Communication-Free Interactions among Rational Agents: A Probabilistic Approach

Jeffrey S. Rosenschein
Computer Science Department
Hebrew University
Jerusalem, Israel

John S. Breese
Rockwell Science Center, Palo Alto Laboratory
444 High Street
Palo Alto, California

Abstract

Recent work on interactions among rational agents has put forward a computationally tractable, deduction-based scheme for automated agents to use in analyzing multiagent encounters. While the theory has defined irrational actions, it has underconstrained an agent's choices: there are many situations where an agent in the previous framework was faced with several potentially rational actions, and no way of choosing among them. This paper presents a probabilistic extension to the previous framework of Genesereth, Ginsberg, and Rosenschein [5] that provides agents with a mechanism for further refining their choice of rational moves. At the same time, it maintains the computational attractiveness of the previous approach.

The probabilistic extension is explicitly representing uncertainty about other players' moves. A three-level hierarchy of rationality is defined, corresponding to ordinal, stochastic, and utility dominance among alternative outcomes. The previous deduction-based formalism is recast in probabilistic terms and is seen to be a particular special case of a more encompassing dominance theory. A technique is presented for using the dominance ideas in interactions with other agents operating under various types of rationality.

1 Introduction

1.1 Interactions Among Rational Agents

Research on artificial intelligence (AI) has begun to concern itself with the design of an autonomous agent operating in real-world environments. Along one dimension, this requires that the agent be capable of dealing with dynamic and incompletely specified situations. It must be able to reason about change, recover from failures, and deal with uncertainty both in the state of the world and in the effects of its own actions.

An equally important capability is the ability to interact flexibly with other agents. There are, in fact, few scenarios where an agent could be expected to operate with *complete* autonomy; almost always there will be others with whom the agent must interact. This will be true whether the agent is operating on a factory floor, building outposts on Mars, or running errands to the corner store. These other agents will in general possess a wide range of reasoning capabilities and thus the agent should be capable of interacting flexibly with agents of different rationality "types".

There has been considerable work in recent years by AI researchers on formalisms for representing interagent beliefs [12, 13, 1, 14, 15, 7], an important component of the reasoning necessary for cooperation. Agents must reason about one another's beliefs to predict activity, provide information, and adapt their own behavior to others' expectations.

Another line of work has been considering the agent interactions themselves as objects about which to reason [4, 5, 17]. In this research, the agents have been defined as operating under the constraints of various *rationality axioms* that restrict their choices in interactions. The effects of various axioms and their relationships to one another have been analyzed.

The current paper continues along this latter line of research. The basic extension proposed is to recognize that reasoning about other agents' actions must deal with uncertainty and to incorporate an explicit mechanism for doing so. Uncertainty is inherent in encounters because of incomplete information about others' objectives, options, and reasoning processes. We are addressing issues of uncertainty against a backdrop of an increased in decision-theoretic concepts of probability and utility theory in AI research [3, 8]. At the same time, we exploit the fact that decision- and game-theorists have been considering the use of Bayesian decision theory in situations of strategic interaction [2, 9, 21, 16]. Though we are not proposing an extension to the concepts of equilibrium proposed by game-theorists, the work in this paper integrates previous studies of rational interaction based on a deductive framework with decision-theoretic ideas and is a first step towards operationalizing recent advances in game-theoretic solution concepts.

1.2 Perspectives on Multiagent Interactions

We examine reasoning about other agents from two different perspectives, the "prescriptive/descriptive" approach and the "jointly prescriptive" approach. Both perspectives have their place in the theory of rational interacting agents, though each leads us to ask different questions about how automated agents should be designed. This paper is focused on prescriptive/descriptive issues, though we make several observations and report results regarding jointly prescriptive methods.

1.2.1 Prescriptive/Descriptive

A "prescriptive/descriptive" approach requires two types of theories [10] to fully capture a multiagent interaction. First, we need a normative theory of what our primary agent *should* do given its values and information. We have a prescriptive theory when we not only define these normative principles for rational behavior, but augment these with a prescription or method for identifying rational moves. This is precisely the approach we take in developing our notion of prescriptive rationality. Second, we require a descriptive theory of other agents. A descriptive theory is useful to the extent it can be used to predict the actions of other agents, and may be based on varying degrees of assumed "rationality" of others.

The "prescriptive/descriptive" approach is basically decision analytic: using our model of interaction, we prescribe a particular course of action for one agent based on the description it has of other agents. This was the approach taken in previous DAI work, where different information about others' rationality would cause an agent to act appropriately. This "prescriptive/descriptive" perspective is central in our design of an agent capable of interacting intelligently, particularly when we will have no control over (and limited information about) the design of the other agents.

1.2.2 Jointly Prescriptive

Of course, if our descriptive theory is the same as our prescriptive theory, i.e., if the best theory one has about other agents is based on introspection regarding one's own reasoning processes, this results in a "jointly prescriptive" approach. "Jointly prescriptive" concerns form the basis for much of modern game theory [11]. These approaches, by and large, develop models of interaction that have certain globally desirable properties, given that all agents subscribe to the same fundamental solution strategies and have common knowledge regarding most aspects of the problem. The "jointly prescriptive" perspective is well-suited to closed systems where the interacting agents are all centrally designed. With total control over their methods of interaction (and hence the ability to engineer away uncertainty regarding others' decision-making strategies), the designer is looking for desirable properties, such as stability and pareto-optimality of solutions.

The jointly prescriptive perspective also has a role to play in competitive interactions. For example, some interaction strategies are known from the game theory literature to be "stable," i.e., if an agent uses this strategy, no opponent can benefit from playing any other strategy. A designer could feel safe in incorporating such a strategy into his agent—he need have no fear of the strategy's presence becoming known, since there is no effective counterstrategy. Thus the identification of stable strategies (which is a jointly prescriptive notion) can be important to any single agent's designer. Similarly, a demonstration of a strategy's stability *and* pareto-optimality might be an effective argument in getting many agents' designers to incorporate it:[1] the best that other agents can do is to "play along," and the overall final results have certain desirable characteristics.

[1]This was, for example, the argument made in [6]

1.3 Assumptions

This paper is concerned with single interactions among agents; though there is a mechanism for using the results of past encounters, there is no explicit concern about future interactions. Each agent is assumed capable of assigning some value to a hypothetical outcome, and (in this paper) we will assume that these assigned payoff values, for all agents, are common knowledge among them all.[2] In addition, once the interaction has been recognized, there is no further communication among the agents; each must decide on its action alone. This is the no-communication scenario used in [5].[3] The agents are assumed to be operating under certain axioms, to be discussed, that control their behavior.

1.4 Overview

In Section 2 we introduce the formal notation for our analysis. In Section 3, various forms of dominance among alternatives are developed. The deduction-based formalism given in [17] is recast in probabilistic terms, and is seen to be a particular special case of a more encompassing dominance theory.

In Section 4 we consider questions relating to the design of an agent using the dominance relations, in the "prescriptive" portion of a "prescriptive/descriptive" approach. Axioms of behavior are given in Section 5 that might describe our agent's opponents,[4] and the ramifications these axioms have on the prescriptive dominance techniques are discussed. In Section 6 we briefly consider, from the "jointly prescriptive" perspective, the global properties of the methods we have outlined.

---

[2] Uncertainty about payoffs can be incorporated into the framework, and is a topic for future research. Also, common knowledge is not always required; for a fuller discussion of how much knowledge is actually needed, see [18].

[3] While this scenario is a simplification of what might be found in real-world encounters, it is a useful starting point for an analysis of interactions. There are also a variety of instances when the assumption that no communication is possible is quite realistic, such as when agents designed in different countries or by different manufacturers unexpectedly encounter one another.

[4] Throughout this paper, our use of the term "opponent" should not be taken in its colloquial sense. When our agent interacts with other agents, we will sometimes refer to them as its opponents, without intending that the agents are necessarily involved in conflict. There may be a convergence of interests among all parties.

2 Notation

We will follow the convention of representing a game as a payoff matrix. Figure 1 is a representation of a two agent encounter.

K

		c	d
J	a	3 1	1 2
	b	2 5	0 1

Figure 1: A Payoff Matrix

A game corresponds to 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 $\overline{S}$. We denote by m_S an element of M_S; this is a collective move (or a "joint" move) for the players in S. To $m_S \in M_S$ and $m_{\overline{S}} \in M_{\overline{S}}$ corresponds 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. The function p thus encodes the payoff matrix in function form.

We denote by $prob_i(m_{\overline{i}} \mid m_i, \xi)$ the probability distribution that agent i has over all the other players making move $m_{\overline{i}}$ (with ξ representing i's knowledge of the world, including his knowledge of other agents). The probability may depend, as seen from this expression, on i's own move m_i.

We could use dual matrices to represent an interaction between agents, with associated probability distributions on their moves. Consider the two matrices in Figure 2.

The left matrix is to be interpreted in the same manner as it was above, namely as defining the payoffs each agent will receive from various outcomes. In addition, each agent is assumed to have a probability distribution on the other's moves, given a move of his own. The second matrix in Figure 2 displays these distributions. For example, if J considers that he will make

K

J		c	d
	a	1, 4	3, 3
	b	2, 1	4, 2

$prob_K(m_J \mid m_K)$

$prob_J(m_K \mid m_J)$		c	d
	a	.4, .2	.6, .5
	b	.7, .8	.3, .5

Figure 2: Payoff and Probability Matrix

move b, he considers that there is a .7 probability that K will make move c, and a .3 probability that K will make move d. Of course, in the probability matrix the columns sum to 1 for K, and the rows sum to 1 for J.

We define a secondary payoff function $pay(i, m_i)$, which gives us the *set* of possible payoffs to i of making move m_i:

$$pay(i, m_i) = \{p(i, \vec{m}) : prob_i(m_{\bar{\imath}} \mid m_i, \xi) > 0\}. \tag{1}$$

The expression $prob_i(m_{\bar{\imath}} \mid m_i, \xi) > 0$ denotes the set of responses "considered possible" to i's move m_i.[5] There are many potential moves that might be expected of other agents, depending on assumptions about them (and their assumptions about you), and similarly, many different subjective probability distributions that one might have over their potential moves. In Section 5 we list several alternate definitions and indicate how each affects the *pay* function or probabilities.

The final element of our notation that needs to be introduced is the notion of a "utility function" over payoffs. The utility function summarizes the agent's attitudes toward uncertain options, while payoffs summarize the agent's valuation under certainty of each possible joint move. The *utility of a joint move* for agent i in our notation is represented as $U_i(p(i, \vec{m}))$; it is a function from the real numbers to the real numbers. We then define the *expected utility* for agent i of a joint move as follows:

$$EU_i(\vec{m}) = \sum_{m_{\bar{\imath}} \in M_{\bar{\imath}}} U_i(p(i, \vec{m}))\, prob_i(m_{\bar{\imath}} \mid m_i, \xi). \tag{2}$$

More generally, the summation can be replaced by an integration. Von Neumann and Morgenstern, in their foundational work [20] , formalized ratio-

[5]Careful readers will note that this expression subsumes the role of the *allowed* function in [17].

nality in terms of axioms that require an agent to behave as if maximizing expected utility.

3 Dominance

A concept essential to this work is *dominance*: when the payoffs resulting from one move are better than those resulting from some other move, for some precise definition of "better," the inferior move is said to be *dominated*. Previous treatments used only one kind of dominance, namely *ordinal dominance*, an "absolute" dominance between the members of two sets. In this paper we consider how two other kinds of dominance, *stochastic* and *utility* dominance, can be combined with the axiomatic approach.

3.1 Ordinal Dominance

Ordinal dominance is straightforward: for nonempty sets $\{\alpha_i\}$ and $\{\beta_j\}$, we say that $\{\alpha_i\}$ is ordinally dominated by $\{\beta_j\}$ (written $\{\alpha_i\} <_o \{\beta_j\}$) if $\alpha_i \leq \beta_j$ for all i, j, and at least one element of $\{\alpha_i\}$ is less than every element of $\{\beta_j\}$. For example, the set $\{5, 3\}$ ordinally dominates the set $\{3, 1\}$, since *every* member of the first is greater than or equal to every member of the second, and in at least one case the relationship is strictly greater than.

3.2 Stochastic Dominance

3.2.1 The Intuition

Before launching into the formal definition of stochastic dominance, we will present the intuitions behind its use.

A *lottery* is defined to be a set of payoffs with associated probabilities. A lottery can be viewed as a state contingent payoff—in an interaction between agents, the payoff is contingent on the (uncertain) move of the opponent.

Stochastic dominance [22] between two alternative lotteries is commonly represented graphically as follows. Consider a graph whose x-axis represents various payoffs, and whose y-axis represents cumulative probabilities (i.e., runs from 0 to 1). For each agent's lotteries, we draw a curve onto this coordinate space whose y position at any x value represents the probability that the agent will receive *less than* that value from that lottery. Each curve begins at the point (p,0) and increases to a maximum of (q,1), where p and q are the minimum and maximum possible payoffs, respectively. If the first lottery's curve lies completely below and to the right of a second lottery's curve (with possible overlap—but no crossing—of the curves permitted), we say that the first lottery stochastically dominates the second. This means that for any given value, the player has a better chance of getting it or less from the second lottery than from the first.

For example, consider an agent that has two lotteries available to him. In the first, he has .2 chance of getting a payoff of 4, a .5 chance of getting a payoff of 6, and a .3 chance of getting a payoff of 7. We draw this lottery's curve as in Figure 3.[6] The curve rises by .2 at 4, rises an additional .5 at 6, and rises an additional .3 at 7.

Figure 3: Agent's First Lottery

Now imagine that there is a second lottery, where he has a .3 chance of getting a payoff of 3, a .2 chance of getting 5, and a .5 chance of getting 6. This second lottery's curve looks like that in Figure 4.

If we now combine these two curves, it is evident that at all points the second curve are above those of the first curve for a given payoff—thus, the second lottery has a higher probability of getting a particular value or less for all values and therefore the first lottery stochastically dominates the second (see Figure 5).

[6]The diagram in Figure 3 is typical of lotteries with discrete moves and payoffs—a step function. We could, just as easily, have a continuous set of payoffs, which would result in a smooth curve in the diagram with no vertical climbs.

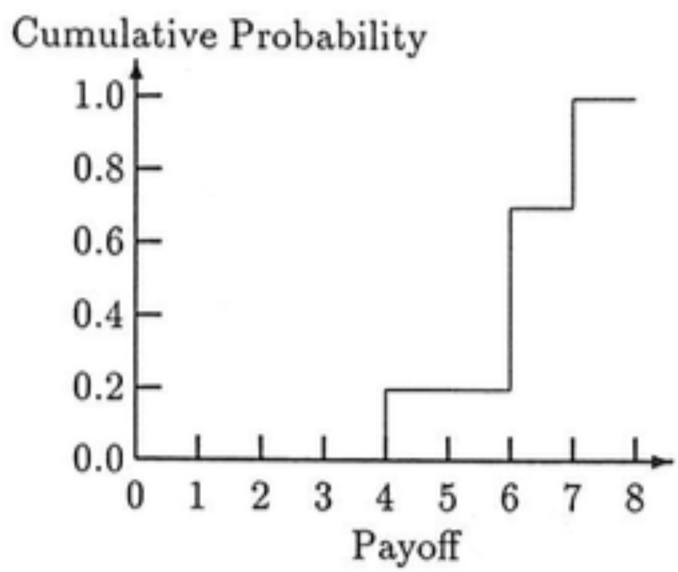

Cumulative Probability
1.0
0.8
0.6
0.4
0.2
0.0
0 1 2 3 4 5 6 7 8
Payoff

Figure 4: Agent's Second Lottery

Figure 5: A Comparison of the Two Lotteries

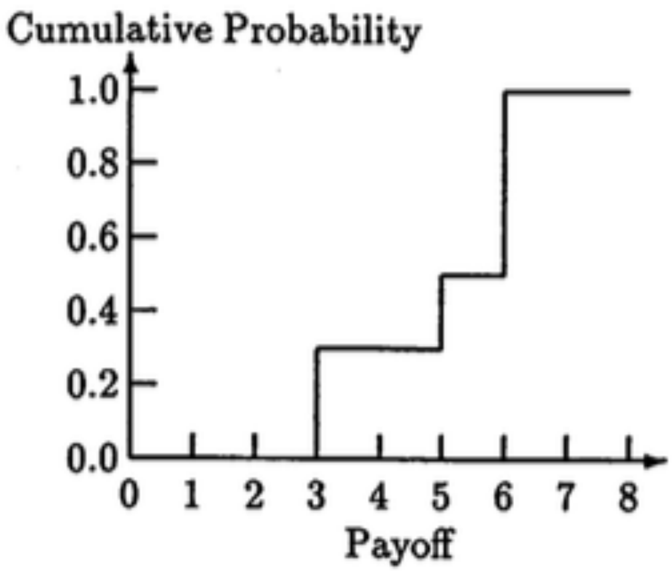
Cumulative Probability
1.0
0.8
0.6
0.4
0.2
0.0
0
1
2
3
4
5
6
7
8
Payoff

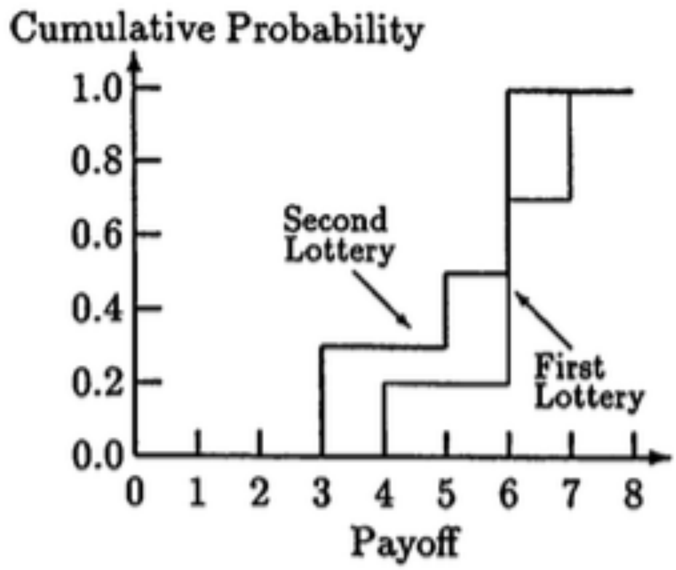
Cumulative Probability
1.0
0.8
0.6
0.4
0.2
0.0
0
1
2
3
4
5
6
7
8
Payoff
Second
Lottery
First
Lottery

Stochastic dominance is relevant in evaluating an agent's choices in the extended payoff matrix below, where purely ordinal considerations leave the agent with an ambiguity. Assume that our agent J is faced with the interaction shown in Figure 6.

		K	
		c	d
J	a	1, 4	3, 3
	b	2, 1	4, 2

$prob_K(m_J \mid m_K)$

$prob_J(m_K \mid m_J)$

	c	d
a	$.5$, $.4$	$.5$, $.4$
b	$.5$, $.6$	$.5$, $.6$

Figure 6: An Uncertain Interaction

If J considers his own potential outcomes, given the probability distribution he assumes over K's moves, he will reason that he has a .5 chance of receiving the value from either column, given any choice of his moves. Thus, if he chooses move a, he faces a .5 chance of getting either 1 or 3; if he chooses move b, he faces a .5 chance of getting either 2 or 4. Although there is no ordinal dominance here, there is stochastic dominance between the two moves, seen as two separate lotteries (Figure 7). Thus, a player who was evalua ting stochastic dominance would realize that move a was dominated.

Figure 7: Stochastic Dominance Between Two Moves' Outcomes

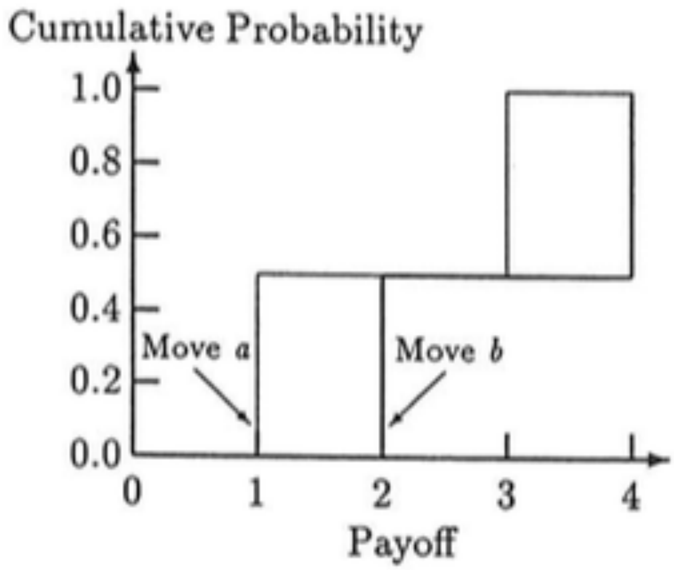
Cumulative Probability
1.0
0.8
0.6
0.4
0.2
0.0
Move a
Move b
0
1
2
3
4
Payoff

3.2.2 Formal Notation for Stochastic Dominance

Since there may be several outcomes with the same payoff to an agent, and since in the probability analysis that the agent performs these outcomes are identical, we would like to "collapse" these identical outcomes in our notation (e.g., combine *all* the chances of getting 3 into a single probability). Thus, we write agent i's subjective probability of getting a certain payoff value, given his choice of move m_i and all his knowledge of the world, as follows:

$$prob_i(v \mid m_i, \xi) = \sum_{\{m_{\bar{i}} \mid p(i,\vec{m})=v\}} prob_i(m_{\bar{i}} \mid m_i, \xi).$$

We describe this as the payoff lottery for i given move m_i. When we have

$$\forall x (\int_{-\infty} [x prob_i(v \mid c_i, \xi)\, dv < \int_{-\infty} [x prob_i(v \mid d_i, \xi)\, dv)$$

we will say that the payoff lottery for i of move d_i is stochastically dominated by the payoff lottery for i of move c_i.[7]

3.3 Utility Dominance

As opposed to ordinal or stochastic dominance, *utility dominance* employs the aggregate measure of "expected utility," which introduces a total order (with equality) over payoff lotteries. The utility function encodes the agent's attitudes toward risky or uncertain payoffs. A utility function that is linear in payoffs will result in expected value decision making.

When the following situation holds,

$$\int_{-\infty} [\infty U(v) prob_i(v \mid d_i, \xi)\, dv < \int_{-\infty} [\infty U(v) prob_i(v \mid c_i, \xi)\, dv, \qquad (3)$$

we say that the expected utility for agent i of move c_i dominates the expected utility for agent i of move d_i. Note that the dominating move is on the larger side of the inequality, in contrast to the definition of stochastic dominance, where the dominating move is on the smaller side of the inequality.

[7] For the reader being newly introduced to stochastic dominance, it might seem odd that c_i's lottery, being everywhere *less than* d_i's lottery, dominates d_i. The curves for better lotteries rise further to the right, and therefore their integrals are smaller up to any given point.

4 Rational Moves—A Prescription for an Agent

In this section we describe a prescriptive theory for rational agents in interactions. The criteria for optimality is that the agent should choose the course of action that maximizes its expected utility. However, we propose a method that makes use of alternative means of screening moves by which an agent can reduce the number of, and data requirements for, expected utility calculations.

4.1 Rationality Using Ordinal Dominance

We will denote by $R_o(p, i)$ the ordinally rational moves for the agent i in the game p. An individual agent i is said to be exhibiting *ordinal rationality* if it makes moves solely from the set $R_o(p, i)$. The following axiom defines a criterion for eliminating a move from $R_o(p, i)$:

$$pay(i, d_i) <_o pay(i, c_i) \Rightarrow d_i \notin R_o(p, i). \tag{4}$$

In other words, if d_i is ordinally dominated by c_i (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 ordinally irrational for i. Note that this does not imply that c_i *is* ordinally rational, since there may be still better moves available.

4.2 Rationality Using Stochastic Dominance

An individual agent i is said to be exhibiting *stochastic rationality* if it makes moves solely from the set $R_s(p, i)$. The following axiom defines a criterion for eliminating a move from $R_s(p, i)$:

$$\forall x(\int_{-\infty}[xprob_i(v \mid c_i, \xi)\, dv < \int_{-\infty}[xprob_i(v \mid d_i, \xi)\, dv) \Rightarrow d_i \notin R_s(p, i). \tag{5}$$

Thus, if d_i is stochastically dominated by any c_i, then d_i is stochastically irrational for agent i. Note again that this does not imply that c_i is stochastically rational—Equation 5 is a rule to exclude moves from R_s, not to prove that they are members.

4.3 Rationality Using Utility Dominance

An agent is utility rational if it seeks to maximize *expected utility*, as defined in Equation 2. The following axiom defines a criterion for eliminating a move from $R_u(p, i)$, the set of moves with maximal expected utility:

$$\int_{-\infty} [\infty U(v) prob_i(v \mid d_i, \xi)\, dv < \int_{-\infty} [\infty U(v) prob_i(v \mid c_i, \xi)\, dv \Rightarrow d_i \notin R_u(p, i). \tag{6}$$

Thus, if the expected utility of d_i is dominated by the expected utility of any c_i, then d_i is utility irrational for agent i. Note once again that this does not imply that c_i is utility rational. However, this definition of rationality differs from the previous ones in that we know there is a unique member of R_u, or a set of equivalent members (i.e., with the same expected utility). Thus this definition can actually be used to narrow the agent's choices to a single move, given the necessary computational resources to find it.

4.4 The Relationship Among Rationalities

The three definitions of rationality are related in the following ways:

$$d_i \notin R_o(p, i) \Rightarrow d_i \notin R_s(p, i) \wedge d_i \notin R_u(p, i) \tag{7}$$

$$d_i \notin R_s(p, i) \Rightarrow d_i \notin R_u(p, i) \tag{8}$$

$$R_u(p, i) \subseteq R_s(p, i) \subseteq R_o(p, i) \tag{9}$$

The fact that ordinal dominance between two moves implies stochastic dominance between the same two moves for any probability distribution is a simple consequence of their definitions. The fact that stochastic dominance implies utility dominance for any monotonic utility function is a well-known result from decision theory. Stochastic dominance is a robust measure of desirability for the agent, since moves can be eliminated no matter what the risk attitude of the agent as encoded in a utility function.

4.5 Using the Dominance Relations

We will exploit the hierarchy of rationalities (as defined in Equation 9) in our automated agent's activity. Ultimately, he would like to identify the set $R_u(p, i)$, but rather than directly trying to find the utility maximizing move, he can prune his search space by eliminating moves from R_o and R_s. Our agent therefore uses the three-level hierarchy of dominance relations, and their related rationality axioms, as follows:

1. Remove ordinally dominated moves. If a single move remains, select it and finish.

2. Assign and/or determine some properties of $prob_i(m_{\bar{i}} \mid m_i, \xi)$, the probabilities of opponents' moves given each of the agent's possible moves. We admit partial information regarding probabilities because this partial information may be sufficient to eliminate irrational moves in steps 3 and 5.

3. Remove stochastically dominated moves. If a single move remains, select it and finish.

4. Assess and apply a utility transformation to the lotteries defined by the remaining moves.

5. Remove utility dominated moves. All remaining moves will have identical expected utilities. Select one and finish.

Using this technique, our agent is able to maintain his commitment to being a utility maximizer, and still reduce the computational burden of computing expected utility. Information regarding the probability distributions of opponents' moves is used effectively. The search space can in many instances be radically pruned using this technique.

4.6 An Example

Consider an agent who is confronted with an encounter represented by the payoff matrix in Figure 8.

Using ordinal dominance, he is immediately able to rule out moves c and d, leaving him with options a and b. While neither of these moves is ordinally dominated, he would still like to choose between them. He assesses the likelihood that his opponent will make any particular move as equiprobable (perhaps his opponent is only able to reason about ordinal dominance, and thus has no dominated moves; see below, Section 5.2.1). He is then able to conclude that move b is stochastically dominated by move a; move a is thus

K

		e	f	g	h
	a	5, 1	6, 2	5, 1	7, 2
	b	4, 2	5, 1	6, 2	7, 1
J	c	4, 1	3, 2	4, 1	3, 2
	d	0, 2	1, 1	2, 2	3, 1

Figure 8: Using Ordinal and Stochastic Dominance

chosen. Had there been no stochastic dominance, he would have proceeded to compute the expected utility of moves a and b.

In general, however, the probability distributions over opponents' moves will not be readily available, and the computational burden of calculating or estimating these probabilities will overshadow the burden of calculating the best move *given* those probabilities. In the next section we describe various approaches where the axiomatic description of opponents allows the agent to deduce properties of the probability distribution for use in the framework described above.

5 Axioms of Rationality—Description

As described above, the second element of our prescriptive/descriptive approach is a descriptive theory of other agents. We will call the agent to whom we are endowing the prescriptive theory the "agent," and the other agents that we are describing as the "opponents." We describe a framework of rationality that allows us to express many levels of rationality that might be operating in opponents. The ultimate purpose of these axioms is to allow the agent to deduce $prob_i(m_{\bar{\imath}} \mid m_i, \xi)$ from more fundamental information.

In the remainder of this section, we describe various classes of rationality that this approach can address, and demonstrate how our three-tiered dominance analysis operates in each situation. Finally, we describe how to incorporate uncertainty about what axioms are present in other agents.

5.1 Minimal rationality

An assumption of minimal rationality corresponds to a situation where the agent has no information regarding the rationality of his opponents. This may include a recognition that other players are engaging in potentially irrational behavior (e. g. that the choices the opponents make are independent of their payoffs).

In this case we have

$$\{m_{\bar{\imath}} : prob_i(m_{\bar{\imath}} \mid m_i, \xi) > 0\} = M_{\bar{\imath}},$$

that is, any combined set of moves by the other agents is possible. One version of minimal rationality implies a commitment to equiprobable moves by the opponents:

$$prob_i(m_{\bar{\imath}} \mid \xi) = prob_i(m'_{\bar{\imath}} \mid \xi)$$

for all opponents' moves $m_{\bar{\imath}}$ and $m'_{\bar{\imath}}$. The effect of this is for the agent to assume that the others will be choosing their moves arbitrarily and ignoring any variation in payoffs.

Minimal rationality does not imply equiprobable assessments. Other information regarding tendencies and biases (aside from explicit consideration of payoffs) that opponents have displayed in the past can form the basis for assigning probabilities. The important point is that the assessment is not based on any explicit model of rationality of opponents. It therefore most closely resembles standard decision making under uncertainty, where uncertainty arises from lack of information and stochastic processes in the environment.

5.2 Separate rationality

In separate rationality the agent explicitly admits the possibility that each opponent is rational (to a greater or lesser degree) and has specific capabilities for reasoning about the moves others, including the agent, will take. Below, we examine several types of rationality that might conceivably be exhibited by opponents.

5.2.1 Ordinally Rational Opponents

If the agent assumes that his opponents are at most ordinally rational, then a successive winnowing process can be used to reduce the payoff matrix to a relevant set (this assumes, as well, that the opponents have knowledge of the agent's ordinal rationality; see [18]). Ordinally dominated moves, for both the agent and opponents, are repeatedly removed. The agent then restricts attention to a *reduced* matrix consisting of all ordinally undominated moves (along with opponents' responses). If the remaining set is a single entry, then there is a unique solution.

If there are multiple entries, then the opponents and the agent are left with an ambiguity—any of the moves not ordinally dominated are equally desirable. The agent can assume that the opponents will choose arbitrarily within the set of remaining moves—considering the opponents minimally rational as in Section 5.1. The agent is permitted to make this inference because the opponents are only capable of reasoning about ordinal dominance; thus further reasoning about the agent by the opponents is impossible (this was the situation exhibited in the example of Section 4.6).

There is a potential subtlety in using the above method for ordinal dominance. Consider a situation where our agent has several ordinally dominated moves; does it matter which is "removed" first from the payoff matrix? As it turns out, the order of removal, both for the agent and his opponents, is irrelevant for ordinal dominance.

5.2.2 Stochastically and Utility Rational Opponents

Here we address the issue of opponents who, like the agent, are capable of engaging in probabilistic reasoning. Since both agent and opponents can reason probabilistically about each other, there is the potential for infinite regress: the agent's choice is dependent on what he believes his opponents will do, which depends on the opponents' beliefs about the agent, and so on.

One weak form of rationality that lends itself to probabilistic reasoning is due to Strait [19]: if the agent prefers one payoff to another, then his opponent will assign a higher probability to the move with that payoff, and similarly for the agent's assessments of the opponents. One consequence of this principle is that the probability distribution over opponents' moves is dependent on the agent's move, i.e., the agent must consider $prob_i(m_{\bar{\imath}} \mid m_i, \xi)$. This dependence is not due to a causal linkage, since we are assuming simultaneous action, but rather results from the agent reasoning about the possibility of the opponents "outguessing" him given a particular move.

The foregoing principle results in a set of constraints on probabilities, given our assumption that the payoffs in the encounter are common knowledge and that all players have monotonic utility functions. There are various methods for dealing with constraints and/or bounds on probability in decision-making situations.

We can strengthen Strait's principle by adding an assumption that the opponents are Bayesian decision makers. We will restrict our attention to the case where there is a single opponent who is capable of screening moves based on both stochastic and utility dominance relationships. Furthermore, the opponents will be assumed to know that the agent is similarly an expected utility maximizer in making choices.[8]

The infinite regress of reasoning alluded to above is a real concern under these assumptions. One way of dealing with the regress is by explicitly modeling (by way of a probability distribution) the number of levels of regress that the agent believes an opponent will reason, and encoding the agent's perception of the opponent's uncertainty at each level. For example, the agent could reason that there is a fifty percent chance that the opponent will reason one level deep, a thirty percent chance two levels deep, a twenty percent chance three levels deep, and zero for all others. This is computationally complex, but is likely to be effective in a world inhabited by computationally limited reasoners.

5.3 Unique rationality

Under unique rationality, the agent assumes that the opponents' moves are fixed in advance, i.e.,

$$prob_i(m_{\bar{\imath}} \mid c, \xi) = prob_i(m_{\bar{\imath}} \mid d, \xi)$$

for all moves c and d to be made by agent i. This can also be expressed as the independence relation,

$$prob_i(m_{\bar{\imath}} \mid m_i, \xi) = prob_i(m_{\bar{\imath}} \mid \xi)$$

[8]This situation more closely resembles the jointly prescriptive theories of game theory.

which states that the agent's probability distribution does not depend on the move the agent makes. Conceptually, the opponents are assumed to have sealed away their moves before the agent makes his choice. This is orthogonal to the question of *how* the opponents will make their choices; thus, unique rationality can be combined with the various forms of separate rationality presented above, or with minimal rationality. The crucial question here is not whether the opponents are *reasoning* about the agent, but whether their moves will actually be dependent on the move made by the agent (as they are in *informed rationality* below). In certain situations, the assumption of unique rationality allows a technique called case analysis to be applied when computing ordinal dominance (see [17]).

5.4 Informed rationality

Under informed rationality, the agent assumes that opponents have perfect information—they know precisely what move the agent is to take. Informed rationality eliminates uncertainty in the encounter when payoffs are common knowledge. The agent's task in this case is to make a choice that maximizes his benefit, given that his opponents will respond omnisciently to his move. This is the situation, for example, when there is a time-lag in the making of choices, and the opponents will be able to actually *respond* to our agent's move.

5.5 Uncertainty about Rationalities

In this section we have sketched various classes of rational opponent that our prescriptively designed agent might encounter, and presented some analysis of how each case is analyzed. In general, though, an agent may be uncertain about what class of opponent he faces in a given encounter. Probability theory provides a solution—assign a probability distribution to the types of agent that might be encountered and form a composite distribution over the opponents' moves based on analysis of each case.

6 The Jointly Prescriptive Issues

Ideally, a set of agents who all use the three-level hierarchy of ordinal, stochastic and utility rationality, with coherent probability distributions, will arrive at stable solutions. In general, however, this cannot be guaranteed. Aumann [2] has shown that a construct termed *correlated equilibria* is the result of interactions between utility-maximizing agents. The equilibrium is a probabilistic notion, a generalization of the mixed randomized strategies developed by game theorists. Each agent selects a definite alternative—the uncertainty in the equilibrium is due to the joint uncertainty of the agents about other agents' moves. The existence of correlated equilibria is based on the existence of a common knowledge prior probability distribution over some underlying state of nature. Differences in probability distributions by the agents are the result of differences in information. Though Aumann has provided a characterization of equilibria, they are inherently uncertain due to the uncertainty of the participants and may in fact admit a wide range of possible solutions. Recently Nau and McCardle [16] have shown that correlated equilibria are consistent with a notion of joint coherency in noncooperative games. This work, however, has not provided an operational procedure for deriving the equilibria based on a single agent's information.

7 Conclusion

The design of automated agents can benefit from the theoretical underpinnings of decision analysis and game theory. Builders of autonomous agents will want to know that their creations are capable of adaptive behavior in the face of various opponents, and can use the "prescriptive/descriptive" aspects of decision analysis to guide their agents' design. The builders of full multiagent systems will want to ensure certain desirable global properties, and can use the "jointly prescriptive" aspects of game theory to choose the agents' built-in strategies.

We have presented a technique that exploits the relationship among ordinal, stochastic, and utility dominance. Combining it with logical axioms that describe opponents, it is particularly suitable for a deductive engine to use in deciding on a move in an interaction. The technique is based on computational considerations, pruning certain moves before performing computationally expensive operations (such as finding expected utility). We have also presented a sampling of rationality axioms that might be useful to an agent's designer, and given some ramifications of their use. This is basically a prescriptive analysis, discussing one way in which an interacting intelligent agent could be built.

References

[1] Douglas E. Appelt. *Planning natural language utterances to satisfy multiple goals*. Tech Note 259, SRI International, Menlo Park, California, 1982.

[2] Robert J. Aumann. Correlated equilibrium as an expression of Bayesian rationality. *Econometrica*, 55(1):1-18, January, 1987.

[3] John S. Breese. *Knowledge Representation and Inference in Intelligent Decision Systems*. PhD thesis, Stanford University, 1987. Also published as Research Report 2, Rockwell International Science Center, Palo Alto Lab, Palo Alto, California, April 1987.

[4] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosens chein. *Solving the Prisoner's Dilemma*. Report No. STAN-CS-84-1032, Computer Science Department, Stanford University, November 1984. Also published as HPP-84-41, Heuristic Programming Project, Computer Science Department, Stanford University, November 1984.

[5] Michael R. Genesereth, Matthew L. Ginsberg, and Jeffrey S. Rosens chein. Cooperation without communication. In *Proceedings of The National Conference on Artificial Intelligence*, pages 51–57, The American Association for Artificial Intelligence, Philadelphia, Pennsylvania, August 1986.

[6] Matthew L. Ginsberg. Decision procedures. In Michael N. Huhns, editor, *Distributed Artificial Intelligence*, pagesc3–28, Morgan Kaufmann Publishers, Inc., Los Altos, California, 1987.

[7] Joseph Y. Halpern and Yoram Moses. *Knowledge and Common Knowledge in a Distributed Environment*. Research Report IBM RJ 4421, IBM

Research Laboratory, San Jose, California, October 1984. Also published in *Proceedings of the Third Annual ACM Conference on Principles of Distributed Computing*, Vancouver, British Columbia, Canada, 1984.

[8] Eric J. Horvitz, John S. Breese, and Max Henrion. Decision theory in expert systems and artificial intelligence. *International Journal of Approximate Reasoning*, 2:247-302,1988.

[9] Joseph B. Kadane and Patrick D. Larkey. Subjective probability and the theory of games. *Management Science*, 28(2):113-120, February, 1982.

[10] JosephcB. Kadane and Patrick D. Larkey. The confusion of is and ought in game theoretic contexts. *Management Science*, 29(12):1365-1379, December, 1983.

[11] R. Duncan Luce and Howard Raiffa. *Games and Decisions, Introduction and Critical Survey.* John Wiley and Sons, New York, 1957.

[12] Kurt Konolige. *A first-order formalization of knowledge and action for a multi-agent planning system.* Tech Note 232, SRI International, Menlo Park, California, Decembe r 1980.

[13] Kurt Konolige. *A Deduction Model of Belief and its Logics.* PhD thesis, Stanford University, 1984.

[14] Robert C. Moore. *Reasoning about knowledge and action.* Tech Notec191, SRI International, Menlo Park, California, 1980.

[15] Robert C. Moore. *A formal theory of knowledge and action.* Tech Note 320, SRI International, Menlo Park, California, 1984. Also in *Formal Theories of the Commonsense World*, Hobbs, J.R., and Moore, R.C. (Eds.), Ablex Publishing Co. (1985).

[16] Robert F. Nau and Kevin F. McCardle. *Coherent Behavior in Noncooperative Games.* Working Paper 8701, The Fuqua School of Business, Duke University, Durham, North Carolina, 1988.

[17] Jeffrey S. Rosenschein. *Rational Interaction: Cooperation Among Intelligent Agents.* PhD thesis, Stanford University, 1986. Also published as STAN-CS-85-1081 (KSL-85-40), Department of Computer Science, Stanford University, October 1985.

[18] Jeffrey S. Rosenschein. The role of knowledge in logic-based rational interaction. In *Proceedings of The IEEE Phoenix Conference on Computers and Communication*, IEEE, Scottsdale, Arizona, March 1988, pp. 497–504. Also published as Technical Memo 1-1987, Rockwell International Science Center, Palo Alto Lab, Palo Alto, California, 1987.

[19] R. Scott Strait. *Decision Analysis of Strategic Interaction*. PhD thesis, Stanford University, 1987. Also published as UCRL-53825, Lawrence Livermore National Laboratory, Livermore, California, December, 1987.

[20] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, 1947.

[21] John Wilson. Subjective probability and the prisoner's dilemma. *Management Science*, 22(1):45-55, January, 1986.

[22] G. A. Whitmore and M. C. Findlay, editors. *Stochastic Dominance: An Approach to Decision Making Under Risk*. D. C. Heath and Company, Lexington, Massachusetts, 1978.