Intention Guided Belief Revision

Timothy William Cleaver
Griffith University
Southport, Australia
Tim.Cleaver@student.griffith.edu.au

Abdul Sattar
Griffith University
Nathan, Australia
a.sattar@griffith.edu.au

Abstract

This paper aims to investigate methodologies to utilize an agent's intentions as a means to guide the revision of its beliefs. For this purpose, we develop a collection of belief revision operators that employ the effect of the revision on the agent's intentions as the selection criteria. These operators are then assessed for rationality against the traditional AGM postulates. There is a large volume of work concerned with classical belief revision, the primary issue of which is the mitigation of the uncertainty inherent in environments in which belief revision is necessary. Traditional approaches attempt to assess the explanatory power of beliefs and utilize this as a heuristic to resolve this ambiguity. We argue that for practical reasoning systems, whose primary focus lies in the maintenance of behavior and not information, an agent's intentions provide a better guide.

Keywords: theories of agency and autonomy

Introduction

Intentional agents are those that form intentions to manage their behaviors. Intentions are mental attitudes that encapsulate a decision to adopt and a commitment to maintain a given behavior. Adopting intentions is beneficial for agents as intentions provide temporal stability of the agent's behaviors and reduce the complexity of decision making by enforcing consistency on its behaviors.

The problem of belief revision(Hansson 1999) is, given an initial belief base and some new datum to integrate into the belief base, how can an agent do this rationally? What makes this problem interesting is that, in general, the properties of the belief base and the input datum together do not determine a unique resulting belief base.

Given that intentions(and beliefs) are costly to adopt and maintain, an agent should consider the effect on its intentions when selecting among potential belief revisions. This differs from traditional approaches that attempt to minimize the change in the agent's beliefs themselves. Doing so is of particular importance to practical reasoning systems (Bratman 1987) where the accumulation and management of knowledge is secondary to the governance of behaviour.

We will consider three different classes of interaction between intention and belief:

1. Each intention that an agent holds is dependent on the agent holding a given belief. Should the agent cease to believe the dependent belief, then the intention must be dropped. (See section).
2. Each intention that an agent holds persists by default, contains a condition which when believed by the agent indicates that the intention has been successful and contains a condition which when believed indicates the failure or invalidity of the intention. (See section).
3. We impose additional structure on the intention base by assuming that there is a priority relation over the intentions. Intentions and beliefs are related as described above. (See section).

In each of these scenarios we will define an intentional agent and formalize the relationship between beliefs and intentions. Based on this relationship we will define a preference ordering over potential belief revisions based on their effects on the intentions of the agent. We will then assess whether this preference ordering is a valid doxastic preference relation. In doing so we assess the rationality of each as utilized within the AGM(Gärdenfors 1988) framework[1].

Maintenance Conditions on Intentions

Example 1. *Tom is an intentional agent and money is tight. He has a large mortgage and an investment account. He believes that if interest rates rise, then his savings will increase, but he will also have to pay more on his mortgage. His beliefs are as follows:*

$$\mathcal{B} = \left\{ \begin{array}{c} \neg\textit{'Savings increase'}, \neg\textit{'Mortgage increase'}, \\ \textit{'Interest rate rise'} \rightarrow \textit{'Savings increase'}, \\ \textit{'Interest rate rise'} \rightarrow \textit{'Mortgage increase'} \end{array} \right\}$$

Tom intends to pay only the minimum back on his mortgage necessary:

$$\mathcal{I} = \{[\textit{'Minimal repayments'} : \texttt{while}\,(\neg\textit{'Mortgage increase'})]\}$$

It is announced that interest rates have been increased. He must now integrate this into his beliefs

[1]Due to the foundational nature of the AGM framework we do not duplicate the fundamentals here, but, refer the reader to (Gärdenfors 1988)

($\mathcal{B} *$ *'Interest rate rise'*). *He has a number of ways of doing so:*

$$\mathcal{B}_1 = \left\{ \begin{array}{c} \text{'Interest rate rise'}, \neg\text{'Mortgage increase'}, \\ \text{'Interest rate rise'} \rightarrow \text{'Savings increase'} \end{array} \right\}$$

$$\mathcal{B}_2 = \left\{ \begin{array}{c} \text{'Interest rate rise'}, \neg\text{'Savings increase'}, \\ \text{'Interest rate rise'} \rightarrow \text{'Mortgage increase'} \end{array} \right\}$$

$$\mathcal{B}_3 = \left\{ \begin{array}{c} \text{'Interest rate rise'} \rightarrow \text{'Mortgage increase'}, \\ \text{'Interest rate rise'} \rightarrow \text{'Savings increase'}, \\ \text{'Interest rate rise'} \end{array} \right\}$$

$$\mathcal{B}_4 = \left\{ \begin{array}{c} \text{'Interest rate rise'}, \neg\text{'Savings increase'}, \\ \neg\text{'Mortgage increase'} \end{array} \right\}$$

We can see that $\mathcal{B}_1$ and $\mathcal{B}_4$ most support Tom's intentions. Thus, one, the other or some combination of both should constitute Tom's selected belief revision.

We now formalize the concept of an intentional agent, the mechanisms by which its intentions relate to its beliefs and a number of auxiliary concepts necessary for establishing the rationality of the preference relations we propose.

Definition 1. *A basic intentional agent is a tuple: $\langle \mathcal{L}_\mathcal{B}, \mathcal{B}, Cn, \mathcal{L}_\mathcal{I}, \mathcal{I} \rangle$ where:*

$\mathcal{L}_\mathcal{B}$ *is the language in which the agent's beliefs are expressed*[2].

$\mathcal{B}$ *is the belief base.*

Cn *is the consequence operator by which the agent derives new beliefs from old*[3].

$\mathcal{L}_\mathcal{I}$ *is the language in which the agent's intentions are expressed.*

$\mathcal{I}$ *are the intentions that agent holds.*

Definition 2. *An agent's intentions ($\mathcal{I}$) are grounded in (conditional with respect to) its beliefs ($\mathcal{B}$) if and only if:*

1. *they are of the form $[\iota : \texttt{while}(\beta)]$ where $\beta \in \mathcal{L}_\mathcal{B}$ and $\iota \in \mathcal{L}_\mathcal{I}$*
2. *for all $[\iota : \texttt{while}(\beta)] \in \mathcal{I}$, if $\mathcal{B}$ is the agent's current belief base and $\beta \in Cn(\mathcal{B})$, then $[\iota : \texttt{while}(\beta)] \in \mathcal{I}'$ where $\mathcal{I}'$ is the current set of intentions revised relative to the agent's new beliefs.*

Definition 3. *The set of intentions retained given a belief base ($\mathcal{B}$), an initial set of intentions ($\mathcal{I}$) and the consequence operator used to derive new beliefs from old (Cn) is the set ($\mathcal{I} \backslash \mathcal{B}$) where: $(\mathcal{I} \backslash \mathcal{B}) = \{[\iota : \texttt{while}(\beta)] \in \mathcal{I} | \beta \in Cn(\mathcal{B})\}$*

Definition 4. *If $\mathcal{B}$ and $\mathcal{B}'$ are belief bases and $\mathcal{I}$ is an intention base then $\left(\sqsubseteq_{\texttt{dif}(\mathcal{I})}\right)$ is a preference relation over belief bases given $\mathcal{I}$ if: $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ if and only if $(\mathcal{I} \backslash \mathcal{B}) \subseteq (\mathcal{I} \backslash \mathcal{B}')$*

Definition 5. *A selection function $\left(\gamma_{\texttt{dif}(\mathcal{I})}\right)$ is a* minimal intention change *selection function given a set of belief bases ($\mathbb{B}$) and an intention base ($\mathcal{I}$) based on a* minimal intention change *relation if and only if: $\gamma_{\texttt{dif}(\mathcal{I})}(\mathbb{B}, \mathcal{I}) = \{\mathcal{B} \in \mathbb{B} | \mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}$ for all $\mathcal{B}' \in \mathbb{B}\}$*

[2] We will assume this to be at least classical sentential logic

[3] We will assume that this includes classical truth functional consequence and demonstrates the properties of deduction and compactness

Definition 6. *A* minimal intention change *belief revision $\left(\mathcal{B} \pm_{\gamma_{\texttt{dif}(\mathcal{I})}} \beta\right)$ given an intention base ($\mathcal{I}$) is the external partial meet of the selected belief bases as selected via a* minimal intention change *selection function $\left(\gamma_{\texttt{dif}(\mathcal{I})}\right)$ where each option is an element of the belief remainder set ($\mathcal{B} \perp \neg\beta$)(Alchourrón & Makinson 1981) union the new belief (β):*

- $\mathcal{B} \pm_{\gamma_{\texttt{dif}(\mathcal{I})}} \beta = \left(\mathcal{B} +_{\gamma_{\texttt{dif}(\mathcal{I})}} \beta\right) \sim_{\gamma_{\texttt{dif}(\mathcal{I})}} \neg\beta$
- $\mathcal{B} +_{\gamma_{\texttt{dif}(\mathcal{I})}} \beta = \mathcal{B} \cup \beta$
- $\mathcal{B} \sim_{\gamma_{\texttt{dif}(\mathcal{I})}} \beta = \cap \gamma_{\texttt{dif}(\mathcal{I})}(\mathcal{B} \perp \beta, \mathcal{I})$

Notice that if the set of intentions is empty, then all belief revisions are equally preferred. If all revisions are equally preferred then the selection function will return all the revisions generated by the remainder set operator. Thus, when there are no intentions the agent will execute a full meet revision (with all the problems of over cautiousness that full meet entails).

Observation 1. *If the set of intentions is empty then all potential revisions are equally preferred by $\sqsubseteq_{\texttt{dif}(\mathcal{I})}$.*

Proof. We need to show that if $\mathcal{I} = \emptyset$ then for all belief bases $\mathcal{B}$ and $\mathcal{B}'$ $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}$. By definition 4 $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ if and only if $(\mathcal{I} \backslash \mathcal{B}) \subseteq (\mathcal{I} \backslash \mathcal{B}')$. If $\mathcal{I} = \emptyset$ then by definition 3 $(\mathcal{I} \backslash \mathcal{B}) = \emptyset$ for all belief bases. Since $\emptyset \subseteq \emptyset$ we have our proof. □

We must now verify that $\sqsubseteq_{\texttt{dif}(\mathcal{I})}$ provides a valid doxastic preference relation. To do so we must verify that it is transitive and weakly, if not completely maximizing.

Theorem 1. *$\sqsubseteq_{\texttt{dif}(\mathcal{I})}$ is transitive.*

Proof. We need to prove that if $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}''$, then $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}''$. By definition 4 if $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ then $(\mathcal{I} \backslash \mathcal{B}) \subseteq (\mathcal{I} \backslash \mathcal{B}')$ and similarly if $\mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}''$ then $(\mathcal{I} \backslash \mathcal{B}') \subseteq (\mathcal{I} \backslash \mathcal{B}''))$. Therefore, $(\mathcal{I} \backslash \mathcal{B}) \subseteq (\mathcal{I} \backslash \mathcal{B}'')$ and $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}''$ as required. □

Theorem 2. *$\sqsubseteq_{\texttt{dif}(\mathcal{I})}$ is weak maximizing.*

Proof. We need to prove that if $\mathcal{B} \subset \mathcal{B}'$, then $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$. If $\mathcal{B} \subset \mathcal{B}'$ then $Cn(\mathcal{B}) \subset Cn(\mathcal{B}')$ by the monotonicity of Cn. Thus, $\{[\iota : \texttt{while}(\beta)] \in \mathcal{I} | \beta \in Cn(\mathcal{B})\} \subseteq \{[\iota : \texttt{while}(\beta)] \in \mathcal{I} | \beta \in Cn(\mathcal{B}')\}$. By definitions 3 and 4 we have $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ as required. □

Unfortunately, we cannot strengthen this to show that $\sqsubseteq_{\texttt{dif}(\mathcal{I})}$ is maximizing.

Theorem 3. *$\sqsubseteq_{\texttt{dif}(\mathcal{I})}$ is not maximizing.*

Proof. We must define $\mathcal{B}$ and $\mathcal{B}'$ such that $\mathcal{B} \subset \mathcal{B}'$ but $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}$. Let us define the set of beliefs relevant for a given set of intentions as $\mathcal{B}_{rel} = \{\beta | [\iota : \texttt{while}(\beta)] \in \mathcal{I}\}$. Let $\mathcal{B}$ be any consistent subset of the relevant beliefs for the agent's current intentions $\mathcal{B} \in \left\{\beta \in 2^{\mathcal{B}_{rel}} | Cn(\beta) \neq \perp\right\}$, $\mathcal{B}'_{\overline{rel}}$ a belief that is not relevant to the agent's current intentions (nor has any as a consequence) $\mathcal{B}'_{\overline{rel}} \in \{\beta | \beta \in \mathcal{L}_\mathcal{B}$ and for all $\beta' \in \mathcal{B}_{rel}, \beta' \notin Cn(\beta)\}$ and let $\mathcal{B}'$

be the same subset of the relevant beliefs as $\mathcal{B}$ plus this element of the belief language that is not in the relevant belief set $\mathcal{B}' = \mathcal{B} \cup \mathcal{B}'_{\overline{rel}}$. We can see from the definitions of $\mathcal{B}$ and $\mathcal{B}'$ that $\mathcal{B} \subset \mathcal{B}'$ and the same number of intentions will be retained given either belief base by definition 3 ($(\mathcal{I} \backslash \mathcal{B}) = (\mathcal{I} \backslash \mathcal{B}')$). Thus, both $(\mathcal{I} \backslash \mathcal{B}) \subseteq (\mathcal{I} \backslash \mathcal{B}')$ and $(\mathcal{I} \backslash \mathcal{B}') \subseteq (\mathcal{I} \backslash \mathcal{B})$ and by definition 4 we have $\mathcal{B} \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\texttt{dif}(\mathcal{I})} \mathcal{B}$ as required. □

Theorem 4. *(Hansson 1999) $\sim_{\gamma_{\texttt{dif}(\mathcal{I})}}$ satisfies the properties of success, inclusion, relevance, uniformity, failure, closure, relative closure, core-retainment, vacuity, extensionality and conjunctive overlap by virtue of being a transitive weakly maximizingly relational partial meet contraction.*

Theorem 5. *(Hansson 1993) $\pm_{\gamma_{\texttt{dif}(\mathcal{I})}}$ satisfies the properties of consistency, inclusion, relevance, success, weak uniformity and pre-expansion by virtue of being an external partial meet revision.*

The relationship between intentions and beliefs as advocated in this section is not the traditional relationship between beliefs and intentions. It does not capture the stability inherent of intentions as each intention is only as stable as the belief it utilizes as its maintenance condition.

Success and Failure Conditions on Intentions

In this section we investigate the case where the beliefs and intentions of an agent are related via success and failure conditions. Intentions persist by default. Each intention has a success and failure condition associated with it. The success condition is a condition on the agent's beliefs that when satisfied indicates that the intention has been successful. At this point the agent can drop the intention and continue in the knowledge that the intention achieved its ends. Similarly, the failure condition of an intention indicates the conditions under which the intention has failed. Thus, the agent may take remedial steps, execute a recovery procedure or accept the failure of its intention as an outright loss and move on to other behaviors. Given that the beliefs of an agent, at a minimum, encapsulate its information regarding the environment, this relationship between intentions and beliefs ensures that the agent retains a realistic perspective on its progress towards its ends.

Example 2. *Tom is informed of another interest rate rise. His potential revisions are the same as those outlined in Example 1. Because of this, Tom decides it is time to save money. He will be successful if he believes his savings has increased and fail if his mortgage increases:*

$$\mathcal{I} = \left\{ \left[\text{'Save money'} : \begin{array}{l} \texttt{suc}(\text{'Savings increase'}), \\ \texttt{fail}(\text{'Mortgage increase'}) \end{array} \right] \right\}$$

If Tom is optimistic, then he will assume the success of his intention in light of uncertainty. Given that $\mathcal{B}_1$ allows Tom to conclude the success of his intention without also supporting its failure, Tom should accept $\mathcal{B}_1$. However, if Tom prefers to identify his potential mistakes early, he may prefer revisions $\mathcal{B}_2$ or $\mathcal{B}_3$. If Tom is cautious, he will attempt to minimize the effects on his intentions and prefer $\mathcal{B}_4$.

The above scenario motivates the following relationship between an agent's intentions and beliefs and its preferences towards the effect of belief revision on its intentions.

Definition 7. *An agent's intentions ($\mathcal{I}$) are grounded in (conditional with respect to) its beliefs ($\mathcal{B}$) if and only if:*

1. *they are of the form $[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')]$ where $\beta \in \mathcal{L}_\mathcal{B}$, $\beta' \in \mathcal{L}_\mathcal{B}$ and $\iota \in \mathcal{L}_\mathcal{I}$*
2. *for all $[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I}$ if $\beta \in Cn(\mathcal{B})$ or $\beta' \in Cn(\mathcal{B})$, then $\iota \notin \mathcal{I}'$ where $\mathcal{I}'$ is the intention base revised to account for the changes in the agent's beliefs.*

An agent can select the revision that:

- maximizes the number of successful intentions, then within those maximally successful intentions minimizes the number of failed intentions $(>_{\texttt{suc}}) > (<_{\texttt{fail}})$. If both the success and failure conditions are satisfied by the agent's beliefs then the intention is considered successful.
- minimizes the number of failed intentions prior to maximizing the successful intentions within the minimally failing intentions $(<_{\texttt{fail}}) > (>_{\texttt{suc}})$. If both the success and failure conditions are satisfied by the agent's beliefs then the intention is considered to have failed.
- minimizes the change required of its intentions in the hope that further information will become available so that it need not decide prematurely $(<_{\texttt{dif}})$

We need to define the sets of intentions rendered a success/failure by a given belief base, the duals of these sets, the set of beliefs relevant for a given intention base, preference relations over belief bases, selection functions based on these preference relations and belief revision operators in terms of these selection functions:

Definition 8. *A successful/failure intention remainder operator ($\backslash_{\overline{\texttt{suc}}}$) / ($\backslash_{\overline{\texttt{fail}}}$) is an infix operator that requires an intention base ($\mathcal{I}$) and a belief base ($\mathcal{B}$) as input and returns the elements of the intention base for which the success/failure condition do not follow from the belief base: $(\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}) = \{[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I} | \beta \notin Cn(\mathcal{B})\}$ and $(\mathcal{I} \backslash_{\overline{\texttt{fail}}} \mathcal{B}) = \{[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I} | \beta' \notin Cn(\mathcal{B})\}$. The dual of these operators are the successful/failure intention operators: $(\mathcal{I} \backslash_{\texttt{suc}} \mathcal{B}) = \{[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I} | \beta \in Cn(\mathcal{B})\}$ and $(\mathcal{I} \backslash_{\texttt{fail}} \mathcal{B}) = \{[\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I} | \beta' \in Cn(\mathcal{B})\}$*

Definition 9. *The conditions relevant to the success/failure of the agent's intentions $\mathcal{I}$: $\mathcal{B}_{\texttt{suc}(\mathcal{I})} = \{\beta | [\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I}\}$ and $\mathcal{B}_{\texttt{fail}(\mathcal{I})} = \{\beta' | [\iota : \texttt{suc}(\beta), \texttt{fail}(\beta')] \in \mathcal{I}\}$*

Definition 10. *$\mathcal{B} \sqsubseteq_{\texttt{suc}(\mathcal{I})} \mathcal{B}'$ if the set of successful intentions given $\mathcal{B}$ is a strict subset of those rendered a success by $\mathcal{B}'$, or, the set of successful intentions is the same for both belief bases and the set of failed intentions among those remaining after removing the successful ones for $\mathcal{B}$ is a superset of those for $\mathcal{B}'$.*

$$\mathcal{B} \sqsubseteq_{\texttt{suc}(\mathcal{I})} \mathcal{B}' = \begin{cases} \top & \text{if } (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}) \supset (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}'), \text{ or} \\ & \text{if } (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}) = (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}') \text{ and} \\ & (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}) \backslash_{\overline{\texttt{fail}}} \mathcal{B} \subseteq (\mathcal{I} \backslash_{\overline{\texttt{suc}}} \mathcal{B}') \backslash_{\overline{\texttt{fail}}} \\ & \mathcal{B}' \\ \bot & \text{otherwise} \end{cases}$$

and conversely:

Definition 11.

$$\mathcal{B} \sqsubseteq_{\mathtt{fail}(\mathcal{I})} \mathcal{B}' = \begin{cases} \top & \textit{if } (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}) \subset (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}'), \textit{ or} \\ & \textit{if } (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}') \textit{ and} \\ & (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B})\backslash_{\overline{\mathtt{suc}}}\mathcal{B} \supseteq (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}')\backslash_{\overline{\mathtt{suc}}} \\ & \mathcal{B}' \\ \bot & \textit{otherwise} \end{cases}$$

Definition 12. *$\mathcal{B} \sqsubseteq_{\mathtt{dif}(\mathcal{I})} \mathcal{B}'$ if both the set of successful and failed intentions for $\mathcal{B}'$ is a subset of those for $\mathcal{B}$.*

$$\mathcal{B} \sqsubseteq_{\mathtt{dif}(\mathcal{I})} \mathcal{B}' = \begin{cases} \top & \textit{if } (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}) \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}') \textit{ and} \\ & (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \\ \bot & \textit{otherwise} \end{cases}$$

Definition 13. *A selection function (γ_*) is an* intention change *selection function given a set of belief bases ($\mathbb{B}$) and an intention base ($\mathcal{I}$) based on an* intention change *relation ($\sqsubseteq_*$)*[4] *if and only if: $\gamma_*(\mathbb{B},\mathcal{I}) = \{\mathcal{B} \in \mathbb{B} | \mathcal{B}' \sqsubseteq_* \mathcal{B} \text{ for all } \mathcal{B}' \in \mathbb{B}\}$*

Definition 14. *An* intention change *belief revision ($\mathcal{B} \pm_* \beta$) given an intention base ($\mathcal{I}$) is the partial meet of the selected belief bases as selected via an* intention change *selection function ($*$) where each option is an element of the belief remainder set ($\mathcal{B}\bot\neg\beta$)(Alchourrón & Makinson 1981) union the new belief (β)*[5]*:*

- $\mathcal{B} \pm_* \beta = (\mathcal{B} +_* \beta) -_* \neg\beta$
- $\mathcal{B} +_* \beta = \mathcal{B} \cup \beta$
- $\mathcal{B} -_* \beta = \cap * (\mathcal{B}\bot\beta, \mathcal{I})$

Notice that whenever the agent does not have any intentions all belief revisions are equally preferred.

Observation 2. *If the set of intentions is empty then all belief bases are equally preferred by $\sqsubseteq_{\mathtt{suc}(\mathcal{I})}$, $\sqsubseteq_{\mathtt{fail}(\mathcal{I})}$ and $\sqsubseteq_{\mathtt{dif}(\mathcal{I})}$.*

Proof. We need to show that if $\mathcal{I} = \emptyset$ then for all belief bases $\mathcal{B}$ and $\mathcal{B}'$, $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$. By definition 10 there are two ways that $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ can hold and vice versa. Only one case need hold for the relation to hold. We will investigate each case individually:

Case 1: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ if $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$. Since $\mathcal{I} = \emptyset$ then $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = \emptyset$ for any $\mathcal{B}$. Since $\emptyset \not\supset \emptyset$ we know the first case does not hold for $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$. Through symmetry we also know that $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$ does not hold via $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})$.

Case 2: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ if $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B} \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}} \mathcal{B}'$. Since $\mathcal{I} = \emptyset$ then $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B} = \emptyset$ for any $\mathcal{B}$. Since $\emptyset \subseteq \emptyset$ then $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ by this case. Through symmetry we also have $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$ as required.

The proofs for $\sqsubseteq_{\mathtt{fail}(\mathcal{I})}$ and $\sqsubseteq_{\mathtt{dif}(\mathcal{I})}$ are analogous. □

Now we must demonstrate that the above relations are transitive and at least weakly, if not, completely maximizing.

[4]where $*$ is either $\mathtt{suc}\ (\mathcal{I})$, $\mathtt{fail}\ (\mathcal{I})$ or $\mathtt{dif}\ (\mathcal{I})$

[5]where $*$ is $\gamma_{\mathtt{suc}(\mathcal{I})}$, $\gamma_{\mathtt{fail}(\mathcal{I})}$ or $\gamma_{\mathtt{dif}(\mathcal{I})}$

Theorem 6. *$\sqsubseteq_{\mathtt{suc}(\mathcal{I})}$ is transitive.*

Proof. We need to show that if $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ then $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$. There are four cases to consider:

Case 1: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$. The transitivity of $\sqsubseteq_{\mathtt{suc}(\mathcal{I})}$ follows from the transitivity of $\supset$.

Case 2: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}' \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}''$. Given definition 10 $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}' \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')\backslash_{\overline{\mathtt{fail}}} \mathcal{B}''$ is only used in the case that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$(otherwise $\mathcal{B}'' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$). Since $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ and $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$ we have $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$ and $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ as required.

Case 3: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B} \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$. Given definition 10 $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B} \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}'$ is used only in the case that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$(otherwise $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$). Since $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ and $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$ we have $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')$ and $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ as required.

Case 4: $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B} \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}} \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}''$ by $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}' \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}'')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}''$. The transitivity of $\sqsubseteq_{\mathtt{suc}(\mathcal{I})}$ follows from the transitivity of $\subseteq$. □

Theorem 7. *$\sqsubseteq_{\mathtt{suc}(\mathcal{I})}$ is not weak maximizing.*

Proof. We need to show that if $\mathcal{B} \subset \mathcal{B}'$ then it need not be the case that $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$. To do so we need to construct a $\mathcal{B}$ and $\mathcal{B}'$ such that $\mathcal{B} \subset \mathcal{B}'$ and $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$. Since $\mathcal{B} \subset \mathcal{B}'$, $Cn(\mathcal{B}) \subset Cn(\mathcal{B}')$ by the monotonicity of Cn. Therefore, by definition 8 we have $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$. Since by definition 10, if we have $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) \supset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ then we have $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$. Therefore, we must construct $\mathcal{B}$ and $\mathcal{B}'$ so that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$ in order to prevent $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$ yet maintain $\mathcal{B} \subset \mathcal{B}'$. Since we require $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$, $\mathcal{B}$ and $\mathcal{B}'$ must be such that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}' \subset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B}$ by definition 10 (it cannot be $\subseteq$ lest it also satisfy the condition for $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}'$). This means that there must be at least one $\beta' \in \mathcal{B}'$ such that: $\beta' \notin \mathcal{B}$ and $[\iota : \mathtt{suc}(\beta), \mathtt{fail}(\beta')] \in \mathcal{I}$ where $\beta \notin \mathcal{B}$. Let us now construct $\mathcal{B}$ and $\mathcal{B}'$. Let $\mathcal{B}$ be any consistent strict subset of the beliefs relevant to the agent's intentions such that it does not imply the failure of any of the agent's intentions:

$$\mathcal{B} \in \left\{ \beta \in 2^{\mathcal{B}_{\mathtt{suc}(\mathcal{I})}} \setminus \mathcal{B}_{\mathtt{suc}(\mathcal{I})} \;\middle|\; \begin{array}{l} Cn(\beta) \neq \bot \text{ and} \\ Cn(\beta) \cap \mathcal{B}_{\mathtt{fail}(\mathcal{I})} = \emptyset \end{array} \right\}$$

Furthermore, let $\mathcal{B}'_{\mathtt{fail}} \in \mathcal{B}_{\mathtt{fail}(\mathcal{I})}$ and $\mathcal{B}' = \mathcal{B} \cup \mathcal{B}'_{\mathtt{fail}}$. We can see directly from the definitions that $\mathcal{B} \subset \mathcal{B}'$ so we must show that $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I})} \mathcal{B}$. Since by construction $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')$, we need to show that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\overline{\mathtt{fail}}}\mathcal{B}' \subset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\overline{\mathtt{fail}}}\mathcal{B}$. Because $\mathcal{B}' = \mathcal{B} \cup \mathcal{B}'_{\mathtt{fail}}$, $\mathcal{B}'_{\mathtt{fail}} \in \mathcal{B}_{\mathtt{fail}(\mathcal{I})}$ and $\mathcal{B} \cap \mathcal{B}_{\mathtt{fail}(\mathcal{I})} = \emptyset$ this also holds as required. □

Theorem 8. *$\sqsubseteq_{\mathtt{fail}(\mathcal{I})}$ is transitive but not weak maximizing.*

Proof. Proofs are analogous to Theorems 6 and 7. □

Theorem 9. $\sqsubseteq_{\mathtt{dif}(\mathcal{I})}$ *is transitive.*

Proof. The proof is analogous to Theorem 1. □

Theorem 10. $\sqsubseteq_{\mathtt{dif}(\mathcal{I})}$ *is not weak maximizing.*

Proof. We need to define $\mathcal{B}$ and $\mathcal{B}'$ such that $\mathcal{B} \subset \mathcal{B}'$ yet $\mathcal{B}' \sqsubseteq_{\mathtt{dif}(\mathcal{I})} \mathcal{B}$. If $\mathcal{B} \subset \mathcal{B}'$ then $Cn(\mathcal{B}) \subset Cn(\mathcal{B}')$ by the monotonicity of Cn. By definitions 12 and 8 we can see that we need to construct $\mathcal{B}$ and $\mathcal{B}'$ such that $\mathcal{B}'$ results in the success or failure of more intentions than $\mathcal{B}$. Thus, let $\mathcal{B} = \emptyset$ and $\mathcal{B}'$ be a subset of the beliefs relevant to the success of the agent's intentions such that this set is consistent and none of the failure conditions of the agent's intentions can be derived from it. Let $\mathcal{B}' \in \left\{\beta \in 2^{\mathcal{B}_{\mathrm{suc}(\mathcal{I})}} \,\middle|\, Cn(\beta) \neq \bot \text{ and } Cn(\beta) \cap \mathcal{B}_{\mathtt{fail}(\mathcal{I})} = \emptyset\right\}$. We can see from the construction that $\mathcal{B} \subset \mathcal{B}'$. Now we must show that these definitions lead us to $\mathcal{B}' \sqsubseteq_{\mathtt{dif}(\mathcal{I})} \mathcal{B}$. From definition 8 we must show $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})$ and $(\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}') \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B})$. Since $\mathcal{B} = \emptyset$ we know that $(\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}) = \mathcal{B}_{\mathtt{fail}(\mathcal{I})}$ and $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}) = \mathcal{B}_{\mathtt{suc}(\mathcal{I})}$. We can see from the construction that $(\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}') = \mathcal{B}_{\mathtt{fail}(\mathcal{I})}$ since $\mathcal{B}'$ contains no beliefs relevant to the failure conditions of the agent's intentions. However, since $\mathcal{B}'$ does contain elements relevant to the success conditions of the agent's intentions we know that $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \subset (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})$. Thus, we have $(\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}') \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})$ and $(\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}') \subseteq (\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B})$ as required.

□

Theorem 11. *(Hansson 1999)* $\sim_{\gamma_{\mathrm{suc}(\mathcal{I})}}$, $\sim_{\gamma_{\mathtt{fail}(\mathcal{I})}}$ *and* $\sim_{\gamma_{\mathtt{dif}(\mathcal{I})}}$ *satisfy the properties of success, inclusion, relevance, uniformity, failure, closure, relative closure, core-retainment, vacuity and extensionality by virtue of being transitive relational partial meet contractions.*

Theorem 12. *(Hansson 1993)* $\pm_{\gamma_{\mathrm{suc}(\mathcal{I})}}$, $\pm_{\gamma_{\mathtt{fail}(\mathcal{I})}}$ *and* $\pm_{\gamma_{\mathtt{dif}(\mathcal{I})}}$ *satisfy the properties of consistency, inclusion, relevance, success, weak uniformity and pre-expansion by virtue of being external partial meet revisions.*

Prioritized Intentions

Although the above formalization captures the relationship between intentions and beliefs in a more flexible way, it fails to address the relationships between intentions themselves. Typically, an intention base is not simply just a set of intentions but has more structure. We will capture this additional structure through the use of a preference relation over intentions.

Example 3. *Tom and his wife have decided to have a baby. Since it will be their first baby they are eligible for a government subsidy called the 'baby bonus' depending on their current economic situation. Tom's beliefs are as in Example 1. As part of intending to have a baby, Tom also intends to gain the 'baby bonus'. Since the 'baby bonus' is worth much more than his investments will earn, he prefers to get the 'baby bonus' over a savings increase. However, should his investments increase in value then that will elevate him over the threshold making him ineligible to receive it:*

$$\mathcal{I} = \left\{ \begin{array}{c} \left[\textit{'Save money'}: \begin{array}{l} \mathtt{suc}(\textit{'Savings increase'}), \\ \mathtt{fail}(\textit{'Mortgage increase'}) \end{array} \right], \\ \left[\textit{'Baby bonus'}: \begin{array}{l} \mathtt{suc}(\textit{'Have baby'}), \\ \mathtt{fail}(\textit{'Savings increase'}) \end{array} \right] \\ \textit{'Save money'} < \textit{'Baby bonus'} \end{array} \right\}$$

Tom is told of another interest rate rise. The effect on his beliefs is as in Example 1. We can see that both $\mathcal{B}_1$ *and* $\mathcal{B}_3$ *are not ideal revisions as they will require Tom to drop his intention of gaining the 'baby bonus'.* $\mathcal{B}_2$ *and* $\mathcal{B}_3$*, if adopted, would result in Tom's having to drop his intention to 'save money'. Thus, the best revision considering Tom's intentions is* $\mathcal{B}_4$*.*

We retain definitions 8, 9, 13 and 14 and the constraints imposed by definition 7 from section but must modify definitions 10, 11 from section to integrate the preference relation over intentions.

Definition 15. *A preferential intentional agent is a tuple:* $\langle \mathcal{L}_\mathcal{B}, \mathcal{B}, Cn, \mathcal{L}_\mathcal{I}, \mathcal{I}, \trianglelefteq, F_n \rangle$ *where:*

$\mathcal{L}_\mathcal{B}$, $\mathcal{B}$, Cn, $\mathcal{L}_\mathcal{I}$ **and** $\mathcal{I}$ *are as Definition 1.*

$\trianglelefteq$ *is a transitive and complete preference order*[6] *over the agent's intentions.*

F_n *is a function that given an intention base and an ordering over intentions produces an ordering*[7] *over all subsets of the intention base.*

Definition 16. $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)} \mathcal{B}'$ *if the set of intentions that* $\mathcal{B}'$ *renders successful is more preferred than the set of intentions rendered successful by* $\mathcal{B}$ *or, the set of intentions rendered successful by* $\mathcal{B}$ *and* $\mathcal{B}'$ *are equally preferred and the set of intentions deemed a failure by* $\mathcal{B}'$ *(disregarding those intentions that have been rendered successful) is less preferable to the set of intentions deemed a failure by* $\mathcal{B}$ *(again disregarding those intentions that had been rendered successful).*

$$\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)} \mathcal{B}' = \begin{cases} \top & \textit{if} \quad (\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B})\, F_n\, (\triangleleft, \mathcal{I})\, (\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B}'), \\ & \textit{or if } ((\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B}')\backslash_{\mathtt{fail}}\mathcal{B}')\ F_n\, (\trianglelefteq, \mathcal{I}) \\ & ((\mathcal{I}\backslash_{\overline{\mathtt{suc}}}\mathcal{B})\backslash_{\mathtt{fail}}\mathcal{B}) \\ \bot & \textit{otherwise} \end{cases}$$

Definition 17.

$$\mathcal{B} \sqsubseteq_{\mathtt{fail}(\mathcal{I},\trianglelefteq)} \mathcal{B}' = \begin{cases} \top & \textit{if } (\mathcal{I}\backslash_{\mathtt{fail}}\mathcal{B}')\, F_n\, (\triangleleft, \mathcal{I})\, (\mathcal{I}\backslash_{\mathtt{fail}}\mathcal{B}), \\ & \textit{or if } ((\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B})\backslash_{\mathtt{suc}}\mathcal{B})\ F_n\, (\trianglelefteq, \mathcal{I}) \\ & ((\mathcal{I}\backslash_{\overline{\mathtt{fail}}}\mathcal{B}')\backslash_{\mathtt{suc}}\mathcal{B}') \\ \bot & \textit{otherwise} \end{cases}$$

Notice, again, that if the intention base is empty then all belief revisions are equally preferred.

Observation 3. *If the set of intentions is empty then all belief bases are equally preferred by* $\sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)}$ *and* $\sqsubseteq_{\mathtt{fail}(\mathcal{I},\trianglelefteq)}$*.*

[6] This preference order is: reflexive, transitive and antisymmetric

[7] The resulting generalized order is also: reflexive, transitive and antisymmetric

Proof. The proof for $\sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)}$ and $\sqsubseteq_{\mathtt{fail}(\mathcal{I},\trianglelefteq)}$ are analagous to Observation 2. □

Theorem 13. $\sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)}$ *is transitive.*

Proof. The proof is analogous to Theorem 6 □

Theorem 14. $\sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)}$ *is not weak maximizing.*

Proof. To see this we need to construct a scenario where $\mathcal{B} \subset \mathcal{B}'$ but it is not the case that $\mathcal{B} \sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)} \mathcal{B}'$. We need to construct $\mathcal{B}$, $\mathcal{B}'$ and $F_n(\trianglelefteq,\mathcal{I})$ such that $\mathcal{B} \subset \mathcal{B}'$ and $(\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B}')\, F_n(\trianglelefteq,\mathcal{I})\, (\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B})$ to show that $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)} \mathcal{B}$. Let

$$\mathcal{B} \in \left\{ \beta \in 2^{\mathcal{B}_{\mathtt{suc}(\mathcal{I})}} \setminus \mathcal{B}_{\mathtt{suc}(\mathcal{I})} \,\middle|\, \begin{array}{l} Cn(\beta) \neq \bot \text{ and} \\ Cn(\beta) \cap \mathcal{B}_{\mathtt{fail}(\mathcal{I})} = \emptyset \end{array} \right\}$$

$$\mathcal{B}'_{\mathtt{suc}} \in \left\{ \beta \in \mathcal{B}_{\mathtt{suc}(\mathcal{I})} \setminus \mathcal{B} \,\middle|\, \begin{array}{l} Cn(\beta \cup \mathcal{B}) \neq \bot \text{ and} \\ Cn(\beta \cup \mathcal{B}) \cap \mathcal{B}_{\mathtt{fail}(\mathcal{I})} = \emptyset \end{array} \right\}$$

$\mathcal{B}' = \mathcal{B} \cup \mathcal{B}'_{\mathtt{suc}}$ and

$$\mathcal{I}' F_n(\trianglelefteq,\mathcal{I})\, \mathcal{I}'' = \begin{cases} \top & \text{if } \mathcal{I}' \supset \mathcal{I}'', \text{ or if there exists an} \\ & \iota \in \mathcal{I} \text{ such that } \iota \notin \mathcal{I}'' \text{ and } \iota \trianglelefteq \iota' \\ & \text{for all } \iota' \in \mathcal{I}'' \text{ such that } \iota' \notin \mathcal{I}', \\ & \text{or if } \mathcal{I}' = \mathcal{I}'' \\ \bot & \text{otherwise} \end{cases}$$

We can see that $F_n(\trianglelefteq,\mathcal{I})$ is a transitive and complete generalized preference relation over $2^{\mathcal{I}}$. It is clear from the construction of $\mathcal{B}$ and $\mathcal{B}'$ that $\mathcal{B} \subset \mathcal{B}'$. Furthermore we know from the construction that $\mathcal{B}'$ renders at least one intention successful beyond the intentions that $\mathcal{B}$ renders successful. Since $(\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B}) \subset (\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B}')$ we have $(\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B}')\, F_n(\triangleleft,\mathcal{I})\, (\mathcal{I}\backslash_{\mathtt{suc}}\mathcal{B})$ by the first clause of definition 16. Thus, we have $\mathcal{B}' \sqsubseteq_{\mathtt{suc}(\mathcal{I},\trianglelefteq)} \mathcal{B}$ as required. □

Theorem 15. $\sqsubseteq_{\mathtt{fail}(\mathcal{I},\trianglelefteq)}$ *is transitive but not weak maximizing.*

Proof. Proofs are analogous to Theorems 13 and 14. □

Theorem 16. *(Hansson 1999)* $\sim_{\gamma_{\mathtt{suc}(\mathcal{I})}}$ *and* $\sim_{\gamma_{\mathtt{fail}(\mathcal{I})}}$ *satisfy the properties of success, inclusion, relevance, uniformity, failure, closure, relative closure, core-retainment, vacuity and extensionality by virtue of being a transitive weakly maximizingly relational partial meet contraction.*

Theorem 17. *(Hansson 1993)* $\pm_{\gamma_{\mathtt{suc}(\mathcal{I})}}$ *and* $\pm_{\gamma_{\mathtt{fail}(\mathcal{I})}}$ *satisfy the properties of consistency, inclusion, relevance, success, weak uniformity and pre-expansion by virtue of being an external partial meet revision.*

Conclusion

(Georgeff & Rao 1995) investigated the effect of belief revision on intentions and intention maintenance. They did so in the context of BDI_{CTL*} logic(Rao & Georgeff 1998; Wooldridge 2000). In order to bootstrap their analysis they presupposed a belief revision function that satisfied the AGM postulates without investigating the details of such a revision. They then analyzed the effect of this revision function on the maintenance of intentions.

In this paper we have investigated the role intention can play in the revision of an intelligent agent's beliefs. To do so we defined a number of ways an agent could relate its beliefs to its intentions. Given these relationships we defined strategies which an agent could utilize to guide its belief revisions. For each strategy we defined a preference relation over the potential revisions an agent could adopt. We then assessed the rationality of the resulting contraction and revision operator according to the traditional AGM methodology.

We showed that when intentions depend on the maintenance of belief conditions, the change minimizing belief change operators proposed satisfied all the basic AGM postulates(see Theorems 4 and 5) in addition to conjunctive overlap which ensures rationality when contracting by conjunctions. We showed that when intentions persist by default but are dropped when success or failure conditions on the intention are satisfied that the three proposed belief contractions and corresponding revisions satisfied all the basic AGM postulates(see Theorems 11 and 12). Finally, we showed that introducing priorities on intentions and redefining the belief change operators to respect these priorities did not affect these results (see Theorems 16 and 17).

Acknowledgements

We would like to hereby thank the Smart Internet Cooperative Research Centre (SIT-CRC) for their funding of this work.

References

Alchourrón, C., and Makinson, D. 1981. Hierarchies of regulation and their logic. *New Studies in Deontic Logic* 125–148.

Bratman, M. 1987. *Intention, Plans and Practical Reason*. Harvard University Press.

Gärdenfors, P. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Bradford Books, The MIT Press.

Georgeff, M., and Rao, A. 1995. The semantics of intention maintenance for rational agents. In Mellish, C., ed., *Proceedings of the Fourteenth International Joint Conference on Artificial Intelligence*, 704–710.

Hansson, S. 1993. Reversing the Levi indentity. *Journal of Philosophical Logic* 22:637–669.

Hansson, S. 1999. *A Textbook of Belief Dynamics: Theory Change and Database Updating*. Kluwer Academic Publishers, first edition.

Rao, A., and Georgeff, M. 1998. Decision procedures for bdi logic. *Journal of Logic and Computation* 8(3).

Wooldridge, M. 2000. *Reasoning About Rational Agents*. MIT Press.