An Introduction to Intention Revision: Issues and Problems

José Martín Castro-Manzano

Instituto de Investigaciones Filosóficas
Universidad Nacional Autónoma de México
Circuito Mario de la Cueva s/n Ciudad Universitaria, 04510 Coyoacán, México
jmcmanzano@hotmail.com
http://www.filosoficas.unam.mx

Abstract. The change of beliefs on the basis of new information has been widely studied; however, the change of other mental states has received less attention, and particularly, intentions. Despite there are philosophical and formal theories about intentions, few of them consider the revision of intentions.We suggest introductory guidelines to define a research program for the revision of intentions regarding that: (i) intentions are intimately related to the beliefs and desires of agents immersed in a dynamic world; (ii) intentions are directly related to planning; and (iii) a reconsideration function is needed.

Key words: Intention, reconsideration, BDI, agents.

1 Introduction

Belief revision is a paradigmatic research program: it is a relatively new area of research that joins two diciplines: computer science and philosophy. Since programmers dealt with databases they faced the problem of updating their information. On the other hand, certain philosophers dealt with the change of information within epistemic structures. So, we can identify, respectively, two important moments in the history of this research program: one in [6]; and the other in [9] and in [12]. A general theory can be found in [1]. This last approach constitutes the core for any program of belief revision.

Thus, although the change of beliefs on the basis of new information has been widely studied with success during the last 25 years, the dynamic process of other mental states has received less attention, and particularly, intentions [10]. Certainly, there are philosophical and formal theories of intention [2], [3], [4], [5], [11], [13] but few of them, if any, consider the possibility of the revision of intentions [10].

In this work we suggest some general and introductory guidelines in order to define a program for intention revision. We think this topic is important because *(i)* intentions are intimately related to the beliefs and desires of the agents immersed in a dynamic world; *(ii)* intentions are related directly with planning; and *(iii)* a function of reconsideration is needed.

The general background of this work assumes the theories of intention as represented by [2], [3], [4]; and the belief revision program as represented by [1].

The rest of the paper is distributed in the next way: in section 2 we describe what do we mean by intention revision and we describe some methodological problems. In section 3 we discuss some issues regarding the problem of representation. In section 4 we adapt and suggest some general postulates for the revision of intentions. Finally, in section 5 we discuss the ideas of this introduction and we give some details about future work.

2   Intention revision

We can study intentions from two general perspectives. One internal, e.g., what are intentions and how do they behave; the other, external, regarding the problems intentions generate, e.g., how do they relate to other mental states and how those relations can be modelled. We will follow this double approach.

2.1   Internal perspective

For our introduction we will need an approach based upon the BDI model of rational agency [16], [13]. This model receives its name from the use of Beliefs, Desires and Intentions in order to model the rationality of agents. Intuitively, the beliefs correspond to the information the agent has about itself and its environment. Desires correspond to the motivational part of the agent, what the agent wants to see as accomplished. Finally, the intentions correspond to the deliberative part and consist in the desires the agent is commited to achieve.

Intentions, as an irreducible component of the BDI model [2], have certain features that, taken together, make them different from beliefs and desires:

- **Pro-activity.** Intentions are pro-active, they move the agent to achieve a goal [2]. In this sense, intentions are conduct-controlling components. It is important to note, however, that intentions are not equal to desires. Both intentions and desires are pro-attitudes, but intentions imply commitment and consistency, while desires do not.
- **Inertia.** Intentions also possess inertia, that is to say, once an intention has been taken, it resists being abandoned. If the intention was adopted and inmediately abandoned, we would have to say the intention was never taken; however, if the reason that generated the intention disappears, it is rational to abandon the intention [2].
- **Admissibility.** Intentions also provide a filter of admissibility. Once an intention has been taken, this constraints the future practical reasonings: while the agent holds a particular intention, the agent will not consider contradictory options. Thus, intentions provide a filter [2], [4].

In this way, we can say that intentions require a notion of commitment (given the principle of pro-activity), a notion of consistency (given the admissibility criteria) and a notion of retractability (given the notion of inertia).

Plans, as long as they are sets of actions, are intentions and in this sense, they share the same properties: they are conduct-controlling, they have inertia and they work as inputs for future practical reasonings [2]. Moreover, plans have certain features:

– **Plans are partial.** Plans are partial, and not complete, because they lack complete information about the state of the world, e.g., the environment is not accesible.
– **Plans are not statical.** Plans cannot be static structures because the environment of the agent is dynamic.
– **Plans are hierarchical.** Plans contain means-ends reasons that have to have an ordered process.

But plans also require the next features:

– **Internal consistency.** Plans must be executable.
– **Strong consistency.** Plans must be consistent with the agent's beliefs.
– **Means-ends coherence.** The means-ends reasoning of the plan must be consistent with the global ends of the plan.

These last features lead us to consider some other problems: that intentions are not isolated mental states [10]. Modifying intentions implies modifying beliefs, and sometimes, modifying beliefs may modify intentions. In this sense, the strong consistency shows us that beliefs and intentions mantain certain relationships: the asymmetry thesis. Bratman [2] considers these relations as principles of rationality:

– **Intention-belief inconsistency.** It is irrational for an agent to intent ϕ and believe at the same time that it will not achieve ϕ.
– **Intention-belief incompletness.** It is rational for an agent to intent ϕ and at the same time not believe that it will achieve ϕ.

Thus, we can say that the notions of consistency and retractability are not exclusive of beliefs; and that the difficulty of considering intention revision lies in the relation between intentions and beliefs.

2.2 External perspective

Based on Bratman, Cohen and Levesque [4] suggested seven ideas -or problems- that a theory of intentions must take into account:

– 1. Intentions pose problems for agents, who need to determine ways of achieving them.
– 2. Intentions provide a filter for adopting other intentions, which must not conflict.
– 3. Agents track the success of their intentions, and are inclined to try again if their attempts fail.
– 4. Agents believe their intentions are possible.

- 5. Agents do not believe they will not bring about their intentions.
- 6. Under certain circumstances, agents believe they will bring about their intentions.
- 7. Agents need not intend all the expected side effects of their intentions.

With these criteria, Cohen and Levesque construct a formal theory of intention based on the notion of persistent goal (according to them, an intention is a form of persistent goal [4]). However, this theory does not deal with the dynamics of intentions [10], [15]. The dynamics of intentions should deal with the problem of how an agent adopts and abandons intentions and what changes these processes produce in other BDI components. The dynamics of intentions requires a theory of intention revision, in the same way the changes in beliefs require a theory of belief revision. So, we modestly add the next postulate to the criteria of Cohen and Levesque:

- 8. Agents can retract their intentions when such intentions present problems.

Broadly speaking, this idea is the one that constitutes the core of intention revision.

2.3 What is intention revision? An example

Let us see, by way of an example, what is intention revision. Assume our agent is immerse in an environment that is inaccesible, non-determistic, episodic, discrete and dynamic [14]. Furthermore, suppose that the agent has certain beliefs and intentions (state α) and that, eventually, desires to achieve certain state of the world (state β) -we represent this situation with the black arrow in figure 1.
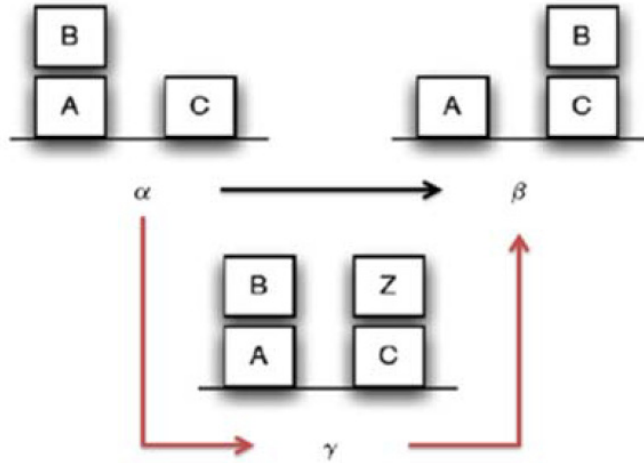

Fig. 1. States of the agent and the environment

In this way, the agent generates an intention of the form $put(B, C)$. Now, given the properties of the environment, let us suppose the agent perceives the state γ -which is denoted by the red arrow- where it is not the case that $free(C)$. Therefore, the intention will fail, the set of intentions would become inconsistent and the goals of the agent will not be achieved.

Let us see this situation in a more precise way. Suppose that we have an agent with a database of intentions -and beliefs- that includes the next data:

- p_1 $!put(x, y)$.
- p_2 $+!put(x, y) : -free(x)$.
- p_3 $+!put(x, y) : -free(y)$.
- p_4 $+!put(x, y) : -!move(x)$.

where $!\phi$ stands for an intention formula and $+\phi$ for an addition of a formula. If the database is equipped with some inference engine, the next formula is required to accomplish the intention:

- p_5 $free(x)$.

Now, suppose that it is the case that x is not free. This means that we have to add the negation of p_5 to the database. But then the set of intentions becomes inconsistent in an intuitive sense. If we want to keep the database consistent, which is a sound methodology, we need to revise the database. This implies that some of the intentions may have to be retracted; however, we do not need to revise the whole set of intentions for that would be an unnecessary lost of time and information. Thus, we have to choose what formulas -i.e., intentions- to retract.

The problem of intention revision, thus, is double: in first place, because it is intimately related to other mental states (like beliefs and desires); and in second place, because logic by itself is not sufficient to determine which intentions should be retracted. These problems lead us to take into account that the change of intentions is associated with changes in beliefs, and that we require extra-logical concepts to deal with these changes.

To complicate the setting even more, beliefs and intentions have certain logical consequences: when retracting intentions we have to choose which consequences (beliefs or intentions) we have to retract.

But to mantain the consistency and the maximum number of intentions accomplished, should we revise all the intentions? The answer is no, because, as we will see, the costs over time and memory would be huge.

2.4 Some methodological problems with intention revision

When dealing with intention revision some methodological problems appear: one related to representation, one related to inference, and finally, one related with a function of selection.

  – **The problem of representation.** How should the intentions be represented? Most databases work with facts and rules of some kind. The language used to represent intentions -together with beliefs and desires- may be related to some logical formalism (for instance, first order logic). This problem is, therefore, double: what language should we use to represent our data? And is this language adequate to relate the BDI components within a context of revision?
  – **The problem of the consequences.** What is the relation between the elements represented as facts and the elements that are inferred? This relation is sensible to the database. In some cases the elements that have been inferred have some special status in comparison with the facts; however, depending on which representation we use we will be able to distinguish these differences.
  – **The problem of the function of selection.** How should we choose which elements to retract? Logic by itself is not sufficient to decide which intentions should be maintained and which intentions should be retracted. We need a heuristic to determine this selection. One idea is that the loss of information should be minimal, for instance, by way of an ordering [7].

3 Models to represent intentional states

We will use a propositional model considering that the elements of the intentional system are propositional formulas. Of course, even with this representation we can have several alternatives. First, we have to pick an appropriate language (for instance, databases may be represented in a Prolog style). In this introduction we will work with a first order language.

We assume that the language L is closed for the operators $\neg$, $\wedge$, $\vee$, $\Rightarrow$ evaluated in a boolean way. We use $\phi, \psi, ...$ as propositional variables in L. The language L not only accepts what is explicitly represented in the database, but also, the consequences of it. Thus, another factor we have to determine is: which logical system should govern the set of intentions? In practice, the answer to this question depends on what mechanism of inference is coupled with the database; however, when doing this theoretical analysis, we will proceed by declaring the general functions of revision. So, for this introduction, we will use a classical propositional logic.

3.1 Sets of intentional states

The easiest way to represent an intentional state is by using well-formed formulas (wff) of L. According to this, we can define a set of intentional states (intentional set, from now on) through a set Σ of wff of L that satisfy the axiom of generalized reflexivity (C): if $\Sigma \vdash \phi$ then $\phi \in \Sigma$. The condition C assures us that Σ is closed under logical consequence. By the properties of classical logic, whenever Σ is inconsistent, then for all ϕ, $\Sigma \vdash \phi$. We will denote this with $\Sigma_\perp$. This means that there is an intentional set that is inconsistent.

There is a very close correspondence between intentional sets and possible worlds models. For any set W_Σ of possible worlds we can define a corresponding intentional set Σ as the set of those sentences that are true in all worlds in W_Σ. From a computational point of view, however, intentional sets are much more tractable than possible worlds models.

3.2 Intentional bases

Nevertheless, we have to consider that some intentions are not basic, but inferred. It is not possible to express this distinction through intentional sets, for the set theoretic representation does not provide markers or flags to indicate which intentions are basic and which are inferred. Moreover, it seems that when we make intentional changes we do not change the whole set of intentions, but a finite subset of it. Formally, this idea can be represented by letting B_Σ be a base for an intentional set Σ if and only if B_Σ is a finite subset of Σ and $Cn(B_\Sigma) = \Sigma$. Then, we introduce the functions for intention revision in bases of intentions (intentional bases from now on). The distinction between intentional set and intentional base allows us to generate and distinguish different structures, e.g., assume two intentional bases B_Σ and B'_Σ such that $Cn(B_\Sigma) = Cn(B'_\Sigma)$ but $B_\Sigma \neq B'_\Sigma$. If we want to implement intentional revision systems, intentional bases are easier to handle than intentional sets.

4 Postulates for intention revision

When dealing with intention revision there are two main strategies that may be followed: to present in an explicit manner the construction of the process of revision or to formulate the general ideas to realize such constructions. The first solution to the problem consists in developing algorithms that compute the functions; the second approach consists in describing the postulates to define the functions to further develop the algorithms.

In this introduction we will follow the second approach. The formulations of the postulates will be given through a series of ideas and conditions. The heuristic behind is similar to the idea of belief revision: the intentional changes should provide *(i)* a maximum of preservation of information (i.e., a minimum change in the intentions) and *(ii)* consistency.

Intention revision should occur when a new piece of information that is inconsistent with the database is added to the system in such a way that the resulting set in inconsistent. But this change is not the only one that may occur. Depending on how intentions are represented and what intentions are accepted, different intentional changes are possible. We can distinguish four intentional changes, three of them similar to belief changes:

– **Expansion.** A new formula ϕ is added to a Σ together with the logical consequences of the addition. The system that results from expanding Σ by a sentence ϕ will be denoted as $\Sigma \oplus \phi$.

- **Revision.** A new formula ϕ that is inconsistent with Σ is added, but in order to mantain consistency in the resulting system, some of the old formulas in Σ have to be deleted. This is denoted by $\Sigma \odot \phi$.
- **Contraction.** Some formula ϕ in Σ is retracted without adding any new facts. In order to mantain the system closed under logical consequence, some members of Σ must be deleted. This will be denoted by $\Sigma \ominus \phi$.
- **Reconsideration.** A new formula ϕ is added to Σ, but eventually such formula has be to contracted or revised. This is denoted by $\Sigma \otimes \phi$.

Expansions are closed under logical consequence (i.e., the expansion of the intentional set with a new formula is $\Sigma \oplus \phi = \{\psi | \Sigma \cup \phi \vdash \psi\}$); however, it is not possible to give a similar characterization of the other changes. The problem of revision, contraction and reconsideration has its roots in the lack of purely logical reasons to accomplish these processes. Thus, we can have different ways to research, specify and verify them.

For the time being, we will assume that the intentional sets model intentional bases. In what follows we will formulate some postulates for intention revision. The motivation behind these postulates (adapted from [1]) is that when we modify our intentions we have to keep to a minimum the change of intentions and we have to maintain consistency. For an agent the obtention of information implies costs and the environment in which is immersed is dynamic, for these reasons the unnecessary loss of information and time have to be avoided. On the other hand, we also require compromise, for the space in memory is not for free. This is an optimization heuristic; and although it is possible to give a quantitative definition of the loss of time or information, it is hard and impractical for our purposes. Instead, we will follow another specification: given that intentions are hierarchical plans [2], we believe that when retracting intentions we must retract the ones with a lesser hierarchy, and given that reconsideration reduces the time of revision, we believe we have to retract intentions on the basis of general rules [8]. In what follows, we will specify the postulates for intention revision considering these ideas.

4.1 Postulates for revision

For intention revision, the first postulate requires closure:

Postulate 1 *($\odot$1) For any formula ϕ and any intentional set Σ, $\Sigma \odot \phi$ is a intentional set.*

The second postulate guarantees that the input sentence is accepted in the revision:

Postulate 2 *($\odot$2) $\phi \in \Sigma \odot \phi$.*

A revision process should occur when the input ϕ contradicts what is already in Σ, that is $\neg\phi \in \Sigma$. However, in order to have the revision function defined for all inputs, we can easily extend it to cover the case when $\neg\phi \notin \Sigma$. Thus, revision is identified with expansion:

Postulate 3 *(⊙3) $\Sigma \odot \phi \subseteq \Sigma \oplus \phi$.*

Postulate 4 *(⊙4) If $\neg\phi \notin \Sigma$, then $\Sigma \oplus \phi \subseteq \Sigma \odot \phi$.*

The purpose of a revision is to produce a new consistent intentional set. Thus $\Sigma \odot \phi$ should be consistent, unless ϕ is logically impossible:

Postulate 5 *(⊙5) $\Sigma \odot \phi = K_\perp$ if and only if $\vdash \neg\phi$.*

We also require equivalence:

Postulate 6 *(⊙6) If $\vdash \phi \Leftrightarrow \psi$, then $\Sigma \odot \phi = \Sigma \odot \psi$.*

The postulates (⊙1) to (⊙6) are the basic postulates for revision. The final two conditions concern composite intention revisions. The idea is that, if $\Sigma \odot \phi$ is a revision of Σ and $\Sigma \odot \phi$ is to be changed by a further formula ψ; such change should be made by expansions of $\Sigma \odot \phi$ whenever possible. The minimal change of Σ to include both ϕ and ψ, that is, $\Sigma \odot \phi \wedge \psi$, ought to be the same as the expansion of $\Sigma \odot \phi$ by ψ, so long as ψ does not contradict the intentions in $\Sigma \odot \phi$:

Postulate 7 *(⊙7) $\Sigma \odot \phi \wedge \psi \subseteq (\Sigma \odot \phi) \oplus \psi$.*

Postulate 8 *(⊙8) If $\neg\psi \notin \Sigma \odot \phi$, then $(\Sigma \odot \phi) \oplus \psi \subseteq \Sigma \odot \phi \wedge \psi$.*

When $\neg\psi \notin \Sigma$, then $(\Sigma \odot \phi) \oplus \psi$ is $\Sigma_\perp$.

4.2 Postulates for contraction

We also need closure:

Postulate 9 *(⊖1) For any formula ϕ and any intentional set Σ, $\Sigma \ominus \phi$ is an intentional set.*

Because $\Sigma \ominus \phi$ is formed from Σ by giving up some intentions, no new intentions should appear:

Postulate 10 *(⊖2) $\Sigma \ominus \phi \subseteq \Sigma$.*

When $\phi \notin \Sigma$, the optimization heuristic requires that nothing has to be retracted:

Postulate 11 *(⊖3) If $\phi \notin \Sigma$, then $\Sigma \ominus \phi = \Sigma$.*

The formula to be contracted should not be a logical consequence of the intentions in $\Sigma \ominus \phi$:

Postulate 12 *(⊖4) If $\nvdash \phi$, then $\phi \notin \Sigma \ominus \phi$.*

From (⊖1) to (⊖4) it follows that if $\phi \notin \Sigma$, then $(\Sigma \ominus \phi) \oplus \phi \subseteq \Sigma$. In other words, if we first retract ϕ and then add ϕ again to the resulting intentional, no intentions are accepted that were not accepted in the original intentional set. The optimization heuristic demands that as many intentions as possible should be kept in $\Sigma \ominus \phi$. So, we need a recovery:

Postulate 13 *(⊖5) If $\phi \in \Sigma$, then $\Sigma \subseteq (\Sigma \ominus \phi) \oplus \phi$.*

This postulate enables us to undo contractions, and although it is controversial, we will assume it for sake of the introduction. The sixth postulate is analogous to $(\Sigma \odot 6)$:

Postulate 14 *(⊖6) If $\vdash \phi \Leftrightarrow \psi$, then $\Sigma \ominus \phi = \Sigma \ominus \psi$.*

These postulates are the basic set of postulates for intention contraction. Again, two further postulates for contractions with respect to conjunctions will be added. The motivations for these postulates are similar to $(\odot 7)$ and $(\odot 8)$.

Postulate 15 *(⊖7) $\Sigma \ominus \phi \cap \Sigma \ominus \psi \subseteq \Sigma \ominus \phi \wedge \psi$.*

Postulate 16 *(⊖8) If $\phi \notin \Sigma \ominus \phi \wedge \psi$, then $\Sigma \ominus \phi \wedge \psi \subseteq \Sigma \ominus \psi$.*

4.3 Postulates for reconsideration

The postulates above are adaptations of [1]. The following are different in some respects. The first postulate requires closure as well:

Postulate 17 *(⊗1) For any formula ϕ and any intentional set Σ, $\Sigma \otimes \phi$ is an intentional set.*

Reconsideration leads to revision [2] or contraction:

Postulate 18 *(⊗2) $(\Sigma \otimes \phi \subseteq \Sigma \odot \phi) \vee (\Sigma \otimes \phi \subseteq \Sigma \ominus \phi)$.*

The purpose of a reconsideration is to produce a new consistent intentional set:

Postulate 19 *(⊗3) $\Sigma \otimes \phi = K_\perp$ if and only if $\vdash \neg\phi$.*

We also require equivalence:

Postulate 20 *(⊗4) If $\vdash \phi \Leftrightarrow \psi$, then $\Sigma \otimes \phi = \Sigma \otimes \psi$.*

This set of postulates is the basic set of postulates for reconsideration. We also have the next ideas. The reconsideration $\Sigma \otimes \phi$ should be done by expansions whenever possible. And the minimal change of Σ to include ϕ and ψ should be the same as the expansion of $\Sigma \otimes \phi$ by ψ.

Postulate 21 *(⊗5) $\Sigma \otimes \phi \wedge \psi \subseteq (\Sigma \otimes \phi) \oplus \psi$.*

Postulate 22 *(⊗6) If $\neg\psi \notin \Sigma \otimes \phi$, then $(\Sigma \otimes \phi) \oplus \psi \subseteq \Sigma \otimes \phi \wedge \psi$.*

Postulate 23 *(⊗7) $\Sigma \otimes \phi \cap \Sigma \otimes \psi \subseteq \Sigma \otimes \phi \wedge \psi$.*

We now display some results regarding intention revision, but first, we require some definitions: an intention ϕ is abandoned if and only if ϕ is retracted from Σ either by a contraction or a revision. And an intention ϕ is continued if and only if $\phi \in (\Sigma \otimes \phi) \oplus \phi$.

The following results are straightforward.

Proposition 1 *The following statements hold:*

- *1. If ϕ is reconsidered, then ϕ is abandoned or continued.*

$$\Sigma \otimes \phi \Rightarrow (\Sigma \odot \phi \vee \Sigma \ominus \phi) \vee ((\Sigma \otimes \phi) \oplus \phi)$$

- *2. Inconsistency of reconsideration results from the inconsistency of intentions.*
$$\Sigma_\perp \Rightarrow \Sigma \otimes \phi = \Sigma_\perp$$

- *3. Reconsidering a consistent Σ with the current intentions does not remove any intention.*
$$(\Sigma \otimes \phi) \oplus \phi = \Sigma$$

- *4. Successful reconsideration produces an intentional set.*

$$\Sigma \otimes \phi = \Sigma$$

Proof. To prove statement 1, assume ϕ is reconsidered. By $\otimes 2$, it follows that $\Sigma \odot \phi \vee \Sigma \ominus \phi$, which means ϕ is abandoned; by addition, ϕ is continued. Statement 2 follows from $\otimes 3$. Statement 3 results from the definition of a continued intention. Statements 4 follows from $\otimes 1$. $\square$

At this point, we have presented some issues and problems of intention revision by isolating intentions from other mental states. In the next proposition we will try to relate intentions and beliefs through the reconsideration function. To see the next results, recall that it is irrational for an agent to intent ϕ and believe at the same time that it will not achieve ϕ: this is intention-belief inconsistency. To avoid this inconsistency, an agent must abandon intentions impossible to achieve. And recall that it is rational for an agent to intent ϕ and at the same time not believe that will achieve $\neg\phi$: this is intention-belief incompletness. To accomplish this property, an agent must continue its intentions.

The following representation theorems require some auxiliar definitions: we say an agent believes ϕ, $\mathsf{BEL}\phi$, if $\phi \in \Sigma$; and an agent has an intention to ϕ, $\mathsf{INTEND}\phi$, if $!\phi \in \Sigma$.

Proposition 2 *The following statements hold:*

- *1. Reconsideration implies intention-belief incompletness.*

$$\vdash \Sigma \otimes \phi \Rightarrow \mathsf{INTEND}\phi \wedge \neg\mathsf{BEL}\neg\phi$$

- *2. Reconsideration avoids intention-belief inconsistency.*

$$\nvdash \Sigma \otimes \phi \Rightarrow \mathsf{INTEND}\phi \wedge \mathsf{BEL}\neg\phi$$

Proof. Assume $\Sigma \otimes \phi$. Furthermore, assume that $\mathsf{INTEND}\phi$ is also given. We have two options: the intention is possible or impossible to achieve.
For statement 1: if the intention is possible to achieve after reconsideration, then $\phi \in (\Sigma \otimes \phi) \oplus \phi$, which means the agent can continue its intention. Thus, $\phi \in \Sigma$.
For statement 2: if the intention is impossible to achieve, then $\Sigma_\perp$, which means

that the reconsideration is inconsistent, but inconsistent reconsideration cannot be the case given ⊗5. □

Intuitively, this means that if an agent reconsiders, such agent is closer to rationality by following the intention-belief incompletness property, because the agent continues intentions possible to achieve. And also, the agent is far from irrationality by avoiding the intention-belief inconsistency, because after reconsideration the agent cannot have inconsistent reconsiderations, and so, the agent has to drop intentions not possible to achieve.

5   Conclusions

Let us sum up some of the main ideas and results of this introduction:

- A) Agents can retract their intentions when such intentions present problems.
- B) If an agent reconsiders an intention, such intention is abandoned or continued.
- C) If an agent reconsiders, such agent is closer to rationality by following the intention-belief incompletness property and by avoiding the intention-belief inconsistency.

So, we have sketched some general and introductory guidelines for intention revision by following the theories of intention and by considering that intentions are not isolated and are related to planning. We are aware these ideas lead to more problems. Some of the open problems we do not want to let to mention are the following:

- How do we relate the topic of this introduction with a non-monotonic logic? Since we can reconsider intentions, we have the possibility to relate our proposal with a non-monotonic consequence relation. Recall from section 2, for instance, that from a state α we want to achieve a state β through some execution of intentions, formally:

$$\frac{\alpha : p_1, \ldots, p_n}{\beta}$$

  but eventually happens that some plan p_i fails, which takes to intention revision. Future work requires the treatment of this situation.
- How do we relate the BDI components and temporal logic with the postulates we have proposed? One of the problems that we have presented is that, although we provide an abstract definition of the revision functions, we do not take into account the role of time within the reasoning process. Another problem is that we have considered intentions in an isolated way; this is necessary nonetheless, since intentions are irreducible componentes of the BDI architecture [2]; however, this is not sufficient. We have to relate the functions to other mental states through bridge rules of the form:

$$\frac{\frac{B_1, \ldots, B_n}{\alpha : p_1, \ldots, p_n}}{\beta}$$

that specify the change of states given certain beliefs (B_i) and intentions. But we also have to construct representation theorems, such as proposition 2, in order to relate different formalisms.

– What is the role of desires within this specification? The BDI architecture also requires desires. We know intentions are the desires the agents have committed to achieve. How does the change of desires affect the intentional changes?

– Which programming language can be adequate to model our proposal? Another problem we have is that our approach in this introduction is closer to an abstract logical specification, but it is far from implementation. Future work requires an integration of this proposal with an implementation.

The introduction we have presented here does not pretend to be exhaustive. On the contrary, we believe that the issues and problems we have showed are complex enough to be solved within the extensions of this work; but also, we believe they are clear enough to open a research program about intention revision.

References

1. Alchourrón, C. E., Gardenfors, P., Makinson, D.: On the logic of theory change: partial meet contraction and revision functions. Journal of Symbolic Logic, 50, 510-530 (1985).
2. Bratman, M.: Intention, Plans, and Practical Reason. Harvard University Press, Cambridge (1987).
3. Bratman, M. E., Israel, D. J., Pollack, M. E.: Plans and resource-bounded practical reasoning. Computational Intelligence, 4, 349-355 (1988).
4. Cohen, P., Levesque, H.: Intention is choice with commitment. Artificial Intelligence 42(3), 213-261 (1990).
5. Dignum, F., Meyer, J.-J. Ch., Wieringa, R. J., Kuiper, R.: A modal approach to intentions, commitments and obligations: Intention plus commitment yields obligation. In M. A. Brown, J. Carmo (eds.), Deontic logic, agency and normative systems, pp. 80-97, Springer-Verlag (1996).
6. Fagin, R., Ullman, J. D., Vardi, M. Y.: On the semantics of updates in databases. Proceedings of Second ACM SIGACT-SIGMOD, Atlanta, 352-365 (1983).
7. Gardenfors, P., Makinson. D.: Revisions of knowledge systems using epistemic entrenchment. In Proceedings of the Second Conference on Theoretical Aspects of Reasoning about Knowledge, M. Vardi (ed.), Los Altos, CA: Morgan Kaufmann (1988).
8. Guerra-Hernández A., Castro-Manzano, J.M., El-Fallah-Seghrouchni, A.: CTLAgentSpeak(L): a Specification Language for Agent Programs. Journal of Algorithms in Cognition, Informatics and Logic, (2009).
9. Harper, W. L.: Rational conceptual change. In PSA 1976, East Lansing, Mich: Philosophy of Science Association (1977).

14

10. Hoek, W. van der, Jamroga, W., Wooldridge, M.: Towards a theory of intention revision. Synthese, Springer-Verlag (2007).
11. Konolige, K., Pollack, M. E.: A representationalist theory of intentions. In Proceedings of International Joint Conference on Artificial Intelligence (IJCAI-93), 390-395, San Mateo: Morgan Kaufmann(1993).
12. Levi, I.: The Enterprise of Knowledge. MIT Press, Cambridge, Massachusetts(1980).
13. Rao, A.S., Georgeff, M.P.: Modelling Rational Agents within a BDI-Architecture. In Huhns, M.N., Singh, M.P., (eds.) Readings in Agents, pp. 317-328. Morgan Kaufmann (1998)
14. Russell, S. J., Norvig, P.: Artificial Intelligence. A modern approach. Prentice Hall, New Jersey, USA (1995).
15. Singh, M.P.: A critical examination of the Cohen-Levesque Theory of Intentions. In Proceedings of the European Conference on Artificial Intelligence (1992).
16. Wooldridge, M.: Introduction to Multiagent Systems. John Wiley and Sons,Ltd. (2001).