Decisions and Games for BD Agents

Mehdi Dastani
Institute of Information and Computing Sciences
Utrecht University
The Netherlands
email: mehdi@cs.uu.nl

Leendert van der Torre
Department of Artificial Intelligence
Vrije Universiteit Amsterdam
The Netherlands
email: torre@cs.vu.nl

Abstract

Strategic games model the interaction between simultaneous decisions of a group of agents. The starting point of strategic games is a set of players (agents) having certain strategies (decisions) and preferences on the game's outcomes. In this paper we do not assume the set of decisions and preferences of agents to be given, but derive them from their mental attitudes. In particular, we introduce a rule-based architecture for agents with beliefs and desires and explain how their decisions and preferences can be derived. We specify groups of such agents, define a mapping from their specification to the specification of the game they play, and use some familiar notions from game theory, such as Pareto efficiency and Nash equilibrium, to characterize the interaction between their decisions. We also discuss a reverse mapping from the specification of games that a group of agent play to the specifications of those agents. This mapping can be used to specify groups of agents that can play a ceratin game.

Introduction

One of the main problems in agent theory is the distinction between formal theories and tools developed for individual autonomous agents, and theories and tools developed for multi agent systems. In the social sciences, this distinction is called the micro-macro dichotomy. The prototypical example is the distinction between classical decision theory based on the expected utility paradigm (usually identified with the work of Neumann and Morgenstern (von Neumann & Morgenstern 1944) and Savage (Savage 1954)) and classical game theory (such as the work of Nash and more recently the work of Axelrod). Whereas classical decision theory is a kind of optimization problem (maximizing the agent's expected utility), classical game theory is a kind of equilibria analysis.

There are several approaches in artificial intelligence, cognitive science, and practical reasoning (within philosophy) to describe the decision making of individual agents and to bring the micro and macro descriptions together. In these approaches, the decision making of individual autonomous agents is described in terms of other concepts than maximizing utility. For example, since the early 40s there is a distinction between classical decision theory and artificial intelligence based on utility aspiration levels and goal based planning (as pioneered by Simon (Simon 1981)). Moreover, in cognitive science and philosophy the decision making of individual agents is described in terms of concepts from folk psychology like beliefs, desires and intentions. In these studies, the decision making of individual agents is characterized in terms of a rational balance between these concepts, and the decision making of a group of agents is described in terms of concepts generalized from those used for individual agents, such as joint goals, joint intentions, joint commitments, etc. Moreover, also new concepts are introduced at this social level, such as norms (a central concept in most social theories).

It is still an open problem how the micro-macro dichotomy of classical decision and game theory is related to the micro-macro dichotomy of these alternative theories. Has or can the micro and macro level be brought together by replacing classical theories by alternative theories? Before this question can be answered, the relation between the classical and alternative theories has to be clarified. Doyle and Thomason (Doyle & Thomason summer 1999) argue that classical decision theory should be reunited with alternative decision theories in so-called qualitative decision theory, which studies qualitative versions of classical decision theory, hybrid combinations of quantitative and qualitative approaches to decision making, and decision making in the context of artificial intelligence applications such as planning, learning and collaboration. Qualitative decision theories have been developed based on beliefs (probabilities) and desires (utilities) using formal tools such as modal logic (Boutilier 1994) and on utility functions and knowledge (Lang 1996). More recently these beliefs-desires models have been extended with intentions or BDI models (Cohen & Levesque 1990; Rao & Georgeff 1991a; 1991b).

In this paper, we introduce a rule based qualitative decision and game theory for agents with beliefs and desires. Like classical decision theory but in contrast to several proposals in the BDI approach (Cohen & Levesque 1990; Rao & Georgeff 1991a), the theory does not incorporate decision processes, temporal reasoning, and scheduling. We also ignore probabilistic decisions and related games with mixed strategies. In particular, we explain how decisions and preferences of individual agents can be derived from their beliefs and desires. We specify groups of agents and discuss the interaction between their decisions and preferences. This enables us to define a mapping from agent specification to the specification of the game that they play. This mapping considers agent decisions as agent strategies and the consequences of decisions as the outcomes of the game. Finally, we discuss the reverse mapping from the specification of a game to the specification of the agents that play the game. This mapping provides the beliefs and desires of agents that can play a certain game.

A qualitative decision and game theory

The qualitative decision and game theory introduced in this section is developed for agents that have conditional beliefs and desires, i.e. the beliefs and desires of agents are sets of rules. The architecture and the behavior of this type of agents are studied in (Broersen *et al.* 2001b; Dastani & van der Torre 2002). Here, we analyze this type of agents from a decision and game theoretic point of view by studying possible decisions of individual agents and the interaction between these decisions. We do so by defining an agent system specification that indicates possible decisions of agents. We show how we can derive agent decisions and preferences from an agent system specification. First we consider the logic of rules we adopt.

Logic of rules

The starting point of any theory of decision is a distinction between choices made by the decision maker and choices imposed on it by its environment. For example, a software upgrade agent (decision maker) may have the choice to upgrade a computer system at a particular time of the day. The software company (the environment) may in turn allow/disallow such a upgrade at a particular time. Let $S = \{\alpha_1, \ldots, \alpha_n\}$ be the society or set of agents, then we therefore assume n disjoint sets of propositional atoms: $A = A_1 \cup \ldots \cup A_n = \{a, b, c, \ldots\}$ (agents' decision variables (Lang 1996) or controllable propositions (Boutilier 1994)) and $W = \{p, q, r, \ldots\}$ (the world parameters or uncontrollable propositions). In the sequel we consider each decision maker as entities consisting of rules. Such a decision maker generates its decisions by applying subsets of rules to its input. Before we proceed some notations will be introduced.

- L_{A_i}, L_W and L_{AW} for the propositional languages built up from these atoms in the usual way, and variables x, y, $\ldots$ to stand for any sentences of these languages.
- Cn_{A_i}, Cn_W and Cn_{AW} for the consequence sets, and $\models_{A_i}$, $\models_W$ and $\models_{AW}$ for satisfiability, in any of these propositional logics.
- $x \Rightarrow y$ for an ordered pair of propositional sentences called a rule.
- $E_R(T)$ for the R extension of T, as defined in Definition 1 below.

In our framework the generation of decisions are formalized based on the notion of extension. In particular, the decision of an agent, which is specified by a set of rules R and has the input T, is the extension calculated based on R and T. This is formalized in the following definition.

Definition 1 (Extension) *Let $R \subseteq L_{AW} \times L_{AW}$ be a set of rules and $T \subseteq L_{AW}$ be a set of sentences. The consequents of the T-applicable rules are:*

$$R(T) = \{y \mid x \Rightarrow y \in R, x \in T\}$$

and the R extension of T is the set of the consequents of the iteratively T-applicable rules:

$$E_R(T) = \cap_{T \subseteq X, R(Cn_{AW}(X)) \subseteq X} X$$

We give some properties of the R extension of T in Definition 1. First note that $E_R(T)$ is *not* closed under logical consequence. The following proposition shows that $E_R(T)$ is the smallest superset of T closed under the rules R interpreted as inference rules.

Proposition 1 *Let*

- $E_R^0(T) = T$
- $E_R^i(T) = E_R^{i-1}(T) \cup R(Cn_{AW}(E_R^{i-1}(T)))$ *for* $i > 0$

We have $E_R(T) = \cup_0^\infty E_R^i(T)$.

This definition of extensions allows inconsistent decisions as illustrated by the following example.

Example 1 *Let $R = \{\top \Rightarrow p, a \Rightarrow \neg p\}$ and $T = \{a\}$, where $\top$ stands for any tautology like $p \vee \neg p$. We have $E_R(\emptyset) = \{p\}$ and $E_R(T) = \{a, p, \neg p\}$, i.e. the R extension of T is inconsistent.*

Of course, allowing inconsistent decisions may not be intuitive. We are here concerned about possible decisions rather than reasonable or feasible decisions. Later we will define reasonable or feasible decisions by excluding inconsistent decisions.

Agent system specification

An agent system specification given in Definition 2 contains a set of agents and for each agent a description of its decision problem. The agent's decision problem is defined in terms of its beliefs and desires, which are considered as belief and desire rules, a priority ordering on the desire rules, as well as a set of facts and an initial decision (or prior intentions). We assume that agents are autonomous, in the sense that there are no priorities between desires of distinct agents.

Definition 2 (Agent system specification) *An agent system specification is a tuple $AS = \langle S, F, B, D, \geq, \delta^0 \rangle$ that contains a set of agents S, and for each agent α_i a finite set of facts $F_i \subseteq L_W$ ($F = \bigcup_{i=1}^n F_i$), a finite set of belief rules $B_i \subseteq L_{AW} \times L_W$ ($B = \bigcup_{i=1}^n B_i$), a finite set of desire rules $D_i \subseteq L_{AW} \times L_{AW}$ ($D = \bigcup_{i=1}^n D_i$), a relation $\geq_i \subseteq D_i \times D_i$ ($\geq = \bigcup_{i=1}^n \geq_i$) which is a total ordering (i.e. reflexive, transitive, and antisymmetric and for any two elements d_1 and d_2 in D_i, either $d_1 \geq_i d_2$ or $d_2 \geq_i d_1$), and a finite initial decision $\delta_i^0 \subseteq L_A$ ($\delta^0 = \bigcup_{i=1}^n \delta_i^0$). For an agent $\alpha_i \in S$ we write $x \Rightarrow_i y$ for one of its rules.*

In general, a belief rule is an ordered pair $x \Rightarrow_i y$ with $x \in L_{AW}$ and $y \in L_W$. This belief rule should be interpreted as 'the agent α_i believes y in context x'. A desire rule is an ordered pair $x \Rightarrow_i y$ with $x \in L_{AW}$ and $y \in L_{AW}$. This desire rule should be interpreted as 'the agent desires y in context x'. It implies that the agent's beliefs are about the world ($x \Rightarrow_i p$), and not about the agent's decisions. These beliefs can be about the effects of decisions made by the agent ($a \Rightarrow_i p$) as well as beliefs about the effects of parameters set by the world ($p \Rightarrow_i q$). Moreover, the agent's desires can be about the world ($x \Rightarrow_i p$, desire-to-be), but also about the agent's decisions ($x \Rightarrow_i a$, desire-to-do). These desires can be triggered by parameters set by the world ($p \Rightarrow_i y$) as well as by decisions made by the agent ($a \Rightarrow_i y$). Modelling mental attitudes such as beliefs and desires in terms of rules results in what might be called conditional mental attitudes (Broersen *et al.* 2001b).

Agent Decisions

The belief rules of the form $L_A \Rightarrow L_W$ can be used to determine the expected consequences of a decision, where a decision δ_i of agent α_i is any subset of L_{A_i} that contains the initial decision δ_i^0. The set of expected consequences of this decision δ_i is the belief extension of $F_i \cup \delta_i$. Moreover, we consider a feasible decision as a decision that does not imply a contradiction.

Definition 3 (Decisions) *Let $AS = \langle\{\alpha_1, \ldots, \alpha_n\}, F, B, D, \geq, \delta^0\rangle$ be an agent system specification. An AS decision profile δ for agents $\alpha_1, \ldots, \alpha_n$ is $\delta = \langle\delta_1, \ldots \delta_n\rangle$ where δ_i is a decision of agent α_i such that*

$$\delta_i^0 \subseteq \delta_i \subseteq L_{A_i} \;\; \text{for } i = 1 \ldots n$$

A feasible decision for agent α_i is δ_i such that

$$E_{B_i}(F_i \cup \delta_i) \text{ is consistent}$$

A feasible decision profile is a decision profile such that

$$E_B(F \cup \delta) \text{ is consistent}$$

The following example illustrates the decisions of a single agent.

Example 2 *Let $A_1 = \{a, b, c, d, e\}$, $W = \{p, q\}$ and $AS = \langle\{\alpha_1\}, F, B, D, \geq, \delta^0\rangle$ with $F_1 = \{\neg p\}$, $B_1 = \{c \Rightarrow q, d \Rightarrow q, e \Rightarrow \neg q\}$, $D_1 = \{\top \Rightarrow a, \top \Rightarrow b, b \Rightarrow p, \top \Rightarrow q, d \Rightarrow q\}$, $\geq_1 = \{b \Rightarrow p > \top \Rightarrow b\}$, and $\delta_1^0 = \{a\}$. The initial decision δ_1^0 reflects that the agent has already decided in an earlier stage to reach the desire $\top \Rightarrow a$. Note that the consequents of all B_1 rules are sentences of L_W, whereas the antecedents of the B_1 rules as well as the antecedents and consequents of the D_1 rules are sentences of L_{AW}. We have due to the definition of $E_R(S)$:*

$E_B(F \cup \{a\}) = \{\neg p, a\}$
$E_B(F \cup \{a, b\}) = \{\neg p, a, b\}$
$E_B(F \cup \{a, c\}) = \{\neg p, a, c, q\}$
$E_B(F \cup \{a, d\}) = \{\neg p, a, d, q\}$
$E_B(F \cup \{a, e\}) = \{\neg p, a, e, \neg q\}$
$\ldots$
$E_B(F \cup \{a, d, e\}) = \{\neg p, a, d, e, q, \neg q\}$

Therefore $\{a, d, e\}$ is not a feasible AS decision profile, because its belief extension is inconsistent. Continued in Example 5.

According to definition 3, the feasibility of a decision profile, which indicates the decisions of individual agents, is formulated in terms of the consistency of the extension that is calculated based on the decision profile. However, there are two different ways to calculate such an extension depending on whether agents can apply rules of each other or not. First, the extension can be calculated by demanding that each agent generates its decisions by applying only its own belief and desire rules. This way of calculating the extension assumes that mental attitudes of agents are private and inaccessible to each other. We will call this notion of feasible decision profile individualistic. Second, the extension can be calculated by allowing agents to apply belief and desire rules of each other. This notion of feasible decision profile, which will be called collective, assumes that mental attitudes of agents are public.

Definition 4 *Feasible decision profiles can be defined in the following two ways:*

- *An individualistic feasible decision profile δ is a decision profile if $E_B(F \cup \delta) = \bigcup_{i=1}^{n} E_{B_i}(F_i \cup \delta_i)$ is consistent.*
- *A collective feasible decision profile δ is a decision profile if $E_B(F \cup \delta) = E_{B_1 \cup \ldots \cup B_n}(F_1 \cup \delta_1 \cup \ldots \cup F_n \cup \delta_n)$ is consistent.*

The following example illustrates individualistic and collective feasible decision profiles.

Example 3 *Let $A_1 = \{a\}$, $A_2 = \{b\}$, $W = \{p\}$ and $AS = \langle\{\alpha_1, \alpha_2\}, F, B, D, \geq, \delta^0\rangle$ with $F_1 = F_2 = \emptyset$, $B_1 = \{a \Rightarrow p\}$, $B_2 = \{b \Rightarrow \neg p\}$, $D_1 = D_2 = \emptyset$, $\geq$ is the identity relation, and $\delta_1^0 = \delta_2^0 = \emptyset$. The decision profile $\delta = \langle\delta_1, \delta_2\rangle = \langle\emptyset, \{b\}\rangle$ is both individualistic and collective since $E_{B_1 \cup B_2}(F_1 \cup \delta_1 \cup F_2 \cup \delta_2) = E_{B_1}(F_1 \cup \delta_1) \cup E_{B_2}(F_2 \cup \delta_2) = \{b, \neg p\}$. But, the decision profile $\delta = \langle\delta_1, \delta_2\rangle = \langle\{a\}, \{b\}\rangle$ is neither individualistic nor collective since $E_{B_1 \cup B_2}(F_1 \cup \delta_1 \cup F_2 \cup \delta_2) = E_{B_1}(F_1 \cup \delta_1) \cup E_{B_2}(F_2 \cup \delta_2) = \{a, p, b, \neg p\}$.*

Finally, the following example illustrates that a decision profile can be individualistic but not collective feasible decision profile.

Example 4 *Let $A_1 = \{a\}$, $A_2 = \{b\}$, $W = \{p, t\}$ and $AS = \langle\{\alpha_1, \alpha_2\}, F, B, D, \geq, \delta^0\rangle$ with $F_1 = F_2 = \emptyset$, $B_1 = \{a \Rightarrow p\}$, $B_2 = \{b \Rightarrow t, p \Rightarrow \neg t\}$, $D_1 = D_2 = \emptyset$, $\geq$ is the identity relation, and $\delta_1^0 = \delta_2^0 = \emptyset$. The decision profile $\delta = \langle\delta_1, \delta_2\rangle = \langle\{a\}, \{b\}\rangle$ is individualistic feasible since $E_{B_1}(F_1 \cup \delta_1) \cup E_{B_2}(F_2 \cup \delta_2) = \{a, p\} \cup \{b, t\} = \{a, p, b, t\}$. However, δ is not a collective feasible decision profile since $E_{B_1 \cup B_2}(F_1 \cup \delta_1 \cup F_2 \cup \delta_2) = \{a, p, b, t, \neg t\}$.*

In the following, we use the definition of individualistic feasible decision profiles.

Agent preferences

In this section we introduce a way to compare decisions. We compare decisions by comparing sets of desire rules that are not reached by the decisions. Since an agent system specification provides only the ordering on individual desire rules, and not ordering on sets of desire rules, we first lift the ordering on individual desire rules to an ordering on sets of desire rules. In this way, we make it possible to compare sets of desire rules using the preference on individual desire rules.

Definition 5 *Let $AS = \langle S, F, B, D, \geq, \delta^0 \rangle$ be an agent system specification, D'_i, D''_i two subsets of D_i, and $D_i \setminus D'_i$ the set of D_i elements which are not D'_i elements. We have $D'_i \succeq D''_i$ if $\forall d'' \in D''_i \setminus D'_i \exists d' \in D'_i \setminus D''_i$ such that $d' > d''$. We write $D'_i \succ D''_i$ if $D'_i \succeq D''_i$ and $D''_i \not\succeq D'_i$, and we write $D'_i \simeq D''_i$ if $D'_i \succeq D''_i$ and $D''_i \succeq D'_i$.*

This definition states that an agent α prefers a set of desire rules D' to another set of desire rules D'' if for each desire rule in D'' there exists a preferable desire rule in D'. Thus, if D' contains only one desire rule which is preferable by α to any desire rule from D'', then α prefers D' to D''. The following propositions show that the priority relation $\succeq$ is reflexive, anti-symmetric, and transitive.

Proposition 2 *Let $D_1 \not\subseteq D_3$ and $D_3 \not\subseteq D_1$. For finite sets, the relation $\succeq$ is reflexive ($\forall D\ D \succeq D$), anti-symmetric $\forall D_1, D_2\ (D_1 \succeq D_2\ \wedge D_2 \succeq D_1) \rightarrow D_1 = D_2$, and transitive, i.e. $D_1 \succeq D_2$ and $D_2 \succeq D_3$ implies $D_1 \succeq D_3$.*

The desire rules are used to compare the decisions. The comparison is based on the set of unreached desires and not on the set of violated or reached desires. A desire $x \Rightarrow y$ is unreached by a decision if the expected consequences of this decision imply x but not y. The desire rule is violated or reached if these consequences imply $x \wedge \neg y$ or $x \wedge y$, respectively.

Definition 6 (Comparing decisions) *Let $AS = \langle S, F, B, D, \geq, \delta^0 \rangle$ be an agent system specification and δ be a AS decision. The unreached desires of decision δ for agent α_i are:*

$U_i(\delta) =$
$\{x \Rightarrow y \in D_i \mid E_B(F \cup \delta) \models x \text{ and } E_B(F \cup \delta) \not\models y\}$

Decision δ is at least as good as decision δ' for agent α_i, written as $\delta \geq^U_i \delta'$, iff

$$U_i(\delta') \succeq U_i(\delta)$$

Decision δ dominates decision δ' for agent α_i, written as $\delta >^U_i \delta'$, iff

$$\delta \geq^U_i \delta' \text{ and } \delta' \not\geq^U_i \delta$$

Thus, a decision δ is preferred by agent α to decision δ' if the set of unreached desire rules by δ is less preferred by α to the set of unreached desire rules by δ', i.e. the unreached desire rules by δ are less important than unreached desire rules by δ'. The following continuation of Example 2 illustrates the comparison of decisions.

Example 5 (Continued) *We have:*
$U(\{a\}) = \{\top \Rightarrow b, \top \Rightarrow q\}$,
$U(\{a, b\}) = \{b \Rightarrow p, \top \Rightarrow q\}$,
$U(\{a, c\}) = \{\top \Rightarrow q\}$,
$U(\{a, d\}) = \{\top \Rightarrow b, d \Rightarrow q\}$,
$U(\{a, e\}) = \{\top \Rightarrow b, \top \Rightarrow q\}$,
$U(\{a, b, c\}) = \{b \Rightarrow p\}$.
...
We thus have for example that the decision $\{a, c\}$ dominates the initial decision $\{a\}$, i.e. $\{a, c\} >^U \{a\}$. There are two decisions for which their set of unreached contains only one desire. Due to the priority relation, we have that $\{a, c\} >^U \{a, b, c\}$.

From Agent System Specifications to Game Specifications

In this subsection, we consider agents interactions based on agent system specifications, their corresponding agent decisions, and the ordering on the decisions as explained in previous subsections. Game theory is the usual tool to model the interaction between self-interested agents. Agents select optimal decisions under the assumption that other agents do likewise. This makes the definition of an optimal decision circular, and game theory therefore restricts its attention to equilibria. For example, a decision is a Nash equilibrium if no agent can reach a better (local) decision by changing its own decision. The most used concepts from game theory are Pareto efficient decisions, dominant decisions and Nash decisions. We first repeat some standard notations from game theory (Binmore 1992; Osborne & Rubenstein 1994).

As mentioned, we use δ_i to denote a decision of agent α_i and $\delta = \langle \delta_1, \ldots, \delta_n \rangle$ to denote a decision profile containing one decision for each agent. δ_{-i} is the decision profile of all agents except the decision of agent α_i. (δ_{-i}, δ'_i) denotes a decision profile which is the same as δ except that the decision of agent i from δ is replaced with the decision of agent i from δ'. $\delta'_i >^U_i \delta_i$ denotes that decision δ'_i is better than δ_i according to his preferences $>^U_i$ and $\delta'_i \geq^U_i \delta_i$ if better or equal. Δ is the set of all decision profiles for agents $\alpha_1, \ldots, \alpha_n$, $\Delta_f \subseteq \Delta$ is the set of feasible decision profiles, and Δ^i is the set of possible decisions for agent α_i.

Definition 7 (Game Specification) *A game specification is a tuple $\langle S, \Delta, (\geq_i) \rangle$ where*

- *S is a set (of agents)*
- *$\Delta = \Delta_1 \times \ldots \times \Delta_n$ where Δ_i is a set of decisions of agent α_i*
- *$(\geq_i)$ is a partial order on Δ (agent i's preferences)*

We now define a mapping $AS2AG$ from an agent specification to a game specification.

Definition 8 (AS2AG) *Let $AS = \langle S = \{\alpha_1, \ldots, \alpha_n\}, F, B, D, \geq, \delta^0 \rangle$ be specification of agent system in S, A_i be the set of AS feasible decisions of agent α_i according to Definition 3, $\Delta = A_1 \times \ldots \times A_n$, and $\geq^U_i$ be the AS preference relation of agent α_i defined on Δ according to definition 6. Then, the game specification of AS is the tuple $\langle S, \Delta, (\geq^U_i) \rangle$.*

We now consider different types of decision profiles which are similar to types of strategy profiles from game theory.

Definition 9 *Let Δ_f be the set of feasible decision profiles. A PS decision profile $\delta = \langle \delta_1, \ldots, \delta_n \rangle \in \Delta_f$ is:*

Pareto decision *if there is no $\delta' = \langle \delta'_1, \ldots, \delta'_n \rangle \in \Delta_f$ for which $\delta'_i >^U_i \delta_i$ for all agents α_i.*

strongly Pareto decision *if there is no $\delta' = \langle \delta'_1, \ldots, \delta'_n \rangle \in \Delta_f$ for which $\delta'_i \geq^U_i \delta_i$ for all agents α_i and $\delta'_j >^U \delta_j$ for some agents α_j.*

dominant decision *if for all $\delta' \in \Delta_f$ and for every agent i it holds: $(\delta'_{-i}, \delta_i) \geq^U_i (\delta'_{-i}, \delta'_i)$ i.e. a decision is dominant if it yields a better payoff than any other decisions regardless of what the other agents decide.*

Nash decision *if for all agents i it holds: $(\delta_{-i}, \delta_i) \geq^U_i (\delta_{-i}, \delta'_i)$ for all $\delta'_i \in \Delta^i_f$*

It is a well known fact that Pareto decisions exist (for finite games), whereas dominant decisions do not have to exist. The latter is illustrated by the following example.

Example 6 *Let α_1 and α_2 be two agents, $F_1 = F_2 = \emptyset$, and initial decisions $\delta^0_1 = \delta^0_2 = \emptyset$. They have the following beliefs en desires:*
$B_{\alpha_1} = \{a \Rightarrow p, \neg a \Rightarrow \neg p\}$
$D_{\alpha_1} = \{\top \Rightarrow p, \top \Rightarrow q\}$
$\geq_{\alpha_1} = \top \Rightarrow p > \top \Rightarrow q > \top \Rightarrow \neg q > \top \Rightarrow \neg p$
$B_{\alpha_2} = \{b \Rightarrow q, \neg b \Rightarrow \neg q\}$
$D_{\alpha_2} = \{\top \Rightarrow \neg p, \top \Rightarrow \neg q\}$
$\geq_{\alpha_2} = \top \Rightarrow \neg p > \top \Rightarrow \neg q > \top \Rightarrow q > \top \Rightarrow p$

Let Δ_f be feasible decision profiles, E_B be the outcomes of the decisions, and $U(\delta_i)$ be the set of unreached desires for agent α_i.

Δ_f	E_B	U_{δ_1}	U_{δ_2}
$\langle a, b \rangle$	$\{p, q\}$	$\emptyset$	$\{\top \Rightarrow \neg p, \top \Rightarrow \neg q\}$
$\langle a, \neg b \rangle$	$\{p, -q\}$	$\{\top \Rightarrow q\}$	$\{\top \Rightarrow \neg p\}$
$\langle \neg a, b \rangle$	$\{\neg p, q\}$	$\{\top \Rightarrow p\}$	$\{\top \Rightarrow \neg q\}$
$\langle \neg a, \neg b \rangle$	$\{\neg p, \neg q\}$	$\{\top \Rightarrow p, \top \Rightarrow q\}$	$\emptyset$

According to definition 5, for α_1 :
$U(\langle a, b \rangle) > U(\langle a, \neg b \rangle) > U(\langle \neg a, b \rangle) > U(\langle \neg a, \neg b \rangle)$
and for α_2 :
$U(\langle \neg a, \neg b \rangle) > U(\langle \neg a, b \rangle) > U(\langle a, \neg b \rangle) > U(\langle a, b \rangle)$.
None of these decision profiles are dominant decisions, i.e. the agents specifications has no dominant solution with respect to their unreached desires.

The following example illustrates a typical cooperation game.

Example 7 $B_1 = \{a \Rightarrow p\ ,\ b \Rightarrow \neg p \wedge q\}$, $D_1 = \{\top \Rightarrow p \wedge q\}$ $B_2 = \{c \Rightarrow q\ ,\ d \Rightarrow p \wedge \neg q\}$ $D_2 = \{\top \Rightarrow p \wedge q\}$. *The agents have a common goal $p \wedge q$, which they can only reach by cooperation.*

The following example illustrates a qualitative version of the notorious prisoner's dilemma, where the selfish behavior of individual autonomous agents leads to global bad decisions.

Example 8 *Let $A_1 = \{a\}$ (α_1 cooperates), $A_2 = \{b\}$ (α_2 cooperates), and AS be an agent system specification with* $D_1 = \{\top \Rightarrow \neg a \wedge b, \top \Rightarrow b, \top \Rightarrow \neg(a \wedge \neg b)\}$, $D_2 = \{\top \Rightarrow a \wedge \neg b, \top \Rightarrow a, \top \Rightarrow \neg(\neg a \wedge b)\}$. *The only Nash decision is $\{\neg a, \neg b\}$, whereas both agents would prefer $\{a, b\}$.*

Starting from an agent system specification, we can derive the game specification and in this game specification we can use standard techniques to for example find the Pareto decisions. However, the problem with this approach is that the translation from an agent system specification to a game specification is computationally expensive. For example, a compact agent representation with only a few belief and desire rules may lead to a huge set of decisions if the number of decision variables is high.

The main challenge of qualitative game theory is therefore whether we can bypass the translation to game specification, and define properties directly on the agent system specification. For example, are there particular properties of agent system specification for which we can prove that there always exists a dominant decision for its corresponding derived game specification? A simple example is an agent system specification in which each agent has the same belief and desire rules.

From Game Specifications to Agent Specifications

In the previous section a mapping from agent system specifications to standard game specifications has been defined. In this section the question is raised whether the notion of agent system specification is expressive enough for game specifications, that is, whether for each possible game specification there is an agent system specification that can be mapped on it. We prove this property in the following standard way:

1. First, we define a mapping from game specifications to agent system specifications;
2. Second, we show that the composite mapping from game specifications to agent system specifications and back to game specifications is the identity relation.

The second step shows that if a game specification GS is mapped in step 1 on agent system specification AS, then this agent system specification AS is mapped on GS. Thus, it shows that there is a mapping from agent system specifications to game specifications for every game specification GS.

Unlike the second step, the composite mapping from agent system specifications to game specifications and back to agent system specifications is not the identity relation. This is a direct consequence of the fact that there are distinct agent system specifications that are mapped on the same game specification. For example, agent system specifications in which the variable names are uniformly substituted by new names.

The mapping from game specifications to agent system specifications consists of the following steps:

1. Copy the set of agents;
2. Introduce a set of decision variables and initial decisions that will generate the set of decisions;

3. Introduce a set of parameters, and for each agent sets of initial facts, beliefs, and desires, such that they generate the preference order for each agent.

We start with the set of decision variables. We introduce for each decision a separate decision variable. Moreover, we assume initial decisions that see to it that the agent selects at least one decision variable, and at most one decision variable.

Definition 10 (Decisions) *Let $\Delta = \{\delta_1, \ldots, \delta_n\}$ be a set of decisions (of some agent). Then we introduce the set of decision variables $A = \{d_1, \ldots, d_n\}$ (for this agent) and the following initial decision (for this agent):*

$$d_1 \vee \ldots \vee d_n$$

$$i \neq j \Rightarrow \neg(d_i \wedge d_j)$$

Moreover, we define the sets of desires.

Definition 11 *Consider a set of decisions with a partial ordering on it (for an agent). For each decision δ, for which we introduced the decision variable d, we introduce the following desire:*

$$\top \Rightarrow \bigvee_{\delta' \geq \delta} d' \wedge \bigwedge_{\delta \not\geq \delta' \text{ and } \delta' \not\geq \delta} \neg d'$$

The following proposition shows that the mapping from game specifications to agent system specifications leads to the desired identity relation for the composite relation.

Proposition 3 *Let GS be a game specification as in Definition 7. Moreover, let AS be an agent system specification such that the decision variables, the initial decisions and the desires are as defined in Definition 10 and 11, and let the set of parameters, the sets of beliefs, the set observations be empty. The mapping from this agent system specification to game specification defined in Definition 8 maps AS to GS.*

Proof. Assume any particular GS. Construct the AS as above. Now apply Definition 6. The unreached desires of decision δ for agent α_i are:

$$U_i(\delta) = \{x \Rightarrow y \in D \mid E_B(F \cup \delta) \models x \text{ and } E_B(F \cup \delta) \not\models y\}$$

The subset ordering on these sets of unreached desires reflects exactly the original ordering on decisions.

The following theorem follows directly from the proposition.

Theorem 1 *The set of agent system specifications with empty set of parameters, beliefs rules and observations is expressive enough for game specifications.*

Proof. Follows from construction in Proposition 3.

The above construction raises the question whether other sets of agent system specifications are complete too, such as for example the set of agent system specifications in which the initial decision is an empty set. Or the set of agent system specifications in which the set of desires contains only desire-to-be desires. We leave these questions for further research.

Concluding remarks

We are using a very simple approach, probably the simplest approach possible, based on a two sorted propositional language. This simplicity is not a disadvantage but an advantage, for the following reasons: it enables us to focus on essentials (compare how Gelfond and Lifschitz (Gelfond & Lifschitz 1993) simple $\mathcal{A}$ language which focussed planning logics on essentials), it is easy to implement or generalize to classical formalisms, and it shows why exactly we need more complex stuff such as causal theories and action languages. Our motivation comes from the analysis of rule-based agent architectures, which have recently been introduced. However, the results of this paper may be relevant for a much wider audience. For example, Dennett argues that automated systems can be analyzed using concepts from folk psychology like beliefs and desires. Our work may be used in the formal foundations of this 'intentional stance' (Dennett 1987).

In this paper we have defined a qualitative decision and game theory in the spirit of classical decision and game theory. The theory illustrates the micro-macro dichotomy by distinguishing the optimization problem from game theoretic equilibria. We also link decision theory to game theory by showing how agent system specification can be mapped to game specification and back. We think that the method of this paper is more interesting than its formal results. The decision and game theory are based on several ad hoc choices which need further investigation. For example, the desire rules are defeasible but the belief rules are not (the obvious extension leads to wishful thinking problems as studied in (Thomason 2000; Broersen *et al.* 2001a)). We hope that further investigations along this line brings the theories and tools used for individual agents and multi agent systems closer together.

There are several topics for further research. The most interesting question is whether belief and desire rules are fundamental, or whether they in turn can be represented by some other construct. Other topics for further research are the development of an incremental any-time algorithm to find optimal decisions, the development of computationally attractive fragments of the logic, and heuristics of the optimization problem.

References

Binmore, K. 1992. *Fun and Games : A Text on Game Theory*.

Boutilier, C. 1994. Toward a logic for qualitative decision theory. In *Proceedings of the KR'94*, 75–86.

Broersen, J.; Dastani, M.; ; and van der Torre, L. 2001a. Resolving conflicts between beliefs, obligations, intentions and desires. In *Symbolic and Quantitative Approaches to Reasoning and Uncertainty. Proceedings of ECSQARU'01*, LNCS 2143, 568–579. Springer.

Broersen, J.; Dastani, M.; Huang, Z.; Hulstijn, J.; and van der Torre, L. 2001b. The BOID architecture: Conflicts between beliefs, obligations, intentions and desires. In *Proceedings of he Fifth International Conference on Autonomous Agents (AA2001)*, 9–16. ACM Press.

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

Dastani, M., and van der Torre, L. 2002. A classification of cognitive agents. In *Proceedings of the 24th Annual Meeting of the Cognitive Science Society (Cogsci'02)*.

Dennett, D. 1987. *The intentional stance*. Cambridge, MA: MIT Press.

Doyle, J., and Thomason, R. summer 1999. Background to qualitative decision theory. *AI magazine* 20(2).

Gelfond, M., and Lifschitz, V. 1993. Representing action and change by logic programs. *JLP* 17:301–321.

Lang, J. 1996. Conditional desires and utilities - an alternative approach to qualitative decision theory. In *In Proceedings of the European Conference on Artificial Intelligence (ECAI'96)*, 318–322.

Osborne, M. J., and Rubenstein, A. 1994. *A Course in Game Theory*. Cambridge, Massachusetts: The MIT Press.

Rao, A., and Georgeff, M. 1991a. Modeling rational agents within a BDI architecture. In *Proceedings of the KR91*.

Rao, A. S., and Georgeff, M. P. 1991b. Deliberation and its role in the formation of intentions. In *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence (UAI- 91)*.

Savage, L. 1954. *The foundations of statistics*.

Simon, H. A. 1981. *The Sciences of the Artificial*. Cambridge, MA: MIT Press, second edition.

Thomason, R. 2000. Desires and defaults: A framework for planning with inferred goals. In *Proceedings KR 2000*, 702–713.

von Neumann, J., and Morgenstern, O. 1944. *Theory of Games and Economic Behavior*. Princeton, NJ: Princeton University Press, 1 edition.