HOMO LUDENS:*
On the play-element in inductive logic

Alessandro Agostini
Divisione SRA, ITC-Irst

August 10, 2001

Abstract

As multi-agent systems are becoming more and more important, my strong belief is that it is important to develop a theory of learning agents in strategic processes, a theory that would explain how an agent's beliefs about the environment (which includes the behavior of others insofar as it affects him) evolve until they have come to agree with the actual properties of the environment. This opens new perspectives on the interaction of game theory and learning theory. Modeling asymmetries in knowledge, in abilities and in perceptions of a cooperative situation by different agents is a fascinating challenge for future research, which models of 'bounded rationality' have begun to tackle. I am convinced that inductive algorithms will become increasingly important for the development of multi-agents systems. Applications range from economics to social and management sciences and software (computable) agents that interact by negotiation with other agents in the World Wide Web.

1 Motivations

My main scientific interest is in the modeling of complex dynamic interactions of strategic decision makers: 'rational' and autonomous agents who repeatedly interact with their environment. To investigate the fundamental challenges of such complex dynamics and systems, I am particularly interested in inductive methods and their application to dynamic processes like coordination, cooperation, delegation, negotiation and teamwork. The key feature of coordination and the like dynamics is that these embrance all the most fundamental characteristics of strategic interactions in a large spectrum of applications and areas involving both individual and coalitional decision making. To model coordination processes it is required modeling the agents' abilities, knowledge and beliefs as well as the environment, which is possibly made unpredictable by the agents behavior itself. This makes coordination and similar strategic dynamics ideal paradigms for modeling, developing and testing general techniques and architectures needed in distributed

*Presented at the 3rd MRG/SRA General Symposium, Levanto (SP), Italy, June 25-27, 2001. Electronic correspondence to: agostini@irst.itc.it multi-agent systems and, we believe, in the emergent field of distributed knowledge management (see e.g. [BBM00]). Regarding these techniques, I am convinced that inductive methods are a very promising approach to handle the learning requirements such systems have to deal with. Learning is a central theme of research in understanding coordination processes, since to coordinate a decision maker has often to discover the strategy adopted by all agents concerned. In this context, I focus on developing inductive representations and algorithms that allow 'rational' decision makers following those representations and algorithms to reliably and efficiently coordinate in several, different strategic situations.

1.1 Games

For everybody interested in the study of communication and learning among 'rational' decision-makers, a game-theoretical standpoint is a challenging, intuitive and rigorous way to start from. Learning is a cognitive process, a "mental action" so well expressed by playing a game. Games involve competition, death and victory, deception and frustration, emotions—exactly as in real life. A game is a natural learning environment, a constructive situation, a "cognitive laboratory for studying deliberation, action, communication and information flow" [vB98, p. 7]. Playing is a basic activity in most cultures [Hui49], so a game is always a social laboratory also. There is, moreover, a rich literature and some interesting tools on games as explanatory paradigms of concepts from logic, language and computation. From the logic side, e.g. [Doe96, Hod97, EF95] give three different ways to present the concept of 'elementary equivalence' in model theory in terms of Ehrenfeucht games. Other games in linguistic include semantic evaluation games (see for instance [HS97] for a survey), dialogue games for validity, and naming games—kind of language games first introduced by Steels [Ste96]. In dialogue games, the validity of some given formula is examined in terms of a two person, perfect information game. Naming games are interactions between two agents, a speaker and a hearer, in which the speaker identifies an object using a name, and the hearer eventually agrees on the name as appropriate for the object identified.[1] A game-theoretic approach in computation is widely used e.g. in game semantics (see for instance introductory [Abr97, AM98]) and program synthesis (see [NYY92] for background and references). More recently, games have been used in modeling pushdown processes and automata, in particular to study the complexity of the model-checking problem [Wal01].

A game is of *coordination* (also "pure coordination") if players's interests perfectly coincide.[2] In a coordination game, each person involved is better off if everybody makes the same decision, and in a coordination equilibrium no one would have been better off if any one player had decided differently. More gener-

[1] For the semantic evaluation games, Tarski's World provides a simple game that students can use when a sentence evaluates in a way they did not expect. See [http://www.csli.stanford.edu/hp/Tarski2.html].

[2] More technically, this means that players have the same payoffs in every cell of the matrix-game. Background material on classical game theory can be found in many textbooks and handbook collections, such as *e.g.* [Bin92, OR94, Osbng].

ally, in such *common interest* games,if E_1 and E_2 are strict Nash equilibria, either E_1 Pareto dominates E_2 or E_2 Pareto dominates E_1.[3] If players could coordinate, they would presumably coordinate on the Pareto-dominant equilibrium, say E_1. In the absence of explicit communication, however, it is not obvious that E_1 will be played. For example, suppose that two people wish to go out together to a concert of music by either Bach or Händel and that they agree on the more desirable concert.[4]

	Bach	*Händel*
Bach	(2,2)	(0,0)
Händel	(0,0)	(1,1)

Figure 1: A coordination game.

Representing the individuals' preferences by payoff functions, we have the game in Figure 1. The game models a situation in which players wish to coordinate on the basis of a mutual interest in reaching a common goal. The game has two Nash equilibria: (*Bach, Bach*) and (*Händel, Händel*). The classical notion of Nash equilibrium, however, does not rule out a coordination state in which the outcome of the game is the inferior equilibrium (*Händel, Händel*).

The foregoing example suggests that players are unlikely to base their decisions entirely on deductions from rationality, and highlights the importance of gathering information from other sources about how coordination problems are solved. In this article, we focus on some coordination games where the agents always move simultaneously and all relevant moves are made by the agents (agents' moving components are functions in mathematical sense; no randomness ever intervenes). These games are all based on pairwise communication, so that the kind of models we discuss is suitable for modeling group situations where communication is not with the whole group, as in an auction, but pairwise, as happens for instance in commercial transactions. Unless the gratuitous assumption is made that agents have a priori common knowledge of their respective beliefs, we deduce that if the agents' beliefs are common knowledge, then an *equilibrium* might obtain by inductive reasoning. Despite we present a glimpse of the formal setting of one of our paradigms of coordination and some sample results, this article should be read as a first attempt to address general work in the formal framework of inductive inference [Ago01] to theory and application in knowledge management and data mining (see for instance [CJLZ99, FPSSU96, FPSS96, BA96, BBM00]).

[3]Let X be a set and suppose $a, b \subseteq X^n$. (X^n denotes the set of vectors of n elements of X.) Let $\preceq$ be a total order on X. Then a **Pareto dominates** b if $b_i \preceq a_i$ for all $i \leq n$.

[4]A different game where two people have conflicting interests is advanced [LR57] and discussed *e.g.* in [OR94].

2 Sam and Sally's coordination problem

> "Sam and Sally like to meet daily in the park, pretending each time that it's yet another chance encounter, walking side by side in shy silence. Each shows up punctually at either noon or 6:00 p.m. hoping the other will have made the same choice. The shifting constraints on their schedules, however, make it hard to predict who will select which time of arrival, and both suffer disappointment when there is mismatch. So both Sam and Sally set about trying to predict the other's choices, desiring to act in concert. Their predictions are based on no more than the history of earlier events...." [MO99b, p. 363]

The game is for "learning to coordinate choices" and of common interest: two agents ("players") want to coordinate by repeatedly showing each other one of two possiblebehaviors. Each player tries to predict the other's behavior, and their predictions are based on no more than the history of earlier events. One player 'learns' the other's behavior just in case his or her own behavior matches the other's forever after. The players' learning is symmetric. To keep matters simple, players face the same two options on each trial, conventionally denoted: 0, 1. A player is therefore a "bit agent", namely, any function from the set of all finite binary sequences into $\{0, 1\}$, where any such sequence is conceived as the history of moves of an opposing player. From a sequence of length n, a player can reconstruct the $2 \times n$ binary matrix that includes his own responses through move n. We call the resulting paradigm: the 01-**coordination paradigm**. It illustrates features of the more complex paradigms of coordination. (See [Ago01] for a discussion on some of these paradigms.)

A criticism

As a model of coordination, the 01-coordination paradigm is limited for at least two reasons. First, the behavior of the 01-agents is represented by bits of information, so that only two possible choices are admitted for agents at any time. For greater realism, it would be better for, say, Sam and Sally to show up in the park more often than at noon and 6:00 p.m., even if the problem of coordination they face becomes presumably harder. This greater realism would be possible if we represent the agents' behavior by a (finitary) first-order language, more precisely by atomic formulas and their negations. Second, the 01-coordination paradigm does not capture either the agent's "beliefs" or the agent's "preferences" over the possible outcomes of their interactions. For example, the story of Sam and Sally that "like to meet daily in the park" [MO99b, p. 363] cannot be formalized within the paradigm. Suppose that Sam prefers to meet Sally in the park on Monday and Friday, but at that famous pub on Thursday, because of live jazz music is played there on Thursday and he is a jazz music lover. (Sally continues to prefer the park every day.) Is then coordination possible between Sam and Sally? Even if both choose identical meeting times, it is quite sure that they will not ever coordinate, because of their different, "uncoordinated" preferences. More generally, the 01-coordination paradigm does not fully capture games where the players have different preferences. The paradigm is a purely behavioral model as agents' predictions are based on no more than the history of early *events*. However, an agent could also use his background knowledge to coordinate. This motivates the search of paradigms where solutions consist of both a "strategy profile"—say a pair $(\Psi(\sigma), \Phi(\sigma))$ for agents Ψ, Φ on finite sequence σ, and a belief system, where a belief system represents the product of all the knowledge concerning the agents and specifies, for each agent at each stage of the coordination process, the beliefs held by the agent who has to move at that coordination stage about the coordination history that occurred.

3 Frank and Gertrud: The passing problem

Consider the following game as reported by Kevin Kelly in [Kel96]:

> "Suppose that Frank and Gertrude have to pass one another in the hall every day. On the first day, Frank thinks Gertrude will move to her right, so he does the same. But Gertrude moves to her left. Frank assumes that her convention is to always move to the left, so he moves to his left. But Gertrude draws the symmetrical inference and moves to her right. Thinking that this has gone far enough, Frank moves to his right again, resolving never to move so that Gertrude can get past. But Gertrude, equally exasperated, makes the same decision! The two die of starvation in the hallway." (p. 267)

Think of passing in the hall as a game for two players, in which both players receive the payoff 1 if they both move right or if they both move left and receive 0 otherwise. Think of the mutual history of Frank and Gertrude as a sequence of successive attempts or trials at passing, whether consecutive attempts occur on the same day (because previous attempts on that day failed) or on successive days (because the last attempt on the preceeding day succeeded). Since they will continue to have "rites of passage" for the unforseeable future, it would certainly be nice if they could eventually reach a state in which they always pass on the first attempt. Suppose that Frank knows in advance that Gertrude will use some strategy in $\mathcal{G}$ and Gertrude knows in advance that Frank will use some strategy in $\mathcal{F}$. Then it would be desirable if Frank could pick a strategy in $\mathcal{F}$ that eventually always passes Gertrude on the first attempt (or, less strictly, after a finite number of failures), no matter what strategy in $\mathcal{G}$ Gertrude happens to be using, and similarly for Gertrude. Under these premises, a question arises: How have to be $\mathcal{F}$ and $\mathcal{G}$ to guarantee us this desiderata? It can be shown that this is possible when $\mathcal{F} = \mathcal{G} =$ the set of all primitive recursive prediction methods.[5] We leave the proof of this result out of the present paper.

[5]Following [Kel96], a **prediction method** is a function $\pi : \omega^* \longrightarrow \omega$, where ω^* denotes the set of all finite sequences of natural numbers and ω denotes the set of all natural numbers.

A criticism

Classical game theorists attempt to explain social phenomena like the kind of coordination problem addressed by Kelly in terms of the concept of equilibrium. An equilibrium for the passing game is a pair of plays (*e.g.*, (left,left)) such that given Frank has played the first, Gertrude would not be inclined to deviate from the second, and given that Gertrude has played the second, Frank would not be inclined to deviate from the first. Hence, (left,left) and (right,right) are the equilibria of the passing game. However, even if we always observe Gertrude and Frank pass each other in an orderly way, by both moving to their right sides, the fact that (right,right) is an equilibrium does *not* explain their success: The traditional approach to analyzing games with multiple equilibria relies on refining Nash's notion of equilibrium until ideally only one survives, and traditional refinements do not accomplish this for coordination games.

4 Deductive vs. inductive approaches

Games are usually divided by game theorists into two general classes: If the players are allowed to enter into binding agreements before the game is played, the game is called 'cooperative'. The games in which there exists no institution or *a priori* knowledge that makes an agreement among players binding are called 'noncooperative'. The most commonly used solution concept in game theory is that of Nash *equilibrium.* The Nash equilibrium is an appealing solution concept for noncooperative games for several reasons. One important virtue of Nash equilibrium is that for games with a finite number of pure strategies and finitely many players, a Nash equilibrium always exists, at least in mixed strategies [Nas51]. A Nash equilibrium also captures a *steady state* of the play of a (repeated) strategic game in which each player holds the correct expectation about the other players' behaviour and acts rationally. It does not attempt to examine the process by which a steady state is reached. As a consequence, game theory has been unsuccessful in explaining *how* players know that a Nash equilibrium will be played. Moreover, the traditional theory is silent on how players know *which* Nash equilibrium is to be played if a game has multiple equally plausible Nash equilibria.

Deductive, or **introspective** theories that attempt to explain equilibrium play directly at the individual decision-making level impose very strong informational assumptions and so are widely recognized as having serious deficiencies. As a consequence, attention has shifted to (purely) inductive, or **behavioral** explanations of equilibrium, motivated by the work of evolutionary biologists. In short, this approach involves exploring the *dynamics* of coordination games. Of course, to do so requires the specification of a dynamic process describing the play of agents involved in such a game. From the opening premises, this means that we must go beyond the Nash equilibrium concept. Two features of the behavioral approach distinguish it from the deductive approach. First, players are not assumed to be so 'rational' as to correctly guess (anticipate) the other players' choices. In other words, in the latter, players are envisaged as potentially very complex machines with very low operating costs, whereas, in the former, their internal complexity is low. Second, the dynamic process that eventually leads to a steady state is specified describing how players adjust their choices over time as they learn from experience about the other players' choices. Therefore, the behavioral approach tries to explain *how* an equilibrium emerges, based on trial-and-error *learning* (instead of introspective-type arguments). Of course, it is a matter of taste if the middle ground between these extremes is more interesting than either extreme. But, as far as we know, only isolated forays have so far been made into this borderline area (see for instance [OR94, Ch. 9] and the reference cited there).

5 Our basic approach

The remarks and criticisms above motivate us in extending the 01-coordination paradigm to a first-order setting. For doing this, we will henceforth identify an "agent" by a pair of objects, where the first element of the pair models the agent's actions and the second element of the pair models the agent's preferences and beliefs. In particular, an agent's action is modeled by a literal on a decidable vocabulary $\mathcal{L}$ together with a countably infinite set $Var = \{v_i \mid i \in N\}$ of variables, while agent's beliefs are modeled by a nonempty collection of structures that interpret $\mathcal{L}$.[6] A paradigm concerning agents of such form will be called of **slow-full coordination**, for short: sf-coordination.

In [Ago01] we develop a new class of models, or *paradigms*, of coordination richer in structure than that considered by Kelly, Montagna and Osherson. Our paradigms are based on the following constitutive components.

(1) **Components of our models of coordination:**

(a) a set of agents, or "players";

(b) a set of interaction-dynamics, or "plays", which provide information about the agents' interaction;

(c) a success criterion that stipulates when and under which conditions the agents can be said to coordinate.

The theory resulting from different paradigms of coordination as drawn by components in (1) is a bag of analytical tools designed to help us understand the phenomena that we observe when decision-makers of the kind evoked in (1)(a) interact. The basic assumptions that underlie the theory are that decision-makers pursue well-defined objectives (they are *rational*) and take into account their knowledge, beliefs, preferences or expectations as well as *other* agents' behavior (they reason *strategically*). The assumption of rationality, however, does not force the agents to be homogeneous; so, it might be the case that a rational agent eventually coordinates with a non-rational agent. *How* agents are rational and reason strategically is yet another question to be investigated. At the present stage of discussion,

[6] The first-order framework we use is standard; see [Ago01, Sec. 1.3] for details.

however, we limit agents to be conceived as systems that examine (partial) evidence coming from other agents' behaviors or empirical data and emit hypotheses and clues. Agents are possibly bounded-resource systems and can fail on some input. What must be explained is precisely how rational agents come to have stable and correct beliefs (on other agents' behavior; on Nature) in the absence of any conclusive evidence about how the other agents (or Nature) are going to choose.

A paradigm of coordination is thus a description—throughout components listed in (1), of the strategic interactions occurring among a set of agents in a given population. Because agent-dependent, these interactions include the constraints on the actions that the agents *can* take and, eventually, the agents' preferences and beliefs. A paradigm of coordination, in contrast, does not specify the actions that the agents actually *do* take. A solution for the decision-makers involved in a paradigm of coordination is a systematic description, in the precise sense of $(1)(c)$, of the outcomes and strategies that may emerge in a family of interactions of the kind in $(1)(b)$. A primary achievement of any paradigm is to suggest reasonable solutions and examine their properties.

A paradigm of slow-full coordination should become clearer by the following intuitive picture of a 2-agent infinitely repeated game, which is intended to help the reader interpret some of the abstract concepts described later.

> We imagine two "knowledge-based" agents, say **A**lfonso and **B**arbara, whose preferences and beliefs are represented by two nonempty classes of structures **A** and **B** that interpret a common, decidable first-order language. To coordinate, **A**lfonso and **B**arbara communicate with each other in order to respectively end up in the limit with a consistent description of a structure, $\mathcal{A} \in \mathbf{A}$ and $\mathcal{B} \in \mathbf{B}$, such that $\mathcal{A}$ is sufficiently close to **B** and $\mathcal{B}$ is sufficiently close to **A**. [...] To dramatize, let us suppose that each agent does not know the preferences of the other, and that agents were never before in a similar situation, so they cannot rely on past experience to solve their coordination problem. [...] To start the game, **A**lfonso is conceived as choosing one member from **A** to be his "actual world", namely, the structure that models **A**lfonso's initial preferences. **A**lfonso's choice is initially unknown to **B**arbara. **A**lfonso then provides a "clue" (**A**lfonso's action) about his chosen world. At the same time [...] **B**arbara does her choice as well, and provides **A**lfonso with a clue about her actual world. [...] Each player may provide "bad clues" in principle, for example to inform the other player of the desire of starting again. [...] **A**lfonso's clues constitute the data upon which **B**arbara will base her hypotheses on **A**lfonso's actual world, that eventually become themselves a clue for **A**lfonso about **B**arbara's world. And so on. Each time **A**lfonso provides a new clue, **B**arbara may produce a new hypothesis, and a new clue for **A**lfonso as well. **A**lfonso and **B**arbara succeed in their coordination game if each player's actual world becomes in the limit "sufficiently close" to the preferences of the other agent." [Ago01, p. 36]

Observe:

1. The game of **A**lfonso and **B**arbara is symmetric, in the sense that both the agents win or lose. However, it is important to observe that the game can be thought to be asymmetric, that is, only one player wins the game. An asymmetric version of the game might be defined by extending agents' preferences to be represented by *ordered* collections of structures. Then, it might be required that only one of the structures built up in the limit by each agent in coordinating is optimal with respect to the "preference relation" of the agent. The agent who stabilizes in the limit on a structure that eventually maximizes the agent's preferences would then be taken to be the winner of the game.[7]

2. Each agent can form his expectation of the other agent's behavior on the basis of his (innate, previously experienced) knowledge as well as on the basis of information coming directly from the other agent's behavior at the time the event occurs. Neither a specific knowledge of the event nor common knowledge or agents' 'rationality' is assumed. The agents choose their actions simultaneously and independently, and have the same language vocabulary as the unique common "communication tool".

3. The paradigm depiced by the game above is a model of a limiting process, in the sense that it concerns the behavior of agents on an infinite subset of their domains. For this reason, we refer to these and similar paradigms as models "in the limit".[8] However, a question about any infinite-horizon paradigm is how it serves to model *finite* situations, like those arising in the actual world. We address such a question here. The fact that the limit is taken to be finite or infinite depends on the use of the model. First, the choice of a finite or infinite limit in the paradigm under analysis depends on the background world each agent concerned is based on. No finite interaction sequence may guarantee success if at least one of the two basic agents is based on a background world with only infinite structures; there is no finite description for an infinite structure. Nevertheless, it is important to stress that *finite* interaction sequences in a sf-coordination paradigm are possible and even necessary to model particular phenomena. Moreover, in applying a limiting paradigm in specific situations we may ask whether a finite-horizon or infinite-horizon is appropriate. Our view is that a limiting paradigm should capture the situation as it is *perceived* by the agents concerned; it should not necessarily aim to describe the situation as depicted by an external observer. This should be enough to justify the use of an infinite-horizon model to describe a situation whose horizon is in some physical sense finite (and vice versa). To illustrate, think again to the situation of Sam and Sally. If after each day Sam and Sally believe that they will continue to encounter in the park for another day, then a paradigm with infinite-horizon is appropriate. In contrast, if Sam and Sally perceive a well-defined final day after which they will stop walking in the park, then a paradigm with finite-horizon is appropriate. Of course, this makes it difficult to decide what model is

[7]The extension of agents to take preference relations on collection of structures into account is not investigated in this paper, but it constitutes foreseeable future work.

[8]Historically, this way of thinking comes from Mark Gold [Gol65, Gol67], that introduced the identification "in the limit" of function and languages.

appropriate if Sam and Sally's perceptions on the future contrast. In a similar event, however, it seems that an infinite-horizon paradigm is more appropriate, as it allow finite representations as a special case.

4. The agents can have a common interest to coordinate. In this sense, one game theorist would say that the paradigm is a *pure* coordination game. In fact, the paradigm is a (finitely, infinitely) repeated pure coordination game with imperfect information. A model of **imperfect information** allows an agent, when taking an action, to have only partial information about the actions taken previously. Thus, the sf-coordination paradigm is a model of imperfect information, since **A**lfonso's choices are *unknown* to **B**arbara, and the vice versa. In other words, the agents choose independently and autonomously what structure to play each time.

5. About the assumption that the agents act simultaneously, we observe that this does not necessarily mean that agents' actions are taken at the same point of time, even if this will be our default interpretation in the sequel. Suppose, for example, that **A**lfonso and **B**arbara are at different geographical locations, in front of terminals, trying to communicate through the Internet. Then the usual net-delay forces **A**lfonso and **B**arbara' acts to be asynchronous, but their actions are still interpreted as simultaneous in our model. In the sf-coordination paradigm it is only important that the agents make decisions independently at any stage of the coordination process, no agent being informed of the choice of the other agent prior to making his own decision.

5.1 Basic components

Some concepts like "agents, "clues", and "success" figure in the foregoing picture of **A**lfonso and **B**arbara coordination game. We formalize them in the next subsection. For this purpose, we step through the components listed in (1) above.[9]

5.1.1 Agents.

We let *SEQ* denote the collection of all finite sequences over the set $\mathcal{L}_{basic}$ of atomic formulas on vocabulary $\mathcal{L}$ and their negations. An agent in a sf-coordination paradigm, or **basic agent**, is a pair $\langle \Psi , \mathbf{A} \rangle$, where Ψ is any mapping of *SEQ* into $\mathcal{L}_{basic}$, and $\mathbf{A}$ is a nonempty (countable) collection of structures. Intuitively, faced with $\sigma \in SEQ$, a basic agent $\langle \Psi , \mathbf{A} \rangle$ believes **action** $\Psi(\sigma)$. Since class $\mathbf{A}$ is interpreted as to represent the agent's "preferences and beliefs", this means that $\Psi(\sigma)$ is expected to be true in some structure in $\mathbf{A}$.[10] For basic agent $\langle \Psi , \mathbf{A} \rangle$, the first component Ψ is called **communication function** or "ability". In particular, if Ψ is total on some $I \subseteq SEQ$, we say that Ψ is a **strategy in** I. We say that

[9]Because of the introductory and motivational charachter of this paper, the rest of this section may be skipped by the reader not interested in a technical exposition. The discussion returns to be plain from Section 6 on.

[10]This is not a strict requirement, however, but the underlying intuition should help the reader in understanding the criterion of success of any slow-full coordination paradigm. By the same interpretation, we require $\mathbf{A}$ to be nonempty, so that agents are assumed to believe something. (It is not a too strong requirement to ask agents to belive $v_0 \doteq v_0$!)

Ψ is a **strategy** if Ψ is a strategy in *SEQ*. The second component $\mathbf{A}$ is called **background world**. We say that $\langle \Psi, \mathbf{A} \rangle$ (or also Ψ) is **based on** $\mathbf{A}$, Ψ is **of** $\langle \Psi, \mathbf{A} \rangle$, and $\langle \Psi, \mathbf{A} \rangle$ **has** Ψ. We write $\mathbf{A}_\Psi$ for $\mathbf{A}$ and $\Psi_\mathbf{A}$ for Ψ (or also $\langle \Psi, \mathbf{A} \rangle$) just in case Ψ is based on $\mathbf{A}$. We refer to $\mathcal{L}$ as the agent's vocabulary, and to $\mathcal{L}_{basic}$ as the agent's language. It is because of the basic language $\mathcal{L}_{basic}$ that agents in the sf-coordination paradigm are termed: "basic". As in the case of 01-agent, a basic agent may have a partial, total, computable or uncomputable communication ability. Then the basic agent is called partial, total, c-computable or c-uncomputable accordingly. We write Λ^b for the class of all basic agents and $\Lambda^b(\mathbf{A})$ for the class of all basic agents based on $\mathbf{A}$.

5.1.2 Dynamics.

Before returning to technical issues, we indulge in a general remark on dynamics. The paradigm of slow-full coordination is designed to examine the logic of long-term interaction. It captures the idea that an agent will take into account the effect of his current behavior on the other agent's future behavior, and aims to explain phenomena like cooperation and threats. For simplicity, at the present stage of our theory development we restrict attention to dynamics based on simultaneous moves, that is, the agents make decisions at the same time, and to pairwise interactions only, that is, interactions involving just two agents.[11]

We now consider formally the interaction between two basic agents. We denote the set $\{0, 1, 2, ...\}$ of natural numbers by N. Let η be an infinite sequence. For $i \in N$, we write $\eta|_i$ for the proper initial sequence of length i in η. We write $|\sigma|$ for the length of a finite sequence, $\emptyset$ for the finite sequence of length zero, σ_i for the ith element of σ, $0 \leq i < |\sigma|$. For every finite or infinite sequence ζ of length $n > k$, $k \in N$, we let $\zeta[k]$ denote the finite sequence $\langle \zeta_0 \cdots \zeta_k \rangle$ and ${}_k|\zeta$ denote the sequence obtained from ζ by deleting its first $k+1$ elements. Thus, $\zeta[k] = \zeta|_{k+1}$, and ${}_0|\zeta = \emptyset$ if $|\zeta| = 1$.

(2) Definition: Let basic agents $\langle \Psi, \mathbf{A} \rangle$ and $\langle \Phi, \mathbf{B} \rangle$ be given.

(a) The **interaction sequence** (or "play") of $\langle \Psi, \mathbf{A} \rangle$ and $\langle \Phi, \mathbf{B} \rangle$ is the infinite sequence

$$D_{\Psi,\Phi} = (\langle \overline{\Psi}_i, \overline{\Phi}_i \rangle \, : \, i \in N),$$

where $\overline{\Psi}_i$ is the ith move of Ψ and $\overline{\Phi}_i$ is the ith move of Φ, defined by induction as follows.

i. $\overline{\Psi}_0 = \Psi(\emptyset)$ and $\overline{\Phi}_0 = \Phi(\emptyset)$.
ii. $\overline{\Psi}_{n+1} = \Psi(\overline{\Phi}[n])$ and $\overline{\Phi}_{n+1} = \Phi(\overline{\Psi}[n])$.

[11] Why 2? One-agent paradigms would degenerate to pure solvability paradigms; it seems that multi-party interaction can be adequately modeled by two-person games, in much the same way that functions with multiple arguments can be reduced to one-place functions and tupling. Of course, this is true in group situation like commercial transictions, where we can assume that communication is pairwise. On the other hand, to model multi-party interaction, as in an auction, pairwise communication seems more problematic.

(*b*) Let $k \in N$ be given. The **interaction sequence** of $\langle \Psi, \mathbf{A} \rangle$ and $\langle \Phi, \mathbf{B} \rangle$ **starting at** k is the infinite sequence

$$^kD_{\Psi,\Phi} = (\langle\, _k|\overline{\Psi}_i \,, _k|\overline{\Phi}_i \,\rangle \,:\, i \in N),$$

where $_k|\overline{\Psi}_i$ is the ith element in $_k|\overline{\Psi}$ and $_k|\overline{\Phi}_i$ is the ith element in $_k|\overline{\Phi}$.

We then say that $_k|\overline{\Psi}_i$ is the ith move of Ψ starting at k and $_k|\overline{\Phi}_i$ is the ith move of Φ starting at k. Observe that $D_{\Psi,\Phi} = \langle\, \overline{\Psi}_0 \,, \overline{\Phi}_0 \,\rangle {}^0D_{\Psi,\Phi}$.

(3) *Remark:* The definition of interaction sequence $^kD_{\Psi,\Phi}$ depends only on the agents' abilities Ψ and Φ; no background worlds are involved. We shall see later in this section that basic agents' background worlds are relevant to determine the criterion of coordination success.

The **response sequence**

$$R(\Psi, \Phi) = (\overline{\Psi}_i \,:\, i \in N)$$

is the (finite or infinite) sequence of moves by basic agent Ψ in response to basic agent Φ, and the **response sequence**

$$R(\Phi, \Psi) = (\overline{\Phi}_i \,:\, i \in N)$$

is the sequence of moves by Φ in response to Ψ. Again, notice that $R(\Psi, \Phi)$ is finite iff at any interaction step $i \in N$, $\Psi(\overline{\Psi}|_i)$ or $\Psi(\overline{\Phi}|_i)$ is undefined. If this is the case, it is immediate to verify that $R(\Phi, \Psi)$ is finite also. Consistently with the notation adopted on infinite sequences within the 01-coordination paradigm, $_k|R(\Psi, \Phi)|_n$ denotes the finite initial sequence in $_k|R(\Psi, \Phi)$ of length n, and $_k|R(\Psi, \Phi)_n$, or also $(_k|R(\Psi, \Phi))_n$ denotes the nth element of $_k|R(\Psi, \Phi)$. Thus, $_k|R(\Psi, \Phi)|_{n+1} = {}_k|R(\Psi, \Phi)|_n \overline{\Psi}_n$ and $R(\Psi, \Phi) = \overline{\Psi}_0 {}^\frown {}_0|R(\Psi, \Phi)$.

5.1.3 Success.

To coordinate basic agents have to stabilize in the limit to a complete description of a structure in their own background world, eventually after a finite number of failures, or disagreements. It is for this reason that we qualified this paradigm of "slow-full" coordination. According to the paradigm, the agents can restart their interaction finitely often, but after the last disagreement they must eventually coordinate. Let us now present all this formally. (The set of elements in any sequence τ is denoted by $rng(\tau)$. Let $\mathcal{S}$ be a structure for $\mathcal{L}$. We write $\mathrm{dom}(S)$ for the domain of $\mathcal{S}$. An assignment h to $\mathcal{S}$ is **complete** if h is a mapping of the variables *Var onto* $\mathrm{dom}(S)$.)

(4) Definition: Let basic agents $\langle \Psi, \mathbf{A} \rangle$ and $\langle \Phi, \mathbf{B} \rangle$ be given. We say that $\langle \Psi, \mathbf{A} \rangle$ ***sf*-coordinates with** $\langle \Phi, \mathbf{B} \rangle$ (written: $\Psi \rightleftharpoons_{sf} \Phi$) just in case for some $s, t \in N$:

(a) $rng({}_s|\overline{\Psi}) = \{\beta \in \mathcal{L}_{basic} \,|\, \mathcal{A} \models \beta[h]\}$ for some $\mathcal{A} \in \mathbf{A}$ and some complete assignment h to $\mathcal{A}$,

(b) $rng({}_t|\overline{\Phi}) = \{\beta \in \mathcal{L}_{basic} \,|\, \mathcal{B} \models \beta[g]\}$ for some $\mathcal{B} \in \mathbf{B}$ and some complete assignment g to $\mathcal{B}$, and

(c) for all $n \in N$, ${}_s|\overline{\Psi}|_n$ is satisfiable in some $\mathcal{B}' \in \mathbf{B}$ and ${}_t|\overline{\Phi}|_n$ is satisfiable in some $\mathcal{A}' \in \mathbf{A}$.

The infinite sequence ${}_s|\overline{\Psi}$ is called an **environment for** $\mathcal{A}$. Note that the definition depends on the interaction sequence $D_{\Psi,\Phi}$, and that structures $\mathcal{A}'$ and $\mathcal{B}'$ depend on n. In case of clause (a), we say that Ψ **enumerates with** Φ ${}_s|\overline{\Psi}$; similarly, in case of clause (b), we say that Φ **enumerates with** Ψ ${}_t|\overline{\Phi}$. So, to sf-coordinate each basic agent outputs in the limit the basic diagram (under some complete assignment) of a structure in his own background world, and such basic diagram must be finitely consistent in some structure of the other agent's background world—generally different from the structure described by him in the limit.

5.2 Main questions

A play between two 'rational' agents will always be described in general terms as follows. Each agent starts the play with his beliefs in mind. At each successive step, each agent's behavior is made known to the other agent. It is important to stress that the agent beliefs are *not explicitly* communicated to the other agent. The learning process then consists of a sequence of temporary equilibria (*i.e.*, pairs of structures, one for each agent), where the agents repeatedly try to estimate the true equilibrium relationship between exogeneous (*i.e.*: communication) and endogenous (*i.e.*: beliefs) variables, using the observed realizations. Since agents update their estimates, the temporary equilibrium relationship between exogeneous and endogenous variables changes over time. Therefore, "learning affects what is to be learned" until some "stabilization" (*convergence*) is eventually reached. From our viewpoint, rational revision is a procedure such that agents' beliefs become eventually stable and correct as the result of an inductive process in the limit.

Let nonempty sets $\Sigma(\mathbf{A})$ and $\Sigma(\mathbf{B})$ of basic agents based on, respectively, $\mathbf{A}$ and $\mathbf{B}$ be given. Let $\{D_{\Psi,\Phi}\} = \{D_{\Psi,\Phi} \,|\, \langle \Psi, \mathbf{A}\rangle \in \Sigma(\mathbf{A}), \langle \Phi, \mathbf{B}\rangle \in \Sigma(\mathbf{B})\}$. A sf-coordination **game** is a triple $\langle \Sigma(\mathbf{A}), \Sigma(\mathbf{B}), \{D_{\Psi,\Phi}\}\rangle$. An **equilibrium** of an sf-coordination game is a pair $(\mathcal{A}, \mathcal{B})$ of structures that satisfies conditions (a)–(c) of the definition above for some pair $\langle \Psi, \mathbf{A}\rangle \in \Sigma(\mathbf{A})$ and $\langle \Phi, \mathbf{B}\rangle \in \Sigma(\mathbf{B})$. Note that an equilibrium of an sf-coordination game is not necessarily unique and does not depend on game stages. A **partial equilibrium** is a pair $(\mathcal{A}', \mathcal{B}')$ of structures that satisfies condition (c) of the definition. Again, note that a partial equilibrium is not unique, but depends on game stages. The basic idea behind the notion of equilibrium is so that an equilibrium should specify not only the basic agents' ability in the limit but also their final beliefs about the coordination history that occurred. The notion of partial equilibrium is less stringent; it specifies the agents' beliefs at *each* stage of the coordination process. Note that equilibrium and partial equilibrium coincide as a special case.

The main idea behind any theory of pure coordination is that whether the agents' interaction sequence is infinite, then mutually desirable outcomes are stable—and we said that the agents have sf-coordinated, or also that an equilibrium of the underlying coordination game exists, if both basic agents believe that communicating a formula inconsistent with the structure chosen by the other agent will terminate the coordination.

Having defined a basic coordination paradigm on pairs of basic agents, namely: the sf-coordination paradigm, we now return to the main concerns of any paradigm of coordination. A fundamental question to be answered about any such paradigm is what types of strategies support mutually desirable "outcomes" for large classes of agents involved in interactive dynamics. More generally, two main problems involve sf-coordination:

(a) Given classes $\mathbf{A}$ and $\mathbf{B}$ of countable structures for the same first-order language on a decidable vocabulary, under which conditions are there basic agents based on, respectively, $\mathbf{A}$ and $\mathbf{B}$ that sf-coordinate?

(b) Let Σ be a collection of basic agents—for example, let Σ a subset of the class of all basic agents based on some $\mathbf{A}$. Under which conditions is there a basic agent that sf-coordinates with each member of Σ?

Let $\langle\, \Sigma(\mathbf{A}), \Sigma(\mathbf{B})\,, \{D_{\Psi,\Phi}\}\,\rangle$ be a sf-coordination game. Let $\Sigma_a = \{\Psi \mid \langle \Psi\,, \mathbf{A}\rangle \in \Sigma(\mathbf{A})\}$ and $\Sigma_b = \{\Psi \mid \langle \Psi\,, \mathbf{B}\rangle \in \Sigma(\mathbf{B})\}$. Then, the question (a) above can be expressed in general game theoretic terms as:

(a′) Is there a communication function $\Psi \in \Sigma_a$ such that for all $\Phi \in \Sigma_b$, $\langle \Psi\,, \mathbf{A}\rangle$ sf-coordinates with $\langle \Phi\,, \mathbf{B}\rangle$, and the vice versa?

In other words, the question is how have to be Σ_a and Σ_b to guarantee sf-coordination. Observe that the answer to question (a) may depend on $\mathbf{A}$, on $\mathbf{B}$, and on the criterion of coordination success adopted, so that some related but different questions may be investigated. Similarly, the answer to question (b) may depend on collection Σ.

To illustrate these questions, we present now some simple results. We rely on the following definition, which generalizes Definition (4) to nonempty collections of basic agents. We call these collections: **coalitions**.

(5) Definition: Let basic agent $\langle \Phi\,, \mathbf{B}\rangle$ and nonempty collection Σ of basic agents be given. We say that $\langle \Phi\,, \mathbf{B}\rangle$ **sf-coordinates with** Σ just in case $\langle \Phi\,, \mathbf{B}\rangle$ sf-coordinates with every member of Σ. In this case, Σ is said to be **sf-tractable**.

Here is an example of sf-tractability. The example shows that there are two infinite groups of agents such that any pair of agents within a group coordinates, and such that whenever an agent coordinates with at least one member of one group he cannot coordinate with any member of the other group.

(6) PROPOSITION: (Uncooperativeness Theorem) There are two infinite coalitions Σ_0 and Σ_1 of basic agents with the following properties.

(a) Every basic agent in Σ_0 sf-coordinates with Σ_0.
(b) Every basic agent in Σ_1 sf-coordinates with Σ_1.
(c) Every basic agent that sf-coordinates with Σ_i, $i = 0, 1$, does not sf-coordinate with any member of Σ_{1-i}.

Proof: Given structure $\mathcal{S}$, environment e for $\mathcal{S}$, and $i = 0, 1$ we define basic agent $\langle \Psi^i_e, \{\mathcal{S}\} \rangle$ as follows.

(a) $\Psi^i_e(\emptyset) = (v_i \doteq v_i)$.
(b) For all $\sigma \in SEQ$ with $|\sigma| > 0$, if $\sigma_0 = (v_i \doteq v_i)$, then $\Psi^i_e(\sigma) = e_{|\sigma|-1}$; otherwise, $\Psi^i_e(\sigma) = \neg (v_0 \doteq v_0)$.

Set $\Sigma_i = \{\Psi^i_e \mid e \text{ is for } \mathcal{S}\}$. Clearly, Σ_0 and Σ_1 satisfy the properties of the proposition. ■

Now we return to the two second main problem mentioned in opening this section. If we limit our investigation to classes of basic agents based on an identical background world, then the problem is the existence of a communication function Ψ, or a strategy, such that following it an agent can coordinate with all members of the class. A sample strategy is provided by the following, simple lemma. It refers to some new terminology and notation. Informally, a "self-centered" basic agent is an agent whose actions completely describe, by interacting with another agent that does not interrupt their interaction sequence, a structure in the background world the agent is based on. Formally, we rely on the following definitions.

(7) DEFINITION: Let $\sigma \in SEQ$ and basic agent $\langle \Psi, \mathbf{A} \rangle$ be given. We define **agent sequence** $\overline{\Psi(\sigma)} \in SEQ$ by induction on the length of σ as follows. *Base*: $\overline{\Psi(\emptyset)} = \Psi(\emptyset)$. Suppose that $\overline{\Psi(\tau)}$ is defined for $\tau \in SEQ$. Given $\beta \in \mathcal{L}_{basic}$, define $\overline{\Psi(\tau\beta)} = \overline{\Psi(\tau)}^\frown \Psi(\tau\beta)$.

Notice that when $\overline{\Psi(\sigma)}$ is defined, $|\overline{\Psi(\sigma)}| > 0$. Also note that the definition does not depend on $\mathbf{A}$. With agent sequences in hand, we set $SEQ_\Psi = \{\overline{\Psi(\sigma)} \mid \sigma \in SEQ\}$.

(8) DEFINITION: Let basic agent $\langle \Psi, \mathbf{A} \rangle$ and nonempty class $\mathbf{K}$ of structures be given.

(a) $\langle \Psi, \mathbf{A} \rangle$ is **K-centered** just in case for all strategies Φ in SEQ_Ψ, $R(\Psi, \Phi)$ is an environment for some structure in $\mathbf{K}$.
(b) $\langle \Psi, \mathbf{A} \rangle$ is **weakly K-centered** just in case for all strategies Φ in SEQ_Ψ, there is $t \in N$ such that ${}_t|R(\Psi, \Phi)$ is an environment for some structure in $\mathbf{K}$.
(c) $\langle \Psi, \mathbf{A} \rangle$ is **(weakly) self-centered** just in case $\langle \Psi, \mathbf{A} \rangle$ is (weakly) $\mathbf{A}$-centered.

We also say that basic agent $\langle \Psi , \mathbf{A} \rangle$ is **world-centered** if $\langle \Psi , \mathbf{A} \rangle$ is (weakly) $\mathbf{W}$-centered for some $\mathbf{W}$. A world-centered basic agent is thus an agent whose actions fully describe some structure, in or out of her beliefs and preference set, by interacting with every agent able to interrupt no part of their interaction sequence.

(9) LEMMA: Every coalition of self-centered basic agents is sf-tractable.

Proof: Let Σ be a nonempty set of self-centered basic agents. We define a total basic agent $\langle \Psi , \mathbf{A} \rangle$ such that:

- (a) $\mathbf{A}$ is the set of structures $\mathcal{S}$ such that there are $\langle \Phi , \mathbf{B} \rangle, \langle \Phi' , \mathbf{B}' \rangle \in \Sigma$ such that $R(\Phi, \Phi')$ is an environment for $\mathcal{S}$.
- (b) $\Psi(\emptyset) = (v_0 \doteq v_0)$.
- (c) For all $\sigma \in SEQ$, if $|\sigma| > 0$ then $\Psi(\sigma) = \sigma_{|\sigma|-1}$.

Observe that $\mathbf{A}$ is nonempty. Thus, when playing with a basic agent $\langle \Phi , \mathbf{B} \rangle \in \Sigma$, Ψ starts by moving $(v_0 \doteq v_0)$ and then copies Φ's last move one step later. It follows that $R(\Psi, \Phi) = (v_0 \doteq v_0)R(\Phi, \Psi)$. Since $\langle \Phi , \mathbf{B} \rangle$ is self-centered, $R(\Psi, \Phi)$ is an environment for some structure $\mathcal{S} \in \mathbf{B}$. From the definition of Σ it follows that $\mathcal{S} \in \mathbf{A}$. So, it is easy to check that $\langle \Psi , \mathbf{A} \rangle$ sf-coordinates with $\langle \Phi , \mathbf{B} \rangle$. ■

An important remark follows. The theory of sf-tractability gives us insights into the structure of behavior when individuals like basic agents interact repeatedly, taking into account the effect of the history of their interaction. Although we regard results like Lemma (9) about the intrinsic form of the strategies that guarantee equilibrium (in the strict sense given by the success criterion of each paradigm) to be the main achievement of the theory, most of our results (see [Ago01]) focus instead on the classes of structures that make coordination between basic agents based on those classes possible. Nevertheless, we stress that in our opinion the main contribution of a "good" theory is the discovery of interesting abilities (strategies) that support coordination with large classes of basic agents. The difficult problem to find interesting necessary and sufficient conditions for sf-coordination and similar paradigms is still open.

6 Learning and Games: passive vs. active players

A paradigm of pure solvability [12] concerns with an infinite game of a learner-scientist Ψ against Nature, where Nature has a null pay-off function. The following picture of scientific inquiry gives an intuitive idea of the game.

[12] Here as in the rest of this article we prefer to use "solvability" instead of "learning" under the assumption that "learnable" denotes a wider class of objects than "solvable" does. In particular, under such assumption we mean learning to be present in both coordination and solvability, in the intuitive sense that one might learn how to coordinate as well as how to solve a given problem. Apart from this somewhat philosophical remark, in this article we shall always mean the two terms as synonymous.

> "First, a class of possible realities is specified in advance; the class is known to both players of the game. Nature is conceived as choosing one member from the class to be the 'actual world'; her choice is initially unknown to the scientist. Nature then provides a series of clues about this reality. These clues constitute the data upon which the scientist will base his hypotheses. Each time Nature provides a new clue, the scientist may produce a new hypothesis. The scientist wins the game if there is sufficient guarantee that his successive conjectures will stabilize to an accurate hypothesis about the reality Nature has chosen." [OdJMW97, p. 740]

So, the scientist wins the game if he makes in the limit the right conjecture on the reality Nature made actual, he loses otherwise.

A criticism

It is a common experience that in order to prove that some problem is *not* solvable, people often use an "adversary argument". By "adversary argument" we can then imagine Nature playing against Ψ and eventually succeed in deceiving him. In contrast to pure solvability, where Nature may be coinceived as a "passive" agent, information by an adversary argument involves "active" players, as it comes from the players' interaction. A similar situation arises in a question-answer system. For example, suppose that a scientist quests about the reality Nature made actual and waits for answer. Again, one might expect that Nature's answer depends on the scientist's question. Of course, mutual interaction is possible only for systems with more than one player and leads it to game-theoretic analyses. As a first consequence, game theory becomes increasingly relevant to the study of solvability issues, with some restrictions that we clarify in a moment. On the other hand, in order to win a game it is often useful for a player to guess the opponent's strategy and predict her next moves. This activity is typical of learning and leads learning theory to become of central importance in game dynamics. Hence, we can at least infer that game-theoretic and learning-theoretic analyses influence each other in some interesting way.

Now we focus on two major differences. First, suppose that Player I (PI) and Player II (PII) start playing chess. Then, learning PII's strategy may be useful for PI, but it is neither necessary nor sufficient for him to win. Indeed, PI could realize that PII can win the game in, say, four moves. If PI's guess is correct, PI lose. However, if PII makes a mistake by moving differently from what her (winning) strategy suggests she to do, then PI failed to predict PII's erroneous move, but at the same time PI have some new chance to win the game. Hence, learning separates from game theory in a first, important sense. Second, suppose that a learner Ψ solves a problem $(\mathbf{W}, \pi)$ in the sense of [Ago01].[13] Then Ψ must stabilize to a sentence in π on *all* environments for *all* structures in $\mathbf{W}$. In other words,

[13]For simplicity, we leave out from the present discussion as many details as possible, and refer the interested reader to [Ago01] for a more technical exposition of the same argument.

Ψ must possess a complete explanation of the problem under investigation. In contrast, success in games only depends on the unique response sequences $R(\Psi, \Phi)$ and $R(\Phi, \Psi)$ of the binary agents Ψ and Φ concerned: Ψ need not to know Φ's ability (strategy, if total) on *all* inputs where it is defined. Again, suppose that Player I and Player II are playing chess and that is PI turn to move. If PI's choice is, say, e2-e4, then PI has no way to predict PII's successive move on say PI's alternative move d2-d4 that follows the same game history, not even after the game has developed further. Of course, it is assumed that PI and PII never played chess together, so that PI cannot rely on previous experiences on how PII moved in identical circumstances, if any. Hence, learning separates from game theory in a second important sense.

To illustrate further on the restrictions above, let suppose that a game between scientist Ψ and Nature is played, where Ψ wins the game just in case Ψ predicts *all* Nature's moves. Such a game has two different interpretations as Nature is conceived to be "active" or "passive" in her interactions with scientist Ψ, the latter being the standard view of Nature in any paradigm of solvability. It is intuitively obvious that the kind of learning of Nature to be coinceived as an "active" agent, that is, an agent whose evidential states and moves strictly depend on the predictions of the other agent in the game (scientist Ψ), is very different from the kind of learning propagated by the pure solvability paradigm. Hence, game theory enters the game of Nature in two ways, say "active" and "passive"—the latter being that of the pure solvability, despite of what kind of agent Nature is coinceived to be.

6.1 Cooperation

The preceding remarks are intended neither to justify the use of game-theoretic analysis in understanding learning issues, nor as definitive or unique motivation for a learning perspective in game theory. Their purpose is merely to motivate exploring the logic of coordination-based learning, and to appreciate the complexity of the interaction of learning theory and game theory, as it involves a wide range of phenomena both natural and social. Agents should then be made able to manage information from different sources. For greater realism, though limited of substantial idealization, these sources should be meant either as "natural" environments or as "social" interactions of the kind evoked in $(1)(b)$ as basic components of our models of coordination. We will refer to both these sources with the term "enumerations". Otherwise, only one-way interaction is possible, that is the interaction between an "active" agent and a "passive" agent, like a scientist and Nature; no mutual interaction game are possible.

In any event, we generally expect a set of agents' behaviors to influence each other—a set $\{scientist, Nature\}$ as a special case, in a way that is not captured by any "active-passive" model of solvability. It is in this sense of lacking that we identified a paradigm of solvability: "pure". Whether agents win or not in their (noncooperative) games against Nature, a real problem in extending solvability models to multi-agent setting is thus determining paradigms of mutual interaction in problem solving—we call these paradigms of **cooperation**.

Coordination and cooperation may be related with an equational *slogan*, that we hope of intuitive and immediate help for the reader.

(10) Cooperation = coordination + solvability.

The dimensions on which the equation is based on are three, roughly: absence of coordination, absence of solvability, presence of both. With the equational slogan (10) in mind, we advance below the basic components that a sound paradigm of cooperation should have.

(11) **Components of our models of cooperation:**

(*a*) a set of agents;

(*b*) a set of possible realities, or "worlds";

(*c*) for each world, a set of interaction-dynamics, or "enumerations", which provide information about the world;

(*d*) a learning context, or "natural", "social" domain;

(*e*) a success criterion that stipulates the conditions under which the agents can be said to coordinate in a learning context (social domain);

(*f*) a success criterion that stipulates the conditions under which the agents can be said to learn a class of realities in a learning context (natural domain).

To summarize, I think it is important to develop a theory of learning in coordination games, a theory that would explain how an agent's beliefs about the environment (which includes the behavior of others insofar as it affects him) evolve until they have come to agree with the actual properties of the environment. This "agreement" means that agents' expectations are stable and accurate, in a precise sense defined for each paradigm in question. All this opens new perspectives on the interaction of game theory and learning theory, and might constitute the core research laboratory for a large overlapping of research areas. I list some of them in the next final section.

7 Directions for research

Many lines of research for exploring fundamental issues in coordination and cooperation at the bordeline of learning theory and game theory lie open. Perhaps the most fundamental questions to be answered about coordination are what strategies characterize and support mutually desirable outcomes for large populations Σ of agents. These questions strictly relate to those in subsection 5.2. The theory of X-tractability and of π-tractability as developed in [Ago01] gives some insights into the structure of behavior when heterogeneous belief-based individuals interact repeatedly, taking into account the effect of the history of their interaction. Of course, the answer to such a question would greatly benefit from a look at the results that have already been obtained for the game-theoretic models of coordination. Various problems and directions remain.

7.1 Belief Revision

At the borderline of learning theory and game theory is the problem of belief changes—how an agent should revise her beliefs upon learning new information. Belief revision is so far only defined for agents whose "starting points", or "background theory", are sets of sentences [Bou98, Fri97, Gär88, Kel96], [KSH97, Lev80, MO97, MO98, Han94, Han99]. On the other hand, we have seen that all our paradigms consider an agent's beliefs to be represented by a "background world", a nonempty collection of structures for a fixed (first-order) language. New paradigms of coordination and cooperation among revision-based agents could then be advanced from our work by answering the following question: is there a natural (justifiable on intuitive grounds), semantic generalization in which revision applies directly to classes of structures? If the answer is "yes", how to extend the mathematics of coordination so as to be able to frame appropriate conditions on the revision operator so that its use has a 'rational' quality to it? The revision operator could then be used to modeling the agents preferences over their beliefs. In particular, coordination and cooperation of two revision-based agents would be investigated into three classes of winning strategies:

(a) strategies that hold constant the background world of each agent and then determine what kind of revision operators allow agents to succeed;

(b) strategies that hold constant the revision operator of each agent and determine what kind of background worlds allow the agents to succeed;

(c) strategies that hold constant both the background worlds and the revision operators, and determine what kind of learning contexts allow the agents to succeed.

In cases (a) and (b), the results will depend on which learning context is in question. In case (c), to discover the "right" learning context will be the result.

7.2 Computability

Most of the applicability of our models depends on how, when and under what limiting costs coordination and cooperation can be proved to be computably tractable. On the reasonable assumption that inductive algorithms will become increasingly important for the development of distributed multi-agent systems, to examine paradigms of inductive inference among belief-based agents that scale up to computable, real world scenarios is therefore desirable. Crucial to this enterprise is thus finding ways to represent nonempty collections of countable structures in a computable manner. It seems therefore necessary to limit attention to *definable* classes of structures via an *r.e.* set of formulas.

7.3 Knowledge Representation

Interaction may well be desirable between a model-theoretic approach to coordination and learning, on the one hand, and issues in knowledge representation, on the other. We can view a (computable) agent as a multiple database DB along with a query function defined on it, and first-order logic can be considered as a query language for DB. Part of the knowledge stored in DB may consist of first-order statements that serve as the axioms of a class of models. To augment its knowledge, DB can wait for external assistance to augment the axiom set, or it can launch its own investigation via an automated system of inductive inference. A speculative note on the latter project is given in [OSW94], that implies a primary set of feasibility questions to be asked to formal and computational learning theory. From the side of the former project, it is therefore interesting to see whether cooperation paradigms have something to offer. Our work on coordination opens a way by modeling the DB requests of assistance as a query-answer system. We have seen the reasons why a paradigm of coordination is suitable in principle for modeling such kind of systems. Answers to feasibility questions along this way will depend on the kind of axioms that form the DB knowledge as well as the way an agent's knowledge is represented. In this perspective, an integration of the theory of contexts [Giu93, GS94, GB97, GS98, GS99, BBG00, GG01, SG00] to the present setting might be a productive direction of research, as inductive paradigms eventually give a dynamics-oriented plus-value to the static view of the compatibility of contexts presented e.g. in [GG01]. First-order logic provides a rich class of database queries. However, some plausible queries are not first-order expressible, so that we may be led to ask for stronger logics.

7.4 Incomplete, "noisy", imperfect, recursive environments

A game between two agents may be seen as concerning alternative construals of available data in empirical inquiry. Similarly to what happens for identification of languages and functions, an environment for a structure may suffer omissions, erroneous intrusions, or both omissions and intrusions. Moreover, a paradigm of solvability often supposes that an input environment can be ordered arbitrarily, and thus requires a learner to succeed regardless of the manner in which the learner data are presented. This is, however, a questionable representation of many natural environments, at least for some; see [JORS99, Ch. 8]. In all these cases, a paradigm of team learning may be applyed. For each team of two agents, one of them may re-interpret an input environment e in different ways by: (a) missing some data from e; (b) introducing some "noisy" data; (c) ordering the data from e according to some additional order structure. When possible, the agent could also be used to enumerate e in a recursive way. Finally, a paradigm of team learning may be used to modeling "multiple environment" learning situations, similarly to the way recursive learning theory does it for language identification. Intuitively, this may be done by considering teams of heterogeneous agents, each playing the role of a source of information on a well-fixed reality to learn.

7.5 Knowledge Management

"The idea is to develop a kind of '*game*' that in a *restricted and controlled environment* enables the creation of the *typical dynamics* of the distributed and autonomous knowledge management. The role of such infrastructure [implemented game] will be to provide a kind of *experimental laboratory* for the *validation* of the research results." [AA.01, p. 8-29, my emphasis]

The aim of this research direction would be to approach by a formal framework similar to that depiced in Section 5 above, and possibly expanded and improved in the fourth directions for research mentioned, some important questions which arise from the previous two citations, among others:

(*a*) What kind of game?

(*b*) What does it mean "restricted and controlled environment"?

(*c*) What are the "typical dynamics" and the "meaningful aspects" of the distributed and autonomous knowledge management?

(*d*) What is a "validation" of the research results by means of an "experimental laboratory" as a game could provide for?

Of course, the foregoing questions can be combined in many ways.

On a final note, it might be possible to extend the work along these lines by deploying the considerable understanding that has accumulated about finite model theory (see *e.g.* [EF95]) and recursive model theory (see *e.g.* [Mil99]) in the last two decennial of the XX century. Perhaps it can be deployed in model-theoretic work on learning via inductive reasoning, and used as a bridge between this quite abstract and technical discipline and real multi-agent scenarios in AI and in economics.

Appendix

This appendix relates with some background and references on the main topics and research areas mentioned in this paper. It can be omitted on the first reading.

Game theory

Game theory generates from the classic [vNM44], and is beautifully summarized and extended by [LR57], now out of date but still full of insightful discussions and examples. Coordination problems that concern common interest agents are usually modeled as noncooperative games with multiple Nash equilibria in which any Pareto-efficient strategy combination is an equilibrium, but players' strategy choices are optimal only when they are based on sufficiently similar beliefs about how the game will be played. Although such games have no incentive [pay-offs] problems as these are normally characterized, playing them often involves real difficulties. See for instance [Coo99] and the references cited there for some experimental recent studies in economics. The traditional approach to analyzing games with multiple equilibria (as common interest games) relies on refining Nash's notion of equilibrium until ideally only one survives, and traditional refinements do not accomplish this for coordination games. As expected, this issue is at the core of the traditional approach of noncooperative game theory. An example of failure of traditional refinements towards a unique Nash equilibrium can be found *e.g.*

in [KMR93] (Section 5), where the riskiness (called *risk-dominance criterion* in [HS88]) of, say, Frank and Gertrude's choice to move to the right relative to the complementary choice to move to the left is taken to be relevant. Of course, the literature on equilibrium selection is immense. We do not attempt to address it in this paper.

Introspective theories of equilibrium selection, also known as "eductive", Binmore [Bin87] and "deductive", [OR94], are widely recognized as having serious deficiencies (see for instance [Bin87, Bin88]). Behavioral theories—called "evolutive" by [Bin87] and "evolutionary" by [KMR93], derive from the inductive, *steady state* interpretation of equilibrium in the view of [Car66]: "The observations we make in everyday life as well as the more systematic observations of science reveal certain repetitions or regularities in the world...The laws of science are nothing more than statements expressing these regularities as precisely as possible." (p. 3) They are historically motivated by the work of evolutionary biologists. We refer the interested reader to the seminal work [MSP73] (also [MS82]), and to [BGM92, HS94] for an introductory survey on evolutionary game theory. A steady state approach for finite games is [FL93]. Some mixed approaches to equilibrium selection based on both introspective and behavioral perspective are reported in [OR94, Ch. 9].

Coordination, cooperation and AI

Coordination and cooperation have received the bulk of attention by Artificial Intelligence (AI). Investigations into the coordination problem can be divided into three general classes: those based on convention, those based on communication, and those based on learning. Some example in the first class was given in [ST92b, ST92a], where social laws [Lew69] are imposed by the system designer (see also [PAM97, ST97]) so that optimal joint action is assured. In the second class, agents' coordination is based on communication (see for instance [Wei93]). This second class might be thought as a special case of the normative class, where the communication language is assumed to be the convention. So, what makes this class different is rather the emphasis it gives to the communicative agents' skills with regard to agents' coordination problems. In this class, it does make sense to speak about *failure messages* that prevent the agents from coordinating (see for instance [WJ95] for some further remarks and references on the influence of *speech act theory* in communication). In the third class, coordination might be learned through repeated interaction; see for instance [Bou96, BGS$^+$91, SSH94, Wei93]. Cooperation has been extensively studied, both in game theory and AI, where fully cooperative problems arise in task distribution as well as within the historically older sub-area of the multi-agent problem solving. An incomplete list is [Dur88, Kra97, PAM97, WJ94, WJ99].

Formal learning theory

From the side of formal learning theory, the field has split down into at least two related parts in the middle of the 1980s. These parts are known at present days as the 'model-theoretic' and the 'recursion-theoretic' tradition. The recursion-theoretic tradition descends from the pioneering studies on inductive inference: [Sol64, Put65, Gol67, BB75]; see [AS83, OSW86] and [JORS99] for surveys. Historically, the work in the recursion-theoretic tradition concerns algorithms for inferring recursive functions from finite samples of their graphs, and has been adapted successfully to characterize abstract languages in the limit. The model-theoretic tradition is more recent and is the leading perspective of this thesis. On the contrary, we do not investigate further the recursion-theoretic arena of solvability, nor shall we consider other interesting but quite different approaches to formal learning in different research fields, such as: the model-theoretic approach of CMU's school, notably [Kel96]; see also [KG89, KG92, KSJ97], where learning is often related to philosophical questions on the nature and methodology of scientific discovery; the "probably approximately correct", or PAC learning, cf. [KV94], where "efficiency" of learning is a primary goal; the "Minimum Description Length principle", a relatively recent method for inductive inference central in statistics, pattern recognition and a branch of Artificial Intelligence called 'machine learning', the fundamental idea of which is that any regularity in a given set of data can be used to 'compress' the data—the more regularities there are in the data, the more them can be compressed (see references listed in [Grü98] and [VL97, LV97] for overviews); "reinforcement learning" and learning in AI (two incomplete lists are [CB97, Bou96, BGS$^+$91, DE97, Mic93] and [SSH94, Sia91, Wei93, Wei95]).

Learning, rationality and knowledge

As far as we know, a "zero-one" paradigm of coordination was first introduced in formal learning theory by [MO99b] by using the tools of recursion theory. [Kel96] advanced a similar paradigm (see the Frank and Gertrude game) without formal development. Flexible agents are "limited strategists", in that they simply adapt their choice to the action taken by the opponent on the preceding move. The kinds of strategies implemented by flexible agents are called *Markov strategies* in [Bic93], because "at each stage a player chooses an action independently of the history of the game except for the immediately preceding action...Agents that use such memoryless strategies will not try to identify complex patterns of play" (p. 237). Flexibility is very similar to the classic TIT FOR TAT; the two strategies are equal if and only if flexible agents' first move coincede. The strategy TIT FOR TAT is simply one of cooperating on the first move and then doing whatever the other agent did on the preceding move. Thus, TIT FOR TAT is a strategy of coordination based upon reciprocity and imitation. In spite of its great simplicity, TIT FOR TAT has been proved successful in many situations and computer applications; see for instance [Axe84, Axe97]; a nice example of strategy success is [Mil87]; an example on the power of strategies based on imitation is [Sch98]. We will encounter and further discuss TIT FOR TAT and some of its variants as the thesis proceeds.

The use of TIT FOR TAT and similar imitative strategies to model agents in the behavioral context of binary coordination and the like is presumably uncontroversial. However, for such a context the emphasis belongs on machines of *low* complexity compared with their environment; while agents involved by the behavioral-deductive approach presented in this article should be understood to have the potential to be of very *high* complexity compared with their environment. See also [Sim77] for a general discussion on complexity. Of course, most of the applicability of deductive models depends on how, when and at which limiting costs such complexity can be proved to be computably tractable.

On the interplay of learning and games the literature is huge, as a consequence of what use of "learning" is made. Some of such literature is cited above. A paradigm of question-answer learning in the context of [Gol65, Gol67]'s "identification in the limit"—where the paradigm is called *black box identification*, is partially reported by [Lun97] to continue the seminal work of [Spe89]. A discussion on finite vs. infinite-horizon limiting processes (games) is presented in [OR94, Ch. 8]. For learning in Chess see for instance [SS92].

The literature on background knowledge and rationality is huge, expecially from game theory and philosophical logic. Some questions which concern us are addressed in [Ago00] and [ER99]. We refer the interested reader to books like [Gär88, Kel96, Lev80, MO98, Han99] for the long debate on the role of belief revision in scientific discovery and its relationship with rational empirical inquiry. [MO97, MO99a] and [KSH97, Kel98] are works on the same topic but from a quite different perspective, the former two articles being model-theoretic in their formal exposition, the latter two articles being set-theoretic. These papers all focus on inductive inference methods in a single agent setting.

What has to be termed an "agent" and what does not is a long debate in Artificial Intelligence and Game Theory; we suggest [FG97] for a critical discussion and [WJ95, Rus97] for respectively a survey and a nice discussion on rationality and rational agents from the one side; [Bin87, Bin88] from the other side, with the promise that we will not even attempt a complete explanation of agency from a so general perspective.

References

[AA.01] AA.VV. Enabling distributed and autonomous management of knowledge. Project EDAMOK Official Proposal, May 2001.

[Abr97] S. Abramsky. Semantics of Interaction: an introduction to Game Semantics. In P. Dybjer and A. Pitts, editors, *Proceedings of the 1996 CLiCS Summer School, Joint Newton Institute*, pages 1–31. Cambridge University Press, 1997. [http://www.dcs.ed.ac.uk/home/samson/coursenotes.ps.gz].

[Ago00] A. Agostini. Notes on formalizing coordination. In E. Lamma and P. Mello, editors, *AI*IA 99: Advanced in Artificial Intelligence*, pages 285–296. Springer-Verlag LNAI 1792, 2000.

[Ago01] A. Agostini. *Paradigms of Coordination and Solvability.* PhD thesis, LOMIT-XI, Dip. di Matematica, Università di Siena, March 2001.

[AM98] S. Abramsky and G. McCusker. Game Semantics. In H. Schwichtenberg and U. Berger, editors, *Logic and Computation: Proceedings of the 1997 Marktoberdorf International Summer School on Computational Logic*. Springer-Verlag, 1998. [http://www.dcs.ed.ac.uk/home/samson/mdorf97.ps.gz].

[AS83] D. Angluin and C. H. Smith. Inductive inference: Theory and Mhetods. *Computing Surveys*, 15(3):237–269, 1983.

[Axe84] R. Axelrod. *The Evolution of Cooperation*. Basic Books, New York, NY, 1984.

[Axe97] R. Axelrod. *The Complexity of Cooperation*. Princeton University Press, Princeton, NJ, 1997.

[BA96] R. J. Brachman and T. Anand. The process of knowledge discovery in databases: A human-centered approach. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 37–57, Menlo Park, CA / Cambridge, MA, 1996. AAAI Press / The MIT Press.

[BB75] L. Blum and M. Blum. Toward a mathematical theory of inductive inference. *Information and Control*, 28(2):125–155, 1975.

[BBG00] M. Benerecetti, P. Bouquet, and C. Ghidini. Contextual Reasoning Distilled. *Journal of Theoretical and Experimental Artificial Intelligence*, 12(3):279–305, 2000.

[BBM00] M. Bonifacio, P. Bouquet, and A. Manzardo. A distributed intelligence paradigm for knowledge management. In S. Staab, editor, *Bringing Knowledge to Business Processes - AAAI Spring Symposium Series 2000*. AAAI Press, 2000.

[BGM92] P. Battigalli, M. Gilli, and M. C. Molinari. Learning and convergence to equilibrium in repeated strategic interactions: An introductory survey. *Ricerche Economiche*, XLVI(3-4):335–378, 1992.

[BGS+91] P. Brazdil, M. Gams, S. Sian, L. Torgo, and W. van de Velde. Learning in distributed systems and multi-agent environments. In Y. Kodratoff, editor, *Machine Learning - European Working Session on Learning*, pages 412–423. Springer-Verlag LNAI 482, 1991.

[Bic93] C. Bicchieri. *Rationality and Coordination*. Cambridge Studies in Probability Induction, and Decision Theory. Cambridge University Press, Cambridge, 1993.

[Bin87] K. Binmore. Modelling rational players - Part I. *Economics and Philosophy*, 3:179–214, 1987.

[Bin88] K. Binmore. Modelling rational players - Part II. *Economics and Philosophy*, 4:9–55, 1988.

[Bin92] K. Binmore. *Fun and Games.* Heath & Company, Lexington, MA, 1992.

[Bou96] C. Boutilier. Planning, learning and coordination in multiagent decision processes. In Yoav Shoham, editor, *Proceedings of the Sixth Conference on Theoretical Aspects of Rationality and Knowledge (TARK-96)*, pages 195–210, San Francisco, CA, 1996. Morgan Kaufmann.

[Bou98] C. Boutilier. A unified model of qualitative belief change: a dynamical systems perspective. *Artificial Intelligence*, 98:281–316, 1998.

[Car66] R. Carnap. *Philosophical Foundations of Physics.* Basic Books, New York, NY, 1966.

[CB97] C. Claus and C. Boutilier. The dynamics of reinforcement learning in cooperative multiagent systems. In *Proceedings of the 15th National Conference on Artificial Intelligence (AAAI-98) and of the 10th Conference on Innovative Applications of Artificial Intelligence (IAAI-98)*, pages 746–752, Menlo Park, CA, 1997. AAAI Press.

[CJLZ99] J. Case, S. Jain, S. Lange, and T. Zeugmann. Incremental concept learning for bounded data mining. *Information and Computation*, 152(1):74–110, 1999.

[Coo99] R. W. Cooper. *Coordination Games.* Cambridge University Press, Cambridge, 1999.

[DE97] W. Davies and P. Edwards. The communication of inductive inferences. In G. Weiß, editor, *Distributed Artificial Intelligence meets Machine Learning*, pages 223–241. Springer-Verlag LNAI 1221, 1997.

[Doe96] K. Doets. *Basic Model Theory.* Studies in Logic, Language and Information. CSLI Publications & FoLLI, Stanford, CA, 1996.

[Dur88] E. H. Durfee. *Coordination of Distributed Problem Solvers.* Kluwer Academic Publishers, Boston, MA, 1988.

[EF95] H-D Ebbinghaus and J. Flum. *Finite Model Theory.* Springer, 1995.

[ER99] D. Easley and A. Rustichini. Choice without beliefs. *Econometrica*, 67(5):1157–1184, 1999.

[FG97] S. Franklin and A. Graesser. Is it an agent, or just a program?: A taxonomy for autonomous agents. In J. P. Müller, M. J. Wooldridge, and N. R. Jennings, editors, *Intelligent Agents III - Agent Theories, Architectures, and Languages*, pages 21–35. Springer-Verlag LNAI 1193, 1997.

[FL93] D. Fudenberg and D. K. Levine. Steady state learning and Nash equilibrium. *Econometrica*, 61(3):547–573, 1993.

[FPSS96] U. M. Fayyad, G. Piatetsky-Shapiro, and P. Smyth. From data mining to knowledge discovery: An overview. In U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors, *Advances in Knowledge Discovery and Data Mining*, pages 1–34, Menlo Park, CA / Cambridge, MA, 1996. AAAI Press / The MIT Press.

[FPSSU96] U. M. Fayyad, G. Piatetsky-Shapiro, P. Smyth, and R. Uthurusamy, editors. *Advances in Knowledge Discovery and Data Mining*, Menlo Park, CA / Cambridge, MA, 1996. AAAI Press / The MIT Press.

[Fri97] N. Friedman. *Modeling Beliefs in Dynamics Systems.* PhD thesis, Stanford University, March 1997.

[Gär88] P. Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States.* The MIT Press, Cambridge, MA, 1988.

[GB97] F. Giunchiglia and P. Bouquet. Introduction to contextual reasoning. An Artificial Intelligence perspective. In B. Kokinov, editor, *Perspectives on Cognitive Science*, volume 3, pages 138–159. NBU Press, Sofia, 1997. Lecture Notes of a course on "Contextual Reasoning" of the European Summer School on Cognitive Science, Sofia, 1996.

[GG01] C. Ghidini and F. Giunchiglia. Local Models Semantics, or contextual reasoning = locality + compatibility. *Artificial Intelligence*, 127(2):221–259, 2001.

[Giu93] F. Giunchiglia. Contextual reasoning. *Epistemologia, special issue on I Linguaggi e le Macchine*, XVI:345–364, 1993. Short version in Proceedings IJCAI-93 Workshop on Using Knowledge in its Context, Chambery, France, 1993, pp. 39–49. Also IRST-Technical Report 9211-20, IRST, Trento, Italy.

[Gol65] E. M. Gold. Limiting recursion. *Journal of Symbolic Logic*, 30(1):28–48, 1965.

[Gol67] E. M. Gold. Language identification in the limit. *Information and Control*, 10:447–474, 1967.

[Grü98] P. Grünwald. *The Minimum Description Length Principle and Reasoning under Uncertainty.* PhD thesis, University of Amsterdam, ILLC, Oct. 1998.

[GS94] F. Giunchiglia and L. Serafini. Multilanguage hierarchical logics, or: how we can do without modal logics. *Artificial Intelligence*, 65(1):29–70, 1994. Also IRST-Technical Report 9110-07, IRST, Trento, Italy.

[GS98] C. Ghidini and L. Serafini. Distributed First Order Logics. In D. Gabbay and M. de Rijke, editors, *Frontiers of Combining Systems 2*, Studies in Logic and Computation, pages 121–140. Research Studies Press, 1998.

[GS99] C. Ghidini and L. Serafini. A Context-Based Logic for Distributed Knowledge Representation and Reasoning. In P. Bouquet, L. Serafini, P. Brezillon, M. Benerecetti, and F. Castellani, editors, *Modelling and Using Context*, pages 159–172, Heidelberg, 1999. Springer-Verlag LNAI 1688. Also IRST-Technical Report 9909-11, IRST, Trento, Italy.

[Han94] S. O. Hansson. Kernel contraction. *Journal of Symbolic Logic*, 59(3):845–859, 1994.

[Han99] S. O. Hansson. *A Textbook of Belief Dynamics.* Kluwer Academic Publishers, Dordrecht, NL, 1999.

[Hod97] W. Hodges. *A Shorter Model Theory.* Cambridge University Press, Cambridge, 1997.

[HS88] J. C. Harsanyi and R. Selten. *A General Theory of Equilibrium in Games.* The MIT Press, Cambridge, MA, 1988.

[HS94] P. Hammerstein and R. Selten. Game Theory and Evolutionary Biology. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume II of *Handbooks in Economics, v. 11*, pages 929–993. Elsevier Science B.V., 1994.

[HS97] J. Hintikka and G. Sandu. Game-theoretical semantics. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 361–410. Elsvier Science B.V., Amsterdam, 1997.

[Hui49] J. Huizinga. *Homo Ludens - A study of the play-element in culture.* International Library of Sociology and Social Reconstruction. Routledge & Kegan Paul Limited, London, 1949.

[JORS99] S. Jain, D. Osherson, J. Royer, and A. Sharma. *Systems That Learn - An Introduction to Learning Theory, 2nd edition.* The MIT Series in Learning, Development, and Conceptual Change, v. 22. The MIT Press, Cambridge, MA, 1999.

[Kel96] K. T. Kelly. *The Logic of Reliable Inquiry.* Oxford University Press, New York, NY, 1996.

[Kel98] K. T. Kelly. The learning power of belief revision. In Itzhak Gilboa, editor, *Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK-98)*, pages 111–124, San Francisco, CA, 1998. Morgan Kaufmann.

[KG89] K. T. Kelly and C. Glymour. Convergence to the truth and nothing but the truth. *Philosophy of Science*, 56(2):185–220, 1989.

[KG92] K. T. Kelly and C. Glymour. Inductive inference and theory-laden data. *Journal of Philosophical Logic*, 21(4):391–444, 1992.

[KMR93] M. Kandori, G. J. Mailath, and R. Rob. Learning, mutation, and long run equilibria in games. *Econometrica*, 61(1):29–53, 1993.

[Kra97] S. Kraus. Negotiation and cooperation in multi-agent environments. *Artificial Intelligence*, 94(1-2):79–97, 1997.

[KSH97] K. T. Kelly, O. Schulte, and V. Hendricks. Reliable belief revision. In M. L. Dalla Chiara, K. Doets, D. Mundici, and J. van Benthem, editors, *Logic and Scientific Methods - Volume One of the Tenth International Congress of Logic, Methodology and Philosophy of Science, Florence, August 1995*, Synthese Library, v. 259, pages 383–398. Kluwer Academic Publishers, 1997.

[KSJ97] K. T. Kelly, O. Schulte, and C. Juhl. Learning theory and the Philosophy of Science. *Philosophy of Science*, 64(2):245–267, 1997.

[KV94] M. J. Kearns and U. V. Vazirani. *An Introduction to Computational Learning Theory.* The MIT Press, Cambridge, MA, 1994.

[Lev80] I. Levi. *The Enterprise of Knowledge.* The MIT Press, Cambridge, MA, 1980.

[Lew69] D. K. Lewis. *Conventions. A Philosophical Study.* Harvard University Press, Cambridge, MA, 1969.

[LR57] R. D. Luce and H. Raiffa. *Games and Decisions.* John Wiley & Sons, New York, 1957.

[Lun97] F. Luna. Learning in a computable setting. applications of gold's inductive inference model. In H. Amman, B. Rustem, and A. Whinston, editors, *Computational Approaches to Economic Problems*, Advances in Computational Economics, v. 6, pages 271–288. Kluwer Academic Publishers, 1997.

[LV97] M. Li and P. M. Vitányi. *An Introduction to Kolmogorov Complexity and Its Applications, 2nd edition.* Springer-Verlag, New York, NY, 1997.

[Mic93] R. S. Michalski. Inferential theory of learning as a conceptual basis for multistrategy learning. *Machine Learning*, 11:111–151, 1993.

[Mil87] M. Milinski. Tit for tat in sticklebacks and the evolution of cooperation. *Nature*, 325:433–435, 1987.

[Mil99] T. S. Millar. Pure Recursive Model Theory. In E. R. Griffor, editor, *Handbook of Computability Theory*, Studies in logic and the fundations of mathematics, v. 140, pages 507–532. Elsevier Science B.V., 1999.

[MO97] E. Martin and D. Osherson. Scientific discovery based on belief revision. *Journal of Symbolic Logic*, 62(4):1352–1370, 1997.

[MO98] E. Martin and D. Osherson. *Elements of Scientific Inquiry.* The MIT Press, Cambridge, MA, 1998.

[MO99a] E. Martin and D. Osherson. Scientific discovery on positive data via belief revision. Submitted to *Journal of Philosophical Logic*, 1999.

[MO99b] F. Montagna and D. Osherson. Learning to coordinate: A recursion theoretic perspective. *Synthese*, 118(3):363–382, 1999.

[MS82] J. Maynard-Smith. *Evolution and the Theory of Games.* Cambridge University Press, Cambridge, 1982.

[MSP73] J. Maynard-Smith and G. R. Price. The logic of animal conflicts. *Nature*, 246:15–18, 1973.

[Nas51] J. F. Nash. Non-Cooperative games. *Annales of Mathematics*, 54:286–295, 1951.

[NYY92] A. Nerode, A. Yakhnis, and S. Yakhnis. Concurrent programs as strategies in games. In Y. Moschovakis, editor, *Logic from Computer Science.* Springer-Verlag, 1992.

[OdJMW97] D. Osherson, D. de Jongh, E. Martin, and S. Weinstein. Formal Learning Theory. In J. van Benthem and A. ter Meulen, editors, *Handbook of Logic and Language*, pages 737–775. Elsevier Science B.V., 1997.

[OR94] M. J. Osborne and A. Rubinstein. *A Course in Game Theory.* The MIT Press, Cambridge, MA, 1994.

[Osbng] M. J. Osborne. *An Introduction to Game Theory.* Oxford University Press, forthcoming.

[OSW86] D. Osherson, M. Stob, and S. Weinstein. *Systems That Learn.* The MIT Series in Learning, Development, and Conceptual Change, v. 4. The MIT Press, Cambridge, MA, 1986.

[OSW94] D. Osherson, M. Stob, and S. Weinstein. Logic and Learning. In S. J. Hanson, G. A. Drastal, and R. L. Rivest, editors, *Computational Learning Theory and Natural Learning Systems - Volume I: Constraints and Prospects*, pages 3–23. The MIT Press, 1994.

[PAM97] E. Plaza, J. L. Arcos, and F. Martín. Cooperative case-based reasoning. In G. Weiß, editor, *Distributed Artificial Intelligence meets Machine Learning*, pages 180–201. Springer-Verlag LNAI 1221, 1997.

[Put65] H. Putnam. Trial and error predicates and a solution to a problem of Mostowski. *Journal of Symbolic Logic*, 30(1):49–57, 1965.

[Rus97] S. J. Russell. Rationality and intelligence. *Artificial Intelligence*, 94(1-2):57–77, 1997.

[Sch98] K. H. Schlag. Why imitate, and if so, how? A boundedly rational approach to multi-armed bandits. *Journal of Economic Theory*, 78:130–156, 1998.

[SG00] L. Serafini and F. Giunchiglia. ML systems: A proof theory for contexts. Technical Report 0006-01, ITC-IRST, Trento, Italy, 2000. To appear in the Journal of Logic Language and Information.

[Sia91] S. S. Sian. *Learning in Distributed Artificial Intelligence.* PhD thesis, Imperial College, UK, 1991.

[Sim77] H. A. Simon. *Models of Discovery.* D. Reidel Publishing Company, Dordrecht, Holland, 1977.

[Sol64] R. J. Solomonoff. A formal theory of inductive inference. *Information and Control*, 7:7–22, 1964.

[Spe89] S. E. Spear. Learning rational expectations under computability constraints. *Econometrica*, 57(4):889–910, 1989.

[SS92] H. A. Simon and J. Schaeffer. The Game of Chess. In R. J. Aumann and S. Hart, editors, *Handbook of Game Theory*, volume I of *Handbooks in Economics, v. 11*, pages 1–17. Elsevier Science B.V., 1992.

[SSH94] S. Sen, M. Sekaran, and J. Hale. Learning to coordinate without sharing information. In *Proceedings of the Twelfth National Conference on Artificial Intelligence (AAAI-94)*, pages 426–431, Menlo Park, CA, 1994. AAAI Press.

[ST92a] Y. Shoham and M. Tennenholtz. Emergent conventions in multi-agents systems: Initial experimental results and observations. In C. Rich, W. Swartout, and B. Nebel, editors, *Proceedings of the Third International Conference on Principles of Knowledge Representation and Reasoning (KR-91)*, pages 225–231, Cambridge, 1992.

[ST92b] Y. Shoham and M. Tennenholtz. On the synthesis of useful social laws for artificial agent societies. In *Proceedings of the Tenth National Conference on Artificial Intelligence*, pages 276–281, San Jose, CA, 1992. AAAI Press.

[ST97] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1-2):139–166, 1997.

[Ste96] L. Steels. Self-organizing vocabularies. In C. Langton, editor, *Proceedings of the V Alife Conference*, Nara, Japan, 1996.

[vB98] J. van Benthem. Logic and games. Lecture notes for a graduate course. ILLC, UvA, Autumn 1998.

[VL97] P. M. Vitányi and M. Li. Minimum Description Length induction, Bayesianism, and Kolmogorov Complexity. *IEEE Transactions on Information Theory*, 1997.

[vNM44] J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior.* John Wiley & Sons, New York, 1944.

[Wal01] I. Walukiewicz. Pushdown Processes: Games and Model-Checking. *Information and Computation*, 164(2):234–263, 2001.

[Wei93] G. Weiß. Learning to coordinate actions in multi-agent systems. In Ruzena Bajcsy, editor, *Proceedings of the Thirteenth International Joint Conference on Artificial Intelligence (IJCAI-93)*, pages 311–316, San Mateo, CA, 1993. Morgan Kaufmann.

[Wei95] G. Weiß. Adaptation and learning in multi-agents systems: Some remarks and a bibliography. In G. Weiß and S. Sen, editors, *Adaptation and Learning in Multi-Agent Systems*, pages 1–21. Springer-Verlag LNAI 1042, 1995.

[WJ94] M. J. Wooldridge and N. R. Jennings. Formalizing the cooperative problem solving process. In *Proceedings of the Thirteenth International Workshop on Distributed Artificial Intelligence (IWDAI-94)*, pages 403–417, Lake Quinhalt, WA, 1994.

[WJ95] M. J. Wooldridge and N. R. Jennings. Agent theories, architectures, and languages: A survey. In M. J. Wooldridge and N. R. Jennings, editors, *Intelligent Agents*, pages 1–39. Springer-Verlag LNAI 890, 1995.

[WJ99] M. J. Wooldridge and N. R. Jennings. The cooperative problem-solving process. *Journal of Logic and Computation*, 9(4):563–592, 1999.