Chapter #

Obligations and Cooperation: Two Sides of Social Rationality

Guido Boella

Dipartimento di Informatica e Centro di Scienza Cognitiva – Università di Torino C.so Svizzera 185 10149 Torino, Italy
email: guido@di.unito.it

Abstract: Even if, at first sight, deciding whether to respect obligations and coordinating cooperation seems two unrelated behaviors, they are both grounded in the *social rationality* of agents: in fact, they depend on the ability of predicting the actions of other agents. In this paper, we present a proposal for managing *anticipatory coordination*, which is used to regulate the *autonomy* of and to *control* agents in a definition of cooperation among agents and in a framework for dealing with obligations.

Key words: autonomous agents, planning, deontic reasoning, cooperation, game theory.

1. INTRODUCTION

Cooperation and obligations[1] constitute situations where the autonomy of an agent and the possibility to control her are of primary importance. In fact, the agent has to make decisions in an environment where her choice directly affects the behavior of other agents - humans or artificial ones - who, in turn, may affect her performance.

In the first case, cooperation, the agent must perform her part of a shared plan, and she must autonomously decide how to perform it, while keeping an eye on the effects her choice has on the part of the shared plan assigned to her partners. Moreover, she has to decide whether to help them by *adopting* some of the goals they want to delegate in order to recover or, more generally, improve the group's performance - and, as a particular case, to decide whether to communicate with the partners. Finally, in order to evaluate the performance of the group, she has not only to plan her own part but also to predict what are the most likely actions of her partners.

In a similar way, in deliberate normative behavior, an agent has to decide whether to fulfill an obligation, knowing that some other agent (the *normative* agent) wants to *delegate* her a goal (the content of the obligation); the goal to be adopted may concern the execution or the exclusion of an action, or the achievement or avoidance of a state. Also in this situation, the agent has to evaluate the outcome of her choices not from a local point of view, but only after having considered the behavior of the other agent: in fact, our agent knows that, if she does not respect a obligation, a reaction of the *normative* agent (the *sanction*) will possibly follow. The reaction is only 'possible', because the *normative* agent is an autonomous agent, too: he has to decide what is

[1] In this paper, we prefer to use the term *obligation* instead of *norm*, since we focus on the behavior of a single agent subjected to an instance of a norm and we do not address the problem of accepting norms as such and of understanding whether the agent is subjected to a general norm.

rational for him to do and perhaps ascertaining violations and sanctioning them is not the best thing to do for him. Hence, our agent must model the autonomous decision process of the other agent.

What have these two scenarios in common? First of all, the agent is facing alternatives among which she has to choose the best one (for her or for some other entity, as the group). Some criterion for deliberating is therefore necessary. In particular, choosing the most promising course of action is important since what is involved is not only the benefit of the agent, but also the benefit of the other ones, since the agent may be considered as partially responsible for the success of the partner's actions.

Second, in choosing what to do, she has not only to consider the alternative ways of action stemming from her private goals, but she must also consider whether to adopt some goals which are proper of other agents: either belonging to the shared plan she participates in, or to some obligation issued by a *normative* agent. In both cases, even if she considers in the planning process also the *adopted* goals, this does not mean that the agent will do something for achieving them, but only that she must inspect them in order to verify if it is rational for her to try to achieve them.

Finally, the evaluation of the outcomes of the agent's choices must not be performed on the basis of the direct effects of the chosen actions. Rather, the resulting situation must be first updated with the predicted actions of the partners or of the *normative* agent involved. In the cooperation case, the interdependence of the agent's actions with those of the partner may cause a decrease/increase of the performance of the overall group. In deontic contexts, the agent's action must be possibly followed by the sanction performed by the *normative* agent if (he believes that) the obligation has been violated.

In this paper, we will describe our definition of cooperation and of deliberate normative behavior in an agent framework whose basic elements are:

- *decision theoretic* planning: the intention of the agent is selected among the plans which achieve her goals, by maximising the utility (in the decision theoretic sense) that their outcomes provide her with.

- *recursive modeling*: the selection of the most promising plan takes into consideration the behavior of the other agents. Game theoretic notions have been introduced in the context of decision theoretic planning in the form of anticipatory coordination. The aim is to improve the agent's own performance, her partners' one or, symmetrically, to worsen another agent's.

- *goal adoption*: the goals of the agent which are given in input to the planner are either her own private goals or goals *adopted* from those that other agents want to delegate, even if those goals are not directly desirable for the agent.

In the following we will describe how the first two issues are dealt with in our proposal and then we will present our definitions of cooperation and normative behavior where *goal adoption* plays a fundamental role.

2. DECISION THEORETIC PLANNING FOR AGENTS

In this Section, we introduce the definition of agent. We assume that an agent has a set of preferences and goals and that she does planning in order to find a (possibly partial) plan which satisfies one or more of these goals and maximizes the agent's utility. The chosen plan constitutes the current (individual) intention of the agent, an intention which respects the specification of [Cohen and Levesque, 1990]. Then, the plan is executed in a reactive manner, i.e., monitoring effects and triggering replanning in case of failure or new information.

When in the following we refer to *states*, we mean sets of attributes which are used to describe the world, together with a probability distribution over the values of the attributes at a certain moment. Moreover, we use also states where the probabilities of the values are not known to represent the total uncertainty of an agent about the current situation. Analogously, probability distributions will be associated to the effects of actions in order to express the fact that an agent is not sure about how the world changes after the execution of an action (either by himself or by someone else).

An agent A is a 6-tuple $\{IB, CG, f, L, KP, CI\}$ where:

- IB are the agent beliefs (including beliefs about other agents and beliefs about the current state of the world);
- CG is the set of current goals of A;
- f is the utility function of A (a function from states to real numbers); it is used to evaluate the outcomes of A's actions. f applies to states expressed as sets of attributes;[2]
- L is a set of tuples representing the obligations known by A of which she is the *bearer*;
- KP is the set of actions which A knows (*action recipes* [Carberry, 1990]);
- CI is the Current Intention to execute a plan (a newly planned plan or the remaining part of the plan A is currently executing). CI is selected within the set CP of candidate plans produced by the planner.

f is based on a set of attributes, each of which is associated with a utility value, and, by means of a combination function, produces the overall desirability of a state on the basis of its description. The action may affect the utility of the resulting state in a positive or negative way; in particular, a decrease in the utility values is used for representing the costs of executing the action in terms of time and resource consumption.

The planner takes as input goals consisting in states or actions (*state goals* and *action goals* in [Castelfranchi, 1998]'s terminology): in case the goal is a state, it is considered as a state to be achieved, so that the planner must find all actions which can contribute to achieving the state; in case the goal is an action, the planner assumes that it is a complex action which has to be executed, so that its (easier) task is to find all possible decompositions of (i.e., ways to carry out) the task and to choose the best of them. The latter activity is called *refinement* of the action.

The set CP of the agent A is produced by the planner starting from the initial state S, and inspecting the KP to find all the action recipes which have among their effects a goal in CG and the recipes which refer to (expand) an action in CG. Then, on the basis of the utility function f, the possible alternatives are examined and the best one (P) is chosen: it becomes the current intention of the agent. The best plan is the one which maximizes the utility:

$$P = argmax_{\{P_i \in CP\}} f(P_i(S))$$

where $P_i(S)$ is the state resulting from the execution of the plan P_i in the state S.

In case of probabilistic effects and uncertain states, $P_i(S)$ is a set of pairs $(S_{i,j}, p_{i,j})$, where $p_{i,j}$ is the probability that $S_{i,j}$ is reached after the execution of P_i. In this case, the best plan is the one which maximises the *expected* utility, as it is usual in Decision Theory:

$$P=argmax_{\{P_i \in CP\}} \sum_{(S_{i,j},p_{i,j}) \in P_i(S)} p_{i,j} * f(S_{i,j})$$

[2] f is based on the theory of *multi-attribute utility functions* developed by [Keeney and Raiffa, 1997]. For a discussion on the usefulness and restrictions of multi-attribute functions see [Haddawy and Hanks, 1998].

However, this decision procedure is not enough for building an agent architecture. As [Schut and Wooldridge, 2000] notice, decision theory is not directly applicable to bounded agents. It must be integrated in a planning framework, since the agent has not at his disposal the plans P_i or they are to many to be considered with respect to the available time.

A recent solution is [Haddawy and Hanks, 1998], where they described a way to relate the notions of goals and planning to that of utility in decision theory. So, in order to develop our approach, we adopted the decision theoretic planner DRIPS described in [Haddawy and Suwandi, 1994], and we built our proposal on the top of it.

DRIPS is a hierarchical planner (with both an abstraction and decomposition hierarchies relating actions) which merges decision theory with standard planning techniques. A utility function is used for evaluating how promising a given plan is with respect to an agent's preferences. In particular, a *multi-attribute* utility function is used which does not compute just the payoff of each possible outcome considered as a whole: the expected utility is computed starting from functions representing the preferences of the agent about single attributes of the world[3]; since some of these attributes appear also among the goals and the action effects, the expected utility depends directly on goal satisfaction.

Moreover DRIPS allows to model different degrees of achievement of a goal (satisfying a goal to a greater extent is considered preferable to satisfying it to a lesser extent) and the success in satisfying one goal component is traded off against the success in satisfying another goal, or against consuming resources.

Therefore, by using multi-attribute utility functions, we can take into account both the individual utility of an agent and the utility of performing a plan for achieving the goal, given the related costs - in terms of resource consumption. Since partial satisfaction of goals is accounted for, it is possible to compare how the alternatives contribute to the satisfaction of the shared goal at different degrees and how different amounts of resources are consumed: the multi-attribute utility represents the rationality criterion of the action.

Finally, DRIPS is based on an anytime algorithm [Zilberstein and Russell, 1993]. That is it always produce that best solution available and the quality of this solution is a function of the time devoted to computing it. As [Helwig and Haddawy, 1996] show it is possible to use the planner for refining partial plans so long as the agent has time to do so. When the time allocated to computation is over, the planner can return the agent the best partial plan constructed so far, so she can execute it.

However, DRIPS deals only with single-agent plans and individual utility. So its basic planning mechanism was extended to deal with *anticipatory coordination*. As we discuss in the next Section.

[3] If some independence assumptions hold (see [Haddawy and Hanks, 1998]).

3. RECURSIVE MODELING

According to [Castelfranchi, 1998], in order to prevent damage or losses of resources and to exploit opportunities, some form of *anticipatory coordination* is needed, based on *mind reading*: «the understanding of the goals and the plan of the other [Castelfranchi, 1998]». In our model, we will express this «co-evolutionary coupling between planning in a multi-agent world and the mind-reading ability» by means of [Goffman, 1970]'s notion of *strategic interaction.*

For Goffman, «individuals, like other objects in this world, affect the surrounding environment in a manner congruent with their own actions and properties. Their mere presence produces signs and marks. Individuals, in brief, exude expressions» (p.4). On the other hand, «all organisms after their fashion make use of information collected from the immediate environment so as to respond effectively to what is going on around them and to what is likely to occur» (p.10).

In fact: "Aware that his actions, expressions, and words will provide information to the observer, the subject incorporates into the initial phases of this activity a consideration of the informing aspects of its later phases, so that the definition of the situation he eventually provides for the observer hopefully will be one he feels from the beginning would be profitable to evoke" (p.12).

In this way, agents engage in «expression games», resorting to their empathy and ability to «take the attitude» (p.13) of the other observer, as G. H. Mead has remarked: «the agent takes the viewpoint of the observer, but he does not 'identify' his interests with it». He does so «only insofar as the observer is engaged in observing him and ready to make decisions on this basis, and only long enough and deep enough to learn from this perspective what might be the best way to control the response of the person who will make it».

So, when an agent considers which course of action to follow, before he takes a decision, he depicts in his mind the consequences of his action for the other involved agents, their likely reaction, and the influence of this reaction on his own welfare. He will adapt his actions to the other agents' reaction before it can even happen. This form of reasoning is what Goffman calls *strategic interaction*.

As an analytical tool for modeling deliberation in situations of strategic interaction, Goffman proposes *Game Theory* [von Neumann and Morgenstern, 1947]. In fact, Game Theory enables one to base a decision not only on the *expected payoff* of an action, but also to model in a structured way the situation of other agents, to predict their rational reaction, and finally, to choose what to do on the basis of the predicted possible final outcomes and their utility for the agent.

Hence, we chose to model this form of reasoning by means of *Game theory,* too. In particular, we integrate game theoretic concepts with planning - as for example, [Gmytrasiewicz and Durfee, 1995], [Ndiaye and Jameson, 1994] and [Boella, 2000] do: we have extended in a game theoretic perspective the planner DRIPS of [Haddawy and Suwandi, 1994] which integrates the decision theoretic notion of *utility* (a numerical representation of preferences) with classical goal-satisfaction planning.

As discussed above, it is not sufficient to take into account the resulting state $P_i(S)$ after an action P_i, but it is also necessary to consider the possible subsequent behavior of the other agents starting from $P_i(S)$. For instance, in a cooperative setting, it may happen that a state very positive for the agent endangers the activity of the interactants, so that the overall (group) goal is harder to achieve. Or that a state that can be achieved easily is then made undesirable by a sanction. Our solution has been to base the evaluation not on $P_i(S)$, but on the states achievable by the interactant B starting from $P_i(S)$ (a kind of one-level lookahead in the spirit of min-max search).

So, in presence of anticipatory coordination, A will select the plan P in the set of candidate plans CP such that:

$$P = argmax_{\{Pi \in CP\}} f(P^{best_{B_{P_i}}}(P_i\ (S)))$$

where $P^{best_{B_{P_i}}}$ is the plan that (according to A's beliefs) will be selected by B when A executes P_i in state S. The decision about $P^{bestB_{P_i}}$ will be taken by B according to some private utility function or by taking into account also the benefit for the group.

In presence of anticipatory coordination, we must take into account actions which have non-deterministic outcomes. In this case, $P_i(S)$ is a set of states with associated probabilities. When B plans her reaction, she will be in a specific state of $P_i(S)$ (since A will have already executed the action he chose). Therefore, A has to simulate B's reaction in each of these states. In this situation, $P^{best_{B_{P_i}}}$ will be a set of (state, probability, plan) tuples (the probability is the one of the state in $P_i(S)$ from which the associated plan has been planned); the above formula must be modified in:

$$P=argmax_{\{Pi\in CP\}} \sum_{(S_{i,j},p_{i,j},P^{B}_{i,j})\in P^{best_B}_{Pi}} p_{i,j}*f(P^{B}_{i,j}(S_{i,j}))$$

Again, the formula above must not assume that the agents *A* and *B* already have at their disposal the set of plans involved. Hence the formula must be translated into a planning algorithm for computing and selecting the plans in an efficient way, as we said in the previous Section.

The construction of a plan by the planner is carried out by an agent *A* in a stepwise fashion: if *A* has to refine an action β^A for achieving goal φ, she first has to find the recipes for β^A (let's say R^y_i, $1\leq i\leq t$) and then she can start refining them. The approach of DRIPS is to expand β^A in all possible ways; then, it proceeds onward and expands the new partial plans. The search goes on in parallel, but the search tree is *pruned* using the utility function (applied to the states resulting from the potential execution of the recipes to the current state *S*, a set of attribute-value pairs); it acts as a heuristic able to exclude some (suboptimal) ways (recipes) to execute an action, while the algorithm allows to find an optimal solution.

In order to implement anticipatory coordination, we had to make the evaluation of the heuristics somewhat more complex: before the evaluation of an action outcome is carried out, the outcome is updated with the effects of the interactant's reaction which the agent tries to predict via a *recursive* modeling of the planning and decision making activity of her interactant about his part of the shared plan.

First of all, *A*'s plans may have non-deterministic outcomes: they are a set of probability-state pairs $S'_i = \{(p'_{i,1}, S'_{i,1}), \ldots, (p'_{i,n}, S'_{i,n})\}$. The *anticipatory coordination* taking into account *B*'s reaction must be applied from each of these different states, and then the expected utility is computed by combining the probability of each $S'_{i,j}$ with that of the outcomes of *B*'s non-deterministic action.

Second, *B*'s point of view must be constructed, starting from the outcomes of each *A*'s alternative R^y_i. In many situations, in fact, *B* is not aware of all the effects of *A*'s action. In this proposal, a STRIPS-like solution to this subjective form of the *frame problem* is adopted; *B*'s knowledge of a state is updated by an action of *A* only with the effects which are explicitly mentioned as believed by *B* in the action definition.[4] Communicative actions with *B* as a receiver are a particular case of actions affecting the partner's beliefs. The attribute values of the initial state that are not affected by changes in the interactant's beliefs are by default propagated. Two different solutions are possible: either the beliefs of the two agents in the initial state are assumed to be the same or the agent has a representation of the beliefs concerning the initial state attributed to the partner.

Since not all states are distinguishable from his point of view, *B*'s beliefs are represented by means of the *Game Theoretic* notion of the *information sets* which abstract away the information not known by *B*:

1. Given $S'_i = \{(p\text{☏}_{i,1}, S\text{☏}_{i,1}), \ldots, (p'_{i,n}, S\text{☏}_{i,n})\}$, create the set of equivalence classes $\hat{S}_i = \{\hat{S}_{i,1}, \ldots, \hat{S}_{i,r}\}$ where each class $\hat{S}_{i,j}$ contains those states of S'_i which are not distinguishable from *B*'s point of view (in practice, which contain the same attribute values).

[4] This solution is clearly a simplification. [Ideki and Hirofumi, 2000] presents a more complex methodology for reasoning on changes in the beliefs of other agents.

2. From each $\hat{S}_{i,j}$ in $\hat{S}_i$, pick up a state S'_{i,j_1} and produce B's point of view $S^B_{i,j}$ as described above.

3. The probability of each $S^B_{i,j}$ is derived as the sum of the probabilities of the states contained in the equivalence classes $\hat{S}_{i,j}$ ($p'_{i,k}$ is the probability of state $S'_{i,k}$ in $\hat{S}_{i,j}$):

$$p^B_{i,j} = \sum_{S'_{i,k} \in \hat{S}_{i,j}} p'_{i,k}$$

4. B's point of view is the set of probability-state pairs ($S^B_i = \{S^B_{i,1}, \ldots, S^B_{i,w}\}$) constructed in the previous steps.

As a consequence of B's different point of view, the outcomes of the simulation cannot be directly used in the evaluation of the expected utility, since they can reflect an incomplete point of view. Hence, A takes the actions chosen by B from each state and re-computes the resulting outcomes when they are executed from her own point of view.

The evaluation of A's alternative R^y_i with outcome S'_i, under the light of *anticipatory coordination* is made in the following way.[5]

1. From S'_i (the result of recipe R^y_i), form the corresponding information sets and B's point of view $S^B_i = \{S^B_{i,1}, \ldots, S^B_{i,r}\}$;

2. On each state $S^B_{i,j}$ in S^B_i:

   (a) Restart the planning process from the perspective of the interactant B and try to solve his current task β^B for achieving a given goal. The result is a set of candidate recipes $\{R^u_1, \ldots, R^u_v\}$.

   (b) For each l ($1 \leq l \leq v$) by means of the recipe R^u_l expand the state $S^B_{i,j}$ obtaining again a set of (probability, state) pairs $S'^B_{i,j,l} = \{(p'^B_{i,j,l,1}, S'^B_{i,j,l,1}), \ldots, (p'^B_{i,j,l,r_l}, S'^B_{i,j,l,r_l})\}$.[6]

   (c) For each l ($1 \leq l \leq v$), the utility function f^B of B is applied to these sets of (probability, state) pairs, and the plan $R^{best}_{i,j}$ which maximizes the following formula is the one selected by agent B for execution in $S^B_{i,j}$ (its outcome is $S^B_{i,j,best_{i,j}} = \{(p'^B_{i,j,best_{i,j},1}, S'^B_{i,j,best_{i,j},1}), \ldots, (p'^B_{i,j,best_{i,j},r_i}, S'^B_{i,j,best_{i,j},r_i})\}$):

$$\sum_{S'^B_{i,j,l,z} \in S'^B_{i,j,l}} p'^B_{i,j,l,z} * f^B(S'^B_{i,j,l,z})$$

3. Expand each state $S'_{i,j}$ in S'_i with the recipe $R^{best}_{i,e}$, where $\hat{S}_{i,e}$ is the equivalence class in $\hat{S}_i$ which $S'_{i,j}$ belongs to; for each j, the result is a set of (probability, state) pairs: $S'^A_{i,j} = \{(p'^A_{i,j,1}, S'^A_{i,j,1}), \ldots, (p'^A_{i,j,qi}, S'^A_{i,j,qi})\}$

4. Given the n initial states $S'_{i,k}$ in S'_i, the probability of each state $S'^A_{i,j,l}$ is $p'_{i,k} * p'^A_{i,j,l}$ (the latter depends on the probability distribution of $R^{best}_{i,j}$ effects). Consequently, the expected utility (according to A's utility function f^A) of the initial states S'_i is:

[5] For the sake of simplicity, in this paper we do not consider the dimension of *uncertainty*, which is necessary for the DRIPS's pruning process: in some cases, B may have no idea about the probability distribution on the different outcomes.

[6] The process is repeated on further refinements of the recipe, as DRIPS does.

$$\sum_{j=0}^{n} \sum_{\{S'^{A}_{i,j,l} \in S'^{A}_{i,j}\}} (p'_{i,k} * p'^{A}_{i,j,l}) * p' * f^{A}(S'^{A}_{i,j,l})$$

Note that the algorithm above is a modification of a two-level min-max algorithm: actually, it is a max-max, since at both levels the best option is selected, although at the second level it is evaluated from B's perspective. As in min-max, *A*, when predicting B's behavior, assumes that her partner is a rational agent, i.e., that he will choose the plan that gets the highest utility. The algorithm scales to further levels of depth, but, as [Gmytrasiewicz and Durfee, 1995] notice, the cognitive load of the agents would become implausible.

In Figure 1, an example of a trace of the algorithm is depicted. In state $S0^A$, the agent A has three alternative action to choose among: a1, a2, a3; note that a3 is a nondeterministic action which has two outcomes according to a probability distribution. However, the choice about what to do is not made by evaluating these outcomes $S1^A$, $S2^{A,}$ $S3'^{A}$, $S3''^{A}$, but by simulating first the behavior of the interactant B. First, for each outcome, B's point of view is constructed $S1^B_A$ from $S1^A$, etc. (assume each outcome corresponds to a different information set of the interactant). From each state, the behavior of B is simulated. Note that, in particular, in the branch corresponding to action a3, B's behavior is simulated in a different way according to each possible outcome of a3. If we focus on the last outcome $S3''^{B}$, we see that B has two feasible alternative b3''1 and b3''2. In A's simulation, according to their respective outcomes and the utility function f^B, B will choose b3''1 (see the dark circle). This information is used by A to update the outcome of the current branch: the utility of a3 must be computed using not $S3''^{A}$, but it must be updated with the effect of executing b3''1 by B. However, A cannot, in general, use the states $S3''^{B}_{A}$, since it is biased by B's perspective on the original state $S3''^{A}$. Hence, the real outcome of b3''1 must be computed by executing it from $S3''^{A}$, thus reaching $S3''1'^{A}$ and $S3''1''^{A}$. Finally, the respective probability of these two states (0.5) must be combined with that of the branch of $S3''^{A}$, i.e. 0.4.

The mechanism described above for the choice of the best recipe is computationally expensive. However, a first step toward a more efficient solution is achieved by exploiting the DRIPS mechanism of pruning the search tree when a partial plan looks unpromising compared to the other hypotheses: this algorithm is applied both when the agent's plans and her partner's ones are devised. Moreover, the use of information sets reduces the number of simulation runs necessary for anticipatory coordination.

Finally, since our proposal admits partial plans, DRIPS can be stopped after it has reached a certain level of detail without expanding the plan completely. In reactive planning, as [Bratman, 1987] noticed, agents limit the search for solutions to partial ones, because working in a dynamic world makes overdetailed plans often useless.

4. GOAL ADOPTION/DELEGATION

[Castelfranchi, 1998] has posed *goal adoption* in the basic ontology of *social rationality*: «an agent *B* performs an action for a given goal, since this action or the goal are included in the plan of another agent *A*».

This notion is basic in case of cooperation, where it is part of the explanation of the commitment of an agent towards the actions of the partner, explanation which does not resort to the notion of intention or *intention-that*

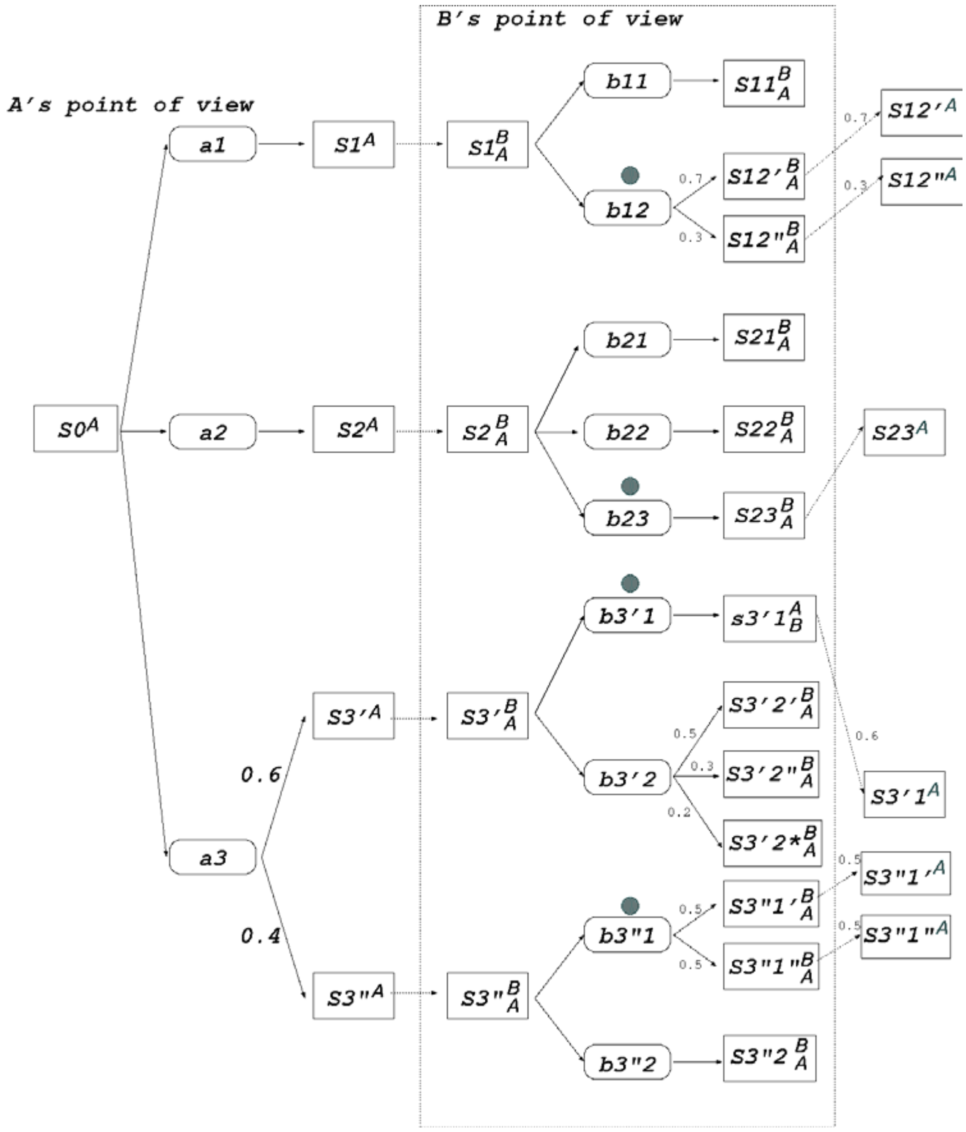

Figure 1. : The recursive modeling of the partner's choice.

towards another agent's action. In case of deliberate normative behavior, goal adoption explains the exogenous origin of the goal of fulfilling a norm.

4.1 The definition of cooperation

A group *GR* composed of agents $G_1, \ldots, G_n$ cooperates to a shared plan α^x for achieving a goal φ (with an associated plan R^x composed of steps $\beta^{x,k}{}_{11}, \ldots, \beta^{x,k}{}_{qm}$), when:[7]

1. each step $\beta^{x,k}{}_l$ has been assigned to an agent G_k in *GR* for its execution;

2. each agent G_k of the group *GR* has the single agent intention to perform her part $\beta^{x,k}{}_l$, an intention relative (in [Cohen and

[7] The notation $\beta^{x,k}{}_l$ refers to the *l*-th step of the plan R^x, a step which has to be executed by agent G_k.

Levesque, 1990]'s sense) to the existence of the group shared plan[8];

3. all the agents of *GR* have the mutual belief that each one (G_k) is engaged in cooperating to φ with *GR* by means of plan R^x;

4. all the agents mutually know that they share a utility function f_{GR} based on a weighed sum of the utility of the goal φ, which the shared plan aims at, and of the resource consumption of the single agents;

5. when an agent G_k becomes aware that a partner G_j has a goal ψ that stems from his intention to do his part $\beta^{x,j}{}_p$, G_k will consider whether to adopt it;

6. each agent G_k remains in the group so long as the group's expected utility of going on in performing $\beta^{x,k}{}_m$ for φ or adopting some of the goals of her partners is greater than the expected utility of doing nothing more for the group.

For what concerns point 5, the goals that are adopted by an agent G_k are the subgoals ψ which G_j has formed while planning how to perform his part and while executing them in a reactive way. Therefore, G_k considers not only the steps of $\beta^{x,b}{}_p$ she may execute to assist G_j in performing his task, but also G_j's goals deriving from his single agent intention to perform $\beta^{x,j}{}_p$: knowing how to perform $\beta^{x,j}{}_p$, achieving its preconditions, monitoring the effects of its execution.
Symmetrically, the agent can rely on the fact that she can delegate a goal of herself if they can be better achieved by her partners.

It must also be observed that at point 5 the term «aware» is used. So, *mind reading*, in the form of plan recognition techniques (e.g., [Carberry, 1990]) can play an important role in helping agents to infer what their partners are doing and therefore improving the cooperation of the group.

Moreover, the control of the group is not necessarily in charge of a single agent, but it is distributed among the agents, as long as they have the relevant information about what the other agents can do. In this way, *mixed initiative* arises as a natural phenomenon in our framework.

In the definition there is no mention of *anticipatory coordination*: this form of reasoning is, indeed, implicit in the planning algorithm of the agents described above. Therefore, *anticipatory coordination* is not by itself a characterizing property of cooperation, but a general property of *socially rational agents*.

[8] This does not mean that G_k has to execute the action by herself: in fact, she could delegate this action to another partner or contract it out; she remains, however, responsible for ensuring the success of her part.

4.2 Cooperation phenomena explained

Models of cooperation are required to explain a number of phenomena ranging from helpful behavior to conflict avoidance via communication for coordination purposes.

Helpful behavior: helpful behavior (i.e., *goal adoption*) is at the basis of cooperation: the agent considers the goals of the partners and, only if it is useful for the group, she (tries to) satisfies them. Helpful behavior should be provided not only when the group cannot otherwise proceed in its plan, but, also, it should not be provided indiscriminately. By exploiting the decision theoretic paradigm, it is possible to keep apart the forms of help that are rational from those that appear to be just a waste of resources. If the effort for satisfying adopted goals conflicts with the ability of the agent to do her own part, then, a decrease in the group utility is obtained, instead of an increase. The limitation of adoption only to «rational cases» rules out the possibility that an agent spends all his time in adopting goals of the other partners.

Communication: As a special case of helpful behavior, it is possible to predict various forms of communication: they arise when the goal to be adopted is a *control goal*; that is, a goal of *B* to know some information, like the truth value of constraints and preconditions of actions, and whether an action succeeded or not. The effect of a communicative action is to make such a goal true (at least with some probability).

In this model, communicating that the shared plan is successful, impossible to execute or irrelevant - a behavior which has a great relevance in group cooperation - is just a special case of help concerning the control goal that an action has been executed. Furthermore, beside communicating the success or failure of the entire shared plan, the members of a group can consider also the success of the part of the shared plan assigned to the single members, since it is also a control goal of them.

Communication allows to «keep the team acting as a unit», but it does not always worth its cost for the group. The agent *A* has to consider the cost of communicating with *B*. If communication is expensive, then it is not convenient for the group to waste resources in kindly communicating, since *B* could discover the desired information in a less expensive manner. The same holds if communication is not reliable (the message can get lost) or slow: there is a probability that communication has not the desired effect (or the receiver gets it too late). In these cases, the utility resulting from a successful communication must be combined with the utility that is achieved when communication fails and the partner wastes the group's resources anyway. Our model is able to cope with all of these aspects: *A* compares the communicative alternative with the other ones, and can foresee what *B* will do in all these situations. If the communicative action succeeds, *B*, e.g., will choose not to go on in his task, since, under the light of the new information, he can evaluate that continuing it does not produce any utility for the group. Moreover, even if an agent decides that is better (for the group) not to communicate, her choice does not disrupt the group: in fact, neither communication nor mutual beliefs are explicitly prescribed by the definition of cooperation.

This approach has the advantage that it does not require the assumption of perfect communication: that is, communication actions do not need to achieve the desired goal necessarily, but they can be faulty. On the contrary, [Cohen and Levesque, 1991] assume that communication never fails, since otherwise the joint intention would be disrupted when an agent fails in notifying to the other partners that he succeeded.

As [Castelfranchi, 1998] noticed, communication at the end of the group activity, as prescribed in [Cohen and Levesque, 1991], is sometimes even irrational. The exemplar case is when the group is in a context of conflict and the internal exchange of information causes the adversary to become aware of important information.
For instance, the group can have the goal not to let other agents know when they achieved their goal: if they are looking for a certain object, communicating that they have found it would mean for the adversary to know who has the object. In this model, the group's goal of preventing the opponents from knowing an information - a side effect of communication - can be included in the group's multi-attribute utility function shared by the members.

No benevolence: group's utility does not mean benevolence, but only a criterion for enhancing the group's performance and provides a means for regulating the agent control: the definition prescribes what an agent in a group can be accounted for by the partner if she fails to conform to it. The definition expresses the rationality of cooperation, and not what it is rational for the individual agent to do from her self-interested point of view. We will return on this issue in Section 6.

Hierarchical groups: the consumption of resources need not be weighed in a uniform way for all members of the group; a sort of hierarchy in the group can be induced by weighing - in the multi-attribute utility function - resource consumption differently depending on the agent who executes an action. For example, if *A*'s communication is more costly than *B*'s one and *A* first succeeds in knowing that the goal holds, she may not communicate this fact to *B*, while *B* would notify her when he becomes aware of that.

Conflict avoidance: since agents share a group utility function and perform *anticipatory coordination*, they will (try to) avoid conflicts with other agents' intentions: performing an action that interferes with the plans of other team members decreases the utility of the whole team.

When *A* considers the possible developments of her partial plan, she examines what effects her action will have on the partners' plans. So also the possible interferences are weighed as any other cost that decreases the group utility: conflicts result in less preferred choices, but they are not necessarily ruled out.

For example, if two partners who are preparing a meal together have only one pan and one of them needs it urgently, she can decide to use it - even if she knows that her partner will need it later - because this is more convenient, for the group, since the pan can be easily washed. On the contrary, if the shared resource is not reusable, like for example eggs, then *A* will use it only if she cannot change his plan without a significant decrease of utility.

Contracting out: one of the requirements of cooperation posed by [Grosz and Kraus, 1996] is that an agent of the group can contract out her task without the contracted agent becoming a member of the group. An agent *A* can delegate a step β^{x,a_l} of her own part to another agent *C* that does not belong to the group *GR*; *A* and *C* will form a new group sharing the goal to perform β^{x,a_l} and a corresponding utility function. In any case, *C* does not become a member of the group *GR*: in fact, he will not necessarily know which are the group's shared goal and utility function. It is possible, therefore, that *C* interferes with *GR*. *C*'s behavior may lead to a decrease in the performance of the group, but he cannot be accounted for interfering with the group activity, since *C* is not requested to know the ongoing shared plan.

Ending cooperation: when all members know that the top-level goal of the group has been achieved, has become impossible or irrelevant, then no more utility can be obtained by any other actions than terminating the group: in fact, termination gets higher utility by saving resources. Therefore, the shared plan is naturally ruled out, without the need of stipulating other explicit conditions for its termination.

Another way for an agent to close cooperation is to opt out from it. This can happen if increasing the utility of the group produces an unacceptable decrease with respect to the agent's private utility. However, as [Grosz and Kraus, 1996] notice, this is not an harmless choice since the other agents can retaliate for being abandoned (for example by not helping her in future situations). When the cost of remaining in the group is greater than the consequences of leaving it, the agent will choose to pursue his private goals.

5. THE DEFINITION OF OBLIGATION

In principle, an obligation is something an agent is obliged to do. In other words, given an initial situation, in any course of events produced by the agent chosen action(s) the obligation must be fulfilled. However, this need not be the most rational way for an agent to act. There can be situations where different obligations contrast with each other, or situations where an obligation cannot be reconciled with the agent's personal desires or goals. In these cases, the agent must evaluate carefully what to do and must decide if the obligation (or, which obligation) must be pursued, and in which way.

For examining the problem of deciding whether to fulfill an obligation, we adopt an approach which is inspired to E. Goffman's work in sociology: we focus our attention on the fact that an obligation involves at least two individuals, both of which have to be modeled as (intelligent) agents: the *bearer* of the obligation, who must respect the obligation, and the *normative* agent, which has posed the obligation, wants that the *bearers* of the obligation fulfill it, and (possibly) will sanction the violators.

While «it is generally acknowledged that norms and normative action emphasize autonomy on the side of *decision*» [Castelfranchi, 1998], no attention has been paid to the fact that norms and obligations are enforced by the *normative* agent, who is an autonomous agent, too. In fact, up to now, the center of attention has been the bearer of the obligation. A remarkable exception [Dignum, 1996] explicitly deals with sanctions, but it does not model the agent who is in charge of monitoring violations and applying the sanction. The first consequence of our approach is that, when he considers whether to fulfill the obligation, the *bearer* of the obligation has to consider explicitly the disadvantage of facing the sanction. [Axelrod, 1986] makes a similar proposal in economics.

According to Goffman, «a norm is that kind of guide to action that is supported by social sanctions» (p.62). And A. Giddens adds: «a sanction is defined as a reaction of others to the behavior of an individual or a group, a reaction having the goal to enforce the respect of a given norm» [Giddens, 1987], p.120.

Hence, from the point of view of sociology, norms come always together with sanctions; since sanctions are actions, they presuppose, in turn, someone to perform them.

As Goffman has noticed, an agent who has to follow some norms can be considered as a player in a game, where the payoffs of his actions depend on the subsequent actions of another social agent, a second player who has the role of making the obligations respected.

In fact, the sanction is not a granted exogenous event, but it is the result of the activity of the *normative* agent. He has the goal of checking the fulfillment of the norm and has a plan for doing so and eventually posing the sanction. But she also has other goals, preferences and obligations as any other agent. The *bearer* of the obligation has to take into account all these facts when she considers the advantage of fulfilling or not the obligation: i.e., she has to model (recursively) also the *normative* agent as an agent.

On the other hand, the recursive modeling of the *normative* agent opens the way to another opportunity for the *bearer* besides a better evaluation of the resulting final state. The *bearer* agent can reason about how the *normative* agent will (decide to) check the fulfillment and will apply the sanction if he discovers a violation. This knowledge can enable an agent to predict when the *normative* agent possibly fails to become aware of a violation and/or enable him to devise a way to induce her to this failure by means of some action.

More specifically, in our model *an obligation consists in a situation where an agent B has a goal φ that another (or more than one) agent B satisfy a goal φ', and who, in case the agent B acts without adopting the goal φ', has to decide whether to perform an action $\beta^{x,n}{}_i$ which (negatively) affects some aspect of the world which (presumably) interests B (as represented by her utility function). Both agents know these facts.*

An obligation Ω is represented as a 4-tuple *{B, B, O, R}* where:

- *B* is an agent who is called the *bearer* of the obligation,
- *B* is an agent called the *normative* agent,

- is the content of the obligation, i.e., the state or action goal which *B* wants to be adopted by *B*,

- *R* is an action (called sanction) which *B* will presumably bring about in case he detects a violation of the obligation.

The content of the obligation *Ω*, *O*, is not necessarily a state (e.g., «the font of the submitted paper should be courier»), but it can be also an action where *B* is the agent (e.g., «the author should send a signed copyright form») or not (e.g., «the head of the department of the author should authorize him to participate to the conference»). Finally, it can be the prescription of not executing an action: «you should not send the submitted paper to other conferences».

Differently from what appears at first sight, this definition covers not only 'institutional' cases, but also other situations like obligations in dialog (see [Ardissono et al., 1999] and [Boella et al., 2000b]); they share the characteristic that new goals are acquired as a consequence of social inputs.

This definition allows BDI (Beliefs, Desires and Intentions) agents to deal with obligations since they are able to manage intentions, to take into account the goals of other agents and their behavior, to devise plans for satisfying goals, and to compare the alternative plans according to their preferences.

Again, in the definition of obligations there is no mention of *anticipatory coordination*. Since our *socially rational* agents use the planning algorithm described in Section 3, they are able to foresee the reaction of the *normative agent* both in case they adopt the content of the norm and they don't.

The *B* counterparty of an agent *B* who is the *bearer* of an obligation cannot be assumed to become immediately acquainted with the (possible) violation of the obligation. According to *B*'s knowledge, this happens: in fact, *B* is assumed to know that *B* has some actions available to check the fulfillment of *O*, that these actions may fail, or can be induced to fail by some *B*'s action, and that just in case of their success, *B* will consider (not necessarily decide) to apply the sanction.

For this reason, in our model, the sanction *R* is represented by a complex action consisting of the monitoring action followed by the very action of affecting the world in a way that is relevant (in one way or the other) for the *bearer* agent.

We have chosen to insert the monitoring action in the definition of an obligation although it would have been sufficient that the sanctioning act had as a precondition the belief of the normative agent that a violation has possibly occurred. But it must be observed that in formal contexts, say laws, not all ways for knowing that a violation occurred are acceptable. For example, the Italian police is not allowed to tape phone calls without the authorisation of an attorney. Illegal recordings do not count as evidence in trials to determine that a violation actually occurred. Therefore, the specification of a norm includes the specification of the possible means for checking violations. In case every means is allowed, the monitoring action will consists of a general action which subsumes all the possible checking actions for a given type of violation.

Finally, it must be noted that the sanction itself can establish another obligation for the violating agent. For example, a policeman can put the obligation of paying a fine for having parked in a no parking area, where the *normative* agent's role is taken by some administration. In turn, the failure to respect the deadline for paying the fine results in another obligation of being brought to trial before another kind of *normative* agent, a judge. And so on in a *crescendo* of more and more negative payoffs for the agent.

However, there is no risk of a regression to infinitum, since this series of obligations usually ends in a sanction which does not pose a new obligation, as Goffman notices: "still steeper penalties should their judgement be rejected, and still deeper penalties should this, in turn, be rejected and so forth, eventually culminating presumably in physically coerced rulings" ([Goffman, 1970], p. 117).

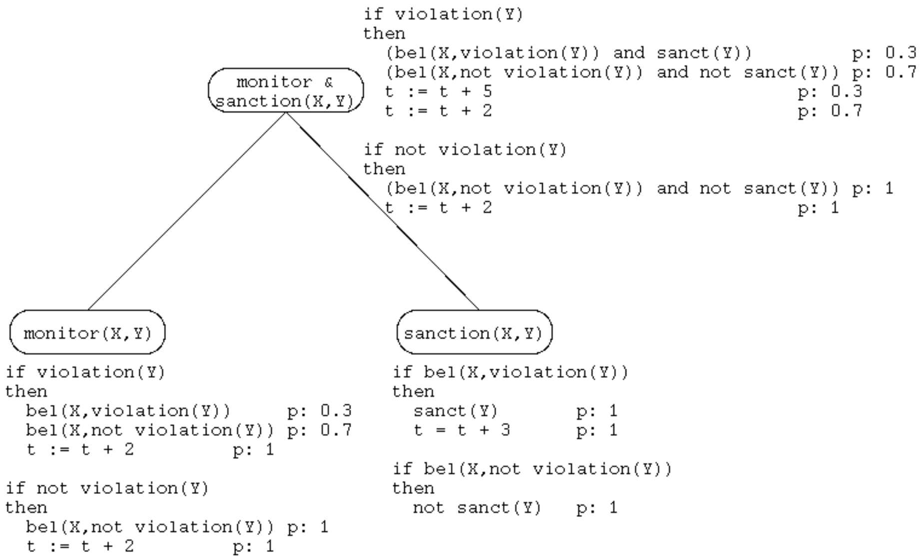

Figure 2. Monitoring and sanctioning an agent.

5.1 The Behavior of the Normative Agent

As stated above, the monitor and sanction action R leads to an actual punishment only if the ‘monitor’ part enables the *normative* agent to *know* (or at least to believe) that there is a violation. In the general case, the *normative* agent is (or may be) unaware of the violation and he has to decide whether to monitor and (possibly) sanction her or to do other things, according to other goals. In order to make his choice, he has to compare the different utilities of his alternatives. Also sanctioning actual violations may provide him with an utility: but why should he choose the monitor and sanction action when he is not aware of the violation? That is, if the effect of this action depends on the actual state of the world which is not known by the agent, how can he evaluate the utility? What is important for this discussion is the role of the action during planning and not during execution by the *normative* agent. During execution, the world can be in one and only one state, and the monitoring action will reflect it. But during the planning process, the *normative* agent has no access to the actual world. The evaluation of the result of an action and of its utility must be made entirely from his belief space. And the *normative* agent may either not know anything about the violation or she may have some *a priori* idea of what happened. We have to find a framework for representing both situations.

In order to deal with this problem, we resort to the probabilistic framework exploited in [Haddawy and Hanks, 1998]. [Haddawy and Hanks, 1998] distinguish two kinds of nondeterminism: first, a world state (as well as an action effect) may not be known for sure, but an agent may know a probabilistic distribution of the values of the attributes which describe the world. E.g., while at work, from our cubicle, we do not know whether it is raining or not, but the weather forecast said that there is a .3 probability of a sunny day. Second, an agent can have no idea of the probability distribution on those values (this does not mean that he believes they are equally likely): the agent is *uncertain* about the actual state of the world. Coming back to the example, if we haven’t listened to the weather forecast we have no idea at all if it is raining or it is a sunny day.

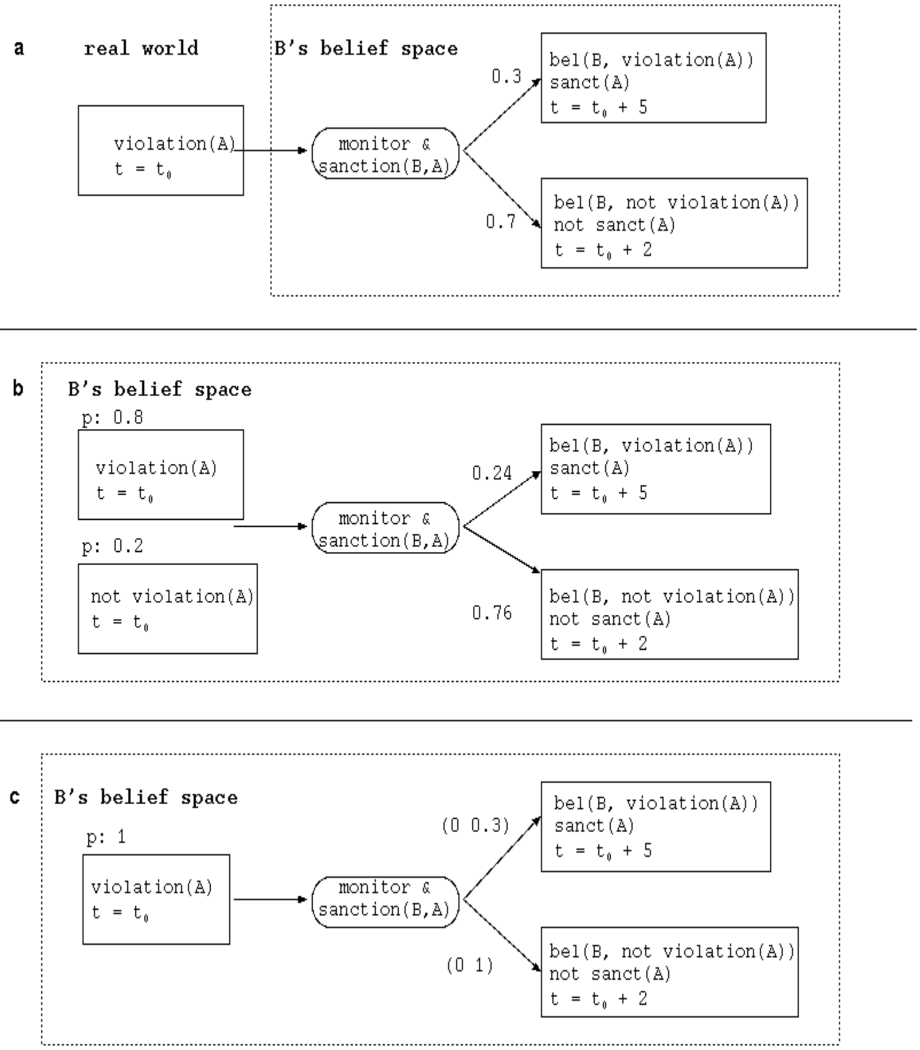

Figure 3. Monitoring and sanctioning an agent.

The representation of action abstraction and uncertainty in the DRIPS planner ([Ha and Haddawi, 1996]) is grounded on the same bases of the Theory of Evidence ([Shafer, 1976]). So we can exploit this theory for representing the situation of the *normative* agent. In of the Theory of Evidence, the certainty about the world is represented by the so-called Basic Probability Assignment (BPA). In a BPA, a probability is assigned to each subset of the universe of possible values of a random variable. As usual, the sum of the BPA values is 1. The assignment of the total mass of probability (1) to the whole universe is intended to mean total ignorance[9] . Although this is just a special case of BPA, it is useful to obtain a concise representation of the amount of knowledge an agent has about the current

[9] This must be contrasted to an assignment of values different from 0 just to singletons, which corresponds to a standard probability assignment, and which is considered as expressing some precise knowledge about the world, i.e. how probable the different outcomes are.

situation. In fact, we assume that just three situation may arise: total ignorance (BPA assigns 1 to the whole universe), knowledge about a probability distribution (BPA assigns values different from 0 just to singletons), certainty about the value of a given attribute (just one singleton has the BPA value 1). In this way, since what is of interest is just the utility of the outcomes, it is possible to focus the attention on the worst and best cases. So, the utility of the different states is evaluated, and the situation of ignorance is represented as a pair whose first element is the state for which the utility is worst and the second element is the state for which the utility is best. The utility of an uncertain state is represented as a pair as well, with the utility of the worst and best case scenario as first and second element. It must be observed that the utility pair can be interpreted as an interval, since the actual utility value necessarily falls within the bounds appearing in the pair.

But also the action of applying the sanction may fail. So, even if the violation has been detected, and *B* has decided to apply the sanction (which he would not do, in case the cost of applying it is greater than the gained utility), the sanctioning action may fail. *B* must (or, at least, we claim that rational agents do) weigh all of these possibilities when he chooses the best way of acting.

As an example of the full representation of an action, in Figure 2, we have depicted the sanctioning action to be performed by the *normative* agent; it is composed of a sequence of two elementary actions, the monitoring action and the (conditional) action of sanctioning the violator. Besides the actions, their (conditional) effects are depicted in a programming language style. The `bel(X, violation(Y))` notation is used to represent the fact that the attribute concerns the view that the agent has about the predicate `violation(Y)`; this must be contrasted with `violation(Y)`, which refers to the actual world (or, during the simulation, the beliefs of the *bearer* agent). The monitoring action says that the agent can *sense* the world and discover whether a violation has occurred with a .3 probability; in contrast, if no violation happened, he knows this fact for sure. The second action, `sanction` , has a deterministic effect conditioned to the fact that the agent knows that a violation happened. The effects besides the top level complex action `Monitor&Sanction` summarize the effects of its decomposition via a sequential abstraction (see [Haddawy and Suwandi, 1994] for details). In particular, these effects say that, if there is a violation, there is a .3 chance that the agent becomes aware of this fact and the violator is sanctioned. The 'monitor' action takes 2 time units, and the possible sanctioning takes 3 time units.

This rather complex representation must be combined with the fact that this action, during the planning phase, must be evaluated starting from a world which can contain probability or uncertainty from the *normative* agent's point of view: i.e., B may or may not know the probability distributions of the outcomes of A's actions.

As a starting point, we show in Figure (a), the notation we adopt to describe the (probabilistic) effect of the execution of the action `Monitor&Sanction` . The hypothesis is that a violation actually occurred: in such a case, during the execution phase, it is sufficient to apply the definition appearing in the upper right corner of Figure 2, using the `violation(Y)` branch.
In 3 (b), the action is evaluated during the planning phase in a situation where the *normative* agent has some idea about the probability of a violation. The .8 probability of a (suspected) violation makes the *normative* agent eager to find out and sanction it. A different distribution could lead to optimistic agents who do not suspect of violations and decide not to monitor for them. In Figure 3 (c), it is depicted the situation where the *normative* agent has no idea whether a violation occurred: no precise probability can be associated with the possible outcomes. In fact the arcs are labelled with pairs: the probability of finding a

violation ranges from 0 (in case no actual violation occurred) to .3 (in case there was a violation, this is the expected probability of finding it); the probability of not finding it, on the contrary, ranges from 0 to 1. Note that the first element of the pair (0 in both cases) is associated with the first element of the pair characterizing the ignorance about the actual state (i.e., no violation) and the same for the second element. So, the uncertainty about the initial situation is propagated to the predicted outcomes of the action.
Uncertainty has the effect that the expected utility of a set of outcomes must be represented as an interval having as lower and upper bounds the lowest and highest expected utility of the various uncertain outcomes. In order to compare utility interval it is possible to use any of the decision criteria developed in decision theory for dealing with uncertainty [von Neumann and Morgenstern, 1947].

5.2 Is it worth fulfilling an obligation ?

So, as we have seen, there are various motivations for an agent to decide not to fulfill an obligation Ω:

1. The agent has adopted the obligation but he cannot do anything for it (i.e., she has no feasible plan).

2. The possible plans which include some actions for fulfilling Ω achieve a lower utility than some other plan (due to the cost of fulfilling the obligation). In particular, this may happen if there is no certainty that the *normative* agent becomes aware of the fulfillment so that he will probably apply the sanction anyway.

3. There is some plan which does not fulfill the obligation but which induces the *normative* agent to believe otherwise.

4. There is some plan which does not fulfill the obligation but which makes the sanction impossible to apply.

5. The *bearer* of the obligation can bribe or menace the *normative* agent so that he does not apply the sanction.

If agents who respect obligations (when they can) are needed, a trade off between autonomy and control must be found; there are various ways to enforce the fulfillment of obligations, besides the *incentive* of sanctions:

1. The content of an obligation Ω can occur also as a preference of the bearer agent: in this way, when adopted, it becomes similar to an intention (but reinforced by the possible sanction): the agent directly achieves an utility from the satisfaction of the obligation (the content of the obligation is a *value* for the agent.)

2. The agent may have the preference not to mislead the *normative* agent: the former agent does not do anything to induce false beliefs in the *normative* agent, e.g., that the obligation is fulfilled when it is not the case. In this way, the agent does not exploit the possibilities described above.

3. The agent has some *social goal* which makes her not prefer situations where he is liable (for example, because she does not want that other agents decrease the trust they have on her).

4. the *normative* agent can overstate his willingness and capability of discovering violations, so that (if he is trusted by the bearers of the obligation) the result of anticipatory coordination favors (i.e., gives higher probability to) states resulting from the application of the sanction.

6. A FURTHER RELATION BETWEEN COOPERATION AND OBLIGATIONS

As we said in Section 4.2, the definition of cooperation does not imply the unrestricted benevolence of the agent. For example, if a partner has some other goals which conflicts with his task in the shared plan, the agent is not required to do anything for him, such as helping him in achieving the private goal; even if, in doing that, the performance of the entire group would be improved. Rather it is the partner who should be blamed for endangering the performance of the group.

From a strictly self-interested point of view, sticking to the definition is not always rational for the agent. If it were, indeed, there would be no need of a separate definition of group rationality. For example, if the agent knows that the goal has been achieved, she generally does not gain anything from communicating this fact to the partners. The agent is not motivated at all to follows the definition of cooperation. Why should she do so?

Rather, the definition provides a notion of accountability for the agent in the group's acting. And this accountability is the indirect motivation for the agent to stick to the definition. This fact shows that the notion of cooperation and obligation are not only based on the same social abilities of the agents, but that the definition of cooperation does make sense only as the content of an obligation the agent is subject to when she enters a group. The sanction of this implicit obligation involving the agents of a group is generally an informal one: the agent's social image is affected by her unreliability. As we discuss in [Boella et al., 2000] also these informal sanctions can be accounted for in our model. Second, the *normative* agents are all the members of the groups which can monitor and sanction all the other partners.

Finally, also in case the agent cooperates correctly since it is a value for her to do so (i.e., it is motivated to do so by some internal reward and not by external sanctions), the definition still plays a role in that it specifies which are the limits of her cooperation.

7. CONCLUSIONS

In this paper, we have shown that *social rationality* is at the basis of cooperation among autonomous agents and deliberate normative behavior even if they seem very different phenomena and they have been traditionally dealt with in very different manners.

In the case of cooperation the autonomy of the agent in performing her part is traded off against the necessity of controlling it under the light of anticipatory coordination. In the case of obligations, we have analysed how to ensure the control of the agent by means of internal and external (i.e., sanctions) incentives and we have shown that the autonomy of the *normative* agent deserves more attention.

The basic elements of social rationality, in our model, are *goal adoption* and *anticipatory coordination*, which are used in an agent framework based on *decision theoretic planning*.

Finally, the present framework has been described in more detail for what concerns cooperation in [Boella, 2000] and [Boella et al., 2000a] and for what concerns obligations in [Boella and Lesmo, 2001a]. In [Boella et al. 2001b], the present model of obligations has been exploited form modeling legal relations.

References

R. Axelrod. An evolutionary approach to norms. *The American Political Science Review*, 80(4):1095–1111, 1986.

G. Boella. *Cooperation among economically rational agents*. PhD thesis, Università di Torino, 2000.

G. Boella, R. Damiano, and L. Lesmo. Cooperation and group utility. In N. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99, Orlando FL)*, pages 319–333. Springer-Verlag, Berlin, 2000a.

G. Boella, R. Damiano, and L. Lesmo. Social goals in conversational cooperation. In *Proc. of Sigdial Workshop*, Honk Kong, 2000b.

G. Boella and L. Lesmo. Deliberate normative agents. In *Social order in MAS*. Kluwer, 2001a.

G. Boella, L. Favali and L. Lesmo. The definition of legal relations in a bdi multiagent framework. In *Proc. of AI*IA*. Springer Verlag, Berlin, 2001b.

M. E. Bratman. Intention, Plans, and Practical Reason, Harvard University Press, Cambridge (MA), 1987.

S. Carberry. *Plan Recognition in Natural Language Dialogue*. MIT Press, 1990.

C. Castelfranchi. Modeling social action for AI agents. *Artificial Intelligence*, 103:157–182, 1998.

P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.

P. Cohen and H. Levesque. Confirmation and Joint Action, In *Proc. 12th IJCAI*, Sydney, 1991.

F. Dignum. Autonomous agents and social norms. In *ICMAS'96 Workshop on Norms, Obligations and Conventions*, 1996.

A. Giddens. *Sociology*. Polity Press, Cambridge, 1989.

P. J. Gmytrasiewicz and E. H. Durfee. Formalization of recursive modeling. In *Proc. of first ICMAS-95*, 1995.

E. Goffman. *Strategic Interaction*. Basil Blackwell, Oxford, 1970.

B. Grosz and S. Kraus. Collaborative plans for complex group action. *Artificial Intelligence*, 86(2):269–357, 1996.

V. Ha and P. Haddawy. Theoretical foundations for abstraction-based probabilistic planning. In *Proc. of 12th Conference on Uncertainty and Artificial Intelligence*. Portland, 1996.

P. Haddawy and S. Hanks. Utility models for goal-directed, decision-theoretic planners. *Computational Intelligence*, 14:392–429, 1998.

P. Haddawy and M. Suwandi. Decision-theoretic refinement planning using inheritance abstraction. In *Proc. of 2nd Int. Conference on Artificial Intelligence Planning Systems*, pages 266–271, Menlo Park, CA, 1994.

James Helwig and Peter Haddaw. An Abstraction-Based Approach to Interleaving Planning and Execution in Partially-Observable Domains, *Plan Execution: Problems and Issues: Papers from the 1996 AAAI Fall Symposium*, AAAI Press, Menlo Park, California 1996.

I.Hideki and K.Hirofumi. Observability-based nested belief computation for multiagent systems and its formalization. In N. Jennings and Y. Lespérance, editors, *Intelligent Agents VI — Proceedings of the Sixth International Workshop on Agent Theories, Architectures, and Languages (ATAL-99)*, Lecture Notes in Artificial Intelligence. Springer-Verlag, Berlin, 2000.

R. L. Keeney and H. Raiffa. *Decision with multiple Objectives: preferences and value tradeoff*. Cambrige University Press, Cambridge (UK), 1976.

A. Ndiaye and A. Jameson. Predictive role taking in dialog: global anticipation feedback based on transmutability. In *Proc. 5th Int. Conf. on User Modeling*, pages 137–144, Kailua-Kona, Hawaii, 1996.

G. Shafer, A Mathematical Theory of Evidence. Princeton University Press, 1976.

M.C. Schut and M. Wooldridge. Intention reconsideration in complex environments. In *Proc. of Agents 2000*, Barcelona, 2000.

J. Von Neumann and O. Morgenstern. *The Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1947.

S. Zilberstein and S.. Russell. Anytime sensing, planning and action: A practical model for robot control. In *Proc.of the Int. Joint Conf. on Artificial Intelligence*, pages 1402--1407, Chambery, France, 1993.