

What You Should Believe: Obligations and Beliefs

Guido Boella¹, Célia da Costa Pereira², Andrea Tettamanzi², Gabriella Pigozzi³, and
Leendert van der Torre³

¹ Università degli Studi di Torino, Dipartimento di Informatica
10149, Torino, C.so Svizzera 185, Italy
guido@di.unito.it

² Università degli Studi di Milano, Dipartimento di Tecnologie dell'Informazione
26013, Crema, via Bramante 65, Italy

{pereira, andrea.tettamanzi}@dti.unimi.it
³ Université du Luxembourg, Computer Science and Communication
L-1359, Luxembourg, rue Richard Coudenhove - Kalergi 6, Luxembourg
{gabriella.pigozzi, leon.vandertorre}@uni.lu

Abstract. This paper presents and discusses a novel approach to indeterministic belief revision. An indeterministic belief revision operator assumes that, when an agent is confronted with a new piece of information, it can revise its belief base in more than one way. We define a rational agent not only in terms of what it believes, as often assumed in belief revision, but also of what it ought or is obliged to do. Hence, we propose that the agent's goals play a role in the choice of (possibly) one of the several available revision options. Properties of the new belief revision mechanism are also investigated.

Keywords. Rational agents, indeterministic belief revision, qualitative decision theory.

1 Introduction

Norms and obligations are increasingly being introduced in Multiagent Systems, in particular to meet the coordination needs of open systems where heterogeneous agents interact with each other. Witness the numerous papers presented at conferences and the organization of workshops like NorMas and COIN in the last years. Introducing norms raise the issue, however, of the interaction between obligations and other mental attitudes like beliefs, goals, and intentions. While the relation between obligations and motivational attitudes is being studied [4,6,5,12,11,19,20,3,10,16], the relation between beliefs and obligations is still unclear. In this paper we study the role of obligations in the task of revising the agent's beliefs under the light of new information. Revising the beliefs can lead to a situation where a choice among different alternatives cannot be made on the basis of the available information. However, obligations and other motivational attitudes can lead a rational agent to choose among the equally likely alternatives, in order not to lose precious opportunities.

For example, suppose that you are a politician who is subject to the obligation to reduce deficit, for example due to a decision of the EU or the IMF, and you believe that

- A) Blocking enrollment leads to a decrease in spending
- B) A decrease of investment in infrastructures leads to a decrease in spending
- C) A decrease in spending leads to a reduction of deficit

Therefore, your plan to meet your obligation is either to block the enrollment, or to decrease investment in infrastructures.

Now, suppose that someone very trustworthy and well-reputed convinces you that blocking enrollment does not lead to a reduction of deficit. Beliefs A and C cannot hold together anymore, and you have to give up one of them.

If you give up A, you still have another possibility to reduce the deficit because you can decrease spending by decreasing investment in infrastructures. However, if you give up C, you do not have any possibility to achieve the reduction of deficit. Indeed,

1. Let us first assume that A is factually wrong, whereas C is true. If you choose to retain (wrong) belief A and to reject C, you will do nothing and you will not succeed in reducing deficit. But, had you kept your belief in C and rejected A, you could have decreased investment in infrastructures in order to decrease spending, and therefore you could have met your obligation to reduce deficit. To conclude, by choosing to maintain A, you risk to miss an opportunity to meet your obligation.
2. Let us now assume that A is actually true and C is wrong. If you choose to keep (wrong) belief C, you will decrease spending, but you will not achieve the goal of reducing deficit. However, even if you had chosen the right revision, i.e., to retain A and reject C, there was no way for you to achieve your goal of reducing deficit. To conclude, by choosing C (wrong), you believed you could achieve a goal when you could not, so you will be disappointed for trying in vain, but at least you tried.

The moral of the story is that, if you are interested only in meeting your obligation (and there are no other goals relevant for you), choosing to maintain C — *even when it is factually wrong*, but you do not know whether it is false or not — is the only rational choice. This is because, independently of C being right or wrong, by choosing that belief you will be better off. Moreover, in one situation — the former — you will be better off if you choose C than if you choose A. Summarizing, you should drop A, because that way, you keep all possibilities to achieve your goal open.

We can formalize the above example, by defining the following atomic propositions:

- b blocking enrollment;
- s decrease spending;
- d reduce deficit;
- i decrease investment in infrastructures.

The belief base before being convinced that blocking enrollment does not lead to a reduction of deficit ($\neg(b \supset d)$) would contain the three formulas $b \supset s$, $i \supset s$, and $s \supset d$. You have to, first of all, reduce deficit, d , and, if possible, not decrease investment in infrastructures, $\neg i$. Adding $\neg(b \supset d)$ to your beliefs would make them inconsistent. Therefore, you have to revise your beliefs by giving up either $b \supset s$ or $s \supset d$. The choice you make may depend on the obligations you can meet in the alternatives: if you give up $b \supset s$, your plan will be to decrease investment in infrastructures, so you will not achieve $\neg i$, but might succeed in achieving d ; if you give up $s \supset d$, your plan will

be to do nothing, so you will certainly not achieve d , but you will fulfill $\neg i$. Depending on the punishment you have for violating your obligation to achieve d or $\neg i$, you could prefer one or the other alternative.

We use the deficit reduction example as a running example throughout the paper.

The choice among belief bases is distinct from other decision problems, due to the possibility of wishful thinking. Consider for example that you have to block the enrollment (b) to decrease spending, and that this obligation is more important than the obligation of reducing deficit (d). What will you do? At least in a naive approach, you could reason by cases as follows. Assume you choose $b \supset s$: in that case you believe that accomplishing the obligation of blocking enrollment leads to a decrease in spending. Assume you choose $s \supset d$: in that case you believe you will achieve the goal of reducing deficit. Since b is more important than d , you choose $b \supset s$.

Instead, the idea of this paper is inspired by the notion of *conventional wisdom* (CW) as introduced by economist John Kenneth Galbraith:

We associate truth with convenience, with what most closely accords with self-interest and personal well-being. ([14])

That is, CW consists of “ideas that are convenient, appealing”. This is the rationale for keeping them. One basic brick of CW could then be the fact that some ideas are maintained because they maximize the goals that the agents (believe) they can achieve. This work may be seen as an initial attempt to formally capture the concept of a CW agent. In the following we provide a logical framework that models how a CW agent revises its beliefs under its obligations.

The paper is structured as follows. In Section 2 we introduce the aim of this paper, the used methodology and particular challenges encountered. In Section 3 we introduce the agent theory we use in our approach, and in Section 4 we introduce an indeterministic belief change operator in this agent theory. In Section 5 we define the choice among beliefs as a decision problem in the agent theory. Section 6 concludes.

2 Aim, Methodology and Challenges

The research problem of this paper is to develop a formal model to reason about the kind of choices among belief bases discussed in the previous section, and to generalize the example above in case of additional beliefs, multiple goals with different importance, conditional obligations, a way to take violated goals into account, and so on.

We use a combination of the framework of belief revision together with a qualitative decision theory. Classical approaches to belief revision assume that, when an agent revises its belief base in view of new input, the outcome is well-determined. This picture, however, is not realistic. When an agent revises its beliefs in the light of some new fact, it often has more than one available alternative. Approaches to belief revision that do not stipulate the existence of a single revision option are called *indeterministic* [18,24]. In this paper we suggest that one possible policy an agent can use in order to choose among available alternatives is to check the effect of the different revisions on the agent’s set of goals.

Moreover, for the qualitative decision theory we are inspired by agent theories such as the BOID architecture [4,6,5,12,11,19,20,3,10], the framework of goal generation in 3APL as developed by van Riemsdijk and colleagues [26], and [8]. In particular, our agent model is based on one of the versions of 3APL, because the belief base in the mental state of a 3APL agent is a consistent set of propositional sentences, just like in the framework of belief revision. However, we do not care about how goals are generated and how their achievability (plan existence) is established. That is because we do not include “goal-adoption rules” or “practical reasoning rules” representing which action to choose in a particular state. We assume that there is a planning module, which would take a set of goals, actions, and an initial world state representation in input and produce a solution plan in output. This planning module might rely on the well-known graphplan algorithm [2], or any other propositional AI planner: as in object-oriented programming, we *encapsulate* the planner within a well-defined interface and overlook the implementation details of how a solution plan is found. This is in line with, on one hand, the BOID architecture [4], where the planning component is kept separate from the remainder of agent deliberation, and, on the other hand with the works of Móra and colleagues describing the relationship between propositional planning algorithms and the process of means-end reasoning in BDI agents. In these works, [21,22], it is shown how the mental state of an agent can be mapped to the STRIPS [13] notation forth and back. This relation has been done on an abstract BDI interpreter named X-BDI [27,7] and augmented with graphplan.

In other words, we model the choice among belief bases essentially as a decision problem, that is, as a choice among a set of alternatives. We do not use classical decision theory (utility function, probability distribution, and the decision rule to maximize expected utility), but a qualitative version based on maximizing achieved goals and minimizing violated goals in an abstract agent theory (see e.g. [9] for various approaches to formalize the decision process of what an agent should do), because such qualitative decision theories include beliefs and therefore are easier to combine with the theory of belief revision. However, what precisely are the alternatives?

An indeterministic belief revision operator associates multiple revision options to a belief base that turns out to be inconsistent as a consequence of a new piece of information. Our revision mechanism selects the revision alternative that allows the agent to maximize its achievable goals. However, it will not always be possible to select exactly one revision alternative. For example, there may be one most important goal set but two revision alternatives that lead the agent to achieve it. In this case, the two belief revision candidates are said to be equivalent. In Section 5.3 we will provide conditions under which a revision for a CW agent is deterministic, that is, when our revision operator can select exactly one revision alternative.

Besides the issue of wishful thinking, another complicating factor when choosing among belief bases in the context of conditional obligation rules, is that a maximization of goals may lead to a meta-goal to derive obligations by choosing revisions where you believe that the condition is true and the obligation applies. However, deriving goals by itself does not have to be desirable. In contrast, it may even be argued that fewer goals are better than more goals, as you risk to violate goals and become unhappy (as in Buddhism). We therefore also take goal violations into account.

3 An Abstract Agent Theory

In this section, we represent the formalism which is used throughout the paper.

3.1 A Brief Introduction to AI Planning and Agent Theory

Any agent, be it biological or artificial, must possess knowledge of the environment it operates in, in the form of, e.g., *beliefs*. Furthermore, a necessary condition for an entity to be an *agent* is that it acts. We shall call the factors that motivate an agent to act *obligations*. For artificial agents, obligations may be the purposes an agent was created for.

Obligations are necessary, not sufficient, conditions for action. When an obligation is met by other conditions that make it possible for an agent to act, that obligation becomes a *goal*.

The reasoning side of acting is known as practical reasoning or deliberation, which may include *planning*. Planning is a process that chooses and organizes actions by anticipating their expected effects with the purpose of achieving as good as possible some pre-stated objectives or goals.

The objective of our formalism is to analyze, not to develop, agent systems. More precisely, our agent must single out the set of goals to be given as input to a traditional planner. That is because the intentions of the agent are not considered. We merely consider beliefs (knowledge the agent has about the world states), obligations (or motivations) and relations (obligation-adopting rules) defining how the obligation base will change with the acquisition of new beliefs and/or new obligations. The goal generation process that underlies this work is very much in line with the work carried out in [25] on *oversubscription planning problems*, in which the main objective is to find the maximal set of desires to be reached in a given period and with a limited quantity of resources, and with goal generation in the BOID architecture [4].

3.2 Beliefs, Obligations, and Goals

The basic components of our language are *beliefs* and *obligations*. Beliefs are represented by means of a *belief base*. A belief base is a finite and consistent set of propositional formulas describing the information the agent has about the world and internal information. Obligations are represented by means of an *obligation base*. An obligation base consists of a set of propositional formulas which represent the situations the agent has to achieve. However, unlike the belief base, an obligation base may be inconsistent, e.g., $\{p, \neg p\}$.

Definition 1 (Belief Base B and Obligation Base O) *Let \mathcal{L} be a propositional language with \top a tautology, and the logical connectives \wedge and \neg with the usual meaning. The agent's belief base B is a consistent finite set such that $B \subseteq \mathcal{L}$. B can also be represented as the conjunction of its propositional formulas. The agent's obligation base is a possibly inconsistent finite set of sentences denoted by O , with $O \subseteq \mathcal{L}$.*

We define two modal operators Bel and Obl such that, for any formula ϕ of \mathcal{L} , $\text{Bel}\phi$ means that ϕ is believed whereas $\text{Obl}\phi$ means that the agent has obligation ϕ . Since the belief and obligation bases of an agent are completely separated, there is no need to nest the operators Bel and Obl.

Definition 2 (Obligation-Adoption Rule) An obligation-adoption rule is a triple $\langle \phi, \psi, \tau \rangle \in \mathcal{L} \times \mathcal{L} \times \mathcal{L}$ whose meaning is: if $\text{Bel}\phi$ and $\text{Obl}\psi$, then τ will be adopted as an obligation as well.

The set of obligation-adoption rules is denoted by R . If $\exists \phi', \psi', \tau'$ such that $\phi \leftrightarrow \phi', \psi \leftrightarrow \psi', \tau \leftrightarrow \tau'$, then $\langle \phi', \psi', \tau' \rangle \in R$.

Goals, in contrast to obligations, are represented by consistent obligation sets. There are various ways to generate candidate goal sets from the obligation adoption rules, as discussed in the remainder of this section.

Definition 3 (Candidate Goal Set) A candidate goal set is a consistent subset of O .

3.3 Mental State Representation

We assume that an agent is equipped with three components:

- belief base $B \subseteq \mathcal{L}$;
- obligation base: $O \subseteq \mathcal{L}$;
- obligation-adoption rule set R .

The mental state S of an agent is completely described by a triple $S = \langle B, O, R \rangle$. In addition, we assume that each agent can be described using a problem-dependent function \mathcal{V} , a goal selection function G , and a belief revision operator $*$, as discussed below.

In our deficit reduction example, we have:

$$\begin{aligned} B &= \{\neg(b \wedge \neg s), \neg(i \wedge \neg s), \neg(s \wedge \neg d)\}, \\ O &= \{d, \neg i\}, \\ R &= \{\langle \top, \top, d \rangle; \langle \top, \top, \neg i \rangle\}. \end{aligned}$$

The semantics we adopt for the belief and obligation operators are standard.

Definition 4 (Semantics of Bel operator) Let $\phi \in \mathcal{L}$, $\text{Bel}\phi \Leftrightarrow B \models \phi$.

Definition 5 (Semantics of Obl operator) Let $\phi \in \mathcal{L}$, $\text{Obl}\phi \Leftrightarrow \exists$ a maximal consistent subset $O' \subseteq O$ such that $O' \models \phi$.

We expect a rational agent to try and manipulate its surrounding environment to fulfill its goals. In general, given a problem, not all goals are *achievable*, i.e. it is not always possible to construct a plan for each goal. The goals which are not achievable or those which are not chosen to be achieved are called *violated goals*. Hence, we assume a problem-dependent function \mathcal{V} that, given a belief base B and a goal set $O' \subseteq O$, returns a set of couples $\langle O^a, O^v \rangle$, where O^a is a maximal subset of achievable goals and O^v is the subset of violated goals and is such that $O^v = O' \setminus O^a$. Intuitively, by considering violated goals we can take into account, when comparing candidate goal sets, what we lose from not achieving certain goals.

3.4 Comparing Goals and Sets of Goals

The aim of this section is to illustrate a qualitative method for goal comparison in the agent theory. More precisely, we define a qualitative way in which an agent can choose among different sets of candidate goals. Indeed, from an obligation base O , several candidate goal sets O_i , $1 \leq i \leq n$, may be derived. How can an agent choose among all the possible O_i ? It is unrealistic to assume that for a rational agent all goals have the same priority. We use the notion of importance of obligations to represent how relevant each goal is for the agent depending, for instance, on the punishment for violating the obligations. The idea is that a rational agent tries to choose a set of candidate goals which contains the greatest number of achievable goals (or the least number of violated goals).

We assume we dispose of a total order \succeq over an agent's obligations. In the example, you have to reduce, in the first place, deficit and, if possible, you should not decrease investments in infrastructures. Therefore, d is more important than $\neg i$, in symbols $d \succeq \neg i$.

The \succeq relation can be extended from goals to sets of goals. We have that a goal set O_1 is more important than another one O_2 if, considering only the goals occurring in either set, the most important goals are in O_1 or the least important goals are in O_2 . Note that \succeq is connected and therefore a total pre-order, i.e., we always have $O_1 \succeq O_2$ or $O_2 \succeq O_1$.

Definition 6 (Equivalent Goals)

A goal ϕ_1 is said equivalent to a goal ϕ_2 , noted $\phi_1 \approx \phi_2$, if and only if ϕ_1 and ϕ_2 are equally important, i.e. $\phi_1 \succeq \phi_2$ and $\phi_2 \succeq \phi_1$.

Definition 7 (Difference Goal set Operator)

Let O_1 and O_2 be two sets of goals. The difference based on the equivalence between goals in O_1 and in O_2 noted $O_1 \setminus \approx O_2$, is defined as follow:

$$O_1 \setminus \approx O_2 = \{\phi_1 \in O_1 \mid \neg \exists \phi_2 \in O_2 \text{ such that } \phi_1 \approx \phi_2\}$$

Definition 8 (Relative Importance of Sets of Goals)

Let $O'_1 = O_1 \setminus \approx O_2$ and $O'_2 = O_2 \setminus \approx O_1$. The goal set O_1 is at least as important as O_2 , denoted $O_1 \succeq O_2$ iff

$$O'_2 = \emptyset \text{ or } \exists \phi_1 \in O'_1, \forall \phi_2 \in O'_2 \phi_1 \succeq \phi_2.$$

In our example, it is easy to verify that $\{d, \neg i\} \succ \{d\} \succ \{\neg i\} \succ \emptyset$. However, we also need to be able to compare the mutual exclusive subsets (achievable and violated goals) of the considered candidate goal, as defined below.

3.5 Comparing Couples of Goal Sets

We propose two methods to compare couples of goal sets.

3.5.1 The Direct Comparison \succeq_D

Given the \succeq_D criterion, a couple of goal sets $\langle O_1^a, O_1^v \rangle$ is at least as important as the couple $\langle O_2^a, O_2^v \rangle$, noted $\langle O_1^a, O_1^v \rangle \succeq_D \langle O_2^a, O_2^v \rangle$ iff $O_1^a \succeq O_2^a$ and $O_1^v \preceq O_2^v$.

\succeq_D is reflexive and transitive but partial. $\langle O_1^a, O_1^v \rangle$ is strictly more important than $\langle O_2^a, O_2^v \rangle$ in two cases:

1. $O_1^a \succ O_2^a$ and $O_1^v \prec O_2^v$, or
2. $O_1^a \succ O_2^a$ and $O_1^v \preceq O_2^v$.

They are indifferent when $O_1^a = O_2^a$ and $O_1^v = O_2^v$. In all the other cases, they are not comparable.

3.5.2 The Lexical Comparison \succeq_{Lex}

Given the \succeq_{Lex} criterion, a couple of goal sets $\langle O_1^a, O_1^v \rangle$ is at least as important as the couple $\langle O_2^a, O_2^v \rangle$ (noted $\langle O_1^a, O_1^v \rangle \succeq_{Lex} \langle O_2^a, O_2^v \rangle$) iff $O_1^a = O_2^a$ and $O_1^v = O_2^v$; or there exists a $\phi \in \mathcal{L}$ such that:

1. $\forall \phi' \succeq \phi$, the two couples are indifferent, i.e., one of the following possibilities holds:
 - a) $\phi' \in O_1^a \cap O_2^a$;
 - b) $\phi' \notin O_1^a \cup O_1^v$ and $\phi' \notin O_2^a \cup O_2^v$;
 - c) $\phi' \in O_1^v \cap O_2^v$.
2. Either of the following holds:
 - a) $\phi \in O_1^a \setminus O_2^a$;
 - b) $\phi \in O_2^v \setminus O_1^v$.

\succeq_{Lex} is reflexive, transitive, but partial.

3.6 Defining the Goal Set Selection Function

In general, given a set of obligations O , there may be many possible candidate goal sets. A rational agent in state $\mathcal{S} = \langle B, O, R \rangle$ will select one precise candidate goal set O' which consists of the most important couple of achievable and violated goals.

Let us call G the function which maps a state \mathcal{S} into the goal set selected by a rational agent in state \mathcal{S} . G is such that $G(\mathcal{S}) = O'$.

4 Situating the Problem: Indeterministic Belief Change

“Most models of belief change are deterministic. Clearly, this is not a realistic feature, but it makes the models much simpler and easier to handle, not least from a computational point of view. In indeterministic belief change, the subjection of a specified belief base to a specified input has more than one admissible outcome.

Indeterministic operators can be constructed as sets of deterministic operations.

Hence, given n deterministic revision operators $*_1, *_2, \dots, *_n, * = \{*_1, *_2, \dots, *_n\}$ can be used as an indeterministic operator.” [17]

Let us consider a belief base B and a new belief β . The revision of B in light of β is simply:

$$B * \beta \in \{B *_1 \beta, B *_2 \beta, \dots, B *_n \beta\}. \quad (1)$$

More precisely, revising the belief base B with the indeterministic operator $*$ in light of the new belief β leads to one of the n belief revision results:

$$B * \beta \in \{B_\beta^1, B_\beta^2, \dots, B_\beta^n\}, \quad (2)$$

where B_β^i is the i -th possible belief revision result.

Applying the operator $*$ is then equivalent to applying one of the virtual operators $*_i$ contained in its definition. While the rationality of an agent does not suggest any criterion to prefer one revision over the others, a defining feature of a CW agent is that it will choose which revision to adopt based on the consequence of that choice. One important consequence is the set of goals the agent will decide to pursue.

In our deficit reduction example, $\beta = \text{Bel}(b \wedge \neg d)$, and

$$B * \beta \in \left\{ \begin{array}{l} B_\beta^1 = \{b \wedge \neg d, \neg(s \wedge \neg d), \neg(i \wedge \neg s)\}, \\ B_\beta^2 = \{b \wedge \neg d, \neg(b \wedge \neg s), \neg(i \wedge \neg s)\} \end{array} \right\}. \quad (3)$$

In the next section we propose some possible ways to tackle the problem of choosing one of the revision options.

5 Belief Revision as a Decision Problem

By considering an indeterministic belief revision, we admit $B * \beta$ to have more than one possible result. In this case, the agent must select (possibly) one among all possible revisions. Among the possible criteria for selection, one is to choose the belief revision operator for which the goal set selection function returns the most important goal set. In other words, selecting the revision amounts to solve an optimization problem.

5.1 Indeterministic State Change

The indeterminism of belief revision influences the obligation-updating process. In fact, the belief revision operator is just a part of the state-change operator, which is indeterministic as well, as a consequence of the indeterminism of belief revision. Therefore, $\mathcal{S}_\beta \in \{\mathcal{S}_\beta^1, \mathcal{S}_\beta^2, \dots, \mathcal{S}_\beta^n\}$, where $\mathcal{S}_\beta^i = \langle B_\beta^i, O_\beta^i, R \rangle$.

Which goal set is selected by an agent depends on G :

$$G(\mathcal{S}_\beta) \in \{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}. \quad (4)$$

In the example, $G(\mathcal{S}_\beta) \in \{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2)\}$, where $G(\mathcal{S}_\beta^1) = \{d\}$ and $G(\mathcal{S}_\beta^2) = \{\neg i\}$. The following table summarizes the possibilities the agent may face when choosing between the two alternative revisions.

reality \rightarrow \downarrow beliefs	$\not\models b \supset s$ $\models s \supset d$	$\models b \supset s$ $\not\models s \supset d$
B_β^1 decrease investment in infrastructures	d is achieved $\neg i$ is not achieved	no obligation is met
B_β^2 do nothing	d is not achieved $\neg i$ is achieved	

A traditional rational agent could not choose one of the $G(\mathcal{S}_\beta^i)$ because they are incomparable. Now, for a CW agent,

$$G(\mathcal{S}_\beta) \in \mathbb{I}\{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}, \quad (5)$$

where $\mathbb{I}(S)$ denotes the most important set of S defined as follows:

Definition 9 (Important Set I) *Given two sets S and X such that $S \subseteq X$, and given an importance relation \succeq over X , the most important set of S is*

$$\mathbb{I}(S) = \{x \in S : \neg \exists x' \in S, x' \succ x\}. \quad (6)$$

5.2 Choosing a Revision

Choosing the most important revision option is not a trivial operation. We can distinguish two situations:

- there is just one most important goal set O' , but more than one alternative options leads to O' ;
- there is no unique most important goal set; that is, there are different goal sets O_1, \dots, O_m , none of which is strictly more important than the others, i.e., for all $i, j \in \{1, \dots, m\}$, $O_i \succeq O_j$.

Definition 10 (Equivalent Belief Revision Candidates) *A belief revision candidate B_β^1 is equivalent to another belief revision candidate B_β^2 (denoted by $B_\beta^1 \approx B_\beta^2$), if and only if $G(\mathcal{S}_\beta^1) \succeq G(\mathcal{S}_\beta^2)$ and $G(\mathcal{S}_\beta^2) \succeq G(\mathcal{S}_\beta^1)$.*

It is easy to verify that \approx is a standard equivalence relation, i.e., reflexive, symmetric, and transitive.

The choice of which revision outcome to adopt may thus be deterministic or indeterministic. It is indeterministic in the two cases presented above. More precisely, the choice depends on the importance relations over the goal sets, which determine the equivalence between revision candidates:

- if $\|\mathbb{I}\{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}\| = 1$, i.e., the equivalent class of an important belief revision is a singleton and, if there is no i, j such that $G(\mathcal{S}_\beta^i) = G(\mathcal{S}_\beta^j)$, the choice of the belief operator is obviously deterministic;
- if $\|\mathbb{I}\{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}\| = 1$, and there is at least a couple i, j such that $G(\mathcal{S}_\beta^i) = G(\mathcal{S}_\beta^j)$, the choice is indeterministic, but also indifferent;
- if $\|\mathbb{I}\{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}\| > 1$, the choice is indeterministic.

It is important to notice that an agent that has to choose between $G(\mathcal{S}_\beta^i)$ and $G(\mathcal{S}_\beta^j)$ is in a different situation than an agent that has to randomly choose among a number of competing revisions. The reason is that a random choice is hardly a rational option. But, when an agent must choose between two revision options, it knows that, no matter which revision it chooses, the outcome does not change. In such a context, a random choice becomes a rational option.

Proposition 1 *Let $*$ be an indeterministic belief operator, and n be the number of possible belief revisions candidate. We have:*

$$1 \leq \|I\{G(\mathcal{S}_\beta^1), G(\mathcal{S}_\beta^2), \dots, G(\mathcal{S}_\beta^n)\}\| \leq n.$$

5.3 Conditions for Determinism of a CW Agent

Traditional indeterministic belief revision approaches allow for the result of belief revision to be indeterminate in the sense that there may be many possible revision alternatives that are equally rational. Our proposal builds on the idea that what an agent wishes to achieve can play a role in the choice of which beliefs to reject and which beliefs to retain. The example we have been using in this paper also tries to capture the intuition that an agent who behaves in this manner is rational. Our richer model can distinguish one revision alternative from the other depending on the effect that each option has on the agent's goal set. Hence, under certain conditions, the choice among several revision alternatives can be reduced to one. This is what we want to investigate now, that is we want to investigate the conditions under which a revision for a CW agent is deterministic even if an indeterministic revision operator is used, i.e., $\|I\{G(\mathcal{S}_\beta^i)\}_{i=1, \dots}\| = 1$ and, for all i, j , $G(\mathcal{S}_\beta^i) \neq G(\mathcal{S}_\beta^j)$.

Observation 1 *$B * \beta$ is deterministic in state $\mathcal{S} = \langle B, O, R \rangle$, iff no two alternative revisions are equivalent, i.e., for all i, j , $B_\beta^i \not\approx B_\beta^j$.*

Proposition 2 *A sufficient condition for no two alternative revisions, B_β^i and B_β^j , being equivalent is that*

1. *for all i, j , $G(\mathcal{S}_\beta^i) \neq G(\mathcal{S}_\beta^j)$;*
2. *the importance relation on goals is strict, i.e., for all $\phi, \phi' \in G(\mathcal{S}_\beta)$, $\phi \neq \phi'$, $\phi \succeq \phi' \Rightarrow \phi' \not\succeq \phi$.*

Proof: From Hypothesis 1 and 2, by applying Definition 8, we obtain $B_\beta^i \not\approx B_\beta^j$. Therefore, no two alternative revisions can be equivalent. \square

6 Conclusions

A new framework, inspired by the concept of *conventional wisdom*, aiming at dealing with indeterminism in belief revision has been proposed. While a traditional agent would not be able to choose among multiple revision candidates in indeterministic belief revision, a CW agent evaluates the effects the different revision options have on its goals and selects the revision which maximizes its achievable goals. Fundamental definitions and properties of such belief revision mechanism have been given.

References

1. Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *J. Symb. Log.*, 50(2):510–530, 1985.
2. Avrim Blum and Merrick L. Furst. Fast planning through planning graph analysis. *Artif. Intell.*, 90(1-2):281–300, 1997.
3. G. Boella, J. Hulstijn, and L. van der Torre. Interaction in normative multi-agent systems. *Electronic Notes in Theoretical Computer Science*, 141(5):135–162, 2005.
4. J. Broersen, M. Dastani, J. Hulstijn, and L. van der Torre. Goal generation in the BOID architecture. *Cognitive Science Quarterly Journal*, 2(3–4):428–447, 2002.
5. J. Broersen, M. Dastani, and L. van der Torre. Realistic desires. *Journal of Applied Non-Classical Logics*, 12(2):287–308, 2002.
6. J. Broersen, M. Dastani, and L. van der Torre. Beliefs, obligations, intentions and desires as components in an agent architecture. *International Journal of Intelligent Systems*, 20:9:893–919, 2005.
7. Michael da Costa Móra, José Gabriel Pereira Lopes, Rosa Maria Vicari, and Helder Coelho. Bdi models and systems: Bridging the gap. In *ATAL*, pages 11–27, 1998.
8. C. da Costa Pereira and A. Tettamanzi. Towards a framework for goal revision. In Wim Vanhoof Pierre-Yves Schobbens and Gabriel Schwanen, editors, *BNAIC-06, Proceedings of the 18th Belgium-Netherlands Conference on Artificial Intelligence*, pages 99–106. University of Namur, 2006.
9. M. Dastani, J. Hulstijn, and L. van der Torre. How to decide what to do? *European Journal of Operational Research*, 160(3):762–784, february 2005.
10. M. Dastani and L. van der Torre. Specifying the merging of desires into goals in the context of beliefs. In *Proceedings of in Information and Communication Technology (EurAsia ICT 2002)*, LNCS 2510, pages 824–831. Springer, 2002.
11. M. Dastani and L. van der Torre. Games for cognitive agents. In *Proceedings of JELIA04*, LNAI 3229, pages 5–17. 2004.
12. M. Dastani and L. van der Torre. What is a normative goal? towards goal-based normative agent architectures regulated agent-based systems. LNAI 2934, pages 210–227. Springer, 2004.
13. Richard Fikes and Nils J. Nilsson. Strips: A new approach to the application of theorem proving to problem solving. *Artif. Intell.*, 2(3/4):189–208, 1971.
14. John K. Galbraith. *The Affluent Society*. Houghton Mifflin, Boston, 1958.
15. P. Gärdenfors. The dynamics of belief systems: Foundations vs. coherence. *Revue Internationale de Philosophie*, 1989.
16. Guido Governatori, Antonino Rotolo, and Vineet Padmanabhan. The cost of social agents. In *AAMAS*, pages 513–520, 2006.
17. Sven Ove Hansson. Logic of belief revision. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Summer 2006.
18. S. Lindström and W. Rabinowicz. Epistemic entrenchment with incomparabilities and relational belief revision. In A. Fuhrmann and M. Morreau, editors, *The Logic of Theory Change*, pages 93–126. 1991.
19. D. Makinson and L. van der Torre. Input-output logics. *Journal of Philosophical Logic*, 29:383–408, 2000.
20. D. Makinson and L. van der Torre. Constraints for input-output logics. *Journal of Philosophical Logic*, 30(2):155–185, 2001.
21. Felipe Rech Meneguzzi, Avelino Francisco Zorzo, and Michael da Costa Móra. Mapping mental states into propositional planning. In *Proceedings of the 3rd International Joint Conference on Autonomous Agents and Multiagent Systems*. ACM Press, 2004.

22. Felipe Rech Meneguzzi, Avelino Francisco Zorzo, and Michael Da Costa Móra. Propositional planning in BDI agents. In *Proceedings of the 2004 ACM Symposium on Applied Computing*, pages 58–63, Nicosia, Cyprus, 2004. ACM Press.
23. Bernhard Nebel. A knowledge level analysis of belief revision. In R. Brachman, H. J. Levesque, and R. Reiter, editors, *Principles of Knowledge Representation and Reasoning: Proceedings of the 1st International Conference*, pages 301–311, San Mateo, 1989. Morgan Kaufmann.
24. Erik J. Olsson. Lindström and Rabinowicz on relational belief revision. In T. Ronnow-Rasmussen, B. Petersson, J. Josefsson, and D. Egonsson, editors, *Hommage à Wlodek. Philosophical Papers Dedicated to Wlodek Rabinowicz*. 2007.
25. David E. Smith. Choosing objectives in over-subscription planning. In Shlomo Zilberstein, Jana Koehler, and Sven Koenig, editors, *Proceedings of the Fourteenth International Conference on Automated Planning and Scheduling (ICAPS 2004)*, pages 393–401, Whistler, British Columbia, Canada, June 3–7 2004. AAAI.
26. M. Birna van Riemsdijk. *Cognitive Agent Programming: A Semantic Approach*. PhD thesis, University of Utrecht, 2006.
27. A.O. Zamberlam, L.M.M. Giraffa, and C.M. Móra. X-bdi: uma ferramenta para programao de agentes bdi (in portuguese). Technical Report 9, PPGCC/PUCRS, 2000.