How to decide what to do?

Mehdi Dastani
Utrecht University
The Netherlands
mehdi@cs.uu.nl

Joris Hulstijn
Utrecht University
The Netherlands
jorish@cs.uu.nl

Leendert van der Torre
CWI Amsterdam
The Netherlands
torre@cwi.nl

Abstract

There are many conceptualizations and formalizations of decision making. In this paper we compare classical decision theory with qualitative decision theory, knowledge-based systems and belief-desire-intention models developed in artificial intelligence and agent theory. They all contain representations of information and motivation. Examples of informational attitudes are probability distributions, qualitative abstractions of probabilities, knowledge, and beliefs. Examples of motivational attitudes are utility functions, qualitative abstractions of utilities, goals, and desires. Each of them encodes a set of alternatives to be chosen from. This ranges from a small predetermined set, a set of decision variables, through logical formulas, to branches of a tree representing events through time. Moreover, they have a way of formulating how a decision is made. Classical and qualitative decision theory focus on the optimal decisions represented by a decision rule. Knowledge-based systems and belief-desire-intention models focus on a model of the representations used in decision making, inspired by cognitive notions like belief, desire, goal and intention. Relations among these concepts express an agent type, which constrains the deliberation process. We also consider the relation between decision processes and intentions, and the relation between game theory and norms and commitments.

1 Introduction

There are several conceptualizations and formalizations of decision making. Classical decision theory [30, 45] is developed within economics and forms the main theory of decision making used within operations research. It conceptualizes a decision as a choice from a set of alternative actions. The relative preference for an alternative is expressed by a utility value. A decision is rational when it maximizes expected utility.

Qualitative variants of decision theory [5, 39] are developed in artificial intelligence. They use the same conceptualization as classical decision theory, but preferences are typically uncertain, formulated in general terms, dependent on uncertain assumptions and subject to change. A preference is often expressed in terms of a trade-off.

Knowledge-based systems [37] are developed in artificial intelligence too. They consist of a high-level conceptual model in terms of knowledge and goals of an application domain, such as the medical or legal domain, together with a reusable inference scheme for a task, like classification or configuration. Methodologies for modeling, developing and testing knowledge-based systems in complex organizations have matured, see [46].

Belief-desire-intention models – typically referred to as BDI models – are developed in philosophy and agent theory [7, 13, 15, 31, 42]. They are motivated by applications like robotic planning, which they conceptualize using cognitive concepts like belief, desire and intention. An intention can be interpreted as a previous decision that constrains the set of alternatives from which an agent can choose, and it is therefore a factor to stabilize the decision making behavior through time.

1.1 Distinctions and similarities

In this paper we are interested in relations among the theories, systems and models that explain the decision-making behavior of rational agents. The renewed interest in the foundations of decision making is due to the automation of decision making in the context of tasks like planning, learning, and communication in autonomous systems [5, 7, 14, 17].

The following example of Doyle and Thomason [24] on automation of financial advice dialogues illustrates decision making in the context of more general tasks. A user who seeks advice about financial planning wants to retire early, secure a good pension and maximize the inheritance of her children. She can choose between a limited number of actions: retire at a certain age, invest her savings and give certain sums of money to her children. Her decision can therefore be modeled in terms of the usual decision theoretic parameters. However, she does not know all factors that might influence her decision. She does not know if she will get a pay raise next year, the outcome of her financial actions is uncertain, and her own preferences may not be clear since, for example, securing her own pension conflicts with her children's inheritance. An experienced decision theoretic analyst therefore interactively guides the user through the decision process, indicating possible choices and desirable consequences. As a result the user may drop initial preferences by, for example, preferring to continue working for another five years before retiring.

The most visible distinction among the theories, systems and models is that knowledge-based systems and beliefs-desire-intention models describe decision making in terms of cognitive attitudes such as knowledge, beliefs, desires, goals, and intentions. In the dialogue example, instead of trying to detail the preferences of the user in terms of probability distributions and utility functions, they try to describe her cognitive state.

Moreover, knowledge-based systems and beliefs-desire-intention models focus less on the definition of the optimal decision represented by the decision rule, but instead also discuss the way decisions are reached. They are therefore sometimes identified with theories of deliberation instead of decision theories [16, 17]. However, as illustrated by the dialogue example, in classical decision theory the way to reach optimal decisions has also been studied in decision theoretic practice called decision analysis.

Other apparent distinctions can be found by studying the historic development of the various conceptualizations and formalizations of decision making. After the introduction of classical decision theory, it was soon criticized by Simon's notion of limited or bounded rationality, and his introduction of utility aspiration levels [49]. This has led to the notion of a goal in knowledge-based systems. The research area of qualitative decision theory developed much more recently out of research on reasoning under uncertainty. It focusses on theoretical models of decision making with potential applications in planning. The research area of belief-desire-intention models developed out of philosophical arguments that – besides the knowledge and goals used in knowledge-based systems – also intentions should be first class citizens of a cognitive theory of deliberation.

The example of automating financial advice dialogues also illustrates some criticism on classical decision theory. According to Doyle and Thomason, the interactive process of preference elicitation cannot be automated in decision theory itself, although they acknowledge the approaches and methodologies available in decision theoretic practice. For example, they suggest that it is difficult to describe the alternative actions to decide on, and that classical decision theory is not suitable to model generic preferences.

A historical analysis may reveal and explain apparent distinctions among the theories, systems and models, but its also hides the similarities among them. We therefore adopt another methodology for our comparison. We choose several representative theories for each tradition, and look for similarities and differences between these particular theories.

1.2 Representative theories

For the relation between classical and qualitative decision theory we discuss the work of Doyle and Thomason [24] and Pearl [39]. For the relation between qualitative decision theory and knowledge-based systems and belief-desire-intention models we focus on the different interpretations of goals in the work of Boutilier [5] and Rao and Georgeff [42]. For the direct relation between classical decision theory and belief-desire-intention models we discuss Rao and Georgeff's translation of decision trees to belief-desire-intention models [41].

Clearly the results of this comparison between representative theories and systems cannot be generalized directly to a comparison between research areas. Moreover, the discussion in this paper cannot do justice to the subtleties defined in each approach. We therefore urge the reader to read the original papers. However, this comparison gives some interesting insights into the relation among the areas, and these insights are a good starting point for further and more complete comparisons.

A summary of the comparison is given in Table 1. In our comparison, some concepts can be mapped easily onto concepts of other theories and systems. For example, all theories and systems use some kind of informational attitude (probabilities, qualitative abstractions of probabilities, knowledge or beliefs) and some kind of motivational attitude (utilities, qualitative abstractions of utilities, goals or desires). Other concepts are more ambiguous, such as intentions. In goal-based planning for example, goals have both a desiring and an intending aspect [22]. Some qualitative decision theories like [5] have been developed as a criticism to the inflexibility of the notion of goal in goal-based planning.

	classical decision theory (CDT)	qualitative decision theory (QDT)	knowledge-based systems (KBS / BDI)
underlying concepts	probability function utility function decision rule	likelihood ordering preference ordering decision criterion	knowledge / belief goal / desire agent type / deliberation
time	(Markov) decision processes	decision-theoretic planning	belief-desire-intention models & systems
multiagent	classical game theory	qualitative game theory	normative systems (BOID)

Table 1: Theories, systems and models discussed in this paper

The table also illustrates that we discuss two extensions of classical decision theory in this paper. In particular, we consider the relation between decision processes and intentions, and the relation between game theory and the role of norms and commitments in belief-desire-intention models. Our discussion of time and decision processes focusses on the role of intentions in Rao and Georgeff's work [42] and our discussion on multiple agents and game theory focusses on the role of norms in a logic of commitments [9].

The relations between the areas may suggest a common underlying abstract theory of the decision making process, but our comparison does not suggest that one approach can be exchanged for another one. Due to the distinct motivations of the areas, and probably due also to the varying conceptualizations and formalizations, the areas have sometimes studied distinct elements of the decision making process. Our comparison therefore not only considers the similarities, but we also discuss some distinctions which suggests ways for further research to incorporate results of one area into another one.

We discuss qualitative decision theory in more detail than knowledge-based systems and belief-desire-intention models, because it is closer to classical decision theory and has been positioned as an intermediary between classical decision theory and the others [24]. Throughout the paper we restrict ourselves to formal theories and logics, and do not go into system architectures or into the philosophical motivations of the underlying cognitive or social concepts.

The layout of this paper is as follows. In Section 2 we discuss classical and qualitative decision theory. In Section 3 we discuss goals in qualitative decision theory, knowledge-based systems and belief-desire-intention models. In Section 4 we compare classical decision theory and Rao and Georgeff's belief-desire-intention model. Finally, in Section 5 we discuss intentions and norms in extensions of classical decision theory that deal with time by means of processes, and that deal with multiple agents by means of game theory.

2 Classical versus qualitative decision theory

In this section we compare classical and qualitative decision theory, based on Doyle and Thomason's introduction to qualitative decision theory [24] and Pearl's qualitative decision theory [39].

2.1 Classical decision theory

In classical decision theory, a decision is the selection of an action from a set of alternative actions. Decision theory does not have much to say about actions – neither about their nature nor about how a set of alternative actions becomes available to the decision maker. A decision is good if the decision maker believes that the selected action will prove at least as good as the other alternative actions. A good decision is formally characterized as the action that maximizes expected utility, a notion which involves both belief and desirability. See [30, 45] for further explanations on the foundations of decision theory.

Definition 1 *Let A stand for a set of alternative actions. With each action, a set of outcomes is associated. Let W stand for the set of all possible worlds or outcomes.*[1] *Let U be a measure of outcome value that assigns a utility $U(w)$ to each outcome $w \in W$, and let P be a measure of the probability of outcomes conditional on actions, with $P(w|a)$ denoting the probability that outcome w comes about after taking action $a \in A$ in the situation under consideration.*

The expected utility $EU(a)$ of an action a is the average utility of the outcomes associated with the action, weighing the utility of each outcome by the probability that the outcome results from the action, that is, $EU(a) = \sum_{w \in W} U(w)P(w|a)$. A rational decision maker will always maximize expected utility, i.e., it selects action a from the set of alternative actions A such that for all actions b in A we have $EU(a) \geq EU(b)$. This decision rule is called maximization of expected utility and typically referred to as MEU.

Many variants and extensions of classical decision theory have been developed. For example, in some presentations of classical decision theory, not only uncertainty about the effect of actions is considered, but also uncertainty about the present state. A classic result is that uncertainty about the effects of actions can be expressed in terms of uncertainty about the present state. Moreover, several other decision rules have been investigated, including qualitative ones, such as Wald's criterion of maximization of the utility of the worst possible outcome. Finally, classical decision theory has been extended in various ways to deal with multiple objectives, sequential decisions, multiple agents and notions of risk. The extensions with sequential decisions and multiple agents are discussed in section 5.1 and 5.2.

Decision theory has become one of the main foundations of economic theory due to so-called representation theorems, such as the famous one by Savage [45]. It shows that each decision maker obeying certain plausible postulates (about weighted choices) acts *as if* he were applying the MEU decision rule with some probability distribution and utility function. Thus, the decision maker does not have to be aware of it and the utility function does not have to represent selfishness. In fact, altruistic decision makers also act as if they were maximizing expected utility. They only use another utility function than selfish decision makers.

[1] Note that outcomes are usually represented by Ω. Here we use W to facilitate our comparison.

2.2 Qualitative decision theory

According to Doyle and Thomason [24, p.58], quantitative representations of probability and utility and procedures for computing with these representations do provide an adequate framework for manual treatment of simple decision problems, but are less successful in more realistic cases. They suggest that classical decision theory does not address decision making in unforeseen circumstances, offers no means for capturing generic preferences, provides little help to decision makers who exhibit discomfort with numeric trade offs, and provides little help in effectively representing decisions involving broad knowledge of the world.

Doyle and Thomason therefore argue for a number of new research issues: formalization of generic probabilities and generic preferences, properties of the formulation of a decision problem, mechanisms for providing reasons and explanations, revision of preferences, practical qualitative decision-making procedures and agent modeling. Moreover, they argue that hybrid reasoning with quantitative and qualitative techniques, as well as reasoning within context, deserve special attention. Many of these issues are studied in artificial intelligence. It appears that researchers now realize the need to reconnect the methods of artificial intelligence with the qualitative foundations and quantitative methods of economics.

First results have been obtained in the area of reasoning under uncertainty, a subdomain of artificial intelligence which mainly attracts researchers with a background in nonmonotonic reasoning. Often the formalisms of reasoning under uncertainty are reapplied in the area of decision making. Typically uncertainty is not represented by a probability function, but by a plausibility function, a possibilistic function, Spohn-type rankings, etc. Another consequence of this historic development is that the area of qualitative decision theory is more mathematically oriented than the knowledge-based systems or the belief-desire-intention community.

The representative example we use in our first comparison is the work of Pearl [39]. A so-called semi-qualitative ranking $\kappa(w)$ can be considered as an order-of-magnitude approximation of a probability function $P(w)$ by writing $P(w)$ as a polynomial of some small quantity ϵ and by taking the most significant term of that polynomial. Similarly, a ranking $\mu(w)$ can be considered as an approximation of a utility function $U(w)$. There is one more subtlety here. Whereas κ rankings are positive, the μ rankings can be either positive or negative. This represents the fact that outcomes can be either very desirable or very undesirable.

Definition 2 *A belief ranking function $\kappa(w)$ is an assignment of non-negative integers to outcomes or possible worlds $w \in W$ such that $\kappa(w) = 0$ for at least one world. Intuitively, $\kappa(w)$ represents the degree of surprise associated with finding a world w realized, and worlds assigned $\kappa(w) = 0$ are considered serious possibilities. Likewise, $\mu(w)$ is an integer-valued utility ranking of worlds. Moreover, both probabilities and utilities are defined as a function of the same ϵ, which is treated as an infinitisimal quantity (smaller than any real number). C is a constant and O is the order of magnitude.*

$$
\begin{array}{rcll}
P(w) & \sim & C\epsilon^{\kappa(w)}, & \\
U(w) & = & O(1/\epsilon^{\mu(w)}), & \textit{if } \mu(w) \geq 0, \\
 & & -O(1/\epsilon^{-\mu(w)}), & \textit{otherwise.}
\end{array} \tag{1}
$$

This definition illustrates the use of abstractions of probabilities and utilities. However, we still have to relativize the probability distribution, and therefore the expected utility, to actions. This is more complex than in classical decision theory, and is discussed in the following section.

2.3 Relation

We first discuss similarities between the set of alternatives and the decision rules to select the optimal action. Then we discuss an apparent distinction between the two approaches.

2.3.1 Alternatives

In classical decision problems the alternative actions typically correspond to a few atomic variables, whereas Pearl assumes a set of actions of the form 'Do(φ)' for every proposition φ. That is, where in classical decision theory we defined $P(w|a)$ for alternatives a in A and worlds w in W, in Pearl's approach we write $P(w|Do(\varphi))$ or simply $P(w|\varphi)$ for any proposition φ. In Pearl's semantics such an alternative can be identified with the set of worlds that satisfy φ, since a valuation function assigns a truth value to every proposition at each world of W. We could therefore also write $P(w|V)$ with $V \subseteq W$.

Consequently, examples formalized in Pearl's theory typically consider much more alternatives than examples formalized in classical decision theory. However, the set of alternatives of both theories can easily be mapped on each other. Classical decision theory also works well with a large number of atomic variables, and the set of alternatives in Pearl's theory can be restricted by adding logical constraints to the alternatives.

2.3.2 Decision rule

Both classical decision theory as presented in Definition 1 and Pearl's qualitative decision theory as presented in Definition 2 can deal with trade-offs between normal situations and exceptional situations. The decision rule from Pearl's theory differs from decision criterion such as 'maximize the utility of the worst outcome'. This qualitative decision rule of classical decision theory has been used in purely qualitative decision theory of Boutilier [5] which is discussed in the following section. The decision criteria from purely qualitative decision theories do not seem to be able to make trade-offs between such alternatives.

The problem with a purely qualitative approach is that it is unclear how, besides the most likely situations, also less likely situations can be taken into account. We are interested in situations which are unlikely, but which have a high impact, i.e., an extremely high or low utility. For example, the probability that your house will burn down is very small, but it is also very unpleasant. Some people therefore decide to take an insurance. In a purely qualitative setting there does not seem to be an obvious way to compare a likely but mildly important effect to an unlikely but important effect. Going from quantitative to qualitative we may have gained computational efficiency, but we seem to have lost one of the useful properties of decision theory.

The ranking order solution proposed by Pearl is based on two ideas. First, the initial probabilities and utilities are neither represented by quantitative probability distributions and utility functions, nor by pure qualitative orders, but by a *semi-qualitative* order in between. Second, the two semi-qualitative functions are assumed to be comparable in a suitable sense. This is called the *commensurability assumption* [26].

Consider for example likely and moderately interesting worlds ($\kappa(w) = 0, \mu(w) = 0$) or unlikely but very important worlds ($\kappa(w) = 1, \mu(w) = 1$). These cases have become comparable. Although Pearl's order of magnitude approach can deal with trade-offs between normal and exceptional circumstances, it is less clear how it can deal with trade-offs between two effects under normal circumstances.

2.3.3 A distinction and a similarity

Pearl explains that in his setting the expected utility of a proposition φ depends on how we came to know φ. For example, if we find the ground wet, it matters whether we happened to find the ground wet (observation) or watered the ground (action). In the first case, finding φ true may provide information about the natural process that led to the observation φ, and we should change the current probability from $P(w)$ to $P(w|\varphi)$. In the second case, our actions may perturb the natural flow of events, and $P(w)$ will change without shedding light on the typical causes of φ. This is represented differently, by $P_\varphi(w)$. According to Pearl, the distinction between $P(w|\varphi)$ and $P_\varphi(w)$ corresponds to distinctions found in a variety of theories, such as the distinction between conditioning and imaging [36], between belief revision and belief update, and between indicative and subjunctive conditionals. However, it does not seem to correspond to a distinction in classical decision theory, although it may be related to discussions in the context of the logic of decision [30]. One of the tools Pearl uses for the formalization of this distinction are causal networks: a kind of Bayesian networks with actions.

A similarity between the two theories is that both suppress explicit reference to time. In this respect Pearl is inspired by deontic logic, the logic of obligations and permissions discussed in Section 5.2. Pearl suggests that his approach differs in this respect from other theories of action in planning and knowledge-based systems, since they are normally formulated as theories of temporal change. Such theories are discussed in the comparison in the following section.

3 Qualitative decision theory versus BDI logic

In this section we give a comparison between qualitative decision theory and knowledge-based systems and belief-desire-intention models, based on their interpretation of beliefs and goals. We use representative qualitative theories that are defined on possible worlds, namely Boutilier's version of qualitative decision theory [5] and Rao and Georgeff's belief-desire-intention logic [41, 43, 44].

3.1 Qualitative decision theory (continued)

Boutilier's qualitative decision theory [5] may be called purely qualitative, because its semantics does not contain any numbers, only more abstract preference relations. It is developed in the context of planning. Goals serve a dual role in most planning systems, capturing aspects of both desires towards states and commitment to pursuing that state [22]. In goal-based planning, adopting a proposition as a goal commits the agent to find some way to accomplish the goal, even if this requires adopting subgoals that may not correspond to desirable propositions themselves [19]. Context-sensitive goals are formalized with basic concepts from decision theory [5, 19, 25]. In general, goal-based planning must be extended with a mechanism to choose which goals must be adopted. To this end Boutilier proposes a logic for representing and reasoning with qualitative probabilities and utilities, and suggests several strategies for qualitative decision making based on this logic.

The MEU decision rule is replaced by a qualitative rule, for example by Wald's criterion. Conditional preference is captured by a preference ordering (an ordinal value function) defined on possible worlds. The preference ordering represents the relative desirability of worlds. Boutilier says that $w \leq_P v$ when w is at least as preferred as v, but possibly more. Similarly, probabilities are captured by a normality ordering $\leq_N$ on possible worlds, which represents their relative likelihood.

Definition 3 *The semantics of Boutillier's logic is based on models of the form*

$$M = \langle W, \leq_P, \leq_N, V \rangle \tag{2}$$

where W is a set of possible worlds (outcomes), $\leq_P$ is a reflexive, transitive and connected preference ordering relation on W, $\leq_N$ is a reflexive, transitive and connected normality ordering relation on W, and V is a valuation function.

Conditional preferences are represented in the logic by means of modal formulas $\mathcal{I}(\varphi|\psi)$, to be read as 'ideally φ if ψ'. A model M satisfies the formula $\mathcal{I}(\varphi|\psi)$ if the the most preferred or minimal ψ worlds with respect to $\leq_P$ are φ worlds. For example, let u be the proposition 'the agent carries an umbrella' and r be the proposition 'it is raining', then $\mathcal{I}(u|r)$ expresses that in the most preferred rain-worlds the agent carries an umbrella. Similar to preferences, probabilities are represented in the logic by a default conditional $\Rightarrow$. For example, let w be the proposition 'the agent is wet' and r be the proposition 'it is raining', then $r \Rightarrow w$ expresses that the agent is wet at the most normal rain-worlds. The semantics of this operator is used in Hansson's deontic logic [27] for a modal operator O to model obligation, and by Lang [33] for a modal operator D to model desire. Whereas in default logic an exception is a digression from a default rule, in deontic logic an offense is a digression from the ideal. An alternative approach represents conditional modalities by so called ‘ceteris paribus’ preferences, using additional formal machinery to formalize the notion of ‘similar circumstances’, see, e.g., [23, 25, 50, 51].

In general, a goal is any proposition that the agent attempts to make true. A rational agent is assumed to attempt to reach the most preferred worlds consistent with its default knowledge. Given the ideal operator and the default conditional, a goal is defined as follows.

Definition 4 *Given a set of facts KB, a goal is any proposition φ such that*

$$M \models \mathcal{I}(\varphi \mid Cl(KB)) \tag{3}$$

where $Cl(KB)$ is the default closure of the facts KB defined as follows:

$$Cl(KB) = \{\varphi \mid KB \Rightarrow \varphi\} \tag{4}$$

3.2 BDI logic

According to Dennett [20], attitudes like belief and desire are folk psychology concepts that can be fruitfully used in explanations of rational behavior. If you were asked to explain why someone is carrying an umbrella, you may reply that he believes it is going to rain and that he does not want to get wet. For the explanation it does not matter whether he actually possesses these mental attitudes. Similarly, we describe the behavior of an affectionate cat or an unwilling screw in terms of mental attitudes. Dennett calls treating a person or artifact as a rational agent the ‘intentional stance’.

> “Here is how it works: first you decide to treat the object whose behavior is to be predicted as a rational agent; then you figure out what beliefs that agent ought to have, given its place in the world and its purpose. Then you figure out what desires it ought to have, on the same considerations, and finally you predict that this rational agent will act to further its goals in the light of its beliefs. A little practical reasoning from the chosen set of beliefs and desires will in most instances yield a decision about what the agent ought to do; that is what you predict the agent will do.” [20, p. 17]

In this tradition, knowledge (K) and beliefs (B) represent the information of an agent about the state of the world. Belief is like knowledge, except that it does not have to be true. Goals (G) or desires (D) represent the preferred states of affairs for an agent. The terms goal and desire are sometimes used interchangeably. In other cases, a desire is treated like a goal, except that sets of desires do not have to be mutually consistent. Desires are long term preferences that motivate the decision process. Intentions (I) correspond to previously made commitments of the agent, either to itself or to others.

As argued by Bratman [7], intentions are meant to stabilize decision making. Consider the following application of a lunar robot. The robot is supposed to reach some destination on the surface of the moon. Its path is obstructed by a rock. Suppose that based on its cameras and other sensors, the robot decides that it will go around the rock on the left. At every step the robot will receive new information through its sensors. Because of shadows, rocks may suddenly appear much larger. If the robot were to reconsider its decision with every new piece of information, it would never reach its destination. Therefore, the agent will adopt a plan until some really strong reason forces it to change it. The intentions of an agent correspond to the set of adopted plans at some point in time.

Belief-desire-intention models, better known as BDI models, are applied in for example natural language processing and the design of interactive systems. The theory of speech acts [3, 47] and subsequent applications in artificial intelligence [1, 14] analyze the meaning of an utterance in terms of its applicability and sincerity conditions and the intended effect. These conditions are best expressed using belief or knowledge, desire or goal, and intention. For example, a question is applicable when the speaker does not yet know the answer and the hearer is expected to know the answer. A question is sincere if the speaker actually desires to know the answer. By the conventions encoded in language, the effect of a question is that it signals the intention of the speaker to let the hearer know that the speaker desires to know the answer. Now if we assume that the hearer is cooperative, which is a reasonable assumption for interactive systems, the hearer will adopt the goal to let the speaker know the answer to the question and will consider plans to find and formulate such answers. In this way, traditional planning systems and natural language communication can be combined. For example, Sadek [8] describes the architecture of a spoken dialogue system that assists the user in selecting automated telephone services like the weather forecast, directory services or collect calls. According to its developers the advantage of the BDI specification is its flexibility. In case of a misunderstanding, the system can retry and reach its goal to assist the user by some other means. This specification in terms of BDI later developed into the standard for agent communication languages endorsed by FIPA. If we want to automate parts of the interactive process of decision making, such a flexible way to deal with interaction is required.

As a typical example of a formal BDI model, we discuss Rao and Georgeff's initial BDI logic [42]. The partial information on the state of the environment, which is represented by quantitative probabilities in classical decision theory and by a qualitative ordering in qualitative decision theory, is now reduced to binary values (0-1). This abstraction of the partial information on the state of the environment models the beliefs of the decision making agent. Similarly, the partial information about the objectives of the decision making agent, which is represented by quantitative utilities in classical decision theory and by qualitative preference ordering in qualitative decision theory, is reduced to binary values (0-1). The abstraction of the partial information about the objectives of the decision making agent, models the desires of the decision making agent. The BDI logic has a complicated semantics, using Kripke structures with accessibility relations for each modal operator B, D and I. Each accessibility relation $\mathcal{B}$, $\mathcal{D}$, and $\mathcal{I}$ maps a world w at a time point t to those worlds, which are indistinguishable with respect to respectively the belief, desire or intention formulas that can be satisfied.

Definition 5 (Semantics of BDI logic [42]) *An interpretation M [2] is defined to be a tuple $M = \langle W, E, T, <, U, \mathcal{B}, \mathcal{D}, \mathcal{I}, \Phi \rangle$, where W is the set of worlds, E is the set of primitive event types, T is a set of time points, $<$ a binary relation on time points, U is the universe of discourse, and Φ [3] is a mapping from first-order entities to elements in U for any given world and time point. A situation is a world, say w, at a particular time point, say t, and is denoted by w_t. The relations $\mathcal{B}$, $\mathcal{D}$[4], and $\mathcal{I}$ map the agent's current belief, desire, and intention accessible worlds, respectively. I.e. $\mathcal{B} \subseteq W \times T \times W$ and similarly for $\mathcal{D}$ and $\mathcal{I}$.*

Again there is a logic to reason about these mental attitudes. We can only represent monadic expressions like $B(\varphi)$ and $D(\varphi)$, and no dyadic expressions like Boutilier's $\mathcal{I}(\varphi|\psi)$. Note that the I modality has been used by Boutilier for ideality and by Rao and Georgeff for intention; we use their original notation since it does not lead to any confusion in this paper. A world at a time point of the model satisfies $B(\varphi)$ if φ is true in all belief accessible worlds at the same time point. The same holds for desire and intention. All desired worlds are equally good, so an agent will try to achieve any of the desired worlds.

Compared to the other approaches discussed so far, Rao and Georgeff introduce a temporal aspect. The BDI logic is an extension of the so-called computational tree logic (CTL^*), which is often used to model a branching time structure, with modal epistemic operators for beliefs B, desires D, and intentions I. The modal epistemic operators are used to model the cognitive state of a decision making agent, while the branching time structure is used to model possible events that could take place at a certain time point and determines the alternative worlds at that time point.

Each time branch represents an event and determines an alternative situation. The modal epistemic operators have specific properties such as closure under implication and consistency (KD axioms). Like in CTL, the BDI logic has two types of formula. The first is called a state formula, and is evaluated at a situation. The second is called a path formula, and is evaluated along a path originating from a given world. Therefore, path formulae express properties of alternative worlds through time.

Definition 6 (Semantics of Tree Branch [42]) *Let $M = \langle W, E, T, <, U, \mathcal{B}, \mathcal{D}, \mathcal{I}, \Phi \rangle$ be an interpretation, $T_w \subseteq T$ be the set of time points in the world w, and $\mathcal{A}_w$ be the same relation as $<$ restricted to time points in T_w. A full path in a world w is an infinite sequence of time points $(t_0, t_1, \ldots)$ such that $\forall i\ (t_i, t_{i+1}) \in \mathcal{A}_w$. A full path can be written as $(w_{t_0}, w_{t_1}, \ldots)$.*

In order to give examples of how state and path formulae are evaluated, let $M = \langle W, E, T, <, U, \mathcal{B}, \mathcal{D}, \mathcal{I}, \Phi \rangle$ be an interpretation, $w, w' \in W$, $t \in T$, $(w_{t_0}, w_{t_1}, \ldots)$ be a full path, and $\mathcal{B}_t^w$ be set of belief accessible from world w at time t. Let B be the modal epistemic operator, $\diamondsuit$ the temporal eventually operator, and φ be a state formula. Then, the state formula $B\varphi$ is evaluated relative to the interpretation M and situation w_t as follows:

$$M, w_t \models B\varphi \Leftrightarrow \forall w' \in \mathcal{B}_t^w\ M, w'_t \models \varphi \quad (5)$$

[2] The interpretation M is usually called model M.

[3] The mapping Φ is usually called valuation function represented by V.

[4] In their definition, they use $\mathcal{G}$ for goals instead of $\mathcal{D}$ for desires.

A path formula $\Diamond\varphi$ is evaluated relative to the interpretation M along a path $(w_{t_0}, w_{t_1}, \ldots)$ as follows:

$$M, (w_{t_0}, w_{t_1}, \ldots) \models \Diamond\varphi \Leftrightarrow \exists k \geq 0 \text{ such that } M, (w_{t_k}, \ldots) \models \varphi \quad (6)$$

3.3 Comparison

As in the previous comparison, we compare the set of alternatives, decision rules, and distinctions particular to these approaches.

3.3.1 Alternatives

Boutilier [5] introduces a simple but elegant distinction between consequences of actions and consequences of observations, by distinguishing between controllable and uncontrollable propositional atoms. Formulas φ built from controllable atoms correspond to actions $Do(\varphi)$. Boutilier does not study the distinction between actions and observations, and he does not introduce a causal theory. His action theory is therefore simpler than Pearl's.

BDI on the other hand does not involve an explicit notion of action, but instead models possible events that can take place. Events in the branching time structure determine the alternative (cognitive) worlds that an agent can reach. Thus, each branch represents an alternative the agent can select. Uncertainty about the effects of actions is not modeled by branching time, but by distinguishing between different belief worlds. So all uncertainty about the effects of actions is modeled as uncertainty about the present state; a well known trick from decision theory we already mentioned in section 2.1.

The problem of mapping the two ways of representing alternatives onto each other is due to the fact that in Boutilier's theory there is only a single decision, whereas in BDI models there are decisions at any world-time pair. If we consider only a single world-time pair, for example the present one, then each attribution of truth values to controllable atoms corresponds to a branch, and for each branch a controllable atom can be introduced together with the constraint that only one controllable atom may be true at the same time.

3.3.2 Decision rules

The qualitative normality and the qualitative desirability orderings on possible worlds that are used in qualitative decision theory are reduced to binary values in belief-desire-intention models. Based on the normality and preference orderings, Boutilier uses a qualitative decision rule like the Wald criterion. Since there is no ordering in BDI models, each desired world can in principle be selected as a goal world to be achieved. However, it is not intuitive to select any desired world as a goal, since a desired world is not necessarily believed to be possible. Selecting a desired world which is not believed to be possible, results in wishful thinking [52] and therefore in unrealistic decision making.

Therefore, BDI proposes a number of constraints on the selection of goal worlds. These constraints are usually characterized by axioms called realism, strong realism or weak realism [11, 44]. Roughly, realism states that an agent's desires should be consistent with its beliefs. Note that this constraint is the same in qualitative decision theories where goal worlds should be consistent with the belief worlds. Formally, the realism axiom states that something which is believed is also desired, or that the set of desire accessible worlds is a subset of the set of belief accessible worlds, i.e.,

$$B(\varphi) \rightarrow D(\varphi) \tag{7}$$

and, moreover, that belief and desire worlds should have identical branching time structure, i.e.,

$$\forall w, v \in W, \forall t \in T \quad \text{if } v \in \mathcal{D}_t^w \text{ then } v \in \mathcal{B}_t^w \tag{8}$$

A set of such axioms to constrain the relation between beliefs, desires, and alternatives determines an agent type. For example, we can distinguish realistic agents from unrealistic agents. BDI systems do not consider decision rules but agent types. Although there are no agent types in classical or qualitative decision theory, there are discussions which can be related to agent types. For example, often a distinction is made between risk neutral, risk seeking, and risk averse behavior.

In Rao and Georgeff's BDI theory, additional axioms are introduced for intentions. Intentions can be seen as previous decisions. These further reduce the set of desire worlds that can be chosen as a goal world. The axioms guarantee that a chosen goal world is consistent with beliefs and desires. The definition of realism therefore includes the following axiom, stating that intention accessible worlds should be a subset of desire accessible worlds,

$$D(\varphi) \rightarrow I(\varphi) \tag{9}$$

and, moreover, that desire and intention worlds should have an identical branching time structure (have the same alternatives), i.e.

$$\forall w \forall t \forall w' \text{ if } w' \in \mathcal{I}_t^w \text{ then } w' \in \mathcal{D}_t^w \tag{10}$$

In addition to these constraints, which are classified as static constraints, there are dynamic constraints introduced in BDI resulting in additional agent types. These axioms determine when intentions or previously decided goals should be reconsidered or dropped. These constraints, called commitment strategies, involve time and intentions and express the dynamics of decision making. The well-known commitment strategies are 'blindly committed decision making', 'single-minded committed decision making', and 'open-minded committed decision making'. For example, the single-minded commitment strategy states that an agent remains committed to its intentions until either it achieves its corresponding objective or does not believe that it can achieve it anymore. The notion of an agent type has been refined and it has been extended to include obligations in Broersen *et al.*'s BOID system [10]. For example, they distinguish selfish agents, that give priority to their own desires, and social agents, that give priority to their obligations.

3.3.3 Two Steps

A similarity between the two approaches is that we can distinguish two steps. In Boutilier's approach, decision-making with flexible goals has split the decision-making process. First a decision is made which goals to adopt, and second a decision is made how to reach these goals. These two steps have been further studied by Thomason [52] and Broersen *et al.* [10] in the context of default logic.

1. First, the agent has to combine desires and resolve conflicts between them. For example, assume that the agent desires to be on the beach, if he is on the beach then he desires to eat an ice-cream, he desires to be in the cinema, if he is in the cinema then he desires to eat popcorn, and he cannot be at the beach as well as in the cinema. Now he has to choose one of the two combined desires as a potential goal: being at the beach with ice-cream or being in the cinema with popcorn.

2. Second, the agent has to find out which actions or plans can be executed to reach the goal, and he has to take all side-effects of the actions into account. For example, assume that he desires to be on the beach, if he will quit his job and drive to the beach, he will be on the beach, if he does not have a job he will be poor, if he is poor then he desires to work. The only desire and thus a potential goal is to be on the beach, the only way to reach this goal is to quit his job, but the side effect of this action is that he will be poor and in that case he does not want to be on the beach but he wants to work.

Now crucially, desires come into the picture two times! First they are used to determine the goals, and second they are used to evaluate the side-effects of the actions to reach these goals. In extreme cases, like the example above, what seemed like a goal may not be desirable, because the only actions to reach the goal have negative effects with much more impact than the original goal.

At first sight, it seems that we can apply classical decision theory to each of these two sub-decisions. However, there is a caveat. The two sub-decisions are not independent, but closely related. For example, to decide which goals to adopt we must know which goals are feasible, and we thus have to take the possible actions into account. Moreover, previously intended actions constrain the candidate goals which can be adopted. Other complications arise due to many factors such as uncertainty, changing environments, etc. We conclude here that the role of decision theory in planning is complex, and that decision-theoretic planning is much more complex than classical decision theory since the interaction between goals and actions in classical decision theory is predefined while in qualitative decision theory this interaction is the subject of reasoning. For more on this topic, see [6].

3.3.4 Goals versus desires

A distinction between the two approaches is that Boutilier distinguishes between ideality statements or desires and goals, whereas Rao and Georgeff do not. In Boutilier's logic, there is a formal distinction between preference ordering and goals expressed by ideality statements. Rao and Georgeff have unified these two notions, which has been criticized by [18]. In decision systems such as [10], desires are considered to be more primitive than goals, because goals have to be *adopted* or *generated* based on desires. Moreover, goals can be based on desires, but also on other sources. For example, a social agent may adopt his obligations as a goal, or the desires of another agent. In many theories desires or candidate goals can be mutually conflicting, but other notions of goals have been considered, in which goals do not conflict. In that case goals are more similar to intentions. There are three main traditions. In the Newell and Simon tradition of knowledge-based systems, goals are related to utility aspiration levels and to limited (bounded) rationality. In this tradition goals have an aspect of desiring as well as an aspect of intending. In the more recent BDI tradition knowledge and goals have been replaced by beliefs, desires and intentions due to Bratman's work on the role of intentions in deliberation process [7]. The third tradition relates desires and goals to utilities in classical decision theory. The problem here is that decision theory abstracts away from the deliberation cycle. Typically, Savage-like constructions only consider the input (state of the world) and output (action) of an agent. Consequently, utilities can be related to both stages in the process, represented by either desires or goals.

3.3.5 Conflict resolution

A similarity between the two logics is that both are not capable of representing conflicts, either conflicting beliefs or conflicting desires.

Although the constraints imposed by the Boutilier's $\mathcal{I}$ operator are rather weak, they are still too strong to represent certain types of conflicts. Consider conflicts among desires. Typically desires are allowed to be inconsistent, but once they are adopted and have become intentions, they should be consistent. Several potential conflicts between desires, including a classification and ways to resolve it, is given in [34]. A different approach to solving conflicts is to apply Reiter's default logic to create extensions. This is recently proposed by Thomason [52] and used in the BOID architecture [10].

Finally, an important branch of decision theory has to do with reasoning about multiple objectives, which may conflict, by means of multiple attribute utility theory [32]. This is also the basis of the theory of 'ceteris paribus' preferences mentioned in previous section. It can be used to formalize conflicting desires. By contrast, all the modal logic approaches above would make conflicting desires inconsistent. Clearly, if we continue to follow the financial advice example of Doyle and Thomason, conflicting desires must be dealt with.

3.3.6 Non-monotonic closure rules

A distinction between the logics is that Rao and Georgeff only present a monotonic logic, whereas Boutilier also presents a non-monotonic extension. The constraints imposed by the $\mathcal{I}$ formulas of Boutilier are relatively weak. Since the semantics of the Boutilier's $\mathcal{I}$ operator is analogous to the semantics of many default logics, Boutilier [5] proposes to use non-monotonic closure rules for the $\mathcal{I}$ operator too. In particular he uses the well-known system Z [38]. Its workings can be summarized as 'gravitation towards the ideal', in this case. An advantage of this system is that it always gives exactly one preferred model, and that the same logic can be used for both desires and defaults. A variant of this idea was developed by Lang [33], who directly associates penalties with desires (based on penalty logic [40]) and who does not use rankings of utility functions but utility functions themselves. More complex constructions have been discussed in [35, 50, 51, 54].

4 Classical decision theory versus BDI logic

In this section we compare classical decision theory to BDI theory. Thus far, we have seen a quantitative ordering in classical decision theory, a semi-qualitative and qualitative ordering in qualitative decision theory, and binary values in BDI. Classical decision theory and BDI thus seem far apart, and the question can be raised how they can be related. This question has been ignored in the literature, except by Rao and Georgeff's translation of decision trees to beliefs and desires in [41]. Rao and Georgeff show that constructions like subjective probability and subjective utility can be recreated in the setting of their BDI logic to extend its expressive power and to model the process of deliberation. The result shows that the two approaches are compatible. In this section we sketch their approach.

4.1 BDI, continued

Rao and Georgeff extend the BDI logic by introducing probability and utility functions in their logic. The intuition is formulated as follows:

> "Intuitively, an agent at each situation has a probability distribution on his belief-accessible worlds. He then chooses sub-worlds of these that he considers are worth pursuing and associates a payoff value with each path in these sub-worlds. These sub-worlds are considered to be the agent's goal accessible worlds. By making use of the probability distribution on his belief-accessible worlds and the payoff distribution on the paths in his goal-accessible worlds, the agent determines the best plan(s) of action for different scenarios. This process will be called *Possible-Worlds (PW) deliberation*. The result of PW-deliberation is a set of sub-worlds of the goal-accessible worlds; namely, the ones that the agent considers best. These sub-worlds are taken to be the intention-accessible worlds that the agent *commits* to achieving." [41, p. 301]

In this extension of the BDI logic two operators for probability and utility are introduced. Formally, if $\varphi_1, \ldots, \varphi_k$ are state formulas, $\psi_1, \ldots, \psi_k$ are state formulas, and $\theta_1, \ldots, \theta_k, \alpha$ are real numbers, then $\theta_1 PROB(\varphi_1) + \ldots + \theta_1 PROB(\varphi_1) \geq \alpha$ and $\theta_1 PAYOFF(\psi_1) + \ldots + \theta_1 PAYOFF(\psi_1) \geq \alpha$ are state formulas. Consequently, the semantics of the BDI logic is extended by adding semantic structures to represent probabilities and utilities.

Definition 7 (Extended BDI models [41]) *The semantics of the extended BDI logic is based on interpretation M of the following form:*

$$M = \langle W, E, T, <, \mathcal{B}, \mathcal{D}, \mathcal{I}, PA, OA, \Phi \rangle \quad (11)$$

where W, E, T, $<$, $\mathcal{B}, \mathcal{D}, \mathcal{I}$ and Φ are as in definition 5 [5]*. PA is a probability assignment function that assigns to each time point t and world w a probability distribution η_t^w* [6]*. Each η_t^w is a discrete probability function on the set of worlds W. Moreover, OA is a utility assignment function that assigns to each time point t and world w a utility function ρ_t^w. Each ρ_t^w is a partial mapping from paths to real-valued numbers.*

Given a state formula φ and a path formula ψ, the semantics of the extended BDI language extends the semantics of the BDI langauge with the following two evaluation clauses for the PROB and PAYOFF expressions.

$$\begin{aligned} M, w_{t_0} \models \text{PROB}(\varphi) \geq \alpha \quad &\Leftrightarrow \quad \eta_{t_0}^w(\{w' \in \mathcal{B}_{t_0}^w \mid M, w'_{t_0} \models \varphi\}) \geq \alpha. \\ M, w_{t_0} \models \text{PAYOFF}(\psi) \geq \alpha \quad &\Leftrightarrow \quad \forall w' \in \mathcal{D}_t^w, \text{ and } \forall x_i \text{ such that } M, x_i \models \psi, \\ &\qquad \text{where } x_i \text{ is a fullpath } (w'_{t_0}, w'_{t_1}, \ldots), \\ &\qquad \text{it is the case that } \rho_{t_0}^w(x_i) \geq \alpha \end{aligned} \quad (12)$$

We do not give any more formal details (they can be found in the cited paper), but we illustrate the logic by an example.

Consider the example illustrated in figure 1. There is an American politician, a member of the house of representatives, who must make a decision about his political career. He believes that he can stand for the house of representatives (*Rep*), switch to the senate and stand for a a senate seat (*Sen*), or retire altogether (*Ret*). He does not consider the option of retiring seriously, and is certain to keep his seat in the house. He must decide to conduct or not conduct an opinion *Poll* the result of which is either a majority approval of his move to the senate (*yes*) or a majority disapproval (*no*). There are four belief-accessible worlds, each with a specific probability value attached. The propositions *win*, *loss*, *yes* and *no* are true at the appropriate points. For example, he believes that he will win a seat in the senates with probability 0.24 if he has a majority approval to his switch and stands for a senate seat. The goal-accessible worlds are also shown, with the individual utility values (payoffs) attached. For example, the utility of winning a seat in the senate if he has a majority approval to his switch is 300. Note that retiring is an option in the belief worlds, but is not considered a goal. Finally, if we apply the maximal expected value decision rule, we end up with four remaining intention worlds, that indicate the commitments the agent should rationally make. The resulting intention-accessible worlds indicate that the best plan of actions is $Poll; ((yes?; Sen) \mid (no?; Rep))$. According to this plan of actions

[5] Note that in this definition of interpretation M they have left out the universe of discourse U.

[6] In the original definition the notation μ_t^w is used instead of η_t^w. The notation is changed here to avoid confusion with the Pearl's notation in which μ is used.

he should conduct a *Poll* followed by (indicated by sequence ; operator) switching to the senate and standing for a senate seat (*Sen*) if the result of the *Poll* is *yes* or (indicated by external choice | operator) not to switch to the senate and standing for a house of representatives seat (*Rep*) if the result of the *Poll* is *no*.

Insert figure 1 about here

Figure 1: Belief, Goal and Intention worlds, using *maxexpval* as decision rule [41]

4.2 Relation between decision theory and BDI

Rao and Georgeff relate decision trees to these structures on possible worlds. They propose a transformation between a *decision tree* and the *goal accessible worlds* of an agent.

A decision tree consists of two types of nodes: one type of nodes expresses agent's choices and the other type expresses the uncertainties about the effect of actions (i.e. choices of the environment). These two types of nodes are indicated respectively by a square and circle in the decision trees as illustrated in figure 2. In order to generate relevant plans (goals), the uncertainties about the effect of actions are removed from the given decision tree (circle in figure 2) resulting in a number of new decision trees. The uncertainties about the effect of actions are now assigned to the newly generated decision trees.

Insert figure 2 about here

Figure 2: Transformation of a decision tree into a possible worlds structure

For example, consider the decision tree in figure 2. A possible plan is to perform *Poll* followed by *Sen* if the effect of the poll is *yes* or *Rep* if the effect of the poll is *no*. Suppose that the probability of *yes* as the effect of a poll is 0.42 and that the probability of *no* is 0.58. Now the transformation will generate two new decision trees: one in which event *yes* takes place after choosing *Poll* and one in which event *no* takes place after choosing *Poll*. The uncertainties 0.42 and 0.57 are then assigned to the resulting trees, respectively. The new decision trees provide two scenarios $Poll; \textit{if yes, then Sen}$ and $Poll; \text{if } no, \text{then } Rep$ with probabilities 0.42 and 0.58, respectively. In these scenarios the effects of events are known. The same mechanism can be repeated for the remaining chance nodes. The probability of a scenario that occurs in more than one goal world is the sum of the probabilities of the different goal worlds in which the scenario occurs. This results in the goal accessible worlds from figure 1. The agent can decide on a scenario by means of a decision rule such as maximum expected utility.

5 Extensions

In this section we first discuss the extension of classical decision theory with time and processes. This extension seems to be related to the notion of intention, as used in belief-intention-desire models of agents. Then we discuss the extension of classical decision theory to game theory. This extension again seems to be related to concepts used in agent theory, namely social norms. Exactly how these notions are related remains an open problem. In this section we mention some examples of the clues to their relation which can be found in the literature.

5.1 Time: processes, planning and intentions

A decision process is a sequence of decision problems. If the next state is dependent on only the current state and action the decision process is said to obey the Markov property. In such a case, the process is called a Markov decision process or MDP. Since intentions can been interpreted as commitments to previous decisions, it seems reasonable to relate intentions to decision processes. However, *how* they should be related to decision processes remains one of the main open problems of BDI theory.

A clue to relate decision processes and intentions may be found in the stabilizing function of intention. BDI researchers [43, 44] suggest that classical decision theories may produce instable decision behavior when the environment is dynamic. Every change in the environment requires the decision problem to be reformulated, which may in turn result in conflicting decisions. For example, a lunar robot may make diverging decisions based on relatively arbitrary differences in its sensor readings.

Another clue to relate decision processes and intentions may be found in commitment strategies to keep, reconsider or drop an intention, because commitment to a previous decision can affect new decisions that an agent makes at each time. Rao and Georgeff discuss blindly committed, single-mindedly committed, and open-mindedly committed agents [43]. According to the first, an agent will deny any change in its beliefs and desires that conflicts with its previous decisions. The second does allow belief changes; the agent will drop previous decisions that conflict with new beliefs. The last strategy allows both desires and beliefs to change. The agent will drop previous decisions that conflict with new beliefs or desires. The process of intention creation and reconsideration is often called the deliberation process.

However, these two clues may only give a partial answer to the question how decision processes and intentions are related. Another relevant question is whether and how the notion of limited or bounded rationality comes into play. For example, do cognitive agents rely on intentions to stabilize their behavior only because they are limited or bounded in their decision making? In other words, would perfect reasoners need to use intentions in their decision making process, or can they do without them?

Another aspect of intentions is related to the role that they play in social interaction. In section 3.2 we discussed the use of intentions to explain speech acts. The best example of an intention used in social interaction is the content of a promise. Here the intention is expressed and made public, thereby becoming a social fact. A combination of public intentions can explain cooperative behavior in a group, using so called joint intentions [55]. A joint intention in a group then consists of the individual intentions of the members of the group to do their part of the task in order to achieve some shared goal.

Note that in the philosophy of mind intentions have also been interpreted in a different way [7]. Traditionally, intentions are related to responsibility. An agent is held responsible for the actions it has willingly undertaken, even if they turn out to involve undesired side-effects. The difference between intentional and unintentional (forced) action, may have legal repercussions. Moreover, intentions-in-action are used to explain the relation between decision and action. Intentions are what causes an action; they control behavior. On the other hand, having an intention by itself is not enough. Intentions must lead to action at some point. We can not honestly say that someone intends to climb Mt. Everest, without some evidence of him actually preparing for the expedition. It seems hard to reconcile these philosophical aspects of intentions, with mere decision processes.

5.2 Multiagent: games, norms and commitments

Classical game theory studies decision making of several agents at the same time. Since each agent must take the other agents' decisions into account, the most popular approach is based on equilibria analysis. Since norms, obligations and social commitments are of interest when there is more than one agent making decisions, these concepts seem to be related to games. However, again it is unclear *how* norms, obligations and commitments can be related to games.

The general idea runs as follows. Agents are autonomous: they can decide what to do. Some behavior will harm other agents. Therefore it is in the interest of the group, to constrain the behavior of its members. This can be done by implicit norms, explicit obligations, or social commitments. Nevertheless, relating norms to game theory is even more complicated than relating intentions to processes, because there is no consensus on the role of norms in knowledge-based systems and in belief-intention-desire models. Only recently versions of BDI have been extended with norms (or obligations) [21] and it is still debated whether and when artificial agents need norms. It is also debated whether norms should be represented explicitly or can remain implicit. Clues for the use of norms have been given in the cognitive approach to BDI, in evolutionary game theory and in the philosophical areas of practical reasoning and deontic logic. Several notions of norms and commitments have been discussed, including the following ones.

Norms as goal generators. The cognitive science approach to BDI [15, 12] argues that norms are needed to model social agents. Norms are important concepts for social agents, because they are a mechanism by which society can influence the behavior of individual agents. This happens through the creation of normative goals, a process which consists of four steps. First the agent has to believe that there is a norm. second, it has to believe that this norm is applicable. Third, it has to decide to accept the norm – the norm now leads to a normative goal – and fourth, it has to decide whether it will follow this normative goal.

Reciprocal norms. The argument of evolutionary game theory [4] is that reciprocal norms are needed to establish cooperation in repeated prisoner's dilemmas.

Norms influencing decisions. In practical reasoning, in legal philosophy and in deontic logic (in philosophy as well as in computer science) it has been studied how norms influence behavior.

Norms stabilizing multiagent systems. It has been argued that obligations play the same role in multiagent systems as intentions do in single agent systems, namely they stabilize its behavior [53].

Here we discuss an example which is closely related to game theory, in particular to the pennies pinching example. This is a problem discussed in philosophy that is also relevant for advanced agent-based computer applications. It is related to trust, but it has been discussed in the context of game theory, where it is known as a non-zero sum game. Hollis [28, 29] discusses the example and the related problem of backward induction as follows.

> A and B play a game where ten pennies are put on the table and each in turn takes one penny or two. If one is taken, then the turn passes. As soon as two are taken the game stops and any remaining pennies vanish. What will happen, if both players are rational? Offhand one might suppose that they emerge with five pennies each or with a six-four split – when the player with the odd-numbered turns take two at the end. But game theory seems to say not. Its apparent answer is that the opening player will take two pennies, thus killing the golden goose at the start and leaving both worse off. The immediate trouble is caused by what has become known as backward induction. The resulting pennies gained by each player are given by the bracketed numbers, with A's put first in each case. Looking ahead, B realizes that they will not reach $(5, 5)$, because A would settle for $(6, 4)$. A realizes that B would therefore settle for $(4, 5)$, which makes it rational for A to stop at $(5, 3)$. In that case, B would settle for $(3, 4)$; so A would therefore settle for $(4, 2)$, leading B to prefer $(2, 3)$; and so on. A thus takes two pennies at his first move and reason has obstructed the benefit of mankind.

Game-theory and backward induction reasoning do not produce the intuitive solution to the problem, because agents are assumed to be rational in the sense of economics and consequently game-theoretic solutions do not consider an implicit mutual understanding of a cooperation strategy [2]. Cooperation results in an increased personal benefit by seducing the other party in cooperation. The open question is how such 'super-rational' behavior can be explained.

Hollis considers in his book 'Trust within reason' [29] several possible explanations why an agent should take one penny instead of two. For example, taking one penny in the first move 'signals' to the other agent that the agent wants to cooperate (and it signals that the agent is not rational in the economic sense). Two concepts that play a major role in his book are trust and commitment (together with norm and obligation). One possible explanation is that taking one penny induces a commitment that the agent will take one penny again in his next move. If the other agent believes this commitment, then it has become rational for him to take one penny too. Another explanation is that taking one penny leads to a commitment of the other agent to take one penny too, maybe as a result of a social norm to share. Moreover, other explanations are not only based on commitments, but also on the trust in the other party.

In [9] Broersen *et al.* introduce a language in which some aspects of these analyses can be represented. They introduce a modal language, like the ones which have seen before, in which they introduce two new modalities. The formula $C_{i,j}(\varphi > \psi)$ means that agent i is committed towards agent j to do φ rather than ψ, and $T_{j,i}(\varphi > \psi)$ means that agent j is more trusted by agent i after executing φ than after executing ψ. To deal with the examples the following relation between trust and commitment is proposed: violations of stronger commitments result in a higher loss of trustworthiness, than violations of weaker ones.

$$C_{i,j}(\varphi > \psi) \rightarrow T_{j,i}(\varphi > \psi) \tag{13}$$

In this paper we only consider the example without communication. Broersen *et al.* also discuss scenarios of pennies pinching with communication.

The set of agents is $G = \{1, 2\}$ and the set of atomic actions $A = \{take_i(1), take_i(2) \mid i \in G\}$, where $take_i(n)$ denotes that the agent i takes n pennies. The following formula denotes that taking one penny induces a commitment to take one penny later on. The notation $[\varphi]\psi$ says that after action φ, the formula ψ must hold.

$$[take_1(1); take_2(1)]\ C_{1,2}(take_1(1) > take_1(2)) \tag{14}$$

The formula expresses that taking one penny is interpreted as a signal that agent 1 will take one penny again on his next turn. When this formula holds, it is rational for agent 2 to take one penny.

The following formula denotes that taking one penny induces a commitment for the other agent to take one penny on the next move.

$$[take_1(1)]C_{2,1}(take_2(1) > take_2(2)) \tag{15}$$

The formula denotes the implications of a social law, which states that you have to return favors. It is like giving a present at someone's birthday, thereby giving the person the obligation to return a present for your birthday.

Besides the commitment operator more complex examples involve also the trust operator. For example, the following formula denotes that taking one penny increases the amount of trust.

$$T_{i,j}((\varphi; take_j(1)) > \varphi). \tag{16}$$

The following formulas illustrate how commitment and trust may interact. The first formula expresses that each agent intends – in the sense of BDI – to increase the amount of trust (long term benefit). The second formula expresses that any commitment to itself is also a commitment to the other agent (a very strong cooperation rule).

$$\begin{array}{lcl} T_{i,j}(\psi > \varphi) & \rightarrow & I_j(\psi > \varphi). \\ C_{j,j}(\psi > \varphi) & \leftrightarrow & C_{j,i}(\psi > \varphi). \end{array} \tag{17}$$

From these two rules, together with the definitions and the general rule, we can deduce:

$$C_{i,j}(take_i(1) > take_i(2)) \leftrightarrow T_{j,i}(take_i(1) > take_i(2)) \qquad (18)$$

In this scenario, each agent is assumed to act to increase its long term benefit, i.e., act to increase the trust of other agents. Note that the commitment of i to j to take one penny increases the trust of j in i and vice versa. Therefore, each agent would not want to take two pennies since this will decrease its long term benefit.

6 Conclusion

In this paper we study how the research areas classical decision theory, qualitative decision theory, knowledge-based systems and belief-desire-intention models are related by discussing relations between several representative examples of each area. We compare the theories, systems and models on three aspects: the way the informational and motivational attitudes are represented, the way the alternative actions are represented, and the way that decisions are reached. The comparison is summarized in table 2.

	CDT	QDT	KBS	BDI
information	probabilities	qualitative probability	knowledge	beliefs
motivation	utilities	qualitative utility	goals	desires
alternatives	small set	do(φ)	decision variable	branches
focus	decision rule	decision rule	deliberation	agent types

Table 2: Comparison

6.1 Similarities

Classical decision theory, qualitative decision theory, knowledge-based systems and belief-desire-intention models all contain representations of information and motivation. The informational attitudes are probability distributions, qualitative abstractions of probabilities and logical models of knowledge and belief, respectively. The motivational attitudes are utility functions, qualitative abstractions of utilities, and logical models of goals and desires.

Each of them has some way to encode a set of alternative actions to be decided. This ranges from a small predetermined set for decision theory, or a set of decision variables for Boutillier's qualitative decision theory, through logical formulas in Pearl's approach and in knowledge-based systems, to branches in a branching time temporal logic for belief-desire-intention models.

Each area has a way of formulating how a decision is made. Classical and qualitative decision theory focus on the optimal decisions represented by a decision rule. Knowledge-based systems and belief-desire-intention models focus on a model of the representations used in decision making, inspired by cognitive notions like belief, desire, goal and intention.

Relations among these concepts express an agent type, which determines the deliberation process.

We also discuss several extensions of classical decision theory which call for further investigation. In particular, we discuss the two-step process of decision making in BDI, in which an agent first generates a set of goals, and then decides how these goals can best be reached. We consider decision making through time, comparing decision processes and the use of intentions to stabilize decision making. Previous decisions, in the form of intentions, influence later iterations of the decision process. We also consider extensions of the theories for more than one agent. In the area of multi-agent systems norms are usually understood as obligations from society, inspired by work on social agents, social norms and social commitments [12]. In decision theory and game theory norms are understood as reciprocal norms in evolutionary game theory [4, 48] that lead to cooperation in iterated prisoner's dilemmas and in general lead to an decrease in uncertainty and an increase in stability of a society.

6.2 Challenges

The renewed interest in the foundations of decision making is due to the automation of decision making in the context of tasks like planning, learning, and communication in autonomous systems [5, 7, 14, 17]. The example of Doyle and Thomason [24] on automation of financial advice dialogues illustrates decision making in the context of more general tasks, as well as criticism on classical decision theory. The core of the criticism is that the decision making process is not formalized by classical decision theory but dealt with only by decision theoretic practice. Using insights from artificial intelligence, the alternative theories, systems and models challenge the assumptions underlying classical decision theory. Some examples have been discussed in the papers studied in this comparison.

1. The set of alternative actions is known beforehand, and fixed.

   As already indicated above, Pearl uses actions $Do(\varphi)$ for any proposition φ. The relation between actions is expressed in a logic, which allows one to reason about effects of actions, including non-desirable side-effects. Boutillier makes a conceptual distinction between controllable and uncontrollable variables in the environment. Belief-desire-intention models use a branching time logic with events to model different courses of action.

2. The user has an initial set of preferences, which can be represented by a utility function.

   Qualitative decision rules studied in classical decision theory as well as Boutilier's purely qualitative decision theory cannot combine preference and plausibility to deliberate over likely but uninfluential events, and unlikely but highly influential events. Pearl's commensurability assumption on the semi-qualitative rankings for preference and plausibility solves this incomparability problem, while retaining the qualitative aspect.

3. The user has an initial set of beliefs which can be represented by a probability distribution.

   The preferences of an agent depend on its beliefs about the domain. For example, our user seeking financial advice may have wrong ideas about taxation, influencing her decision. Once she has realized that the state will not get all her savings, she may be less willing to give to charity for example. This dependence of preference on belief is dealt with by Pearl, by Boutillier and by BDI models in different ways. Pearl uses causal networks to deal with belief revision, Boutillier selects minimal elements in the preference ordering, given the constraints of the probability ordering, and in BDI models realism axioms restrict models.

4. Decisions are one-shot events, which are independent of previous decisions and do not influence future decisions.

   This assumption has been dealt with by (Markov) decision processes in the classical decision theory tradition, and by intention reconsideration and planning in knowledge-based systems and BDI.

5. Decisions are made by a single agent in isolation.

   This assumption has been challenged by the extension of classical decision theory called classical game theory. In multi-agent systems belief-desire-intention models are used. Belief-desire-intention logics allow one to specify beliefs and desires of agents about other agents' beliefs and desires, etc. Such nested mental attitudes are crucial in the application of interactive systems. In larger groups of agents, we may need social norms and obligations to restrict the possible behavior of individual agents. In such theories agents are seen as autonomous; socially unwanted behavior can be forbidden, but not be prevented. By contrast, in game theory agents are programmed to follow the rules of the 'game'. Agents are not in a position to break a rule. The set of alternative actions must now also include potential violations of norms, by the agent itself or by others.

Our comparison has resulted in a list of similarities and differences between the various theories of decision making. The differences are mostly due to varying conceptualizations of the decision making process, and a different focus in its treatment. For this reason, we believe that the elements of the theories are mostly complementary. Despite the tension between the underlying conceptualizations, we found several underlying similarities. We hope that our comparison will stimulate further research into hybrid approaches to decision making.

Acknowledgments

Thanks to Jan Broersen and Zhisheng Huang for many discussions on related subjects in the context of the BOID project.

References

[1] J. Allen. and G. Perrault. Analyzing intention in dialogues. *Artificial Intelligence*, 15(3):143–178, 1980.

[2] R. Auman. Rationality and bounded rationality. *Games and Economic behavior*, 21:2–14, 1986.

[3] J. L. Austin. *How to do things with words.* Harvard University Press, Cambridge Massachusetts, 1962.

[4] R. Axelrod. *The evolution of cooperation.* Basic Books, New York, 1984.

[5] C. Boutilier. Towards a logic for qualitative decision theory. In *Proceedings of the Fourth International Conference on Knowledge Representation and Reasoning (KR'94)*, pages 75–86. Morgan Kaufmann, 1994.

[6] C. Boutilier, T. Dean, and S. Hanks. Decision-theoretic planning: structural assumptions and computational leverage. *Journal of Artificial Intelligence Research*, 11:1–94, 1999.

[7] M. E. Bratman. *Intention, plans, and practical reason.* Harvard University Press, Cambridge Massachusetts, 1987.

[8] P. Bretier and D. Sadek. A rational agent as the kernel of a cooperative spoken dialogue system: implementing a logical theory of interaction. In *Intelligent Agents III: Proceedings of the ECAI'96 Workshop on Agent Theories, Architectures and Languages (ATAL'96)*, LNCS 1193, pages 189–204. Springer, 1997.

[9] J. Broersen, M. Dastani, Z. Huang, and L. van der Torre. Trust and commitment in dynamic logic. In *Proceedings of The First Eurasian Conference on Advances in Information and Communication Technology (EurAsia ICT 2002)*, LNCS 2510, pages 677–684. Springer, 2002.

[10] J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly*, 2(3-4):428–447, 2002.

[11] J. Broersen, M. Dastani, and L. van der Torre. Realistic desires. *Journal of Applied Non-Classical Logics*, 12(2):287–308, 2002.

[12] C. Castelfranchi. Modelling social action for AI agents. *Artificial Intelligence*, 103(1-2):157–182, 1998.

[13] P. Cohen and H. Levesque. Intention is choice with commitment. *Artificial Intelligence*, 42(2-3):213–261, 1990.

[14] P. Cohen and C. Perrault. Elements of a plan-based theory of speech acts. *Cognitive Science*, 3:177–212, 1979.

[15] R. Conte and C. Castelfranchi. Understanding the effects of norms in social groups through simulation. In G. N. Gilbert and R. Conte, editors, *Artificial Societies: the computer simulation of social life*. University College London Press, London, 1995.

[16] M. Dastani, F. de Boer, F. Dignum, and J-J.Ch. Meyer. Programming agent deliberation: An approach illustrated using the 3apl language. In *Proceedings of the Second International Conference on Autonomous Agents and Multiagent Systems (AAMAS'03)*. ACM Press, 2003.

[17] M. Dastani, F. Dignum, and J-J.Ch. Meyer. Autonomy and agent deliberation. In *Proceedings of The First International Workshop on Computatinal Autonomy - Potential, Risks, Solutions (Autonomous 2003)*, 2003.

[18] M. Dastani and L. van der Torre. Specifying the merging of desires into goals in the context of beliefs. In *Proceedings of The First Eurasian Conference on Advances in Information and Communication Technology (EurAsia ICT 2002)*, LNCS 2510, pages 824–831. Springer, 2002.

[19] T. Dean and M. P. Wellman. *Planning and control*. Morgan Kaufmann, 1991.

[20] D. Dennett. *The intentional stance*. MIT Press, Cambridge Massachusetts, 1987.

[21] F. Dignum. Autonomous agents with norms. *Artificial Intelligence and Law*, 7:69–79, 1999.

[22] J. Doyle. A model for deliberation, action and introspection. Technical Report AI-TR-581, MIT AI Laboratory, 1980.

[23] J. Doyle, Y. Shoham, and M. P. Wellman. The logic of relative desires. In *Sixth International Symposium on Methodologies for Intelligent Systems*, Charlotte, North Carolina, 1991.

[24] J. Doyle and R. Thomason. Background to qualitative decision theory. *AI magazine*, 20(2):55–68, 1999.

[25] J. Doyle and M. P. Wellman. Preferential semantics for goals. In *Proceedings of the Tenth National Conference on Artificial Intelligence (AAAI'91)*, pages 698–703, 1991.

[26] D. Dubois and H. Prade. Possibility theory as a basis for qualitative decision theory. In *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence (IJCAI'95)*, pages 1924–1930. Morgan Kaufmann, 1995.

[27] B. Hansson. An analysis of some deontic logics. *Nous*, 3:373–398, 1969.

[28] M. Hollis. Penny pinching and backward induction. *Journal of Philosophy*, 88:473–488, 1991.

[29] M. Hollis. *Trust within reason.* Cambridge University Press, Cambridge, 1998.

[30] R. C. Jeffrey. *The logic of decision.* McGraw-Hill, New York., 1965.

[31] N. R. Jennings. On agent-based software engineering. *Artificial Intelligence*, 117(2):277–296, 2000.

[32] R. L. Keeney and H. Raiffa. *Decisions with multiple objectives: preferences and value trade-offs.* John Wiley and Sons, New York, 1976.

[33] J. Lang. Conditional desires and utilities - an alternative approach to qualitative decision theory. In *Proceedings of the Twelth European Conference on Artificial Intelligence (ECAI'96)*, pages 318–322. John Wiley and Sons, New York, 1996.

[34] J. Lang, L. van der Torre, and E. Weydert. Hidden uncertainty in the logical representation of desires. In *Proceedings of Eighteenth International Joint Conference on Artificial Intelligence (IJCAI'03)*, to appear.

[35] J. Lang, E. Weydert, and L. van der Torre. Utilitarian desires. *Autonomous Agents and Multi-Agent Systems*, 5(3):329–363, 2002.

[36] D. Lewis. *Counterfactuals.* Basil Blackwell, Oxford, 1973.

[37] A. Newell. The knowledge level. *Artificial Intelligence*, 18(1):87–127, 1982.

[38] J. Pearl. System Z: a natural ordering of defaults with tractable applications to nonmonotonic reasoning. In *Proceedings of Theoretical Aspects of Reasoning about Knowledge (TARK'90)*, pages 121–135. Morgan Kaufmann, 1990.

[39] J. Pearl. From conditional ought to qualitative decision theory. In *Proceedings of the Ninth Conference on Uncertainty in Artificial Intelligence (UAI'93)*, pages 12–20. John Wiley and Sons, New York, 1993.

[40] G. Pinkas. Reasoning, nonmonotonicity and learning in connectionist network that capture propositional knowledge. *Artificial Intelligence*, 77(2):203–247, 1995.

[41] A. S. Rao and M. P. Georgeff. Deliberation and its role in the formation of intentions. In *Proceedings of the Seventh Conference on Uncertainty in Artificial Intelligence (UAI'91)*, pages 300–307. Morgan Kaufmann, 1991.

[42] A. S. Rao and M. P. Georgeff. Modeling rational agents within a BDI architecture. In *Proceedings of Second International Conference on Knowledge Representation and Reasoning (KR'91)*, pages 473–484. Morgan Kaufmann, 1991.

[43] A. S. Rao and M. P. Georgeff. BDI agents: from theory to practice. In *Proceedings of the First International Conference on Multi-Agent Systems (ICMAS'95)*, pages 312–319. AAAI Press, 1995.

[44] A. S. Rao and M. P. Georgeff. Decision procedures for BDI logics. *Journal of Logic and Computation*, 8(3):293–342, 1998.

[45] L. J. Savage. *The foundations of statistics*. John Wiley and Sons, New York, 1954.

[46] G. Schreiber, H. Akkermans, A. Anjewierden, R. de Hoog, N. Shadbolt, W. van de Velde, and B. Wielinga. *Knowledge engineering and management: the CommonKADS methodology*. The MIT Press, Cambridge Massachusetts, 1999.

[47] J. Searle. *Speech acts: an essay in the philosophy of language*. Cambridge University Press, Cambridge, 1969.

[48] Y. Shoham and M. Tennenholtz. On the emergence of social conventions: modeling, analysis, and simulations. *Artificial Intelligence*, 94(1-2):139–166, 1997.

[49] H. A. Simon. A behavioral model of rational choice. *Quarterly Journal of Economics*, pages 99–118, 1955.

[50] S.-W. Tan and J. Pearl. Qualitative decision theory. In *Proceedings of the Thirteenth National Conference on Artificial Intelligence (AAAI'94)*, pages 928–933. AAAI Press, 1994.

[51] S.-W. Tan and J. Pearl. Specification and evaluation of preferences under uncertainty. In *Proceedings of the Fourth International Conference on Knowledge Representation and Reasoning (KR'94)*, pages 530–539. Morgan Kaufmann, 1994.

[52] R. H. Thomason. Desires and defaults: a framework for planning with inferred goals. In *Proceedings of Seventh International Conference on Knowledge Representation and Reasoning (KR'00)*, pages 702–713. Morgan Kaufmann, 2000.

[53] L. van der Torre. Contextual deontic logic: normative agents, violations and independence. *Annals of Mathematics and Artificial Intelligence*, 37(1-2):33–63, 2003.

[54] L. van der Torre and E. Weydert. Parameters for utilitarian desires in a qualitative decision theory. *Applied Intelligence*, 14:285–301, 2001.

[55] M. J. Wooldridge and N. R. Jennings. The cooperative problem-solving process. *Journal of Logic and Computation*, 9(4):563–592, 1999.

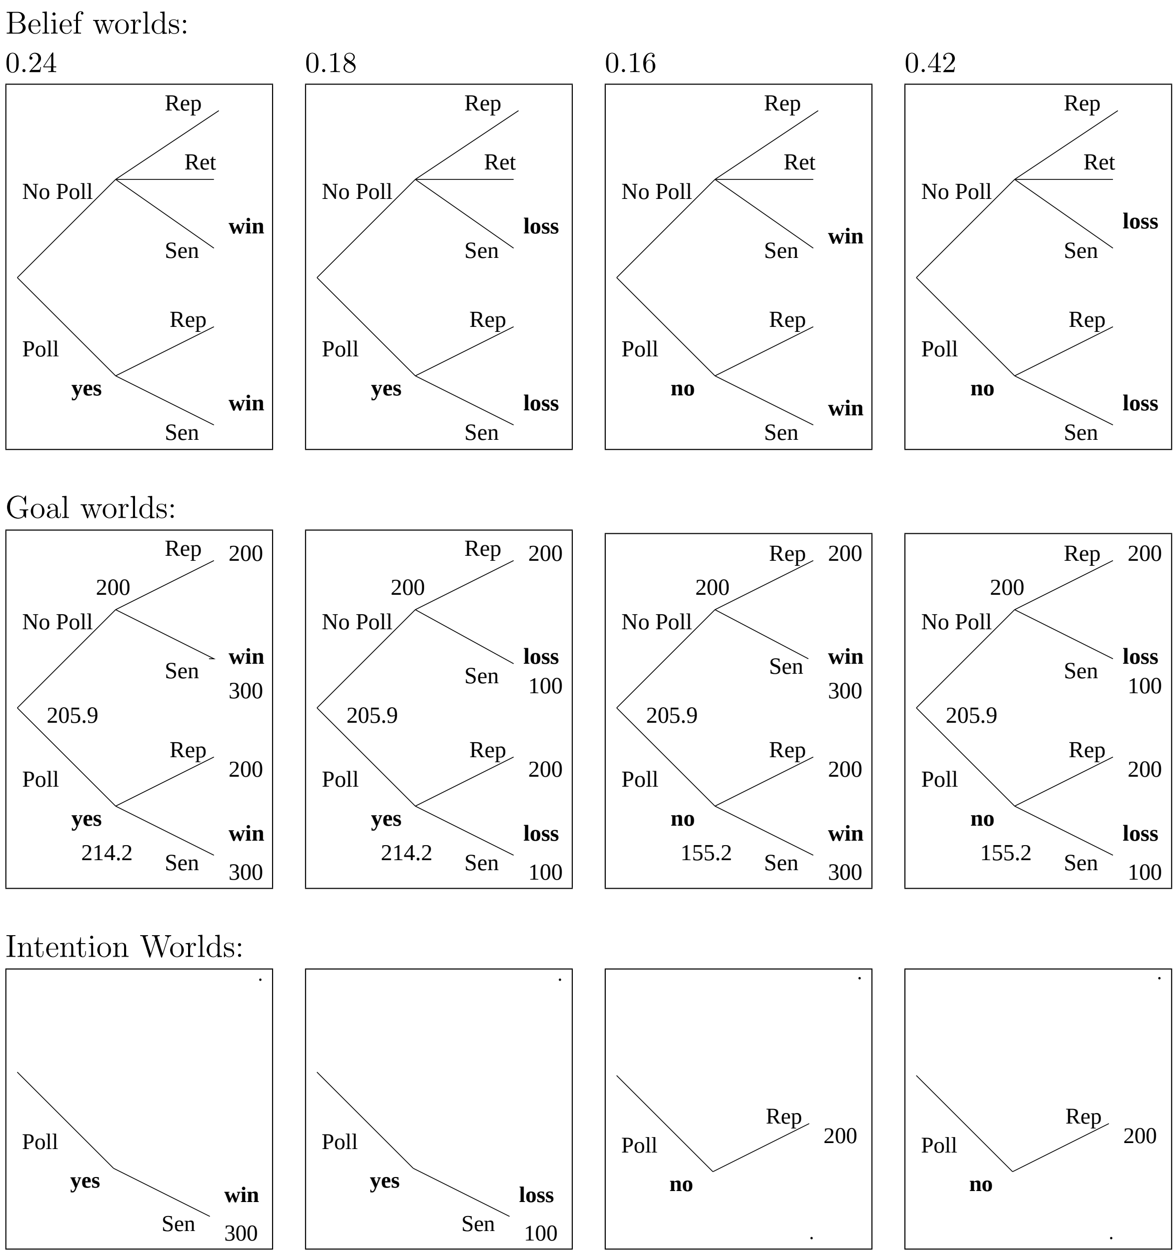

Figure 1: Belief, Goal and Intention worlds, using *maxexpval* as decision rule [41]

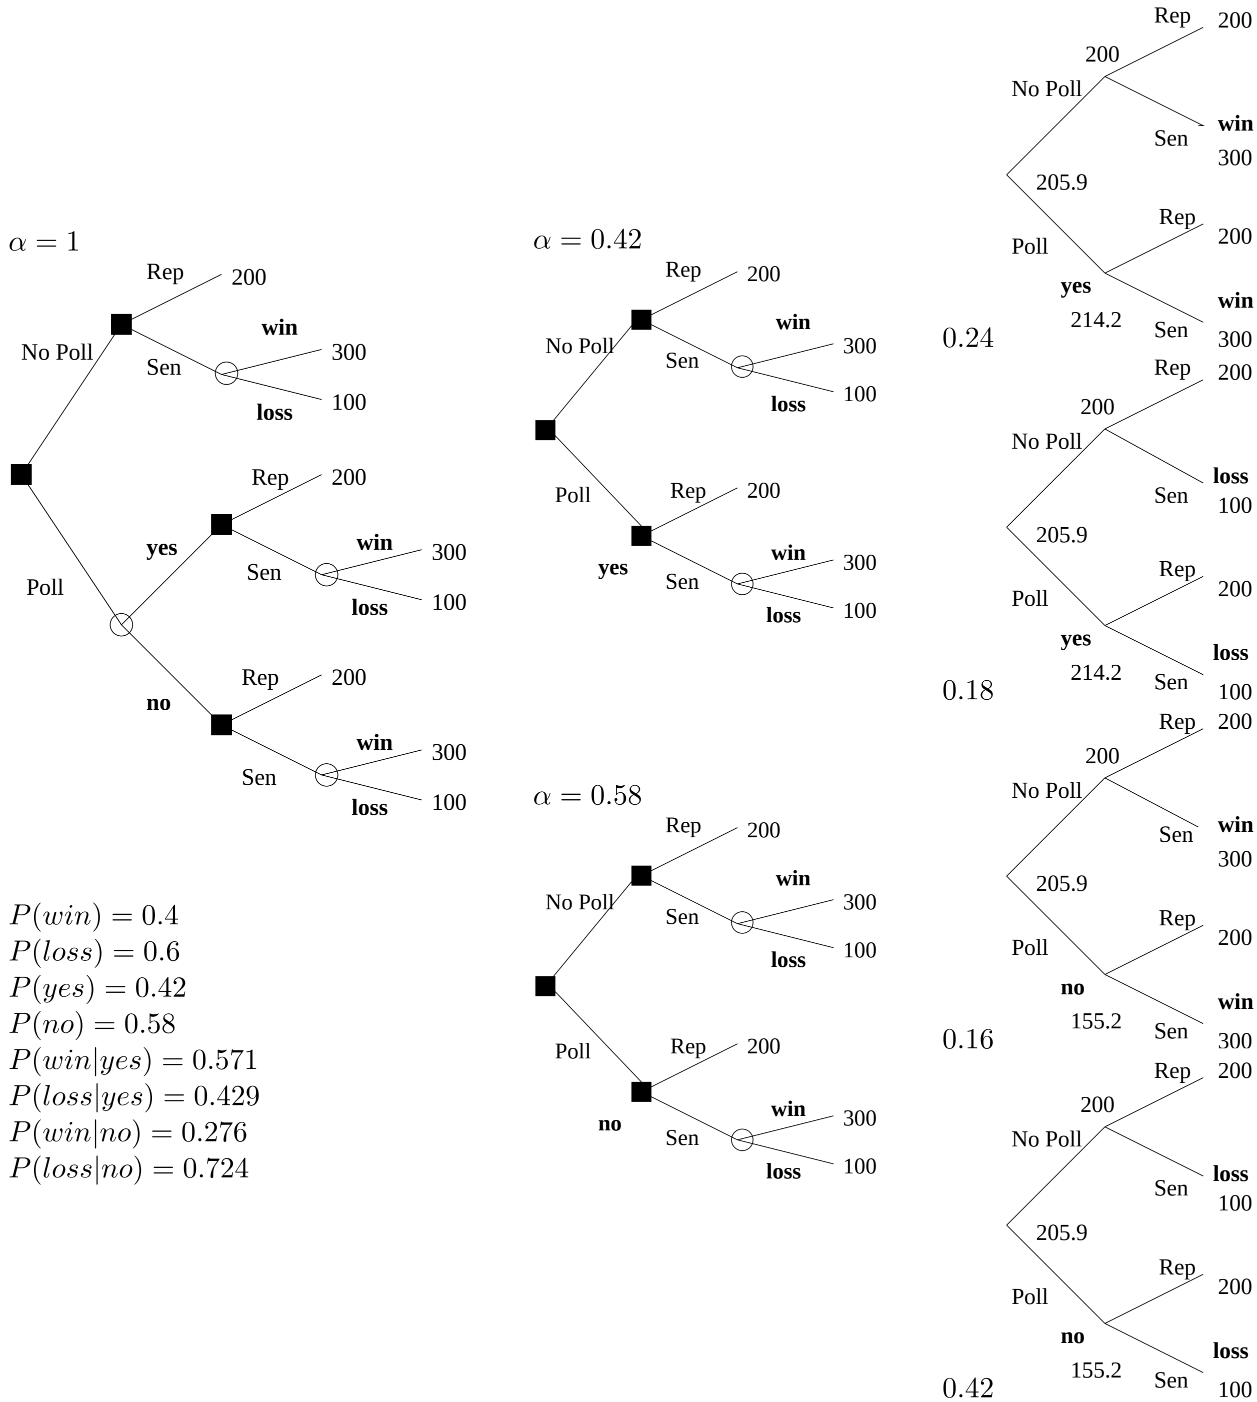

Figure 2: Transformation of a decision tree into a possible worlds structure