

Report 84-44
Stanford -- KSL

Scientific DataLink

Deals Among Rational Agents.
Jeffrey S. Rosenschein, Michael R. Genesereth,
Dec 1984

card 1 of 1

Stanford Heuristic Programming Project
Report No. HPP-84-44

December 1984

Deals Among Rational Agents

Jeffrey S. Rosenschein
Michael R. Genesereth

COMPUTER SCIENCE DEPARTMENT
Stanford University
Stanford, California 94305

Deals Among Rational Agents

Abstract

A formal framework is presented that models communication and promises in multi-agent interactions. This framework generalizes previous work on cooperation without communication, and shows the ability of communication to resolve conflicts among agents having disparate goals. Using a deal-making mechanism, agents are able to coordinate and cooperate more easily than in the communication-free model. In addition, there are certain types of interactions where communication makes possible mutually beneficial activity that is otherwise impossible to coordinate.

§1. Introduction

1.1 Artificial Intelligence and the Multi-Agent Paradigm

Research in artificial intelligence has focused for many years on the problem of a single intelligent agent. This agent, usually operating in a relatively static domain, was designed to plan, navigate, or solve problems under certain simplifying assumptions, most notable of which was the absence of other intelligent entities.

The presence of multiple agents, however, is an unavoidable condition of the real world. People must plan actions taking into account the potential actions of others, which might be a help or a hindrance to their own activities. In order to reason about others' actions, a person must be able to model their beliefs and desires.

The artificial intelligence community has only lately come to address the problems inherent in multi-agent activity. A community of researchers, working on distributed artificial intelligence (DAI), has arisen. Even as they have begun their work, however, these researchers have added on a new set of simplifying assumptions that severely restrict the applicability of their results.

1.2 Benevolent Agents

Virtually all researchers in DAI have assumed that the intelligent agents in their domain have identical or non-conflicting goals. Work has thus proceeded on the question of how these agents can best help one another in carrying out their common tasks [3, 4, 6, 7, 24], or how they can avoid interference while using common resources [10, 11]. The

rationale for studying multiple agent interaction stems from a desire for increased system efficiency or increased capabilities. For example, it is hoped that if a group of agents carry out a task cooperatively, the task will take less time than if it were performed by a single agent.

Of course, when there is no conflict, there is no need to study the wide range of interactions that can occur among intelligent agents. All agents are fundamentally assumed to be helping one another, and will trade data and hypotheses as well as carry out tasks that are requested of them. We call this aspect of the paradigm the *benevolent agent assumption*.

1.3 Interactions of a More General Nature

In the real world, agents are not necessarily benevolent in their dealings with one another. Each agent has its own set of desires and goals, and will not necessarily help another agent with information or with actions. Of course, while conflict among agents exists, it is not total. There is often potential for compromise and mutually beneficial activity. Previous work in distributed artificial intelligence, bound as it has been to the benevolent agent assumption, has generally been incapable of handling these types of interactions.

Intelligent agents capable of interacting even when their goals are not identical would have many uses. For example, autonomous land vehicles (ALV's), operating in a combat environment, can be expected to encounter both friend and foe. In the latter case there need not be total conflict, and in the former there need not be an identity of interests. One can imagine ALV's from two different battalions, or from two NATO allies, meeting and having different goals (though of course a certain amount of cooperation would be called for). Even encounters with the enemy may have the potential for mutually beneficial compromise (for example, each may find it advantageous to temporarily avoid combat).

Other domains in which general interactions are prevalent are resource allocation and management tasks. An automated secretary [12], for example, may be required to coordi-

nate a schedule with another automated (or human) secretary, while properly representing the desires of its owner. The full capability to negotiate, to compromise and promise, would be highly desirable in these types of encounters.

Finally, even in situations where all agents in theory have a single goal, the complexity of interaction might be better handled by a framework that recognizes and resolves sub-goal conflict in a general manner. For example, robots involved in the construction of a space station are fundamentally motivated by the same goal; in the course of construction, however, there may be many minor conflicts caused by occurrences that cannot fully be predicted (e.g., fuel running low, drifting of objects in space). The building agents, each with a different task, could then negotiate with one another and resolve conflict.

1.4 Game Theory's Model and Extensions

In modeling the interaction of agents with potentially diverse goals, we borrow the simple construct of game theory, the payoff matrix. Consider, for example, the following matrix:

	A	B
A	3/1	2
B	2/5	0/1

The first player is assumed to choose one of the two rows, while the second simultaneously picks one of the two columns. The row/column outcome determines the payoff to each; for example, if the first player picks row B and the second player picks column A, the first player receives a payoff of 2 while the second receives a payoff of 5. If the choice results in an identical payoff for both players, a single number appears in the square (e.g., the A/B payoff above is 2 for both players). Payoffs designate utility to the players of a particular joint move [18].

Game theory addresses the issues of what moves a rational agent will make, given that other agents are also rational. We wish to remove the *a priori* assumption that other agents will necessarily be rational, while at the same time formalizing the concept of

rationality in various ways (and with greater precision than is generally done in the game theory literature).

Our model in this paper allows communication among the agents in the interaction, and allows them to make binding promises to one another. We will consider a variety of assumptions about the rationality of the agents, as to how they decide on both moves and promises to make. The formalism handles the case of agents with disparate goals as well as the case of agents with identical goals.

§2. Notation

We expand on the notation developed in [8]. For each game there is a set P of players and, for each player $i \in P$, a set M_i of possible moves for i . For $S \subset P$, we denote $P - S$ by \bar{S} , and write \bar{i} instead of $\{i\}$ (so $\bar{i} = P - \{i\}$). We write M_S for $\prod_{i \in S} M_i$.

We denote by m_S an element of M_S ; this is a joint move for the players in S . To $m_S \in M_S$ and $m_{\bar{S}} \in M_{\bar{S}}$ correspond an element \vec{m} of M_P . The payoff function for a game is a function

$$p : P \times M_P \rightarrow \mathbb{R}$$

whose value at (i, \vec{m}) is the payoff for player i if move \vec{m} is made.

Each agent is able to specify a set of joint moves (i.e., elements of M_P) that specify outcomes the agent is willing to accept; this set is called an *offer group*. If any move or moves offered by one agent are likewise offered by all other agents, this set of moves constitutes the *deal* (i.e., the deal is the intersection of all the agents' offer groups). In practice, a single element of the deal set will be selected by a fair arbiter, and the result of the selection communicated to all agents. At that point, the agents are all compelled to carry out their part of the move. Of course, if the deal set has only one member, no arbiter is needed.

We now define a secondary payoff function $pay(i, m_i, D_i)$, the set of possible payoffs

to i of making move m_i and suggesting offer group P_i :

$$\text{pay}(i, m_i, P_i) = \begin{cases} \{p(i, \vec{d}) : \vec{d} \in P_i \wedge \exists O_i [O_i \in \text{allowed}_o(i, P_i) \wedge \vec{d} \in O_i]\}, & \text{if such a } \vec{d} \\ & \text{exists;} \\ \{p(i, \vec{m}) : \vec{m} \in \text{allowed}_m(i, m_i)\}, & \text{otherwise.} \end{cases}$$

$\text{allowed}_m(i, m_i)$ is the set of moves that other agents might potentially make while i makes move m_i , and $\text{allowed}_o(i, D_i)$ is the set of deals that other agents might make while i suggests offer group D_i . Our formalism implicitly separates offer groups from moves (i.e., there will be no effect on moves by offer groups or vice versa). Intuitively, this reflects simultaneously revealing one's move and offer group, with one's eventual action determined by others' offer groups (that is, only if there is no agreement will you have to carry out your move). Future work might investigate the situation where offers are made before moves are chosen, and may thus affect them.

For nonempty sets $\{\alpha_i\}$ and $\{\beta_j\}$, we write $\{\alpha_i\} < \{\beta_j\}$ if $\alpha_i < \beta_j$ for all i, j (and say that $\{\beta_j\}$ *strictly dominates* $\{\alpha_j\}$). Likewise, we write $\{\alpha_i\} \leq \{\beta_j\}$ for nonempty sets $\{\alpha_i\}$ and $\{\beta_j\}$ if $\alpha_i \leq \beta_j$ for all i, j and the inequality is strict in at least one case. We then say that $\{\beta_j\}$ *dominates* $\{\alpha_j\}$.

Finally, we define $p(S, d_S)$ as $\{p(i, \vec{d}) : i \in S \wedge d_S \in M_S\}$. These are the possible payoffs to a group S of players of making move d_S .

2.1 Rational Moves

We will denote by $R_m(p, i)$ the set of rational moves for agent i in game p . We use the following definition to constrain what moves are elements of $R_m(p, i)$ (i.e., what moves are rational):

$$\text{pay}(i, d_i, \emptyset) < \text{pay}(i, c_i, \emptyset) \Rightarrow d_i \notin R_m(p, i). \quad (1)$$

In other words, if, when no binding agreement will be reached, every possible payoff to i of making move d_i is less than every possible payoff to i of making move c_i , then d_i is irrational for i . Of course, this does not imply that c_i is rational, since better moves may still be available.

In general, it will not be possible to fully specify the value of $pay(i, m_i, \emptyset)$ for all m_i , since there is not full information as to the moves that the other agents will make. Instead, we use (1) to show that some moves are not rational.

2.2 Rational Offer Groups

We define a *rational offer group* in a way analogous to how we defined a rational move above. We denote by $R_o(p, i)$ the set of rational offer groups for agent i in game p , and characterize a rational offer group by the following constraint on $R_o(p, i)$'s members:

$$\exists m_i [pay(i, m_i, P_i) < pay(i, m_i, O_i)] \Rightarrow P_i \notin R_o(p, i). \quad (2)$$

In other words, if for some move m_i every possible payoff resulting from offer group P_i is less than every possible payoff resulting from offer group O_i , then P_i is not a rational offer group.

There is one additional constraint on members of $R_o(p, i)$: rational offer groups specify (through the function p) a continuous range of payoffs that are acceptable to an agent. Intuitively, a rational offer group must reflect the notion of "monotonic satisfaction"—if a rational agent is satisfied with a particular payoff, he will be satisfied with one of equal or greater value (this is a fundamental meaning of "utility"). Formally, we write

$$[p(i, \vec{r}) \leq p(i, \vec{s}) \wedge \vec{r} \in O_i] \Rightarrow \vec{s} \in O_i \quad (3)$$

for all $O_i \in R_o(p, i)$ and moves \vec{r} and \vec{s} . For a particular game and player, a rational offer group can thus be unambiguously specified by any of its members with the lowest payoff.

In general, there may be more than one rational offer group for an agent in a game. If full information were available to an agent about the offers others were going to make (along with their "backup moves"), it would be trivial to determine $R_o(p, i)$. In practice such information is not available, but a rational agent i may be able to discover *some* rational offer group, i.e., some offer group provably in $R_o(p, i)$.

2.3 Rational Moves and Offer Groups for a Set of Players

We also wish to define the rational moves and the rational offer groups available to a set of players. For $S \subset P$, we denote by $R_m(p, S)$ the rational moves for the group S in

the game p . It follows that the members of $R_m(p, S)$ are elements of M_S . We assume that

$$R_m(p, S) \subset R_m(p, S') \times M_{S-S'} \text{ for } S' \subset S.$$

This states that no rational move for a set can require irrationality on the part of a subset. An obvious consequence of this assumption is that

$$R_m(p, S) \subset \prod_{i \in S} R_m(p, i).$$

A move that is rational for a group of players is thus rational for each player in the group.

Similarly, we denote by $R_o(p, S)$ the set of rational deals for S in the game p (that is, the members of $R_o(p, S)$ are sets of elements from M_P). It is the "crossproduct-intersection" of rational offer groups for the individual agents:

$$R_o(p, S) = \{O : O = \bigcap_{j \in S} O_j \wedge O_j \in R_o(p, j)\}.$$

2.4 Rationality Assumptions

The value of $pay(i, m_i, D_i)$ will depend, of course, on the values of $allowed_m(i, m_i)$ and $allowed_o(i, D_i)$ (i.e., the moves and the deals that other agents can make). In order to constrain the value of pay , we now define each of the *allowed* functions ($allowed_m$ is defined as in [8]).

1. **Minimal move rationality:** $allowed_m(i, m_i) = M_i$. Each player assumes that the others may be moving randomly.
2. **Separate move rationality:** $allowed_m(i, m_i) \subset R_m(p, \bar{i})$. Each player assumes that the others are moving rationally.
3. **Unique move rationality:** For all m_i and m'_i , $allowed_m(i, m_i) = allowed_m(i, m'_i)$ and $|allowed_m(i, m_i)| = 1$. Each player assumes that the others' moves are fixed in advance. This may be combined with separate rationality.

The assumptions above do not fully specify what is or is not a rational move. Rather, they help constrain the set of rational moves by allowing us to prove that certain moves

are not rational. We now define analogous assumptions regarding deals other agents might be making:

1. **Minimal deal rationality:** $allowed_o(i, D_i) \subset \mathcal{P}(M_P)$, where $\mathcal{P}(M_P)$ denotes the power set of M_P . Each player assumes that the others may be making random deals.
2. **Separate deal rationality:** $allowed_o(i, D_i) \subset R_o(p, \bar{i})$. Each player assumes that the others are making rational deals.
3. **Unique deal rationality:** $allowed_o(i, D_i) = allowed_o(i, E_i)$ and $|allowed_o(i, D_i)| = 1$ for all D_i and E_i . Each player assumes that the others' offers are fixed in advance. This may be combined with separate deal rationality.

We will refer to the combination of separate and unique move rationality as individual move rationality, and to the combination of separate and unique deal rationality as individual deal rationality. As in [8], any move that can be proven irrational under the assumption of minimal move rationality will be similarly irrational under the other move rationality assumptions. Analogously, any offer group that can be proven irrational under the assumption of minimal deal rationality will be irrational under the other deal rationality assumptions.

§3. Rational Deal Characteristics

With our notational conventions defined, we now prove several characteristics of $R_o(p, i)$. We will henceforth use \bar{s} to denote any move that gives agent i his highest payoff.

Theorem 1 (Existence of a non-null rational offer group). $|R_o(p, i)| \geq 1$.

Proof. If $R_o(p, i)$ were empty then i would do best by making no offers and relying on his move to generate his payoff. But $pay(i, m_i, \bar{s})$ will be greater than or equal to $pay(i, m_i, \emptyset)$ for all m_i (since \bar{s} will either be matched by other agents, increasing i 's payoff, or will not be matched, and will therefore be harmless since it doesn't affect other's moves). Thus the offer group $\{\bar{s}\}$ would also be in $R_o(p, i)$, guaranteeing it to have at least one non-null member. \square

It follows directly from the definition of a rational offer group (3) that all non-empty members of i 's set of rational offer groups include \bar{s} . Together with Theorem 1, this implies that it is always rational for an agent to include in his offer group the move that gives him his highest payoff.

In addition, an agent can often restrict his offers to those whose payoffs are higher than that which he can get by making the null offer, relying on his move to give him this payoff.

Theorem 2 (Lower bound). *Assuming unique deal rationality, if for any move m_i and joint move $\vec{d} \neq \bar{s}$,*

$$p(i, \vec{d}) \leq \text{pay}(i, m_i, \emptyset),$$

$$\exists O_i [O_i \in R_o(p, i) \wedge \vec{d} \notin O_i].$$

Proof. There are two cases:

1. $p(i, \vec{d}) < \text{pay}(i, m_i, \emptyset)$: The only way for \vec{d} to be in some rational offer group P_i is for the \vec{d} deal not to be accepted (otherwise $\text{pay}(i, m_i, P_i)$ would be dominated by the offer group $\text{pay}(i, m_i, O_i)$ where $O_i = \{\vec{c} : p(i, \vec{c}) > p(i, \vec{d})\}$). But if \vec{d} is not accepted, then it is equivalent to another offer group that includes only those moves with payoffs higher than \vec{d} . This smaller offer group will then also be in $R_o(p, i)$.
2. $p(i, \vec{d}) = \text{pay}(i, m_i, \emptyset)$: Assume that \vec{d} is in some rational offer group P_i . If \vec{d} is not accepted, or is accepted along with other offers, then $\text{pay}(i, m_i, O_i) > \text{pay}(i, m_i, P_i)$ where $O_i = \{\vec{c} : p(i, \vec{c}) > p(i, \vec{d})\}$, so there is another rational offer group (namely O_i) without \vec{d} . If \vec{d} is the only accepted offer, then $\text{pay}(i, m_i, \bar{s}) = \text{pay}(i, m_i, P_i)$ (where \bar{s} is the move that gives i his highest payoff), since \bar{s} will not be accepted anyway and therefore $\text{pay}(i, m_i, \bar{s}) = \text{pay}(i, m_i, \emptyset)$. Again, there is a rational offer group that does not include \vec{d} . \square

Note that Theorem 2 will not hold for \bar{s} (i.e., the joint move that gives i his highest payoff) since that would contradict Theorem 1 (Theorem 2's proof makes implicit use of the fact that $\vec{d} \neq \bar{s}$ in its construction of the dominating offer group O_i). Note also that

Theorem 2 will *not* hold under minimal deal rationality. Imagine that a perverse opponent chooses his offer group as follows:

1. If you include in your offer group deals with low payoff (for you), he will accept the deal with your best payoff;
2. If you don't offer that low deal he will accept no deals and you will have to rely on your move to get a payoff.

Under these circumstances (fully consistent with minimal deal rationality), it might be to your advantage to offer a low-payoff deal, since that might be the only way to get your maximal payoff.

3.1 Restricted Case Analysis

The consequences of Theorem 2 will differ, of course, based on assumptions about $allowed_m$ since these will affect $pay(i, m_i, \emptyset)$ for any given m_i . Consider the following payoff matrix:

	A	B
A	1/4	0/5
B	3/2	2/7

It is shown in [8] that, assuming minimal move rationality (potentially random or even malevolent moves by other agents), the row agent can still use "restricted case analysis" to constrain his move to B . If unique deal rationality can be assumed then the offer group consisting solely of move B/A (i.e., bottom left corner) is *guaranteed* by Theorems 1 and 2 to be a rational offer group. Of course, there may be other rational offer groups, for example the offer $\{B/B, B/A\}$, depending on what deals the other player can offer.

We formalize part of the above discussion:

Corollary 3 (Restricted case analysis). *Assuming minimal move rationality and unique deal rationality, if for some c_i and d_i , for all c_i and d_i ,*

$$p(i, \vec{d}) < p(i, \vec{c}),$$

then there exists an $O_i \in R_o(p, i)$ such that no \vec{d} is in O_i .

Proof. Follows from Lemma 3 in [8] and Theorem 2. \square

3.2 Case Analysis and Iterated Case Analysis

Restrictions on rational offer groups analogous to Corollary 3 apply for case analysis and iterated case analysis under the assumptions of unique and individual move rationality, respectively. The case analysis situation is represented in the following payoff matrix, seen from the row player's perspective:

	A	B
A	4/1	2
B	3/5	0/1

The row player need only assume that the other player's move will not be affected by his own move (i.e., unique move rationality) to realize that making move A is in all circumstances superior to making move B. As long as unique deal rationality can also be assumed, there is a guaranteed rational offer group consisting only of move A/A.

Corollary 4 (Case analysis). *Assuming unique move rationality and unique deal rationality, if for some c_i and d_i , for all c_i and d_i with $c_i = d_i$,*

$$p(i, \vec{d}) < p(i, \vec{c}),$$

then there exists an $O_i \in R_o(p, i)$ such that no \vec{d} is in O_i .

Proof. Follows from Lemma 4 in [8] and Theorem 2. \square

Similarly, if the column player can assume that the row player is rational and making moves independent of the column player's moves (i.e., individual move rationality), then he can prove that move B is optimal in the above matrix (since the row player will play A). With unique deal rationality, he has a guaranteed rational offer group of A/B.

The effect of Theorems 1 and 2 is to show us that there is always a rational offer group that includes an agent's highest payoff outcome, and includes no outcomes below or

equal to what he could achieve without deals. Below, we consider other constraints on an agent's rational offer groups.

§4. The Group Rationality Theorem

The work in [8] and [9] was concerned with the formalization of cooperative behavior, given certain constraints about the agents participating in an interaction. Using our notation, a desirable general result would have been

$$\text{pay}(P, \vec{d}, \emptyset) < \text{pay}(P, \vec{c}, \emptyset) \Rightarrow \vec{d} \notin R_m(p, P), \quad (4)$$

that is, if any joint move for all players is dominated by any other, then the dominated joint move is not rational for them. This result could not be proven, and the inability to do so stemmed directly from the lack of communication inherent in the model. Without at least minimal communication (e.g., self-identification), there is no way to coordinate on a universally perceived best move when several such moves exist.

We are now able to derive an important result about $R_o(p, P)$ very similar to the elusive non-communication result in (4).

Theorem 5 (Group offers). *Assuming individual deal rationality,*

$$p(P, \vec{d}) < p(P, \vec{c}) \Rightarrow \exists O_i [O_i \in R_o(p, i) \wedge \vec{d} \notin O_i]$$

for all $i \in P$.

Proof. There are two possible cases:

1. $\forall O_j [O_j \in R_o(p, \vec{i}) \Rightarrow \vec{d} \notin O_j]$: Since \vec{d} will not be a consummated deal, if P_i is any offer group containing \vec{d} then $\text{pay}(i, m_i, O_i) \geq \text{pay}(i, m_i, P_i)$ where $O_i = \{\vec{c} : p(i, \vec{c}) > p(i, \vec{d})\}$. Along with Theorem 1, this shows the existence of a non-null rational offer group without \vec{d} .
2. $\exists O_j [O_j \in R_o(p, \vec{i}) \wedge \vec{d} \in O_j]$: All other agents are rational (by assumption), and any rational offer group that includes \vec{d} also includes \vec{c} (3); thus, if P_i is any offer group containing \vec{d} , then $\text{pay}(i, m_i, O_i) > \text{pay}(i, m_i, P_i)$ where $O_i = \{\vec{c} : p(i, \vec{c}) > p(i, \vec{d})\}$.

This, along with Theorem 1, shows the existence of a rational offer group without \bar{d} .

□

Because of Theorem 5, a rational agent interacting with other rational agents knows that he need not offer a move that is dominated for all players—doing so cannot increase his payoff. If the other rational agents also know that all agents are rational, they too will realize that they can refrain from offering a move that is dominated for all players. Higher levels of knowledge [13], such as their knowing that all agents know that all agents are rational, are not needed. In addition, because of the definition of rational offer groups (3), the agents can refrain from offering any moves with smaller payoffs, since those groups would necessarily include the dominated move.

§5. Examples

We will now examine the consequences of our rational offer theorems in several additional types of games.

5.1 Best Plan

The best plan scenario is reflected in the following matrix:

	A	B
A	7	4
B	5	6

All agents recognize that there is a single best move; how will their offer groups reflect this? From Theorem 1, a rational agent knows that he can safely offer the move that gives him his best payoff (i.e., move A/A), even assuming minimal deal rationality on the part of other players (though the theorem is noncommittal as to whether other moves can or should be included with it). All players can also rule out move A/B using Theorem 2 if unique deal rationality holds (since A/B yields the lowest payoff). If there is an assumption of individual deal rationality, Theorem 5 can guarantee each agent that the offer group consisting solely of A/A is rational. Communication thus allows coordination on the best plan under more intuitive assumptions about the interaction than those used in [8].

5.2 Breaking Symmetries—Multiple Best Plan

Our rational offer group theorems allow us to solve the "Multiple Best Plans" case that could not be solved in [8]. The following matrix illustrates the scenario:

	A	B
A	-1	2
B	2	-1

Assuming minimal deal rationality, an agent can rationally offer B/A and A/B . In addition, assuming unique deal rationality an agent knows that he can rationally not offer A/A and B/B (since they are lowest yield moves). This analysis can be done by both agents if they are rational and operating under the unique deal assumption. Their offer sets will overlap on the multiple best outcomes; selection of a single alternative from the multiple agreements then occurs.

5.3 Prisoner's Dilemma

The prisoner's dilemma is represented by the following matrix:

	C	D
C	3	0/5
D	5/0	1

Each agent most desires to play D while the opponent plays C , then to play C along with the opponent, then to play D along with the opponent, and least of all to play C while the opponent plays D . The dilemma comes about because case analysis implies that it is always better to play D ; both players choosing D , however, is less desirable *for both* than if they had chosen C . The dilemma has received much attention within the philosophy and game theory literature [2, 5, 22, 26]. In the usual presentation of the prisoner's dilemma, playing C is called "cooperating," and playing D is called "defecting." With the presence of communication, in fact, there is no dilemma:

Corollary 6 (Prisoner's Dilemma). *If all players know that all players are operating under the assumption of individual deal rationality, agents will cooperate in the prisoner's dilemma.*

Proof. The first player knows that it is rational to offer D/C (since it is rational even under minimal rationality, Theorem 1); he also knows it is irrational to offer C/D (from Theorem 2, since individual deal rationality includes unique deal rationality). By Theorem 5, there is a rational offer group without D/D . Now he knows that the other agent will not offer D/C (since the other agent is assumed rational and operating under the assumption of unique deal rationality, Theorem 2). Since D/C will certainly not be met, $pay(i, D, \{D/C\}) \leq pay(i, D, \{D/C, C/C\})$. Thus, the offer group $\{D/C, C/C\}$ is rational. The second agent will, if rational and working under the same assumptions, come to the same conclusion. The deal C/C will be struck, and the agents avoid the D/D trap. \square

§6. Extending the Model

For certain types of interactions, the model presented above (i.e., the various assumptions and theorems about rational moves and deals) does not specify rational activity in sufficient detail. We can extend the model in a variety of ways to handle these cases, and at the same time capture a wider range of assumptions about the interaction. In this section, we briefly present some of the extensions that might be made to our original model.

6.1 Similar bargainers

Consider the following payoff matrix (equivalent to game 77 in Rapoport and Guyer's taxonomy [23])

	A	B
A	3	2
B	5/0	0/5

Assuming separate deal rationality, the first player can assume that B/A should be in a rational offer group of his, and that B/B should not be. What else can be said about what constitutes a rational offer group in this game? There are three choices, namely $\{B/A\}$, $\{A/A, B/A\}$, and $\{A/B, A/A, B/A\}$. In order to decide among the choices, we would like to make more assumptions about the "bargaining tendencies" of the other agent

(since, in fact, some agents might be tougher deal-makers than others). We will ignore what value the agents might place on making a particular move in the absence of a deal, since the payoff is underdetermined.

Let us define two offer groups O_i and O_j to be *similar* if and only if they both have the same lower boundary for what deals are included or not included. $similar(O_i, O_j)$ is true if and only if

$$\exists n[p(i, \vec{d}) > n \Leftrightarrow \vec{d} \in O_i \wedge p(j, \vec{c}) > n \Leftrightarrow \vec{c} \in O_j]$$

for some number n .

It might seem that, since both players' payoffs are designated in numbers, it is reasonable to compare their utilities (e.g., player A values a payoff of 4 more than player B values a payoff of 3). In fact, this should not be taken for granted, and utility theory does not ordinarily allow such a comparison to be made. Nevertheless, if we use the similar bargainers definition, we implicitly assume some meaningful measure for comparing inter-personal utility.

One assumption to use in deciding upon rational offer groups is now that the other agent will accept deals that you would accept; that is, $O_j \in R_o(p, j) \Leftrightarrow O_i \in R_o(p, i)$ where $similar(O_i, O_j)$.

Under this assumption, we can decide what deal is rational in the above game. Player 1 reasons that if he offers $\{B/A\}$, player 2 (who is a similar bargainer) will offer only $\{B/B\}$. There will be no match. In the same way, if it would be rational for player 1 to offer $\{A/A, B/A\}$ then player 2 will offer $\{A/A, B/B\}$, with an agreement on A/A and a payoff of $\{3\}$ for both. If player 1 offers $\{A/B, A/A, B/A\}$ then player 2 will offer $\{A/B, A/A, B/B\}$ and there will be agreement on A/A and on A/B , with a payoff of $\{2, 3\}$ for both. Since $\{3\}$ dominates $\{2, 3\}$, agents who assume common knowledge [13] of the similar bargainer assumption should choose the rational offer group that yields agreement on A/A .

6.2 Stochastic Model—The Game of Chicken

Note, however, the following payoff matrix (commonly known as the game of chicken [23]):

	A	B
A	3	2/5
B	5/2	1

Two agents, even if they assume individual deal rationality and the similar bargainers assumption, will be faced with the following choices: a payoff of {3} or a payoff of {2, 3, 5}. According to our definitions, neither of these sets dominates the other, and it is not clear how to decide between them.

If, however, we extend the model to include a probabilistic choice from within the agreement set, it is clear that the latter agreement set dominates the former (with an expected value of 3.33 versus 3). This would correspond to the fair arbiter mentioned earlier tossing an n -sided coin to decide among the n members of an agreement set. The usual understanding of "utility" also supports viewing the payoffs in this way.

A further stochastic extension to our model would allow moves themselves to be specified probabilistically (e.g., A with probability .5, and B with probability .5). In the game theory literature, this is the distinction between pure strategies and mixed strategies [18]. An analysis of this model is beyond the scope of the present discussion.

6.3 Conjunctive Offers—Battle of the Sexes

In the game of chicken example presented above, there was an added complexity that was temporarily ignored: the possibility of "defection." If one agent reasons that the other agent will accept all payoffs above 2, it to the first agent's benefit to only offer moves of payoff 5 (this is analogous to the prisoner's dilemma, with the same potential that both players will use identical reasoning and no agreement will be reached). A similar problem can be seen in the so-called battle of the sexes matrix, seen below.

	A	B
A	-1	1/2
B	2/1	-1

One approach to solving this problem is to allow "composite" offers, for example, an offer consisting of a conjunct of several moves (the conjunct must be matched exactly in order for a deal to occur). Thus, the offer consisting of $A/B \wedge B/A$ can consistently be made by both agents without the potential of defection (and with an expected utility of 1.5 for each). This notion can be extended to general logical offers consisting of disjuncts, conjuncts and negations of joint moves. The battle of the sexes can thus be uniquely solved with the assumption of similarity in bargaining, if conjunctive offers are allowed.

6.4 Repeated Interactions

The technique of allowing conjunctive offers, presented above, in fact transforms the original game into a new supergame that can then be approached using our previous analysis. A similar technique can be used when agents will be participating in a fixed number of interactions, each with known payoff characteristics.

Imagine that there are n games, p_1, p_2, \dots, p_n . To each game j corresponds, for player i , a set of moves M_i^j . We now construct a new game whose moves for each player are the crossproduct of his moves in the individual games:

$$\prod_{j=1}^n M_i^j,$$

and whose outcomes are likewise crossproducts of the individual games' outcomes. Of course, the player's utility from any of these supergame outcomes is some function of his utilities from its constituent outcomes; summing the utilities seems reasonable, although utility theory does not actually sanction such an approach.

6.5 Specifying Partial Outcomes

We would like to generalize the analysis of deal-making and promises to the case where agents need not specify complete outcomes (i.e., a joint move for all players) in their offer group. In this section we will briefly note some consequences of generalizing the types of offers that can be made.

Let us say that each agent is able to specify an *offer group of partial moves* that specify outcomes the agent is willing to accept. If a move which is offered by one agent

is likewise accepted by all other agents in the relevant set, those agents are all required to carry out the move. Multiple agreements again result in a set of moves; in practice (as above), a single element of this set would be selected by a fair arbiter. A partial move r_S^i for agent i is therefore a vector that specifies a move for some set S of agents. i must be an element of S , but may be its only element. Formally, if r_S^i is a partial move,

$$r_S^i \subset M_i \times M_{S-i} = M_S.$$

The payoff vector p is defined over partial moves as

$$p(i, r_S^i) = \{p(i, \vec{r}) : \vec{r} = r_S^i \times M_{\bar{S}}\}.$$

We propose a new restriction on offer groups in the current model that did not exist previously: there can be no "contradictions" among an agent's offers. That is, we must assume that there are no two partial moves that are specified in the offer group such that, if both were accepted, the agent could only satisfy one. This problem did not arise above because all agents were bound by every offer. Thus, it was impossible that two different agents could be promised contradictory moves by an agent without themselves taking part in accepting all of those moves. Once the assumption is made that an agent will not promise contradictory partial moves, the analysis is similar to that developed above.

A further complication, however, arises in specifying the secondary payoff function $pay(i, m_i, D_i)$ within this new model; in particular, what is i 's payoff when a deal has been reached by a set of agents that does not include him? There are several possible alternatives, including i 's being forced to make move m_i , or being allowed to alter his move based on the agreement reached by the others. Analysis of the consequences these assumptions have on strategies will be a subject of future research, and is outside the scope of this paper.

There is, however, another method for handling the offer of partial moves that *does* allow our analysis to be used directly. We could define the offer of a partial move as implicitly including its extension in the total move space (i.e., if r_S^i is offered explicitly,

the implicit offer includes all joint moves in M_P that include r_S^i). Then, of course, our original analysis is completely appropriate. Further work might also consider the situation (with total or partial move offers) where deals consist of unions of offers rather than their intersection.

§7. Previous Work

The subject of interacting rational agents has been addressed within the field of artificial intelligence as well as in the discipline of game theory. Here we will briefly review relevant contributions from these two areas, and contrast our present approach with previous efforts.

7.1 Work in Artificial Intelligence

As mentioned above, researchers in distributed artificial intelligence have begun to address the issues arising in multi-agent interactions. Lesser and Corkill [4] have performed empirical studies to determine cooperation strategies with positive characteristics (such as, for example, what types of data should be shared among distributed processors). They are solely concerned with groups of agents who share a common goal, but have acknowledged the benefit even under this assumption of having agents demonstrate "skepticism" (i.e., not being distracted by others' information).

Georgoff [10, 11] has developed a formal model to combine separate plans of independent agents. The primary concern is to avoid destructive interference caused by simultaneous access to a shared resource. The model used assumes that the agents have separate goals, but that these goals do not directly oppose one another. Cooperative action is neither required nor exploited, except insofar as it allows agents to keep out of each other's way.

Other notable efforts include Smith's work on the contract net [7], Malone's work extending the contract net model using economic theory [19], and the theoretical work on knowledge and belief of Appelt, Moore, Konolige, Halpern and Moses [1, 14, 15, 16, 17, 20, 21].

The current work extends these previous models of interaction by allowing a fuller range of goal disagreements among agents. By using a framework that captures total and partial goal conflicts, it allows investigation into compromise, promises and cooperative action. None of these could be handled using previous schemes.

This paper considers the communication scenario in ways similar to the manner in which previous work [8, 9] investigated cooperation among rational agents when no communication occurs. Below we briefly note the advantages that were gained when communication and promises were added to the interaction model.

The best plan interaction was handled in our framework by assuming individual deal rationality. Because in the no-communication case this scenario could not be solved using individual move rationality, other assumptions were introduced: *informed rationality* in [8] and *common rationality* in [9]. Informed rationality, in our notation, constrained *allowed_m* in a way that assumed each player would respond rationally to the others' moves, whatever they might be. For a fixed m_S , the *restricted game* was defined to be

$$p|_{m_S} : P \times M_{\bar{S}} \rightarrow \mathbb{R} \quad p|_{m_S}(i, m_{\bar{S}}) = p(i, \vec{m}).$$

Intuitively, this was the game where the players in S were assumed to make the move M_S . Informed rationality was defined as follows:

- **Informed rationality:** $allowed_m(i, m_i) = R(p|_{m_i}, \vec{i})$. Each player assumes that all others will respond rationally to whatever move he makes.

It should be noted in passing that an assumption of common knowledge of rationality will also allow for a unique solution to the best plan case, though this has not been previously pursued in the literature.

To solve the prisoner's dilemma, even more assumptions had to be introduced. The interested reader is referred to [8] and [9] for full details.

Even using a variety of assumptions, previous work could not handle the multiple best plan case, where there are several outcomes all equally recognized as best by all players. To break the symmetry, some communication is needed, though this communication can

be as simple as self-identification and reliance on a common rule (e.g., agent with lowest name performs lowest ordered action). We were able to solve the multiple best plan case as easily as the best plan case, as well as handle several novel interactions through extensions to the model, such as the similar bargainers assumption.

7.2 Game Theory

Game theory has focused on a variety of interactions, and sought to characterize the types of actions that rational agents will take in each. Many of the same questions that come up in our work have been addressed by game theoreticians. Their approach, however, has left a great many important issues unexamined. Consider the following quote from the classic game theory text, [18]:

Though it is not apparent from some writings, the term "rational" is far from precise, and it certainly means different things in the different theories that have been developed. Loosely, it seems to include any assumption one makes about the players maximizing something, and any about complete knowledge on the part of the player in a very complex situation... [*Games and Decisions*, p. 5]

As another example, consider the following best plan interaction:

	A_2	B_2
A_1	4	1/2
B_1	3/1	2/3

It was demonstrated above that the best plan case can only be solved under particular definitions of rationality. Rapoport and Guyer, however, writing in [23], put forward the following *assumption* regarding agents' behavior (citing the similarity with [25]):

(A_3). If a game has a single Pareto equilibrium, the players will choose the strategy which contains it...

Our assumption (A_3) says that A_1A_2 is the natural outcome, which, of course, is dictated by common sense... we shall refer to this as a *prominent solution*. [*Taxonomy of 2 x 2 Games*]

In short, game theory has been willing to take for granted certain types of behavior without carefully formalizing its definitions of rationality, nor its assumptions of inter-agent knowledge.

These questions are particularly important in the field of artificial intelligence. We are not interested in characterizing game matrices: we want to characterize agent rationality and explore the consequences of various assumptions. The goal is to be able to implement intelligent agents whose strategies of behavior will be provably rational.

§8. Conclusion

In real world domains, intelligent agents will inevitably need to interact flexibly. Previous work has not modeled the full range and complexity of agents' varied goals. The benevolent agent assumption, which assumes that agents have identical or non-conflicting goals, has permeated previous approaches to distributed AI.

This paper has presented a framework for interaction that explicitly accounts for communication and promises, and allows multiple goals among agents. The model provides a unified solution to a wide range of problems, including the types of interactions discussed in [8] and [9]. Through the use of communication and binding deals, agents are able to coordinate their actions more effectively, and handle interactions that were previously problematical. By extending the communication model even further, a wider variety of interactions can be handled.

Acknowledgement

The authors wish to thank Matt Ginsberg, who has played an invaluable role in the development of our ideas on cooperation among rational agents. This research has been supported by DARPA under NAVELEX grant number N00039-83-C-0136.

References

- [1] Appelt, D.E., Planning natural language utterances to satisfy multiple goals, Tech Note 259, SRI International, Menlo Park, California (1982).
- [2] Axelrod, R., *The Evolution of Cooperation*, Basic Books, Inc., New York (1984).
- [3] Cammarata, S., McArthur, D., and Steeb, R., Strategies of cooperation in distributed problem solving, *IJCAI-83*, Karlsruhe, West Germany (1983) 767-770.
- [4] Corkill, D.D., and Lesser, V.R., The use of meta-level control for coordination in a distributed problem solving network, *IJCAI-83*, Karlsruhe, West Germany (1983) 748-756.
- [5] Davis, L., Prisoners, paradox and rationality, *American Philosophical Quarterly* 14 (1977).
- [6] Davis, R., A model for planning in a multi-agent environment: steps toward principles for teamwork, Working Paper 217, MIT AI Lab (1981).
- [7] Davis, R., and Smith, R. G., Negotiation as a metaphor for distributed problem solving, *Artificial Intelligence* 20 (1983) 63-109.
- [8] Genesereth, M. R., Ginsberg, M. L., and Rosenschein, J. S., Solving the prisoner's dilemma, HPP Report 84-41 (1984).
- [9] Genesereth, M. R., Ginsberg, M. L., and Rosenschein, J. S., Cooperation without communication, HPP Report 84-36 (1984).
- [10] Georgeff, M., Communication and interaction in multi-agent planning, *AAAI-83*, Washington, D.C. (1983) 125-129.
- [11] Georgeff, M., A Theory of action for multi-agent planning, *AAAI-84*, Austin, Texas (1984) 121-125.
- [12] Goldstein, I. P., Bargaining Between Goals, A.I. Working Paper 102, Massachusetts Institute of Technology Artificial Intelligence Laboratory (1975).
- [13] Halpern, J. Y., and Moses, Y., Knowledge and common knowledge in a distributed environment, IBM Research Report IBM RJ 4421 (1984).
- [14] Konolige, K., A first-order formalization of knowledge and action for a multi-agent planning system, Tech Note 232, SRI International, Menlo Park, California (1980).
- [15] Konolige, K., Circumscriptive ignorance, *AAAI-82*, Pittsburgh, Pennsylvania (1982) 202-204.
- [16] Konolige, K., A deductive model of belief, *IJCAI-83*, Karlsruhe, West Germany (1983) 377-381.
- [17] Konolige, K., *A Deduction Model of Belief*, Ph.D. Thesis, Stanford University (1984).
- [18] Luce, R.D., and Raiffa, H., *Games and Decisions, Introduction and Critical Survey*, John Wiley and Sons, New York (1957).

- [19] Malone, T.W., Fikes, R.E., and Howard, M.T., Enterprise: a market-like task scheduler for distributed computing environments, Working paper, Cognitive and Instructional Sciences Group, Xerox Palo Alto Research Center (1983).
- [20] Moore, R.C., Reasoning about knowledge and action, Tech Note 191, SRI International, Menlo Park, California (1980).
- [21] Moore, R.C., A formal theory of knowledge and action, Tech Note 320, SRI International, Menlo Park, California (1984). Also to appear in *Formal Theories of the Commonsense World*, Hobbs, J.R., and Moore, R.C. (Eds.), Ablex Publishing Co. (1984).
- [22] Parfit, D., *Reasons and Persons*, Clarendon Press, Oxford (1984).
- [23] Rapoport, A. and Guyer, M., A taxonomy of 2 x 2 games, *Yearbook of the Society for General Systems Research XI* (1966) 203-214.
- [24] Rosenschein, J.S., and Genesereth, M.R., Communication and cooperation, HPP Report 84-5 (1984).
- [25] Schelling, T.C., *The Strategy of Conflict*, Oxford University Press, New York (1963).
- [26] Sowden, L., That there is a dilemma in the prisoners' dilemma, *Synthese* 55 (1983) 347-352.

Copyright © 1985 by KSL and
Comtex Scientific Corporation

FILMED FROM BEST AVAILABLE COPY