

**INFORMATION AND BELIEFS
IN GAME THEORY**

(Bernard WALLISER, ENPC, EHESS, Paris)

September 30, 2006

Game theory is devoted to the study of strategic interactions between several rational players. It is concerned by information since players have some uncertainty about their environment and compensate for it by direct observation or by transmission from other players. It is concerned by beliefs since players form expectations about their environment by relying on representations about it as well as representations about others' representations. Information and beliefs are strongly linked since acquired information modifies the prior beliefs and uncertain beliefs ask for more information. Hence, information and beliefs became together a central topic of game theory and even appear as the main tool allowing the players to coordinate. Moreover, since game theory grounds economic theory, their study leads directly to the 'economics of information' (Macho-Stadler et alii, 2001) and to the 'economics of knowledge' (Foray, 2004).

The aim of this paper is not to summarize the huge technical literature concerning information and beliefs in game theory. Many chapters of game theory textbooks are devoted to the study of information, with a lot of illustrative examples (Osborne-Rubinstein, 1994; Aumann-Hart, 1992, 1994, 2002). Many proceedings of specialized conferences such as LOFT (1994 to 2006) and TARK (1986 to 2005) are devoted to the treatment of beliefs in game theory, even if a synthesis is not already available. The aim of the paper is rather to proceed to a rational reconstruction of information and beliefs in game theory. It shows in what terms the problems were initially conceptualized, how some formal tools were imported and considered to be adequate, and in what ways these problems were then considered as solved. Problems and solutions will be presented in a rather informal way, but various toy examples (used at length in game theory) will be mentioned.

First section recalls how information and beliefs were historically introduced in game theory and induced an evolution of the usual equilibrium notions. Second section presents what type of information is needed and gathered by a player and describes the formal tools used by the modeller in order to formalize it. Third section considers in which way such information is treated by a player, and especially how it contributes to the revision of its basic or crossed beliefs. Fourth section examines how information is captured by a specific equilibrium notion and how it is assessed by a player by means of an 'information value'. Fifth section analyzes information as a strategic item that a player has interest or not to diffuse, and examines if the equilibrium state reveals or not the players' information. Sixth section is devoted to the way

information is diffused among the players and examines if it leads or not to a common knowledge of information. Seventh section considers the learning process of players endowed with bounded rationality and examines if it converges or not towards a usual equilibrium state.

1. Ontology of game theory

Game theory is intended to describe formally how rational players engage in strategic interactions. Interactions are strategic in the sense that the consequences of a player's action depend on other players' action. Hence, each player follows a decision process taking into account his expectations about the others' behaviour. Such expectations are grounded on prior beliefs of the player as well as on information progressively gathered by him. Hence, in the history of game theory, information and beliefs became progressively more and more explicit. But game theory developed its own formalisms even if they appeared later to be similar to foreign ones.

1. Classical game theory

Any game is defined by two kinds of actors, a set of players (acting as genuine decision-makers) and nature (acting mechanically) involved in strategic interactions. Each player is assumed to choose an action independently by a rational decision process. More precisely, he is endowed with three 'determinants': opportunities (delimitating his action set), beliefs (representing his view of his material and social environment) and preferences (evaluating the expected consequences of his actions). He combines these three determinants into a 'choice rule' which reflects two distinct forms of rationality (Walliser, 1989). Instrumental rationality refers to the adequacy of pursued objectives to available means and cognitive rationality refers to the adequacy of designed beliefs to available information. Nature, which summarizes the common material environment, is assumed to take states by some passive process, characterized by a 'generation law'. It is assumed to be neutral with regard to the players, taking its states in a deterministic way independently of players' actions.

In the origin (von Neumann-Morgenstern 1944; Nash, 1950), the players were considered to play simultaneously. Their joint actions (or 'profile of actions') lead to common consequences which are evaluated differently by the players. Such a one-shot game is structurally expressed by the modeller under the 'normal form' which makes the determinants more precise. A game matrix is obtained by considering every combination of an action for each player and of a state of nature. In each cell of the matrix, the payoff (or the utility) obtained by each player is indicated (nature is passive and has no payoffs). According to the payoff structure, different types of games are distinguished, illustrated for two-player games without nature. In 'twin games', the players have the very same interests, hence their payoffs are identical in each cell. In 'zero-sum games', the players have contradictory interests, hence their payoffs are opposite in each cell. In 'symmetric games', the players play analogous roles in the sense that they have same sets of actions and, when they exchange their actions, they exchange their payoffs too.

In a normal form game, the two basic notions of ‘dominance’ and ‘best response’ give rise to three main equilibrium concepts. First, for some player, an action dominates another if it gives a higher payoff whatever the other’s action. Hence, an action of a player is dominant if it dominates all other actions of that player and an action is dominated if some other action of that player dominates it. A ‘dominant equilibrium’ state is just a profile of dominant actions. A ‘sophisticated equilibrium’ state is a profile of actions which survive ‘iterated elimination of dominated actions’. In that procedure, the dominated actions of each player are first deleted, the same is done on the remaining game and so on. A dominant equilibrium seldom exists, while sophisticated equilibria are often numerous. Second, an action of some player is a best response to the other’s one if no other action gives a higher payoff to it. A ‘Nash equilibrium’ state is just a profile of actions where each action of some player is a best response to the other’s action. Appearing as a fixed point, a Nash equilibrium (in pure actions) may not exist, be unique or be multiple.

1.2. Introduction of time and uncertainty

Around year 1970 (Selten, 1975), a time dimension was introduced by enlarging the game theory framework. Players were considered to play sequential actions which induce at each move or globally at the end of the game some common consequences. Such a game was structurally expressed by the modeller under the ‘extensive form’ which again makes the determinants more precise. A game tree is obtained by considering at each node the possible moves for the player or nature playing at this node. Hence, one move from one actor conditions the further moves of the other ones. For each path in the tree formed of successive actions and states and leading to a terminal node, the utility obtained by each player is indicated. A special case obtains when the players play a (finite or infinite) sequence of the same basic game. Each player receives then a payoff at each period, these payoffs being aggregated thanks to a discount rate (which depreciates the payoff of each period with regard to the preceding one).

At the same period (Harsanyi, 1967), uncertainty was made explicit in the normal form of a game. It was first related to nature since a player does not know, when he plays, what this state is. Usually, the modeller assumes that nature takes its states randomly, according to some prior probability law, moreover assumed to be known by the players from the beginning of the game. It was further related to the determinants of a player which are badly known by the other one, and even by him. In that case, the modeller assumes that the determinants of a player can be summarized in a ‘type’ of that player, that this type can be treated as a state of nature (nature attributes a type to each player at the beginning of the game), and that a probability distribution exists on the types, again known by the players. Moreover, uncertainty was introduced in the extensive form of a game. It is essentially related with past actions of a player which are not perfectly observed by another one and even by him. The modeller assumes that the nodes of a player associated to the non discriminated actions of another player are gathered in some ‘information set’.

The equilibrium notions were extended in order to integrate time and uncertainty. As for time, it can first be observed that, by introducing the notion of a ‘strategy’ (defined as the action played by a player at each relevant node), the extensive form game can be transformed into a

normal form, hence a Nash equilibrium can be defined. A ‘subgame perfect equilibrium’ is then defined as an action path which is a Nash equilibrium in the game and all its subgames. When the game is finite, it is obtained by ‘backward induction’. This procedure states that a player at a terminal node chooses his best action, a player at a penultimate node chooses his best action, knowing what the next player will do, and so on till the root node. A subgame perfect equilibrium is a refinement of a Nash equilibrium, but it always exists and is generally unique. As for uncertainty, a ‘Bayesian equilibrium’ state is defined as a Nash equilibrium of the game where the players are decomposed in as much agents as possible types and choose their best action, in average on other’s types. A Bayesian equilibrium is a Nash equilibrium of an extended game, and it frequently exists and is even multiple.

3. Introduction of beliefs and learning

Around year 1980 (Aumann, 1976, 1999), beliefs were introduced explicitly in the game theoretical framework. The usual framework formalizes essentially the opportunities and the preferences of the players, respectively as action sets and payoffs. But beliefs do not appear in a formal way, even if observations and even expectations already appear as basic beliefs. Beliefs are representations made by a player about his environment, essentially nature and other agents. In fact, the players have crossed beliefs since each player holds beliefs about others’ beliefs. Moreover, in dynamic games, the players have prior beliefs that they modify when receiving messages. Beliefs were represented semantically in a specific framework which revealed further to be identical to that of epistemic logics, more precisely a possible world framework. Moreover, strong assumptions were implicitly introduced: truth (what a player knows is true), positive introspection (a player knows what he knows), negative introspection (a player knows what he does not know). Hence, beliefs were expressed by ‘information partitions’ in a possible world space. When some world is the actual one, a player ignores in what world of the corresponding partition cell he is.

In the same period, learning rules were introduced and form the basis of a new research program labelled as ‘evolutionist game theory’. The classical program considered that the players are strongly rational and optimize their behaviour. The new framework integrates the fact that the players are boundedly rational and just react to their environment in a utility-improving direction. The models were first based on biological analogies where the players are directed by mutation and selection mechanisms. They consequently abandon all their reasoning faculties and adopt zero rationality. Hence, more psychological features were introduced where the players revise their beliefs or adapt their behaviour rules to past observations. Moreover, random elements were introduced in the players’ encounters, observations, expectations and decisions. In these models, the work of crossed beliefs is replaced by the work of repeated experience. Technically, since the players have a purely reactive behaviour, their actions will follow some stochastic process.

The equilibrium notions were completely revisited by these original views. The main problem about an equilibrium notion is that it is not constructively defined. An equilibrium state is a state such that if the players find themselves in it, they have no (cognitive or preferential) incentive to deviate from it. But the modeller exhibits no concrete process leading the players to such an equilibrium state. Especially, he does not make precise how an equilibrium state is

selected in case of multiplicity. On one hand, when introducing beliefs, it was possible to study under what conditions players come to an equilibrium state by their sole reasoning. These conditions constitute the ‘epistemic justifications’ of the usual equilibrium notions (see § 6.3.). On the other hand, when introducing learning, it was possible to study under what conditions a dynamic process asymptotically leads to an equilibrium state. These conditions constitute the ‘evolutionist justifications’ of the usual equilibrium notions (see § 7.3.).

2. Individual gathering of information

In game theory, players are endowed with beliefs conceived as mental representations of their surrounding environment. These beliefs are considered as imperfect and incomplete with regard to the modeller’s model acting as a reference. They are formalized in the same terms than the model itself, but with truncated relations between partial variables. Information is then conceived as some message received by the player from outside by observation or communication. It gives further details about some actual phenomenon or about some law followed by the system. It is formalized in a very simple way, as a more precise specification of a variable or as a signal correlated to the actual value of the variable.

2.1. Sources of uncertainty

Game theory assumes that the modeller, like God, has a complete and perfect model of the system under study, including the players’ beliefs. This model is represented by a set of generic relations between observable or non observable variables and generates alternative paths of the system. Besides, game theory assumes that each player has the same general ontology than the modeller, as concerns the basic entities and their interactions. But his model is more imperfect, incomplete than the modeller’s one, and possibly even false. Uncertainty is associated to some event of structural feature of the system. It concerns the value of some variable or the specification of some relation, especially some parameter. To compensate for uncertainty, information makes more precise some badly known element of the system. It concerns some statement about the actual value of a variable or the true specification of a relation.

According to the modeller’s model, each player faces nine basic types of uncertainty obtained by crossing two independent criteria about the uncertain features. On one hand, considering the concerned entities, a player may be uncertain about nature (physical uncertainty), about the other players (actorial uncertainty) or even about himself (personal uncertainty). On the other hand, concerning the entities’ attributes, he may be uncertain about past events (factual uncertainty), about atemporal structural features (structural uncertainty) or about future events (strategic uncertainty). More precisely, the (past or future) events considered are the nature’s states or the player’s actions while the (atemporal) structural features are the nature’s generation law or the player’s determinants, generally summarized in players’ types having their own generation law.

Likely, each player is able to receive information from various sources before or during the play of the game. The player has some prior information about the game structure, especially about the other’s determinants. He becomes naturally informed about some past events, since

he observes realized states of nature, watches implemented actions of another player or feels his own utility obtained by a past action. Information received by a player is believed by him with some degree of credibility. But information is distinguished from belief in the following way. Information is considered as a flux coming from the environment of the player by observation or communication. Belief is a stock anchored in the player himself and transformed by outside information or by inside restructuring. Moreover, information is generally an elementary item while a belief is more structured.

2.2. Structure of uncertainty

Uncertainty affected to some (discrete or continuous) relevant variable can be expressed semantically under two main forms. In the probabilistic form, it is expressed by a probability distribution defined over its values. In the set-theoretic form, it is expressed by stating that the variable belongs to some subset of values. The two cases cannot be easily compared, even if certainty appears as a limit case of both (probability 1 on some value or unique value). Probabilities are objective when the modeller considers the phenomenon as really stochastic and when it informs the player of the actual probabilities, the last endorsing them. Such probabilities are frequencies (in a sequence of similar events) or proportions (in a population of objects). Probabilities are subjective when the player forms a personal probability about the phenomenon, whether or not stochastic. Such probabilities are logical (expressing a relation between variables) or decisional (expressing the willingness to bet of a player). The last ones may be revealed by the modeller from player's actions, under strong conditions.

Uncertainty may also be expressed under a hierarchical form, meaning that a player is uncertain either about his own or about other's uncertainty. On one hand, a player may have some second order uncertainty on his first order uncertainty. The first order uncertainty is generally objective (basic uncertainty) while the second order uncertainty is always subjective (ambiguity). Each level may be expressed in a probabilistic or a set-theoretic form, leading to original two-level belief structures: a probability distribution on probability distributions, a probability distribution on sets defining Dempster-Shafer belief functions (Shafer, 1976), a set of probability distributions defining multi-prior belief functions (Gilboa-Schmeidler, 1989). On the other hand, a player may have a second order uncertainty on the first order uncertainty of another player. Here again, the second order uncertainty is a subjective one while the first is of any kind. The levels are generally expressed either both in a set-theoretic form or both in a probabilistic form.

In game theory, uncertainty is expressed in a way adapted to the uncertain element as illustrated by some examples. First, uncertainty of a player about nature's states can be expressed in a more or less precise way (Knight, 1921). Bernoullian uncertainty happens when the player forms a probability distribution over the states. Knightian uncertainty happens when the player just knows the set of possible states without weighting their respective occurrences. Radical uncertainty happens when a player does not even know the set of possible states, due to 'unexpected contingencies'. Second, uncertainty of a player about the other's player type (and even about his own type) is usually expressed in a probabilistic form (the players' types are in fact treated as initial states of nature). Third,

uncertainty of a player about another player's past actions is expressed in a set-theoretic way, by 'information sets' gathering all nodes he cannot discriminate

2.3. Structure of information

Information is generally represented by a 'message' received by a player. Such a message concerns a given variable or parameter or some other variable correlated to it and called a 'signal'. For instance, when checking the existence of gas in the soil, a gas company may search for salt since presence of salt is correlated with the presence of oil. Hence, information is modular in the sense that it appears as a 'psychical quantum' understandable independently of another piece of information or prior beliefs. Information is endowed with a truth value since a direct observation is correct or not while a communicated item is true or not. Information is moreover unambiguously interpreted since it does not depend on its material support or its language. More precisely, it is univocally interpreted by each player in the sense that he knows what variable is concerned. Likely, it is interpreted in the same way by all players since they agree on a structural representation of the system.

Information is expressed in different forms, but it always depends on the actual value of the relevant variable. A message is set-theoretic when it indicates, for each actual value, a subset of possible values. A message is probabilistic when it indicates, for each actual value, a probability distribution over the values. In the last case, the player receives, in a set of possible signals, some specific signal and, knowing the probability of that signal conditional on the actual value, he computes a probability distribution on the real value. When usual properties are attributed to the message (truth, positive introspection, negative introspection), it is defined either as a partition on the set of values or as a unique probability distribution on the values. In many applications, a player receives both public information characterized by a prior objective probability distribution and private information characterized by a subjective partition, specific to each player.

With regard to information, two extreme types of action are considered for any player. An 'operational action' aims at producing some material consequences. An 'informational action' aims at gathering some additional information. In fact, a same action may simultaneously provide information and produce a material impact. More precisely, information can be obtained in three ways, corresponding to different forms of experimentation. First, in exogenous experimentation, a player buys some information at some cost from specialized instances. Second, in passive experimentation, a player obtains information as a by-product of an operational action, generally for free. Third, in active experimentation, a player obtains information by deviating voluntary from some operational action. He loses some utility (or incurs some cost) at short term, compensated by the advantage brought by the information in enlightening operational actions at long term.

3. Individual treatment of information

A player is willing to acquire information since it may help to modify his beliefs in the direction of truth. Information gathering is considered by the modeller as a specific action complementary or substitutive to a material action. The fundamental mental operation on

information is belief revision, which can be implemented for any message, true or false. Belief revision is a prototypical reasoning operation, which is realized according to some rules independently of choice. Moreover, information is helpful if a player ‘knows more’ with his final belief than with his initial one. In order to check this, the modeller is able to formalize the fact that some belief is more accurate than another.

3.1. Individual belief revision

Consider first a single decision-maker who holds an initial belief over states of nature. The decision-maker receives a message on the actual state, which may be compatible or contradict the initial belief. Two main revision contexts are generally considered. In a ‘revising context’, the agent just receives some additional information about a fixed environment. In an ‘updating context’, the player receives information about how an evolving environment has changed. A third context, the ‘focusing context’, happens when a specific instance of the environment is sorted out and information is given about it. Formally, the focusing context can be reduced to a revising context if a ‘projection principle’ is satisfied. For instance, a merchant may sell different products according to future meteorological conditions (tempest, rain, cloudy, little sun, high sun) crudely assessed. In a revising context, he learns that it never rains in the region. In an updating context, he learns that a depression affects globally the meteorological conditions. In a focusing context, he learns that it does not rain this day.

The revision process, conveniently represented by a revision rule in semantics, is sustained by an axiom system in syntax. For set-theoretic beliefs, in a revising context, the AGM axiom system (Alchourron-Gärdenfors-Makinson, 1985) designs a ‘conditioning rule’. In the possible worlds space, nested coronas more and more distant from the first one constituted by the initial belief are defined. The final belief is just the intersection of the message with the first corona intersecting it. Especially, when the initial belief and the message are compatible, the final belief resumes to their intersection. In an updating context, the KM axiom system (Katsuno-Mendelzon, 1992) designs an ‘imaging rule’ in a similar way. When dealing with probabilistic beliefs, the usual revision rule used in game theory is Bayes rule, restricted to the case where initial belief and message are compatible. The posterior probability distribution is obtained by adjusting homothetically the prior distribution to the worlds not excluded by the message. In fact, Bayes rule is relevant only in a revising context and can only be justified by very strong axioms (Walliser- Zwirn, 2002).

Belief revision is strongly related to other reasoning modes attributed to a player in game theory and appears as his central reasoning mode. A correspondence can be established in syntax between their respective axiom systems and in semantics between their corresponding rules. For instance, non monotonic reasoning is isomorphic to belief revision in a revising context (Kraus-Lehman-Magidor, 1990). Likely, abductive reasoning is isomorphic to some reverse belief revision process, always in a revising context (Walliser-Zwirn-Zwirn, 2005). Finally, counterfactual reasoning is isomorphic to belief revision, but in an updating context (Stalnaker, 1968). Other reasoning modes attributed to players remain unrelated to belief revision, for instance analogical reasoning (used in case-based reasoning) or taxonomical reasoning (used for game categorization).

3.2. Collective belief revision

In a game, each player treats his available information in order to reduce sequentially the different types of uncertainty he faces. First, he uses his acquired information in order to reduce factual uncertainty by implementing directly some belief revision process. For instance, when he gets a message about the actual state of nature, he revises his prior belief about the nature's generation law accordingly. Second, he uses his factual beliefs in order to reduce structural uncertainty by implementing some abductive process. For instance, when observing some actions of another player, he tries to reveal the other's preferences (and/or beliefs) by making some case about him. Third, he uses structural beliefs in order to reduce strategic uncertainty by implementing some counterfactual reasoning. For instance, when knowing the other's determinants and assuming that he is rational, he infers more or less precisely what action he will choose in a game tree, even in nodes which may not be reached by the equilibrium path.

Consider now the different players, each of them endowed with a hierarchical belief structure in a set-theoretic form. Such a structure expresses the crossed beliefs of the players about the states of nature ('I know that you know... that p'). The players receive a message, defined both by its content and its status. The 'content' of the message expresses as usual what information is given to the players. For instance, one defines a material message (about the state of nature) or an epistemic message (about a player's belief on the state of nature). The 'status' of the message indicates to whom the message is diffused. For instance, one defines a public message (the message is sent to all players and this is common belief), a private message (the message is sent to one player, the other knowing that the first received a message, but not its content, and all this is common belief) or a secret message (one player receives a message, the other being unaware of this). But a lot of other types of messages are conceivable (private believed public, quasi secret). Such a two-dimensional message can conveniently be expressed by an auxiliary belief structure, called the message structure.

Moreover, a 'specification message' is a message which does not contradict the initial belief, hence does not surprise the player. However, the initial belief and the final belief can include errors, as illustrated by a secret message which turns a true belief into a false one. In the framework of dynamic logics, a multi-agent belief revision rule combines in a precise way the initial belief structure and the message structure in order to obtain a final belief structure (Baltag-Moss, 2004). The syntactic counterpart of the revision rule can be expressed by a few axioms, the main one being some kind of *modus ponens* (Billot-Vergnaud-Walliser, 2006). It assesses that a player believes a proposition in the final belief when he learns a message and believes that the message entails that proposition. The case of a 'rectification message', which contradicts the initial belief, is harder to deal with and needs again to introduce an order between possible worlds.

3.3. Accuracy orders on belief structures

Two belief structures can be compared by defining 'accuracy orders', which express that one structure is more informative than another. For set-theoretic structures, in semantics, stronger and stronger accuracy orders can be defined (Billot-Vergnaud-Walliser, 2006). In semantics, a

relation is defined between corresponding worlds and conditions are expressed on the accessibility relations in these worlds. For instance, a belief structure is collectively more accurate than another if, in two corresponding worlds, the accessibility domain of the first is always included in the accessibility domain of the second. This just means that, in any world, each player considers fewer worlds as accessible in the second than in the first. These accuracy orders have again well defined syntactical counterparts. For probabilistic structures, an accuracy order is less obvious to define. However, a probability distribution can be defined as less accurate than another if it is a ‘mixture’ of it.

A specific order on belief structures can be constructed when these structures are defined in a group of players and concern some material proposition. In syntax, the proposition is ‘distributed belief’ when the players may deduce it by gathering their beliefs. The proposition is ‘individual belief’ when one player at least believes it. The proposition is ‘shared belief’ when all players believe it. The proposition is shared belief at order k when all players believe it, believe that the others believe it and so on till level k . The proposition is ‘common belief’ when it is shared belief at any order (Lewis, 1969). This hierarchical definition of common belief is weaker than a circular definition (Barwise, 1988). The last states that a proposition is common belief if everybody believes that it is common belief. All these operators have well defined semantic counterparts. Especially, in semantics where players have information partitions, the ‘common belief partition’ is the finest partition of the coarsened mutual partitions.

Coming back to belief revision, the various types of messages can be partially ordered according to their relative accuracy. For instance, a public message is collectively more accurate than any other message, a private message as well as a null message. A fundamental result links accuracy orders to belief revision. It states that the accuracy order is preserved when carried from the message to the final belief (Billot-Vergnaud-Walliser, 2006). More precisely, for a given initial belief structure, if some message is more accurate than another, the corresponding final structure is itself more accurate than the other final structure. For instance, the final belief obtained by a public message is collectively more accurate than the initial belief. Such a condition, satisfied by a specification message, precisely states that the (true) message has improved the player’s knowledge.

4. Collective impact of information

Player’s information acts not directly on his selected actions, but only indirectly through his revised beliefs. Since player’s beliefs are affected by uncertainty, the choice rules and equilibrium notions are adapted to it. Information gets evaluated no more essentially with regard to its contribution to truth, but to its utility in decision. The value of information is precisely introduced in order to check if more accurate information is also more efficient. As expected, the answer is positive when considering individual decision-making against probabilistic uncertainty. But surprisingly, it may be negative in a game context since information has a complex impact on crossed beliefs.

4.1. Choice under uncertainty

When a single decision-maker is confronted to nature, cognitive rationality resumes to belief revision and instrumental rationality to utility maximization. For a static choice against nature, the relevant choice rule is the maximization of expected utility. When the law of nature is probabilistic and the decision-maker is aware of these objective probabilities, he computes the choice rule with these probabilities (von Neumann-Morgenstern, 1944). When the law of nature is unknown to him, he applies the same choice rule, but with subjective probabilities (Savage, 1954). For a dynamic choice against nature, the same choice rule is again assumed to be relevant, but associated with a new principle, the ‘backward induction principle’. The decision-maker proceeds in the ‘decision tree’ (the game tree reduced to one player and nature) from the terminal nodes to the root node. At a nature’s node, he computes the expected utility over states endowed with progressively revised probabilities. At a decision-maker’s node, he retains the action which maximizes expected utility.

The choice rules receive cognitive justifications under the form of an axiom system relying on preferences about strategies. This axiom system can be extended from Bernoullian uncertainty to Knightian uncertainty, from static choice to dynamic choice. They receive moreover pragmatic justifications, especially the Dutch book argument stating that if a decision-maker does not maximize his expected utility, he can be confronted to a sequence of choices at the end of which he always loses. They receive finally evolutionist justifications, stating that in competition with others, a decision-maker who does not maximize his expected utility will be eliminated. In other respects, two main interpretations are generally given to the choice rules. The instrumental interpretation considers that the decision-maker behaves as if he maximized expected utility (like a billiard player). The realist interpretation considers that the decision-maker consciously maximizes his expected utility (like a poker player).

When several players are involved in an uncertain and dynamic game, a specific equilibrium notion called ‘perfect Bayesian equilibrium’ is stated. It extends simultaneously a Bayesian equilibrium and a subgame perfect equilibrium. At each information set of the game tree, the decision-maker associates a probability distribution to the constituting nodes (expressing his beliefs about the past moves). The two usual rationality principles are then applied. Instrumental rationality states that, at each node, the player chooses the action which maximizes his expected utility, for future nodes already optimized (backward induction procedure). Cognitive rationality states that the probability distribution, at each node, is adjusted along Bayes rule, with respect to the information gathered about past moves (Bayes conditioning procedure). Hence, uncertainty about the player’s action is treated in the same way that uncertainty about the nature’s state, even if the first is endogenous and the second exogenous.

4.2. Information value in individual decision-making

Consider first a decision-maker who, before choosing an operational action, proceeds to an exogenous experimentation. More precisely, he is endowed with a prior probability distribution about the states of nature and may buy exogenously a partitional message. As for any item, he chooses to buy the message if its (exogenous) cost is less than its value. By definition, the information value brought by the message is the difference of actor’s expected utility before receiving the message and after receiving it. A fundamental theorem (Blackwell,

1951) states that the information value is always positive for a strongly rational decision-maker. The decision-maker cannot be worse after receiving the message than before receiving it. However, the information value may nevertheless be negative in two unusual cases: the decision-maker adopts another choice rule than expected utility maximization, the message is not a partitional one. Similar results are obtained with probabilistic messages.

Consider now that the decision-maker is involved in a sequential choice in which he acquires progressively some information by active experimentation. More precisely, he faces a repeated choice process with an infinite horizon, receives payoffs at each period and aggregates them thanks to a discount factor. He faces then a typical trade-off between exploration and exploitation. Exploration consists in gathering as much information as possible, exploitation consists in using at best the available information. In fact, this trade-off is automatically solved by computing the dynamic optimal choice. Intuitively, he will proceed to much exploration at the beginning of the process and much exploitation at its end. Moreover, when the discount rate increases, he makes more exploration at the beginning of the process. When the discount rate tends to 1, he proceeds only to exploration in the first periods till acquiring the information he wants, then he shifts to pure exploitation.

For instance, consider that the decision-maker is confronted in a casino to a ‘two-armed bandit’. He may use at successive periods one of two levers with a random effect. Each lever gives him a payoff of 1 with a given probability and a payoff of 0 with the complementary probability. He faces structural uncertainty since the probability of winning with each lever is unknown to him. He is only endowed with a second-order probability distribution about the probability of positive payoff. The optimal behaviour can be proved to be a deterministic index rule (Gittins, 1989). In all periods, the decision-maker associates an index to each lever, which depends on the prior probability, on past experiences and on the discount rate. In simple cases, the index aggregates an exploration value and an exploitation value of each lever. Within a period, the decision-maker chooses the lever with greatest index, observes the payoff he obtains and adapts the index consequently. After some delay, the decision-maker always uses the same lever, even if he has a (small) probability of being wrong.

4.3. Information value in games

In a one-shot game involving nature, the players receive from outside a set-theoretic message of any content and status about its actual state. The information value of the message for some player is again the differential utility he gets at some Bayesian equilibrium state computed before and after receiving the message. In fact, different notions of information value, more and more averaged, can be defined and assessed. The actual value is the utility differential really improved by the player, but it is only known by the modeller. The ex post value is the utility differential measured with the final beliefs, and it is computable by the player after he got the message. The ex ante value is the utility differential measured with the initial beliefs, and it is computable by the player before receiving the message, hence allows him to decide to acquire it or not.

Contrary to decision-making against nature, the ex ante information value in games may well be negative for some player (Kamien-Taumann-Zamir, 1990). More precisely, when one of

two players receives a message, all combinations of signs of information value may be obtained: both may become better, both may become worse, the receiver may become better and the other worse, the receiver may become worse and the other better. However, under technical assumptions, the information value is always positive for the receiver of a message for some types of messages and some classes of games. The first case corresponds to a secret message in any game (Neyman, 1991). The player finds himself in a similar position than a single decision-maker. The second case corresponds to a private message in a zero-sum game (Gossner-Mertens, 2001). The player receives a message such that its impact is opposite to the impact of the message on the other player. The third case corresponds to a public message in a pure coordination game. The players act as a team and can only be beneficial from the message.

In a repeated game involving nature, the players are again confronted to the exploration-exploitation dilemma. The uncertain item may be the nature's law, but also the distribution of players' types. The dilemma is far trickier than for an individual decision-maker and receives no general solution, but it can be solved in specific situations. For instance, consider an investor in some product, who is confronted to a high or low demand with some prior probability. His investment can be realized in one step or decomposed in two successive steps, the second option involving an additional cost with regard to the first option, but allowing him to observe the demand after the first step. Hence, at the first period, he can adopt an irreversible option (invest completely) or a flexible option (invest partially, observe the demand and complete the investment if and only if demand is high). It can be shown that, in the choice process, the flexible option has to be given some 'bonus', which is precisely equal to the information value given by the message about the demand.

5. Individual providing of information

If a player may receive information from another player, he may also provide information to another player. Transmission of information is again a voluntary action, but is realized directly or through a material action. Information becomes then a strategic item which is delivered only if it is in the player's interest. Such a phenomenon is precisely studied in specific classes of games with asymmetric information such as 'signalling games'. The sender who has some information has interest to transmit it to the receiver who acts materially only in specific contexts. These results are based on equilibrium concepts which assume that the players are able to simulate each other, hence form self-fulfilling expectations.

5.1. Asymmetry of information

Many concrete situations involve prior asymmetric information between players. Such an asymmetry may be reduced if a player provides his information to a second one by direct transmission or through an action (which reflects more or less some information he has). However, depending on how the second player is expected to use such information, hence on the utility differential it will induce on him, the first player has an interest or not to provide it. When information is directly communicated, he may provide it, abstain to provide it or even distort it. When information transits through an action, he may implement the intended action, render the action fuzzy or even try to transmit biased information. Two situations are usually

distinguished, according to the item concerned by information. ‘Moral hazard’ happens when a player has no interest to publicize the action he implements. ‘Adverse selection’ happens when a player has no interest to diffuse a state of nature he privately knows (for instance his type).

Reflecting the moral hazard situation, the ‘agency game’ considers two players (principal and agent) facing nature (Grossman-Hart, 1983). The agent first performs some action of interest for the principal. Nature provides then a signal related both to that action and to its actual state to the principal. The principal finally observes the signal (and knows its dependence on the action), but not the agent’s action itself, and gives a retribution to the agent according to the signal. The payoff of both players depends on the agent’s action and on the principal’s retribution, and eventually on the signal itself. The problem faced by the principal is to induce the agent to act in his own interest by a ‘contract’ fixing an adequate retribution. In a perfect Bayesian equilibrium state, the agent takes an action which departs more or less from the action he would have taken if this action were observable by the principal.

For instance, on a car insurance market, the insured may or not implement some self-protective action, nature provides a signal as to the occurrence of an accident (which includes a random element) and the insurer just observes the accidents and pays a reimbursement linked to the premium. However, the insurer may adapt the premium and the reimbursement to the number of past accidents or even to variables more correlated to the self-protective actions (driver’s age). Likely, in a firm, the employee is able to modulate his effort rate at work, nature indicates the production of the firm (which also depends on other random factors) and the employer just observes the production and gives a wage to the employee related to the production. Here again, the employer may adjust the wage not only to its final profit, but to variables more correlated with the effort rate.

5.2. Signalling games

Reflecting the adverse selection situation, a ‘signalling game’ considers two players (sender and receiver) facing nature (Rasmusen, 1989). Nature first defines a state according to some probability distribution which is common belief. The sender observes the actual state and sends one of two signals to the receiver about that state, eventually a mixed one (probability distribution over the pure signals). The receiver observes the signal, but not the state of nature, and implements an action, eventually a mixed one (probability distribution over the pure actions). The payoff of both players depends in general on the state, on the signal as well as on the action. However, in ‘cheap talk’, the players’ payoff does not depend on the signal, hence each player may talk freely to the other by expressing what is in its interest. Finally, the receiver revises his beliefs (the probability of the state conditional to the signal) according to public information.

Two contrasted types of perfect Bayesian equilibrium states may appear. In a ‘separating equilibrium’, a different signal is transmitted by the sender in each state of nature, hence the receiver is able to reveal from the signal the actual state. The sender has an interest to transmit his private information and the receiver learns it perfectly. In a ‘pooling equilibrium’, the same signal is transmitted by the sender for all states of nature, hence the receiver is not able

to reveal anything. The sender has no interest to transmit his information and the receiver learns nothing. Some hybrid equilibrium states may also be available for which the message transmitted by the sender is a mixed one. The equilibrium state which actually occurs is conditioned on the main parameters of the problem (probability of states, cost of actions). For given values of the parameters, one or more equilibrium states may happen simultaneously.

For instance, on a health insurance market, nature defines the actual risk of illness of the insured, the insured gives a signal to the insurer in form of the degree of insurance he wants to buy and the insurer fixes the insurance premium according to the cover degree. In a separating equilibrium, the insurer is able to differentiate high risk insured (asking for full cover) and low risk insured (asking for partial cover). Again, the insurer may condition the treatment reimbursement to variables correlated to the health (age, gender). Likely, on a car market (Akerlof, 1976), nature fixes the quality of the car (with a probability commonly known), the sender knows the quality of the car and proposes some price, the buyer observes only the price and accepts or not to transact. In many cases, only a pooling equilibrium happens, the one where nobody transacts. However, when the price is exogenously fixed, a transaction always takes place when information about the quality of the car stays private, but fails when it becomes public. This is a specific case of negative value of information for both players.

5.3. Self-fulfilling expectations

The previous examples show that the players make some expectations about the others' behaviour and these expectations are realized at the equilibrium state. In fact, such self-fulfilling expectations can happen at two levels. At the first level, they concern directly some action of another player. By definition of an equilibrium, these expectations have to be fulfilled at the equilibrium state. At the second level, they concern a relation between some structural variable (especially the other's type) and an action. This relation has again to be satisfied at the equilibrium state, but it may not hold out of equilibrium. However, in both cases, no process is described showing how the expectations are computed by the players, hence how the equilibrium state is concretely achieved. Moreover, if many self-fulfilling expectations are possible, hence if many equilibrium states are possible, no process describes how one is selected.

A self-fulfilling expectation, in its structural form, is made explicit in the 'job market game' (Spence, 1973). Nature attributes to an employee a strong or weak (exogenous) ability, acting as his type. The employee may acquire a high or low education level at a cost which is inversely proportional to its ability. The employer gives a wage to the employee according to his assessment concerning his type. It is fixed according to a prior belief which relates causally the ability of the employee to his education level. More precisely, the employer believes that a highly educated employee acquires a strong ability and conversely. In fact, a reversed causality actually holds since, for the modeller, education follows from ability. If the equilibrium is separating, such a belief becomes self-fulfilled. The belief considered to be true by the employer contributes to realize what it asserts (except for the direction of causality).

Two more illustrations can be given in a dynamic setting. First, the 'simplified poker game' exemplifies the notion of 'bluff'. Nature distributes a high or low card, the first player stakes

or not, the second player asks to see or not (if the first stakes). Of the two forms of bluff theoretically possible for the first player (to stake with a low card, not to stake with a high card), only the first appears at equilibrium. Second, the ‘repeated entry game’ exemplifies the notion of ‘reputation’ astutely formalized. Nature defines the hard or soft type of a monopolist, the incumbent enters or not, the monopolist is aggressive or pacific (if the incumbent enters). In a one-shot game, a soft monopolist is always pacific while a hard one is always aggressive. In a repeated game, even a soft monopolist may be aggressive in order to acquire a reputation of being hard. At equilibrium, the monopolist is aggressive in the first periods and the incumbent keeps out, but later the monopolist becomes pacific with some probability and the incumbent enters as soon as the monopolist is no more aggressive once.

6. Collective diffusion of information

At a collective level, communication of information is the main device able to ensure efficient coordination between the players. An equilibrium state appears no longer as an equilibrium in actions, but becomes an equilibrium in beliefs. A first question is whether some private information becomes shared belief and even common belief through their interactions. It is answered in classical puzzles which were developed outside game theory, but are easily reinterpreted in game theory. A related question is whether hyper-intelligent players can coordinate on some equilibrium state by their sole reasoning. Unexpectedly, if some equilibrium notions can easily be justified, this is not the case for Nash equilibrium.

6.1. Communication between players

The fundamental question is how information diffuses among players, with regard to their more or less convergent interests. Assume that the information of players concerns only states of nature. As usual, each player is endowed both with public information (prior probability distribution on states) and private information (information partition, signal correlated to the actual state). He combines these two sources of information in a Bayesian way. Moreover, the players exchange sequentially some information either by direct communication or through their actions. A first problem is to examine if they achieve asymptotically a shared belief which gathers in some sense all their private beliefs (homogenization of beliefs). A second problem is to examine if this shared belief becomes even a common belief (homogenization of crossed beliefs). However, even for similar agents, their actions may become identical even if their beliefs are not.

A general result about communication is the ‘not agreeing to disagree’ theorem (Aumann, 1976). In a dynamic version, two actors share a common prior probability over a set of states, receive initially some private information about the nature’s state and announce sequentially and publicly their posterior probability about some specific event. The players get no specific utility from their announcement, which means that their announcements have no strategic component. It can be shown that their beliefs converge in a finite number of steps to a common posterior probability of the event. The result follows from the pre-coordination of the players by their common prior probability, reflecting some common culture. To be sure, the result no longer holds when the players have different priors due for instance to different past experiences.

An application of the preceding theorem is the ‘no-trade’ result (Milgrom-Stokey, 1982). Consider two risk-averse agents who are able to trade together, an exchange contingent on the actual state of nature. They share a common prior probability and receive partitional private information over the states of nature. The players have opposite interests with two more restrictive conditions. At a Bayesian equilibrium state, it appears that no trade will actually take place. The reason is that any player thinks that if the other is willing to trade, he possesses some private information in his favour. Hence, this information is in his disfavour and he will abstain. However, if the players do not share a common prior (for instance, if one has initially an optimistic belief and the other a pessimistic one about the state of nature), trade may take place.

6.2. Usual puzzles

In the ‘three hats problem’ (), three boys have a white or red hat on their head, actually all three are red. Each boy observes the others’ hat but not his own one. At successive periods, he has to say if he knows the colour of his hat. Before the initial period, an observer says that one hat at least is red, such information being already shared knowledge but becoming then common knowledge. Each boy gets a positive utility if he is right, a zero utility if he does not answer and a negative one if he is false; hence, his utility function depends only on his own action. In a Bayesian equilibrium state, each player gives no answer at the first two periods and answers rightly at the third (as can be shown by iteration on the number of boys). In that case, the players converge towards a common belief about the colours of their hats. Technically, the possible worlds are finite, since they are materially constituted by the possible combinations of hats and associated beliefs about them. With finite worlds, shared belief increases from one level at each period and necessarily becomes common belief.

In the ‘Byzantine generals problem’ (Rubinstein, 1989), two allied generals have to choose to attack or not a common enemy. One general observes the situation which may be good or bad (with some probability). If the situation is good, he sends a message to the other, but the message has a small probability of being lost. Hence, the second general sends a confirming counter message which has the same probability of being lost and so on. The payoffs are such that a general gets a high disutility when attacking alone, and gets a smaller disutility when they attack together in a bad situation. In a perfect Bayesian equilibrium state, the generals never attack. In fact, even if at least one message is sent in both ways, the shared belief that the situation is good never becomes common belief. Technically, the game involves an infinite number of possible worlds, expressing either that the situation is bad or that the situation is good and n messages exactly arrived. Since common belief is here necessary for a joint attack, this never happens. But some conventions (attack if two messages are sent) make nevertheless an attack possible.

In the ‘two restaurants problem’ (), two restaurants where one is slightly better than the other are situated in a same street. The customers come sequentially in order to choose a restaurant. Public information consists in a common prior probability distribution over the quality of the restaurants. Private information gives to each customer a signal correlated with the relative quality of the restaurants. Additional (public) information comes from the observation of the

previous choices of the customers. The payoffs just give a better utility to the best restaurant, hence consider the customers as payoff-independent without externalities. In a perfect Bayesian equilibrium, after some period, all customers go to the same restaurant, with (small) probability that it is the wrong one. The fact that one restaurant is the best one becomes common belief, even if it may be the wrong one. Technically, the possible worlds reduce to two for their material parts, but the beliefs about them are here truly probabilistic.

6.3. Epistemic justifications of equilibria

An extended reasoning process followed by hyper-intelligent players leads to ‘epistemic justifications’ of static equilibrium notions (Walliser, 2006). With the only assumptions of common knowledge of the game structure and of the players’ rationality, the relevant equilibrium notion is ‘sophisticated equilibrium’, obtained by iterated elimination of dominated strategies. With one additional assumption that players play independently, the relevant notion is ‘rationalizable equilibrium’, where each strategy is a best response to others’ strategies, considered as best responses, and so on. With another additional assumption that the players have a common prior, the relevant notion is ‘correlated equilibrium’, where a correlator chooses probabilistically an issue of the game and indicates to each player what it should do. Surprisingly, the two alternative assumptions taken together are not enough to justify a Nash equilibrium. To obtain it, it must moreover (rather heroically) be stated (Aumann-Brandenburger, 1995) that the strategies of the players become shared belief (for two players) or common belief (for more players).

Dynamic equilibrium notions are apparently easier to justify cognitively. For an extensive form game without uncertainty, a ‘subgame perfect equilibrium’ obtains when common knowledge of rationality applies at each node. The problem is that a player may observe a deviation from the equilibrium path during the play of the game and needs to interpret it. According to a standard result (Aumann, 1995), a deviation just cannot happen under the main assumption and the subgame perfect equilibrium is perfectly justified. However, when such a deviation is counterfactually considered as possible, a player may wonder what assumption sustaining the equilibrium is not satisfied (Reny, 1993; Binmore, 1997). The relevant equilibrium notion depends on that precise assumption. For instance, considering the ‘trembling hand’ assumption (any player may deviate from the intended action with some exogenous probability) preserves the subgame perfect equilibrium (Selten, 1975). At the opposite, lack of common belief of rationality enlarges the set of possible equilibria.

In both cases, the results were obtained by introducing more and more epistemic logics in classical game theory. In that respect, the state space of the system, which already included the nature’s state, has to be extended to the players’ strategies. Conversely, the selection of some equilibrium state in case of multiplicity is treated in a more informal way. Some ‘selection principles’ are exhibited which are more or less homogenous with the former ‘implementation principles’. A first path is to consider that some states are ‘culturally’ salient hence are conjointly selected as ‘focal points’ (Schelling, 1960). But salience refers to cultural traits which are not included in the game structure and are essentially context-dependent. A second path is to assume that some selection rules are grounded on various properties of equilibrium states (Pareto optimality, simplicity, symmetry) and act as common ‘conventions’

among agents. But the origin of such conventions is not made explicit and may be history-dependent.

7. Information and learning

On one hand, a player gathers limited information since he faces search costs when he voluntary acquires it. On the other hand, a player is endowed with bounded rationality since he faces computation costs when he treats it. Hence, learning models are introduced where the players supply for the lack of cognitive capacities by repeated experience. However, if the strong rationality model is unique, there exists a large spectrum of bounded rationality models. The strategic dimension of game theory is lost since the players just react to past information without expectation loops. Nevertheless, the usual equilibrium notions can easily be justified as asymptotic states of such processes.

7.1. Bounded rationality

Players are more and more considered as endowed with bounded rationality (Rubinstein, 1994), related to limited capacities of an actor for treating information (Simon, 1982). Bounded rationality was initially associated with instrumental rationality. A first primitive model is the ‘satisficing model’ (Simon, 1957) where an actor chooses the first action which entails results above some aspiration levels on partial criteria. A second primitive model is the ‘stochastic choice model’ (Luce, 1959) where an actor chooses an action with a probability proportional to its utility. But it is difficult to state precisely the link of these rules with limited reasoning capacities. Hence, bounded rationality is better associated with cognitive rationality. Two further models inspired by Artificial Intelligence are the ‘automaton model’ (an actor computes his intended action with a finite number of inner states) and the ‘complexity model’ (the actor faces complexity constraints in his computation). More subtly, limited reasoning may be related to some violations of the cognitive rationality axioms. Especially, the actor may lack ‘logical omniscience’ in the sense that he is not able to deduce all consequences of what he knows.

Of course, the usual equilibrium notions may easily be extended to boundedly rational agents keeping the idea of a fixed point of players’ actions. But bounded rationality is more naturally expressed in learning processes at work in situations where the same players play sequentially the same basic game (a static or a dynamic one). Learning just means that a player behaves by adjusting his actions to what he observed in the past in order to perform better. Due to limited information and bounded rationality, he no more takes into account the strategic dimension of the game. He generally assumes that the other players are not influenced by his action and have even a stationary strategy (even if he knows that they learn as he does). Moreover, he holds little prior structural information and relies essentially on factual information. His behaviour rule is globally fixed and only his action may change, but it may change even if he is faced to the same situation. Of course, a second order learning process may act on the behaviour rule, but the two levels can often formally be collapsed into one.

Five principles are involved in evolutionist game theory, only the first being common to classical game theory (Fudenberg-Levine, 1998). They all introduce some form of

randomness. First, the ‘utility principle’ indicates that each player has a given action set and receives at each period some utility from any combination of players’ actions. Second, the ‘interaction principle’ describes what players meet at each period, the matches being situated in some ‘interaction neighbourhood’ and involving some ‘encounter randomness’. The ‘information principle’ describes the information gathered by a player, such an information being limited to some ‘information neighbourhood’ and implying some ‘sampling randomness’. The ‘evaluation principle’ describes how a player treats his past information in order to build ‘indices’ for the future, eventually introducing some ‘computation randomness’. The ‘decision principle’ indicates how a player uses the preceding indices in order to obtain the chosen strategy, eventually introducing some ‘behaviour randomness’.

7.2. Basic learning processes

A first learning process, ‘belief-based learning’, is grounded on a belief revision procedure. Each player first observes the other’s past actions. He then revises his belief about the other’s behaviour and forms an expectation about the other’s future action. In doing so, he assumes that the other follows a stationary mixed strategy. He finally takes the action which maximizes his expected utility, hence ensuring exploitation behaviour. In order to add exploration behaviour, he may alternatively and with a small probability implement a random action. Especially, in the ‘fictitious play model’, a player considers the past frequency of the other’s action, transforms it into a future probability of other’s action and implements a best response to it. For instance, on the road, a driver observes if the others drive more often right or left and adopts the side followed by the majority.

A second learning process, ‘reinforcement learning’, is grounded on reinforcement of best actions. Each player first contents with observing the past utility obtained with its own actions. He then computes a ‘performance index’ for each action by aggregating the past utilities. In doing so, he assumes that the performance of each action is stationary, even if it concretely evolves. He finally chooses an action with a probability increasing with the past performance. This rule already incorporates the exploration-exploitation dilemma since the players use more and more the best performing actions without eliminating totally any other one. Especially, with the ‘basic reinforcement model’ (Roth-Erev, 1995), a player computes the cumulated utility obtained by each action and chooses an action with a probability which is proportional to that index. For instance, on the road, a driver observes the accidents he had when driving right or left and drives on the side with the less accidents.

A third process, an ‘evolutionary process’, is no more a learning process, but shares some features with it even if inspired by biology (Weibull, 1995). Each player belongs to a subpopulation of players using the same fixed (pure or mixed) strategy and interacts randomly with players from another subpopulation or from the whole population. He observes nothing (except in order to implement his strategy) and he even computes nothing. He reproduces according to his utility, assimilated to the biological notion of ‘fitness’ and this ensures exploitation. Moreover, some mutants in small quantity may be introduced randomly in the population in order to ensure exploration. Especially, in ‘replicator dynamics’, the players of some subpopulation reproduce proportionally to the average utility they get from their

interactions, without mutation. For instance, on the road, a driver dies if he meets another one driving on the other side and duplicates if he meets another one driving on the same side.

7.3. Evolutionist justifications of equilibria

Of course, these learning or evolution rules can be mixed in different ways. Different players may use different rules, a same player may adapt the rule to the context or to the present stage of a game, the rules may be combined in hybrid ones. Moreover, it is easy to observe that the evolutionary process is isomorphic to reinforcement learning. The proportion of players playing some action is replaced by the probability that a player chooses some action. Hence, if an evolutionary process based on biological analogies was historically an important framework to produce original results, it is more and more replaced by the two more realistic learning processes. However, since these learning processes appear themselves as too passive and past oriented, learning processes tend to endow the players with more elaborated cognitive activity (categorization of the choice frame, analogical reasoning in the choice process).

The modeller may be interested by the transitory behaviour of the learning system since it is such a transitory behaviour which is essentially observable. Nevertheless, he is essentially interested by its asymptotic behaviour, characterized by various notions of convergence. The modeller looks for insights not only about the attractors of the process, but also about its speed of convergence. He studies not only the final distribution of strategies, but the emergence of some spatial or qualitative regularities (segmentation of the agents, construction of permanent links). The main problem is that the learning processes are numerous and lead to dispersed results, especially influenced by the type of randomness introduced. The only result which is valid for almost all processes is their capacity to eliminate the (strongly) dominated strategies.

The asymptotic properties of the learning processes provide an ‘evolutionist justification’ to the equilibrium notions when it can be shown that the system converges to a corresponding equilibrium state. For a basic static game, many processes lead to some Nash equilibrium, in mixed strategies and even in pure strategies. Some processes even lead to ‘refinements’ of Nash equilibrium. For a basic dynamic game, many processes lead to a subgame perfect equilibrium. However, some processes may converge towards other states, for instance Pareto-optimal states which are not equilibrium states. A good news is that the selection problem is no more relevant in this dynamical view. The system always evolves and, if it converges, it converges towards a well-defined equilibrium state (at least probabilistically). But this state not only depends on the initial conditions, but also on the history of the game (‘path-dependency’).

* *
*

Game theory was able to internalize information and beliefs and derive strong statements by making minimal and empirically naïve assumptions about them. Information is just an item

which makes the player's belief more precise with regard to the modeller's model. However, game theory imported two frameworks from outside, namely probability theory and epistemic logics which helped to give a formal account of information and beliefs. From these, original consequences could be deduced as concerns the value of information or the diffusion of information. Conversely, other developments like the Shannon theory of communication or the Kolmogorov theory of complexity were ignored. The reason is just that it was not possible to deduce interesting conclusions from them. However, game theory stays open to new tools, for instance able to deal with the 'meaning' of information or to the 'acceptance' of beliefs.

In its applications, game theory was used to explore other phenomena involving heavily information. For instance, in economics, various institutions are justified through game theory by informational principles. Since Hayek (1973), the competitive market is seen essentially as providing prices which are a good summary of what agents have to know about scarcity and desirability of goods. Some institutions sustaining the market may help when information is imperfect or incomplete, especially trust or money. Finally, some institutions are able to replace the market by providing a more local treatment of information, such as auction mechanisms. In other respects, in linguistics, speech acts are analyzed in a game-theoretical framework, either classical or evolutionary (Rubinstein, 2000; Benz-Jaeger-van Rooy, 2005). They deal with syntactic structures, evolution of meaning or pragmatic communication.

References

- Akerlof, G. (1970): The market for lemons: quality uncertainty and the market mechanism, *Quarterly Journal of Economics*, 84(3), 488-500.
- Alchourron, C.E. - Gärdenfors, P. - Makinson, D. (1985): On the logic of theory change: partial meet contraction and revision functions, *Journal of Symbolic Logic*, 50, 510-530.
- Aumann, R.J. (1976): Agreeing to disagree, *The Annals of Statistics*, 4(6), 1236-39.
- Aumann, R.J. (1995): Backwards induction and common knowledge of rationality, *Games and Economic Behavior*, 17, 138-146.
- Aumann, R.J. (1999): Interactive epistemology, *International Journal of Game Theory*, 28(3), 263-314.
- Aumann, R. - Brandenburger, A. (1995): Epistemic conditions for Nash equilibrium, *Econometrica*, 63 (5), 1161-80.
- Aumann, R.J. - Hart, S. (1992, 1994, 2002): *Handbook of game theory with economic applications*, 3 volumes, Elsevier.
- Baltag, A. - Moss, L.S. (2004): Logics for epistemic programs, *Synthese*, 139(2), 165-224.
- Barwise (1988): Three views of common knowledge, in M. Vardi ed., *Proceedings of the TARK Conference*, Morgan Kaufmann.

- Benz, A. - Jaeger, G. - van Rooij, R eds.(2005): *Game theory and pragmatics*,
- Billot, A.- Vergnaud, J.C. - Walliser, B. (2006) : Multiplayer belief revision and accuracy orders, *Proceedings of the LOFT Conference*.
- Binmore, K. (1987): Modeling rational players, *Economics and Philosophy*, 3, 9-55 ; 4, 179-214.
- Binmore, K. (1997): Rationality and backward induction, *Journal of Economic Methodology*, 4, 23-41.
- Binmore, K. - Brandenburger, A. (1990): Common knowledge and game theory, in *Essays in the Foundations of Game Theory*, Blackwell, 105-150.
- Blackwell, D. (1951): Comparison of experiments, *Proceedings of the second Berkeley symposium on mathematical statistics and probability*, University of California Press, 93-102.
- Foray, D. (2004): *The economics of knowledge*, MIT Press.
- Fudenberg, D. - Levine, D. (1998): *The theory of learning in games*, MIT Press.
- Gilboa, I. - Schmeidler, D. (1989): Maxmin expected utility with non unique prior, *Journal of Mathematical Economics*, 18, 141-153.
- Gittins, J. (1989): *Multi-armed bandits allocation indices*, Wiley.
- Gossner, O. – Mertens, J.F. (2001): The value of information in zero-sum games, mimeo.
- Grossman, S. - Hart, O. (1983): An analysis of the principal-agent problem, *Econometrica*, 51 (1), 42-64.
- Harsanyi, J.C. (1967): Game with incomplete information played by Bayesian players, *Management Science*, 159-82, 320-34, 486-502.
- Kamien, M.- Taumann, Y.- Zamir, S. (1990): On the value of information in a strategic conflict, *Games and Economic Behavior*, 2, 129-53.
- Katzuno, A. - Mendelzon, A. (1992): Propositional knowledge base revision and nonmonotonicity, in P. Gärdenfors ed.: *Belief revision*, Cambridge University Press.
- Knight, F. (1921): *Risk, uncertainty and profit*, Kelley.
- Kraus, D. - Lehmann, D. - Magidor, M. (1990): Non monotonic reasoning, preferential models and cumulative logics, *Artificial Intelligence*, 44, 167-208.

Lewis, D. (1969): *Conventions: a philosophical study*, Harvard University Press.

LOFT (1994, 96, 98, 2000, 02, 04, 06): Proceedings of the Conferences ‘Logic and the foundations of game and decision theory’.

Luce, R. A. (1959): *Individual choice behaviour,:a theoretical analysis*, John Wiley.

Macho-Stadler, I. – Perez-Castillo, D. – Watt R. (2001): *An introduction to the economics of information: incentives and contracts*, Oxford University Press.

Milgrom, P. – Stokey, N. (1982): Information, trade and common knowledge, *Journal of Economic Theory*.

Nash, J. (1950): Equilibrium points in N-person games, *Proceedings of the National Academy of Sciences (USA)*.

Neyman, A. (1991): The positive value of information, *Games and Economic Behavior*, 3, 350-55.

Osborne, M.J. - Rubinstein, A. (1994): A course in game theory, MIT Press.

Rasmusen, E. (1989): Games and Information, Blackwell.

Reny, P.J. (1993): Common belief and the theory of games with imperfect information, *Journal of Economic Theory*, 59, 257-274.

Roth, A. - Erev, I. (1995): Learning in extensive form game, *Games and Economic Behavior*, 8, 164-212.

Rubinstein, A. (1989): The electronic mail game: strategic behavior under almost common knowledge, *American Economic Review*, 79(3), 385-391.

Rubinstein, A. (1994): *Models of bounded rationality*, MIT Press.

Rubinstein, A. (2000): *Economics and language. Five essays*, Cambridge University Press.

Savage, L.J. (1954): *The foundations of statistics*, John Wiley.

Schelling, T. (1960): *The strategy of conflict*, Harvard University Press.

Selten, R. (1975): Reexamination of the perfectness concept for equilibrium points in in extensive games, *International Journal of Game Theory*, 4(1), 25-55.

Shafer, G. (1976): *A mathematical theory of evidence*, Princeton University Press.

Simon, H. (1957): *Models of man, social and rational*, John Wiley.

- Simon, H. (1982): *Models of bounded rationality*, MIT Press.
- Spence, A.M. (1973): Job market signalling, *Quarterly Journal of Economics*, 87(3), 355-74.
- Stalnaker, R.C. (1968): A theory of conditionals, in N. Rescher ed. *Studies in Logical Theory*, Blackwell.
- Stalnaker, R. C. (1996): Knowledge, belief and counterfactual reasoning in games, *Economics and Philosophy*, 12, 133-163.
- TARK (1986, 88, 90, 92, 94, 96, 98, 2001, 03, 05): Proceedings of the Conferences 'Theoretical aspects of reasoning about knowledge', Morgan Kaufmann.
- von Hayek, F. (1973): *Law, legislation and liberty*, University of Chicago Press.
- von Neumann, J. – Morgenstern, O. (1944): Theory of games and economic behaviour,
- Walliser, B. (1989): Instrumental rationality and cognitive rationality, *Theory and Decision*, 27, 7-36.
- Walliser, B. (2006): Justifications of game equilibrium notions, in R. Arena and A. Festre eds. *Knowledge, beliefs and economics*, Edward Elgar.
- Walliser, B. - Zwirn, D. (2002): Can Bayes rule be justified by cognitive rationality principles? , *Theory and Decision*, 53, 95-135.
- Walliser, B.- Zwirn, D.- Zwirn, H. (2005): Abductive logics in a belief revision framework, *Journal of Logic, Language and Information*, 14, 87-117.
- Weibull, J. (1995): *Evolutionary game theory*, MIT Press.

