Rational Communication in Multi-Agent Environments

Piotr J. Gmytrasiewicz* and Edmund H. Durfee**
*Computer Science and Engineering
University of Texas at Arlington, TX 76013
**Department of Electrical Engineering and Computer Science
University of Michigan, Ann Arbor, Michigan 48109
piotr@cse.uta.edu, durfee@.umich.edu

January 12, 2000

Abstract

We address the issue of rational communicative behavior among autonomous self-interested agents that have to make decisions as to what to communicate, to whom, and how. Following decision theory, we postulate that a rational speaker should design a speech act so as to optimize the benefit it obtains as the result of the interaction. We quantify the gain in the quality of interaction in terms of the expected utility, and we present a framework that allows an agent to compute the expected utilities of various communicative actions. Our framework uses the Recursive Modeling Method as the specialized representation used for decision-making in a multi-agent environment. This representation includes information about the agent's state of knowledge, including the agent's preferences, abilities and beliefs about the world, as well as the beliefs the agent has about the other agents, the beliefs it has about the other agents' beliefs, and so on. Decision-theoretic pragmatics of a communicative act can be then defined as the transformation the act induces on the agent's state of knowledge about its decision-making situation. This transformation leads to a change in the quality of interaction, expressed in terms of the expected utilities of the agent's best actions before and after the communicative act. We analyze decision-theoretic pragmatics of a number of important kinds of communicative acts and investigate their expected utilities using examples. Finally, we report on the agreement between our method of message selection and messages that human subjects choose in various circumstances, and show an implementation and experimental validation of our framework in a simulated multi-agent environment.

[0]This research was supported, in part, by the Department of Energy under contract DG-FG-86NE37969, by the National Science Foundation under grant IRI-9015423, by the PYI award IRI-9158473, by ONR grant N00014-95-1-0775 and by the National Science Foundation CAREER award IRI-9702132.

1 Introduction

This paper follows the tradition of cognitive science and related fields [5, 11, 30], according to which the fundamental function of communication is to confer some advantage to the speaker by influencing what the hearer(s) knows and intends to do. The contribution of this paper is to propose a well-defined mechanism that realizes this function in autonomous, self-interested artificial agents that have to decide what to communicate, to whom, and how. We treat decisions about communication just like decisions about any other action, and employ decision-theoretic techniques to select the action with the highest expected utility. As in the case of any action, the expected utility is determined by the desirability of the expected outcome. However, unlike a physical action that changes the physical state of the world, a communicative action changes the state of knowledge of the agents involved. Our framework is aimed at representing the changes in the state of knowledge about the decision-making situation the agents are involved in, quantifying the benefits of communicative actions that bring them about, and allowing rational communicative behavior that executes the communicative actions with the highest expected utility.

Our approach is knowledge-based and relies on a general purpose knowledge base (KB), in our case implemented as a system of classes of objects and their instantiations. To facilitate effective communication, the agent's KB has to include information about the possible states of knowledge, abilities and preferences of the other agent(s) present in the environment [10, 19, 47]. For the purpose of decision-theoretic calculations we use the formalism of the Recursive Modeling Method (RMM) [16, 17]. The advantage of RMM, when used for expected utility calculation, is that is able to succinctly represent the content of the agent's KB, including its preferences, abilities, and beliefs about the physical world, as well as the agent's beliefs about the other agents, their preferences and abilities, their beliefs about the world and about other agents, their beliefs about others' beliefs, and so on. The need for considering the nestedness of the agents' beliefs for communication has been widely recognized in the linguistics and AI literatures before [2, 4, 6, 8, 18, 19, 31, 35, 38, 39], while research in cognitive science [11, 47] yielded evidence of nested mental models used by humans for purpose of communication. Clearly, without a model of the other agents' mental states it would be impossible to properly assess the impact of a communicative act.

We should note that the RMM representation is not intended as a general knowledge representation formalism to be used for multi-agent interactions; this is left to a general purpose knowledge base. Rather, the RMM representation (the payoff matrices and the probabilities) are assembled from the information contained in the KB, and used for the specific purpose of computing expected utilities of alternative courses of action. The alternative actions (physical or communicative) are generated by a symbolic planning module that may be domain specific and that uses information in the KB. The computation of expected utility of alternative physical actions is detailed in [16, 17], while this paper addresses the computation of expected utility of alternative communicative acts.

With each communicative act we identify its *decision-theoretic (DT) pragmatics*, defined as the transformation of the state of knowledge about the decision-making situation the act brings about. We model DT pragmatics using the RMM representation to compute the utility of the communicative act.[1] The transformation in the speaker's decision-making situation, as represented by RMM's recursive model structure, may change the intentions of the other agents, and thus change the expected utilities of the original agent's alternative actions. We define the change of the expected utility brought about by a communicative action as the expected utility of this action itself. By evaluating the alternative communicative acts in this way, the speaker can select and send the highest utility message—the message that causes the greatest gain in the expected utility of the speaker's action.[2] The agent should not communicate if there is no message that results in an increase of the agent's expected utility. We should further note that, in this paper, our approach to computing the values of communicative acts is myopic, i.e., we simplify the analysis by considering only the immediate effects and benefits of such acts.

DT pragmatics of a communicative act differs from its pragmatic meaning, usually defined [6, 18, 39] as the change of the state of knowledge brought about by the act. Imagine two agents engaged together in assembling a bicycle from parts scattered about a garage. A communicative act "The front wheel is in the southwest corner of the garage", uttered by one of the agents, has the pragmatic meaning of changing the other agent's beliefs about the location of the front wheel, if it did not know the location before. This act also changes the decision-making situation the agents are in: The other agent is now in the better position to get the front wheel and complete the bicycle assembly, and the time saved could be of benefit to both agents. The above communicative act, therefore, is endowed with both decision-theoretic pragmatics, as well as pragmatic meaning. But a communicative act "The temperature in the center of Alpha Centauri is 5800 K", uttered in the same situation, with its pragmatic meaning of changing the hearer's state of knowledge about the temperature of the neighboring star, does nothing to the decision-making situation the agents are facing. Therefore, the DT pragmatics of the second communicative act is the identity transformation, and its expected utility is zero.[3]

Clearly, DT pragmatics of a communicative act is uniquely determined by its pragmatic meaning. However, our purpose for defining it separately is that the decision-theoretic calculation of the value of an alternative, but not yet executed, communicative act is substantially simpler if it is performed using the compiled representation of the KB and DT pragmatics. In other words, instead of projecting the effects of a candidate communicative act using the full-blown representation of the KB, it is simpler for the speaker to use compiled representation of this information (in RMM, payoff matrices and selected probabilities). Once the

[1] Since our formalism is probabilistic, it naturally handles cases when the meaning of a message is itself uncertain.

[2] The notion of the utility of a message we use here differs from the notion of the value of information considered in decision theory [34, 38]. The latter expresses the value of information to its recipient. We, on the other hand, consider the value of a message to its sender, since, of course, it is the sender that makes the decision of if, and what, to communicate. The two notions coincide in two special cases, when the preferences of the agents perfectly coincide, and when a speaker requests information from the hearer by asking a question.

[3] Note that the same message may have nontrivial DT pragmatics, as well as a considerable value, in a different decision-making situation.

expected utilities of alternative acts have been computed and the best one has been executed, the speaker updates its KB according to how it expects the executed act to change the state of knowledge of the hearer(s), i.e., according to the act's pragmatic meaning. Thus, the decision-theoretic computations take place before a communicative act is executed, and are performed by the speaker, from the perspective of the speaker's state of knowledge.

Our approach builds on, and also complements, related work on meaning of speech acts [2, 8, 20, 29, 35, 40, 41]. While the above work has concentrated on the communicative acts from the perspective of the agents' declarative knowledge base, usually in the form a set of sentences in predicate calculus (with some additions), we look at the decision-making level instead. Our decision-theoretic approach leads to important differences, however. First, our agents are selfish utility maximizers that autonomously make decisions in the absence of pre-existing protocols. Thus, we do not assume that the agents are cooperatively disposed towards each other, although we certainly do not forbid it. The consequence of not making this assumption reveals itself when we discuss, for example, requests and questions. It turns out that, if exchanged among self-interested (and, in this paper, myopic) agents, these acts seem to lose much of the function we are accustomed to ascribing to them. For example, agents cannot be assumed to automatically answer the questions posed to them; they may not want to reveal information only because another agent wants to know it. Similarly, requests do not impact on the decision-making process of selfish agents and are not guaranteed to invoke a response.[4] Second, while our present implementation uses KQML [29], we are not forced to make the assumption that the agents know and understand KQML. In fact, the speaker is free to implement the acts in any language without a guarantee that the hearer will decode them properly: Our framework can accommodate the probability that the hearer does not understand the messages received, in which case they simply are devoid of value. We give examples of this later in the paper.

In the following sections we first formalize the notions of DT pragmatics and value of communicative acts mentioned above, and then consider examples of various types of communicative acts. Our strongest results address the communicative acts that agents can use to share information about their environment (we call these modeling messages), acts used to express the current intention of the speaker (intentional messages), and acknowledging messages. Other types of communicative acts follow, but are treated in a preliminary form and are intended to present a point of departure for further investigation. We then present results on the agreement between our method of message selection and messages that humans choose, and show an experimental validation of our framework in a simulated multi-agent environment. We close with a brief review of related work and plans for future research.

[4]The agents' selfishness does not preclude that various forms of cooperative behavior will emerge in a society of repeatedly interacting agents [1, 3]. In this paper we do not go into issues of repeated interactions and non-myopic evaluations they necessitate, leaving it to future work.

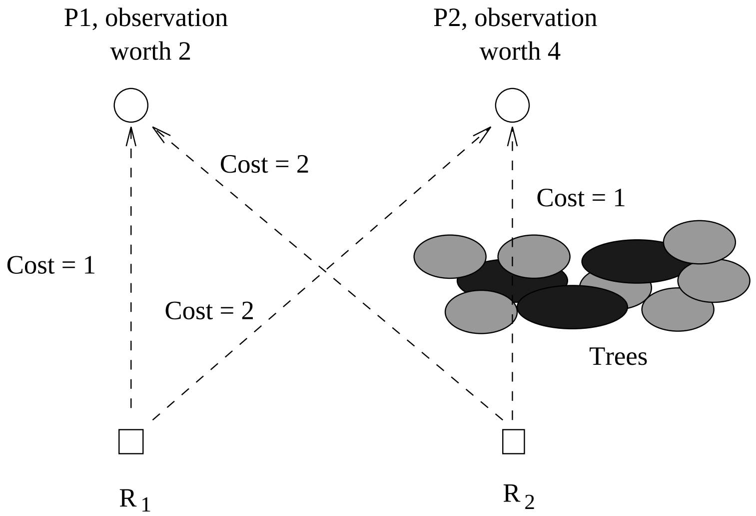

Figure 1: Example Scenario of Interacting Agents

2 Value of Communication – Basic Approach

The value of communication stems from the changes in the beliefs of the agents, and from the improved coordination and overall quality of the interaction that results. We now briefly describe a simple interaction between two agent, and present a compiled representation of a state of knowledge of one of them that we will use during further discussion of communication.

2.1 Simple Example of Interacting Agents

Consider the example of interaction depicted in Figure 1 (described in detail in [17]). It involves two agents, R_1 and R_2, engaged in a common mission of gathering information. We take the perspective of R_1 (R_1 will be the speaker in later examples), who can detect two possible observation points, P1 and P2, allowing observations worth 2 and 4, respectively[5]. Point P1 is closer to R_1 and P2 is closer to R_2, and the costs of getting to the points are assumed to be 1 or 2, as indicated in Figure 1. As we mentioned, this information resides in the agent's general purpose KB. R_1 has to make a decision as to whether to pursue the observation from P1 (we'll label this option a_1^1), from P2 (a_2^1), or do neither and just sit still (a_3^1), and would like to do so in a way the maximizes the total value of information obtained by both agents, since it's a joint mission, reduced by its own cost. We assume that these two factors are the only ones that determine R_1 expected utility in this case. Note that the expected utilities of R_1's actions depend what it expects R_2 to do. If R_2 observes from P2 then R_1 is best off observing from P1 for the total payoff of $2 + 4 - 1 = 5$, i.e, total value of observations minus R_1's own cost. But if R_2 decides to observe from P2 or do nothing at

[5]Say, because P2 is higher or allows observation of a more interesting area.

all, then it's best for R_1 to observe from P2. The expected payoffs of alternative behaviors of R_1 can be assembled into a payoff matrix, on top of the structure in Figure 2. Note that, as we mentioned earlier, the information contained in the payoff matrix is a compilation of information contained in R_1's KB.

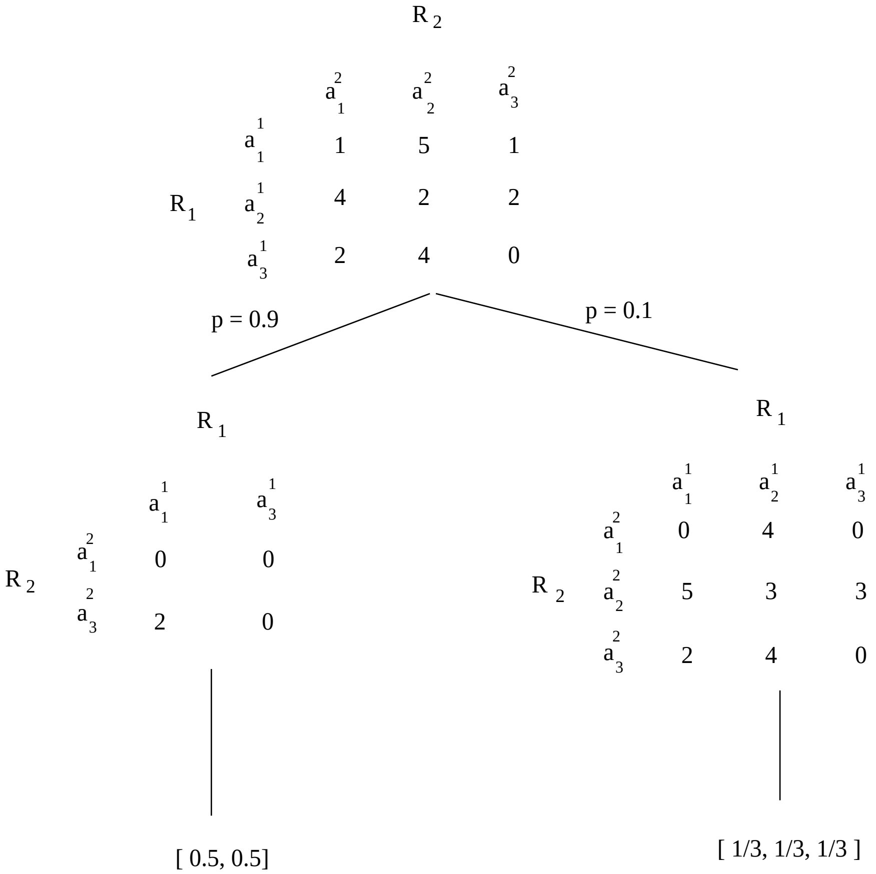

Figure 2: Recursive Model Structure depicting R_1's Decision-Making Situation in Example 1.

Given the situation, it is clear that R_1 should try to coordinate its actions with R_2. To coordinate, R_1 needs to predict what R_2 will do.[6] The difficulty in predicting R_2 action, in this case, is two-fold. First, there are trees located between R_2 and point P2, so it's likely that R_2 is not even aware of P2. Second, it seems that to predict R_2's behavior R_1 may need to figure out what R_2 expects of R_1, which in turn may depend on what R_2 thinks R_1 expects of R_2, and so on. A solution to the first problem we suggest is for R_1 to maintain two models of R_2, one for each possibility of its knowing about P2 or not. These two models, again represented as payoff matrices, are depicted on the second level of structure in Figure 2, which we call a recursive model structure. The predictions of R_2's behavior generated by each of these models can be combined with weights equal to probabilities associated with the models to yield the

[6]See [17] for more formal definition and more details.

overall prediction of R_2's behavior. This is called Bayesian model averaging; see [17, 23] for further details.

To handle the issue of predicting the other agent's action, while the other agent attempting to do the same, we suggest a knowledge-based approach. Intuitively, instead of attempting to guess what the other agent will do, based on what its guess is as to what the original agent will do, etc., the agent should simply represent all of the information is has about the other agent, about what the other agent knows about the original agent, and so on. We argue (see discussion in [17] and references therein) that in realistic situations the information the agent has is finite and has to terminate at some finite level of nesting. Thus, the representation of this information is a finitely nested hierarchy of models that can be processed bottom-up.[7]

For the purpose of the current example we assumed that the agent R_1 knows that R_2 has no information it can use to model R_1. That means that the recursive model structure representing R_1's decision-making situation in this scenario, depicted in Figure 2, terminates at the leaves with, what we call, no-information models.[8] Thus, R_2's lack of any information about R_1 is represented as uniform probability distributions on the third level of the structure. They precisely correspond to R_2's lack of knowledge about R_1, since they contain no information about R_1's action. The two models that R_1 has of R_2's decision-making situation, on the second level in Figure 2, reflect R_1's uncertainty as to R_2's being able to see point P2. In this case we assumed that R_1, given the density of the foliage between R_2 and P2, assigns a probability 0.1 to R_2's being able to see through the trees, and a probability of 0.9 to it not being able to see P2. We call them *modeling probabilities*. In general, modeling probabilities are associated with alternative models, or branches, on any level of the recursive model structure.

The bottom-up solution of the structure in Figure 2 amounts to computing the expected behaviors of agents given what they, in turn, expect of other agents. In the right branch, for example, given that R_2 assigns equal probabilities of $\frac{1}{3}$, the expected utilities of R_2's actions can be computed as: $\frac{1}{3}(0+4+0) = \frac{4}{3}$, $\frac{1}{3}(5+3+3) = \frac{11}{3}$, and $\frac{1}{3}(2+4+0) = \frac{6}{3}$, for the consecutive alternatives. Thus, if R_2 can see P2, its best alternative is a_2^2, i.e., to pursue the observation from P2. Analoguous analysis of the other model shows that if R_2 cannot see P2 then its a_2^2 is best and it will remain stationary. These two predictions can be probabilistically mixed with weights equal to 0.9 and 0.1 yielding an overall estimate of what actions R_1 can expect $R2$ to perform, which will be called the *intentional probability distribution*, or the conjecture. In this case, the intentional probability distribution is: $p_{R_2}^{R_1} = [0, 0.1, 0.9]$, i.e., R_1 is certain R_2 will not pursue observation from P1, estimates that there is 10% probability that R_2 will observe from P2, and that there is 90% probability that R_2 will stay put. The best choice for R_1 is then to pursue its option a_2^1, that is to move toward point P2 and make an observation from there, with its expected utility of 2.

[7]In [17] we describe a dynamic programming implementation of this precess, also called a decision-theoretic approach to game theory. The bottom-up version is simplified and sufficient here, however.

[8]Structure in Figure 2 is a simplified version of recursive structures presented in [17].

2.2 Defining the Value of Communication

We now define the notions needed to compute the expected utilities of alternative messages in a general setting. As we mentioned, the expected utilities are computed by the speaker agent before the best communicative act is executed. We will use the example scenario when we describe examples of various types of communicative acts.

We first define the space, $\mathbf{RMS_{R_i}}$, of the recursive model structures, RMS_{R_i}, representing the decision-making situation of the speaker agent, R_i. An example is depicted in Figure 2. Further, we define the set, $\mathbf{M}$, as the set of the communicative acts agent R_i can perform. We will assume for simplicity that this set is finite and it consists of alternatives generated, for example, by a communication planning module. The elements of the set $\mathbf{M}$ are communicative acts that differ in the content of the communicated information, but also differ in the way this content is encoded (the language used), and in the communication medium used for its transmission. Thus, a message in English over a phone constitutes a different communicative act from a message in German transmitted over email, even if the two messages translate into the same content.[9]

We now formally define the decision-theoretic pragmatics.

Definition 1: *DT pragmatics is a function* $Prag_{DT} : \mathbf{RMS_{R_i}} \times \mathbf{M} \rightarrow \mathbf{RMS_{R_i}}$.

DT pragmatics of a communicative act M, $Prag_{DT}(RMS_{R_i}, M) = RMS_{R_i}^M$, is the transformation M induces on the recursive model structure, RMS_{R_i}, of the speaker agent R_i. We write $RMS_{R_i} \xrightarrow{M} RMS_{R_i}^M$ as an intuitive notation for M's decision-theoretic pragmatics. We will call RMS_{R_i} the prior structure, and $RMS_{R_i}^M$ the projected structure, in the context of M.

We should remark that the set of the communicative acts can, in principle, encompass all of the courses of action as generated by a planning module. Some of these actions can be physical in nature, which means that physical actions may be treated as communicative actions. Also, as we mentioned, we have not made any commitments as to the agents sharing a communication language, and we will not attempt to define a language in this paper. While our implementation uses KQML [29, 41], our framework is general and we will consider communicative acts as changing the decision-making situation of the agents, but without detailing how these acts are actually executed.

As we described in the preceeding section, R_i can analyze its recursive model structure and use dynamic programming to solve it to yield the intentional probability distributions, or conjectures, of the agents that it interacts with. In what follows, we define $p_{\{-R_i\}}$ as a shorthand for R_i's prior conjecture, i.e., the expected behaviors of all agents except R_i, obtained by solving the prior structure RMS_{R_i}. Let us denote as X the rational choice of

[9]This distinction may seem counterintuitive since content is the same, but note that the pragmatics of these two communicative acts may be quite different if, for example, the recipient does not speak German.

the agent R_i based on the prior conjecture, and its expected utility as $U_{p_{\{-R_i\}}}(X)$ (in the example interaction above, the R_1's prior conjecture is $[0, 0.1, 0.9]$, while its best action, X, is to observe from P2, with its utility of 2.) Similarly, $p^M_{\{-R_i\}}$ will stand for R_i's projected conjecture, obtained by solving the projected structure $RMS^M_{R_i}$. Further, denote the rational choice of the agent R_i based on the projected conjecture as Y (which may or may not be the same as X), with its expected utility of $U_{p^M_{\{-R_i\}}}(Y)$. We can now define the (myopic) expected utility of a communicative act.

Definition 2: *Expected utility*[10] *of the communicative act M is the difference between the payoff the agent expects before and after executing the act:*

$$U(M) = U_{p^M_{\{-R_i\}}}(Y) - U_{p_{\{-R_i\}}}(X). \tag{1}$$

Definition 3: *A trivial message is a message, M, whose DT pragmatics is $Prag_{DT}(RMS_{R_i}, M) = RMS_{R_i}$, i.e., for which the prior structure and the projected structure are identical.*

Thus, trivial messages are ones that do not change the decision-making situation of the speaker. An example of a trivial message is one in a communication language that the speaker knows the hearer cannot understand.[11]

Corollary 1: *The expected utility of a trivial message is zero.*

This corollary follows directly from definitions 2 and 3. We have $U_{p^M_{\{-R_i\}}}(Y) = U_{p_{\{-R_i\}}}(X)$ for any trivial message M since the utilities are calculated based on the prior and projected structures which are identical.

The uselessness of trivial messages is intuitive, but there may be examples of nontrivial messages that are useless as well. Further, some nontrivial messages may have expected utilities that are negative, in contrast to the usual decision-theoretic notion of the value of information, which is never negative [38]. We give examples of these in the following sections.

It may be useful to classify communicative acts into types. Some of the types we consider have close correspondents in speech act theory [2, 46], and in various kinds of performatives considered in KQML [29].

[10] We should remark that we are using a concept of cardinal utility, so that the difference between utilities is well defined. See, for example [22] for discussion.

[11] The message that is not understood is trivial under usual circumstances. Under some circumstances such a message may be of value, as in the "German soldier" example described by Searle, but it is clearly not trivial in the light of Definition 3 in this paper.

3 Modeling Messages

Our modeling communicative acts update the hearer's and the speaker's model of the multi-agent world. The close correspondents of these type of communicative acts in speech act theory are the inform, assert, and tell acts.

Definition 4: *Modeling communicative acts are ones that contain information about the modeling probabilities, which represent R_i's beliefs about the other agents in the environment.*

The modeling probabilities in the above definition are the probabilities associated with different models, or branches, in the recursive model structure. They are rigorously defined in [17].

Consider again the example of interaction depicted in Figure 1, and the recursive model structure representing R_1's decision-making situation in this scenario in Figure 2, also depicted on the left in Figure 3. Assume, for the time being, that both agents can understand and generate communicative acts in English. Consider what would happen if R_1 were to send a message, M_1, stating "There is an observation point P2, twice as high as P1, behind the trees." Assuming that R_1 estimates that the M_1 is certain to reach R_2, the decision-theoretic pragmatics of M_1 is as depicted in Figure 3.

The DT pragmatics of M_1 illustrates the changes in modeling probabilities: the probability of the first model, $p_1^{R_1}$, according to which R_2 does not know about point P2, changed from 0.9 to zero (it is, therefore, not included in Figure 3), and the probability of the second model, $p_2^{R_1}$, increased to 1, since R_1 can be sure that R_2 will know about the point P2 as a result of the message having been sent. According to definition 4, therefore, message M_1 is, under the assumed circumstances, a modeling message.

The projected structure, on the right in Figure 3, can be easily solved, showing that R_1 would be sure that R_2 would observe from point P2, taking action a_2^2. Thus, the projected conjecture that R_1 ascribes to R_2 is $p_{R_2}^{M_1;R_1} = [0, 1, 0]$. The best alternative for R_1 according to projected structure is to make an observation from P1, but the expected payoff has increased to 5. Thus, by sending the message M_1 to R_2, R_1 was able to increase the expected utility it gets from the interaction from 2 to 5. As defined in Equation 1, the utility of sending the message M_1 is $U(M_1) = 5 - 2 = 3$. This illustrates how our approach implements the fundamental function of communication, which is to confer some advantage to the speaker by influencing what the hearer knows and intends to do.

The above analysis assumes that R_2 is guaranteed to receive and properly decode the content of R_1's communicative act. However, it may be that R_2 does not understand English, or that R_1 used an unreliable communication channel. As we mentioned, R_1's attempt to transmit the content above would then, formally, constitute a different communicative act, $M_{1.1}$. $M_{1.1}$ also has a different, although still well defined, DT pragmatics. Let us represent the imperfections in $M_{1.1}$'s transmission by assigning a probability, p_c ($0 \leq p_c \leq 1$), to R_2's

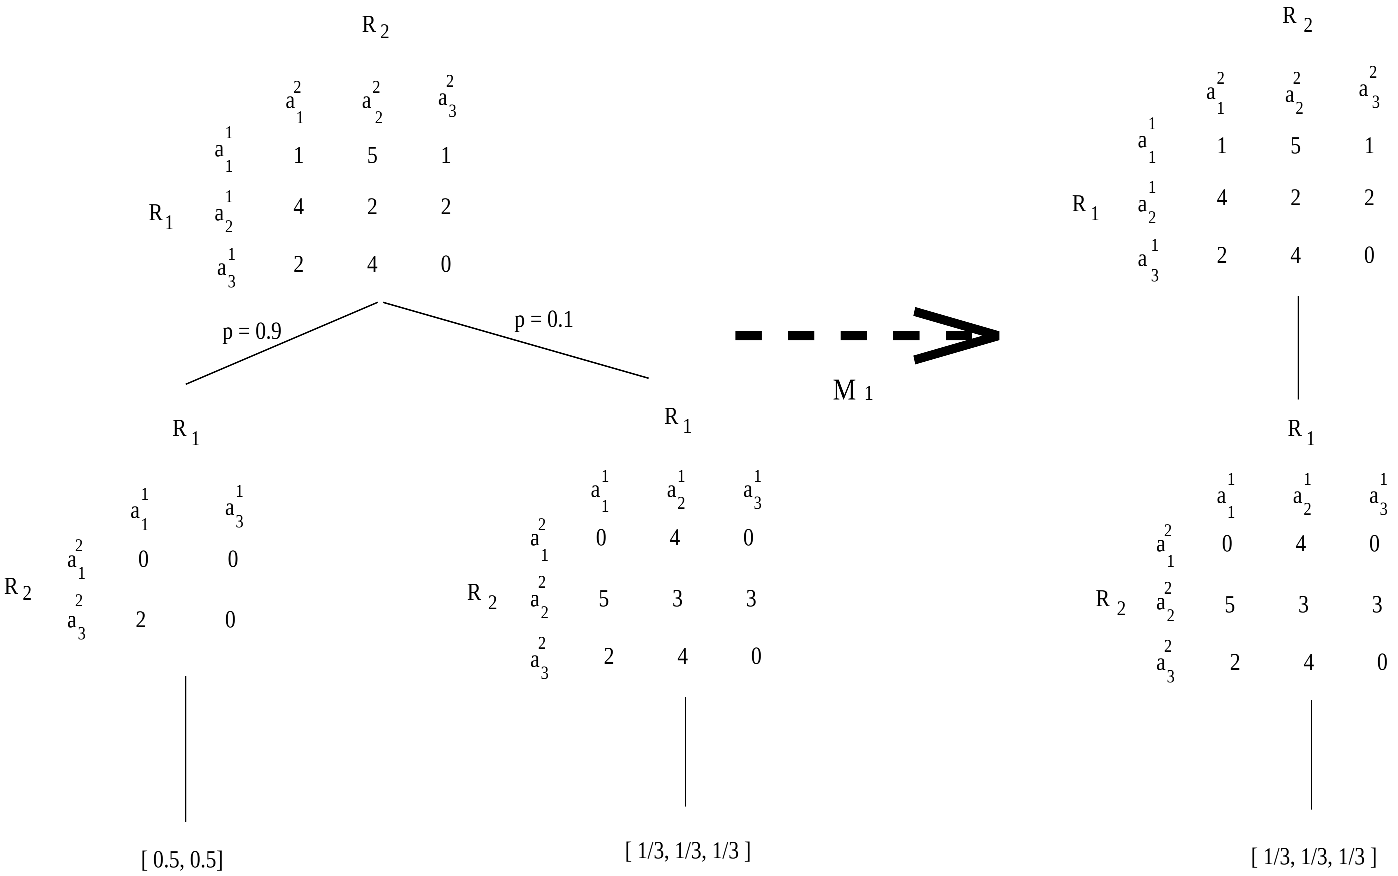

Figure 3: Decision-Theoretic Pragmatics of Modeling Message M_1.

properly receiving and understanding the content of $M_{1.1}$.[12] Then, DT pragmatics of $M_{1.1}$ is as depicted in Figure 4.

Solving the projected structure in Figure 4 reveals that the intentional probability distribution describing R2's action is $[0, 0.1 + 0.9p_c, 0.9 - 0.9p_c]$. The expected utilities of R_1's alternatives can now be computed as:

$$\begin{aligned} u^{R_1}_{a^1_1} &= 1.4 + 3.6p_c \\ u^{R_1}_{a^1_2} &= 2 \\ u^{R_1}_{a^1_3} &= 0.4 + 3.6p_c \end{aligned}$$

From the above we can see that, if $p_c > 1/6$, the value of the message $M_{1.1}$ depends on p_c as: $U(M_{1.1}) = 3.6p_c - 0.6$. If $p_c < 1/6$, R_1 would prefer to choose a^1_2 and observe from P2,

[12]For simplicity we neglect here the various ways in which $M_{1.1}$ can be misunderstood.

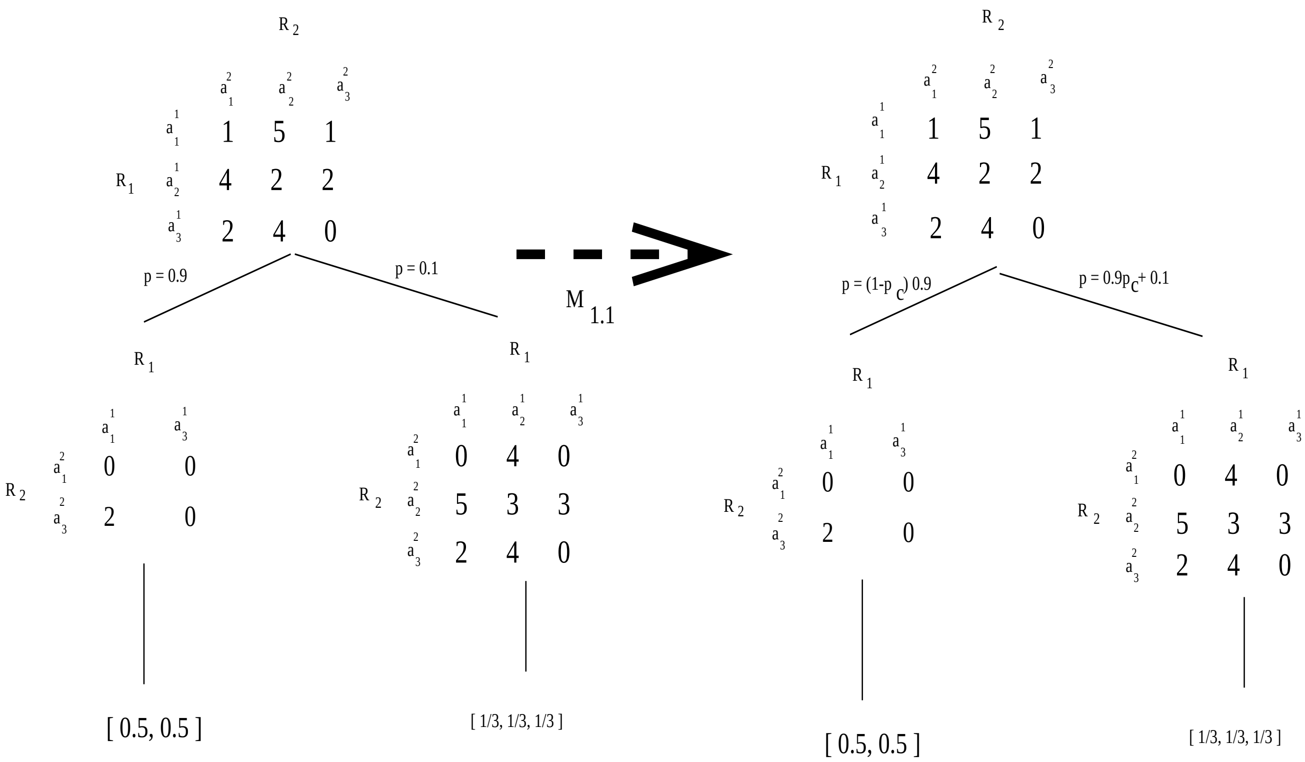

Figure 4: DT Pragmatics of an Unreliable Modeling Message $M_{1.1}$.

with its payoff of 2. This is the same as without communication, so, if $p_c < 1/6$, the expected utility of$M_{1.1}$ is zero.

4 Intentional Communicative Acts

The purpose of intentional communicative acts is to inform other agents about the speaker's current intentions.[13,14] These acts loosely correspond to promise acts in the speech act theory, but they do not imply the notion of commitment. For example, an agent may inform another agent of its current intention to perform some action, but, say in view of newly acquired information, it is free to change its intention. Note, however, that it would be in this agent's best interest to inform the other agent about the change of plans by sending another intentional message.

[13]Here, intention is the current best behavioral alternative considered by the agent.

[14]Actually, the speaker may want to inform another agent about what it thinks other agents will do, as well as about what it thinks others expect other agents to do, and so on. All of these acts are of intentional kind according to our definition.

Definition 5: *Intentional communicative acts contain information about the intentional probabilities* $p_{R_k}^{R_i,R_j}$*, that represent* R_i*'s belief about an agent* R_j *expectation as to another agent's,* R_k*, actions.*

In the simplest and most intuitive case, the speaker declares its own intentions and, in the above definition, R_i and R_k are the same speaker agent. In general, the intentional probabilities are conjectures the agents use to describe the expected actions of other agents. They are rigorously defined in [16, 17].

For the purpose of the present discussion we assume that the truthfulness of these messages is guaranteed (as we mentioned the agents' intentions may change, but it is in their best interest to inform others of such changes. See also [15] for cases involving lying). Thus, a hearer can use an intentional message to predict what the speaker will do. In modeling the hearer, therefore, the speaker can truncate the projected recursive structure, because it knows that the hearer's conjecture of the speaker's actions correspond exactly to the content of the intentional message.

Let us take again the interaction in Figure 1, and suppose that R_1 considers using a perfect communication channel to transmit a message "I will observe from point P2" to R_2. Let us denote this communicative act as M_2. DT pragmatics of M_2 is as depicted in Figure 5.

If R_2 is not aware of the point P2 and receives M_2, it models R_1 as pursuing its a_3^1 option, i.e., doing something other than observing from P1 (labeled a_1^1),[15] and so, for R_2, the options a_1^2 and a_3^2 are equally good and equally likely. If R_2 can see the point P2 and receives M_2, its options a_1^2 and a_3^2 are also equally good. Thus, the new overall intentional probability distribution over R_2's options is $p_{R_2}^{M_2;R_1} = [0.5, 0, 0.5]$. Now, the expected utility of R_1's action a_2^1 increased to 3. According to Equation 1, the utility of M_2 is $U(M_2) = 3 - 2 = 1$.

The above analysis assumes that R_2 is guaranteed to receive and properly decode M_2. If the reliability of the communication used is characterized by the probability p_c instead, the intentional probabilities R_1 ascribes to R_2 will be:

$$p_c[0.5, 0, 0.5] + (1 - p_c)[0, 0.1, 0.9] = [0.5p_c, 0.1 - 0.1p_c, 0.9 - 0.4p_c].$$

The expected utilities of R_1's alternatives can now be computed as:

$$\begin{aligned} u_{a_1^1}^{R_1} &= 1.4 - 0.4p_c \\ u_{a_2^1}^{R_1} &= 2 + p_c \\ u_{a_3^1}^{R_1} &= 0.4 + 0.6p_c \end{aligned}$$

[15] In human communication, this message would likely be interpreted as a conjunction of two messages. First one saying that there is a point P2 out there (modeling), and the second saying that R_1 will do something about it (intentional). In this discussion we treat M_2 purely as an intentional message.

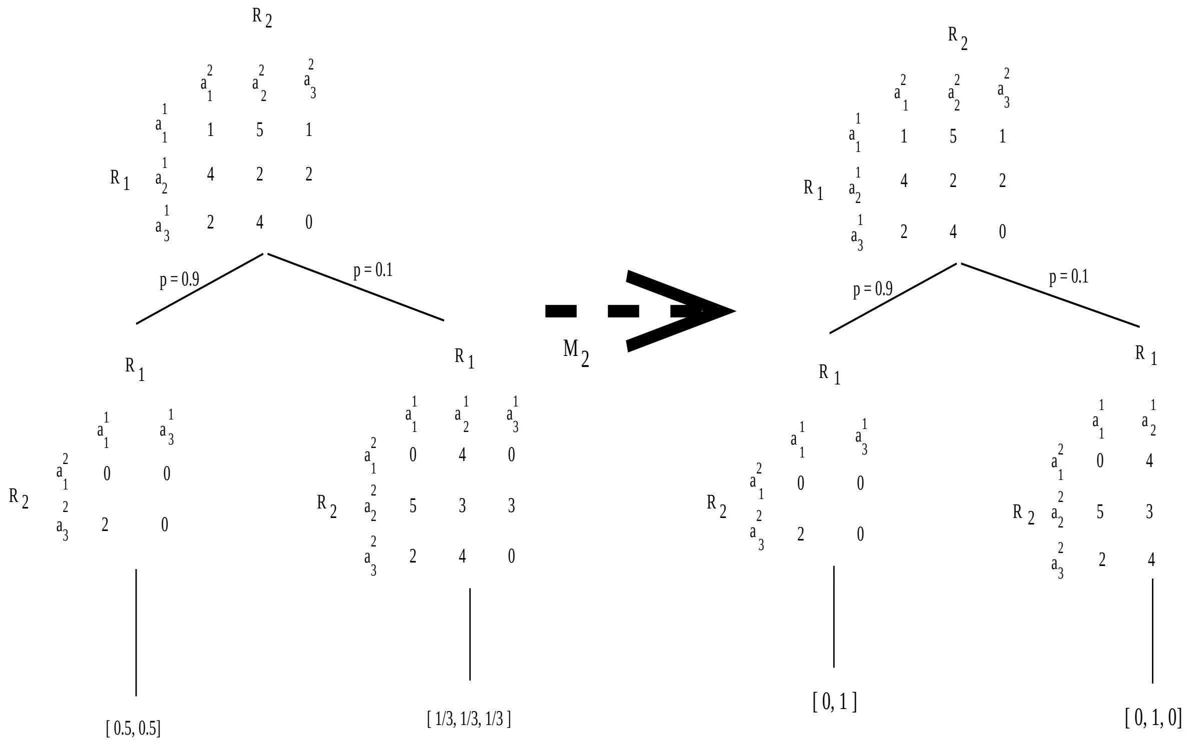

Figure 5: DT Pragmatics of Intentional Message M_2.

Thus, the expected utility of sending this message over an unreliable channel is equal to the probability p_c.

5 Modeling vs. Intentional Acts – Discussion

Note that, assuming reliable communication, the intentional message M_2 from the preceding section was less valuable than the modeling message considered in Section 3. That is an intuitive result; in the situation in which an agent is quite unlikely (with only a 10% chance) to be aware of an important feature of the environment, in our example point P2, human speakers would also tend to favor the message telling the agent about this feature.

But let us now look at a different situation, depicted in Figure 6, in which there is no uncertainty about how R_2 sees the environment—R_2 is sure to see P2. However, now there is uncertainty about the agents' properties. Assume that R_1 knows that observation point P1 also provides a sunny spot where a vehicle with a solar array can re-charge its batteries. While R_1 has no solar array on board, it believes that R_2 might, and that re-charging is worth 2 to it. Thus, if R_2 has a solar array, it will value going to P1 as 4 (2 for observing plus 2 for re-charging), but, if it has no array, then R_2 will value P1 the same as R_1 does (at 2).

Moreover, R_1 is camouflaged to resemble a bush (R_2 is clearly an agent), and R_1 estimates that it is equally likely that its disguise will fool R_2 as it is that R_2 will correctly recognize

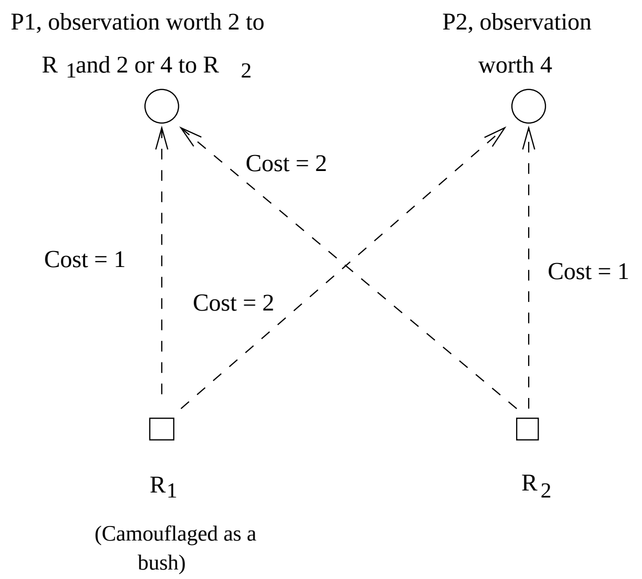

Figure 6: **Another Scenario of Interacting Agents**

R_1 as a robotic vehicle without a solar array. In the case where the disguise fools R_2, it will assume that R_1 will stay still.

The recursive model structure representing R_1's state of knowledge is depicted on the left in Figure 7. It can be solved using dynamic programming and results in R_1's expected utility equal to 3. Since R_1 knows that R_2 can see both observation points, a modeling message that tells R_2 about any of the observation points is bound to be useless. Now, however, an intentional communicative act may be more appropriate; say that R_1 considers the content "I will observe from P1" sent over a perfect channel, which we will call M_3. M_3 's DT pragmatics is the transformation in Figure 7, and it is similar to DT pragmatics of M_2 before.

It is easy to see that R_1 would conclude that, after receiving M_3, R_2 would choose to make an observation from point P2, resulting in the expected utility of observing from P1 for R_1 as 5. The value of M_3 is the difference between 5, expected after sending the message, and 3 expected before, i.e., 2.

Interestingly, there is a modeling message that R_1 could also consider in this case. It is M_4 with content "I am not a bush". DT pragmatics of M_4 is simple: R_2 would have just one correct model of R_1's decision-making situation. The expected utility of this message turns out to be 0.5 (we invite the reader to check this result), thus, as we would expect, M_4 is viable and beneficial in this case, although not the best.

This example illustrates how, using our approach, rational communicative agents can tailor the types of messages they send to each other based on the particular situation at hand. In some situations it's best to talk about the features of the environment, but in other cases it's best to talk about the agents' intentions.

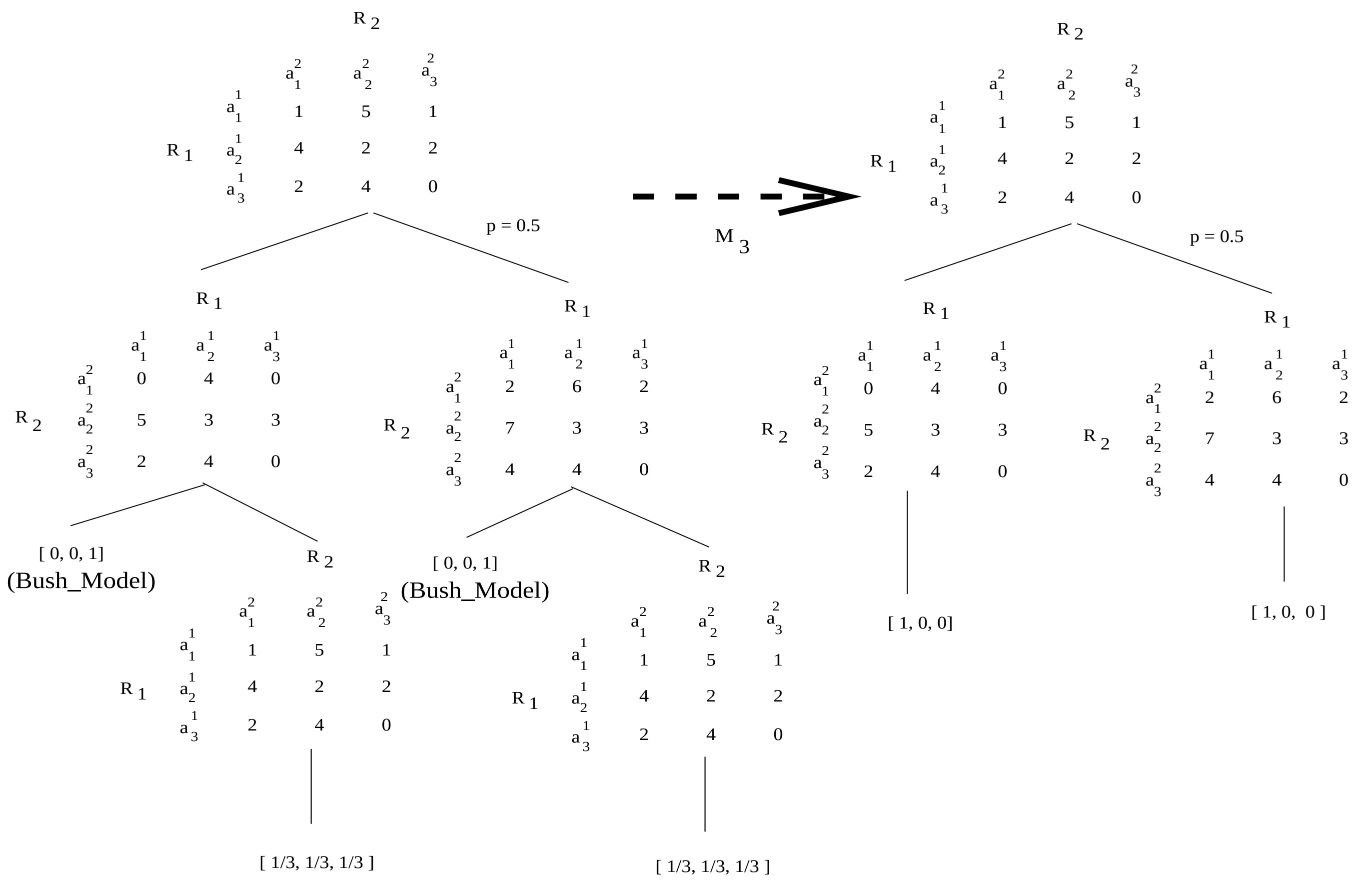

Figure 7: DT Pragmatics of Intentional Message M_3.

6 Acknowledging Acts

We now turn to acknowledging messages. We assume that acknowledging messages can be sent to confirm the reception of other messages. Consider the recursive model structure depicted in Figure 8. It represents information similar to what R_1 had in the example before, on the left in Figure 3, but now R_1 knows more: it knows that, if R_2 can see point P2, it will know that R_1 sees it also, and it will model R_1's payoff matrix correctly. Further, R_1 knows that, if R_2 can see P2, it will be aware that R_1 is uncertain whether R_2 can see P2, and that results in the branching on the third level of nesting.

Now let us imagine that R_1 has received a modeling message, M_5, from R_2 informing it that R_2 can see point P2, sent over an imperfect communication channel, and can acknowledge having received it. After having received M_5, R_1 can update its recursive model structure by eliminating the left first-level branch, resulting in the structure depicted on the left in Figure 9. Let us note that the branching on the third level remains, as it corresponds to R_2's thinking that R_1 may still not know that R_2 can see P2. The probability of this branching has been updated to include the probability that R_1 thinks R_2 assigns to the reliability of the communication channel used: $p_c^{R_1,R_2}$.

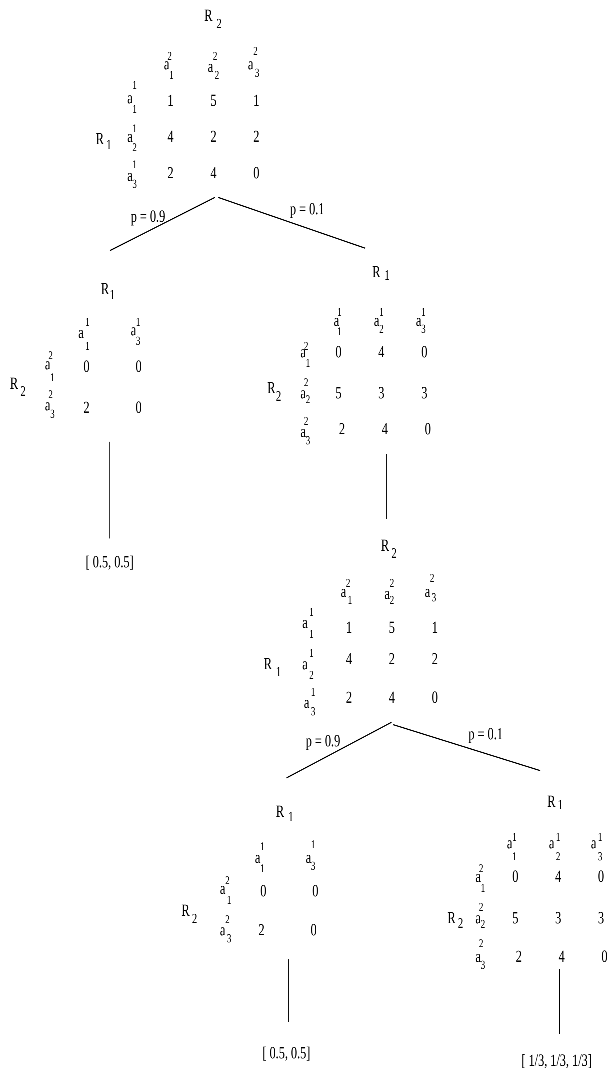

Figure 8: R_1's Recursive Model Structure before it Received Message M_5.

The solution to this structure depends on the value of $p_c^{R_1,R_2}$. If $p_c^{R_1,R_2} < 1/6$ then the intentional probability R_1 would think R_2 ascribes to R_1 would be $p_{R_1}^{R_1,R_2} = [0,1,0]$; if $p_c^{R_1,R_2} = 1/6$ then $p_{R_1}^{R_1,R_2} = [0.5, 0.5, 0]$; and if $p_c^{R_1,R_2} > 1/6$, we have $p_{R_1}^{R_1,R_2} = [1,0,0]$. The first of these cases, propagated upward, gives $p_{R_2}^{R_1} = [0.5, 0, 0.5]$, and the last two result in $p_{R_2}^{R_1,} = [0,1,0]$. The first of these, in turn, points to a_2^1 as R_1's best choice with the utility of 3, while the second makes a_1^1 rational with an expected payoff of 5.

The pragmatic meaning of acknowledging the receipt of M_5, called M_6, sent over a reliable communication channel, is depicted in Figure 9. The effect of this message, if received, would be to eliminate the possibility for R_2 that R_1 does not know that R_2 can see point P2, i.e., to modify the modeling probabilities on the third level of nesting. In general, one can consider further acknowledgments, that is, acknowledgments that previously sent acknowledgments

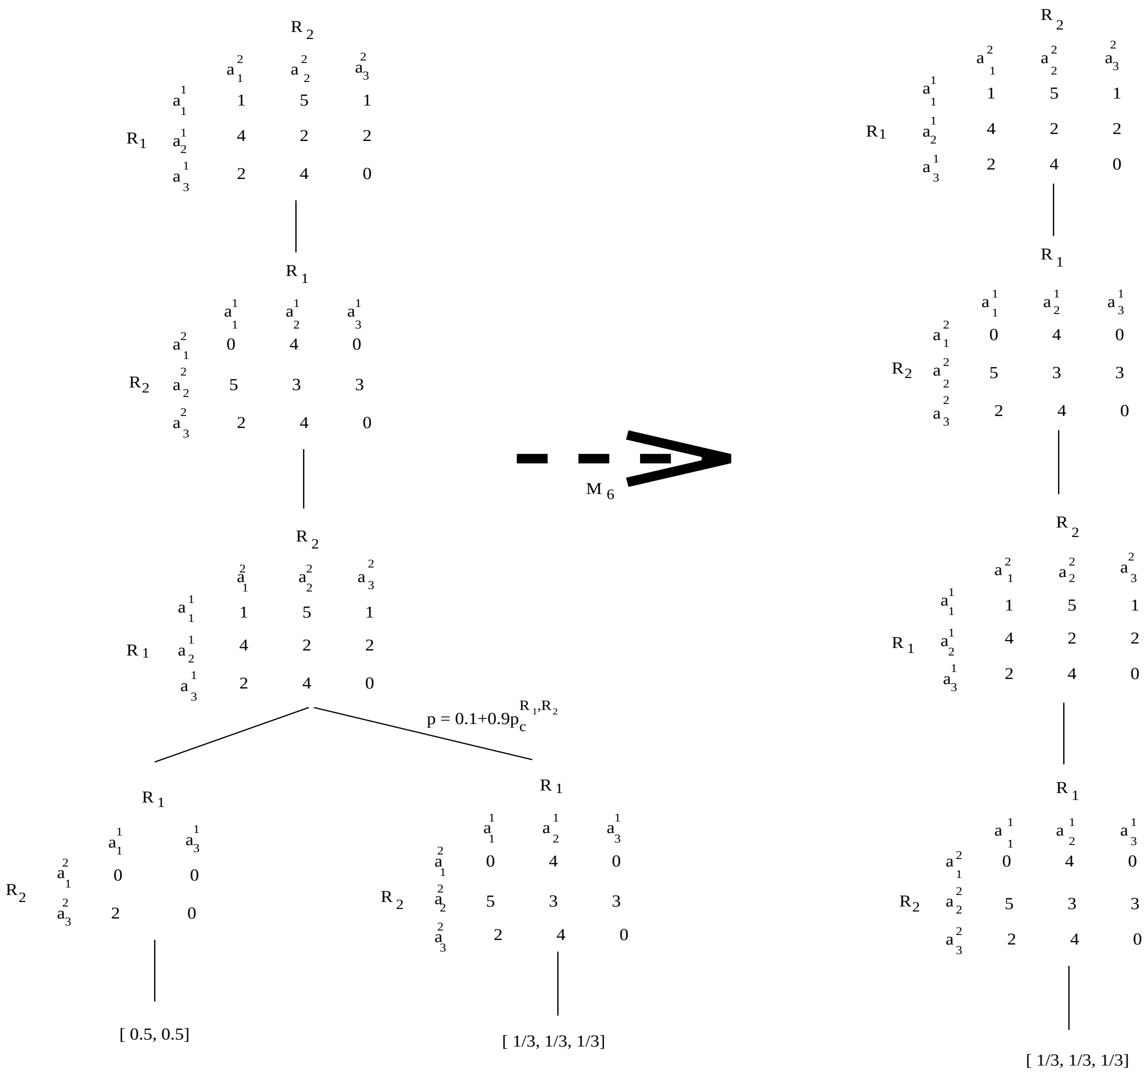

Figure 9: DT Pragmatics of Acknowledging Message M_6.

have been received, and so on. All of them modify the modeling probabilities on deeper levels of the recursive model structure. These considerations lead to the following observation:

Observation 1: The communicative acts acknowledging the receipt of previously sent messages are modeling messages, as previously defined.

The projected structure of M_6, on the right in Figure 9, can be easily solved and results in R_1's anticipating that R_2 will observe from P2: $p_{R_2}^{R_1,'} = [0,1,0]$, which makes a_1^1 R_1's best option with the expected utility of 5.

We can now see that the value of the acknowledging message M_6 depends on how R_1 thinks

R_2 evaluates the reliability of the communication channel $p_c^{R_1,R_2}$. If $p_c^{R_1,R_2} < 1/6$, then the value of M_6 is 2, otherwise M_6 turns out to be useless. This confirms our intuitions about acknowledging messages: If R_1 thinks that R_2 thinks that the original message had a good chance of getting though, then the acknowledgment is not needed; otherwise it is in R_1's best interest to acknowledge.

An interesting situation arises if the acknowledgment is sent over an unreliable communication channel or encoded in a language that may be unknown to R_2, characterized together by the probability p_c, as before. The result is that for $p_c^{R_1,R_2} < 1/6$, the value of such acknowledgment is actually negative for $p_c < 0.5$ (the intentional distribution that R_1 would ascribe to R_2 becomes then $p_{R_2}^{R_1,'} = [0.5(1-p_c), p_c, 0.5(1-p_c)]$). This illustrates the point mentioned before: There is no guarantee that the communicative acts considered will have nonnegative expected utilities. This is in contrast to the theory of value of information [38], and it exemplifies a counterintuitive fact that communicating truthful messages among agents engaged in a cooperative interaction may be ill advised, if the reliabilities of transmission are low.

7 Questions

The ability to ask questions features prominently in human communicative behavior, and they correspond to the *ask* act in the speech act theory. In our approach, however, the following issue arises: Why should a fully autonomous and myopically selfish agent pay any attention to other agents' requests for information? As we have shown above, the computation of the utilities of messages and actions is performed exclusively from the point of view of a given agent, and the fact that another agent would like to receive some information does not enter into these calculations. Therefore, there is no guarantee that questions will be answered among myopic selfish agents.

One way to explain cooperative communicative behavior is to view it as dictated by some interaction protocol. One such protocol may be a social convention, say, humans use to answer truthfully a question about the current time asked by a stranger on a street. Such protocols may be useful to rational individuals since they could save time needed for deliberation. However, since our approach in this paper in orthogonal to the issue of protocol design, we will not pursue this avenue here.

Another approach is to relax the assumption that agents are myopic and allow them to consider long-term effects of their interactions. As related research in game theory shows [1, 3], a key to cooperative behavior among selfish agents is the agents' anticipation that they will interact repeatedly. When the agents look at the long-range repercussions of their behavior during repetitious interactions it is rational for them to be cooperative in order to elicit cooperative behaviors from others during future encounters. While we have studied repeated interactions within the RMM framework before [14], we will leave a further formal study for future work.

Below, we provide an informal analysis from the point of view of non-myopic agents to give a flavor of what extensions this could allow. We also found it useful to view questions not as requests (or demands) for information, but as declarations of lack of knowledge.

Let us consider the scenario depicted in Figure 1 again, but let us modify it slightly such that now the stand of trees between R_2 and P2 consists of only a single tree. Assume that R_1 now expects R_2 to be very likely to see P2 (say, with probability .99). Repeating the solution process we went through in Section 2.1, in this case R_1 would compute that R_2 will stand still with probability 0.01 (if it does not see P2), and will observe from P2 with probability 0.99. This gives R_1 expected utilities of 4.96 for observing from P1 (a_1^1), 2 for observing from P2 (a_2^1), and 3.96 for staying still (a_3^1).

As detailed in Section 3, if R_1 were to send R_2 a modeling message M_1 to ensure that R_2 knew about P2, then R_1 could increase its expected utility to 5 (assuming correct message transmission). Thus, in this case, the message would increase the expected utility by $5.0 - 4.96 = 0.04$. Assuming that it costs nothing to send a message, sending M_1 would be beneficial. But let us assume that sending a message costs something, say 0.1. Now the utility of the message minus the cost of the communicative action is $0.04 - 0.1 = -0.06$, so sending M_1 would not pay for R_1. Intuitively, R_1 is sufficiently sure that R_2 has the relevant information it needs, so it does not pay to transmit information that is very likely to be redundant.

However, now imagine that R_1 receives from R_2 a message, M_7, declaring R_2's ignorance: "I cannot see through the tree". The immediate pragmatics of M_7, which plays a role similar to that of a question, is to cause R_1 to transform its recursive model structure leading it now to believe that R_2 only knows about P1 (formally, therefore,M_7 is a modeling message). Based on this, R_1 conjectures that R_2's behavior is described by the intentional probability distribution of [0, 0, 1], meaning that R_2 will sit still. In turn, now R_1 expects utilities of 1 for a_1^1, 2 for a_2^1, and 0 for a_3^1. Its best action now is a_2^1, with an expected payoff of 2; the message it received caused it to revise its expected utility downward, from 4.96 to 2.

But now R_1 should reconsider sending the message M_1, informing R_2 about P2, providing the information and, in effect, answering the question. As before, successfully sending the message leads R_1 to expect a payoff of 5, which is much better than its current expectation of 2. The utility of the message minus the cost of sending it is $3 - 0.1 = 2.9$, and R_1, being rational, will respond to R_2's M_7 with M_1, and inform it about point P2. That means that M_7, while declaring ignorance, would elicit an informative response from R_1, living up to our interpretation of it as an effective question that R_2 asked R_1.

The above establishes that declarations of lack of information can elicit an informative response from an autonomous self-interested agent, in effect functioning as questions. Now, let us turn to asking the question in the first place. Looking at the situation from R_2's perspective, all it sees in its environment are R_1, P1, and a tree. Based on prior knowledge (for example, that observation points commonly come in pairs), R_2 might hypothesize that with probability of , say 0.4, there is another observation point hidden behind the tree. If it assumes that this observation point will be worth 2 like P1,[16] and that it will cost 1 to get there, the expected utility for R_2 to go toward the hoped-for observation point is the probability the point is there times the worth of the point, minus the cost of going there: $(0.4 \times 2) - 1 = -0.2$. This negative expected utility means that it would be irrational for R_2

[16]Note that, in reality, this is an underestimate.

to act on the hunch that another observation point might be behind the tree.

But here is when asking a question, stated as a declaration of ignorance, can help. To compute the value of the question R_2 has to look not only at immediate effects, but also at possible responses. Say that R_2 believes that there is a 0.4 probability that it will receive an affirmative answer (using the prior knowledge above), and in that case it goes to the point to gain a payoff of 1 (since the expected worth is 2 and the expected cost of going to the point is 1). With 0.6 probability, it will receive a negative answer, and will stay still, gaining no additional payoff from its actions. Since it is currently expecting to gain nothing from its actions, the expected utility of asking the question is the expected improvement to its payoff (0.4×1) minus the cost of sending the message (0.1). Asking the question thus has a utility of 0.3. Let us note that the value of asking the question computed above coincides with the expected value of information obtained as a result, which is the usual notion of information value [34, 38], as we would expect.

We can summarize the above in the following:

Observation 2: Among selfish agents communicative acts corresponding to questions are acts declaring ignorance on the part of the speaker. Their value is closely related to the expected value of information considered in decision theory.

8 Statements of Propositional Attitudes

We take statements of propositional attitudes to be ones like "I know what is behind the trees," or "He knows what is in the jar." The semantics of propositional attitude statements have received much attention in the AI literature [21, 27]. Here, we would like to illustrate how the utility of statements of this kind can be computed within the framework we provide, given some straightforward postulates of what they intend to convey.

Recall the example scenario considered before (Figure 1) in which agent R_1 did not know whether the other agent, R_2, could see the observation point P2 through the trees. Let us assume that R_2 can, in fact, see the point P2 behind the trees. Is it valuable for R_2 to let R_1 know, for example by sending the message "I know what is behind the trees" over, say, a prefect communication channel? The answer can be arrived at by considering how R_2 could model the impact of such a communicative act, call it M_8, on its decision-making situation, as depicted in Figure 10.

The decision-theoretic pragmatics of M_8 is that it removes R_1's uncertainty as to whether R_2 knows what is behind the trees. Thus, it is a modeling act, as defined in Section 3. The solution of both of the recursive model structures in Figure 10 is straightforward. Before M_8 is sent, R_2 would expect R_1 to observe from P2, and R_2 would expect a benefit of 4. If M_8 is sent, on the other hand, R_1 would observe from P1, and R_2 could make the observation from P2 and obtain a benefit of 5. The value of the message to R_2 is $U(M_8) = 5 - 4 = 1$.

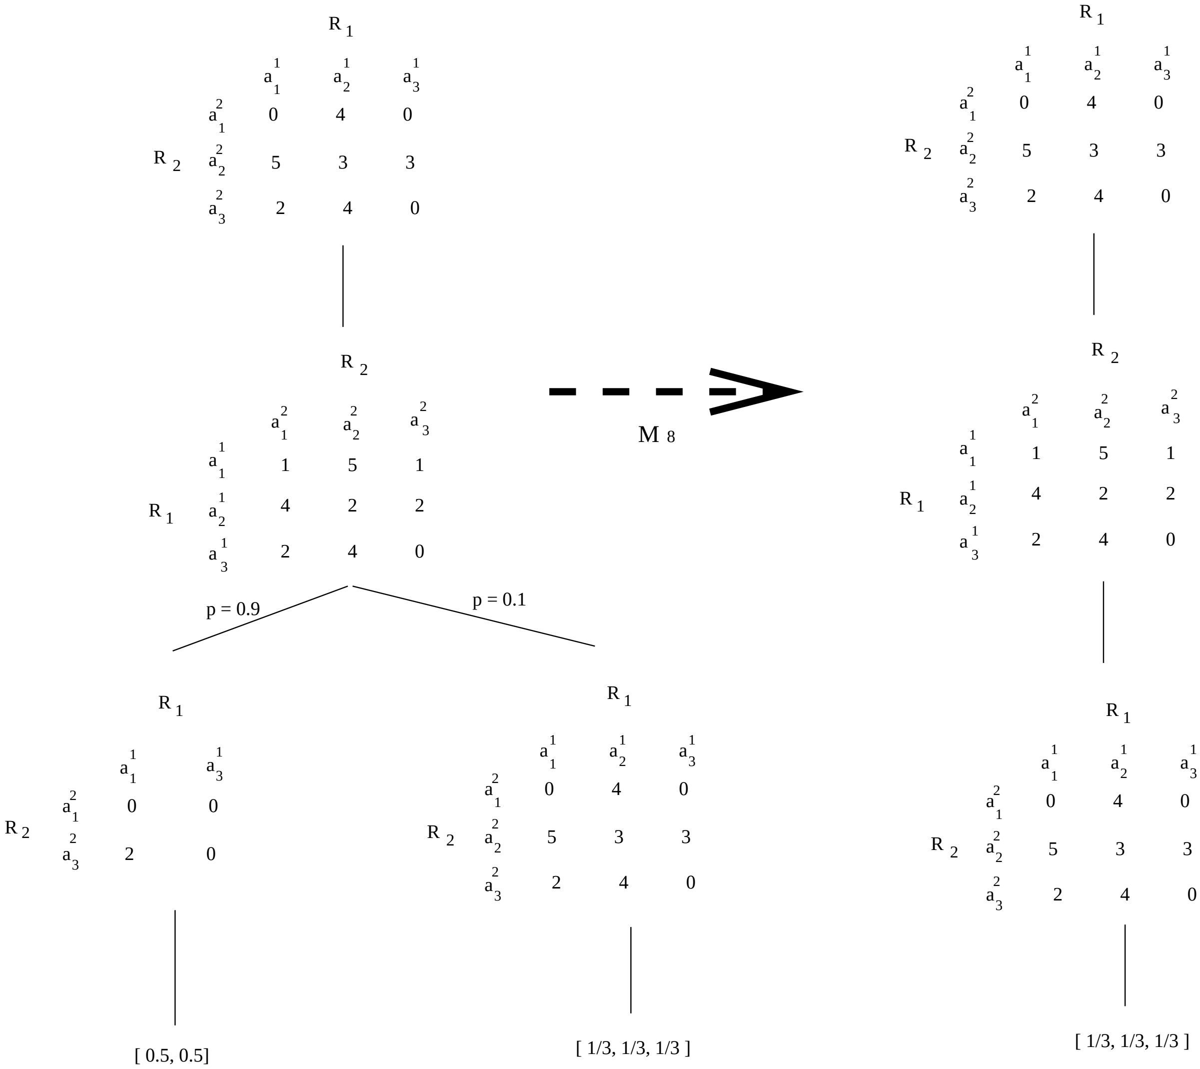

Figure 10: R_2**'s Model of DT Pragmatics of** M_8**.**

This case illustrates the possibility that the agents' states of knowledge are not consistent. It may be that R_1 believes that R_2 does not know anything about R_1, which may be known to R_2, and not true. Clearly, inconsistencies of this sort may happen frequently, and the agents have to be able to take them into consideration while effectively communicating.

9 Toward the Value of Imperatives

The consideration of imperative messages, for instance an order for R_2 issued by R_1, "Observe from P2!", raises issues similar to those considered while analyzing questions. Apart from non-myopic agents that interact repeatedly and agents that follow a cooperative interaction protocol dictating that orders be followed, there seems to be no reason why an autonomous agent should ever pay any attention to the orders given by others. In the majority of situations this indeed is the case, and persuasion should be more effective than issuing an order. By persuasion we mean informative statements, like the modeling message M_1 considered before. They are valuable since they inform an agent about relevant circumstances, and lead to desirable courses of action as a side effect.

However, there are circumstances in which a pure and unexplained order makes sense. These are circumstances in which one accounts for the costs (time, effort) of decision-making. In a nutshell, it is rational for an agent to obey an order from another agent if the default utility of obeying is greater than the utility of independent decision-making, including the cost of that decision-making. Intuitively, when a rational agent gets a message "Duck!" it might be better for it to rapidly follow this instruction than to respond more slowly (too late) after assessing the situation and deciding for itself on the proper action.

In its simplest form, an imperative can be postulated to transform an agent's recursive model structure from which it derives its strategy into simply a strategy to follow. That is, the agent discards its deliberative mechanisms and immediately obeys the command. Because the decision to do so involves a tradeoff of costs and benefits of using the decision-theoretic reasoning, an imperative causes reasoning at the meta-level. Since RMM has not yet addressed meta-level reasoning issues to any great extent, the decision as to when to follow the command cannot at this time be reduced to operational terms. Clearly, the deciding factors will involve comparing the expected payoff of following the command (what does the command imply about the likely decisions of others, and what kind of payoffs are likely to be received based on these implications) against the expected payoff of deeper reasoning (what better decisions might be made, how much better they might be, and how much it will cost to find them). Our ongoing research is delving into meta-level issues, so as to eventually capture such reasoning about imperatives.

10 Experiments

This section describes some of our experiments of coordination with communication in the air defense domain, in which two defense batteries have to coordinate their actions of intercepting multiple incoming threats. First, we show that RMM's decision-theoretic message selection in most cases agrees with selections chosen by human subjects in four simple defense scenarios. Then, we show results of scaled-up defense episodes in which RMM agents perform slightly better than the human subjects.

10.1 RMM vs. Human Message Selection in Simple Scenarios

In the simple scenarios below, we will consider optimal communicative behavior of Battery1 (triangle on the left in our scenarios) only, and assume that Battery2 is silent but can receive messages. Further, for simplicity, in all of the anti-air defense scenarios considered below Battery1 is assumed to have a choice of six communicative behaviors, generated by a communication planning module:

No Comm.: No communication
$M1$: I'll intercept Missile A.
$M2$: I'll intercept Missile B.
$M3$: I have both long and short range interceptors.
$M4$: There is a missile A, whose position and warhead size are P_A and W_A, respectively.
$M5$: There is a missile B, whose position and warhead size are P_B and W_B, respectively.

We wanted to investigate how RMM agents rank the messages in the above list, and whether there is an agreement between the communicative behavior advocated by RMM and human communicative behavior. As human subjects we used 32 CSE and EE graduate students. Each of them was presented with a scenario, and was given a description of what was known and what was uncertain in that scenario. The students were then asked to indicate which of the six messages was the most appropriate in each case, and which one was the second choice.[17]

10.1.1 Scenario 1

Consider the scenario depicted on the left in Figure 11. Here, the defense batteries face an

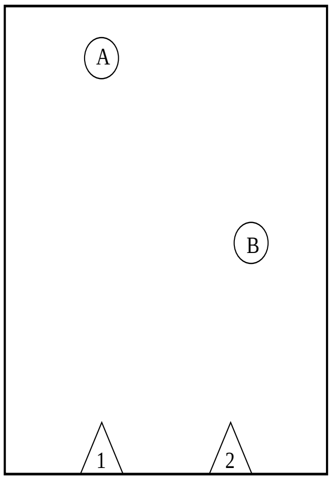

Missile	Position	Warhead
A	(7,3)	500
B	(14,10)	100

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	4	8
M1: I'll shoot down A.	7.31	24	5
M2: I'll shoot down B.	-1.26	3	10
M3: I have both ammo.	0	0	3
M4: A - (7,3), 500	0	1	1
M5: B - (14,10), 100	0	0	5

Figure 11: Scenario 1, and summary of results.

attack by missiles A and B. A has a larger warhead size than B, but it is farther from the defended territory. The state of Battery1's knowledge before communication is summarized as a two-level recursive model structure on the left in Figure 12. Assume that Battery1

[17] We would expect that anti-air specialists, equipped with a modern defense doctrine, could perform better than our subjects. However, the defense doctrine remains classified and was not available to us at this point.

assigns the probability of 0.9 to Battery2's being fully operational (having both long and short range interceptors and thus being able to target both missiles), and the probability of 0.05 to Battery2's being incapacitated (in which case it cannot do anything). The remaining probability of 0.05 is assigned to the No-Information model representing all of the possible remaining unknown cases. In this scenario Battery1 is assumed to have no more information. In particular, Battery1 does not know what action Battery2 expects of Battery1. This is represented by another No-Information model on the lowest level in Figure 12 (see [17] for detailed discussion).

Figure 12 depicts DT pragmatics of message $M1$, "I will intercept missile A." It illustrates

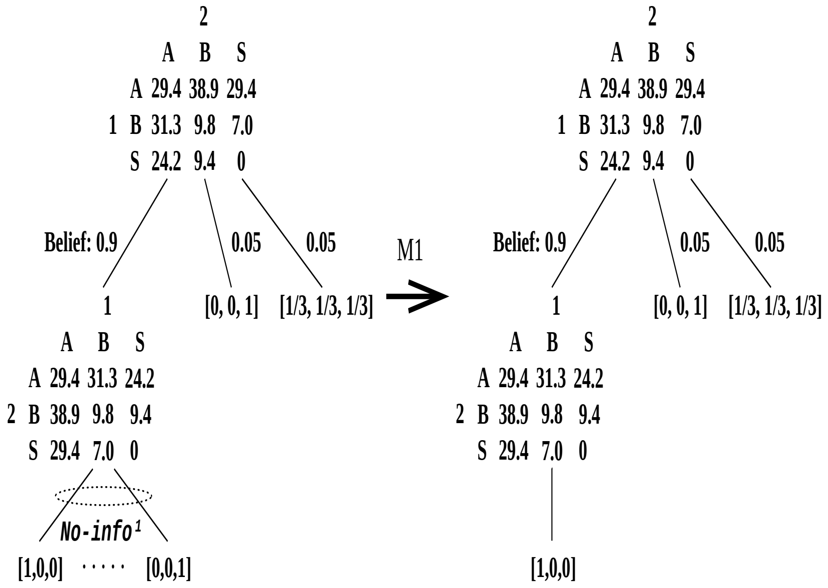

Figure 12: DT Pragmatics of message $M1$ in Scenario 1.

that, as a result, Battery1 expects Battery2 to know that Battery1 will select missile A, which is represented by $[1, 0, 0]$, probability distribution that Battery2 would use to describe Battery1's actions.

To compute the value of communication according to Equation 1, we solve both model structures in Figure 12 and compare results. Before communication (the left part of Figure 12), Battery1 computes that if Battery2 is operational then the probability distribution over Battery2's possible actions A, B, and S is $[0.85, 0.15, 0.0]$ (this result was obtained using logic sampling discussed in [17]). Using dynamic programming one can now easily compute that Battery1's best option is to shoot at missile A, with an expected utility $U_p(A)$ of $30.83 (= 0.78 \times 29.4 + 0.15 \times 38.9 + 0.07 \times 29.4)$.

After sending the message $M1$ (the right part of Figure 12), the probability distribution over Battery1's actions at Level 2 is $[1, 0, 0]$. Thus, if Battery2 is fully operational, it will choose to shoot at missile B, i.e., the probability distribution over Battery2's actions becomes [0,1,0]. This probability distribution is combined with the model of Battery2 being incapacitated and with the third No-Information model: $(0.9 \times [0, 1, 0] + 0.05 \times [0, 0, 1] + 0.05 \times [1/3, 1/3, 1/3]) =$

$[0.02, 0.92, 0.06]$. The resulting distribution is Battery1's overall expectation of Battery2's actions, given all of the remaining uncertainty. The combined probability distribution describing Battery2's actions is used to compute the expected utility of Battery1's action of shooting A. We have: $U_{p^{M1}}(A) = 0.02 \times 29.4 + 0.92 \times 38.9 + 0.06 \times 29.4 = 38.14$. According to Equation 1, the expected utility of the intentional communicative act $M1$, $U(M1)$, is $7.31 (= 38.14 - 30.83)$.

The expected utilities of the other messages are computed analogously, and the results are shown in Figure 11. As expected, some of the messages have no value in this situation, and their computed expected utility is zero, since they do not convey anything useful and novel. Note that message $M2$ has a negative expected utility; it is a bad idea for Battery1 to announce its intention to shoot at missile B in this scenario.

The results of human choices are also summarized in Figure 11. Twenty four, out of thirty two, i.e., 75% of the subjects chose message $M1$ as the best in this situation, while five subjects judged it as a second best. This shows a considerable agreement between RMM's calculations and selections of the human subjects.

10.1.2 Scenario 2

In this scenario, in Figure 13, the warhead of missile B is larger than that of A. However, due

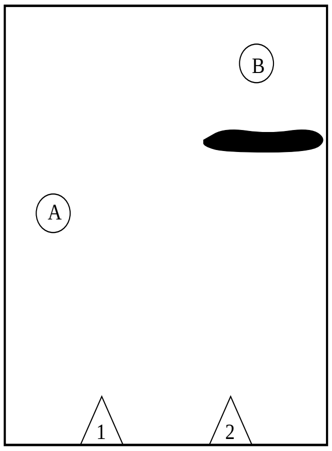

Missile	Position	Warhead
A	(4,9)	100
B	(16,3)	400

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	0	3
M1: I'll shoot down A.	0	1	7
M2: I'll shoot down B.	0.09	9	11
M3: I have both ammo.	0	0	3
M4: A - (4,9), 100	0	0	1
M5: B - (16,3), 400	3.64	22	7

Figure 13: Scenario 2, and summary of results.

to bad weather (rain or cloud), Battery1 thinks that Battery2 is unlikely to detect Missile B, and assigns to it the probability of only 0.1. Again, we allow for further uncertainty: Battery2 can be incapacitated by enemy fire with probability 0.05, and all of the other possibilities about which Battery1 has no information about Battery2 are grouped into a No-Information model with probability 0.05.

In this case, the most interesting message is a modeling message $M5$, "There is a missile B, with position (16, 3) and size 400." The result of this message would be to eliminate the possibility that Battery2 does not know about Missile B, and its DT pragmatics is depicted in Figure 14.

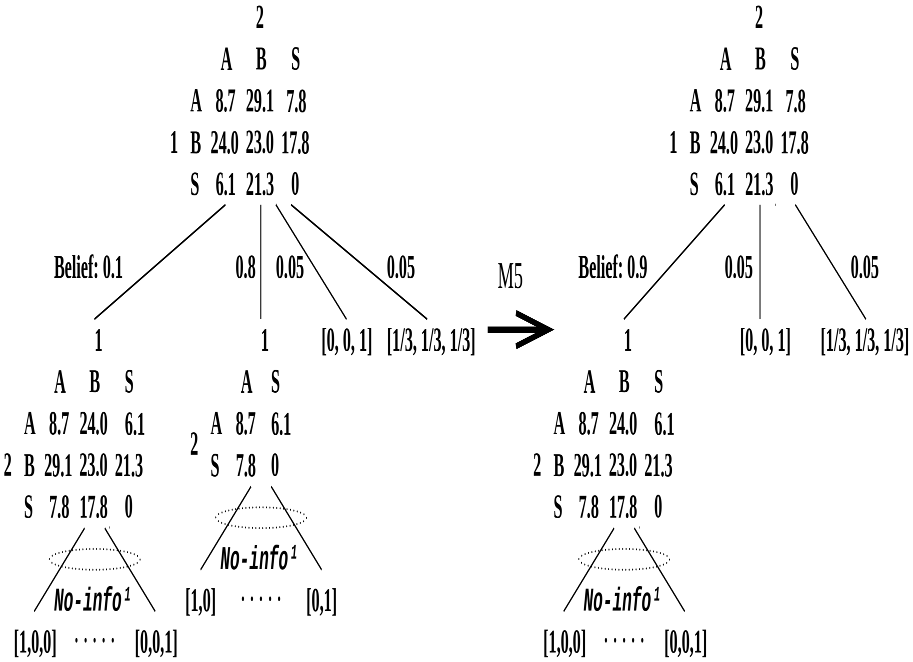

Figure 14: DT pragmatics of $M5$ in scenario 2.

The recursive model structures in Figure 14 can be solved using dynamic programming as before. Before communication, Battery1 would assign a distribution of $[0.82, 0.12, 0.06]$ to Battery2's actions, and Battery1 could get the expected utility of 23.42. After sending message $M5$, the probability that Battery2 will shoot at B increases to 0.9, the new distribution is $[0.04, 0.90, 0.06]$, and the expected utility of Battery1, $U_{p^{M5}}(A)$, is: $U_{p^{M5}}(A) = 0.04 \times 8.7 + 0.90 \times 29.1 + 0.06 \times 7.8 = 27.06$. Therefore, the expected utility of the communicative act $M5$, $U(M5)$, is $3.64 (= 27.06 - 23.42)$.

The expected utilities of the other communicative alternatives are computed similarly and the results are in Figure 13. The results of experiments with human subjects are also summarized. This time, 22 (69%) out of 32 subjects chose message $M5$ as the best one. 9 of our subjects picked message $M2$ as the best, which the RMM agent rated as second best in this case. We think that high number of students choosing message $M2$ is due to the subjects not thinking it would be possible for Battery2 to shoot at missile B just based on knowing its position (without visual contact), which we allowed for during the experiments between RMM agents.

10.1.3 Scenario 3

If we allow more uncertainty in a defense scenario, the decision-theoretic message selection will be more complicated. For scenario 3, (see Figure 15), we assumed that Battery1 is uncertain whether Battery2 has any short range interceptors left. If Battery2 has only long range interceptors, it will be unable to attack missile B, and can only attempt to shoot down missile A.

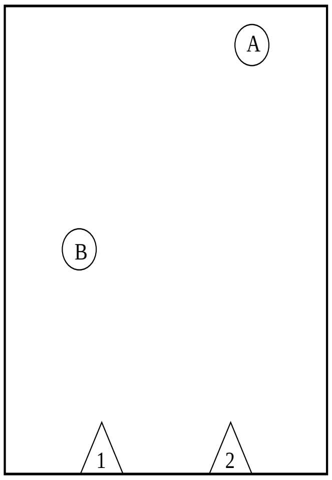

Missile	Position	Warhead
A	(14,2)	120
B	(6,10)	100

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	1	3
M1: I'll shoot down A.	-1.47	0	3
M2: I'll shoot down B.	2.63	27	3
M3: I have both ammo.	2.02	3	19
M4: A - (14,2), 120	0	1	4
M5: B - (6,10), 100	0	0	0

Figure 15: Scenario 3, and summary of results.

From Battery1's point of view, therefore, Battery2's decision making situation is modeled as one of four cases: Battery2 has both short and long range interceptors; Battery2 has only long range interceptors; Battery2 has been damaged or incapacitated; and a No-Information model.

In this scenario DT pragmatics of $M2$, "I will intercept missile B," is depicted in Figure 16. Another viable message is message $M3$, with DT pragmatics as depicted in Figure 17. This transformation shows that if Battery1 sends the message "I have both long and short range interceptors," Battery2 will include the fact that Battery1 has both munitions in its modeling of Battery1 on the third level of the model structure, and will solve these models to arrive at Battery1's target selection.

Before communication, Battery1's best option is to shoot down missile B, with an expected utility of 12.76. If $M1$ were to be sent, Battery1 could expect a utility of 11.29, which results in $M1$'s value being negative. After sending $M2$ Battery1 could expect $15.39(= 0.92 \times 15.9 + 0.02 \times 9.8 + 0.06 \times 9.4)$, while $M3$ would result in an expected utility of $14.78(= 0.82 \times 15.9 + 0.11 \times 9.8 + 0.07 \times 9.4)$.

These results, together with selections of the human subjects, are summarized in Figure 15. Twenty seven out of our thirty two subjects (84%) agreed with the RMM calculation and

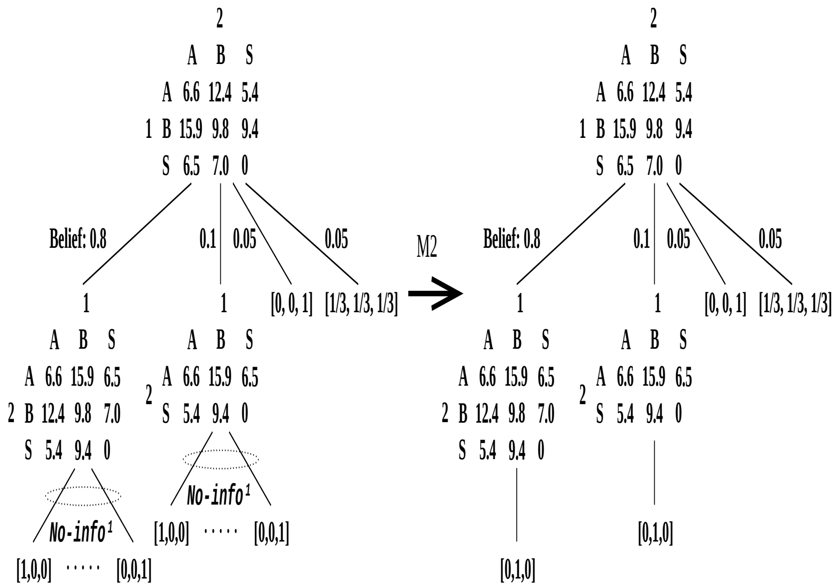

Figure 16: DT pragmatics of $M2$ in scenario 3.

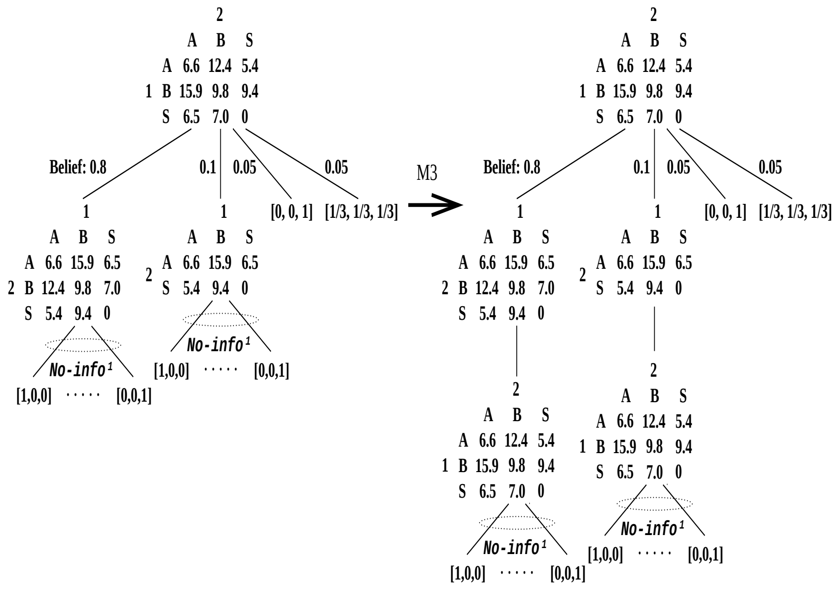

Figure 17: DT pragmatics of $M3$ in scenario 3.

picked $M2$ as the best message in this case. Further, 19 subjects rated $M3$ as their second choice.

10.1.4 Scenario 4

Our fourth scenario was intended as a test of a case when communication is not necessary, for example, because even without communication it is clear what the agents should do. In general, since communication may be expensive, a rational agent should not communicate if there is no gain from communication. In scenario 4 we set the warhead size of missile B as bigger than that of missile A. The altitudes of missiles are the same, as depicted in Figure 18. Also, Battery1 assigns the probability 0.05 to Battery2 being incapacitated by enemy fire,

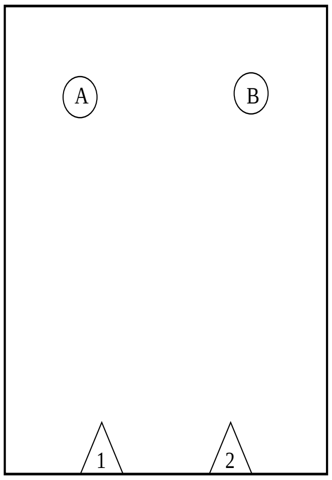

Missile	Position	Warhead
A	(6,5)	100
B	(14,5)	500

Messages	RMM U(M)	Human 1st	Human 2nd
No Communication	0	11	10
M1: I'll shoot down A.	0	19	9
M2: I'll shoot down B.	-3.31	2	5
M3: I have both ammo.	0	0	1
M4: A - (6,5), 100	0	0	2
M5: B - (14,5), 500	0	0	5

Figure 18: Scenario 4, and summary of results.

with the No-Information model having the probability of 0.05. Thus, Battery1 believes that Battery2 has both interceptors and can see Missile A and B with probability 0.9.

For this scenario, the expected utilities of intentional messages are less than or equal to the expected utility of no communication, as depicted in Figure 18. The modeling messages are useless as well, since they do not change the decision-making situation in any beneficial way. The results of experiments with the human subjects suggests that the majority of them were indifferent between communicating and not communicating in this case. There is, however, a slight preference the human subjects seem to exhibit for communicating useless messages. This can be accounted for by two factors: First it was probably difficult for our subjects to understand that the values of messages that make perfect sense under usual conditions are useless in this particular situation; and second, human communicative behavior may be driven not strictly by the values of messages we have been computing here. People may be used to talking just to stay in touch and maintain a social bond, even when the situation does not demand it.

10.2 Experiments in Scaled-Up Defense Episodes

In the scaled-up experiments we ran, two defense units were faced with an attack by seven incoming missiles. Therefore, the RMM agents used an 8×8 payoff matrix to represent the agents' decision-making situations. For all settings, each defense unit was initially assumed to have the following uncertainties (beliefs) in its knowledge base: the other battery is fully functional and has both long and short range interceptors with probability 60%; the other battery is operational and has only long range interceptors with probability 20% (in this case, it can shoot down only distant missiles, which are higher than a specific altitude.); the other battery has been incapacitated by enemy fire with probability 10%, and the remaining 10% was assigned to the No-Information model. Further, each battery has no deeper nested knowledge about the other agent's beliefs.

The warhead sizes of missiles were $470, 410, 350, 370, 420, 450$, and 430 unit for missiles A through G, respectively. In these experiments each of the two defense units was assumed to be equipped with three interceptors, if they were not incapacitated. Thus, they could launch one interceptor at a time, and did it three times during a course of one defense episode.

We set up 100 scenarios for RMM team and 20 scenarios for human team. We allowed for one-way communication between defense units before each salvo. If both agents wanted to send messages, the speaker was randomly picked in the RMM team, and the human team flipped a coin to determine who would be allowed to talk. The listener was silent and could only receive messages. For uniformity, in all of the anti-air defense scenarios, each battery was assumed to have a choice of the following communicative behaviors: "No communication," "I'll intercept Missile A." through "I'll intercept Missile G.", "I have both long and short range interceptors.", "I have only long range interceptors.", or "I'm incapacitated.". As human subjects, we used 20 CSE and EE graduate students.

Among the various episodes we ran we will consider two illustrative examples to examine the coordination achieved by RMM and the human team in more detail. In these examples, each defense unit was fully functional and has both long and short range interceptors.

Table 1: Performance analysis for scenario 1.

Scenario 1	No Communication	
Agent	RMM	Human
Targets	F,A;E,B;D,C	D,B;E,A;F,C
Total damage	721.8	731.2
Scenario 1	**Communication**	
Agent	RMM	Human
Message 1	has(B2,both_ammo)	intercept(B2,A)
Message 2	intercept(B2,B)	intercept(B2,C)
Message 3	intercept(B1,E)	intercept(B1,G)
Targets	F,A;G,B;E,C	F,A;E,C;G,B
Total damage	630.0	647.2

The result of interaction in scenario 1, in Figure 19 is presented in Table 1. Without communication, the RMM batteries 1 and 2 shot at threat F and A, respectively, during the first salvo; at E and B, respectively, during the second salvo; and at D and C, respectively, during the third and final salvo, as depicted in Table 1. The total damage sustained by the RMM team in this encounter was 721.8. The choices made by a human team without

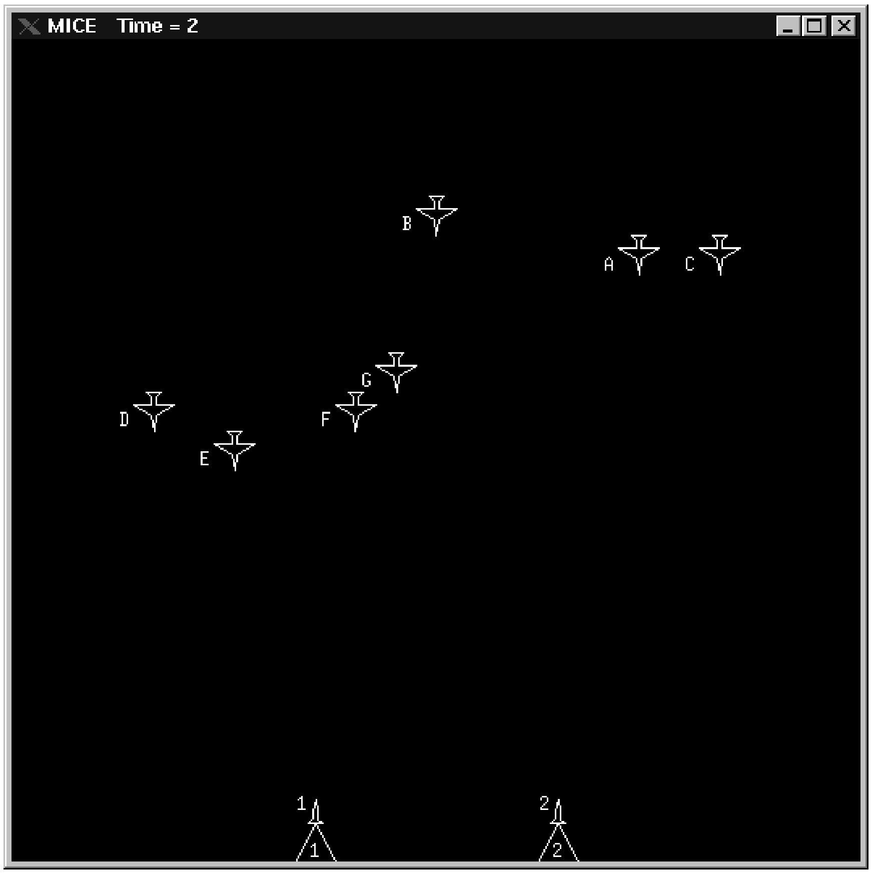

Figure 19: Defense Scenario 1.

communication is similarly displayed in the upper right corner of Table 1; the damage suffered by the human team was 731.2.

The lower portion of Table 1 illustrates what happened when the agents could exchange messages. Before the first salvo Battery2 was allowed to talk. The RMM agent in charge of Battery2 sent the message was "I have both long and short range interceptors," and shot at target A. Upon receiving the message, Battery1 controlled by RMM intercepted F. In case of the human team, Battery2's initial message was "I will intercept target A", and the human team also shot at targets F and A during the first salvo. The messages exchanged and the firings in the following two salvos are also shown in the lower portion of Table 1. As expected, the performance with communication was better than one without communication for both teams; the RMM suffered damages of 630.0, while the human team scored 647.2.

The difference of total damage in RMM and human teams with and without communication shows the benefit of communication. In this scenario, the expected utilities of communicative acts executed by the RMM team were $U(has(B2, both\_ammo)) = 18.04$, $U(intercept(B2, B))$ $= 41.05$, and $U(intercept(B1, E)) = 32.49$, which sums up to 92.03. This amount is closely related to the benefit of the communication, i.e., $91.8 (= 721.8 - 630.0)$. This shows that, in this scenario, the expected utilities of the messages transmitted were an adequate estimate of the benefits of communication actually realized. As the results show, the human team's score was very similar to that of the RMM team, but humans chose different communicative acts. For example, the message chosen by the human player before the first salvo "I will intercept target A" has the expected utility of 12.17. This message is, therefore, useful, although slightly suboptimal from the decision-theoretic point of view.

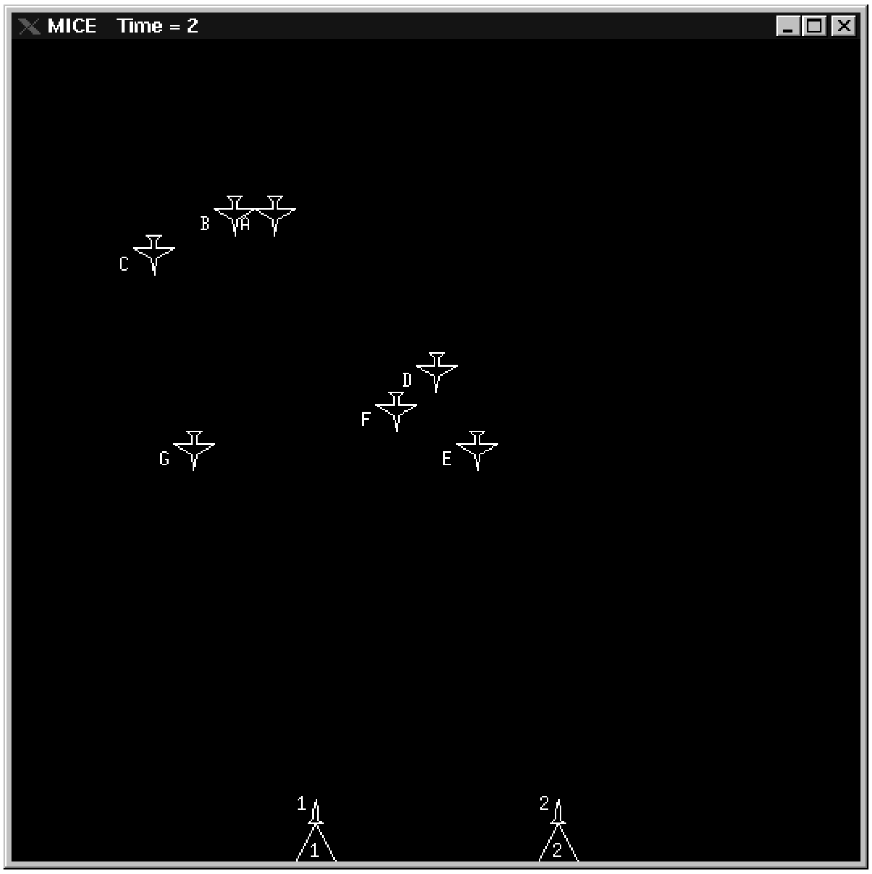

Figure 20: Defense Scenario 2.

In scenario 2, in Figure 20, the targets are clustered in front of Battery1, unlike in scenario 1, in which targets are scattered. In this case, communication is more important and the utility of the messages are greater because it is more likely that targets could be intercepted redundantly without communication, resulting in greater overall damage incurred. Intuitively, as most targets head toward Battery1, the target that Battery2 selects as the biggest threat is likely to be also the most threatening to Battery1. As shown in Table 2, when communication is available, redundant target selection was prevented and the total expected damages were drastically reduced. In this scenario the sum of the the expected utilities of the three messages sent by the RMM team is 246.57, while the benefits of communication actually obtained is 356.4.

Table 2: Performance analysis for scenario 2.

Scenario 2	No Communication	
Agent	RMM	Human
Targets	F,E;B,D;G,G	A,E;F,F;G,D
Total damage	1020.3	1001.8
Scenario 2	**Communication**	
Agent	RMM	Human
Message 1	intercept(B2,E)	has(B2,both_ammo)
Message 2	intercept(B1,G)	intercept(B1,B)
Message 3	intercept(B1,B)	intercept(B1,G)
Targets	A,E;G,F;B,D	A,E;B,F;G,D
Total damage	663.9	708.4

The summary of all of the experimental runs we performed is shown in Table 3. Table 3 presents the average number of selected targets and the average total expected damage by RMM agents after 100 trials and by human agents after 20 trials. We focus on the performances of three different teams: RMM-RMM, RMM-human, and human-human team. To see whether the differences in results obtained by the teams are not due to chance, we performed an analysis of variance (ANOVA) on the basis of total expected damage. For this purpose we randomly chose 20 trials of the RMM team among the 100 trials available.

In Case I of Table 3, both defense units are fully functional and there is no communication. Since the computed value $f = 3.45$ in ANOVA exceeds $3.15 (= f_{.05,2,57})$, we know that the three teams are not all equally effective at the 0.05 level of significance, i.e., the differences in their performance are not due to chance with probability 0.95. Results in Case II include communication between fully functional defense units. ANOVA test shows that the differences in performance are not due to chance with probability 0.95 again. When the communication was available, the performances achieved by three teams were improved, with the RMM team performing slightly better than the other teams.

The experiments in Case III and IV intended to examine how fully functional battery can cope with the situation in which the other battery has only long range interceptors. Like Case I and II, the experimental results of Case III and IV proved the effectiveness of communication. The ANOVA test indicates that the observed differences in the performance of four teams for the target selection are significant at the 0.01 level of significance, and are not due to chance with probability 99%. In Case V, the battery controlled by human or RMM cannot take advantage of communication, because the other battery is incapacitated or behaves randomly. The huge value of $f = 122.01$ tells that the performances of four different teams are clearly distinct.

Table 3: The performances of RMM and human team

Cases	Team (B1-B2)	No of selected targets	Total expected damage	ANOVA
Case I (B2:both ammo, w/o comm.)	RMM-RMM	5.95 ± 0.21	717.01 ± 110.71	f=3.45
	RMM-Human	5.70 ± 0.47	797.39 ± 188.35	
	Human-Human	5.75 ± 0.44	800.45 ± 147.69	
Case II (B2:both ammo, w/ comm.)	RMM-RMM	6.00 ± 0.00	652.33 ± 58.97	f=3.43
	RMM-Human	6.00 ± 0.00	717.93 ± 94.45	
	Human-Human	6.00 ± 0.00	710.20 ± 100.92	
Case III (B2:only long, w/o comm.)	RMM-RMM	5.83 ± 0.37	852.01 ± 160.79	f=6.96
	RMM-Human	5.75 ± 0.44	862.70 ± 120.19	
	Human-RMM	5.40 ± 0.50	895.92 ± 127.32	
	Human-Human	5.30 ± 0.47	997.32 ± 145.15	
Case IV (B2:only long, w/ comm.)	RMM-RMM	5.88 ± 0.32	787.42 ± 110.26	f=4.58
	RMM-Human	5.85 ± 0.36	842.50 ± 131.26	
	Human-RMM	5.75 ± 0.44	815.67 ± 133.60	
	Human-Human	5.80 ± 0.41	908.08 ± 103.25	
Case V (B2:incap. or random)	RMM-Incap.	3.00 ± 0.00	1742.22 ± 64.45	f=122.01
	Human-Incap.	3.00 ± 0.00	1786.86 ± 87.94	
	RMM-Random	4.86 ± 0.58	1079.52 ± 210.64	
	Human-Random	4.85 ± 0.59	1115.57 ± 228.94	

Note) For all of cases, Battery1 is fully functional.
$f_{.05,2,57} = 3.15$, $f_{.01,3,76} = 4.13$

11 Related Work

As we mentioned, our approach follows the tradition of cognitive science [5, 11, 30], which postulates that the function of communication is to confer some advantage to the speaker by influencing what the hearer knows and intends to do. For instance, MacLennan ([30], page 636) states that "If we want genuine meaning and original intentionality, then communication must have real relevance to the communicators." Further, cognitive scientists were able to confirm the role and importance of models, including nested models, of other agents used for effective communication, and how the ability to form and process these models sets humans apart from other primates (see [11] and references therein). For example, adult humans can reliably (with 5-15% error rates) reason about models of others nested up to four levels deep. Monkeys (and some autistics) lack these abilities, which manifests itself in much poorer communicative abilities.

The issue of modeling agents for the purpose of communication has also received considerable attention in the field of human-computer interaction [43]. Fischer [12], for example, stresses the importance of a system's model of the user and the user's model of the system during man-machine communication. This leads to the nesting of models represented explicitly by our framework. Uncertainty-based decision-theoretic frameworks for HCI have been used by Jameson and associates in [25, 26].

In earlier work in AI, Cohen and Levesque provide a logical formalization of the concepts of intention and commitment [7], and apply it to issues of communication [8]. In this formalization, the concept of rationality takes a prominent place but, as the authors remark ([7], page 40), it is not the formalization that the agents should actually use as a guide to their action. Rather, it is an attempt to formally describe a rational agent's behavior. We agree with this view. For example, the fact that a rational agent will not attempt to achieve goals that it believes are already true (in either action or communication) is just a manifestation of the expected utility of such attempts to a rational agent being zero.[18] Cohen and Levesque, as well as Perrault in [35], also 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 the discussion in [17]). In a similar vein is the work of Grosz and Sidner [20], Pollack [37], and Meyden [44]. As we mentioned before, our research differs, but also complements this previous work. We concentrate on the notion of decision-theoretic rationality as a normative paradigm in communicative behavior, and define the decision-theoretic version of pragmatics of communicative acts. This enables the agents to rationally choose which communicative behavior should be executed in a situation at hand.

Other work in AI include efforts on semantics of KQML [29, 41], which is closely related to earlier work of Cohen and Levesque, but it does not include the notion of value central to our approach. While work on communication in negotiation is reported in [28, 48, 49].

Recently, Tambe [42] suggested decision-theoretic communication selectivity to establish mutual belief among agents in a team. This approach is similar to ours in that the focus is on whether or not an agent should transmit a given message (fact) to others. Tambe uses a decision tree containing reward and penalty values that are domain knowledge. However, obtaining these values when the environment is dynamic, not fully observable or when the status of another agent is not fully predictable, is problematic. Our framework, in contrast, does not rely on pre-determined rewards and penalties, but computes them based on possibly incomplete and uncertain models of other agents. Tambe's work is further closely related to work by Cohen [8] and Grosz [20] mentioned above.

Communication among rational agents has also been of interest for the researchers in game theory, usually viewed a part of pre-game "cheap talk" [9], or as threat games [32] (see also discussion in [13] and references therein). These approaches, however, concentrate on the influence of pre-play communication on equilibria of the resulting game. Thus, it assumes common knowledge and takes the global view of the multi-agent system, as opposed to our taking an agent-centered and decision-theoretic view. Additionally, this work does not explicitly represent the state of knowledge of the agent, and does not represent the change of this state due to a communicative act, which is central to our approach. A qualitative approach, more closely related to ours and including dishonesty, has been presented by Myerson in [31]. In a similar vein, Parikh [33] used game-theoretic insights for disambiguation.

Other relevant work includes value of information approaches by Horvitz and associates [36], building on earlier work by Howard [24]. Horvitz's approach is related in that the decision-making situation of the agent includes uncertainty, and is represented as an influence diagram.

[18] See [45] for arguments showing how the notion of utility generalizes the notion of a goal.

These diagrams contain an agent's alternative actions, utility, and states that influence utility, which are the elements represented by the payoff matrices that we use. The main difference is that, as we mentioned, the value of information from the hearer's point of view, which the above work is computing, need not correspond to the value of transmitting the message from the speaker's point of view, which is what we compute. In other words, we concentrate on the rational communicative acts executed by a speaker, while the value-of-information approach concentrates on an agent's active search for information.

12 Conclusions

We address the issue of rational communicative behavior among autonomous intelligent agents that make decisions as to what to communicate, to whom, and how. We treat communicative actions as aimed at increasing the efficiency of interaction among agents. We postulate that a rational speaker design a speech act so as to maximally increase the benefit obtained as the result of the interaction. We quantify the gain in the quality of interaction as the expected utility, and we present a framework that allows an agent to compute the expected utility of various communicative actions. Our approach uses the Recursive Modeling Method as a convenient compilation of available information residing in the knowledge base pertaining to the agent's decision-making situation. We then define the decision-theoretic pragmatics of a speech act as the transformation it induces on the agent's decision-making situation. This transformation leads to a change in the quality of the interaction, expressed in terms of the benefit to the agent. We analyze the decision-theoretic pragmatics of a number of important communicative acts, and investigate their expected utilities using examples. Our approach to computing the expected utilities of communicative acts accommodates and quantifies the realistic possibility that the agents do not share a communication language, and the possibility that the communication channel is unreliable.

We considered a number of types of speech acts and evaluated their value for the agents involved using examples. Included were intentional, acknowledging and modeling messages, and, in a preliminary form, questions, messages stating propositional attitudes, and imperatives. It turns out that in the society of purely autonomous agents intent on maximizing their own benefit, questions may best be implemented as declarations of ignorance, and imperatives make sense only when individual decision-making is both redundant and costly.

We have conducted a number of experiments on the air defense domain and our results validate the reasonableness of our approach. The expected utilities of messages we have considered largely coincide with human assessment of what messages are most appropriate, which indicates that our method is psychologically plausible. We verified the usefulness of our method in numerous scenarios by measuring the increase in the quality of interaction that has been achieved due to rational communicative behavior. We compared the performance of our automated agents with that of communicating humans and showed that our agents are competent in their interaction and communication, frequently outperforming humans in our experimental domain.

In our future work we will investigate techniques that can be used to compile the results of full-blown RMM method into situation/communication pairs, to be used to urgent situations. This naturally gives rise to the establishment of protocols. The combinatorics of the search for the best communicative act we consider here can be greatly reduced by restricting the types of messages available and by clustering collections of communications together into "standard" interactions (question-answer, or announce-bid-award, for example). As clusters grow and are reapplied to different (but similar) situations, they become stored plans of communication resembling protocols. Thus, while substantial research is needed to adequately operationalize this process of going from computing utilities of individual actions based on transforming a nested modeling structure, all the way up to representing generalized protocols for guiding dialogues, this research path looks both important (to build systems that can survive and establish new protocols based on first-principle methods when known protocols fail) and promising (since we have, in this paper, described how to derive primitives for such protocols, such as query-response, within this framework).

Another strand of our future work will be devoted to the issues of emerging cooperativeness of communicating agents that interact repetitively. Previous work in game theory [1], as well as experimental work [3] shows how cooperation emerges among selfish repeatedly interacting utility maximizers. Our work will be aimed at the emergence of cooperative communicative behavior, so closely related to the status of questions, requests and orders in the human society.

Acknowledgments

We would like to gratefully acknowledge the many comments and help of Professor Jeffrey S. Rosenschein form the Department of Computer Science at the Hebrew University of Jerusalem, Israel. We also acknowledge the help of our student, Sanguk Noh from the University of Texas at Arlington's CSE department, for his invaluable help in implementing and experimenting with RMM.

References

[1] Robert J. Aumann and Sergiu Hart. Repetition as a paradigm for cooperation in games of incomplete information. Unpublished technical report, Hebrew University, Israel, 1981.

[2] J. L. Austin. *How to do Things with Words*. Clarendon Press, 1962.

[3] Robert Axelrod. *The Evolution of Cooperation*. Basic Books, 1984.

[4] Afzal Ballim and Yoric Wilks. *Artificial Believers*. Earlbaum Associates, Inc., 1991.

[5] Gordon M. Burghardt. Ontogeny of communication. In Thomas A. Sebeok, editor, *How Animals Communicate*, pages 71–97. Indiane University Press, 1977.

[6] Herbert H. Clark. *Arenas of Language Use*. The University of Chicago Press, Chicago, 1992.

[7] P. R. Cohen and H. J. Levesque. Persistence , intention and commitment. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication.* MIT Press, 1990.

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

[9] V. P. Crawford and J. Sobel. Strategic information transmission. *Econometrica*, 50(6):1431–1452, 1982.

[10] D. Dennett. Intentional systems. In D. Dennett, editor, *Brainstorms.* MIT Press, 1986.

[11] Robin Dunbar. Theory of mind and the evolution of language. In James R. Hurford, Michael Studdert-Kennedy, and Chris Knight, editors, *Approaches to the Evolution of Language*, pages 92–110. Cambridge University Press, 1998.

[12] Gerhard Fischer. The imortance of models in making complex systems comprehensible. In M. J. Tauber and D. Ackerman, editors, *Mental Models and Human-Computer Interaction.* North Holland, 1991.

[13] Drew Fudenberg and Jean Tirole. *Game Theory.* MIT Press, 1991.

[14] Piotr Gmytrasiewicz. *A Decision-Theoretic Model of Coordination and Communication in Autonomous Systems.* PhD thesis, University of Michigan, December 1991.

[15] Piotr J. Gmytrasiewicz and Edmund H. Durfee. Toward a theory of honesty and trust among communicating autonomous agents. *Group Decision and Negotiation*, 2:237–258, 1993.

[16] Piotr J. Gmytrasiewicz and Edmund H. Durfee. A rigorous, operational formalization of recursive modeling. In *Proceedings of the First International Conference on Multiagent Systems, ICMAS'95*, pages 125–132, July 1995.

[17] Piotr J. Gmytrasiewicz and Edmund H. Durfee. Rational coordination in multi-agent environments. *Accepted for publication, AAMAS Journal, 1999*, 1999. Available in postscript from http://www-cse.uta.edu/ piotr/piotr.html.

[18] H. P. Grice. Meaning. *Philosophical Review*, (LXVI):377–388, 1957.

[19] H. P. Grice. Utterer's meaning, sentence-meaning and word meaning. *Foundations of Language*, (4):225–245, 1968.

[20] B. J. Grosz and C. Sidner. Plans for discourse. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication.* MIT Press, 1990.

[21] Joseph Y. Halpern and Yoram Moses. A guide to the modal logics of knowledge and belief. Technical Report 74007, IBM Corporation, Almaden Research Center, 1990.

[22] J. Harsanyi. Bayesian decision theory and utilitarian ethics. *American Economic Review*, (68):223–228, 1978.

[23] David Heckerman and David Maxwell Chickering. A comparison of scientific and engineering criteria for bayesian model selection. Technical Report MSR-TR-96-12, Microsoft Research, Microsoft Corporation, Redmond, WA, 1996.

[24] R. A. Howard. Information value theory. *IEEE Transactions on Systems Science and Cybernetics*, 2:22–26, 1966.

[25] Anthony Jameson. But what will the listener think? belief ascription and image maintenance in dialog. In A. Kobsa and W. Wahlster, editors, *User Models in Dialog Systems*. Springer-Verlag, 1989.

[26] Anthony Jameson, Ralph Shaefer, Joep Simons, and Thomas Weis. Adaptive provision of evaluation-oriented information: Tasks and techniques. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, pages 1886–1893, August 1995.

[27] Kurt Konolige. *A Deduction Model of Belief*. Morgan Kaufmann, 1986.

[28] Sarit Kraus and Jonathan Wilkenfeld. The function of time in cooperative negotiations. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, August 1991.

[29] Y. Labrou and T. Finin. A semantics approach for KQML – a general purpose communication language for software agents. In *Proceedings of the Third International Conference on Information and Knowledge Management*, November 1994.

[30] Bruce MacLennan. Synthetic ethology: An approach to the study of communication. In C. G. Langton, C. Taylor, J. D. Farmer, and S. Rasmussen, editors, *Artificial Life II, SFI Studies in the Sciences of Complexity*, pages 631–658. Addison-Wesley, 1991.

[31] Roger B. Myerson. Incentive constraints and optimal communication systems. In *Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge*, pages 179–193, March 1988.

[32] Roger B. Myerson. *Game Theory: Analysis of Conflict*. Harvard University Press, 1991.

[33] Prashant Parikh. A game-theoretic account of implicature. In *Proceedings of the Conference on Theoretical Aspects of Reasoning about Knowladge*, pages 85–93. Morgan Kaufman, August 1992.

[34] Judea Pearl. *Probabilistic Reasoning in Intelligent Systems: Networks of Plausible Inference*. Morgan Kaufman, 1988.

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

[36] Kim L. Poh and Eric J. Horvitz. A graph-theoretic analysis of information value. In *Proceedings of the Twelfth Conference on Uncertainty in Artificial Intelligence (UAI-96)*, pages 427–435, August 1996.

[37] Martha Pollack. Plans as complex mental attitudes. In P. R. Cohen, J. Morgan, and M. E. Pollack, editors, *Intentions in Communication*. MIT Press, 1990.

[38] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*. Prentice Hall, 1995.

[39] Stephen Schiffer. *Meaning*. Clarendon Press, 1972.

[40] Munindar Singh. Towards a formal theory of communication for multiagent systems. In *Proceedings of the Twelfth International Joint Conference on Artificial Intelligence*, August 1991.

[41] Ira A. Smith and Philip R. Cohen. Toward a semantics for an agent communications language based on speech acts. In *Proceedings of the National Conference on Artificial Intelligence*, pages 24–31, Portland, OR, August 1996.

[42] Milind Tambe. Agent architectures for flexible, practical teamwork. In *Proceedings of the Fourteenth National Conference on Artificial Intelligence*, pages 22 – 28, 1997.

[43] M. J. Tauber and D. Ackerman. *Mental Models and Human-Computer Interaction 2*. North Holland, 1991.

[44] Ron van der Meyden. Mutual belief revision (preliminary report). In *Proceedings of the International Conference on Principles of Knowledge Representation and Reasoning*, pages 595–606, May 1994.

[45] Michael P. Wellman. The preferential semantics for goals. In *Proceedings of the National Conference on Artificial Intelligence*, pages 698–703, July 1991.

[46] Terry Winograd and Fernando Flores. *Understanding Computers and Cognition: A New Foundation for Design*. Ablex Publishing, 1986.

[47] Robert Worden. The evolution of language from social intelligence. In James R. Hurford, Michael Studdert-Kennedy, and Chris Knight, editors, *Approaches to the Evolution of Language*, pages 148–166. Cambridge University Press, 1998.

[48] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and task sharing among autonomous agents in cooperative domains. In *Proceedings of the Eleventh International Joint Conference on Artificial Intelligence*, pages 912–917, August 1989.

[49] Gilad Zlotkin and Jeffrey S. Rosenschein. Negotiation and conflict resolution in non-cooperative domains. In *Proceedings of the National Conference on Artificial Intelligence*, pages 100–105, July 1990.