A Rigorous, Operational Formalization of Recursive Modeling

Piotr J. Gmytrasiewicz* and Edmund H. Durfee†

*Department of Computer Science and Engineering
University of Texas at Arlington, Arlington, TX 76019-0015

†Department of Electrical Engineering and Computer Science
University of Michigan
Ann Arbor, Michigan 48109
piotr@cse.uta.edu, durfee@caen.engin.umich.edu

Abstract

We present a formalization of the Recursive Modeling Method, which we have previously, somewhat informally, proposed as a method that autonomous artificial agents can use for intelligent coordination and communication with other agents. Our formalism is closely related to models proposed in the area of game theory, but contains new elements that lead to a different solution concept. The advantage of our solution method is that always yields the optimal solution, which is the rational action of the agent in a multi-agent environment, given the agent's state of knowledge and its preferences, and that it works in realistic cases when agents have only a finite amount of information about the agents they interact with.

Introduction

Since its initial conceptual development several years ago (Gmytrasiewicz, Durfee, & Wehe 1991a; 1991b), the Recursive Modeling Method (RMM) has provided a powerful decision-theoretic underpinning for coordination and communication decisionmaking, including decisions about synchronized plans (Gmytrasiewicz & Durfee 1992), about knowledge-oriented actions (Gmytrasiewicz & Rosenschein 1993), about honesty and trust among self-interested agents (Gmytrasiewicz & Durfee 1993 to appear), and about the principled adoption of protocols (Durfee, Gmytrasiewicz, & Rosenschein 1994). Initially developed for coordinating robots in a nuclear reactor, RMM has more recently been identified as a prime candidate for realizing agile manufacturing systems (Gmytrasiewicz, Huang, & Lewis 1995), because it can unify operations research techniques (used for decades in factory automation and scheduling) and the agent-oriented paradigm emerging in artificial intelligence. However, RMM has had to become more theoretically mature as a precursor to such applications; its earlier formulation glossed over crucial details and introduced *ad hoc* mechanisms, particularly to break out of cyclic reasoning at increasing levels of recursion.

This paper provides a rigorous formulation of RMM, moving beyond our earlier description by providing a clearer semantics and stronger solution method. Our new formulation exposes the details necessary for operationalizing RMM. It also strengthens RMM's theoretical foundation, and in particular allows us to highlight the differences between RMM, which has been developed from a distinctly AI perspective to address the question of what an agent should do in an arbitrary multi agent situation, and game-theoretic methods, which have focused more squarely on characterizing multi agent equilibria. While our representation resembles some of those used in game theory (Harsanyi 1967; Mertens & Zamir 1985; Aumann & Brandenburger 1994; Brandenburger 1992; Brandenburger & Dekel 1993), our emphasis on the practical limitations on knowledge and on decision-theoretic principles applied to arbitrary situations distinguishes our work from game theory.

General Form of the Recursive Modeling Method

RMM views a multi agent situation from the perspective of an agent that is individually trying to decide what physical and/or communicative actions it should take right now. To make decisions, RMM uses a decision-theoretic paradigm of rationality, where an agent attempts to maximize its expected utility given its beliefs. The central elements of RMM are, thus, a formalized representation of the agent's state of beliefs, together with a procedure that operates on this representation and solves it, yielding a rational choice of action in a multi agent environment. The basic building block of RMM's representation is a payoff matrix, which succinctly and completely expresses the agent's beliefs about its environment, its capabilities, and its preferences. The agent's payoff matrix, therefore, precisely depicts the decision-making situation that the agent finds itself in. In order to solve its own decisionmaking situation, the agent needs an idea of what the other agents are likely to do. It can arrive at it by representing what it knows about the other agents' decision-making situations, thus modeling them in terms of their own payoff matrices. The fact that the other agents could also be modeling others, including the original agent, leads to a recursive nesting of the models. The whole recursive model structure represents, therefore, all of the information the agent has that is relevant to its current decision-making situation: its own view of the world, abilities and preferences, what it knows about the views, abilities and desires of the other agents, what it knows about how they model other agents, and so on. Based on its practical limitations in acquiring knowledge, an agent can only build a finite nesting of models, and RMM exploits this limitation in its solution method, which applies decision-theoretic principles to the model beginning at its leaves. We now look at the representation and solution method more formally.

We first formally define the payoff matrix, which is the basic building block of RMM's modeling structure. Following the definition used in game theory (Rasmusen 1989), we define the payoff matrix P_{R_i} of an agent R_i as a triplet (R, A, U). R is a set of agents in the environment, labeled R_1 through R_n $(n \geq 1)$. R will be taken to include all possible decision-making agents impacting the welfare of the agent R_i. A is a set of sets A_j, where $A_j = \{a_1^j, a_2^j \cdots\}$ represents the alternative actions of agent R_j. We will call A_j the decision space of the agent R_j. The alternative actions of an agent will be taken to include all of the environmental activities (including no activity) that the agent is equipped for and could possibly undertake, and which influence the payoffs of agent R_i.[1]

Finally, U is a payoff function that assigns a number (payoff) to each of the possible combinations of the actions of all of the agents: $U : A_1 \times A_2 \times ... \times A_n \longrightarrow \mathbf{R}$, where $\mathbf{R}$ is the set of real numbers. Intuitively, any purposeful agent has reason to prefer some actions (that further its purposes in the current situation) to others (Wellman 1991). Our ability to represent agents' preferences over actions as payoffs follows directly from the axioms of utility theory, which postulate that ordinal preferences among actions in the current situation can be represented as cardinal, numeric values (Chernoff & Moses 1959; Doyle 1992). We represent R_i's payoff associated with R_i taking action a_m^i while the other agents take their actions (R_1 takes action a_k^1, ... R_n takes action a_l^n) as $u^{R_i}_{a_k^1 \cdots a_m^i \cdots a_l^n}$.

In summary, then, the payoff matrix defined by (R, A, U) has a dimensionality equal to the number of agents n, has indices along each dimension defined by the corresponding set in A, and has entries in the matrix defined by the payoff function U. A payoff matrix so defined succinctly represents all relevant parameters of the decision-making situation an agent finds itself in. It includes the agent's abilities (alternative actions A_j), its beliefs about the other agents are present and about the way the collective actions of the agents are expected to influence the agent's preferences. However, it does not include the agent's beliefs about the decision-making situations the other agents are in – these are contained in the recursive model. We now define the recursive model structure RMS_{R_i} of the agent R_i, as the following pair:

$$RMS_{R_i} = (P_{R_i}, RM_{R_i}), \tag{1}$$

where P_{R_i} is R_i's payoff matrix as just defined, and RM_{R_i} is the recursive model that R_i uses to model the other $n-1$ agents in the environment.[2] A recursive model RM_{R_i} is defined as follows:

$$RM_{R_i} = ((p_1^{R_i}, M^{(R_i,1)}_{\{-R_i\}})..(p_\alpha^{R_i}, M^{(R_i,\alpha)}_{\{-R_i\}})..(p_m^{R_i}, M^{(R_i,m)}_{\{-R_i\}}), \tag{2}$$

where $M^{(R_i,\alpha)}_{\{-R_i\}}$ is one of R_i's m alternative models of the other agents (the set of other agents—all agents except R_i—is represented by the subscript $\{-R_i\}$). Agent R_i, thus, represents its subjective belief that each of these alternatives is correct by assigning probabilities $p_\alpha^{R_i}$ to them. These probabilities, which we will call *modeling* probabilities, have to sum up to unity: $\sum_{\alpha=1}^m p_\alpha^{R_i} = 1$. Each of the alternative models of the other agents consists simply of the list of models of each of them:

$$M^{(R_i,\alpha)}_{\{-R_i\}} = (M^{(R_i,\alpha)}_{R_1}..M^{(R_i,\alpha)}_{R_{i-1}}, M^{(R_i,\alpha)}_{R_{i+1}}..M^{(R_i,\alpha)}_{R_n}). \tag{3}$$

The models $M^{(R_i,\alpha)}_{R_j}$ come in three possible forms:

$$M^{(R_i,\alpha)}_{R_j} = \begin{cases} IM^{(R_i,\alpha)}_{R_j} & \text{– the Intentional model,} \\ NM^{(R_i,\alpha)}_{R_j} & \text{– the No-Information model,} \\ SM^{(R_i,\alpha)}_{R_j} & \text{– the Sub-Intentional model.} \end{cases} \tag{4}$$

The Intentional model corresponds to R_i modeling R_j as a rational agent, as advocated, for instance, by Daniel Dennett in (Dennett 1986). It is defined as:

[1] We would like to remark that these alternatives can be on any level of abstraction and can represent single primitive actions, detailed plans, or high–level strategies with contingencies.

[2] In game theory this construct is also called the agent's *theory*.

$$IM^{(R_i,\alpha)}_{R_j} = RMS^{(R_i,\alpha)}_{R_j}, \tag{5}$$

that is, it is the recursive model structure that agent R_i ascribes to agent R_j. This structure, as defined in Equation 1, further consists of the payoff matrix that R_i ascribes to R_j in this model, $P^{(R_i,\alpha)}_{R_j}$, and the recursive model $RM^{(R_i,\alpha)}_{R_j}$ that R_i thinks R_j uses to solve its decision-making situation.

The No-Information model, $NM^{(R_i,\alpha)}_{R_j}$, corresponds to R_i's having no information based on which it could predict R_j's behavior. Therefore, the No-Information model is defined as assigning an equiprobable distribution to R_j's choosing to act on any of its alternative behaviors, $a^j_k \in A_j$, as enumerated in R_i's payoff matrix: $p^{(R_i,\alpha)}_{R_j} = [p^{(R_i,\alpha)}_{a^j_1}, p^{(R_i,\alpha)}_{a^j_2}, ...] = [\frac{1}{|A_j|}, \cdots, \frac{1}{|A_j|}]$, where $|A_j|$ is the number of alternative actions that are ascribed to R_j at the given level of modeling. The uniform probability distribution of the No-Information model contains no information (Neapolitan 1990) and thus precisely represents R_i's lack of knowledge in this case.

The Sub-Intentional model corresponds to treating the agent as not being rational. In realistic situations it may, for example, be unclear whether the agent is, in fact, a rational agent, or whether it is, say, a piece of furniture. As Dennett proposes (Dennett 1986), the intentional stance is not the only one possible and useful. There is the *design* stance, which predicts behavior using functionality (such as how the functions of a console controller board's components lead to its overall behavior (Hamscher 1986)), and the *physical stance*, which predicts behavior using the description of the physical state of what is being modeled along with knowledge of the laws of nature (like in the qualitative model of a bouncing ball (Forbus 1980)). For the purpose of this paper, we will assume that an agent can incorporate well-studied techniques such as model-based reasoning and qualitative physics to make predictions about the behavior of sub-intentional entities, resulting in a probability distribution over their alternative behaviors, as enumerated in the agent's payoff matrix. Thus, if there is a possibility that an agent may, in fact, be a piece of furniture, the Sub-Intentional model would likely suggest that the object will simply stay put.

The definition of the recursive model structure and the Intentional model are recursive, but we would like to argue that the recursion is bound to end due to practical limitations in attaining infinite knowledge. In other words, in representing its knowledge of its own decision-making situation, and the situations of the other agents, and of what the other agents know about others, and so on, an agent is bound to run out of information at some level of nesting, in which case the recursion would terminate with a No-Information model. Of course, some recursive branches can also terminate with Sub-Intentional models that do not lead to further recursion.

We now go on to describe the RMM's solution method. The method traverses the recursive model structure, propagating the information bottom up, and arriving at an assignment of expected utilities to the agent's alternative actions, based on all of the relevant information the agent has. Our rational agent can then choose an action with the highest expected utility.

The utility of the m-th element, a^i_m, of R_i's set of alternative actions can be evaluated as:

$$u^{R_i}_{a^i_m} = \sum_{a^1_k \in A_1} .. \sum_{a^{i-1}_l \in A_{i-1}} \sum_{a^{i+1}_o \in A_{i+1}} .. \sum_{a^n_p \in A_n} \{p^{R_i}_{a^1_k} .. p^{R_i}_{a^n_p} u^{R_i}_{a^1_k \cdots a^i_m \cdots a^n_p}\} \tag{6}$$

where $p^{R_i}_{a^j_k}$ represents the probability R_i assigns to agent R_j's intending to act on the k-th element of R_j's set of alternative actions A_j, which we will refer to as an *intentional* probability. $u^{R_i}_{a^1_k \cdots a^i_m \cdots a^n_p}$ is R_i's payoff as an element of its payoff matrix, P_{R_i}. Equation 6 expresses the expected utility for an action as a sum of the utilities of each of the joint actions that include that action, weighted by the probability of each joint action. We assumed that the agents take their actions independently, as should be the case for autonomous agents.

R_i should attempt to determine the intentional probabilities $p^{R_i}_{a^j_k}$, i.e., arrive at some idea of what the others are likely to do, by using its modeling knowledge contained in the recursive model RM_{R_i}. As Equations 2 and 3 show, R_i can have a number of alternative models $M^{(R_i,\alpha)}_{R_j}$ of R_j, each of them estimated to be correct with the modeling probability $p^{R_i}_\alpha$. If we label the predicted probability of behavior a^j_k of R_j resulting from a model $M^{(R_i,\alpha)}_{R_j}$ as $p^{(R_i,\alpha)}_{a^j_k}$, we can express the overall intentional probability $p^{R_i}_{a^j_k}$ as the following probabilistic mixture:

$$p^{R_i}_{a^j_k} = \sum_\alpha p^{R_i}_\alpha \times p^{(R_i,\alpha)}_{a^j_k}. \tag{7}$$

If the model $M^{(R_i,\alpha)}_{R_j}$ is in the form of the Sub-Intentional model, then the probabilities $p^{(R_i,\alpha)}_{a^j_k}$ can be derived by whatever techniques (model-based, qualitative physics, etc.) R_i has for reasoning about such agents. If R_i has no modeling information at this point at all, the No-Information model is used, yielding the probabilities $p^{(R_i,\alpha)}_{a^j_k} = \frac{1}{|A_j|}$ specified in this model.

If, on the other hand, R_i has assumed an intentional stance toward R_j in its model $M_{R_j}^{(R_i,\alpha)}$, i.e., if it is modeling R_j as a rational agent, then it has to model the decision-making situation that agent R_j faces, as specified in Equations 4, 5, and 1, by R_j's payoff matrix $P_{R_j}^{(R_i,\alpha)}$ and its recursive model $RM_{R_j}^{(R_i,\alpha)}$. R_i can then identify the intentional probability $p_{a_k^j}^{(R_i,\alpha)}$ as the probability that the k-th alternative action is of the greatest utility to R_j in this model:

$$p_{a_k^j}^{(R_i,\alpha)} = Prob(u_{a_k^j}^{(R_i,\alpha),R_j} = Max_{k'}(u_{a_{k'}^j}^{(R_i,\alpha),R_j})). \quad (8)$$

$u_{a_{k'}^j}^{(R_i,\alpha),R_j}$ is the utility R_i estimates that R_j assigns to its alternative action $a_{k'}^j$ in this model, and it can be further analogously to Equation 6, but based on based on R_i's model of R_j. This calculation would contain the probabilities, which represent what R_i thinks R_j assigns to agent R_k, $p_{a_p^k}^{(R_i,\alpha),R_j}$, and they can in turn be expressed in terms of the models that R_i thinks R_j has of the other agents in the environment, contained in $RM_{R_j}^{(R_i,\alpha)}$, and so on.

The intentional stance R_i uses to model R_j is formalized in Equation 8. It states that agent R_j is an expected utility maximizer and, therefore, its intention can be identified as the course of action that has the highest expected utility, given R_j's beliefs about the world and its preferences. What the intentional stance does not specify, however, is how R_j will make its choice if it finds that there are several alternatives that provide it with the maximum payoff. Consequently, using the principle of indifference once more, R_i will assign an equal, nonzero probability to R_j's option(s) with the highest expected payoff, and zero to all of the rest. Formally, we can construct the set of R_j's options that maximize its utility:

$$Amax_j^{(R_i,\alpha)} = \{a_k^j \in A_j^{(R_i,\alpha)} \wedge u_{a_k^j}^{(R_i,\alpha),R_j} = Max_{k'}(u_{a_{k'}^j}^{(R_i,\alpha),R_j})\} \quad (9)$$

Then, the probabilities are assigned according to the following:

$$p_{a_k^j}^{(R_i,\alpha)} = \begin{cases} \frac{1}{|Amax_j^{(R_i,\alpha)}|} & \text{if } a_k^j \in Amax_j^{(R_i,\alpha)} \\ 0 & \text{otherwise.} \end{cases} \quad (10)$$

Simple Example

We will consider briefly a particular example of a multi-agent reconnaissance in a hazardous environment, such as a nuclear or a chemical plant in an emergency situation, or a toxic waste facility. The example is simple but we would like to argue that the method we are proposing is general and could be used in *any* multi-agent situation, given *any* finite state of agent's knowledge about the environment and the other agents. Let us imagine an autonomous outdoor robotic vehicle, called R_1 (see Figure 1), that is attempting to coordinate its actions with another robotic vehicle, R_2. The mission is to gather information while minimizing cost (fuel and/or time consumed), thus nearby locations with high elevation are good candidates for observation points. For R_1, two points, P1 and P2, are worth considering. P2 has a higher elevation and would allow twice as much information to be gathered as P1, and R_1 estimates that the information gathered from P1 and P2 have the expected values of 2 and 4, respectively. Assume that R_1 has three alternative courses of action: observing form P1, observing form P2, or doing neither and sitting still, labeled a_1^1, a_2^1, and a_3^1. Say that the expected cost (time or energy) to R_1 of pursuing each of these is proportional to the distance traveled, yielding a cost of 1 for a_1^1, 2 for a_2^1, and 0 for a_3^1. We further assume in this example that each of the robots can make only one observation, and that each of them benefits from *all* information gathered (no matter by which robot), but incurs cost only based on its own actions.[3] The relevant alternatives of R_2 are a_1^2 through a_3^2, and correspond to R_2's taking the observation from point P1, P2, and staying put or doing something else, respectively. In the framework we propose, the above information is represented as the payoff matrix, which is depicted at the top of Figure 2.

In order to arrive at the rational decision as to which of its options to pursue, R_1 has to predict what R_2 will do. That depends on the decision-making situation R_2 is facing. In the scenario considered there are trees between R_2 and P2, so R_2 may be unaware of P2. R_1, therefore, has to deal with two alternative models of R_2's decision-making situation. If R_2 is unaware of P2, then it will not consider combinations of actions involving P2, i.e., a_2^1 or a_2^2, and its payoff matrix is 2 × 2. But, if it can see through the trees, the matrix is 3 × 3, as depicted on the second level in Figure 2. There, R_1, having some knowledge about the sensors available to R_2 and assessing the density of the foliage between R_2 and P2, assigns a probability for R_2 seeing through the trees as 0.1. For simplicity we consider here the case when R_1 has no knowledge about how it might be modeled by R_2, and a No-information model assigning a uniform probability distribution over R_1's set of actions on this level is used. It indicates that R_1

[3] These assumptions are only for the purpose of keeping our example simple. In no way do they limit the applicability of the method we propose.

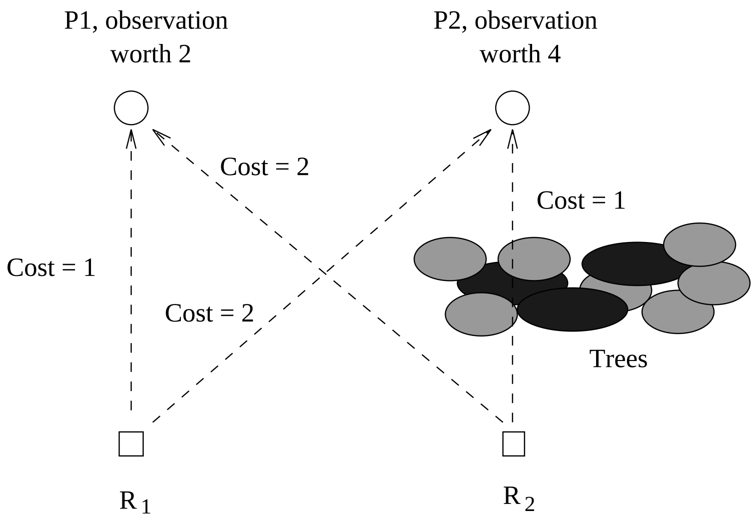

Figure 1: Example Scenario of Interacting Agents

could undertake any one of alternative actions a_1^1, a_2^1, and a_3^1, with probabilities $[\frac{1}{3}, \frac{1}{3}, \frac{1}{3}]$. Thus, R_1's knowledge about the decision-making situation that it faces in this simple example can be fully represented as depicted in Figure 2.

As we mentioned, our approach to the solution is to arrive at the rational choice of R_1's action bottom-up, from the leaves of the recursive model structure toward the top matrix of the agent. We show this for our simple example here without referring to the theoretical components in Section . On the second level of the model structure, in the model of R_2 corresponding to it not seeing through the trees, the expected utilities of R_2 actions are 0 and 1 $(= 2\times 0.5+0\times 0.5)$, for a_1^2 and a_3^2, respectively. This indicates that R_2 would stay put in this case. If R_2 can see through the trees, on the other hand, the calculation shows that R_2's preferred action would be a_2^2, i.e., to observe from P2. The conclusion from both of the models on the second level of modeling is, therefore, that R_2 would pursue a_3^2 if it cannot see through the trees, and a_2^2, if it can, which have the estimated probabilities of 0.9 and 0.1, respectively. This results in a probability distribution over R_2's actions a_1^2, a_2^2, and a_3^2 being $[0, 0.1, 0.9]$. The above result can be propagated up and used to solve the matrix on the top of Figure 2. The expected utilities of R_1's actions can be computed straightforwardly as; $u^{R_1}_{a_1^1} = 0\times 1+.1\times 5+.9\times 1 = 1.4$, $u^{R_1}_{a_2^1} = 0\times 4+.1\times 2+.9\times 2 = 2$, $u^{R_1}_{a_3^1} = 0 \times 2 + .1 \times 4 + .9 \times 0 = 0.4$. Thus, the best choice for R_1 is to pursue its option a_2^1, that is, to move toward point P2 and make an observation from there. It is the rational coordinated action given R_1's state of knowledge, since the computation included all of the information R_1 has about agent R_2's expected behavior. Intuitively, this means that R_1 believes that R_2 is so unlikely to go to P2 that R_1 should go there itself.

Relation to Other Work

Our approach in constructing the nested models in RMM is closely related to work in game theory devoted to games of incomplete information by Harsanyi[4], Aumann, and others (Aumann & Brandenburger 1994; Brandenburger & Dekel 1993; Harsanyi 1967; Mertens & Zamir 1985). In this work, the state of an agent's knowledge, or its epistemic state, is represented as an agent's type, consisting, roughly speaking, of a payoff matrix and the probability distribution over the possible types of the other agents. This structure, called an interactive belief system in (Aumann & Brandenburger 1994), represents the state of an agent's knowledge, very much like RMM's recursive model structure does. The differences between RMM and the above approaches is the inclusion of the No-Information and Sub-Intentional models in RMM, and in our postulate that the recursive model structure must terminate, while the interactive belief system studied in (Aumann & Brandenburger 1994) and similar structures in (Harsanyi 1967; Mertens & Zamir 1985) are infinite. This seemingly superficial difference, however, leads to significant differences in the solution concepts employed.

The solution concept traditionally employed in game theory is that of an equilibrium.[5] However, the appro-

[4]Nobel Prize in Economics, 1994.

[5]An equilibrium, for example in a two-person game, is defined as a pair of strategies (s, t) such that s is an optimal

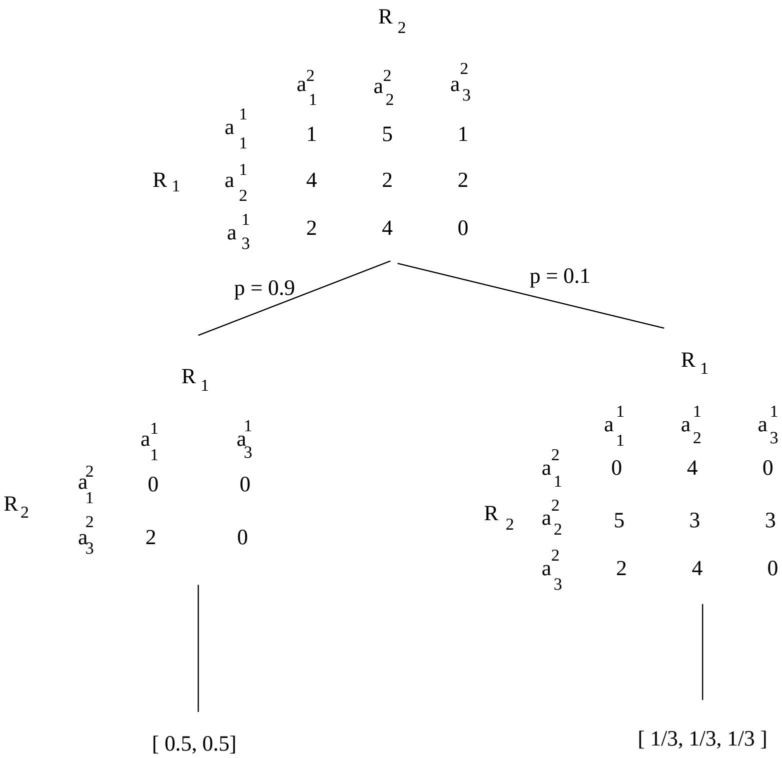

Figure 2: Recursive Model Structure depicting R_1's Knowledge in the Example Scenario.

priateness of this concept as a directive and a predictor for actions of rational agents is in doubt, mainly due to the fact that the players have to be assumed to have common knowledge of certain facts[6] to make it viable. This matter is a subject of considerable debate in the field of game theory itself (Aumann & Brandenburger 1994; Bicchieri 1993; Binmore 1982; Brandenburger 1992; Geanakoplos 1992; Reny 1988; Tan & Werlang 1988). Binmore (Binmore 1982) and Brandenburger (Brandenburger 1992) both point out that the weakness of the assumption of common knowledge points directly to the need for an explicit model of the decision-making of the agents involved, given their states of knowledge. This is exactly our approach in RMM. This modeling is not needed if one wants to talk only of the possible equilibria. Further, Binmore points out that the common treatment in game theory of equilibria without any reference to the equilibrating process that achieved the equilibrium[7] accounts for the inability of predicting which particular equilibrium is the right one and will actually be realized, if there happens to be more than one candidate.

Our method, remaining true to artificial intelligence's knowledge-level perspective (Newell 1981), postulates that a rational agent should apply all of its relevant knowledge to further its current goals. In recent game-theoretic literature (Aumann & Brandenburger 1994; Brandenburger 1992), this approach is also advocated, and called the decision-theoretic approach to game theory. Aumann and Brandenburger introduce the concept of an interactive belief system, similar to our recursive model structure in RMM, but without its additional elements of No-Information and Sub-Intentional models. As we mentioned, the No-Information model represents the situation when an agent has run out of the modeling knowledge on a certain level of nesting, and is a necessary property for agents only endowed with finite amounts of information about the world.[8] How is the case of finite information handled in the interactive belief systems employed in (Aumann & Brandenburger 1994; Brandenburger 1992)? Intuitively, these formulations have to "blow up" on the level below which no information is available, and therefore all possible types are present with equal probabilities, and all of the pos-

response to t and t is an optimal response to s.

[6]In general, the structure of the game and the rationality of the players are assumed to be common knowledge.

[7]Binmore compares it to trying to decide which of the roots of the quadratic equation is the "right" solution without reference to the context in which the quadratic equation has arisen.

[8]The finiteness of information practical agents are endowed with is unavoidable, since there are no practical means for agents to acquire infinite information about the outside world.

sibilities again on the next deeper level, and so on, ad infinitum. Solutions of such structures have not been proposed in (Aumann & Brandenburger 1994; Brandenburger 1992), but we would like to observe that, since they contain no information, their solutions must be completely uninformative. In this sense, our No-Information models faithfully summarize such structures, and provide for the ease of finding solutions, as they did in the simple example considered in this paper.

Apart from the related work in game theory on which we have concentrated above, a large amount of work in AI should be mentioned. Due to space limitations, we selectively mention the most germane publications. Horvitz's and Russell's work (Horvitz, Cooper, & Heckerman 1989; Russell & Wefald 1991) is concentrated on reasoning rationally under time pressure. Closely related is the concept of practical rationality in (Pollock 1989). Our approach also follows Dennett's formulation of the intentional stance (Dennett 1986), and his idea of the ladder of agenthood (see (Narayanan 1988) for a succinct discussion), the first five levels of which we see as actually embodied in RMM. Cohen and Levesque (Cohen & Levesque 1990a) formalize intention and commitment, and apply it to issues of communication (Cohen & Levesque 1990b). These authors, as well as Perrault in (Perrault 1990), analyze the nestedness of beliefs so important in issues of communication, but rely on a notion of common belief, the justifiability of which (as with common knowledge) we find problematic (see also (Bicchieri 1993) for related discussion). Shoham's agent-oriented programming (AOP) (Shoham 1993) takes more of a programming-language perspective. However, while Shoham has proposed it as an extension, decision-theoretic rationality has not yet been included in AOP. The models investigated by Fagin and colleagues in (Fagin, Halpern, & Vardi 1991) are related to ours, and include a no-information extension (like the No-Information model in RMM, but containing infinite recursion) to handle the situation where an agent runs out of knowledge at a finite level of nesting. However, they do not elaborate on any decision mechanism that could use their representation (presumably relying on centrally designed protocols), while all of our representation is specifically geared toward decision-making in the absence of a pre-determined protocol. Another related work on nested belief, with an extensive formalism, is one by Ballim and Wilkes (Ballim & Wilks 1991) and by Korf (Korf 1989). The applications of game-theoretic techniques to the problem of decision-making in multi-agent domains have also received attention in the Distributed AI literature, for example in (Rosenschein & Breese 1989; Rosenschein & Genesereth 1985; Zlotkin & Rosenschein 1989; 1990a; 1990b). An approach to communication similar to ours, although more qualitative, is presented in (Mayerson 1988). In a very similar spirit, Parikh gives an interesting utility-based approach to disambiguating certain kinds of implicatures (Parikh 1992).

References

Aumann, R. J., and Brandenburger, A. 1994. Epistemic conditions for Nash equilibrium. Working paper, Harvard Business School, Harverd.

Ballim, A., and Wilks, Y. 1991. *Artificial Believers*. Earlbaum Associates, Inc.

Bicchieri, C. 1993. *Rationality and Coordination*. Cambridge University Press.

Binmore, K. 1982. *Essays on Foundations of Game Theory*. Pitman.

Brandenburger, A., and Dekel, E. 1993. Hierarchies of beliefs and common knowledge. *Journal of Economic Theory* 59:189–198.

Brandenburger, A. 1992. Knowledge and equilibrium in games. *Journal of Economic Perspectives* 6:83–101.

Chernoff, H., and Moses, L. E. 1959. *Elementary Decision Theory*. New York: John Wiley.

Cohen, P. R., and Levesque, H. J. 1990a. Persistence , intention and commitment. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication*. MIT Press.

Cohen, P. R., and Levesque, H. J. 1990b. Rational interaction as the basis for communication. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication*. MIT Press.

Dennett, D. 1986. Intentional systems. In Dennett, D., ed., *Brainstorms*. MIT Press.

Doyle, J. 1992. Rationality and its role in reasoning. *Computational Intelligence* 8:376–409.

Durfee, E. H.; Gmytrasiewicz, P.; and Rosenschein, J. 1994. The utility of embedded communications and the emergence of protocols. In *Proceedings of 13th International Distributed Artificial Intelligence Workshop*.

Fagin, R. R.; Halpern, J. Y.; and Vardi, M. Y. 1991. A model-theoretic analysis of knowledge. *Journal of the ACM* (2):382–428.

Forbus, K. 1980. Spatial and qualitative aspects of reasoning about motion. In *Proceedings of the First Annual National Conference on Artificial Intelligence*.

Geanakoplos, J. 1992. Common knowledge. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowladge*, 255–315. Morgan Kaufman.

Gmytrasiewicz, P. J., and Durfee, E. H. 1992. Decision-theoretic recursive modeling and the coordinated attack problem. In *Proceedings of the First International Conference on Artificial Intelligence Planning Systems — AIPS92*, 88–95.

Gmytrasiewicz, P. J., and Durfee, E. H. 1993, to appear. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation.*

Gmytrasiewicz, P. J., and Rosenschein, J. S. 1993. The utility of embedded knowledge-oriented actions. In *Proceedings of the Twelveth International Workshop on Distributed Artifcial Intelligence.*

Gmytrasiewicz, P. J.; Durfee, E. H.; and Wehe, D. K. 1991a. The utility of communication in coordinating intelligent agents. In *Proceedings of the National Conference on Artificial Intelligence*, 166–172.

Gmytrasiewicz, P. J.; Durfee, E. H.; and Wehe, D. K. 1991b. A decision-theoretic approach to coordinating multiagent interactions. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, 62–68.

Gmytrasiewicz, P. J.; Huang, H. H.; and Lewis, F. L. 1995. Combining operations research and agent-oriented techniques for agile manufacturing system design. In *To Appear in the Proceedings of the IASTED International Conference on Robotics and Manufacturing.*

Hamscher, W. C. 1986. Modeling digital circuits for troubleshooting. *Artificial Intelligence* 51:1991.

Harsanyi, J. C. 1967. Games with incomplete information played by 'bayesian' players. *Management Science* 14(3):159–182.

Horvitz, E. J.; Cooper, G. F.; and Heckerman, D. E. 1989. Reflection and action under scarce resources: Theoretical principles and theoretical study. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 1121–1127.

Korf, R. E. 1989. Multi-agent decision trees. In *Proceedings of the 1989 Distributed AI Workshop*, 293–306.

Mayerson, R. B. 1988. Incentive constraints and optimal communication systems. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, 179–193.

Mertens, and Zamir. 1985. Formulation of bayesian analysis for games with incomplete information. *International Journal of Game Theory* 14:1–29.

Narayanan, A. 1988. *On Being a Machine.* Ellis Horwood.

Neapolitan, R. E. 1990. *Probabilistic Reasoning in Expert Systems.* John Wiley and Sons.

Newell, A. 1981. The knowledge level. *AI Magazine* 2(2):1–20.

Parikh, P. 1992. A game-theoretic account of implicature. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowladge*, 85–93. Morgan Kaufman.

Perrault, C. R. 1990. An application of default logic to speech act theory. In Cohen, P. R.; Morgan, J.; and Pollack, M. E., eds., *Intentions in Communication.* MIT Press.

Pollock, J. L. 1989. *How to Build a Person: A Prolegomenon.* MIT Press.

Rasmusen, E. 1989. *Games and Information.* B. Blackwel.

Reny, P. J. 1988. Extensive games and common knowledge. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowladge*, 395. Morgan Kaufman.

Rosenschein, J. S., and Breese, J. S. 1989. Communication-free interactions among rational agents: A probablistic approach. In Gasser, L., and Huhns, M. N., eds., *Distributed Artificial Intelligence*, volume 2 of *Research Notes in Artificial Intelligence.* Pitman. 99–118.

Rosenschein, J. S., and Genesereth, M. R. 1985. Deals among rational agents. In *Proceedings of the Ninth International Joint Conference on Artificial Intelligence*, 91–99. (Also published in *Readings in Distributed Artificial Intelligence*, Alan H. Bond and Les Gasser, editors, pages 227–234, Morgan Kaufmann, 1988.).

Russell, S., and Wefald, E. 1991. Principles of metareasoning. *Artificial Intelligence* 49:361–395.

Shoham, Y. 1993. Agent-oriented programming. *Artificial Intelligence* 60(1):51–92.

Tan, T. C., and Werlang, S. R. 1988. A guide to knowledge and games. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge.*

Wellman, M. P. 1991. The preferential semantics for goals. In *AAAI91*, 698–703.

Zlotkin, G., and Rosenschein, J. S. 1989. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, 912–917.

Zlotkin, G., and Rosenschein, J. S. 1990a. Blocks, lies, and postal freight: Nature of deception in negotiation. In *Proceedings of the 1990 Distributed AI Workshop.*

Zlotkin, G., and Rosenschein, J. S. 1990b. Negotiation and conflict resolution in non-cooperative domains. In *Proceedings of the National Conference on Artificial Intelligence*, 100–105.