Revising Beliefs and Intentions in Stochastic Environments

Nima Motamed[1] , **Natasha Alechina**[2,1] , **Mehdi Dastani**[1] and **Dragan Doder**[1]
[1]Utrecht University, The Netherlands
[2]Open University, The Netherlands
{n.motamed, n.a.alechina, m.m.dastani, d.doder}@uu.nl

Abstract

The development of autonomous agents operating in dynamic and stochastic environments requires theories and models of how beliefs and intentions are revised while taking their interplay into account. In this paper, we initiate the study of belief and intention revision in stochastic environments, where an agent's beliefs and intentions are specified in a decidable probabilistic temporal logic. We then provide general Katsuno & Mendelzon-style representation theorems for both belief and intention revision, giving clear semantic characterizations of revision methods.

1 Introduction

Autonomous agents operating in dynamic environments need a method for keeping track of their beliefs, updating these as new information is received. But an agent also carries *intentions* they are committed to bring about, and these also need to be updated on the basis of new information or deliberation. While there is an enormous amount of literature on the former topic of *belief revision*, the topic of *intention revision* received, in comparison, less attention [Cohen and Levesque, 1990; Wooldridge, 2000; van der Hoek *et al.*, 2007; Castelfranchi and Paglieri, 2007; Lorini and Herzig, 2008; Shoham, 2009; Grant *et al.*, 2010; Icard *et al.*, 2010; van Ditmarsch *et al.*, 2011; Shapiro *et al.*, 2012; van Zee and Doder, 2016; van Zee *et al.*, 2020].

As observed by many of the aforementioned authors, intention and beliefs are deeply intertwined, and their revisions need to be studied together. For instance, for an agent to be able to intend something, it does not suffice if it is merely the case that their beliefs and the intention are separately consistent on their own: they need to *cohere* together, in the sense that the agent's beliefs should be consistent with the achievement of the intention [van Zee *et al.*, 2020].

In prior work on intention revision, such interactions between belief and intention have only been studied for agents acting in deterministic environments. But for many practical applications, agents act in *stochastic* environments, where their actions have uncertain outcomes.

In this paper, we present the *first* framework for the joint revision of beliefs and intentions in stochastic environments. We define beliefs and intentions in a decidable probabilistic temporal logic interpreted on Markov Decision Processes (MDPs), giving the ability to express beliefs about the uncertainty in actions' outcomes, as well as complex temporal intentions. We propose a set of rationality postulates for revision operators in this logic. Using these rationality postulates, we also obtain semantic understanding of revision through Katsuno & Mendelzon-style representation theorems [Katsuno and Mendelzon, 1991].

Working with stochastic environments and highly expressive beliefs and intentions about probability and time brings new challenges compared to prior work in the field, both technical and conceptual. In terms of technical challenges, giving a representation theorem is complicated by the fact that in general, there are infinitely many MDPs satisfying certain beliefs, due to the presence of probability, and that we have no general guarantees on the expressibility of sets of MDPs. These features are vital to the original theorems of [Katsuno and Mendelzon, 1991]. To overcome this challenge, we employ results and methods from [Falakh *et al.*, 2023] to still obtain representation theorems. But to make our work more practically applicable, we go further by defining two novel postulates which additionally provide representation theorems for a specific class of operators that are conceptually simple and generally *computable*, properties which are not guaranteed by the results of [Falakh *et al.*, 2023].

In terms of conceptual challenges, stochastic environments also require more nuanced approaches to the interplay between beliefs and intentions than is considered in prior work. Let us illustrate this with the following example (which we will frequently revisit throughout this paper).

Consider an agent named Theo, who intends to attend an obscure, underground concert tomorrow. He only has \$5, and admission costs precisely \$5. He believes that if he goes to the ticket office today, he can almost certainly buy a ticket since no one knows the band and they are therefore unlikely to be sold out. Theo is also aware that the concert venue sometimes sends out free tickets for concerts in the mail, as a promotion. So Theo also believes that if he stays at home, he may receive tickets, albeit with a small probability.

Now, for simplicity assume that Theo's favorite brand of beer costs \$5, both at shops and at the concert venue. How would Theo revise his intentions with the new intention to also buy a drink tomorrow? His beliefs do in one sense support the possibility of achieving both intentions, since he could luck out and get free tickets, attend the concert, and buy the drink there. So perhaps the revision should just be adding the new intention on top of his original one. But should rational agents make decisions based on unlikely outcomes? It would be more reasonable to say that an agent's intentions are coherent only if they can achieve them with reasonably high probability, with the definition of being 'reasonably high' as a choice to be made. So it may be more reasonable that when revising with this new intention, Theo drops his old intention of attending the concert.

To see how the nuance of probabilities also influences the interplay between beliefs and intentions, note that accepting intentions should revise the agent's beliefs: if he intends to attend the concert, Theo should also believe he no longer has \$5 after the concert, as the only way to be reasonably sure about his attendance is to spend it on tickets.

Similarly, revising beliefs should also be able to trigger intention revision. Say Theo gets a call from his friend Stan, who tells him that the band, which Theo thought to be obscure, has suddenly become an international sensation overnight. Theo revises his beliefs about him being able to buy tickets at the ticket office, as they are now highly likely to be sold out. And since it so unlikely to get tickets, no matter what he does, it may be reasonable to drop his intention of attending the concert.

2 PLBP+

An agent's beliefs and intentions naturally concern themselves with probability ("I believe speeding is *likely* to be dangerous") and time ("I intend to apply for jobs *until* I succeed"). Therefore we choose to express beliefs and intentions in an appropriate probabilistic temporal logic. We choose to work with an extended version of the Probabilistic Logic of Bounded Policies (PLBP), introduced by [Motamed *et al.*, 2023]. We refer to this extension as PLBP+. We now explain the syntax and semantics of the logic.

We fix a countably infinite set Prop of *propositional variables* and a finite set $\mathscr{A}$ of *actions* throughout the paper. To each action $a \in \mathscr{A}$ we associate a *precondition* pre_a, which is a conjunction of literals over Prop, and a finite nonempty list Post_a of possible *postconditions*, which are also conjunctions of literals. Intuitively, the precondition of an action is precisely what must be satisfied so that the action can be executed, and a postcondition of an action is one of its possible outcomes, i.e. that which is made true after executing the action. We refer to the ith postcondition of a as $\mathsf{post}_{a,i}$.

Before we can define the stochastic environments that serve as models of the logic, some notation is required. We write $f \colon X \rightharpoonup Y$ to denote that f is a *partial* function from X to Y, and write $f(x){\downarrow}$ to denote that f is defined on input $x \in X$. Given a set X, some $x \in X$ and $n \geqslant 0$, we write X^n_x for the set of all sequences in X of length n starting with x (i.e. with $X^0_x = \emptyset$ and $X^1_x = \{x\}$), and we write $X^{\leqslant n}_x = \bigcup_{k=0}^n X^k_x$. Finally, we write $\Delta(X)$ for the set of all finitely supported (discrete) probability distributions on X, i.e. those probability distributions $D : X \to [0, 1]$ such that $D(x) > 0$ for only finitely many x.

Definition 1 (MDP)**.** A *Markov decision process* (MDP) over $\mathscr{A}$ is a tuple $\mathbb{M} = (S, P, V)$, where S is the set of *states*, $P \colon S \times \mathscr{A} \rightharpoonup \Delta(S)$ is the *partial probabilistic transition function*, and $V \colon S \to 2^{\mathsf{Prop}}$ is the *valuation*. We often abbreviate $P(s, a)$ by $P_{s,a}$. These are required to satisfy the following conditions. First, for all $s \in S$, there is some $a \in \mathscr{A}$ such that $s \Vdash \mathsf{pre}_a$, where $\Vdash$ here is the standard propositional satisfaction relation. Second, $P_{s,a}{\downarrow}$ iff $s \Vdash \mathsf{pre}_a$. And third, given $P_{s,a}{\downarrow}$, we have that (i) for all $t \in S$ such that $P_{s,a}(t) > 0$, there is a unique $\mathsf{post}_{a,i}$ in Post_a such that $t \Vdash \mathsf{post}_{a,i}$, and (ii) for all $\mathsf{post}_{a,i}$ in Post_a there is at most one t such that $P_{s,a}(t) > 0$ and $t \Vdash \mathsf{post}_{a,i}$.

A *pointed MDP* (pMDP) is a pair $\mathbb{P} = (\mathbb{M}, s)$ where $\mathbb{M}$ is an MDP, and s is a state in $\mathbb{M}$. We denote the set of all pMDPs $\mathbb{P}$ by pMDP.

A key feature of PLBP+ is the ability to reason about *n-step policies*, which specify how the agent will act for n time steps.

Definition 2 (n-step policies)**.** Given a pMDP $\mathbb{P} = (\mathbb{M}, s)$ and $n \geqslant 0$, an *n-step $\mathbb{P}$-policy* is a function $\pi \colon S^{\leqslant n}_s \to \mathscr{A}$ such that $s_k \Vdash \mathsf{pre}_{\pi(s_1\cdots s_k)}$ for all $s_1 \cdots s_k \in S^{\leqslant n}_s$.

Given n-step π with $n \geqslant 1$ and $t \in S$, define the *pushforward of π to t* to be the $(n-1)$-step $(\mathbb{M}, t)$-policy π_t given by putting $\pi_t(\mathbf{t}) = \pi(s\mathbf{t})$.

We will drop $\mathbb{P}$ when the pMDP is clear from context, speaking just of n-step policies. Or conversely, when the n is not relevant within some context, we just speak of $\mathbb{P}$-policies.

Definition 3 (Paths and path distributions)**.** For an n-step policy π, we define $\mathsf{Paths}(\pi)$ as the set of all *paths* $s_0 \cdots s_n \in S^{n+1}$ such that $s_0 = s$, and $P(s_k, \pi(s_0 \cdots s_k))(s_{k+1}) > 0$ for all $0 \leqslant k < n$.

Given an n-step policy π, its *path distribution* is the probability distribution $\mu_\pi \in \Delta(\mathsf{Paths}(\pi))$ defined as $\mu_\pi(s_0 \cdots s_n) = \prod_{0 \leqslant k < n} P(s_k, \pi(s_0 \cdots s_k))(s_{k+1})$. This extends to sets of paths in the standard way, i.e. $\mu_\pi(X) = \sum_{\mathbf{w} \in X} \mu_\pi(\mathbf{w})$.

Definition 4 (Syntax)**.** The *syntax of PLBP+* is inductively defined by the following grammar:

$$\begin{aligned}
\varphi &::= p \mid \varphi \wedge \varphi \mid \neg\varphi \mid \Diamond^n \alpha,\\
\alpha &::= \mathsf{do}_a \mid \mathsf{Pr}_{\bowtie r}\Phi \mid \alpha \wedge \alpha \mid \neg\alpha\\
\Phi &::= \varphi \mid \alpha \mid \Phi \wedge \Phi \mid \neg\Phi \mid \mathsf{X}\Phi,
\end{aligned}$$

where $p \in \mathsf{Prop}$, $n \geqslant 0$, $a \in \mathscr{A}$, $r \in \mathbb{Q} \cap [0, 1]$, and $\bowtie \in \{\geqslant, \leqslant\}$. Formulas φ are referred to as *state formulas*, formulas α are referred to as *policy formulas*, and formulas Φ are referred to as *path formulas*.

State formulas are interpreted over states, policy formulas over policies, and path formulas over policies and paths. The state formula $\Diamond^n \alpha$ should be read as "there exists an n-step policy for which α holds". The policy formula do_a is read as "the current policy chooses action a now", and $\mathsf{Pr}_{\bowtie r}\Phi$ as "under the current policy, Φ holds with probability $\bowtie r$". The path formula $\mathsf{X}\Phi$ is read as "at the next time on this path and policy, Φ holds". Note that [Motamed *et al.*, 2023] introduce syntactic restrictions in formulas so that the amount of nested X-operators never exceeds the length of a path on which it is interpreted. We allow arbitrary path formulas to be interpreted over arbitrary paths, closer to how logics such as LTLf approach temporal logic over finite traces [De Giacomo and Vardi, 2013].

Definition 5 (Semantics)**.** Given an MDP $\mathbb{M}$, the *semantics of PLBP+* is defined via simultaneous induction over state, policy and path formulas. We omit the semantics of Boolean operators as it is standard. For state formulas, it is defined for states s as

- $\mathbb{M}, s \Vdash p$ iff $p \in V(s)$,
- $\mathbb{M}, s \Vdash \Diamond^n \alpha$ iff there exists an n-step $(\mathbb{M}, s)$-policy π such that $\mathbb{M}, \pi \Vdash \alpha$.

For policy formulas, it is defined for an $(\mathbb{M}, s)$-policy π as

- $\mathbb{M}, \pi \Vdash \mathsf{do}_a$ iff $\pi(s) = a$,
- $\mathbb{M}, \pi \Vdash \mathsf{Pr}_{\bowtie r}\Phi$ iff $\mu_\pi(\{\mathbf{w} \in \mathsf{Paths}(\pi) \mid \mathbb{M}, \pi, \mathbf{w} \Vdash \Phi\}) \bowtie r$.

For path formulas, it is defined for policies π and paths $\mathbf{w} = s_0 \cdots s_n \in \mathsf{Paths}(\pi)$ as

- $\mathbb{M}, \pi, \mathbf{w} \Vdash \varphi$ iff $\mathbb{M}, s_0 \Vdash \varphi$,
- $\mathbb{M}, \pi, \mathbf{w} \Vdash \alpha$ iff $\mathbb{M}, \pi \Vdash \alpha$,
- $\mathbb{M}, \pi, \mathbf{w} \Vdash \mathsf{X}\Phi$ iff $n \geqslant 1$ and $\mathbb{M}, \pi_{s_1}, s_1 \cdots s_n \Vdash \Phi$.

For a state formula φ, we write $[\![\varphi]\!] := \{\mathbb{P} \in \mathsf{pMDP} \mid \mathbb{P} \Vdash \varphi\}$. We say $\varphi \Vdash \psi$ if $[\![\varphi]\!] \subseteq [\![\psi]\!]$, and $\varphi \equiv \psi$ if $[\![\varphi]\!] = [\![\psi]\!]$. For a path formula Φ, we write $[\Phi]$ for the set of all triples $(\mathbb{M}, \pi, \mathbf{w})$ consisting of an MDP $\mathbb{M}$, a $(\mathbb{M}, s)$-policy for some state s, and a path $\mathbf{w} \in \mathsf{Paths}(\pi)$ such that $\mathbb{M}, \pi, \mathbf{w} \Vdash \Phi$. Given a $(\mathbb{M}, s)$-policy π, we write $[\Phi]_\pi := \{\mathbf{w} \in \mathsf{Paths}(\pi) \mid \mathbb{M}, \pi, \mathbf{w} \Vdash \Phi\}$. We say $\Phi \vDash \Psi$ if $[\Phi] \subseteq [\Psi]$, and $\Phi \cong \Psi$ if $[\Phi] = [\Psi]$.

As is usual, we define the other Boolean connectives $\vee$ and $\rightarrow$ as abbreviations, as well as the propositional constants $\top$ and $\bot$ for verity and falsity. Similarly, we can also define a universally quantified version of the $\Diamond^n$-modality by putting $\Box^n := \neg\Diamond^n\neg$. And we define $\mathsf{Pr}_{\bowtie r}\Phi$ for symbols $\bowtie \in \{=, \neq, >, <\}$ by appropriate abbreviations, e.g. $\mathsf{Pr}_{>r}\Phi := \neg\mathsf{Pr}_{\leqslant r}\Phi$. Note that we can also define more complicated (bounded) temporal operators, such as the bounded *until* $\Phi \; \mathsf{U}^n \; \Psi := \bigvee_{0 \leqslant k \leqslant n}(\mathsf{X}^k\Psi \wedge \bigwedge_{0 \leqslant \ell < k} \mathsf{X}^\ell \Phi)$, where X^k is an abbreviation for X k-many times. And finally, observe that we may also define a general $\Diamond$-modality with the semantics of there being an n-step policy for *some* n. This is done by observing that for each α there is a number $\mathsf{td}(\alpha)$ (the *temporal depth* of α) such that if there is an n-step policy for which α holds, then there is an n-step policy with $n \leqslant \mathsf{td}(\alpha)$ for which α holds. This $\mathsf{td}(\alpha)$ is intuitively the amount of nested X-operators and do_a-propositions appearing in α. We can then define $\Diamond\alpha := \bigvee_{0 \leqslant n \leqslant \mathsf{td}(\alpha)} \Diamond^n \alpha$.

Satisfiability and therefore entailment for PLBP+ state formulas is decidable, placing the problem in the class 2EXPSPACE of problems decidable with a deterministic algorithm using space double exponential w.r.t. the input.

Theorem 1 (Satisfiability of PLBP+)**.** *The satisfiability problem for PLBP+ state formulas (i.e. given a state formula φ, decide whether $[\![\varphi]\!] \neq \emptyset$) lies in 2EXPSPACE.*

Proof sketch. The proof is an easy adaptation of the proof for PLBP. Using the unravelling method of modal logic, we observe that PLBP+ possesses the finite model property, with satisfiable formulas having models of of size up to a bound calculatable from the formula. We iterate over sets S of size up to this bound. We then define a first-order logic formula α over the signature of real-closed field (RCF; see [Chang and Keisler, 2012] for a definition), such that α is valid in RCF iff there exists a valuation and probabilistic transition function over S such that the resulting pMDP satisfies φ. Since α is of double exponential size w.r.t. φ, and is in existential prenex form, it follows from existential RCF being decidable in polynomial space [Canny, 1988] that the overall procedure runs in double exponential space. □

Having discussed the logic we base our work on, we can now introduce how we define beliefs and intentions. We identify beliefs with state formulas: the agent's beliefs are about what the environment (pMDP) looks like, independent of how the agent chooses to act. As an example, the formula $\Box^1(\mathsf{do}_a \rightarrow \mathsf{Pr}_{=.5}\mathsf{X}\mathsf{post}_{a,i})$ expresses the belief that after executing a now, we get the ith postcondition of a afterwards with probability 50%. Similarly, we identify intentions with path formulas: the agent intends to bring about certain outcomes (paths) throughout time, as in e.g.[Rao and Georgeff, 1991; Khan and Lespérance, 2012; Leask, 2021]. As an example, the formula $\mathsf{do}_a \; \mathsf{U}^n \; \mathsf{Pr}_{>r}\Phi$ expresses the intention to keep executing a until a point from which Φ will hold with probability $> r$, for a maximum of n steps. Note that this gives us a very rich notion of intention, going far beyond the atomic action intentions ("I intend to perform a at time t") of [Shoham, 2009; van Zee *et al.*, 2020].

Example 1 (Running example)**.** Coming back to our example from Section 1 about Theo and the concert, we can rewrite the beliefs and intentions therein as PLBP+ formulas (fixing some probabilities for the terms 'likely' and 'unlikely'). The relevant propositional variables are $\mathsf{Prop} = \{\mathsf{five}, \mathsf{conc}\}$, with the former denoting Theo having \$5, and the latter denoting Theo being at the concert. The example concerns the actions stay of staying at home, goTick of going to the ticket office to buy a ticket, and buyDr of buying his favorite drink. The relevant preconditions are $\mathsf{pre}_{\mathsf{goTick}} = \mathsf{five}$ (it costs \$5 to buy tickets), and $\mathsf{pre}_{\mathsf{buyDr}} = \mathsf{five}$ (Theo needs to have \$5 to buy drinks). The relevant postconditions are $\mathsf{Post}_{\mathsf{goTick}} = \{\mathsf{conc} \wedge \neg\mathsf{five}, \neg\mathsf{conc} \wedge \mathsf{five}\}$ (Theo either gets to attend and spends \$5, or he doesn't and keeps it), $\mathsf{Post}_{\mathsf{stay}} = \{\mathsf{conc} \wedge \mathsf{five}, \neg\mathsf{conc} \wedge \mathsf{five}\}$ (whether he gets a ticket or not, he keeps his money), and $\mathsf{Post}_{\mathsf{buyDr}} = \{\neg\mathsf{five}\}$ (buying a drink is deterministic: Theo pays \$5). His original beliefs are modelled by the state formula $\varphi_1 = \mathsf{five} \wedge \Box((\mathsf{do}_{\mathsf{goTick}} \vee \mathsf{do}_{\mathsf{stay}}) \wedge \mathsf{do}_{\mathsf{goTick}} \rightarrow \mathsf{Pr}_{\geqslant .9}\mathsf{X}\mathsf{conc} \wedge \mathsf{do}_{\mathsf{stay}} \rightarrow \mathsf{Pr}_{\leqslant .01}\mathsf{X}\mathsf{conc})$, i.e. Theo has \$5, can try to buy a ticket or stay at home, and with probability at least 90% the former gets him to the concert, while at most 1% the latter does. The intentions of attending the concert and buying drinks are respectively $\Phi_1 = \mathsf{X}\mathsf{conc}$ and $\Phi_2 = \mathsf{X}\mathsf{do}_{\mathsf{buyDr}}$. The statement that the band has become popular (i.e. tickets are likely to be sold out) is $\varphi_2 = \Box(\mathsf{do}_{\mathsf{goTick}} \rightarrow \mathsf{Pr}_{\leqslant 0.01}\mathsf{X}\mathsf{conc})$ (there's only a 1% chance Theo gets to buy tickets).

3 Coherence and Beliefs

In this section, we introduce our proposal for defining the 'coherence' of beliefs and intentions, as well as what it means to believe something on the basis of certain intentions.

As mentioned in Section 1, it is not sufficient that an agent's beliefs and intentions are separately consistent: we need joint consistency. This joint consistency between beliefs and intentions is what has been referred to as *coherence* by [van Zee *et al.*, 2020]. What we want, is a requirement on beliefs and intentions, stating that the agent's beliefs should somehow be consistent with the possible achievement of their intentions. Note that this is effectively the Strong Consistency Principle proposed by [Bratman, 1987] in his influential theory of intentions.

Coherence is defined in prior work as stating that there exists some model of the agent's beliefs, in which the agent has a policy which achieves the intention. Translating this definition to our setting, we are asking given a belief φ and an intention Φ, whether there exists some $\mathbb{P} \in [\![\varphi]\!]$, and a $\mathbb{P}$-policy π, such that π 'achieves' Φ. What does it mean that π 'achieves' Φ? Contrary to the prior work, in our setting actions have uncertain outcomes: some of the paths produced by π may satisfy Φ, while others do not. It is sensible to define the notion of achievement with regards to the probability with which these paths satisfy Φ, i.e. $\mu_\pi([\Phi]_\pi)$.

We therefore propose to take a parametric, application-dependent approach to coherence, in which we take some value $0 \leqslant \theta < 1$ with the intuition being that the agent only considers policies that achieve intentions with probability greater than θ. The value of θ will depend on the specific problem domain or type of agent we are considering at the time.

Definition 6 (θ-coherence)**.** Let $0 \leqslant \theta < 1$ be a rational number. A state formula φ and path formula Φ are *θ-coherent* if there exists a pMDP $\mathbb{P} \in [\![\varphi]\!]$ and a $\mathbb{P}$-policy π such that $\mu_\pi([\Phi]_\pi) > \theta$.

Note that this definition is equivalent to stating that the formula $\varphi \wedge \Diamond \mathsf{Pr}_{>\theta}\Phi$ is satisfiable. Therefore, by Theorem 1 we get that it is actually decidable whether beliefs and intentions are θ-coherent.

Example 2 (Running example, continued)**.** Our discussion in Section 1 effectively observed that Theo's strong beliefs φ_1 and intentions $\Phi_1 \wedge \Phi_2$ *are* 0-coherent (where $\Phi_1 = \mathsf{X}\mathsf{conc}$ and $\Phi_2 = \mathsf{X}\mathsf{do}_{\mathsf{buyDr}}$), due the possibility of choosing the action stay. But note that this combined intention is *not* e.g. 0.05-coherent with φ_1, since in any model of φ_1, any policy will assign at most probability 0.01 to $\Phi_1 \wedge \Phi_2$. We do have that Φ_1 alone *is* 0.05-coherent with φ_1 - the same holds for Φ_2.

While beliefs as we have discussed them so far have been about what the environment looks like, independent of how the agent actually chooses to act, an essential part of a theory of belief and intention is that since an agent is committed to performing intentions, they actually in some sense 'believe' that the outcomes of those intentions will hold. To model this, [van der Hoek *et al.*, 2007] propose to separate so-called *strong* beliefs from *weak* beliefs. The former are what we have considered so far: independent of how the agent actually acts. The latter are believed on the assumption the agent performs their intentions. The separation prevents scenarios like the Little Nell Problem [McDermott, 1982].

Since weak beliefs are contingent on the performance of intentions, we propose to consider these to be path formulas as well: we weakly believe certain properties (like outcomes of intended actions) will hold throughout time. This is similar to the approach to weak beliefs of [van Zee *et al.*, 2020]. Trying to use their approach naively, we could choose to define the weak beliefs obtained from our beliefs φ and intentions Φ to be those path formulas Ψ such that for all $\mathbb{P} \in [\![\varphi]\!]$ and $\mathbb{P}$-policies π, we have for all $\mathbf{w} \in \mathsf{Paths}(\pi)$ that if $\mathbb{M}, \pi, \mathbf{w} \Vdash \Phi$, then also $\mathbb{M}, \pi, \mathbf{w} \Vdash \Psi$.

But this approach is problematic, as our running example of Theo and the concert shows. By our naive definition of weak beliefs, Theo will *not* weakly believe that he will have \$0 tomorrow, because of the existence of the policy in which he stays at home today. One could certainly argue that it is more practical and rational that Theo should in fact believe that, assuming he manages to attend the concert, he will be out of money. The problem here, as with coherence, is that the policy of staying at home achieves the intention with too low a probability. Therefore, we also parameterize the definition of weak beliefs by the same θ stating when a policy should be considered as 'achieving' an intention, removing from consideration policies which only achieve the intention by sheer luck.

Definition 7 (θ-weak beliefs)**.** Let $0 \leqslant \theta < 1$ be a rational number. Given a state formula φ and path formula Φ, the set $\mathsf{WB}_\theta(\varphi, \Phi)$ of *θ-weak beliefs* of φ and Φ is defined as the set of all path formulas Ψ such that for all $\mathbb{P} \in [\![\varphi]\!]$ and $\mathbb{P}$-policies π such that $\mu_\pi([\Phi]_\pi) > \theta$, it holds for all $\mathbf{w} \in \mathsf{Paths}(\pi)$ that if $\mathbb{M}, \pi, \mathbf{w} \Vdash \Phi$, then $\mathbb{M}, \pi, \mathbf{w} \Vdash \Psi$.

Note that $\Psi \in \mathsf{WB}_\theta(\varphi, \Phi)$ if and only if $\varphi \wedge \mathsf{Pr}_{>\theta}\Phi \wedge \Phi \vDash \Psi$, which is again if and only if $\varphi \Vdash \Box(\mathsf{Pr}_{>\theta}\Phi \rightarrow \mathsf{Pr}_{=1}(\Phi \rightarrow \Psi))$. By Theorem 1, we can therefore decide whether a given formula Ψ is in $\mathsf{WB}_\theta(\varphi, \Phi)$.

Example 3 (Running example, continued)**.** We have that indeed $\mathsf{X}\neg\mathsf{five} \notin \mathsf{WB}_0(\varphi_1, \Phi_1)$, while $\mathsf{X}\neg\mathsf{five} \in \mathsf{WB}_{0.05}(\varphi_1, \Phi_1)$ and $\mathsf{XX}\neg\mathsf{five} \in \mathsf{WB}_{0.05}(\varphi_1, \Phi_2)$: Theo 0.05-weakly believes he will not have his money tomorrow if he intends to attend the concert tomorrow. And he also 0.05-weakly believes he will not have his money the day after tomorrow if he intends to buy a drink tomorrow.

We see directly from the definitions that the satisfiability of weak beliefs is equivalent to the coherence of the corresponding strong beliefs and intentions.

4 Representation for Revision Operators

We now move our attention to the revision of beliefs and intentions. Similar to prior work [van Zee *et al.*, 2020], we follow the approach of revising an agent's strong beliefs φ and intentions Φ with new strong beliefs ψ and intentions Ψ, by first revising the strong beliefs independently of the intentions, and afterwards revising the intentions in a way that takes the revised strong beliefs into account. We explore this joint revision in Section 4.3.

We will start by discussing belief revision operators, and providing a representation theorem for them. By virtue of our setup, we will afterwards see that almost all of the machinery we develop here will also apply to intention revision operators. We conclude this section by discussing what our setup tells us about the interplay between belief and intention.

Throughout this section, we write Bel for the set of state formulas, and Int for the set of path formulas.

4.1 Belief Revision

Definition 8 (Belief revision operator)**.** A *belief revision operator* is a function $\circ\colon \mathsf{Bel} \times \mathsf{Bel} \to \mathsf{Bel}$, with $\varphi \circ \psi$ being read as the result of *revising* φ *by* ψ.

Conceptually, we wish for belief revision operators to satisfy the intuition that they are revising φ by ψ by 'minimally' changing φ until ψ can be added without problems. But the original approach of [Katsuno and Mendelzon, 1991] with finite-signature propositional logic does not apply to our setting: there are *infinitely* many pMDPs due to e.g. the presence of probability in our models, and generally speaking, there exists no formula φ that expresses even a single pMDP $\mathbb{P}$ (i.e. such that $[\![\varphi]\!] = \{\mathbb{P}\}$), since we could always consider pMDPs which differ from $\mathbb{P}$ only at times which are 'outside the scope' of φ. Luckily, the approach of [Falakh *et al.*, 2023] still allows us to obtain a representation theorem by adding some requirements on the semantic side, which they refer to a *min-completeness* and *min-expressibility*.

First some preliminaries on orders are required. A relation $\preceq \subseteq X \times X$ is a *total preorder* if it is reflexive, transitive and connected (i.e. $x \preceq y$ or $y \preceq x$ for all $x, y \in X$). We write $x \prec y$ iff $x \preceq y$ and $y \not\preceq x$. We say $x \in U \subseteq X$ is $\preceq$*-minimal in* U if $x \preceq y$ for all $y \in U$, and write $\min(U, \preceq)$ for the set of all such $\preceq$-minimal elements in U.

Definition 9 (pMDP assignments)**.** A *pMDP assignment* is a function $\preceq_{(-)}\colon \mathsf{Bel} \to 2^{\mathsf{pMDP}\times\mathsf{pMDP}}$ such that $\preceq_\varphi$ is a total preorder for all φ. The assignment is *faithful* if the following three conditions hold: (i) If $\mathbb{P} \Vdash \varphi$ and $\mathbb{P}' \Vdash \varphi$, then $\mathbb{P} \preceq_\varphi \mathbb{P}'$ and $\mathbb{P}' \preceq_\varphi \mathbb{P}$. (ii) If $\mathbb{P} \Vdash \varphi$ and $\mathbb{P}' \nVdash \varphi$, then $\mathbb{P} \prec_\varphi \mathbb{P}'$. (iii) If $\varphi \equiv \psi$, then $\preceq_\varphi = \preceq_\psi$. The assignment is *min-complete* if $\min([\![\psi]\!], \preceq_\varphi) \neq \emptyset$ for all φ and ψ with $[\![\psi]\!] \neq \emptyset$. It is *min-expressible* if for all φ and ψ there exists $\psi^{\preceq,\varphi} \in \mathsf{Bel}$ such that $[\![\psi^{\preceq,\varphi}]\!] = \min([\![\psi]\!], \preceq_\varphi)$.

The rationality postulates used by [Falakh *et al.*, 2023] are precisely those of [Katsuno and Mendelzon, 1991], and are equivalent to the original presentation of the AGM postulates by [Alchourrón *et al.*, 1985]. Though Katsuno & Mendelzon's work was concerned with finite-signature propositional logic, we emphasize that the rationality postulates are applicable in much more general settings: citing [Segerberg, 1999], "AGM is not really logic; it is a theory about theories".

The postulates are as follows.

(B1) $\varphi \circ \psi \Vdash \psi$.

(B2) If $\varphi \wedge \psi$ is satisfiable, then $\varphi \circ \psi \equiv \varphi \wedge \psi$.

(B3) If ψ is satisfiable, then $\varphi \circ \psi$ is satisfiable.

(B4) If $\varphi_1 \equiv \varphi_2$ and $\psi_1 \equiv \psi_2$, then $\varphi_1 \circ \psi_1 \equiv \varphi_2 \circ \psi_2$.

(B5) $(\varphi \circ \psi) \wedge \chi \Vdash \varphi \circ (\psi \wedge \chi)$.

(B6) If $(\varphi\circ\psi)\wedge\chi$ is satisfiable, then it holds that $\varphi\circ(\psi\wedge\chi) \Vdash (\varphi \circ \psi) \wedge \chi$.

The postulate (B1) expresses the success of revision: the new information must be incorporated by the revision. The postulate (B2) expresses that when the new information is consistent with our beliefs, revision should be the same as just adding the new information. The postulate (B3) expresses that we maintain consistency whenever we add consistent information to our beliefs, and the postulate (B4) expresses that revision should not care about syntax. Finally, the postulates (B5) and (B6) ensure that revision can be thought of as based on a notion of *minimal change*, with the outcome of revision $\varphi \circ \psi$ giving those models of ψ that are 'closest' to models of φ. Given such a interpretation, we would naturally have that a model of ψ which is closest to φ and satisfies χ is automatically also the model of $\psi \wedge \chi$ closest to φ, which is exactly what (B5) states. And similarly, we would never have that there exist models $M, N \in [\![\psi \wedge \chi]\!]$ such that M is closest to φ amongst $[\![\psi \wedge \chi]\!]$ but N is *strictly* closer to φ than M amongst $[\![\psi]\!]$, which is exactly what (B6) rules out.

Example 4 (Running example, continued)**.** Consider Theo's original strong beliefs φ_1 and the new piece of information $\varphi_2 = \Box(\mathsf{do}_{\mathsf{goTick}} \to \mathsf{Pr}_{\leqslant 0.01}\mathsf{X}\mathsf{conc})$ (which he gains by learning the band is actually popular). Then the postulates tell us already several things about how the revision $\varphi_1 \circ \varphi_2$ will look, for any belief revision operator $\circ$ satisfying them. For example, (B3) tells us that $\varphi_1 \circ \varphi_2$ must be satisfiable, since φ_2 clearly is. We get that $\varphi_1 \circ \varphi_2$ cannot be equivalent to $\varphi_1 \wedge \varphi_2$, as that formula is not satisfiable: Theo has to drop some beliefs. But whatever he drops, (B1) tells us that $\varphi_1 \circ \varphi_2 \Vdash \varphi_2$: he comes to believe that if he goes to the ticket office, the probability that tickets are available is at most 1%.

The following general representation result, Theorem 2, expresses that belief revision minimally changes beliefs to accommodate new information.

Theorem 2 (General representation for PLBP+)**.** *The following statements hold:*

- *For every min-complete, min-expressible and faithful pMDP assignment $\preceq_{(-)}$ there exists a belief revision operator $\circ$ satisfying (B1)-(B6) such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.*
- *For every belief revision operator $\circ$ satisfying (B1)-(B6) there exists a min-complete, min-expressible and faithful pMDP assignment $\preceq_{(-)}$ such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.*

Proof. This follows from the results of [Falakh *et al.*, 2023], as we now explain. Consider PLBP+ state formulas as what they refer to as a *base logic*: PLBP+ state formulas are the *sentences*, pMDPs are the *worlds*, the satisfaction relation $\Vdash$ is used, the *bases* are the singletons $\{\varphi\}$ of state formulas, and the *abstract union* of $\{\varphi\}$ and $\{\psi\}$ is $\{\varphi \wedge \psi\}$. Then what they refer to as *base change operators*, are precisely what we refer to as *belief revision operators*.

[Delgrande *et al.*, 2018] introduce a postulate (Acyc), which (applied to PLBP+ state formulas as a base logic) states the following:

- For any φ and $\psi_0, \ldots, \psi_n$ with $[\![\psi_i \wedge (\varphi \circ \psi_{(i+1) \bmod (n+1)})]\!] \neq \emptyset$ for all i, it holds that $[\![\psi_0 \wedge (\varphi \circ \psi_n)]\!] \neq \emptyset$.

By Theorem 10.13 of [Falakh *et al.*, 2023], we get that the following statements hold:

- For every min-complete, min-expressible and faithful pMDP assignment $\preceq_{(-)}$, there exists a belief revision operator $\circ$ satisfying (B1)-(B6) and (Acyc) such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.
- For every belief revision operator $\circ$ satisfying (B1)-(B6) and (Acyc), there exists a min-complete, min-expressible and faithful pMDP assignment $\preceq_{(-)}$ such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.

We see that in fact, PLBP+ state formulas form what they call a *disjunctive* base logic, since PLBP+ state formulas contain disjunctions. Therefore, their Corollary 10.17 applies and tells us that any belief revision operator $\circ$ satisfying (B1)-(B6) also automatically satisfies (Acyc), which completes our proof. □

While Theorem 2 gives us a representation theorem for belief revision operators, it does not give us much in the way of *practical* belief revision. It is not clear from the start how we can come up with min-complete and min-expressible pMDP assignments. And even if we find such assignments, we have no guarantee in general that the resulting belief revision operator $\circ$ is actually *computable*. Therefore we propose *finitely expressible* pMDP assignments, which are a more specific class of pMDP assignments that are both easily understood from an intuitive, conceptual point of view, as well as computable under very mild assumptions. We start with some technical preliminaries.

Definition 10 (Finitariness)**.** A total preorder $\preceq$ on a (possibly infinite) set X is *finitary* if the set $\{[x]_\preceq \mid x \in X\}$ is finite, where $[x]_\preceq = \{y \in X \mid x \preceq y \text{ and } y \preceq x\}$. Given a finitary total preorder $\preceq$ and $x \in X$, we write $\mathsf{level}_\preceq(x) := |\{[y]_\preceq \mid y \in X \text{ and } y \preceq x\}|$. Writing $\mathsf{height}(\preceq) := |\{[x]_\preceq \mid x \in X\}|$, we define the *level sets* $X_{\preceq,1}, \ldots, X_{\preceq,\mathsf{height}(\preceq)}$ by putting $X_{\preceq,k} := \{x \in X \mid \mathsf{level}_\preceq(x) = k\}$.

The idea is that a finitary total preorder may be defined on an infinite set, but it divides the set into finitely many 'levels' of equivalent elements.

Definition 11 (FEF pMDP assignments)**.** A pMDP assignment $\preceq_{(-)}$ is *finitely expressible* if $\preceq_\varphi$ is finitary for all $\varphi \in \mathsf{Bel}$, and for all φ there exist $\varphi^1, \ldots, \varphi^{\mathsf{height}(\preceq_\varphi)} \in \mathsf{Bel}$ such that $[\![\varphi^k]\!] = \mathsf{pMDP}_{\preceq_\varphi,k}$, i.e. the k-th level set of pMDP according to $\preceq_\varphi$. We refer to finitely expressible and faithful pMDP assignments as *FEF pMDP assignments* for short.

Finitely expressible pMDP assignments divide the space of pMDPs into finitely many levels for every φ, with every level expressible by formulas of PLBP+. And indeed, finite expressibility implies both min-completeness and min-expressibility.

FEF pMDP assignments are conceptually simple, because they are in fact the same as specifying for every φ a finite sequence of *weakenings* of φ, which corresponds intuitively to the idea that in revision we keep removing information step-by-step from our beliefs until they can accommodate the new information. We make this correspondence precise in Proposition 1.

Definition 12 (Belief weakening maps)**.** A *belief weakening map* is a function $\langle - \rangle \colon \mathsf{Bel} \to \mathsf{Bel}^+$ (where we write Bel^+ to denote the set of all non-empty finite sequences of $\varphi \in \mathsf{Bel}$) assigning to all $\varphi \in \mathsf{Bel}$ a finite sequence $\langle\varphi\rangle_1 \Vdash \cdots \Vdash \langle\varphi\rangle_n$, such that $\langle\varphi\rangle_1 \equiv \varphi$.

Proposition 1. *The following statements hold:*

- *For every FEF pMDP assignment $\preceq_{(-)}$ there exists a belief weakening map $\langle - \rangle$ such that $|\langle\varphi\rangle| = \mathsf{height}(\preceq_\varphi) + 1$ and $[\![\langle\varphi\rangle_k]\!] = \bigcup_{1 \leqslant \ell \leqslant k} \mathsf{pMDP}_{\preceq_\varphi,\ell}$ for all φ and k.*
- *For every belief weakening map $\langle - \rangle$, there exists an FEF pMDP assignment $\preceq_{(-)}$ such that $\mathsf{height}(\preceq_\varphi) \leqslant |\langle\varphi\rangle| + 1$ and $[\![\langle\varphi\rangle_k]\!] = \bigcup_{1 \leqslant \ell \leqslant k} \mathsf{pMDP}_{\preceq_\varphi,\ell}$ for all φ and k.*

Proof. Take an FEF pMDP assignment $\preceq_{(-)}$. Since it is finitely expressible, we have that for all φ there exist $\varphi^1, \ldots, \varphi^{\mathsf{height}(\preceq_\varphi)}$ such that $[\![\varphi^k]\!] = \mathsf{pMDP}_{\preceq_\varphi,k}$. Defining the belief weakening map by putting $\langle\varphi\rangle_1 := \varphi$ and $\langle\varphi\rangle_{k+1} := \bigvee_{1 \leqslant \ell \leqslant k} \varphi^\ell$ for $1 \leqslant k \leqslant \mathsf{height}(\preceq_\varphi)$, we indeed get the required property.

Now take a belief weakening map $\langle - \rangle$. Given φ and $\mathbb{P}$, write

$$\langle\varphi, \mathbb{P}\rangle := \begin{cases} \min\{1 \leqslant \ell \leqslant |\langle\varphi\rangle| \mid \mathbb{P} \Vdash \langle\varphi\rangle_\ell\} & \text{if } \mathbb{P} \Vdash \bigvee\langle\varphi\rangle \\ \infty & \text{otherwise} \end{cases},$$

i.e. the number of the first weakening of φ that $\mathbb{P}$ satisfies if possible, and otherwise ∞. We define the pMDP assignment by putting $\mathbb{P} \preceq_\varphi \mathbb{P}'$ iff $\langle\varphi, \mathbb{P}\rangle \leqslant \langle\varphi, \mathbb{P}'\rangle$. It is immediate that this assignment produces total preorders, since it is defined by comparing numbers. Faithfulness follows from the observations that (i) $\langle\varphi, \mathbb{P}\rangle = 1$ if $\mathbb{P} \Vdash \varphi$, (ii) $\langle\varphi, \mathbb{P}\rangle > 1$ if $\mathbb{P} \nVdash \varphi$, and (iii) the definition being completely syntax-invariant.

Finitariness follows from $\langle\varphi, \mathbb{P}\rangle$ taking *at most* $|\langle\varphi\rangle| + 1$ values. If we do not have $\langle\varphi\rangle_k \equiv \langle\varphi\rangle_\ell$ for any $k \neq \ell$ (i.e. the weakenings are semantically distinct), if $\bigvee\langle\varphi\rangle$ is not a tautology (i.e. the value $\langle\varphi, \mathbb{P}\rangle = \infty$ can appear), and if φ is satisfiable, then $\mathsf{height}(\preceq_\varphi) = |\langle\varphi\rangle| + 1$. If any of those assumptions do not hold (e.g. if $\bigvee\langle\varphi\rangle$ *is* a tautology), then the height may be lower, though never higher.

To see that the assignment is finitely expressible, we recursively define for $1 \leqslant k \leqslant \mathsf{height}(\preceq_\varphi)$ the formulas φ^k by putting

$$\varphi^k := \langle\varphi\rangle_{\max\{\ell;\langle\varphi\rangle_\ell \Vdash \bigvee_{1 \leqslant m < k} \varphi^m\}+1} \wedge \neg \bigvee_{1 \leqslant m < k} \varphi^m,$$

where the maximum is taken to be 0 if it would not be defined.

The idea of this construction is that it takes care of the possibilities that (i) $\varphi \equiv \bot$, and that (ii) $\langle\varphi\rangle_k \equiv \langle\varphi\rangle_\ell$ for some $k \neq \ell$. □

We now argue that the belief revision operators resulting from FEF pMDP assignments (in the sense of Theorem 2) are *computable* if the underlying belief weakening map is computable. To compute $\varphi \circ \psi$, first check if ψ is satisfiable (which is decidable by Theorem 1): if it is not, we safely output $\bot$. Then compute $\langle\varphi\rangle_1, \ldots, \langle\varphi\rangle_n$, and iterate over all $1 \leqslant k \leqslant n$, checking whether $\psi \wedge \langle\varphi\rangle_k$ is satisfiable. For the first such k, output the formula $\psi \wedge \langle\varphi\rangle_k$. If no such k exists, output ψ.

We now show that we can actually give a representation theorem for FEF pMDP assignments. For this we introduce a novel postulate (Bω), expressing the finitariness of the revision operator. In the postulate, we define a *belief partition* to be a sequence $\varphi^1, \ldots, \varphi^n \in \mathsf{Bel}$ of some length n, such that (i) $\varphi^1 \vee \cdots \vee \varphi^n$ is a tautology, (ii) φ^k is satisfiable for all k, and (iii) $\varphi^k \wedge \varphi^\ell$ is unsatisfiable for all $k \neq \ell$.

(Bω) For all φ there exists a belief partition $\varphi^1, \ldots, \varphi^n$, such that for all ψ, there exists k with $\varphi \circ \psi \equiv \psi \wedge \varphi^k$.

Theorem 3 (Representation for FEF pMDP assignments)**.** *The following statements hold:*

- *For every FEF pMDP assignment $\preceq_{(-)}$ there exists a belief revision operator $\circ$ satisfying (B1)-(B6) and (Bω), such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.*
- *For every belief revision operator $\circ$ satisfying (B1)-(B6) and (Bω) there exists an FEF pMDP assignment $\preceq_{(-)}$ such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$.*

Proof. Take an FEF pMDP assignment $\preceq_{(-)}$. Since finite expressibility entails min-expressibility, we apply the construction used by [Falakh *et al.*, 2023] in their Definition 6.3: we let $\varphi \circ \psi := \psi^{\preceq,\varphi}$ (which we defined in our Definition 9). As observed by [Falakh *et al.*, 2023] in the proof of their Theorem 6.4, it follows from the min-completeness (and *min-retractiveness*, which is entailed by our assignment producing total preorders) of $\preceq$ (which we get again from it being finitely expressible) that $\circ$ satisfies (B1)-(B6). And by definition of min-expressibility, we get that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$. It remains to show that $\circ$ also satisfies (Bω). It actually suffices to see that the formulas $\varphi^1, \ldots, \varphi^k$ we get from $\preceq$ being finitely expressible, form a belief partition. Since $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$, it follows from finitariness that for all φ and ψ there is some k such that $[\![\varphi \circ \psi]\!] = [\![\psi]\!] \cap \mathsf{pMDP}_{\preceq,k} = [\![\psi \wedge \varphi^k]\!]$.

Take a belief revision operator $\circ$ satisfying (B1)-(B6) and (Bω). By Corollary 10.17 of [Falakh *et al.*, 2023], we know that $\circ$ satisfies their postulate (Acyc) as well. Therefore, by their Theorem 10.13, there exists a faithful pMDP assignment $\preceq$ such that $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$. Fix φ. Since $\circ$ satisfies (Bω), there exists a belief partition $\varphi^1, \ldots, \varphi^n$ such that for all ψ, there is some k with $\varphi \circ \psi \equiv \psi \wedge \varphi^k$.

In the following, we introduce the notation $\mathsf{pMDP}/{\preceq_\varphi} = \{[\mathbb{P}]_{\preceq_\varphi} \mid \mathbb{P} \in \mathsf{pMDP}\}$.

For the finite expressibility of $\preceq$, we will show that $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\} = \mathsf{pMDP}/{\preceq_\varphi}$, which is precisely what finite expressibility means. We note that it suffices to show that (i) for all $1 \leqslant k \leqslant n$, there is some $X \in \mathsf{pMDP}/{\preceq_\varphi}$ such that $[\![\varphi^k]\!] \subseteq X$, and (ii) the previous correspondence is injective, i.e. for all $X \in \mathsf{pMDP}/{\preceq_\varphi}$ there is at most one $1 \leqslant k \leqslant n$ with $[\![\varphi^k]\!] \subseteq X$. The reason this would suffice is as follows. By definition of belief partitions, the sets $[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]$ form a partition of pMDP. And clearly, $\mathsf{pMDP}/{\preceq_\varphi}$ is also a partition. So (i) states that $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\}$ is a refinement of $\mathsf{pMDP}/{\preceq_\varphi}$. But in fact, we can see that $\mathsf{pMDP}/{\preceq_\varphi}$ is also a refinement of $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\}$, and thus, since refinement is an antisymmetric relation, we are able to conclude that $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\} = \mathsf{pMDP}/{\preceq_\varphi}$. To see that $\mathsf{pMDP}/{\preceq_\varphi}$ is also a refinement of $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\}$, take any $X \in \mathsf{pMDP}/{\preceq_\varphi}$, and suppose for the sake of contradiction that there is no k such that $X \subseteq [\![\varphi^k]\!]$. Since $\{[\![\varphi^1]\!], \ldots, [\![\varphi^k]\!]\}$ is a partition and it must be that $X \neq \emptyset$, this means that there exist distinct $\ell \neq m$ such that $X \cap [\![\varphi^\ell]\!] \neq \emptyset \neq X \cap [\![\varphi^m]\!]$. But by (i), we have that $[\![\varphi^\ell]\!] \subseteq X^\ell$ and $[\![\varphi^m]\!] \subseteq X^m$ for some $X^\ell, X^m \in \mathsf{pMDP}/{\preceq_\varphi}$, and by (ii), we must have that $X^\ell \neq X^m$. Therefore, we have that $X \cap X^\ell \neq \emptyset \neq X \cap X^m$, which contradicts $\mathsf{pMDP}/{\preceq_\varphi}$ being a partition, since a cell cannot intersect two distinct cells (it only intersects itself).

To see that (i) holds, take any φ^k. By definition of belief partitions, φ^k is satisfiable. Therefore, by postulate (B3), we get that $\varphi \circ \varphi^k$ is satisfiable as well. By (Bω), it must be that $\varphi \circ \varphi^k \equiv \varphi^k \wedge \varphi^l$ for some ℓ. But ℓ must be equal to k, since φ^k and φ^ℓ for distinct $k \neq \ell$ are mutually unsatisfiable, which would contradict $\varphi \circ \varphi^k$ being satisfiable. Therefore, $\varphi \circ \varphi^k \equiv \varphi^k$. Since $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$ for all ψ, we get in particular that $[\![\varphi^k]\!] = \min([\![\varphi^k]\!], \preceq_\varphi)$. But that means that for all $\mathbb{P}, \mathbb{P}' \in [\![\varphi^k]\!]$, we have that $\mathbb{P} \preceq_\varphi \mathbb{P}'$ and vice-versa. So by definition of $\mathsf{pMDP}/{\preceq_\varphi}$, we get that $\varphi^k \subseteq X$ for some $X \in \mathsf{pMDP}/{\preceq_\varphi}$.

To see that (ii) holds, take any $X \in \mathsf{pMDP}/{\preceq_\varphi}$, and suppose for the sake of contradiction that there exist $k \neq \ell$ such that $[\![\varphi^k]\!] \subseteq X$ and $[\![\varphi^\ell]\!] \subseteq X$. By definition of $\mathsf{pMDP}/{\preceq_\varphi}$, we get that for all $\mathbb{P} \in [\![\varphi^k]\!]$ and $\mathbb{P}' \in [\![\varphi^\ell]\!]$, it holds that $\mathbb{P} \preceq_\varphi \mathbb{P}'$ and vice-versa. Consider the formula $\varphi^k \vee \varphi^\ell$. As this is satisfiable, we have that $\varphi \circ (\varphi^k \vee \varphi^\ell)$ is as well, by (B3). By (Bω), we know that $\varphi \circ (\varphi^k \vee \varphi^\ell) \equiv (\varphi^k \vee \varphi^\ell) \wedge \varphi^m$ for some m. It must be that $m \in \{k, \ell\}$, as otherwise $(\varphi^k \vee \varphi^\ell) \wedge \varphi^m$ would be unsatisfiable by definition of belief partitions. So either $\varphi \circ (\varphi^k \vee \varphi^\ell) \equiv \varphi^k$ or $\varphi \circ (\varphi^k \vee \varphi^\ell) \equiv \varphi^\ell$ - assume w.l.o.g. that it is the former. Since $[\![\varphi \circ \psi]\!] = \min([\![\psi]\!], \preceq_\varphi)$ for all ψ, we thus get in particular that $[\![\varphi^k]\!] = \min([\![\varphi^k \vee \varphi^\ell]\!], \preceq_\varphi) = \min([\![\varphi^k]\!] \cup [\![\varphi^\ell]\!], \preceq_\varphi)$. As we know by definition of belief partitions that $[\![\varphi^k]\!] \neq [\![\varphi^\ell]\!]$, this contradicts our earlier finding that for all $\mathbb{P} \in [\![\varphi^k]\!]$ and $\mathbb{P}' \in [\![\varphi^\ell]\!]$, it holds that $\mathbb{P} \preceq_\varphi \mathbb{P}'$ and vice-versa. □

4.2 Intention Revision

Unlike belief revision operators, intention revision operators do not act solely on intentions, but actually perform revision in the context of strong beliefs.

Definition 13 (Intention revision operator)**.** An *intention revision operator* is a function $\bullet \colon \mathsf{Int} \times \mathsf{Bel} \times \mathsf{Int} \to \mathsf{Int}$, with $\Phi \bullet_\varphi \Psi$ being read as the result of *revising Φ by Ψ under the strong beliefs φ*.

There are a couple of desiderata for an intention revision operator $\bullet$. First, since we are *revising* Φ with a new intention Ψ, the outcome of revision $\Phi \bullet_\varphi \Psi$ should in some sense 'entail' Ψ. Or at least, we should *believe* (given our strong beliefs φ) that Ψ is achieved whenever $\Phi \bullet_\varphi \Psi$ is achieved. In other words, Ψ should be *weakly believed* to hold given φ and $\Phi \bullet_\varphi \Psi$. As we may want to be able to neglect outcomes produced by policies that achieve $\Phi \bullet_\varphi \Psi$ with negligible probability, we again speak of θ-weak beliefs per our discussion in Section 3. We follow the traditional AGM approach of revision, which always accepts the new datum we are revising by.

Fixing a θ throughout this section, our desideratum becomes the following postulate:

(I1) $\Psi \in \mathsf{WB}_\theta(\varphi, \Phi \bullet_\varphi \Psi)$.

Second, if $\Phi \wedge \Psi$ and φ are θ-coherent, then revising Φ by Ψ w.r.t. φ should produce something whose achievement is 'equivalent' to that of $\Phi \wedge \Psi$: we only drop intentions when necessary. Similar considerations apply as with the first desideratum: the equivalence need only be up to the agent's weak beliefs. We formalize this as stating that $\Phi \wedge \Psi \in \mathsf{WB}_\theta(\varphi, \Phi \bullet_\varphi \Psi)$ and $\Phi \bullet_\varphi \Psi \in \mathsf{WB}_\theta(\varphi, \Phi \wedge \Psi)$ - or equivalently, the following postulate:

(I2) If φ and $\Phi \wedge \Psi$ are θ-coherent, then $\mathsf{WB}_\theta(\varphi, \Phi \bullet_\varphi \Psi) = \mathsf{WB}_\theta(\varphi, \Phi \wedge \Psi)$.

Third, intention revision should produce actionable intentions when possible. If φ and Ψ are θ-coherent, then φ and $\Phi \bullet_\varphi \Psi$ should also be. Our postulate becomes:

(I3) If φ and Ψ are θ-coherent, then φ and $\Phi \bullet_\varphi \Psi$ are θ-coherent.

At this point it is first important to note that our desiderata so far are conceptually highly similar to the postulates (B1)-(B3) of belief revision. This similarity can in fact be made precise: we define a new 'entailment' relation on path formulas, $\vDash^\theta_\varphi$, given by putting $\Phi \vDash^\theta_\varphi \Psi$ iff $\Psi \in \mathsf{WB}_\theta(\varphi, \Phi)$, and $\Phi \cong^\theta_\varphi \Psi$ iff $\mathsf{WB}_\theta(\varphi, \Phi) = \mathsf{WB}_\theta(\varphi, \Psi)$. The desiderata are then precisely the postulates (B1)-(B3), with $\vDash^\theta_\varphi$ replacing $\Vdash$, $\cong^\theta_\varphi$ replacing $\equiv$, and θ-coherence w.r.t. φ replacing satisfiability. The next three postulates are similarly analogues of (B4-B6).

(I4) If $\mathsf{WB}_\theta(\varphi, \Phi_1) = \mathsf{WB}_\theta(\varphi, \Phi_2)$ and $\mathsf{WB}_\theta(\varphi, \Psi_1) = \mathsf{WB}_\theta(\varphi, \Psi_2)$, then $\mathsf{WB}_\theta(\varphi, \Phi_1 \bullet_\varphi \Psi_1) = \mathsf{WB}_\theta(\varphi, \Phi_2 \bullet_\varphi \Psi_2)$.

(I5) $\Phi \bullet_\varphi (\Psi \wedge \Omega) \in \mathsf{WB}_\theta((\Phi \bullet_\varphi \Psi) \wedge \Omega)$.

(I6) If φ and $(\Phi \bullet_\varphi \Psi) \wedge \Omega$ are θ-coherent, then $(\Phi \bullet_\varphi \Psi) \wedge \Omega \in \mathsf{WB}_\theta(\Phi \bullet_\varphi (\Psi \wedge \Omega))$.

The postulate (I4) is stating that intention revision is syntax-invariant, though again up to the agent's weak beliefs. Postulates (I5) and (I6), similar to (B5) and (B6) play a role in ensuring the representation of Theorem 4 goes through.

Example 5. We consider Theo's original intention Φ_1 of attending the concert tomorrow, and revise with the intention Φ_2 of buying a drink tomorrow. The postulates again tell us some things about the revision $\Phi_1 \bullet_{\varphi_1} \Phi_2$. By (I3), we have that since Φ_2 is 0.05-coherent w.r.t. φ_1, it must be that $\Phi_1 \bullet_{\varphi_1} \Phi_2$ is also 0.05-coherent w.r.t. φ_1. Therefore $\Phi_1 \bullet_{\varphi_1} \Phi_2 \not\cong^{0.05}_{\varphi_1} \Phi_1 \wedge \Phi_2$ (as we know $\Phi_1 \wedge \Phi_2$ not to be 0.05-coherent w.r.t. φ_1), and since $\cong \subseteq \cong^{0.05}_{\varphi_1}$, we get that $\Phi_1 \bullet_{\varphi_1} \Phi_2 \not\cong \Phi_1 \wedge \Phi_2$, i.e. the revision must drop some of his intentions. But whatever it drops, (I1) tells us that we still have $\Phi_2 \in \mathsf{WB}_{0.05}(\varphi_1, \Phi_1 \bullet_{\varphi_1} \Phi_2)$ - Theo 0.05-weakly believes he will buy drinks tomorrow. And $\Phi_1 \notin \mathsf{WB}_{0.05}(\varphi_1, \Phi_1 \bullet_{\varphi_1} \Phi_2)$, as we saw that $\Phi_1 \bullet_{\varphi_1} \Phi_2 \not\cong^{0.05}_{\varphi_1} \Phi_1 \wedge \Phi_2$: the revised intention does not contain the intention of attending the concert.

Now, we start on a representation theorem for intention revision. Instead of intention revision using orderings on paths, like one might think, we need to take the strong beliefs φ and the threshold θ into account in order to produce θ-coherent intentions. The orderings are on what we call *θ-bundles*, which are to be thought of as sufficiently (according to θ) likely outcomes of a policy.

Definition 14 (θ-bundles)**.** A *θ-bundle* in a pMDP $\mathbb{P} = (\mathbb{M}, s)$ is a triple $\mathbb{B} = (\mathbb{M}, \pi, W)$ consisting of a $\mathbb{P}$-policy π and a set $W \subseteq \mathsf{Paths}(\pi)$ such that $\mu_\pi(W) > \theta$. The set of θ-bundles in $\mathbb{P}$ is denoted $\mathsf{Bun}_\theta(\mathbb{P})$. For $\varphi \in \mathsf{Bel}$, we write $\mathsf{Bun}_\theta(\varphi) = \bigcup_{\mathbb{P} \in [\![\varphi]\!]} \mathsf{Bun}_\theta(\mathbb{P})$.

Now, for a θ-bundle $\mathbb{B} = (\mathbb{M}, \pi, W)$ and $\Phi \in \mathsf{Int}$, we define $\mathbb{B} \vDash \Phi$ iff $\mathbb{M}, \pi, \mathbf{w} \Vdash \Phi$ for all $\mathbf{w} \in W$. The seemingly odd notation is justified by the following observation: $\Psi \in \mathsf{WB}_\theta(\varphi, \Phi)$ iff for all $\mathbb{B} \in \mathsf{Bun}_\theta(\varphi)$ it holds that $\mathbb{B} \vDash \Phi$ implies $\mathbb{B} \vDash \Psi$. In other words, we can think of our earlier 'entailment relation' $\vDash^\theta_\varphi$ of path formulas (defined as $\Phi \vDash^\theta_\varphi \Psi$ iff $\Psi \in \mathsf{WB}_\theta(\varphi, \Phi)$) as actually being derived in the standard way from the *satisfaction relation* $\vDash \subseteq \mathsf{Bun}_\theta(\varphi) \times \mathsf{Int}$.

This means that, fixing $\varphi \in \mathsf{Bel}$, PLBP+ path formulas $\Psi \in \mathsf{Int}$ and θ-bundles $\mathbb{B} \in \mathsf{Bun}_\theta(\varphi)$ form an instance of *base logics* in the sense of [Falakh *et al.*, 2023], in which θ-coherence w.r.t. φ is the notion of satisfiability. We can utilize this powerful fact to give definitions and a representation theorem that *fully mirror* those we gave for belief revision: the definitions and theorems here are the same as in Section 4.1, in which occurences of $\Vdash$ are replaced by $\vDash^\theta_\varphi$ and pMDPs are replaced by θ-bundles $\mathbb{B} \in \mathsf{Bun}_\theta(\varphi)$.

In the following, we write $[\Phi]^\theta_\varphi$ for the set of all $\mathbb{B} \in \mathsf{Bun}_\theta(\varphi)$ such that $\mathbb{B} \Vdash \Phi$.

Definition 15 (θ-bundle assignments)**.** A *θ-bundle assignment* is a Bel-indexed family $\preceq_{(-,-)}$ of functions $\preceq_{\varphi,(-)} : \mathsf{Int} \to 2^{\mathsf{Bun}_\theta(\varphi) \times \mathsf{Bun}_\theta(\varphi)}$ such that $\preceq_{\varphi,\Phi}$ is a total preorder for all Φ. The assignment is *faithful* if: (i) If $\mathbb{B} \vDash \Phi$ and $\mathbb{B}' \vDash \Phi$, then $\mathbb{B} \preceq_{\varphi,\Phi} \mathbb{B}'$ and $\mathbb{B}' \preceq_{\varphi,\Phi} \mathbb{B}$. (ii) If $\mathbb{B} \vDash \Phi$ and $\mathbb{B}' \nvDash \Phi$, then $\mathbb{B} \prec_{\varphi,\Phi} \mathbb{B}'$. (iii) If $\Phi \cong^\theta_\varphi \Psi$, then $\preceq_{\varphi,\Phi} = \preceq_{\varphi,\Psi}$.

The assignment is *finitely expressible* if $\preceq_{\varphi,\Phi}$ is finitary for all φ and Φ, and for all φ and Φ there exist $\Phi^{\varphi,1}, \ldots, \Phi^{\varphi,\mathsf{height}(\preceq_{\varphi,\Psi})} \in \mathsf{Int}$ such that $[\Phi^{\varphi,k}]^\theta_\varphi = (\mathsf{Bun}_\theta(\varphi))_{\preceq_{\varphi,\Phi},k}$, i.e. the k-th level set of $\mathsf{Bun}_\theta(\varphi)$ according to $\preceq_{\varphi,\Phi}$. We refer to finitely expressible and faithful θ-bundle assignments as *FEF θ-bundle assignments* for short.

We have an analogue of Proposition 1, showing how FEF θ-bundle assignments are precisely *θ-intention weakening maps*, assigning to every φ and Φ a sequence of path formulas $\langle\Phi\rangle^\varphi_1 \vDash^\theta_\varphi \cdots \vDash^\theta_\varphi \langle\Phi\rangle^\varphi_n$ such that $\langle\Phi\rangle^\varphi_1 \cong^\theta_\varphi \Phi$. And using a new postulate (Iω), we can give a representation theorem for FEF θ-bundle assignments. In the following, a *θ-intention partition for* φ is a sequence $\Phi^1, \ldots, \Phi^n \in \mathsf{Int}$ such that (i) $\Phi^1 \vee \cdots \vee \Phi^n$ is a tautology w.r.t. $\vDash^\theta_\varphi$, (ii) φ and Φ^k are θ-coherent for all k, and (iii) φ and $\Phi^k \wedge \Phi^\ell$ are not θ-coherent for all $k \neq \ell$.

(Iω) For all φ and Φ there exists a θ-intention partition $\Phi^1, \ldots, \Phi^n$ w.r.t. φ, such that for all Ψ, there exists k with $\Phi \bullet_\varphi \Psi \cong^\theta_\varphi \Psi \wedge \Phi^k$.

Theorem 4 (Representation for intention revision)**.** *The following statements hold:*

- *For every FEF θ-bundle assignment $\preceq_{(-,-)}$ there exists an intention revision operator $\bullet$ satisfying (I1)-(I6) and (Iω), such that $[\Phi \bullet_\varphi \Psi]^\theta_\varphi = \min([\Psi]^\theta_\varphi, \preceq_{\varphi,\Phi})$.*
- *For every intention revision operator $\bullet$ satisfying (I1)-(I6) and (Iω) there exists an FEF θ-bundle assignment $\preceq_{(-,-)}$ such that $[\Phi \bullet_\varphi \Psi]^\theta_\varphi = \min([\Psi]^\theta_\varphi, \preceq_{\varphi,\Phi})$.*

Proof sketch. The proof is identical to that of Theorem 3: for every φ, we now consider path formulas with the entailment relation $\vDash^\theta_\varphi$ to form a base logic in the sense of [Falakh *et al.*, 2023], interpreted over θ-bundles in $\mathsf{Bun}_\theta(\varphi)$. □

We conclude this section by stating that intention revision operators satisfying (I1)-(I6) and (Iω) are computable if the underlying θ-intention weakening maps are computable, by a conceptually identical algorithm as for belief operators. To compute $\Phi \bullet_\varphi \Psi$, first check if φ and Ψ are θ-coherent: if not, output $\bot$. Then compute $\langle\Phi\rangle^\varphi_1, \dots, \langle\Phi\rangle^\varphi_n$, and iterate over $1 \leqslant k \leqslant n$, checking whether φ and $\Psi \wedge \langle\Phi\rangle^\varphi_k$ are θ-coherent. For the first such k, output $\Psi \wedge \langle\Phi\rangle^\varphi_k$. If no such k exists, output Ψ.

4.3 Joint Revision of Beliefs and Intentions

In this section, we present the fundamental interaction between beliefs and intentions that motivates the current work: revision of beliefs may trigger revision of intentions, and vice-versa. We capture this interaction through the following definition in which joint revision proceeds by first revising strong beliefs, and then revising intentions w.r.t. the revised strong beliefs.

Definition 16 (Belief-intention revision operator)**.** Given a belief revision operator $\circ$ and an intention revision operator $\bullet$, their *belief-intention revision operator* is the function $*\colon (\mathsf{Bel} \times \mathsf{Int})^2 \to \mathsf{Bel} \times \mathsf{Int}$ defined as $(\varphi, \Phi) * (\psi, \Psi) := (\varphi \circ \psi, \Phi \bullet_{\varphi\circ\psi} \Psi)$.

Note that we define the joint revision operator through independent belief and intention revision operators. We could have taken joint revision operators as primitive and provided postulates for them capturing the desired interplay, but the definition above simplifies presentation and follows earlier lines of work [van Zee *et al.*, 2020].

In joint revision, belief revision may trigger intention revision, since the intention revision's notion of coherence depends on the new strong belief $\varphi\circ\psi$. And vice-versa, intention revision triggers belief revision, though key is the separation between strong and weak beliefs: only the latter depend on intentions, as seen in Definition 7. As argued earlier, intention revision should not change *strong* beliefs.

The postulates for belief and intention revision tell us certain properties of these joint belief-intention revision operators, which we illustrate through our running example.

Example 6 (Running example, continued)**.** First, note that intention revision changes weak beliefs. We saw in Example 3 that $\mathsf{X}\neg\mathsf{five} \in \mathsf{WB}_{0.05}(\varphi_1, \Phi_1)$: when intending to attend the concert, Theo 0.05-weakly believes he will be out of money tomorrow. Now consider the revision $(\varphi_1, \Phi_1) * (\top, \Phi_2) = (\varphi_1, \Phi_3)$ (i.e. we just wish to revise with the intention to buy a drink). By the postulates for intention revision, we get that $\mathsf{Xfive} \in \mathsf{WB}_{0.05}(\varphi_1, \Phi_3)$, since $\Phi_3 \vDash^{0.05}_\varphi \Phi_2 \vDash^{0.05}_\varphi \mathsf{Xfive}$. So after revision, we weakly believe that Theo *does* have \$5 tomorrow: otherwise, he could not buy a drink tomorrow.

Now, note that belief revision actually changes intentions. Consider the revision $(\varphi_1, \Phi_1) * (\varphi_2, \top) = (\varphi_1 \circ \varphi_2, \Phi_1 \bullet_{\varphi_1\circ\varphi_2} \top)$. We know that since φ_2 is satisfiable, $\varphi_1 \circ \varphi_2$ must also be by our belief revision postulates. So we have that $\top$ is 0.05-coherent w.r.t. $\varphi_1 \circ \varphi_2$ ($\top$ always is as long as the strong beliefs are satisfiable). We can derive by our intention revision postulates that $\Phi_1 \bullet_{\varphi_1\circ\varphi_2} \top$ must also be 0.05-coherent w.r.t. $\varphi_1 \circ \varphi_2$. But we know that φ_2 and Φ_1 are not 0.05-coherent. Therefore, it cannot be that $\Phi_1 \bullet_{\varphi_1\circ\varphi_2} \top \vDash^{0.05}_{\varphi_1} \Phi_1$. This means that by revising with the belief that the band is popular, Theo is forced to abandon his intention to attend the concert.

5 Conclusion

We have initiated the study of belief and intention revision in stochastic environments, expressed in an appropriate probabilistic temporal logic. We give rationality postulates for revision, and prove representation theorems both for revision operators in general, as well for operators that are generally computable using the novel postulates (Bω) and (Iω).

In future work, we plan to extend our framework to allow for *iterated* revision, à la [Darwiche and Pearl, 1997]. We also wish to develop a representation theorem with a modified version of (I1), which would allow success of revision only given coherent new intentions, as is also the case in the work of [van Zee *et al.*, 2020]. Such revision could follow the approach as taken in the *consistent AGM revision* of [Hansson, 2023a; Hansson, 2023b], where revision with inconsistent statements maintains the original beliefs, or the more nuanced approach of [Konieczny *et al.*, 2010], in which revision by a formula does not need to immediately entail it, but instead only increase its plausibility every time it is revised with.

References

[Alchourrón *et al.*, 1985] Carlos E. Alchourrón, Peter Gärdenfors, and David Makinson. On the logic of theory change: Partial meet contraction and revision functions. *The Journal of Symbolic Logic*, 50(2):510–530, 1985.

[Bratman, 1987] Michael Bratman. *Intention, Plans, and Practical Reason*. Cambridge, MA: Harvard University Press, Cambridge, 1987.

[Canny, 1988] John Canny. Some algebraic and geometric computations in PSPACE. In *Proceedings of the Twentieth Annual ACM Symposium on Theory of Computing STOC'88*, page 460–467. ACM, 1988.

[Castelfranchi and Paglieri, 2007] Cristiano Castelfranchi and Fabio Paglieri. The role of beliefs in goal dynamics: prolegomena to a constructive theory of intentions. *Synthese*, 155(2):237–263, 2007.

[Chang and Keisler, 2012] Chen Chung Chang and Howard Jerome Keisler. *Model Theory*. Dover Publications, 3 edition, 2012.

[Cohen and Levesque, 1990] Philip R Cohen and Hector J Levesque. Intention is choice with commitment. *Artificial intelligence*, 1990.

[Darwiche and Pearl, 1997] Adnan Darwiche and Judea Pearl. On the logic of iterated belief revision. *Artificial Intelligence*, 89(1):1–29, 1997.

[De Giacomo and Vardi, 2013] Giuseppe De Giacomo and Moshe Y. Vardi. Linear temporal logic and linear dynamic logic on finite traces. In *Proceedings of the Twenty-Third International Joint Conference on Artificial Intelligence*, IJCAI '13, page 854–860. AAAI Press, 2013.

[Delgrande *et al.*, 2018] James P. Delgrande, Pavlos Peppas, and Stefan Woltran. General belief revision. *J. ACM*, 65(5), sep 2018.

[Falakh *et al.*, 2023] Faiq Miftakhul Falakh, Sebastian Rudolph, and Kai Sauerwald. AGM belief revision, semantically, 2023. arXiv:2112.13557.

[Grant *et al.*, 2010] John Grant, Sarit Kraus, Donald Perlis, and Michael J. Wooldridge. Postulates for revising BDI structures. *Synth.*, 175(Supplement-1):39–62, 2010.

[Hansson, 2023a] Sven Ove Hansson. A basis for AGM revision in bayesian probability revision. *J. Philos. Log.*, 52(6):1535–1559, 2023.

[Hansson, 2023b] Sven Ove Hansson. Iterated AGM revision based on probability revision. *J. Log. Lang. Inf.*, 32(4):657–675, 2023.

[Icard *et al.*, 2010] Thomas Icard, Eric Pacuit, and Yoav Shoham. Joint revision of belief and intention. In *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning*, KR'10, page 572–574. AAAI Press, 2010.

[Katsuno and Mendelzon, 1991] Hirofumi Katsuno and Alberto O. Mendelzon. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3):263–294, 1991.

[Khan and Lespérance, 2012] Shakil M. Khan and Yves Lespérance. Logical foundations for a rational BDI agent programming language (extended version). In Louise A. Dennis, Olivier Boissier, and Rafael H. Bordini, editors, *Programming Multi-Agent Systems - 9th International Workshop, ProMAS 2011*, volume 7217 of *Lecture Notes in Computer Science*, pages 3–21. Springer, 2012.

[Konieczny *et al.*, 2010] Sébastien Konieczny, Mattia Medina Grespan, and Ramón Pino Pérez. Taxonomy of improvement operators and the problem of minimal change. In *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning*, KR'10, page 161–170. AAAI Press, 2010.

[Leask, 2021] Sam Leask. *GROVE - a computationally grounded model for rational intention revision in BDI agents*. PhD thesis, University of Nottingham, UK, 2021.

[Lorini and Herzig, 2008] Emiliano Lorini and Andreas Herzig. A logic of intention and attempt. *Synthese*, 163(1):45–77, 2008.

[McDermott, 1982] Drew McDermott. A temporal logic for reasoning about processes and plans. *Cognitive Science*, 6(2):101–155, 1982.

[Motamed *et al.*, 2023] Nima Motamed, Natasha Alechina, Mehdi Dastani, Dragan Doder, and Brian Logan. Probabilistic temporal logic for reasoning about bounded policies. In Edith Elkind, editor, *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, pages 3296–3303, August 2023. Main Track.

[Rao and Georgeff, 1991] Anand S. Rao and Michael P. Georgeff. Modeling rational agents within a bdi-architecture. In *Proceedings of the 2nd International Conference on Principles of Knowledge Representation and Reasoning (KR'91)*, pages 473–484, 1991.

[Segerberg, 1999] Krister Segerberg. *Two Traditions in the Logic of Belief: Bringing them Together*, pages 135–147. Springer Netherlands, Dordrecht, 1999.

[Shapiro *et al.*, 2012] Steven Shapiro, Sebastian Sardiña, John Thangarajah, Lawrence Cavedon, and Lin Padgham. Revising conflicting intention sets in BDI agents. In Wiebe van der Hoek, Lin Padgham, Vincent Conitzer, and Michael Winikoff, editors, *International Conference on Autonomous Agents and Multiagent Systems, AAMAS 2012*, pages 1081–1088, 2012.

[Shoham, 2009] Yoav Shoham. Logical theories of intention and the database perspective. *Journal of Philosophical Logic*, 38(6):633–647, 2009.

[van der Hoek *et al.*, 2007] Wiebe van der Hoek, Wojciech Jamroga, and Michael Wooldridge. Towards a theory of intention revision. *Synthese*, 155(2):265–290, 2007.

[van Ditmarsch *et al.*, 2011] Hans van Ditmarsch, Tiago de Lima, and Emiliano Lorini. Intention change via local assignments. In Mehdi Dastani, Amal El Fallah Seghrouchni, Jomi Hübner, and João Leite, editors, *Languages, Methodologies, and Development Tools for Multi-Agent Systems*, pages 136–151, Berlin, Heidelberg, 2011. Springer.

[van Zee and Doder, 2016] Marc van Zee and Dragan Doder. Agm-style revision of beliefs and intentions. In *ECAI*, volume 285 of *Frontiers in Artificial Intelligence and Applications*, pages 1511–1519. IOS Press, 2016.

[van Zee *et al.*, 2020] Marc van Zee, Dragan Doder, Leendert van der Torre, Mehdi Dastani, Thomas Icard, and Eric Pacuit. Intention as commitment toward time. *Artificial Intelligence*, 283, 2020.

[Wooldridge, 2000] Michael John Wooldridge. *Reasoning about Rational Agents*. MIT Press, 2000.