Deciding Between Conflicting Influences

Andreas Schmidt Jensen

Department of Applied Mathematics and Computer Science,
Technical University of Denmark, Matematiktorvet, Building 303B,
DK-2800 Kongens Lyngby, Denmark

Abstract. This paper investigates an approach of decision making internally in an agent in which a decision is based on both preference and expectation. The approach uses a logic for qualitative decision theory proposed by Boutilier in order to express such notions. To make readily use of this we describe a simple method for generating preference and expectation models that respect certain rules provided by the agents.

1 Introduction

Agents taking part in a multi-agent system are usually seen as intelligent entities that autonomously are able to bring about (from their own perspectives) desirable states. In a fixed setting with a controlled number of agents and globally desirable states, the designer will in many cases be able to implement the agents such that their own desirable states coincide with the globally desirable states. In open societies agents often come from different sources and their desires cannot as such be assumed to match the global desires. A suggestion is to impose an organization upon the agents which is able to influence the actions of the agent towards the desires of the organization.

When agents are constrained by an organization, their own goals may conflict with those of the organization. In some of the previous work towards resolving such conflicts the solution has often been to order desires and obligations a priori, such that an agent either prefers desires over obligations or obligations over desires. This results in agents that are always selfish or always social. We argue in this paper that such distinction can be too hard; even a selfish agent could in some cases benefit by preferring certain obligations over its desires. We consider an approach on how to resolve such conflicts which is based on work in the area of qualitative decision theory by Boutilier [3], where the expected consequences of bringing about a state are considered. We show that this result in agents that are not always either social or selfish, but are instead able to decide based on the consequences of bringing about the states.

In order to make the approach readily useful we furthermore describe a simple method for generating models for preference and expectation, based on basic rules specified by the agents, such as “I prefer to drive to work when it is raining”. This will help building agents that are able to reason using our approach.

The paper is organized as follows: In section 2 we discuss the issues that arise when an agent has to make a decision between conflicting influences. In section 3 we present a new approach on how to solve such conflicts without having to put the agents into the categories "selfish" or "social". We present a method for generating models that conform with the agent's preferences in section 4. In section 5 we discuss a case in which agents have conflicting desires and show that our method enables them to choose using both their own preference and the expected consequence of bringing about each influence. Finally we conclude our work and discuss future research directions in section 6.

2 Conflicting influences

Agents entering an environment will be subject to influences from multiple sources; their own desires, requests from others, and obligations from an organization. In the well-known BDI model an agent's desires become intentions, when the agent commits to bringing about these desires. One could argue that if an agent wants to accept requests from other agents, or if it wants to adhere to the obligations of an organization, these influences are merely desires as well; the agent simply desires to do so. The incentives for doing so are however not clear, since there should be different reasons for committing to actual desires and to requests or obligations disguised as desires. For example, if an agent has a desire to move a box from A to B, it typically *wants* to do so. However if the agent wants to pay a bill before its due date, this "desire" has more likely arisen from the fact that the agent does not want to pay a fine, rather than being an actual desire to pay the bill. In such a situation the desire may actually be an obligation or a request to pay the bill, which means that the agent should reason differently, since the actual desire is to avoid paying a fine.

Furthermore, consider an agent that receives an undesirable request from an agent that it desires to help. It may choose to commit to the task even though the task itself is not desirable, since the desire to help the other agent is stronger than the desire to not perform the task (the consequence of *not* helping the other agent might be a bad reputation). Similarly if an agent is obligated to perform certain tasks for an organization, it should not only be able to consider whether the task is desirable, but also weigh this against the penalty for violating the obligation.

In this paper we call all the propositions that the agent has to take into account when making a decision "decision influences" since they influence how the agent chooses to act. It naturally has to take into account its desires, since it would be irrational to ignore them, but the consequence of not reasoning about e.g. obligations might be intolerable so these influence the agent as well. This also means that the agent is not supposed to be reasoning explicitly about whether it should commit to bringing about an arbitrary obligation or desires, since they are merely considered influences. Several approaches are proposed on how to let agents choose between such specific influences [1,4,5,6], so we briefly discuss how our solution differs.

In [4] conflicts between beliefs, obligations, intentions and desires are discussed, with a focus on a distinction between *internal* conflicts, e.g. contradictory beliefs, conflicting obligations and *external* conflicts, such as a desire which is in conflict with an obligation. The solution proposed, the BOID architecture, impose a strict ordering between beliefs, obligations, intentions and desires, such that the order of derivation determines the agent's attitude. Thus different agent types emerge; an agent deriving desires before beliefs is a wishful thinking agent, while an agent deriving obligations before desires is social.

We believe this ordering is too strong; if an agent is social it will always choose obligations over desires, and vice versa for selfish agents. This might not always be appropriate. For instance, a selfish agent might desire not to go to work, but if the consequence of not fulfilling the obligation of going to work is severe (i.e. getting fired), even a selfish agent should consider this consequence before deciding not to go to work.

Dignum *et al.* suggests that *"both norms and obligations should be explicitly used as influences on an agent's behavior"* [5]. In their approach obligations (and norms) are represented using Prohairetic Deontic Logic [8], a preference-based dyadic deontic logic which allows for contrary-to-duty obligations (obligations holding in a sub-ideal context). Furthermore they propose modified BDI-interpreter in which selected events are augmented with potential deontic events which, put simply, are obligations and norms that may become applicable when a plan is chosen. For instance, if agent a has an obligation to perform a task for another agent b, and a does not intend to do so he ought to inform b about this. The modified interpreter generates a number of options depending on these potential events and a relevant plan is chosen based on the agent's attitude.

In [6] it is argued that the preference orderings induced by desires, obligations and norms should be combined into a single ordering. It is noted that a common way to do so is to allow that a single preference ordering determine the aggregate ordering, such that the agent might always put obligations over norms and norms over desires, similarly to the BOID architecture. Another approach is also discussed in which the orderings are mapped into a common scale, such that very desirable situations could outweigh the cost of violating certain obligations. Such ordering should be quite dynamic since, as noted, obligations towards a trusted agent should become less important if that agent becomes less trustworthy. Simple rules are presented to deal with few alternatives, but it is noted that the situation is more complex if an agent has to choose between three or more alternatives and none of the three orderings agree on a preferred alternative. A simple rule which orders the alternatives in a fixed order results in a very simple-minded agent and it is suggested that the consequences of different situations is considered, however this is not investigated further.

The approach presented in this paper compares the influences by taking the agent's preferences into account such that more preferred influences will be chosen over less preferred influences. When influences are equally preferred there are two possibilities, either let the agent choose arbitrarily between the influences or compare them using a different approach. When considering only desires the former option might be acceptable since if the agent wants to do two things and one is not more preferred than the other, it might seem reasonable to choose one arbitrarily. On the other hand, when influences include those originating from other agents or an organization the latter solution is more suitable. Otherwise since only the agent's preference is considered, no reasoning about violating obligations or ignoring requests occur.

We suggest that the agent should reason about the expected *consequences* of choosing to pursue a decision influence and furthermore that this reasoning should be based on *tolerance*. When two decision influences are equally preferable, the agent should consider whether the consequences of bringing about one influence are more tolerable than the consequences of bringing about the other. We define a proposition as being tolerated when its negation is not preferred (e.g. working is tolerated if staying at home is not preferred over working). The reason for using tolerance instead of preference in the case of consequences is that the agent should not need to desire the consequences of bringing about a state. The consequences are side effects which may not be desired in the same way as the actual influence is. Of course, if a consequence is preferable, then clearly it is also tolerable but the opposite need not be the case. This means that even though none of the expected consequences are actually preferable we are still able to compare them. Finally, if the expected consequences are equally tolerable then the agent is allowed to choose arbitrarily between the influences.

Note that our approach does not incorporate the notion of an organization as such; the focus is on the propositions that may influence the agent's reasoning, such as the obligations toward an organization. As a result, we model consequences as expectations from the environment, that is, which possible world is the most expected, which is the second most and so on. This means that if the consequences of the violation of an obligation are specified in an organizational model, these consequences are in our approach modeled such that worlds, in which the violation has occurred *and* the consequences are in action are more expected than those where the violation has occurred without resulting in any consequences. This will be evident in the example in section 5 where all expected consequences are incorporated into the same model.

3 Modeling Influence and Consequence

We base our work on the Logic for Qualitative Decision Theory (QDT) by Boutilier [2,3] by extending the notion of preference to allow multiple modalities in order to represent the preference of individual agents. The semantics and axiomatization are those of QDT and we define a few new abbreviations to be used in the decision making.

The basic idea behind the QDT model is as follows. An agent has the ultimate desire of achieving the goals to which it is committed. This can be modeled by a possible worlds-model in which the agent has achieved its goal when it is in a world where those goals hold. The most preferred world in an ideal setting is the world in which all of the agent's goals are achieved. However, such world is often unreachable, for a number of reasons: the agent could have contradicting goals, other agents could prevent the agent from achieving all of its goals, an organization could impose obligations which contradict the agent's goals, etc. By ordering the worlds in a preference relation it is possible to choose the most preferred world(s) in a sub-ideal situation.

In our approach we require that the consequence of bringing about a state should be taken into account. If the consequence of pursuing a personal desire is to be fired from your workplace, it might not be reasonable to do so even though the desire was more preferred than the obligations from work. We briefly describe QDT below with the semantics to include the notion of multiple agents before moving on to modeling the expected consequence of bringing about a state.

A QDT model is of the form:

$$M = \langle W, Ag, \leq_P^1, \ldots, \leq_P^n, \leq_N, \pi \rangle$$

where W is the non-empty set of worlds, Ag is the set of agents, $\leq_P^i$ is the transitive, connected preference ordering for each agent[1], $\leq_N$ is the transitive, connected normality ordering, and π is the valuation function. The normality ordering is used to model how likely each world is, e.g. it is normally cold when it is snowing.

The semantics are given as follows:

$$\begin{aligned}
M, w \models p &\iff p \in \pi(w) \\
M, w \models \neg\varphi &\iff M, w \not\models \varphi \\
M, w \models \varphi \wedge \psi &\iff M, w \models \varphi \wedge M, w \models \psi \\
M, w \models \Box_P^i \varphi &\iff \forall v \in W, v \leq_P^i w, M, v \models \varphi \\
M, w \models \overleftarrow{\Box}_P^i \varphi &\iff \forall v \in W, w <_P^i v, M, v \models \varphi \\
M, w \models \Box_N \varphi &\iff \forall v \in W, v \leq_N w, M, v \models \varphi \\
M, w \models \overleftarrow{\Box}_N \varphi &\iff \forall v \in W, w <_N v, M, v \models \varphi
\end{aligned}$$

We can define the other operators $(\vee, \rightarrow, \Diamond, \overleftarrow{\Diamond})$ as usual. Finally we can talk about a formula being true in all worlds or some worlds: $\overleftrightarrow{\Box}_P^i \varphi \equiv \Box_P^i \varphi \wedge \overleftarrow{\Box}_P^i \varphi$ and $\overleftrightarrow{\Diamond}_P^i \varphi \equiv \Diamond_P^i \varphi \vee \overleftarrow{\Diamond}_P^i \varphi$, respectively (similarly for normality). The following abbreviations are defined:

(1) $I(\psi \mid \varphi) \equiv \overleftrightarrow{\Box}_P^i \neg\varphi \vee \overleftrightarrow{\Diamond}_P^i (\varphi \wedge \Box_P^i(\varphi \rightarrow \psi))$ (Conditional preference)

(2) $\varphi \leq_P^i \psi \equiv \overleftrightarrow{\Box}_P^i (\psi \rightarrow \Diamond_P^i \varphi)$ (Relative preference)

(3) $T(\psi \mid \varphi) \equiv \neg I(\neg\psi \mid \varphi)$ (Conditional tolerance)

(4) $\varphi \Rightarrow \psi \equiv \overleftrightarrow{\Box}_N \neg\varphi \vee \overleftrightarrow{\Diamond}_N (\varphi \wedge \Box_N(\varphi \rightarrow \psi))$ (Normative conditional)

[1] We adopt the notion by Boutilier and others that we prefer minimal models, so $v \leq_P^i w$ denotes that according to agent i, v is at least as preferred as w.

The abbreviations state that (1) ψ is ideally true if φ is true, (2) φ is at least as preferred as ψ, (3) ψ is tolerable given φ and (4) that ψ normally is the case when φ is.

In order to make decisions as motivated above, we define the following abbreviations, which allow us to specify different kinds of relative preference, and relative tolerance.

$$\varphi \not\leq_P^i \psi \equiv \neg(\varphi \leq_P^i \psi) \qquad \text{(Not as preferred)}$$

$$\varphi <_P^i \psi \equiv (\varphi \leq_P^i \psi \wedge \psi \not\leq_P^i \varphi) \qquad \text{(Strictly preferred)}$$

$$\varphi \approx_P^i \psi \equiv (\varphi \leq_P^i \psi \wedge \psi \leq_P^i \varphi) \vee (\varphi \not\leq_P^i \psi \wedge \psi \not\leq_P^i \varphi) \qquad \text{(Equally preferred)}$$

$$\varphi \leq_{T(\gamma)}^i \psi \equiv (T(\varphi \mid \gamma) \wedge \neg T(\psi \mid \gamma)) \vee ((T(\varphi \mid \gamma) \leftrightarrow T(\psi \mid \gamma)) \wedge (\varphi \leq_P^i \psi \vee \varphi \approx_P^i \psi)) \qquad \text{(Relative tolerance)}$$

Thus φ is at least as tolerable as ψ w.r.t γ when either φ is tolerable given γ and ψ is not, or both φ and ψ are tolerable given γ (or both are not), and φ is at least as preferred as ψ, or they are equally preferable. This means that even if neither is tolerable, they are still comparable.

3.1 Making a decision

We now show how the extended QDT-logic can be used to decide between conflicting desires and obligations. We define a model for decision making as follows:

$$\mathcal{M}_C = \langle M, F, C, B \rangle,$$

where

- M is an extended QDT-model as defined above,
- F is for each agent the set of influences,
- C is for each agent the set of controllable propositions[2],
- B is the belief base for each agent.

We define the set of potential consequences $C'(i)$ for an agent i such that if $\varphi \in C(i)$ then $\varphi, \neg\varphi \in C'(i)$.

Definition 1 (Expected consequences). *Given an agent i, its belief base $B(i)$ as a conjunction of literals, the set of potential consequences $C'(i)$ and a literal φ. The expected consequences of bringing about φ, denoted $EC_i(\varphi)$, is given by:*

$$EC_i(\varphi) = \bigwedge C_\varphi \textit{ for all } C_\varphi \in \{C_\varphi \mid (B(i) \wedge \varphi \Rightarrow C_\varphi) \textit{ where } C_\varphi \in C'(i)\}$$

i.e. the conjunction of all literals C_φ that are normally consequences of bringing about φ given the current belief base. If there are no expected consequences, then $EC_i(\varphi) = \top$.

[2] A controllable proposition is, roughly, a proposition which the agent is able to influence, directly or indirectly, by an action. E.g. *snow* is not controllable and cannot be a consequence of an action, whereas *work* is.

Consider an agent i, and a normality ordering in which we have that

$$A \wedge \alpha \Rightarrow B, \qquad A \wedge \neg\alpha \Rightarrow C, \quad D \wedge \neg\alpha \Rightarrow E,$$

and belief base $B(i) = \{\alpha\}$. Then we have that $EC_i(A) = \{B\}$ and $EC_i(D) = \emptyset$. If $B(i) = \{\neg\alpha\}$, then $EC_i(A) = \{C\}$ and $EC_i(D) = \{E\}$.

An agent i can make a decision by selecting from the set of potentially conflicting influences, $F(i)$, the most preferred influences having the most tolerable consequences.

Definition 2 (Decision). *Given an agent i, the set of influences $F(i)$ and the expected consequences $EC_i(\varphi)$ for all $\varphi \in F(i)$, we can get the set of best influences (the decision) the agent should choose from, denoted by the function $Dec : Ag \to 2^{F(i)}$ as follows:*

$$\begin{aligned} Dec(i) = \{\varphi \mid\ & \varphi \in F(i),\ \textit{and} \\ & \textit{for all } \psi \in F(i), \psi \neq \varphi,\ \textit{either} \\ & \varphi <^i_P \psi,\ \textit{or} \\ & \varphi \approx^i_P \psi \textit{ and } EC(\varphi) \leq^i_{T(\varphi \vee \psi)} EC(\psi)\} \end{aligned}$$

Given a model $\mathcal{M}$, an agent i can then choose an arbitrary literal from $Dec(i)$, since all of these are equally preferred and with equally tolerable consequences.

If there are no expected consequences of bringing about a certain proposition, i.e. if $EC(\varphi) = \emptyset$, then we consider it tolerable, since we do not expect any consequences. Therefore for all other consequences, γ, we have to consider $\top \leq^i_{T(C)} \gamma$ and $\gamma \leq^i_{T(C)} \top$. Note that $T(\top \mid \varphi)$ is true iff φ is true in any world[3]. Furthermore, $\top \leq^i_P \varphi$ is always true, and $\varphi \leq^i_P \top$ is true iff φ is true in all worlds. Thus it is possible to make a decision even if some obligations or desires have no known consequences.

Lemma 1. *Given an agent i and expressions φ, ψ, and γ, the following relation holds for relative tolerance:*

$$\neg(\varphi \leq^i_{T(\gamma)} \psi) \to (\psi \leq^i_{T(\gamma)} \varphi)$$

Proof. We assume $\neg(\varphi \leq^i_{T(\gamma)} \psi)$ and prove that $(\psi \leq^i_{T(\gamma)} \varphi)$. We divide the proof into two parts based on the definition of relative tolerance:

$$\neg(T(\varphi \mid \gamma) \wedge \neg T(\psi \mid \gamma)) \tag{1}$$

$$\neg((T(\varphi \mid \gamma) \leftrightarrow T(\psi \mid \gamma)) \ \wedge\ (\varphi \leq^i_P \psi \vee \varphi \approx^i_P \psi)) \tag{2}$$

1. When (1) does not hold, then we have that either $T(\varphi \mid \gamma) \leftrightarrow T(\psi \mid \gamma)$ or $\neg T(\varphi \mid \gamma) \wedge T(\psi \mid \gamma)$ holds. In the latter case we have that $\psi \leq^i_{T(\gamma)} \varphi$ by the definition of relative tolerance. Otherwise they are equally tolerable and we have to consider the second case.

[3] Since $T(\top \mid \varphi) \equiv \overset{\leftrightarrow}{\Diamond}{}^i_P \varphi \wedge \overset{\leftrightarrow}{\Box}{}^i_P(\neg\varphi \vee \Diamond^i_P(\varphi \wedge \top)$.

2. When (2) does not hold, then either $\neg(T(\varphi \mid \gamma) \leftrightarrow T(\psi \mid \gamma))$ or $\neg(\varphi \leq^i_P \psi \vee \varphi \approx^i_P \psi)$. If the former is the case, then one is tolerated and the other is not. Since (1) does not hold, we have that $\neg T(\varphi \mid \gamma) \wedge T(\psi \mid \gamma)$ and therefore $\psi \leq^i_{T(\gamma)} \varphi$. If the latter is the case then we have that $\neg(\varphi \leq^i_P \psi) \wedge \neg(\varphi \approx^i_P \psi)$. In that case we have that $\psi <^i_P \varphi$ and therefore $\psi \leq^i_{T(\gamma)} \varphi$.

□

Proposition 1. *Given an agent i, a non-empty set of influences $F(i)$ and the expected consequences $EC_i(\varphi)$ for all $\varphi \in F(i)$, the set of decisions is always non-empty.*

Proof. If $|F(i)| = 1$ then $|Dec(i)| = 1$ as well, since there are no $\psi \neq \varphi$ in $F(i)$. If $F(i)$ contains more than one influence we consider two arbitrary influences φ and ψ. We want to show that of such two influences, either one or both are chosen. If $\varphi <^i_P \psi$ then $\varphi \in Dec(i)$. If $\varphi \approx^i_P \psi$ and $EC(\varphi) \leq^i_{T(\varphi \vee \psi)} EC(\psi)$ then $\varphi \in Dec(i)$. We proceed by showing that if neither holds, then ψ is more preferred (or its consequences are more tolerable) and it is therefore chosen.

1. If $\neg(\varphi <^i_P \psi)$ then either $\varphi \nleq^i_P \psi$ or $\psi \leq^i_P \varphi$. If both are the case, then $\psi <^i_P \varphi$ so $\psi \in Dec(i)$. Otherwise we have $\varphi \leq_P \psi \wedge \psi \leq_P \varphi$ or $\varphi \nleq_P \psi \wedge \psi \nleq_P \varphi$, i.e. $\varphi \approx^i_P \psi$ which is considered in the second case.
2. Either $\varphi \approx^i_P \psi$ does not hold or $EC(\varphi) \leq^i_{T(\varphi \vee \psi)} EC(\psi)$ does not hold. In the former case we then have that either $\varphi <^i_P \psi$ or $\psi <^i_P \varphi$, which means that $\varphi \in Dec(i)$ or $\psi \in Dec(i)$ respectively. In the latter case we have that $\neg(EC(\varphi) \leq^i_{T(\varphi \vee \psi)} EC(\psi))$. By lemma 1 we then have $EC(\psi) \leq^i_{T(\psi \vee \varphi)} EC(\varphi)$, and therefore $\psi \in Dec(i)$.

Thus when deciding between two arbitrary influences, at least one will be chosen. If the influences are equally preferred and tolerable then both will be chosen. □

4 Generating models

If an agent has certain preferences, they are usually not described as a model shown above. Rather will they be expressions such as "I prefer that it does not rain" or "When it rains, I want to stay inside". In order to utilize such preferences in the decision procedure above, a transformation is required. In the following we present a method which will generate a QDT-model that respects non-contradictory rules specified by the agent.

The model is initialized using the set of possible atoms, $\mathcal{L}$. We then create a model containing a world for each set in $2^{\mathcal{L}}$, where each set either contains the atom or its negation. For instance, given $\mathcal{L} = \{a, b\}$, the initial model will be $2^{\mathcal{L}} = \{\{a, b\}, \{\neg a, b\}, \{a, \neg b\}, \{\neg a, \neg b\}\}$. Certain worlds may be deemed impossible, such as raining from a clear sky. These situations are specified as prohibitions in the environment using expressions, e.g. $\neg a \wedge b$. All worlds which entails such an expression are then removed from the initial set of worlds, yielding the set of possible worlds, W.

Algorithm 1 Rule application

```
function APPLY((φ, ψ), W, ≤)
    max ← max(≤)
    for all w ∈ W do
        if w ⊨ φ ∧ ¬ψ then W_c ← w
        if (w ⊨ φ ∧ ψ) and ¬∃w′(w′ ∈ W ∧ (w′, w) ∈ lock) then
            o(w) = max + 1
            W_s ← w
    if W_s = ∅ then return ⊥
    for all w ∈ W_s, w′ ∈ W_c do lock(w, w′)
    return ⊤
```

An ordering, $\leq$, is the result of a mapping from a world to a natural number, the o-value, denoted $o\ : W \rightarrow \mathcal{N}$, such that worlds with higher numbers are more favored. Worlds can have the same o-value if they are equally favored. The maximum o-value of an ordering $\leq$ is denoted $max(\leq)$.

Each agent specify a set of rules of the form (φ, ψ), where φ and ψ are standard propositional formulas. A rule is to be understood as follows. Worlds w, in which $w \models \varphi \wedge \psi$, are favored over worlds w', where $w' \models \varphi \wedge \neg\psi$. Thus a rule can be roughly interpreted as the conditionals for preference and normality. In the following we propose a method for generating preference and normality orderings which respect such rules by utilizing this interpretation. The generic definition of the conditional operators is

$$\textit{if } \varphi \textit{ then } \psi \equiv \overleftrightarrow{\Box}\neg\varphi \vee \overleftrightarrow{\Diamond}(\varphi \wedge \Box(\varphi \rightarrow \psi)).$$

From this definition it is clear that there are two ways to ensure that a rule (φ, ψ) is respected. Either (a) φ is never true or (b) in the most favored world(s) where φ is true, ψ is also true. Clearly (a) is easily achieved; we simply remove all worlds where φ is true. This is however probably not what was intended by the agent, since the rules are most likely specified such that favored situations are actually also possible situations. We therefore require that the method does not remove any worlds from W. The method should ensure that after the application of a rule we have $M \models (\varphi, \psi)$. Another natural requirement is that previously applied rules still hold after application of a new rule. If this is not possible, we say that the new rule contradicts previously applied rule, and it is discarded.

We propose using a *locking* mechanism in which the ordering between two worlds can be locked, such that if $lock(w_1, w_2)$ then it must always be the case that $w_1 \leq w_2$. We can use this to e.g. lock the ordering between worlds $w_1 = \{\varphi, \psi\}$ and $w_2 = \{\varphi, \neg\psi\}$, such that if a rule (φ, ψ) is applied, we create a lock $lock(w_1, w_2)$ such that w_1 is always favored over w_2. Then if a rule $(\varphi, \neg\psi)$ is applied, the ordering cannot be changed so that w_2 is favored over w_1 because it would result in the previously applied rule no longer being respected (since ψ would not be entailed by the most favored world where φ holds).

A rule is applied using the function *apply* : $(\mathcal{R}, \leq) \rightarrow \{\top, \bot\}$ (algorithm 1). Applying a rule (φ, ψ) is done by finding all worlds in which both φ and ψ

holds (the sought worlds) and all worlds in which φ and $\neg\psi$ holds (the contradictory worlds). The sought worlds are given an o-value of $max(\leq) + 1$ and all contradictory worlds are locked in relative position to the sought worlds.

A rule (φ, ψ) cannot be applied if there is no world w in which $w \models \varphi \wedge \psi$ or for all such worlds a lock $lock(\_\_, w)$ exists.

Proposition 2. *Given an initial ordering $\leq$, a set of rules $\mathcal{R} = \{r_1, \ldots, r_n\}$ where each r_i is of the form (φ_i, ψ_i), the result of successfully applying rules r_1 to r_i, $0 < i \leq n$ is an ordering which respects rules $\{r_1, \ldots, r_i\}$.*

Proof. When $i = 1$ no previous rules have been applied, so we only have to show that the model respects rule r_1 after successful application. We have $o(w) = 1$ for all worlds w. Applying r_1 can only fail if no worlds entail $\varphi_1 \wedge \psi_1$ or all entailing worlds are locked. Since $lock = \emptyset$ initially, only the former can be the case. But then the rule would describe an impossible world and cannot be applied. Otherwise, after applying r_1, it is entailed by the model, since for all worlds w where $w \models \varphi_1 \wedge \psi_1$ we have $o(w) = 2$ and the o-value of all other worlds is unchanged. Thus the worlds entailing r_1 are most preferred so the rule it self is entailed by the model.

When $i > 1$ we assume that all rules up to and including r_{i-1} have been applied successfully. We therefore have

$$M \models (\varphi_1, \psi_1) \wedge \cdots \wedge (\varphi_{i-1}, \psi_{i-1}).$$

Let $(w^{\varphi}_{\psi}, w^{\varphi}_{\neg\psi}) = \{(w, w') \mid w \models \varphi \wedge \psi \text{ and } w' \models \varphi \wedge \neg\psi\}$ be the set of locks between worlds with contradictory consequents of a rule (φ, ψ). Before applying r_i the set $lock$ then contains

$$lock = (w^{\varphi_1}_{\psi_1}, w^{\varphi_1}_{\neg\psi_1}) \cup \cdots \cup (w^{\varphi_{i-1}}_{\psi_{i-1}}, w^{\varphi_{i-1}}_{\neg\psi_{i-1}})$$

Rule r_i can then be applied if there is at least one world w in which $w \models r_i$ which is not the second entry of a pair in $lock$ (i.e. there is a world entailing r_i which is not locked by another world). If there is no such world then either the rule describes an impossible world and should be rejected, or a previously applied rule contradicts it, which also means it should be rejected. Otherwise the rule will be successfully applied resulting in a model entailing all rules up to and including r_i:

$$M \models (\varphi_1, \psi_1) \wedge \cdots \wedge (\varphi_i, \psi_i),$$

and a new $lock$ set: $lock' = lock \ \cup \ (w^{\varphi_i}_{\psi_i}, w^{\varphi_i}_{\neg\psi_i})$. Assuming that the rule is successfully applied we know that for all w in which $w \models r_i$ we have $o(w) = max(\leq)+1$. Clearly r_i is then entailed by the model. We then have to show that all rules up to r_i are still entailed as well.

Consider rule r_j where $0 < j < i$. Rule r_j was entailed by the model before applying r_i. Therefore there are worlds w_j where $w_j \models \varphi_j \wedge \psi_j$ and no lock of it exists, and w'_j where $w'_j \models \varphi_j \wedge \neg\psi_j$, and for all such worlds we have that $o(w_j) > o(w'_j)$ and $(w_j, w'_j) \in lock$. Thus all worlds contradicting r_j are locked relative to those entailing it. If $w'_j \cap w_s \neq \emptyset$ then some of the sought worlds are

Algorithm 2 Model generation

function GENERATE($\mathcal{L}, \mathcal{P}, \mathcal{R}$)
    $W_0 \leftarrow init(\mathcal{L})$
    $W \leftarrow clean(W_0, \mathcal{P})$
    $\leq \leftarrow o(W)$
    $\mathcal{R}' \leftarrow sort(\mathcal{R})$
    for all $(\varphi, \psi) \in \mathcal{R}'$ **do**
        $apply((\varphi, \psi), W, \leq)$
    return $\leq$

locked by r_j, but since w_s only contains unlocked worlds, this cannot be the case. Therefore no worlds w'_j will be given a higher o-value than any w_j world. Furthermore, since w'_j contains all the worlds that could invalidate r_j, clearly r_j is still entailed after applying r_i.

Thus the procedure respects previously applied rules when applying new rules. □

Even though a successful application of a set of rules can be done, we still need to touch upon how to maximize the number of successful applications of rules. Note that the use of a locking mechanism decreases the number of worlds that can be moved around every time a rule is successfully applied. Therefore, by minimizing the number of worlds being locked in each iteration, we maximize the number of rules that can be applied. The function $s : \mathcal{R} \rightarrow \mathcal{N}$ gives each rule a score, where rules with many propositions and operators receive higher scores than rules with few.

$$
\begin{aligned}
&s((\top, \psi)) &&= s(\psi) - 1\\
&s((\varphi, \psi)) &&= s(\varphi) + s(\psi)\\
&s(\varphi \wedge \psi) &&= s(\varphi) + s(\psi) + 1\\
&s(\varphi \vee \psi) &&= s(\varphi) + s(\psi) + 1\\
&s(\neg\varphi) &&= s(\varphi) + 1\\
&s(\top) &&= 0\\
&s(p) &&= 1
\end{aligned}
$$

By applying the highest valued rules (the most specialized) first, we ensure that as few worlds as possible are locked. Notice that that rules where the antecedent is $\top$ will be penalized as they are very general, whereas $\top$ in the consequent is ignored.

The algorithm $generate : (\mathcal{L}, \mathcal{R}) \rightarrow \leq$ then works as follows (algorithm 2). Generate an initial model using $init(\mathcal{L})$. Sort rules descending according to their s-value using $sort(\mathcal{R})$. Each rule in $\mathcal{R}$ is then applied using $apply((\varphi, \psi), \leq)$. Finally, the ordering $\leq$ which respects all successfully applied rules is returned.

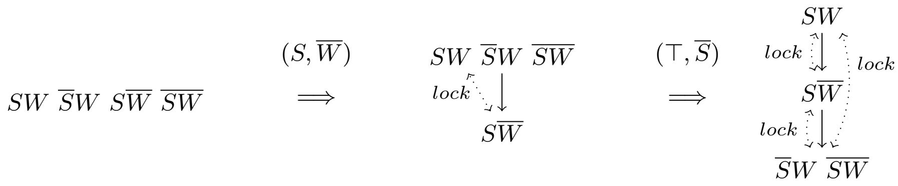

Fig. 1. Generation of Alice's preferences.

5 Case study

We consider a situation in which agents are normally expected to go to work, but during snowy weather, they are not expected to go to work. The agent Alice prefers that it does not snow, but when it snows she wants to stay at home. Thus we have the following rules for expectations of the environment and preferences of the agent:

$$\mathcal{R}_{Env} = \{(\top, \mathit{work}), (\mathit{snow}, \neg\mathit{work})\}$$
$$\mathcal{R}_{Alice} = \{(\top, \neg\mathit{snow}), (\mathit{snow}, \neg\mathit{work})\}.$$

In the following we let S abbreviate *snow* and W *work*. We denote negation using an overline, e.g. $\overline{S}$ when it is not snowing and conjunction is implicitly present when propositions are written next to each other, e.g. SW when it is snowing and the agent is working. From the description above it is clear that $\mathcal{L} = \{W, S\}$. The orderings $\leq_P$ and $\leq_N$ are then generated using the algorithms described above. Figure 1 shows how Alice's preference ordering is generated using her rules.

This situation is however not very interesting, since even when it is not snowing, there is no expected consequence of *not* going to work. We therefore add the possibility of *getting fired* (F) and of *leaving early* (E). Alice's rules are updated accordingly:

$$\mathcal{R}_{Alice} = \{(\top, \overline{S}), (S, \overline{W}), (\top, \overline{F}), (W, E)\}.$$

Thus she does not want to get fired, and in situations where she chooses to go to work, she prefers to leave early. The rules of the environment are updated to conform to this change; if it snows, one can stay home without getting fired, but this is not the case when it does not snow.

$$\mathcal{R}_{Env} = \{(\top, W), (S, \overline{FW}), (\overline{SW}, F), (\top, \overline{E}), (W, \overline{F})\}.$$

Note also that agents are not expected to leave early, and will normally not be fired if they work.

It should be evident that certain worlds are not possible given the new propositions; an agent will not be working if it is fired, and if it is not working, it will not leave early. We therefore express the following prohibitions to be used for cleaning the set of possible worlds.

$$\mathcal{P} = \{FW, E\overline{W}\}.$$

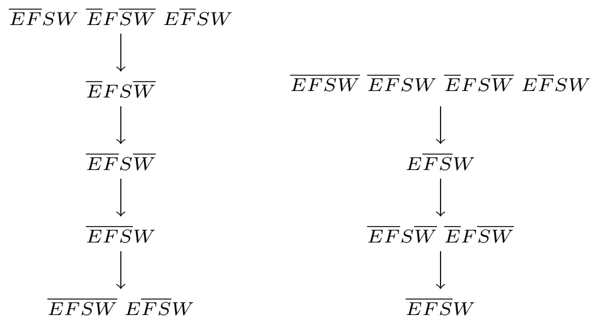

(a) Alice's preferences (b) Expectation

Fig. 2. The preference and normality orderings generated using the rules and prohibitions specified by the environment and Alice.

Thus the set of possible worlds W is then reduced to those worlds where none of the prohibitions above are entailed. The preference and normality orderings resulting from these rules are shown in figure 2(a) and 2(b).

5.1 Making a decision

Given the model generated above, Alice is now able to decide between her influences. Say Alice has a desire to stay at home, but an obligation towards her employer to go to work. Her influences are then $F(a) = \{W, \overline{W}\}$, where a denotes Alice. Furthermore, let us consider two cases; one where it snows, and one where it does not.

a) We have that $B(a) = \{S\}$ so all worlds in which it does not snow can be ignored since Alice believes it snows. This leaves us with four possible worlds. Alice's most preferred world in which it snows is $\overline{EF}S\overline{W}$. This world is more preferred than any other world, thus $Dec(a) = \{\overline{W}\}$.

b) Alice does not believe that it snows, so we have that $B(a) = \{\overline{S}\}$. In this case Alice's most preferred worlds are $\overline{EFSW}$ and $E\overline{FS}W$. As these worlds are equally preferable, she has to consider the expected consequences of either influence. Looking at the normality ordering (figure 2(b)), we realize that $EC(W) = \overline{EFS}$ and $EC(\overline{W}) = \overline{E}F\overline{S}$. From the Alice's preferences it should be clear that the expected consequences of going to work are more preferred than those of staying home (i.e. not getting fired is preferred), thus $Dec(a) = \{W\}$.

Note that at this point we have not labeled Alice as "social" or "selfish". Her preference and the expected consequences are taking into account, and this leads to the results above. When she chooses to go to work, this does not mean that she is strictly social. She might very well have a (selfish) desire to leave early, which could be chosen if the consequences of doing so are tolerable.

6 Conclusion

We have argued that conflicts are prone to arise when agents interact in open societies and enact roles in an organization, since their own desires may be in conflict with obligations towards other agents or the obligations of the role(s) they are enacting. We have discussed why obligations along with desires should be considered influences on the agent's behavior rather than being seen as desires being imposed onto the agent by other entities. Since obligations do not (necessarily) represent propositions the agent wants to achieve, such influences should only be pursued if their consequences can be tolerated by the agent.

Our approach to resolving such influence conflicts, which is based on qualitative decision theory, is an attempt to let the agent reason about the influences without taking into account that one influence is a desire, and another is an obligation, since such bias results in labeling the agent "selfish" or "social" in advance. This approach works by taking the consequence of bringing about a state into consideration, thus letting the agent take its preferences into account, without choosing something that results in an intolerable state. We have argued that this indeed lets the agents reach a decision without strictly preferring desires over obligations or vice versa.

To make the procedure readily available we furthermore have developed a simple method which can generate models to be used in the reasoning process by the use of expressions describing the agent's preferences. By use of a simple locking mechanism, the method generates models which respect non-contradictory rules specified by the agent, such that it is possible to make a decision among a set of influences. The simple nature of the method also allows us to generate the models on the fly, such that if the agent's preferences change during execution a new model can be generated. Since the method works by generating all possible states, it may prove to be inefficient in more complex cases. It would be natural to look into how this can be optimized, e.g. by considering smaller sets of propositions relevant to each rule, which would then be combined into a preference order.

A reasonable direction for further research would be to investigate how to integrate the procedure into an existing agent programming language such as GOAL [7]. In GOAL the choice of committing to different goals and performing actions is quite simple; a program consists of a list of rules which are either evaluated in linear or random order. This means that either the preference ordering is specified a priori, or it is not specified at all. While requiring a specification of the rules of the agent, it should be possible to integrate the procedure in GOAL to evaluate the rules using the decision procedure such that the most preferred and most tolerable goals are pursued.

Furthermore the non-propositional case should be investigated such that more complex reasoning about the agent's preferences can be done. For instance it should be possible for the agent to prefer being at home, *at(home)*, compared to other places such as work, while still being able to express that being at the zoo is more preferred than being at home.

References

1. Natasha Alechina, Mehdi Dastani, and Brian Logan. Programming Norm-Aware Agents. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, AAMAS '12, pages 1057–1064, Richland, SC, 2012. International Foundation for Autonomous Agents and Multiagent Systems.
2. Craig Boutilier. Conditional logics of normality: a modal approach. *Artif. Intell.*, 68(1):87–154, July 1994.
3. Craig Boutilier. Toward a Logic for Qualitative Decision Theory. In *In Proceedings of the KR'94*, pages 75–86. Morgan Kaufmann, 1994.
4. Jan Broersen, Mehdi Dastani, Joris Hulstijn, Zisheng Huang, and Leendert van der Torre. The BOID Architecture – Conflicts Between Beliefs, Obligations, Intentions and Desires . In *Proceedings of the Fifth International Conference on Autonomous Agents*, pages 9–16. ACM Press, 2001.
5. F. Dignum, D. Morley, E. A. Sonenberg, and L. Cavedon. Towards Socially Sophisticated BDI Agents. In *Proceedings of the Fourth International Conference on MultiAgent Systems (ICMAS-2000)*, ICMAS '00, pages 111–118, Washington, DC, USA, 2000. IEEE Computer Society.
6. Frank Dignum, David Kinny, and Liz Sonenberg. From Desires, Obligations and Norms to Goals. *COGNITIVE SCIENCE QUARTERLY*, 2:2002, 2002.
7. Koen V. Hindriks. Programming Rational Agents in GOAL. *Multi-Agent Programming: Languages, Tools and Applications*, 2:119–157, 2009.
8. Leendert van der Torre and Yao-Hua Tan. Contrary-to-duty reasoning with preference-based dyadic obligations. *Annals of Mathematics and Artificial Intelligence*, 27(1-4):49–78, November 1999.