratefully acknowledged. Support to one of the authors (G.D.) by the Designated Research Centre of the ESRC at the Science Policy Research Unit (SPRU), University of Sussex, and to both by the Italian Ministry of Education (Progetti M.P.I. 40%), is gratefully acknowledged.

References

- Akerlof G (1970) The market for lemons: qualitative uncertainty and the market mechanism. *Quarterly Journal of Economics* 84: 488–500
- Alchian AA (1951) Uncertainty, evolution and economic theory. *Journal of Political Economy* 57: 211–221
- Amendola M, Gaffard JL (1988) *The innovative choice*, Oxford, Basil Blackwell
- Arbib MA (1969) *Theories of abstract automata*, Englewood-Cliffs, NJ, Prentice-Hall
- Arrow KJ (1951) Alternative approaches to the theory of choice in risk-taking situations. *Econometrica* 19: 404–437
- Arrow KJ (1951a) *Social choice and individual values*. New York, Wiley
- Arrow KJ (1970) *Essays in the theory of risk-bearing*. Amsterdam, North-Holland
- Arrow KJ (1971) Exposition of the theory of choice under uncertainty. In: McGuire and Radner
- Arrow KJ (1982) Risk perception in psychology and economics. *Economic Inquiry* 20: 1–9
- Arrow KJ (1983) Innovation in large and small firms. In: Ronen
- Arrow KJ, Hahn F (1971) *General competitive analysis*. San Francisco, Holden-Day
- Arthur WB (1985) Competing technologies and lock-in by historical events. The dynamics of allocation under increasing returns. Stanford, Stanford University, CEPR Discussion Paper
- Arthur WB (1988) Competing technologies. In: Dosi, Freeman, Nelson, Silverberg, Soete
- Atkinson AB, Stiglitz JE (1969) A new view of technological change. *Economic Journal*: Sept. 1979: 573–578
- Baumol WJ, Quandt RE (1964) Rules of thumb and optimally imperfect decisions. *American Economic Review*
- Cass D, Shell K (1983) Do sunspots matter? *Journal of Political Economy* 91: 193–227
- Colombo MG, Mariotti S (1986) Flexible automation, idle times and economies of scope. Milan, Politecnico, Department of Electronics, Rapporto Interno no 86-009
- Coricelli F, Dosi G (1988) Coordination and order in economic change and the interpretative power of economic theory. In: Dosi, Freeman, Nelson, Soete, Silverberg
- Cutland NJ (1980) *Computability. An introduction to recursive function theory*. Cambridge, MA, Cambridge University Press
- David PA (1975) *Technical choice, innovation and economic growth*. Cambridge, MA, Cambridge University Press
- Davies M (1958) *Computability and unsolvability*. New York, McGraw-Hill
- Debreu G (1959) *Theory of value*. New York, Wiley

- Diamond P, Rothschild M (eds) (1978) *Uncertainty in economics*. Orlando, FL, Academic Press
- Dosi G (1982) Technological paradigms and technological trajectories. A suggested interpretation of the determinants and directions of technical change. *Research Policy* 11: 147–162
- Dosi G (1984) *Technical change and industrial transformation*. London, Macmillan
- Dosi G (1988) Sources, procedures and microeconomic effects of innovation. *Journal of Economic Literature*
- Dosi G, Orsenigo L (1988) Structures, performance and change in evolutionary environments. In: Dosi, Freeman, Nelson, Soete, Silverberg
- Dosi G, Freeman C, Nelson R, Soete L, Silverberg G (eds) (1988), *Technical change and economic theory*, London, Francis Pinter
- Dosi G, Teece D, Winter S (1991) Toward a theory of corporate conference. In: Dosi G, Giannetti R and Toninelli PA (eds), *Technology and the enterprise in a historical perspective*. Oxford, Oxford University Press (Forthcoming)
- Egidi M (1986) The generation and diffusion of new market routines. Presented at the International Conference on Innovation Diffusion, Venice, 17–21 March
- Friedman M (1953) *Essays in positive economics*. Chicago, University of Chicago Press
- Gaffard JL (1986) *Mutations technologiques et choix strategiques des entreprises*. Nice, University of Nice, LATAPSES, mimeo
- Hahn F (1984) *Equilibrium and macroeconomics*. Oxford, Basil Blackwell
- Hahn F (1985) *Money, growth and stability*. Oxford, Basil Blackwell
- Heiner RA (1983) The origin of predictable behaviours. *American Economic Review*
- Heiner RA (1988) Imperfect decisions, routinized behaviour and inertial technical change. In: Dosi, Freeman, Nelson, Soete, Silverberg
- Hey JD (1982) Search for rules for search. *Journal of Economic Behaviour and Organisation* 3:65–81
- Hirshleifer J, Riley J (1979) The analytics of uncertainty and information: A survey. *Journal of Economic Literature* 17: 1375–1421
- Hogart RM, Reder MW (eds) (1987) *Rational choice*. Chicago, The University of Chicago Press
- Holland JH (1975) *Adaptation in a natural and artificial systems*. Ann Arbor, MI, Univ. Michigan Press
- Holland JH, Holyoak KJ, Nisbett RE, Thaeard PR (1986) *Induction: Processis of inference learning and discovering*, Cambridge, MA, MIT Press
- Hopcroft J, Ullman JD (1969) *Formal language in their relation to automata*. New York, Addison-Wesley
- Hopcroft J, Ullman JD (1979) *Introduction to automata theory, languages, and computation*. New York, Addison-Wesley
- Kahneman D, Tversky A (1979) Prospect theory: an analysis of decision under risk. *Econometrica* 47: 263–291
- Kahneman D, Tversky A (1986) Rational choice and the framing of decisions. In: Hogard and Reder (1987)
- Katzner DW (1986) Potential surprise, potential confirmation, and probability. *Journal of Post-Keynesian Economics* 9: 58–78
- Lewis AA (1985) On Effectively computable realisations of choice functions. *Mathematical Social Science* 10
- Lewis AA (1986) *Structure and complexity. The use of recursion theory in the foundations of neo-classical mathematical economics and the theory of games*. Ithaca, Cornell University, Dept. of Mathematics
- Machlup F (1946) Marginal analysis and empirical research. *American Economic Review* 36: 519–554
- Marschak J (1950) Rational behaviour, uncertain prospects and measurable utility. *Econometrica* 43: 111–141
- McGuire CB, Radner R (eds) (1971) *Decision and organisation*. Amsterdam, North-Holland
- Merin A (1985) *Elementary social relations, duality principles and model paradigmata. Socio-morphic structures of english grammar*. Brighton, University of Sussex, Dept of Linguistics, DPhil Thesis
- Minsky M (1967) *Computation. Finite and infinite machines*. Englewood Cliffs, NJ, Prentice-Hall
- Nelson RR (1981) The role of knowledge in R&D efficiency. *Quarterly Journal of Economics*

- Nelson RR, Winter SG (1982) *An evolutionary theory of economic change*. Cambridge (MA), The Belknap Press of Harvard University Press
- von Neumann J, Morgenstern O (1953) *Theory of games and economic behaviour*. Princeton, Princeton University Press, 3rd ed
- Newell A, Simon HA (1972) *Human problem solving*. Englewood Cliffs, NJ, Prentice-Hall
- Nilsson NJ (1971) *Problem solving methods in artificial intelligence*. New York, McGraw Hill
- Nilsson NJ (1980) *Principles of artificial intelligence*. Palo Alto, Tioga Publishing Company
- Radner R (1968) Competitive equilibrium under uncertainty. *Econometrica* 36: 31–58
- Ronen J (ed) (1983) *Entrepreneurship*. Lexington (MA), DC Heath
- Savage LJ (1954) *The foundations of statistics*. New York, Wiley
- Shackle GLS (1969) *Decision, order and time in human affairs*. Cambridge, Cambridge University Press
- Sharples M et al. (1989) *Computers and thought*. Cambridge (MA), MIT Press
- Shelling TC (1978) *Micromotives and macrobehaviours*. New York, Norton
- Silverberg G (1987) Technical progress, capital accumulation and effective demand: A self-organisation model. In: Batten D, Casti J (ed) *Economic evolution and structural change*. Berlin/New York, Springer-Verlag
- Silverberg G (1988) Modelling economic dynamics and technical change: mathematical approaches to self-organisation and evolution. In: Dosi, Freeman, Nelson, Soete and Silverberg (1988)
- Simon HA (1957) *Models of man*. New York, Wiley
- Simon HA (1959) Theories of decision making in economics. *American Economic Review* 49: 253–383
- Simon HA (1973) The structure of ill-structured problems. *Artificial Intelligence* 4
- Simon HA (1979) Rational decision making in business organisations. *American Economic Review* 69: 493–513
- Simon HA (1981) *The science of artificial*. Cambridge (MA), MIT Press, 2nd ed
- Sternberg RJ (ed) (1984) *Mechanisms of cognitive development*. New York, Freeman and Co
- Stiglitz J (1984) Information and economic analysis: A perspective. *Economic Journal, Conference Papers*
- Stiglitz J (1986) Learning to learn. Presented at the Conference on the Economic Theory of Technology Policy, CEPR, London, September 1986
- Taylor M, Taylor P (1986) The surprise game: An exploration of constrained relativism. Coventry, Warwick University, Institute for Management, Aston-Warwick Scale Papers n 4
- Teece D (1982) Toward an economic theory of multi-product firms. *Journal of Economic Behaviour and Organisation*
- Teece D (1986) Profiting from innovation. *Research Policy*
- Trakhtenbrot BA (1963) *Algorithms and automatic computing machines*. Boston, DC Heath (English trans. from Russian original)
- Turing AM (1936) On computable numbers, with an application to the Entscheidungsproblem. *Proceedings of the London Mathematical Society* 42: 230–265
- Tversky A, Kahneman D (1974) Judgement under uncertainty: heuristics and biases. *Science* 185: 1124–1131
- Tversky A, Kahneman D (1981) The framing of decisions and the psychology of choice, *Science* 221: 453–458
- Vercelli A (1986) Technological flexibility, financial fragility and the recent revival of Schumpeterian entrepreneurship. Siena, University of Siena. Presented at the International Conference on Technological and Social Factors in Long Term Fluctuations, December 1986
- Williams EF, Findlay MC (1986) Risk and the role of failed expectations in an uncertain world. *Journal of Post-Keynesian Economics* 9: 32–47
- Williamson O (1975) *Markets and hierarchies*. New York, Free Press
- Williamson O (1985) *The economic institutions of capitalism*. New York, Free Press
- Winter SG (1982) An essay on the theory of production. In: Hymans SH (ed) *Economics and the world around it*. Ann Arbor, Michigan University Press
- Winter SG (1986) Adaptive behaviour and economic rationality. Comments on Arrow and Lucas. In Hogart and Reder
- Winter SG (1988) *Competition and selection*. The New Palgrave