Games for Cognitive Agents

Mehdi Dastani[1] and Leendert van der Torre[2]

[1] Utrecht University mehdi@cs.uu.nl
[2] CWI torre@cwi.nl

Abstract. Strategic games model the interaction among simultaneous decisions of agents. The starting point of strategic games is a set of players (agents) having strategies (decisions) and preferences on the game's outcomes. In this paper we do not assume the decisions and preferences of agents to be given in advance, but we derive them from the agents' mental attitudes. We specify such agents, define a mapping from their specification to the specification of the strategic game they play. We discuss a reverse mapping from the specification of strategic games that agents play to a specification of those agents. This mapping can be used to specify a group of agents that can play a strategic game, which shows that the notion of agent system specification is expressive enough to play any kind of game.

1 Introduction

There are several approaches in artificial intelligence, cognitive science, and practical reasoning (within philosophy) to the decision making of individual agents. Most of these theories have been developed independently of classical decision theory based on the expected utility paradigm (usually identified with the work of Neumann and Morgenstern [12] and Savage [10]) and classical game theory. 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 [11]). Qualitative decision theories have been developed based on beliefs (probabilities) and desires (utilities) using formal tools such as modal logic [1]. Also, these beliefs-desires models have been extended with intentions or BDI models [3, 9]. 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, new concepts are introduced at this social level, such as norms (a central concept in most social theories). We are interested in the relation between AI theories of decision making, and their classical counterparts.

We introduce a rule based qualitative decision theory for agents with beliefs and desires. Like classical decision theory but in contrast to several proposals in the BDI approach [3, 9], the theory does not incorporate decision processes, temporal reasoning, and scheduling. We also ignore probabilistic decisions. 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. The problems we address are: 1) How can we map the specification of the agent system to the specification of the strategic game that they play? This mapping considers agent decisions as agent strategies and decision profiles (a decision for each agent) as the outcomes of the strategic game. 2) How can we map the specification of a strategic game to the specification of the agent system that plays the game? This mapping provides the mental attitudes of agents that can play a strategic game. We show that the mapping which is composed by a mapping from the specification of a strategic games to the specification of an agent system and back is the identity relation, while the mapping composed of a mapping from the specification of an agent system to the specification of a strategic game and back is not necessarily the identity relation.

The layout of this paper is as follows. In section 2 we introduce the rule based qualitative decision theory. In section 3 we define a mapping from the specification of the agent system to the specification of the strategic game they play. In section 4, we discuss the reverse mapping from the specification of a strategic game to the specification of the agent system that plays the game.

2 Agents, decisions, and preferences

The specification of agent systems introduced in this section is developed for agents that have conditional beliefs and desires. The architecture and the behavior of this type of agent is studied in [2]. Here, we analyze this type of agent 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 individual agents and possible decision profiles (i.e., multiagent decisions). We show how we can derive agent decision profiles and preferences from an agent system specification. In this section, we first define the specification of multiagent systems. Then, we study possible and feasible individual and multiagent decisions within an agent system specification. Finally, we discuss agents' preference ordering defined on the set of individual and multiagent decisions.

2.1 Agent system specification

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 an upgrade at a particular time. Therefore, we assume n disjoint sets of propositional atoms $A = A_1 \cup \ldots \cup A_n$ with typical elements $a, b, c, \ldots$ (agents' decision variables [7] or controllable propositions [1]) and a set of propositional atoms W with typical elements $p, q, r, \ldots$ (the world parameters or uncontrollable propositions) such that $A \cap W = \emptyset$. In the sequel, the propositional languages that are built up from $A_i,\ A,\ W$, and $A \cup W$ atoms are denoted by L_{A_i}, L_A, L_W, and L_{AW}, respectively. Finally, we use variables $x,\ y, \ldots$ to stand for any sentences of the languages L_{A_i}, L_A, L_W, and L_{AW}.

An agent system specification given in Definition 1 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 formalized as belief and desire rules, a preference ordering on the powerset of the set of desire rules, a set of facts, and an initial decision (or prior intentions). An initial decision reflects that the agent has already made a decision (intention) in an earlier stage. One may argue that it is not realistic to define the preference ordering on the power set of the set of desire rules since this implies that for each agent its preference on all combinations of its individual desires should be specified beforehand. As it is explained elsewhere [4, 5], it is possible to define the preference ordering on the set of desire rules and then lift this ordering to the powerset of the set of desire rules. We have chosen to define the preference ordering on the powerset of the set of desire rules to avoid additional complexity which is not related to the main focus of this paper. The preference ordering is also assumed to be a preorder (i.e. reflexive, transitive, and complete). Again, although this assumption is quite strong for realistic applications, we can use it since the main claim of this paper is not a theory for realistic applications. Finally, we assume that agents are autonomous, in the sense that there are no priorities between desires of distinct agents.

Definition 1 (Agent system specification). *An agent system specification is a tuple* $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ *that contains a set of agents* $S = \{\alpha_1, \ldots, \alpha_n\}$*, and for each agent* α_i *a finite set of facts* $F_i \subseteq L_W$ $(F = \langle F_1, \ldots, F_n \rangle)$*, a finite set of belief rules* $B_i \subseteq L_{AW} \times L_W$ $(B = \langle B_1, \ldots, B_n \rangle)$*, a finite set of desire rules* $D_i \subseteq L_{AW} \times L_{AW}$ $(D = \langle D_1, \ldots, D_n \rangle)$*, a relation* $\geq_i$ *on the powerset of* D_i*, i.e.* $\geq_i \subseteq Pow(D_i) \times Pow(D_i)$ $(\geq = \langle \geq_1, \ldots, \geq_n \rangle)$ *which is reflexive, transitive, and complete, and a finite initial decision* $\lambda_i^0 \subseteq L_{A_i}$ $(\lambda^0 = \langle \lambda_1^0, \ldots, \lambda_n^0 \rangle)$.

In general, a belief rule is an ordered pair $x \Rightarrow y$ with $x \in L_{AW}$ and $y \in L_W$. This belief rule should be interpreted as 'the agent believes y in context x'. A desire rule is an ordered pair $x \Rightarrow y$ 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 ($x \Rightarrow p$), and not about the agent's decisions. These beliefs can be about the effects of decisions made by the agent ($a \Rightarrow p$) as well as beliefs about the effects of parameters set by the world ($p \Rightarrow q$). Moreover, the agent's desires can be about the world ($x \Rightarrow p$, desire-to-be), but also about the agent's decisions ($x \Rightarrow a$, desire-to-do). These desires can be triggered by parameters set by the world ($p \Rightarrow y$) as well as by decisions made by the agent ($a \Rightarrow y$). Modelling mental attitudes such as beliefs and desires in terms of rules can be called modelling conditional mental attitudes [2].

2.2 Agent decisions

In the sequel we consider each agent from an agent system specification as a decision making agent. A decision λ of the agent α_i is any consistent subset of L_{A_i} that contains the initial decision λ_i^0.

Definition 2 (Decisions). *Let* $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ *be an agent system specification,* L_{A_i} *be the propositional language built up from* A_i*, and* $\models_{A_i}$ *be satisfiability in propositional logics* L_{A_i}*. An* AS *decision* λ *is a decision of the agent* α_i *such that* $\lambda_i^0 \subseteq \lambda \subseteq L_{A_i}$ $\&$ $\lambda \not\models_{A_i} \bot$*. The set of possible decisions of agent* α_i *is denoted by* Λ_i*.*

The set of possible decisions Λ_i of an agent α_i contains logically equivalent decisions. Two decisions $\lambda, \lambda' \in \Lambda_i$ are logically equivalent, denoted as $\lambda \equiv \lambda'$, if and only if for all models M of L_{A_i} : $M \models \lambda$ iff $M \models \lambda'$, where $M \models \lambda$ iff $\forall x \in \lambda\ M \models x$.

Definition 3 (Non-equivalent Decisions). *Let Λ_i be the set of possible decisions of agent α_i. A set of possible logically non-equivalent decisions of agent α_i, denoted as $\widetilde{\Lambda}_i$, is a subset of Λ_i such that:* $\forall \lambda \in \Lambda_i\ \exists \lambda' \in \widetilde{\Lambda}_i\ \ \lambda \equiv \lambda'\ \&\ \forall \lambda, \lambda' \in \widetilde{\Lambda}_i\ \ \lambda \not\equiv \lambda'$.

The decisions of an agent depend on the believed consequences of those decisions. The consequences are generated by applying its belief rules to its input facts together with those decisions. In our framework the decisions are formalized based on the notion of extension.

Definition 4 (Belief Extension). *Let Cn_A, Cn_W and Cn_{AW} be the consequence sets for theories from $L_A,\ L_W$, and L_{AW}, respectively, and $\models_A$, $\models_W$ and $\models_{AW}$ be satisfiability, in propositional logics $L_A,\ L_W$, and L_{AW}, respectively. Let B_i be a set of belief rules of agent α_i, $\lambda_i \in \Lambda_i$ be one of its possible decisions, and $F_i \subseteq L_A$ be its set of facts. The belief consequences of $F_i \cup \lambda_i$ of agent α_i are: $B_i(F_i \cup \lambda_i) = \{y \mid x \Rightarrow y \in B_i, x \in F_i \cup \lambda_i\}$ and the belief extension of $F_i \cup \lambda_i$ is the set of the consequents of the iteratively B_i-applicable rules: $E_{B_i}(F_i \cup \lambda_i) = \bigcap_{F_i \cup \lambda_i \subseteq X, B_i(Cn_{AW}(X)) \subseteq X} X$.*

We give some properties of the belief extension of facts and possible decisions in Definition 4. First note that $E_{B_i}(F_i \cup \lambda_i)$ is *not* closed under logical consequence. The following proposition shows that $E_{B_i}(F_i \cup \lambda_i)$ is the smallest superset of $F_i \cup \lambda_i$ closed under the belief rules B_i interpreted as inference rules.

Proposition 1. *Let $E^0_{B_i}(F_i \cup \lambda_i) = F_i \cup \lambda_i$ and $E^j_{B_i}(F_i \cup \lambda_i) = E^{j-1}_{B_i}(F_i \cup \lambda_i) \cup B_i(Cn_{AW}(E^{j-1}_{B_i}(F_i \cup \lambda_i)))$ for $j > 0$. We have $E_{B_i}(F_i \cup \lambda_i) = \cup_{j=0}^{\infty} E^j_{B_i}(F_i \cup \lambda_i)$.*

The following example illustrates that extensions can be inconsistent.

Example 1. Let $B_i = \{\top \Rightarrow p, a \Rightarrow \neg p\}$, $F_i = \emptyset$, and $\lambda_i = \{a\}$, where $\top$ stands for any tautology like $p \vee \neg p$. We have $E_{B_i}(\emptyset) = \{p\}$ and $E_{B_i}(F_i \cup \lambda_i) = \{a, p, \neg p\}$, which means that the belief extension of $F_i \cup \lambda_i$ is inconsistent.

Although decisions with inconsistent belief consequences are not feasible decisions, we consider them, besides decisions with consistent consequences, as possible decisions. Feasible decisions are defined by excluding decisions that have inconsistent belief consequences.

Definition 5 (Feasible Decisions). *Let $AS = \langle S, F, B, D, \geq, \lambda^0\rangle$ be an agent system specification, Λ_i and $\widetilde{\Lambda}_i$ be the set of possible decisions and a set of possible logically non-equivalent decisions for agent $\alpha_i \in S$, respectively. The set of feasible decisions of agent α_i, denoted by Λ^f_i, is the subset of its possible decisions Λ_i that have consistent belief consequences, i.e., $\Lambda^f_i = \{\lambda_i \mid \lambda_i \in \Lambda_i\ \&\ E_{B_i}(F_i \cup \lambda_i)$ is consistent $\}$. A set of logically non-equivalent feasible decisions of agent α_i, denoted by $\widetilde{\Lambda}^f_i$, is the subset of a set of possible non-equivalent decisions $\widetilde{\Lambda}_i$ that have consistent belief consequences, i.e. $\widetilde{\Lambda}^f_i = \{\lambda_i \mid \lambda_i \in \widetilde{\Lambda}_i\ \&\ E_{B_i}(F_i \cup \lambda_i)$ is consistent $\}$.*

The following example illustrates the decisions of a single agent.

Example 2. Let $A_1 = \{a, b, c, d\}$, $W = \{p, q\}$ and $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ with $S = \{\alpha_1\}$, $F_1 = \emptyset$, $B_1 = \{b \Rightarrow q, c \Rightarrow p, d \Rightarrow \neg p\}$, $D_1 = \{b \Rightarrow p, d \Rightarrow \neg q\}$, $\geq_1 = \emptyset < \{b \Rightarrow p\} < \{d \Rightarrow \neg q\} < \{b \Rightarrow p, d \Rightarrow \neg q\}$, and $\lambda_1^0 = \{a\}$. Note that the consequents of all B_1 rules are sentences of L_W. We have due to the definition of $E_{B_1}(F_1 \cup \lambda_1)$, for example, the following logically non-equivalent decisions.
$E_{B_1}(F_1 \cup \{a\}) = \{a\}$, $E_{B_1}(F_1 \cup \{a, b\}) = \{a, b, q\}$,
$E_{B_1}(F_1 \cup \{a, c\}) = \{a, c, p\}$, $E_{B_1}(F_1 \cup \{a, d\}) = \{a, d, \neg p\}$,
$E_{B_1}(F_1 \cup \{a, b, c\}) = \{a, b, c, p, q\}$, $E_{B_1}(F_1 \cup \{a, b, d\}) = \{a, b, d, \neg p, q\}$,
$E_{B_1}(F_1 \cup \{a, c, d\}) = \{a, c, d, p, \neg p\}$, $E_{B_1}(F_1 \cup \{a, b, c, d\}) = \{a, b, c, d, p, \neg p, q\}, \ldots$
Therefore $\{a, c, d\}$ and $\{a, b, c, d\}$ are infeasible AS decisions, because their belief extensions are inconsistent. Continued in Example 4.

2.3 Multiagent decisions

In the previous subsection, we have defined the set of decisions, sets of logically non-equivalent decisions, the set of feasible decisions, and sets of logically non-equivalent feasible decisions for one single agent. In this section, we concentrate on multiagent decisions, which are also called decision profiles, and distinguish various types of multiagent decisions.

Definition 6. *Let $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ be an agent system specification where $S = \{\alpha_1, \ldots, \alpha_n\}$. Let also Λ_i and $\widetilde{\Lambda}_i$ be the set of possible decisions and a set of logically non-equivalent decisions for agent $\alpha_i \in S$, respectively. The set of possible decision profiles and a set of logically non-equivalent AS decision profiles are $\Lambda = \Lambda_1 \times \ldots \times \Lambda_n$ and $\widetilde{\Lambda} = \widetilde{\Lambda}_1 \times \ldots \times \widetilde{\Lambda}_n$, respectively. An AS decision profile (i.e., a multiagent decision) λ is a tuple $\langle \lambda_1, \ldots, \lambda_n \rangle$, where $\lambda_i \in \Lambda_i$ for $1 \leq i \leq n$.*

According to definition 5, the feasibility of decisions of individual agents is formulated in terms of the consistency of the extension that is calculated based on the decision and its own facts and beliefs. In a multiagent setting the feasibility of decisions of a single agent depends also on the decisions of other agents. For example, if an agent decides to open a door while another agent decides to close it, then the combined decision can be considered as infeasible. In order to capture the feasibility of multiagent decisions, we consider the feasibility of decision profiles which depends on whether agents' beliefs, facts, and decisions are private or public.

Definition 7. *Let $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ where $S = \{\alpha_1, \ldots, \alpha_n\}$. A decision profile $\lambda = \langle \lambda_1, \ldots, \lambda_n \rangle$ is feasible if $E_B(F \cup \lambda)$ is consistent. Below, eight ways to calculate $E_B(F \cup \lambda)$ are distinguished.*
1. $E_{B_1 \cup \ldots \cup B_n}(F_1 \cup \lambda_1 \cup \ldots F_n \cup \lambda_n)$, i.e., public beliefs, facts, and decisions
2. $\bigcup_i E_{B_1 \cup \ldots \cup B_n}(F_1 \cup \ldots F_n \cup \lambda_i)$, i.e., public beliefs and facts, private decisions
3. $\bigcup_i E_{B_1 \cup \ldots \cup B_n}(F_i \cup \lambda_1 \cup \ldots \cup \lambda_n)$, i.e., public beliefs and decisions, private facts
4. $\bigcup_i E_{B_1 \cup \ldots \cup B_n}(F_i \cup \lambda_i)$, i.e., public beliefs, private facts and decisions
5. $\bigcup_i E_{B_i}(F_1 \cup \lambda_1 \cup \ldots F_n \cup \lambda_n)$, i.e., public facts and decisions, private beliefs
6. $\bigcup_i E_{B_i}(F_i \cup \lambda_1 \cup \ldots \cup \lambda_n)$, i.e., public decisions, private beliefs and facts
7. $\bigcup_i E_{B_i}(F_1 \cup \ldots F_n \cup \lambda_i)$, i.e., public facts, private beliefs and decisions
8. $\bigcup_i E_{B_i}(F_i \cup \lambda_i)$, i.e., private beliefs, facts, and decisions

Given one of these definitions of $E_B(F \cup \lambda)$, the set of feasible decisions and a set of logically non-equivalent feasible decisions profiles are denoted by Λ^f and $\widetilde{\Lambda}^f$, respectively.

Another way to explain these definitions of the feasibility of decision profiles is in terms of communication between agents. The agents communicate their beliefs, facts, or decisions and through this communication decision profiles become infeasible. The following example illustrates the feasibility of decision profiles according to the eight variations.

Example 3. Let $A_1 = \{a\}, A_2 = \{b\}, W = \{p\}$ and $AS = \langle\{\alpha_1, \alpha_2\}, F, B, D, \geq, \lambda^0\rangle$ with $F_1 = F_2 = \emptyset$, $B_1 = \{b \Rightarrow p\}$, $B_2 = \{a \Rightarrow \neg p\}$, $D_1 = D_2 = \emptyset$, $\geq$ is the universal relation, and $\lambda_1^0 = \lambda_2^0 = \emptyset$. Note that the only belief of each agent is about the consequence of the decisions that can be taken by the other agent. The following four AS decision profiles are possible; the numbers associated to the following decision profiles indicate according to which definitions of $E_B(F \cup \lambda)$ the decision profile λ is feasible: $\langle\emptyset, \emptyset\rangle : 1 \ldots 8$ $\langle\{a\}, \emptyset\rangle : 1 \ldots 8$ $\langle\emptyset, \{b\}\rangle : 1 \ldots 8$ $\langle\{a\}, \{b\}\rangle : 7, 8$
Since the consequence of decisions that can be taken by each agent is captured by the belief of the other agent, the decision profile $\langle\lambda_1, \lambda_2\rangle = \langle\{a\}, \{b\}\rangle$ is only feasible when the two agents do not communicate their beliefs and decisions, i.e.,
7. $E_B(F \cup \lambda) = E_{B_1}(F_1 \cup F_2 \cup \lambda_1) \bigcup E_{B_2}(F_1 \cup F_2 \cup \lambda_2) = \{a\} \bigcup \{b\} = \{a, b\}$.
8. $E_B(F \cup \lambda) = E_{B_1}(F_1 \cup \lambda_1) \bigcup E_{B_2}(F_2 \cup \lambda_2) = \{a\} \bigcup \{b\} = \{a, b\}$.
In all other cases, the decision profile $\langle\{a\}, \{b\}\rangle$ will be infeasible since $E_B(F \cup \lambda) = \{a, p, b, \neg p\}$ is inconsistent.

In general, agents may communicate their beliefs, facts, and decisions only to some, but not all, agents. The set of agents to which an agent communicates its beliefs, facts, and decisions depend on the communication network between agents. In order to define the feasibility of decision profiles in terms of communication, one need to restrict the union operators in various definition of $E_B(F \cup \lambda)$ to subsets of agents that can communicate with each other. In this paper, we do not study this further requirement.

Various definitions of $E_B(F \cup \lambda)$, as proposed in definition 7, are related to each other with respect to the public-/privateness of beliefs, facts, and decisions.

Definition 8. *A definition $\mathcal{D}$ of $E_B(F \cup \lambda)$, as given in definition 7, is more public than another definition $\mathcal{D}'$ of $E_B(F \cup \lambda)$, written as $\mathcal{D}' \sqsubseteq_p \mathcal{D}$, if and only if all aspects (i.e., beliefs, facts, and decisions) that are public in $\mathcal{D}'$ are also public in $\mathcal{D}$.*

The definition results a lattice structure on the eight definitions of $E_B(F \cup \lambda)$, as illustrated in Figure 1. The top of the lattice is the definition of $E_B(F \cup \lambda)$ according to which beliefs, facts, and decisions of agents are communicated to each other, and the bottom of the lattice is the definition of $E_B(F \cup \lambda)$ according to which beliefs, facts, and decisions of agents are not communicated. This lattice shows that the more-public-than relation is the same as subset relation on the public aspects.

Proposition 2. *Let $\mathcal{D}$ and $\mathcal{D}'$ be two definitions of $E_B(F \cup \lambda)$ such that $\mathcal{D}' \sqsubseteq_p \mathcal{D}$. The feasibility of decision profiles persists under the $\sqsubseteq_p$ relation, i.e., for all decision profiles λ and λ' if the decision profile λ is feasible w.r.t. the definition $\mathcal{D}$, then it is also feasible w.r.t. the definition $\mathcal{D}'$.*

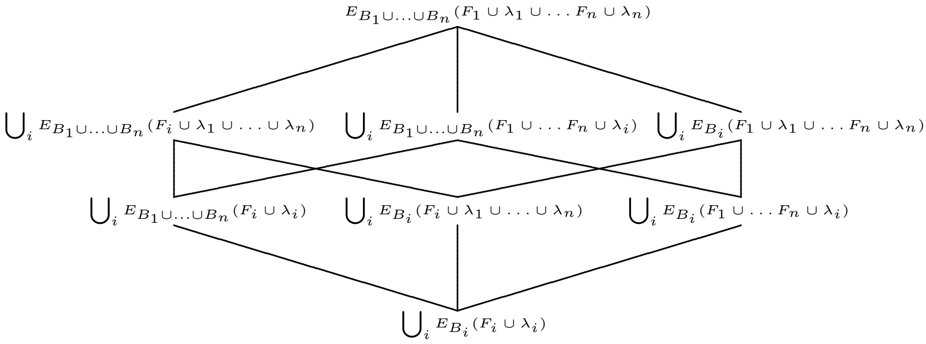

Fig. 1. Different definition for $E_B(F \cup \lambda)$ and their internal structures.

This proposition states that if a decision is feasible when aspects are public, the decision remains feasible when public aspects become private. The following proposition states that communication is relevant for the feasibility of decisions only if the agent system specification consists of more than one agent.

Proposition 3. *Let $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$ be an agent system specification where $|S| = 1$, i.e., there exists only one agent. If a decision λ is feasible according one definition of $E_B(F \cup \lambda)$, then it is feasible according to all definitions of $E_B(F \cup \lambda)$.*

In this sequel, we proceed with the definition of $E_B(F \cup \lambda)$ where agents' facts, beliefs, and decisions are assumed to be private, i.e., $E_B(F \cup \lambda) = \bigcup_i E_{B_i}(F_i \cup \lambda_i)$.

2.4 Agent preferences

In this section we introduce a way to compare decisions. Decisions are compared by sets of desire rules that are not reached by the decisions. A desire $x \Rightarrow y$ is unreached by a decision if the expected consequences of the decision imply x but not y[3].

Definition 9 (Comparing decision profiles). *Let $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$, λ be a AS decision profile, and $E_B(F \cup \lambda) = \bigcup_i E_{B_i}(F_i \cup \lambda_i)$. The unreached desires of λ for agent α_i are: $U_i(\lambda) = \{x \Rightarrow y \in D_i \mid E_B(F \cup \lambda) \models x$ and $E_B(F \cup \lambda) \not\models y\}$. Decision profile λ is at least as good as decision profile λ' for agent α_i, written as $\lambda \geq_i^U \lambda'$, iff $U_i(\lambda') \geq_i U_i(\lambda)$. Decision profile λ is equivalent to decision profile λ' for agent α_i, written as $\lambda \approx_i^U \lambda'$, iff $\lambda \geq_i^U \lambda'$ and $\lambda' \geq_i^U \lambda$. Decision profile λ dominates decision profile λ' for agent α_i, written as $\lambda >_i^U \lambda'$, iff $\lambda \geq_i^U \lambda'$ and $\lambda' \not\geq_i^U \lambda$.*

Thus, a decision profile λ is preferred by agent α_i to decision profile λ' if the set of unreached desire rules by λ is less preferred by α_i to the set of unreached desire rules by λ'. Note that the set of unreached desire rules for an infeasible decision is the whole set of desire rules such that all infeasible decision profiles are equally preferred for each agent. The following continuation of Example 2 illustrates the comparison of decisions.

Example 4 (Continued). In example 2, the agent system specification consists of only one agent such that each decision profile consists only of decisions of one agent. Below

[3] The comparison can also be based on the set of violated or reached desires. The desire rule is violated or reached if these consequences imply $x \wedge \neg y$ or $x \wedge y$, respectively.

are some decision profiles and their corresponding sets of unreached desire rules:
$U_1(\langle\{a\}\rangle) = \emptyset$, $U_1(\langle\{a, b\}\rangle) = \{b \Rightarrow p\}$, $U_1(\langle\{a, c\}\rangle) = \emptyset$,
$U_1(\langle\{a, d\}\rangle) = \{d \Rightarrow \neg q\}$, $U_1(\langle\{a, b, c\}\rangle) = \emptyset$, $U_1(\langle\{a, b, d\}\rangle) = \{b \Rightarrow p\ ,\ d \Rightarrow \neg q\}$,
$U_1(\langle\{a, c, d\}\rangle) = \emptyset$, $U_1(\langle\{a, b, c, d\}\rangle) = \emptyset$, $\ldots$
We thus have for example that the decision profile $\langle\{a, c\}\rangle$ dominates the decision profile $\langle\{a, b\}\rangle$, $\langle\{a, b\}\rangle$ dominates $\langle\{a, d\}\rangle$, and $\langle\{a, d\}\rangle$ dominates $\langle\{a, b, d\}\rangle$, i.e., $\{a, c\} >^U_1 \{a, b\} >^U_1 \{a, d\} >^U_1 \{a, b, d\}$. Moreover, $\langle\{a\}\rangle$, $\langle\{a, c\}\rangle$, and $\langle\{a, b, c\}\rangle$ are equivalent, i.e., $\langle\{a\}\rangle \approx \langle\{a, c\}\rangle \approx \langle\{a, b, c\}\rangle$.

In Example 2, the decision profiles $\langle\{a, c, d\}\rangle$ and $\langle\{a, b, c, d\}\rangle$ are infeasible and therefore have the whole set of desire rules as the set of unreached desire rules. A problem with defining the unreached based ordering on infeasible decisions is that the unreached based ordering cannot differentiate between some feasible decisions, e.g., $\langle\{a, b, c\}\rangle$, and infeasible decisions, e.g., $\langle\{a, c, d\}\rangle$.

The following proposition states that the agent's ordering on decision profiles based on unreached desire rules is a preference ordering.

Proposition 4. *Let $AS = \langle S, F, B, D, \geq, \lambda^0\rangle$ be an agent system specification. For each agent, the ordering on decision profiles based on the set of unreached desire rules is a preorder.*

3 From Agent System Specifications to Game Specifications

In this subsection, we consider interactions between agents based on agent system specifications, their corresponding agent decisions, and the ordering on the decisions as explained in previous subsections. 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 [8].

In the sequel, we consider only strategic games, which is described in [8] as follows: *A strategic game is a model of interactive decision making in which each decision maker chooses his plan of action once and for all, and these choices are made simultaneously*. For strategic games, we use δ_i to denote a decision (strategy) of agent α_i and $\delta = \langle\delta_1, \ldots, \delta_n\rangle$ to denote a decision (strategy) profile containing one decision (strategy) for each agent. We use δ_{-i} to denote the decision profile of all agents except the decision of agent α_i, and (δ_{-i}, δ'_i) to denote a decision profile which is the same as δ except that the decision of agent i from δ is replaced with the decision δ'_i. Finally, we use Δ to denote the set of all possible decision profiles for agents $\alpha_1, \ldots, \alpha_n$ and Δ_i to denote the set of possible decisions for agent α_i.

Definition 10 (Strategic Game Specification [8]). *A strategic game specification is a tuple $\langle N, \Delta, \geq^{gs}\rangle$ where $N = \{\alpha_1, \ldots, \alpha_n\}$ is a set of agents, $\Delta \subseteq \Delta_1 \times \ldots \times \Delta_n$ for Δ_i is a set of decisions of agent α_i, and $\geq^{gs} = \langle\geq^{gs}_1, \ldots, \geq^{gs}_n\rangle$, for $\geq^{gs}_i$ is the preference order of agent α_i on Δ. The preference order $\geq^{gs}_i$ is reflexive, transitive, and complete.*

We now define a mapping $AS2AG$ from the specification of an agent system to the specification of a strategic game. This mapping is based on the correspondence between logically non-equivalent feasible decisions from an agent system specification and agents' decisions in strategic game specifications, and between agents' preference orderings in both specifications. We have chosen to make a correspondence based on the logically non-equivalent decisions (and not possible decisions) since otherwise the resulting strategic game specification contain logically equivalent decisions profiles.

Definition 11 (AS2AG). *Let $S = \{\alpha_1, \ldots, \alpha_n\}$, $AS = \langle S, F, B, D, \geq, \lambda^0 \rangle$, $\widetilde{\Lambda}^f$ be a set of AS logically non-equivalent feasible decision profiles according to Definition 7 and for a given definition of $E_B(F \cup \lambda)$ (i.e., $\bigcup_i E_{B_i}(F_i \cup \lambda_i)$), and $\geq_i^U$ be the AS preference order of agent α_i defined on $\widetilde{\Lambda}^f$ according to definition 9. The strategic game specification of AS is $GS = \langle N, \Delta, \geq^{gs} \rangle$ if there exists a bijjective function g from AS to GS such that $g : S \rightarrow N$, $g : \widetilde{\Lambda}^f \rightarrow \Delta$, and $g : \geq^U \rightarrow \geq^{gs}$.*

Example 5. Let α_1 and α_2 be two agents, $F_1 = F_2 = \emptyset$, and initial decisions $\lambda_1^0 = \lambda_2^0 = \emptyset$. They have the following beliefs en desires: $B_1 = \{a \Rightarrow p\}$, $D_1 = \{\top \Rightarrow p, \top \Rightarrow q\}$, $\emptyset \geq_1 \{\top \Rightarrow q\} \geq_1 \{\top \Rightarrow p\} \geq_1 \{\top \Rightarrow p, \top \Rightarrow q\}$, $B_2 = \{b \Rightarrow q\}$, $D_2 = \{q \Rightarrow \neg p, p \Rightarrow \neg q\}$, $\emptyset \geq_2 \{q \Rightarrow \neg p\} \geq_2 \{p \Rightarrow \neg q\} \geq_2 \{p \Rightarrow \neg q, \Rightarrow \neg p\}$. Let λ be feasible decision profile according to any of the eight formulations of $E_B(F \cup \lambda)$ as proposed in definition 7, and $U_i(\lambda)$ be the set of unreached desires for agent α_i and decision profile λ.

λ	$E_B(F \cup \lambda)$	$U_1(\lambda)$	$U_2(\lambda)$
$\langle \emptyset, \emptyset \rangle$	$\emptyset$	$\{\top \Rightarrow p, \top \Rightarrow q\}$	$\emptyset$,
$\langle \emptyset, \{b\} \rangle$	$\{q\}$	$\{\top \Rightarrow p\}$	$\{q \Rightarrow \neg p\}$
$\langle \{a\}, \emptyset \rangle$	$\{p\}$	$\{\top \Rightarrow q\}$	$\{p \Rightarrow \neg q\}$
$\langle \{a\}, \{b\} \rangle$	$\{p, q\}$	$\emptyset$	$\{p \Rightarrow \neg q, q \Rightarrow \neg p\}$

According to definition 9 the preference ordering over possible decision profiles for agent α_1 is $\langle \{a\}, \{b\} \rangle \geq_1^U \langle \{a\}, \emptyset \rangle \geq_1^U \langle \emptyset, \{b\} \rangle \geq_1^U \langle \emptyset, \emptyset \rangle$, and for agent α_2 is $\langle \emptyset, \emptyset \rangle \geq_2^U \langle \emptyset, \{b\} \rangle \geq_2^U \langle \{a\}, \emptyset \rangle \geq_2^U \langle \{a\}, \{b\} \rangle$. Consider now the following mapping from AS to GS based on the bijective function g defined as $g(\alpha_i) = \alpha_i \;\; \forall \alpha_i \in S, g(\lambda) = \lambda \;\; \forall \lambda \in \widetilde{\Lambda}^f$, and $g(\geq_i^U) = \geq_i^U$ for $i = 1, 2$. The resulting strategic game specification of AS is $GS = \langle N, \Delta, \geq^{gs} \rangle$, where $N = \{g(\alpha_i) \mid \alpha_i \in S\}$, $\Delta = \{g(\lambda) \mid \lambda \in \widetilde{\Lambda}^f\}$, and $\geq^{gs} = \langle g(\geq_1^U), g(\geq_2^U) \rangle$.

We now use the mapping from AS to GS and consider different types of decision profiles which are similar to types of decision (strategy) profiles from game theory.

Definition 12. *[8] Let $\widetilde{\Lambda}^f$ be a set of logically non-equivalent feasible decision profiles that is derived from the AS specification of an agent system and GS be the strategic game specification of AS based on the mapping g. A feasible AS decision profile $\lambda \in \widetilde{\Lambda}^f$ is* **Pareto decision** *if $g(\lambda) = \delta$ is a pareto decision in GS, i.e., if there is no $\delta' \in \Delta$ for which $\delta_i' >_i^{gs} \delta_i$ for all agents $\alpha_i \in N$. A feasible AS is* **strongly Pareto decision** *if $g(\lambda) = \delta$ is a strongly Pareto decision in GS, i.e., if there is no $\delta' \in \Delta$ for which $\delta_i' \geq_i^{gs} \delta_i$ for all agents α_i and $\delta_j' >_j^{gs} \delta_j$ for some agents α_j. A feasible AS is* **weak dominant decision** *if $g(\lambda) = \delta$ is a weak dominant decision in GS, i.e., if for all $\delta' \in \Delta$ and for every agent α_i it holds: $(\delta_{-i}', \delta_i) \geq_i^{gs} (\delta_{-i}', \delta_i')$. A feasible AS is* **strong**

dominant decision *if for all* $\delta' \in \Delta$ *and for every agent* α_i *it holds:* $(\delta'_{-i}, \delta_i) >_i^{gs} (\delta'_{-i}, \delta'_i)$*. Finally, a feasible* AS *is* **Nash decision** *if* $g(\lambda) = \delta$ *is a Nash decision in* GS*, i.e., if for all agents* α_i *it holds:* $(\delta_{-i}, \delta_i) \geq_i^{gs} (\delta_{-i}, \delta'_i)$ *for all* $\delta'_i \in \Delta_i$*.*

It is a well known fact that Pareto decisions exist (for finite games), whereas dominant decisions do not have to exist. Consider the strategic game specification GS which is derived from the agent system specification AS in Example 5. None of the decision profiles in GS are dominant decisions.

Starting from an agent system specification, we can derive the strategic game specification and in this game specification we can use standard techniques to for example find the Pareto decisions. The problem with this approach is that the translation from an agent system specification to a strategic 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. A challenge of qualitative game theory is therefore whether we can bypass the translation to strategic 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? An example is the case in which the agents have the same individual agent specification, because in that case the game reduces to single agent decision making. In this paper we do not pursue this challenge, but we consider the expressive power of agent specifications.

4 From Game Specifications to Agent Specifications

In this section the question is raised whether the notion of agent system specification is expressive enough for strategic game specifications, that is, whether for each possible strategic game specification there is an agent system specification that can be mapped on it. We prove this property in the following way. First, we define a mapping from strategic game specifications to agent system specifications. Second, we show that the composite mapping from strategic game specifications to agent system specifications and back to strategic game specifications is the identity relation. The second step shows that if a strategic game specification GS is mapped in step 1 on agent system specification AS, then this agent system specification AS can be mapped on GS. Thus, it shows that there is a mapping from agent system specifications to strategic game specifications for every strategic game specification GS.

Unlike the second step, the composite mapping from agent system specifications to strategic 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 strategic game specification. For example, agent system specifications in which the variable names are uniformly substituted by new names. The mapping from strategic game specifications to agent system specifications consists of the following steps. 1) The set of agents from strategic game specification is the set of agents for the agent system specification. 2) For each agent in the agent system specification a set of decision variables is introduced that will generate the set of decisions of the agent in the strategic game specification. 3) For each agent in the agent system specification a set of desire rules and a preference ordering on the powerset of the set of desire rules are introduced, such that they generate the preference order on decision profiles for the agent in the strategic game specification.

According to definition 10, the ingredients of the specification of strategic games are agents identifiers, agents decisions, and the preference ordering of each agent on decision profiles. For each decision $\delta_i \in \Delta_i$ of each agent α_i in a strategic game specification GS we introduce a separate decision variable d_i for the corresponding agent α_i in the agent system specification AS. The set of introduced decision variables for agent α_i is denoted by A_i. The propositional language L_{A_i}, which is based on A_i, specifies the set of possible AS decisions for agent α_i. However, L_{A_i} specifies more AS decisions than GS decisions since it can combine decision variables with conjunction and negation. In order to avoid additional decisions for each agent, we design the initial decisions for each agent such that possible decisions are characterized by one and only one decision variable. In particular, the initial decision λ_i^0 of agent α_i is specified as follows: $\lambda_i^0 = \{\bigvee_k d^k\ ,\ d \rightarrow \neg d' \mid d^k \in A_i\ \&\ d \neq d'\ \&\ d, d' \in A_i\}$. Moreover, the set of decision profiles of strategic game specifications is a subset of all possible decision profiles, i.e., $\Delta \subseteq \Delta_1 \times \ldots \times \Delta_n$. This implies that we need to exclude AS decision profiles that correspond with the excluded decision profiles in GS. For this reason, we introduce belief rules for relevant agents to make excluded decision profile infeasible in AS. For example, suppose in a strategic game specification with two agents α_1 and α_2 the decision profile is excluded, i.e., $(\delta_1, \delta_2) \notin \Delta$, and that d_i is the introduced decision variable for δ_i for $i = \{1, 2\}$. Then, we introduce a new parameter $p \in W$ and add one belief formula for each agent as follows: $d_1 \Rightarrow p \in B_1\ \ \&\ \ d_2 \Rightarrow \neg p \in B_2$. Note that in this example the decision profile (d_1, d_2) in AS is not a feasible decision profile anymore.

We use the preference ordering $\geq_i^{gs}$ of each agent α_i, defined on the set of decision profiles Δ of the strategic game specification, and introduce a set of desire rules D_i together with a preference ordering $\geq_i$ on the powerset of D_i for agent α_i. The set of desire rules for an agent and its corresponding preference ordering are designed in such a way that if the set of unreached desire rules for a decision in the agent system specification are more preferred than the set of unreached desire rules for a second decision, then the first decision is less preferred than the second one (see definition 9).

Definition 13. *Let* $GS = \langle N, \Delta, \geq^{gs}\rangle$ *be a strategic game specification where* $N = \{\alpha_1, \ldots, \alpha_n\}$ *is the set of agents,* $\Delta \subseteq \Delta_1 \times \ldots \times \Delta_n$ *is the set of decision profiles,* $\geq^{gs} = \langle \geq_1^{gs}, \ldots, \geq_n^{gs}\rangle$ *consists of the preference orderings of agents on* Δ*, and* $\bigvee \neg d_{-i} = \neg d_1 \vee \ldots \vee \neg d_{i-1} \vee \neg d_{i+1} \ldots \vee \neg d_n$*. Then,* $AS = \langle S, \emptyset, B, D, \geq, \lambda^0\rangle$ *is the agent system specification derived from* GS*, where* $B = \langle B_1, \ldots, B_n\rangle, D = \langle D_1, \ldots, D_n\rangle, \lambda^0 = \langle \lambda_1^0, \ldots, \lambda_n^0\rangle, \geq = \langle \geq_1, \ldots, \geq_n\rangle$*, and* $S = N$*,* $A_i = \{d \mid d$ *is a decision variable for* $\delta \in \Delta_i\}$*,* $W = \{p_1, \ldots, p_n\}$ *with parameter for each infeasible decision profile,* $\lambda_i^0 = \{\bigvee_k d^k,\ d \rightarrow \neg d' \mid d^k \in A_i\ \&\ d \neq d'\ \&\ d, d' \in A_i\}, (d_i \Rightarrow p) \in B_i \& (d_j \Rightarrow \neg p) \in B_j\ \forall (\delta_i, \delta_j) \notin \Delta\ \&\ d_i \in A_i\ \&\ d_j \in A_j\ \&\ p \in W$*, and* $D_i = \{d_i \Rightarrow \bigvee \neg d_{-i} \mid \forall \delta = \langle \delta_1, \ldots, \delta_n\rangle \in \Delta\}$*. The preference relation* $\geq_i$ *defined on* $Pow(D_i)$ *is characterized as follows: 1)* $s \geq_i \emptyset$ *for all* $s \in Pow(D_i)$*, 2)* $\{d_i \Rightarrow \bigvee \neg d_{-i}\} \geq_i \{d_i' \Rightarrow \bigvee \neg d_{-i}'\}$ *iff* $\delta' \geq_i^{gs} \delta$*, and 3)* $s' \geq_i s$ *for all* $s, s' \in Pow(D_i)\ \&\ |s| \leq 1\ \&\ |s'| > 1$.

In this definition, the set D_i is designed in the context of ordering on decisions based on unreached desire rules, as it is defined in Definition 9. In particular, for agent α_i we define for each decision variable a desire rule that will be unreached as the only desire rule by exactly one decision profile. This is the construction of D_i in this definition. Then, we use the reverse of the preference ordering from the strategic game specification, which was defined on decision profiles, to order the unreached (singletons) sets of desire rules. Since each decision profile has exactly one unreached desire rule, the preference ordering $\geq_i^{gs}$ on decision profiles can be directly used as a preference ordering on unreached sets of desire rules, which in turn is the reverse of the preference ordering $\geq_i$ defined on the powerset of the set D_i of desire rules. This is realized by the last three items of this definition. The first item indicates that any desire rule is more preferred than no desire rule $\emptyset$. The second item indicates that if a decision profile δ' is preferred over a decision profile δ according to the unreached desire rules, then the desire rules that are unreached by δ are preferred over the desire rules that are unreached by δ'. Finally, the last item guarantees that the characterized preference ordering $\geq_i$ is complete by indicating that all other sets of of desire rules are preferred over the singletons of desire rules (sets consisting of only one desire rule) and the empty set of desire rules.

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

Proposition 5. *Let GS be a strategic game specification as in Definition 10. Moreover, let AS be the derived agent system specification as defined in Definition 13. The application of the mapping from agent system specification to strategic game specification, as defined in Definition 11, maps AS to GS.*

Proof. Assume any particular GS. Construct the AS as above. Now apply Definition 9. The unreached desires of decision δ for agent α_i are $U_i(\delta) = \{x \Rightarrow y \in D \mid E_B(F \cup \delta) \models x$ 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 facts is expressive enough for strategic game specifications.*

Proof. Follows from construction in Proposition 5.

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 set of desires contains only desire-to-be desires. We leave these questions for further research.

5 Concluding remarks

In this paper we introduce agent system specifications based on belief and desire rules, and we show how various kinds of strategic games can be played (depending on whether the beliefs, desires and decisions are public or private), and we show how for each possible strategic game an agent specification can be defined that plays that game. The agent system specification we propose is relatively simple, but the extension of the results to more complex agent system specifications seems straightforward. We believe that such results give new insights in the alternative theories which are now developed in artificial intelligence, agent theory and cognitive science.

Our work is typical for a line of research knows as qualitative decision theory which aims at closing the gap between on the one hand classical decision and game theory, and on the other hand alternative theories developed in artificial intelligence, agent theory and cognitive science. Our main result is in our opinion not the particular technical results of this paper, but their illustration how the classical and alternative theories can use each others' results. 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 [6] 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'.

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

1. C. Boutilier. Toward a logic for qualitative decision theory. In *Proceedings of the KR'94*, pages 75–86, 1994.
2. J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly. Special issue on 'Desires, goals, intentions, and values: Computational architectures'*, 2(3-4):428–447, 2002.
3. P.R. Cohen and H.J. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42:213–261, 1990.
4. M. Dastani and L. van der Torre. Decisions and games for BD agents. In *Proceedings of The Workshop on Game Theoretic and Decision Theoretic Agents, Canada*, pages 37–43, 2002.
5. M. Dastani and L. van der Torre. What is a normative goal? Towards Goal-based Normative Agent Architectures. In *Regulated Agent-Based Systems, Postproceedings of RASTA'02*, pages 210–227. Springer, 2004.
6. D. Dennett. *The intentional stance*. MIT Press, Cambridge, MA, 1987.
7. J. Lang. Conditional desires and utilities - an alternative approach to qualitative decision theory. In *In Proceedings of the European Conference on Artificial Intelligence (ECAI'96)*, pages 318–322, 1996.
8. Martin J. Osborne and Ariel Rubenstein. *A Course in Game Theory*. The MIT Press, Cambridge, Massachusetts, 1994.
9. A. Rao and M. Georgeff. Modeling rational agents within a BDI architecture. In *Proceedings of the KR91*, pages 473–484, 1991.
10. L. Savage. *The foundations of statistics*. Wiley, New York, 1954.
11. H. A. Simon. *The Sciences of the Artificial*. MIT Press, Cambridge, MA, 1981.
12. J. von Neumann and O. Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 1 edition, 1944.