

Rational Revision of Group Intentions

Nima Motamed¹, Natasha Alechina^{2,1}, Mehdi Dastani¹, Dragan Doder¹

¹Utrecht University, The Netherlands,

²Open University, The Netherlands

{n.motamed, n.a.alechina, m.m.dastani, d.doder}@uu.nl

Abstract

In systems such as group calendars or collaborative platforms, agents make group commitments to future actions that must adapt as new facts or constraints emerge. We develop a formal framework for revising such group intentions in systems where coalitions adopt shared, temporally extended intentions represented in a logic based on Alternating-Time Temporal Logic with strategy contexts. After formulating coherence criteria for systems of group intentions, we establish representation theorems in the style of Katsuno and Mendelzon, showing that revision operators satisfy rationality postulates precisely when they can be represented by preorders on strategy profiles. These results extend classical revision theory by covering non-total preorders and a logic of higher expressive power. Altogether, the framework lays the groundwork for principled revision of group intentions in systems where both coordination and change are essential.

1 Introduction

Intentions, as commitments to future courses of action, lie at the heart of planning and coordination in autonomous systems. When an agent acquires new information, changes their priorities, or reasons further about their abilities, their existing intentions may need to be revised. Having a clear account of how those revisions should proceed is essential for designing agents whose behaviour remains coherent and predictable rather than ad-hoc or brittle.

In the single-agent setting, Shoham (2009) conceptualises this problem by treating intentions as entries in a personal database of commitments. From that perspective, revising intentions is a matter of inserting, deleting, or rearranging database entries while preserving overall consistency. A line of work has developed frameworks for revision of such single-agent intention databases in both deterministic and stochastic settings (Icard, Pacuit, and Shoham 2010; Van Zee et al. 2020; Motamed et al. 2024), adapting techniques from classical belief revision theory and providing postulates for rational revision.

When several agents collaborate, intentions no longer concern isolated courses of action but shared commitments that bind whole groups. A logistics network, for instance,

succeeds only if warehouses, vehicles, and routing algorithms collectively intend a feasible delivery schedule; a swarm of exploration drones must agree not merely on individual waypoints but on a combined trajectory that satisfies energy and safety constraints for the group as a whole. In such multiagent settings, changes to one group’s intentions can ripple through to other groups. Accordingly, the important task is not merely the revision of one group’s isolated intention, but rather the revision of an interconnected system of intentions for different groups of agents, in a way that leaves the resulting system still feasibly executable and consistent with environmental constraints.

There is existing work on re-planning, revising, scheduling, and recognising intentions in a multiagent setting, e.g. (Cawsey et al. 1993; Dunin-Keplicz and Verbrugge 2002; Zhang, Nguyen, and Kowalczyk 2007; Dann et al. 2023; Zhang et al. 2023). But a principled framework that describes how intention revision methods should operate, in particular with interconnected systems of group intentions, is still lacking. **In this paper** we close the gap by introducing such a framework that (i) represents the intentions of groups of agents in a logic based on Alternating-Time Temporal Logic with strategy contexts (Brihaye et al. 2009), (ii) defines two coherence notions that capture when a system of intentions is jointly consistent and adheres to constraints, and (iii) specifies rational revision operators, axiomatised by postulates and characterised by a representation theorem in the style of Katsuno and Mendelzon (1991). This theorem is of particular interest, as it provides a semantic grounding for intention revision in terms of *minimal change* via preorders on strategy profiles, and uses an innovative proof method, adding to the literature on revision theory.

2 Logical Preliminaries

Before we define our notion of group intention, we need to first introduce the logic in which we formulate and interpret them. We choose a fragment of *alternating-time temporal logic with strategy contexts* (Brihaye et al. 2009) only using the next-time operator X as temporal operator, and enriched with primitive formulas expressing the execution of actions. We refer to this logic as $XATL_{sc,do}$. Similar extensions of temporal logics with explicit actions were also pursued in other work (Herzig, Lorini, and Walther 2013; Motamed et al. 2023). The chosen logic possesses precisely the

expressive power needed to express and reason about group intentions.

We fix a finite set Ag of agents, a finite set Act of actions, and a finite set Prop of propositional variables. We refer to $\alpha = (\alpha^i)_{i \in \text{Ag}} \in \text{Act}^{\text{Ag}}$ as *joint actions*, containing an action $\alpha^i \in \text{Act}$ for each agent $i \in \text{Ag}$. Moreover, we use the following bits of notation. Given $A \subseteq \text{Ag}$, we write $\bar{A} = \text{Ag} \setminus A$. And given a partial function $f: X \rightarrow Y$ and $x \in X$, we write $f(x) \downarrow$ to denote that $f(x)$ is defined. We denote by $\text{dom}(f)$ the set of all $x \in X$ for which $f(x) \downarrow$.

Definition 2.1 (Concurrent game structure). A *concurrent game structure* (CGS) is a tuple $\mathbb{G} = (S, M, T, L)$, where S is a finite set of states, $M: S \times \text{Ag} \rightarrow 2^{\text{Act}} \setminus \{\emptyset\}$ is the availability function specifying which actions are available to which agent at which state, $T: S \times \text{Act}^{\text{Ag}} \rightarrow S$ is a partial transition function, and $L: S \rightarrow 2^{\text{Prop}}$ is a labelling by propositional atoms. We require that $T(s, \alpha) \downarrow$ iff $\alpha^i \in M(s, i)$ for all $i \in \text{Ag}$.

As is standard with the semantics of ATL, transitions are only defined for joint actions that are available. This is important for our applications: reasoning about the executability of intentions is an essential part of their revision.

A *path* from state s is an infinite sequence $\pi = s_0 \alpha_1 s_1 \dots \in (S; \text{Act}^{\text{Ag}})^\omega$ of states and joint actions, such that $s_0 = s$, $T(s_n, \alpha_{n+1}) \downarrow$ and $T(s_n, \alpha_{n+1}) = s_{n+1}$ for all $n < \omega$. A *history* from s is defined analogously as a *finite* sequence $\eta \in S; (\text{Act}^{\text{Ag}}; S)^*$ of states and joint actions, ending in a state, such that again $s_0 = s$, $T(s_n, \alpha_{n+1}) \downarrow$ and $T(s_n, \alpha_{n+1}) = s_{n+1}$ for all n . We write $\text{Path}(s)$ and $\text{Hist}(s)$ for sets of paths and histories from s , respectively.

In a CGS, agents can adopt strategies that independently and concurrently pick actions to perform at each state, given the history of states and joint actions. We refer to a collection of such strategies for a subset of agents as a *strategy context*.

Definition 2.2 (Strategy context, profile, outcome). A (strategy) *context* from s is a partial function $\sigma: \text{Ag} \rightarrow \text{Act}^{\text{Hist}(s)}$, such that for all $i \in \text{dom}(\sigma)$ we have that $\sigma(i)(s_0 \alpha_1 \dots s_n) \in M(s_n, i)$. We write $\text{Cxt}_A(s)$ for the set of strategy contexts σ from s with $\text{dom}(\sigma) = A$. If $\text{dom}(\sigma) = \text{Ag}$, then we refer to σ as a *strategy profile*, and write $\text{Prof}(s) = \text{Cxt}_{\text{Ag}}(s)$ for the set of all strategy profiles from s .

An *outcome* of a strategy context σ from s is a path $s_0 \alpha_1 s_1 \dots \in \text{Path}(s)$ such that for all $n < \omega$ and $i \in \text{dom}(\sigma)$, it holds that $\sigma(i)(s_0 \alpha_1 \dots s_n) = \alpha_{n+1}^i$. The set of outcomes of σ is denoted by $\text{Out}(\sigma)$. Note that if σ is a strategy profile, then necessarily $\text{Out}(\sigma)$ is a singleton, and we then refer to the unique path in $\text{Out}(\sigma)$ as $\text{out}(\sigma)$.

Strategy contexts only specify the actions taken for a subset of agents. The defining feature of ATL with strategy contexts is the ability to quantify over ways to extend contexts with strategies for other agents. This process of extension is referred to as *merging*.

Definition 2.3 (Merge). The *merge* of strategy contexts σ and σ' is the strategy context $\sigma \times \sigma'$ defined by letting $(\sigma \times \sigma')(i) = \sigma'(i)$ if $i \in \text{dom}(\sigma')$, and $(\sigma \times \sigma')(i) = \sigma(i)$ if $i \in \text{dom}(\sigma) \setminus \text{dom}(\sigma')$. If $\text{dom}(\sigma) \cap \text{dom}(\sigma') = \emptyset$ and

$\text{dom}(\sigma) \cup \text{dom}(\sigma') = \text{Ag}$, we also write (σ, σ') to denote the merge $\sigma \times \sigma'$.

In other words, in the merge $\sigma \times \sigma'$, we overwrite all agents' strategies in σ with those from σ' , except for those agents who do not have a specified strategy in σ' . If the two contexts' domains do not overlap and cover all agents, we write (σ, σ') for the merge, which is borrowed from game theory where it is commonly used in similar ways.

We can now finally start introducing $\text{XATL}_{\text{sc}, \text{do}}$. The syntax consists of *context formulas* φ (interpreted over contexts) and *path formulas* Φ (interpreted over paths), defined by the following grammar.

$$\begin{aligned} \varphi &::= p \mid \varphi \vee \varphi \mid \neg \varphi \mid \langle A \rangle \varphi \mid \langle A \rangle \varphi \mid \Diamond \Phi \\ \Phi &::= \varphi \mid \text{do}_{i,a} \mid \Phi \vee \Phi \mid \neg \Phi \mid X\Phi. \end{aligned}$$

Here, $p \in \text{Prop}$, $A \subseteq \text{Ag}$, $i \in \text{Ag}$ and $a \in \text{Act}$.

The operator $\langle A \rangle$ quantifies existentially over strategy contexts defined for A and merges them with the current context. The operator $\langle A \rangle$ releases the context for A : it forgets the strategies of agents in A . The operator $\Diamond$ quantifies existentially over the outcomes of the current context. Path formulas are interpreted with X being the next-time operator which moves to the next point in a path (and moves the strategy context accordingly). The proposition $\text{do}_{i,a}$ says that agent i does action a at the current time on a path.

We can define other Boolean connectives as abbreviations in the standard way. Furthermore, we can define the duals $[A] = \neg \langle A \rangle \neg$ and $\Box = \neg \Diamond \neg$.

Remark 2.4. Note that our presentation of the syntax deviates slightly from the original presentation of ATL with strategy contexts, in that we have separate operators $\langle A \rangle$ (for quantification over contexts) and $\Diamond$ (for quantification over outcomes). This is innocent, as these separate operators are also definable in the original syntax of ATL with strategy contexts.

As examples of formulas, consider $\langle \text{Ag} \rangle \Diamond X \text{do}_{i,a}$, saying some profile has an outcome where i does a next, and $\langle \text{Ag} \rangle \bigwedge_i (\{i\}) \Box X p$, saying there is a strategy profile achieving p next, immune to unilateral deviation.

We now define the formal semantics of our logic $\text{XATL}_{\text{sc}, \text{do}}$. Given a CGS $\mathbb{G}$, a state $s \in \mathbb{G}$, a context σ from s and a path $\pi \in \text{Path}(s)$, we specify through simultaneous induction over context formulas φ and path formulas Φ the satisfaction relations $\mathbb{G}, s, \sigma \models \varphi$ and $\mathbb{G}, \pi, \sigma \models \Phi$.

A bit of notation is required before the semantics of context formulas: given a partial function $f: X \rightarrow Y$ and $X' \subseteq X$, we write $f|_{X'}$ for the *restriction* of f to X' : the partial function $f|_{X'}: X' \rightarrow Y$ with $\text{dom}(f|_{X'}) = \text{dom}(f) \cap X'$ and $f|_{X'}(x) = f(x)$ for $x \in \text{dom}(f|_{X'})$.

$$\begin{aligned} \mathbb{G}, s, \sigma \models p &\iff p \in L(s) \\ \mathbb{G}, s, \sigma \models \varphi \vee \psi &\iff \mathbb{G}, s, \sigma \models \varphi \text{ or } \mathbb{G}, s, \sigma \models \psi \\ \mathbb{G}, s, \sigma \models \neg \varphi &\iff \mathbb{G}, s, \sigma \not\models \varphi \\ \mathbb{G}, s, \sigma \models \langle A \rangle \varphi &\iff \exists \sigma' \in \text{Cxt}_A(s): \mathbb{G}, s, \sigma \times \sigma' \models \varphi \\ \mathbb{G}, s, \sigma \models \langle A \rangle \varphi &\iff \mathbb{G}, s, \sigma|_{\bar{A}} \models \varphi \\ \mathbb{G}, s, \sigma \models \Diamond \Phi &\iff \exists \pi \in \text{Out}(\sigma): \mathbb{G}, \pi, \sigma \models \Phi \end{aligned}$$

For the semantics of path formulas we are again required to introduce some notation: given a context σ from s , as well as a joint action α with $T(s, \alpha) \downarrow$, we define σ_α to be the context from $s_0 = T(s, \alpha)$ obtained by letting $\sigma_\alpha(i)(s_0 \alpha_1 \dots s_n) = \sigma(i)(s \alpha s_0 \alpha_1 \dots s_n)$ (in other words, we ‘push’ σ ahead one step in time via α). Also, given $\pi = s_0 \alpha_1 s_1 \dots \in \text{Path}(s)$ we write $\text{hd}(\pi) = s_0$, $\text{act}(\pi) = \alpha_1$ and $\text{tl}(\pi) = s_1 \alpha_2 \dots \in \text{Path}(s_1)$.

$$\begin{aligned} \mathbb{G}, \pi, \sigma \Vdash \varphi &\iff \mathbb{G}, \text{hd}(\pi), \sigma \Vdash \varphi \\ \mathbb{G}, \pi, \sigma \Vdash \text{do}_{i,a} &\iff \text{act}(\pi)^i = a \\ \mathbb{G}, \pi, \sigma \Vdash \Phi \vee \Psi &\iff \mathbb{G}, \pi, \sigma \Vdash \Phi \text{ or } \mathbb{G}, \pi, \sigma \Vdash \Psi \\ \mathbb{G}, \pi, \sigma \Vdash \neg \Phi &\iff \mathbb{G}, \pi, \sigma \not\Vdash \Phi \\ \mathbb{G}, \pi, \sigma \Vdash X\Phi &\iff \mathbb{G}, \text{tl}(\pi), \sigma_{\text{act}(\pi)} \Vdash \Phi \end{aligned}$$

We will often write $\mathbb{G}, s \Vdash \varphi$ to mean $\mathbb{G}, s, \epsilon \Vdash \varphi$, where ϵ is the empty strategy context (i.e. $\text{dom}(\epsilon) = \emptyset$). Given a strategy profile σ , we will also often write $\mathbb{G}, \sigma \Vdash \Phi$ instead of $\mathbb{G}, \text{out}(\sigma), \sigma \Vdash \Phi$ (the idea being that since strategy profiles uniquely determine a path, we can directly evaluate path formulas on them).

It is simple to see that the satisfiability problem (and therefore also the validity and entailment problems) of $\text{XATL}_{\text{sc}, \text{do}}$ is decidable: the logic possesses a bounded model property, and one can enumerate all models up to a certain size since the sets of agents and actions are fixed. While complexity analysis is beyond the topic of this paper, since our focus is on studying and characterising revision operators, the decidability is conceptually useful as it also allows for checking coherence, as defined in Section 3.

3 Intention Systems and Coherence

Having introduced the syntax and semantics of $\text{XATL}_{\text{sc}, \text{do}}$, we can now spell out what we mean by group intentions in the style of Shoham (2009). We consider each group’s intention to be a path formula, describing actions and outcomes the coalition will perform or bring about. A collection of such formulas for each coalition constitutes what we call a *basic intention system*. The adjective ‘basic’ is used here, as we will later introduce a more expressive construction.

Definition 3.1 (Basic intention system). Let $\text{Coal} = 2^{\text{Ag}} \setminus \{\emptyset\}$ be the set of *coalitions*. A *basic intention system* is a coalition-indexed family $I = \{I_A\}_{A \in \text{Coal}}$ of path formulas I_A . The set of all basic intention systems is denoted BIS.

Note that we can always represent coalitions without intentions via trivial path formulas, i.e. $I_A = \top$ (where $\top$ is any tautology like $p \vee \neg p$).

Our interest lies in the *dynamics* of intention systems. When coalitions adopt new commitments, the resulting intentions must remain coherent with each other and the known environmental facts. Suppose two overlapping coalitions A and B share an agent a and have individually feasible intentions I_A and I_B . If the facts require a to act differently to satisfy I_A than to satisfy I_B , then a cannot do both, and so the overall system becomes incoherent.

To obtain a definition of coherence, it is natural to generalize definitions from the single-agent setting. There, it is

defined, following the *Strong Consistency Principle* of Bratman (1987), as the existence of a model consistent with the agent’s beliefs in which the agent has a strategy that achieves all intended goals (Van Zee et al. 2020; Motamed et al. 2024). In multiagent systems, we replace the single strategy by a *strategy profile*. We say that a profile σ *weakly satisfies* a basic intention system when the outcome of σ fulfills every coalition’s intention; the use of the terminology ‘weakly’ will become clear when we introduce a stronger notion of satisfaction later.

Definition 3.2 (Weak satisfaction & coherence). Given $(\mathbb{G}, s)$ and a strategy profile $\sigma \in \text{Prof}(s)$, a basic intention system I is *weakly satisfied* by σ (denoted $\mathbb{G}, \sigma \Vdash_w I$) if $\mathbb{G}, \sigma \Vdash I_A$ for all $A \in \text{Coal}$. Given a context formula φ encoding *constraints* on the environment, I is *weakly coherent* w.r.t. φ if there exists $(\mathbb{G}, s)$ with $\mathbb{G}, s \Vdash \varphi$ and such that there is $\sigma \in \text{Prof}(s)$ with $\mathbb{G}, \sigma \Vdash_w I$.

While it is an immediate generalisation of single-agent coherence, weak coherence assumes that a coalition can rely on outside agents for their intentions, which is not realistic if the agents do not form one single cooperative entity. To avoid such an assumption, we also provide a stronger satisfaction notion: a coalition’s intentions must be achievable through its own strategic capabilities.

Definition 3.3 (Strong satisfaction & coherence). Given $(\mathbb{G}, s)$ and a strategy profile $\sigma \in \text{Prof}(s)$, a basic intention system I is *strongly satisfied* by σ (denoted $\mathbb{G}, \sigma \Vdash_s I$) if for all $A \in \text{Coal}$ and contexts $\sigma'_A \in \text{Cxt}_A(s)$ it holds that $\mathbb{G}, (\sigma|_A, \sigma'_A) \Vdash I_A$. Given constraints φ , I is *strongly coherent* w.r.t. φ if there exists $(\mathbb{G}, s)$ with $\mathbb{G}, s \Vdash \varphi$ and such that there is $\sigma \in \text{Prof}(s)$ with $\mathbb{G}, \sigma \Vdash_s I$.

In other words, a profile strongly satisfies a basic intention system if every coalition’s part of the profile is *sufficient* for achieving their intention, regardless of what the other agents do. Note that strong coherence implies weak coherence.

Example 3.4. Alice (a) and Bob (b) are finishing up a joint article. Three steps are left before submission: one of the proofs needs to be completed, an introduction has to be written, and the authors need to have a final discussion together. Potentially, agents may also wish to rest. We model this scenario with three actions `writeProof`, `writelnIntro`, and `discuss`, as well as a propositional variable `done` representing that submission is completed.

We describe three possible basic intention systems that the agents may adopt. In the first, I^1 , we have $I_a^1 = \text{do}_{a, \text{writeProof}}$, $I_b^1 = \text{do}_{b, \text{writelnIntro}}$, and $I_{a,b}^1 = \text{XXdone}$. In the second, I^2 , the intentions are $I_a^2 = \text{do}_{a, \text{writeProof}}$, $I_b^2 = \text{do}_{b, \text{writelnIntro}} \wedge X\text{do}_{b, \text{discuss}}$, and $I_{a,b}^2 = \text{XXdone}$. Finally, in the third, I^3 , the intentions are $I_a^3 = \text{do}_{a, \text{writelnIntro}}$, $I_b^3 = \text{do}_{b, \text{writelnIntro}}$, and $I_{a,b}^3 = \text{XXdone}$.

Take as constraints φ the statement that the agents can write the proof and introduction, that discussion requires the proof and introduction to be written, and that the paper is done after discussion. I.e. $\varphi = \Box(X\text{do}_{\text{discuss}} \rightarrow (\text{do}_{\text{writeProof}} \wedge \text{do}_{\text{writelnIntro}}) \wedge \text{XXdone} \leftrightarrow X\text{do}_{\{a,b\}, \text{discuss}})$, where we use abbreviations $\text{do}_x = \text{do}_{a,x} \vee \text{do}_{b,x}$ and

$\text{do}_{\{a,b\},x} = \text{do}_{a,x} \wedge \text{do}_{b,x}$ for actions x . We then see that I^1 is strongly and, as a consequence, weakly, coherent w.r.t. φ : there is a strategy profile (in some CGS) in which Alice writes the proof, Bob writes the introduction, and they then discuss to finish, with every coalition's strategies being sufficient for their intentions. However, for I^2 , we only get weak coherence: Bob intends to write the intro and discuss after, even though that requires Alice to first write the proof, which Bob cannot unilaterally enforce. Finally, for I^3 we also lose weak coherence, since the agents cannot finish the paper if no one writes the proof.

Strong satisfaction removes the cooperation assumptions made by its weaker counterpart by preventing coalitions from relying on outside agents. On the other hand, in some cases, it may be natural to allow some mild, rational assumptions on how others act. One can think of alternative, intermediate notions of coherence that achieve this. For example, we could require each coalition to guarantee its goal if the others behave consistently with their own intentions. One could also demand that coalitions act in a way that actively brings about their intention (to prevent situations where a coalition intends something that is under a disjoint coalition's control), adapting existing notions of group responsibility in CGSs (Yazdanpanah et al. 2019; Gladyshev et al. 2025). Such middle-ground notions still block weak coherence's assumptions while allowing limited and reasonable interdependence. We choose to focus on weak and strong coherence, leaving such notions to future work.

Our main object of study is a *revision operator*: an operator $\bullet$ that, given environmental constraints φ and two (for now, basic) intention systems I and J , returns $I \bullet_\varphi J$, being the outcome of revising I by J w.r.t. φ . We will require several rationality postulates, in particular: if J is coherent (for the chosen notion of coherence) with respect to φ , then so is $I \bullet_\varphi J$. However, in multiagent settings, coherence can often be restored in several *incomparable* ways, and a principled operator should not arbitrarily select one.

Consider agents a, b with intentions $I_a = \text{do}_{a,x}$ and $I_b = \text{do}_{b,x}$. A new directive for their joint coalition $\{a, b\}$ arrives: exactly one of them must perform x , i.e. the system J with $J_{\{a,b\}} = (\text{do}_{a,x} \wedge \neg \text{do}_{b,x}) \vee (\neg \text{do}_{a,x} \wedge \text{do}_{b,x})$. Adding J to I breaks coherence. Minimal repair requires dropping either I_a or I_b ; nothing distinguishes the two options. Following other work on indeterminate belief revision (Lindström and Rabinowicz 1989; De Rijke 1994), we let the operator return *both* repaired systems. In the Katsuno–Mendelzon view, this corresponds to minimising w.r.t. a nontotal preorder, whose minima may contain several incomparable elements.

Producing several outcomes amounts to returning a *disjunction* of basic intention systems. Because no single basic system K can, in general, capture the disjunction “ I or J ,” we generalize to *intention systems*: finite disjunctions of basic intention systems.

Definition 3.5 (Intention system). An *intention system* $\mathcal{I} \subseteq \text{BIS}$ is a finite nonempty subset of basic intention systems. We denote the set of all intention systems by IS .

Given intention systems $\mathcal{I}$ and $\mathcal{J}$, we define the *conjunction* $\mathcal{I} \sqcap \mathcal{J} = \{\{I_A \wedge J_A\}_{A \in \text{Coal}} \mid I \in \mathcal{I}, J \in \mathcal{J}\}$ and the

disjunction $\mathcal{I} \sqcup \mathcal{J} = \mathcal{I} \cup \mathcal{J}$.

Given $(\mathbb{G}, s)$ and $\sigma \in \text{Prof}(\sigma)$, we say that an intention system $\mathcal{I}$ is *weakly* (resp. *strongly*) *satisfied* by σ , denoted $\mathbb{G}, \sigma \Vdash_w \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s \mathcal{I}$) if there is $I \in \mathcal{I}$ with $\mathbb{G}, \sigma \Vdash_w I$ (resp. $\mathbb{G}, \sigma \Vdash_s I$). Given constraints φ , $\mathcal{I}$ is *weakly* (resp. *strongly*) *coherent* w.r.t. φ if there is $(\mathbb{G}, s)$ with $\mathbb{G}, s \Vdash \varphi$ and such that there is $\sigma \in \text{Prof}(s)$ with $\mathbb{G}, \sigma \Vdash_w \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s \mathcal{I}$).

In what follows, we therefore work with revision operators on intention systems. This captures the intended indeterminacy: revising I by J under φ may return several incomparable options, collected as one intention system. For homogeneity, we consider our revision operators to also take in intention systems, though we note that we will give postulates later on that decompose revision of intention systems into revision of basic ones, making this a harmless choice.

Definition 3.6 (Intention system revision operator). Denote by Constr the set of constraints (context formulas). An *intention system revision operator* is a function $\bullet : \text{Constr} \times \text{IS} \times \text{IS} \rightarrow \text{IS}$, taking in constraints φ and intention systems $\mathcal{I}$ and $\mathcal{J}$, and outputting $\mathcal{I} \bullet_\varphi \mathcal{J}$, to be read as the result of revising $\mathcal{I}$ with $\mathcal{J}$ relative to φ .

If the intention systems $\mathcal{I}$ or $\mathcal{J}$ are singletons $\{I\}$ or $\{J\}$, we drop the set brackets and write e.g. $I \bullet_\varphi J$.

4 A General Representation Theorem

We now turn to postulates for revision operators and their semantic characterisation. As we will explain, we achieve this via a more general theorem, which we prove in this section.

Following the approach of Katsuno and Mendelzon (1991) for belief revision over finite-signature propositional logic, we seek a representation theorem: an operator satisfies the postulates iff it selects, among the models that satisfy the new information, those that are *minimal* in a preorder induced by the original intention system, encoding how close models are to satisfying the latter. This expresses revision as satisfying a principle of *minimal change*.

Motamed et al. (2024) carry the idea over to intention revision: they read coherence as logical consistency and invoke logic-agnostic results like those of Delgrande, Peppas, and Woltran (2018) and Falakh, Rudolph, and Sauerwald (2023). These approaches, however, rely on *total* preorders, whereas indeterminate revision requires *non-total* ones. Several incomparable minimal sets, and hence several admissible outcomes, can coexist. Existing theorems thus do not apply.

In order to obtain representation theorems for intention system revision operators, we return to the original proof technique of Katsuno and Mendelzon (1991): we isolate the ingredients that do not rely on totality, and obtain a *general, logic-agnostic* representation theorem that allows non-total preorders, and covers both weak and strong coherence in a single statement. Two additional parameters are introduced: a *disjunctive basis* of “primitive” formulas (for intention system revision, the basic intention systems) on which the operator is specified, and a *frame* that restricts admissible preorders to encode extra-logical priorities. With these tools in

place, we set up the abstract framework that underlies the concrete instantiation in Section 5.

We begin our development by defining the basic setup.

Definition 4.1 (Logics). A *logic* is a triple $\mathbb{L} = (\mathcal{F}, \mathcal{M}, \text{Mod})$, where $\mathcal{F}$ is a nonempty set of *formulas*, $\mathcal{M}$ is a nonempty set of *models*, and $\text{Mod}: \mathcal{F} \rightarrow 2^{\mathcal{M}}$ is the *semantics*. The logic is *conjunctive* if there exists an operator $\wedge: \mathcal{F} \times \mathcal{F} \rightarrow \mathcal{F}$ such that $\text{Mod}(\varphi \wedge \psi) = \text{Mod}(\varphi) \cap \text{Mod}(\psi)$, referred to as *conjunction*. It is *disjunctive* if there exists an operator $\vee: \mathcal{F} \times \mathcal{F} \rightarrow \mathcal{F}$ such that $\text{Mod}(\varphi \vee \psi) = \text{Mod}(\varphi) \cup \text{Mod}(\psi)$, referred to as *disjunction*. And the logic is *full* if for every model $M \in \mathcal{M}$ there is a formula $\varphi \in \mathcal{F}$ with $M \in \text{Mod}(\varphi)$.

A formula φ is *satisfiable* if $\text{Mod}(\varphi) \neq \emptyset$, and two formulas $\varphi, \psi \in \mathcal{F}$ are *mutually unsatisfiable* if $\text{Mod}(\varphi) \cap \text{Mod}(\psi) = \emptyset$. Given $\varphi, \psi \in \mathcal{F}$, we say φ *entails* ψ if $\text{Mod}(\varphi) \subseteq \text{Mod}(\psi)$, and say φ is *equivalent* to ψ if $\text{Mod}(\varphi) = \text{Mod}(\psi)$.

A *disjunctive basis* for a disjunctive logic $\mathbb{L}$ is a set $\mathcal{B} \subseteq \mathcal{F}$ such that for all satisfiable $\varphi \in \mathcal{F}$, there exist $\beta_1, \dots, \beta_n \in \mathcal{B}$ for which φ is equivalent to $\beta_1 \vee \dots \vee \beta_n$.

A formula φ is an *atom* if it is satisfiable and for all $\psi \in \mathcal{F}$, either φ entails ψ , or φ and ψ are mutually unsatisfiable. We write $\mathcal{A}_{\mathbb{L}}$ for the set of all atoms of $\mathbb{L}$. The logic $\mathbb{L}$ is *atomistic* if it is disjunctive, full, and $\mathcal{A}_{\mathbb{L}}$ is a disjunctive basis for $\mathbb{L}$.

Noting that $\vee$ is associative up to equivalence, we will write the expression $\varphi_1 \vee \dots \vee \varphi_n$ without bracketing.

What we refer to as an atom, is referred to by Katsuno and Mendelzon (1991) as a ‘complete’ formula. We instead borrow terminology from lattice theory, where one speaks of atoms and atomistic lattices, used similarly to the way we use the terminology.

Next, we define the main object of study: revision operators. Note that unlike in Definition 3.6, these are the standard binary operators of revision theory.

Definition 4.2. An $\mathbb{L}$ -*revision operator* is a function $\bullet: \mathcal{F} \times \mathcal{F} \rightarrow \mathcal{F}$.

Before introducing assignments, the semantic object of the representation theorem, some formal preliminaries are needed. Given a set X , write $\text{Ord}(X)$ for the set of preorders on X . Given a preorder $\leq$, we denote its strict part by $<$, defined by $x < y$ iff $x \leq y$ and $y \not\leq x$. Given $\leq \in \text{Ord}(X)$ and $A \subseteq X$, we write $\min(A, \leq) = \{x \in A \mid \neg \exists y \in A. y < x\}$ for the set of all $\leq$ -minimal elements of A .

The representation theorems relate revision operators satisfying postulates to assignments of preorders to formulas. In the final theorem, these assignments will only be defined on a disjunctive basis (corresponding to basic intention systems in the setting of group intentions). Therefore, the definition is parameterized by a subset of formulas.

Definition 4.3 ($\mathcal{K}$ -assignment). Take $\mathcal{K} \subseteq \mathcal{F}$. A $\mathcal{K}$ -*assignment* is a function $\preceq_{(-)}: \mathcal{K} \rightarrow \text{Ord}(\mathcal{M})$ assigning to each formula $\varphi \in \mathcal{K}$ a preorder $\preceq_{\varphi}$ on $\mathcal{M}$. We say $\preceq$ is *faithful* if for all $\varphi, \psi \in \mathcal{K}$, it holds that (i) if $M, N \in \text{Mod}(\varphi)$, then $M \not\prec_{\varphi} N$; (ii) if $M \in \text{Mod}(\varphi)$ and $N \notin \text{Mod}(\varphi)$, then $M \prec_{\varphi} N$; and (iii) if φ is equivalent to ψ , then $\preceq_{\varphi} = \preceq_{\psi}$.

We say $\preceq$ is *logical* if whenever M and N satisfy precisely the same formulas (i.e. $M \in \text{Mod}(\psi)$ iff $N \in \text{Mod}(\psi)$ for all $\psi \in \mathcal{F}$), then $M \preceq_{\varphi} N$ and $N \preceq_{\varphi} M$ for all $\varphi \in \mathcal{K}$.

Given $\mathcal{K}$ -assignments $\preceq^1$ and $\preceq^2$, we say $\preceq^1$ *min-extends* $\preceq^2$ if whenever $M \prec_{\varphi}^2 N$, it holds that $N \in \min(\text{Mod}(\psi), \preceq_{\varphi}^1)$ implies that $M \in \min(\text{Mod}(\psi), \preceq_{\varphi}^1)$ for all $\varphi \in \mathcal{K}$ and $\psi \in \mathcal{F}$ with $M \in \text{Mod}(\psi)$.

Given an $\mathbb{L}$ -revision operator $\bullet$, a disjunctive basis $\mathcal{B}$, and a $\mathcal{B}$ -assignment $\preceq$, we say $\bullet$ and $\preceq$ are *compatible* if for all $\varphi, \psi \in \mathcal{F}$ with φ equivalent to $\beta_1 \vee \dots \vee \beta_n$ (where $\beta_i \in \mathcal{B}$), it holds that $\text{Mod}(\varphi \bullet \psi) = \bigcup_{1 \leq i \leq n} \min(\text{Mod}(\psi), \preceq_{\beta_i})$.

Let us unpack some of these definitions. Faithfulness, as in the original work of (Katsuno and Mendelzon 1991), captures that the assignment provides a measure of how close models are to satisfying a formula: $M \preceq_{\varphi} N$ is to be read as saying that M is closer to being a model of φ than N . We introduce the notion of logicity to capture that an assignment should not be able to make finer distinctions than the logic itself: if a distinction is made between models, then this distinction has to do with some difference in the models’ logical properties. The notion of min-extension is interpreted as follows. If $\preceq^1$ min-extends $\preceq^2$, then this says that whenever $\preceq^2$ expresses a strict preference of one model over another, that the process of minimization (corresponding to revision) has to reflect that preference. Finally, compatibility captures the essence of the representation theorem, stating when a revision operator is semantically described by an assignment (in which the disjunctive basis is respected).

We can now define *frames*, which we use to add extra-logical constraints to revision operators. These are families of preorders of formulas, corresponding semantically to ordered partitions of the set of models. Intuitively, the ordering of cells in a partition captures that revision should prioritize satisfying certain formulas over another. More formally, we will ask that our assignments min-extend the frame interpreted as an assignment in the natural way.

Definition 4.4 ($\mathcal{K}$ -frame). Given $\mathcal{K} \subseteq \mathcal{F}$, a $\mathcal{K}$ -*frame* is a $\mathcal{K}$ -indexed family $\Xi = (\Xi_{\varphi}, \sqsubseteq_{\varphi})_{\varphi \in \mathcal{K}}$, in which $\Xi_{\varphi} \subseteq \mathcal{F}$ is a set of formulas, and $\sqsubseteq_{\varphi} \in \text{Ord}(\Xi_{\varphi})$ is a preorder. We require that all formulas in Ξ_{φ} are satisfiable, that all pairs $\psi, \chi \in \Xi_{\varphi}$ are either equivalent or mutually unsatisfiable, and that $\bigcup_{\psi \in \Xi_{\varphi}} \text{Mod}(\psi) = \mathcal{M}$. We say Ξ is *faithful* if for all $\varphi \in \mathcal{K}$ and $\psi, \chi \in \Xi_{\varphi}$, it holds that

1. Either φ and ψ are mutually unsatisfiable, or ψ entails φ .
2. If ψ and χ both entail φ , then $\psi \not\sqsubseteq_{\varphi} \chi$.
3. If ψ entails φ but χ does not, then $\psi \sqsubseteq_{\varphi} \chi$.

We can now begin presenting postulates. In the following, take the logic $\mathbb{L}$ to be conjunctive and disjunctive, and let $\mathcal{B}$ be a disjunctive basis for $\mathbb{L}$ with Ξ a $\mathcal{B}$ -frame. All formulas are assumed to be arbitrary formulas from $\mathcal{F}$, unless specified otherwise. The postulates (R1)-(R5) and (R7) are directly from Katsuno and Mendelzon (1991) and use the same numbering: (R6) is not used to allow non-totality. The postulate (R8B) is a slightly tweaked version of their postulate (R8), modified to account for the disjunctive basis as required in the representation theorem. The postulates (R Ξ)

and (RDL) are novel, expressing min-extension of the frame and revision respecting disjunctions, respectively.

- (R1) $\varphi \bullet \psi$ entails ψ .
- (R2) If $\varphi \wedge \psi$ is satisfiable, then $\varphi \bullet \psi$ is equivalent to $\varphi \wedge \psi$.
- (R3) If ψ is satisfiable, then $\varphi \bullet \psi$ is also satisfiable.
- (R4) If φ is equivalent to φ' , and ψ is equivalent to ψ' , then $\varphi \bullet \psi$ is equivalent to $\varphi' \bullet \psi'$.
- (R5) $(\varphi \bullet \psi) \wedge \chi$ entails $\varphi \bullet (\psi \wedge \chi)$.
- (R7) If $\varphi \bullet \psi$ entails ψ' and $\varphi \bullet \psi'$ entails ψ , then $\varphi \bullet \psi$ is equivalent to $\varphi \bullet \psi'$.
- (R8B) For $\beta \in \mathcal{B}$, $(\beta \bullet \psi) \wedge (\beta \bullet \psi')$ entails $\beta \bullet (\psi \vee \psi')$.
- (RΞ) For $\beta \in \mathcal{B}$ and $\mu, \nu \in \Xi_\beta$ with $\mu \sqsubset_\beta \nu$, if $(\beta \bullet \psi) \wedge \nu$ is satisfiable, then $\mu \wedge \psi$ entails $\beta \bullet \psi$.
- (RDL) $(\varphi \vee \varphi') \bullet \psi$ is equivalent to $(\varphi \bullet \psi) \vee (\varphi' \bullet \psi)$.

Before presenting the representation theorem, we first introduce some notation. Given a $\mathcal{K}$ -frame Ξ , define the $\mathcal{K}$ -assignment $\preceq^\Xi$ by letting $M \preceq^\Xi N$ iff there exist $\psi, \chi \in \Xi_\varphi$ with $M \in \text{Mod}(\psi)$, $N \in \text{Mod}(\chi)$ and $\psi \sqsubseteq_\varphi \chi$. It is easily verified that $\preceq^\Xi$ is indeed a preorder. We say a $\mathcal{K}$ -assignment $\preceq$ respects Ξ if it min-extends $\preceq^\Xi$.

Theorem 4.5 (General representation theorem). *Let $\mathbb{L}$ be a conjunctive, disjunctive and atomistic logic, and let $\mathcal{B}$ be a disjunctive basis for $\mathbb{L}$ with a faithful $\mathcal{B}$ -frame Ξ . The following statements hold:*

- Every faithful and logical $\mathcal{B}$ -assignment respecting Ξ is compatible with some $\mathbb{L}$ -revision operator satisfying (R1) – (R5), (R7), (R8B), (RΞ), and (RDL).
- Every $\mathbb{L}$ -revision operator satisfying (R1) – (R5), (R7), (R8B), (RΞ), and (RDL) is compatible with some faithful and logical $\mathcal{B}$ -assignment respecting Ξ .

Proof sketch. The idea is to emulate the proof of Katsuno and Mendelzon (1991, Theorem 5.2) for atomistic logics. For the first half, using the logicity of the assignment, the set of minimal models of a formula must be described by a disjunction of atoms, allowing us to define our operator. We then verify that the postulates hold, taking the disjunctive basis into account. For the second half, note that in the original proof of Katsuno and Mendelzon (1991, Theorem 5.2), the assignment is defined using the observation that for every set of models there exists a formula that is precisely satisfied by that set. The assignment is then defined by ordering a pair of models on the basis of which of the two models are contained in a revision by the formula expressing the pair. In our setting, we generally cannot capture every set of models via a formula. Instead, we work with the atoms that models are contained in, ordering models on the basis of which of their atoms are consistent with a revision by the disjunction of the atoms. $\square$

We observe in the following proposition that, to a degree, a symmetric analogue of (RDL), stating that we have disjunctive factoring on the right of the revision operator, is a consequence of our other postulates. For intention system revision, this together with (RDL) will mean that intention system revision operators are fully determined by their actions on basic intention systems.

Proposition 4.6. *If $\bullet$ satisfies (R1) – (R5), (R7), (R8B) and (RDL), then for all $\varphi, \psi, \psi' \in \mathcal{F}$ with φ equivalent to $\beta_1 \vee \dots \vee \beta_n$, there exist nonempty $K_i \subseteq \{\psi, \psi'\}$ for $1 \leq i \leq n$ such that the formula $\varphi \bullet (\psi \vee \psi')$ is equivalent to $\bigvee_{1 \leq i \leq n} \bigvee_{\chi \in K_i} \beta_i \bullet \chi$.*

Proof sketch. The proof applies (Falahk, Rudolph, and Sauerwald 2023, Proposition 11.7), taking the disjunctive basis into account. $\square$

5 Representation Theorem for Intention System Revision

In this section, we formulate rationality postulates for intention system revision and prove a Katsuno–Mendelzon–style representation theorem instantiating the approach from Section 4. The key move, which follows the single-agent treatment of Motamed et al. (2024), is to read both weak and strong coherence in $\text{XATL}_{\text{sc}, \text{do}}$ as satisfiability in new logics (in the sense of Definition 4.1) whose (i) formulas are intention systems, (ii) models are strategy profiles, and (iii) truth clause incorporates the constraints φ . The representation theorem then expresses that revision works via minimal change: to revise old intentions by some new ones, we select those profiles satisfying the new intentions that are minimally close to satisfying the old ones.

These constructed logics must satisfy atomisticity as in Definition 4.1, but unbounded time prevents that: an intention can always be extended with claims about later moments that are neither entailed nor contradictory, disallowing atoms. In order to recover atomisticity, we follow Van Zee et al. (2020) by considering revision operators up to a finite time horizon. This change is innocent, since few applications allow intentions referring to an unbounded length of time.

Let us begin by defining the basic setup of bounding time. Given an intention system $\mathcal{I}$, we define its *temporal depth* $\text{td}(\mathcal{I})$ to be the maximal nesting depth of the X -operators appearing in any formula in $\mathcal{I}$.

Definition 5.1 (Time-bounded semantics, $\mathbb{L}_w^{n, \varphi}$ and $\mathbb{L}_s^{n, \varphi}$). Let $n \geq 0$. Given $(\mathbb{G}, s)$ and $(\mathbb{G}', s')$, together with strategy profiles σ and σ' (from s and s' respectively), we write $\mathbb{G}, \sigma \equiv_w^n \mathbb{G}', \sigma'$ (resp. $\mathbb{G}, \sigma \equiv_s^n \mathbb{G}', \sigma'$) iff for all intention systems $\mathcal{I}$ with $\text{td}(\mathcal{I}) \leq n$ it holds that $\mathbb{G}, \sigma \Vdash_w \mathcal{I}$ iff $\mathbb{G}', \sigma' \Vdash_w \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s \mathcal{I}$ iff $\mathbb{G}', \sigma' \Vdash_s \mathcal{I}$). We also write $\mathbb{G}, \sigma \Vdash_w^n \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s^n \mathcal{I}$) if there is $(\mathbb{G}', \sigma')$ with $\mathbb{G}, \sigma \equiv_w^n \mathbb{G}', \sigma'$ (resp. $\mathbb{G}, \sigma \equiv_s^n \mathbb{G}', \sigma'$) such that $\mathbb{G}', \sigma' \Vdash_w \mathcal{I}$ (resp. $\mathbb{G}', \sigma' \Vdash_s \mathcal{I}$).

We say that an intention system $\mathcal{I}$ is *weakly* (resp. *strongly*) *coherent* w.r.t. constraints φ up to n , if there exist $\mathbb{G}, s$ and a strategy profile σ from s such that $\mathbb{G}, s \Vdash \varphi$ and $\mathbb{G}, \sigma \Vdash_w^n \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s^n \mathcal{I}$). An intention system $\mathcal{I}$ *weakly* (φ, n)-*entails* $\mathcal{J}$ if $\mathbb{G}, \sigma \Vdash_w^n \mathcal{I}$ implies $\mathbb{G}, \sigma \Vdash_w^n \mathcal{J}$. *Strong* (φ, n)-*entailment* is defined analogously. Similarly, $\mathcal{I}$ and $\mathcal{J}$ are *weakly* (resp. *strongly*) (φ, n)-*equivalent* if they weakly (resp. strongly) (φ, n)-entail each other.

Fix constraints φ . We define the logics (in the sense of Definition 4.1) $\mathbb{L}_w^{n, \varphi}$ and $\mathbb{L}_s^{n, \varphi}$ by taking intention systems as formulas, pairs $(\mathbb{G}, \sigma)$ of CGSs and strategy profiles (with σ from s) as models, and defining the semantics in $\mathbb{L}_w^{n, \varphi}$ (resp.

$\mathbb{L}_s^{n,\varphi}$ of $\mathcal{I}$ to be the set of all $(\mathbb{G}, \sigma)$ such that $\mathbb{G}, s \Vdash \varphi$ and $\mathbb{G}, \sigma \Vdash_w^n \mathcal{I}$ (resp. $\mathbb{G}, \sigma \Vdash_s^n \mathcal{I}$).

Weak (resp. strong) coherence with respect to φ up to horizon n is exactly satisfiability in $\mathbb{L}_w^{n,\varphi}$ (resp. $\mathbb{L}_s^{n,\varphi}$). Entailment and equivalence are therefore defined relative to the same pair (n, φ) : environmental constraints (e.g., “ a must always be followed by b ”) alter what an intention system implies, while times beyond n are rightly ignored.

The following theorem states these logics using weak and strong satisfaction have the properties we require to instantiate the representation theorem from Section 4.¹

Theorem 5.2. *The logics $\mathbb{L}_w^{n,\varphi}$ and $\mathbb{L}_s^{n,\varphi}$ are conjunctive, disjunctive and atomistic.*

Proof sketch. It is clear that the logics are conjunctive and disjunctive using the operations $\sqcap$ and $\sqcup$ on intention systems defined in Definition 3.1. Showing that the logics are atomistic is more complicated. The method consists of first establishing that the equivalence relations $\equiv_w^n$ and $\equiv_s^n$ coincide with an appropriate notion of n -step bisimilarity for strategy contexts; in other words, we show a (bounded) Hennessy-Milner theorem (see e.g. Goranko and Otto (2007) for an introduction to bisimulations in modal logic). We then show that n -step bisimilarity is equivalent to the existence of winning strategies in certain Ehrenfeucht-Fraïssé games. The existence of such strategies is logically definable via intention systems. It follows that that every equivalence class of $\equiv_w^n$ and $\equiv_s^n$ are expressed by intention systems under the appropriate semantics. In particular, the construction of these intention systems shows that there are finitely many equivalence classes. By definition of $\equiv_w^n$ and $\equiv_s^n$, every intention system must then be equivalent to a disjunction of the constructed intention systems under the appropriate semantics. $\square$

Showing atomisticity makes use of an innovative proof method. For readers familiar with the terminology, we note that the proof works by showing a Hennessy-Milner theorem for an appropriate notion of bisimulation: full details are in the supplementary material.

Now, we can begin presenting postulates for intention system revision operators, obtained by instantiating the work from Section 4. In the following, we drop the words ‘weak’ and ‘strong’: effectively we get two sets of postulates, obtained by inserting these adjectives uniformly throughout. Additionally, we fix a parameter $n \geq 0$.

- (I1) $\mathcal{I} \bullet_\varphi \mathcal{J} (\varphi, n)$ -entails $\mathcal{J}$.
- (I2) If $\mathcal{I} \sqcap \mathcal{J}$ is coherent w.r.t. φ up to n , then $\mathcal{I} \bullet_\varphi \mathcal{J}$ is (φ, n) -equivalent to $\mathcal{I} \sqcap \mathcal{J}$.
- (I3) If $\mathcal{J}$ is coherent w.r.t. φ up to n , then $\mathcal{I} \bullet_\varphi \mathcal{J}$ is also coherent w.r.t. φ up to n .
- (I4) If $\mathcal{I}$ is (φ, n) -equivalent to $\mathcal{I}'$, and $\mathcal{J}$ is (φ, n) -equivalent to $\mathcal{J}'$, then $\mathcal{I} \bullet_\varphi \mathcal{J}$ is (φ, n) -equivalent to $\mathcal{I}' \bullet_\varphi \mathcal{J}'$.

¹We note that the alternative coherence conditions we described briefly in Section 3 do not satisfy the required properties, meaning that a representation theorem for them would require a completely different proof technique.

- (I5) $(\mathcal{I} \bullet_\varphi \mathcal{J}) \sqcap \mathcal{K}$ entails $\mathcal{I} \bullet_\varphi (\mathcal{J} \sqcap \mathcal{K})$.
- (I7) If $\mathcal{I} \bullet_\varphi \mathcal{J} (\varphi, n)$ -entails $\mathcal{J}'$ and $\mathcal{I} \bullet_\varphi \mathcal{J}' (\varphi, n)$ -entails $\mathcal{J}$, then $\mathcal{I} \bullet_\varphi \mathcal{J}$ is (φ, n) -equivalent to $\mathcal{I} \bullet_\varphi \mathcal{J}'$.
- (I8-BI) For basic intention systems I , it holds that $(I \bullet_\varphi \mathcal{J}) \sqcap (I \bullet_\varphi \mathcal{J}') (\varphi, n)$ -entails $I \bullet_\varphi (\mathcal{J} \sqcup \mathcal{J}')$.
- (IDL) $(\mathcal{I} \sqcup \mathcal{I}') \bullet_\varphi \mathcal{J}$ is (φ, n) -equivalent to $(\mathcal{I} \bullet_\varphi \mathcal{J}) \sqcup (\mathcal{I}' \bullet_\varphi \mathcal{J})$.

These postulates instantiate those from Section 4. (I1) ensures success of revision, (I2) ensures revision does not modify intentions unnecessarily, (I3) demands coherence when possible, and (I4) gives syntax-invariance. (I5), (I7) and (I8-BI) are technical guarantees that ensure that revision semantically works through minimal change, with (I7) and (I8-BI) allowing non-totality, the latter using basic systems. (IDL) allows us to decompose revisions of intention systems to those of basic intention systems.

At this point, we could already present our representation theorem, but we add one more optional parameter allowing extra conditions of interest to the multiagent setting. One may want revision to respect coalition-sensitive priorities: for example, (i) keep the intentions of as many coalitions as possible, or (ii) guarantee that a designated core coalition is disturbed only as a last resort. We explain how such priorities are incorporated into revision operators via a postulate.

Consider a strict preorder $\sqsubseteq$ on $2^{\text{Coal}} \setminus \{\emptyset\}$, the non-empty subsets of Coal, which encodes priorities; we read $\mathbb{A} \sqsubseteq \mathbb{B}$ as saying that preserving exactly the intentions of $\mathbb{A}$ takes precedence over $\mathbb{B}$, with incomparable pairs being indifferent in revision. We fold $\sqsubseteq$ into the framework by building a frame (Definition 4.4), and using that frame in postulate (R Ξ), forcing revision to honor the chosen priority order.

Let us make precise what it means for a revision operator to behave according to the priority order. Consider a basic intention system I . For nonempty $\mathbb{A} \subseteq \text{Coal}$, define basic intention systems $\text{keep}_w^\mathbb{A}(I) = \{I_A \mid A \in \mathbb{A}\} \cup \{\neg I_A \mid A \notin \mathbb{A}\}$ and $\text{keep}_s^\mathbb{A}(I) = \{I_A \mid A \in \mathbb{A}\} \cup \{(\neg \bar{A}) \langle \bar{A} \rangle \square \neg I_A \mid A \notin \mathbb{A}\}$. Then a strategy profile weakly satisfies $\text{keep}_w^\mathbb{A}(I)$ (resp. strongly satisfies $\text{keep}_s^\mathbb{A}(I)$) iff $\mathbb{A}$ is the greatest set of coalitions for which the profile weakly (resp. strongly) satisfies the basic intention system $\{I_A \mid A \in \mathbb{A}\} \cup \{\top \mid A \notin \mathbb{A}\}$. In other words, $\text{keep}_w^\mathbb{A}(I)$ (resp. $\text{keep}_s^\mathbb{A}(I)$) expresses that a profile precisely weakly (resp. strongly) satisfies the intentions in I for coalitions in $\mathbb{A}$.

Using $\text{keep}_w^\mathbb{A}(I)$ and $\text{keep}_s^\mathbb{A}(I)$, we can give the following postulate expressing that a revision operator respects the priority ordering. Let $x \in \{w, s\}$ denote ‘weak’ or ‘strong’ (and as before, consider the chosen adjective to apply uniformly to coherence and entailment).

- (IPrio) For basic intention systems I and J , if $(I \bullet_\varphi J) \sqcap \text{keep}_x^\mathbb{B}(I)$ is coherent w.r.t. φ up to n , then for every $\mathbb{A} \sqsubseteq \mathbb{B}$, $J \sqcap \text{keep}_x^\mathbb{A}(I) (\varphi, n)$ -entails $I \bullet_\varphi J$.

Let us read out what this means. The postulate states that if the revision outcome $I \bullet_\varphi J$ admits profiles that preserve precisely $\mathbb{B}$ ’s intentions from I , then every profile preserving a more important set $\mathbb{A}$ of coalitions’ intentions from I

that the revision could also have picked (because they satisfy J) must also be admitted. This is precisely what we want: among profiles that accommodate J , those that preserve a more important family of coalitions should never be rejected while a less important family is accepted.

The connection to frames and the $(R\Xi)$ postulate is immediate using the constructed $\text{keep}_x^A(I)$. We say a BIS-assignment (as in Definition 4.3) respects $\sqsubseteq$ if it respects the corresponding frame. Note that BIS-assignments are mappings from basic intention systems to preorders over strategy profiles whose underlying state satisfies the constraints.

Instantiating our representation theorem in Theorem 4.5, making use of the required properties shown in Theorem 5.2 we can now conclude with the following representation theorem for intention system revision operators.

Theorem 5.3 (Intention system representation theorem). *Fix constraints φ , $n \geq 0$, let $\mathbb{L} \in \{\mathbb{L}_w^{n,\varphi}, \mathbb{L}_s^{n,\varphi}\}$, and take a strict preorder $\sqsubseteq$ on $2^{\text{Coal}} \setminus \{\emptyset\}$. The following hold:*

- Every faithful and logical BIS-assignment respecting $\sqsubseteq$ is compatible with some intention system revision operator satisfying (I1) – (I5), (I7), (I8-BI), (IPrio), and (IDL).
- Every intention system revision operator satisfying (I1) – (I5), (I7), (I8-BI), (IPrio), and (IDL) is compatible with some faithful and logical BIS-assignment respecting $\sqsubseteq$.

Note that by (IDL) and Proposition 4.6, revision reduces to that of basic intention systems, in the sense that the revision of $\mathcal{I}$ by $\mathcal{J}$ is equivalent to a disjunction of revisions of basic intention systems from $\mathcal{I}$ and $\mathcal{J}$. And also note that if we use a trivial, empty priority preorder, then (IPrio) trivializes, giving a representation theorem omitting priorities.

It is important to observe that Theorem 5.3 is constructive and operational: given any assignment meeting the required conditions, the theorem explicitly constructs the corresponding intention system revision operator. We work out an example of an intention system revision operator obtained in this fashion. Intuitively, this operator resolves conflicts between the original and the new intentions by progressively trimming parts of the original intentions as far in the future as possible, until all conflicts are resolved.

Fix $n \geq 0$ and constraints φ . Consider the logic $\mathbb{L}_w^{n,\varphi}$. Given $I \in \text{BIS}$, a coalition $A \in \text{Coal}$, a CGS $\mathbb{G}$, and a strategy profile σ from state s , we define the *failure depth* $\text{fd}_I^A(\sigma)$ as, if it exists, the least $k < n$ such that looking at the k -step prefix $\eta = s_0\alpha_1s_1 \dots \alpha_k s_k$ of the outcome $\text{out}(\sigma) \in \text{Path}(s)$, it holds that $\mathbb{G}, \pi, \sigma \not\models I_A$ for all paths $\pi \in \text{Path}(s)$ that extend η . Intuitively, the failure depth counts how many steps in time we must move along the outcome of σ before I_A is falsified. If it does not exist, this means that I_A cannot be falsified within n steps.

We now define a BIS-assignment. Let $(\mathbb{G}, \sigma) \preceq_I (\mathbb{G}', \sigma')$ iff either (i) $\mathbb{G}, \sigma \models_w^n I$, or (ii) $\mathbb{G}, \sigma \not\models_w^n I$, $\mathbb{G}', \sigma' \not\models_w^n I$ and $\text{fd}_I^A(\sigma) \geq \text{fd}_I^A(\sigma')$ for all $A \in \text{Coal}$ (noting that the failure depths are defined since neither profile weakly satisfies I up to n). The first clause is necessary for faithfulness, as in those cases the failure depth will not exist. For an intuition, note that $(\mathbb{G}, \sigma) \prec_I (\mathbb{G}', \sigma')$ precisely when σ never fails I earlier than σ' , while σ' does fail I earlier for some coalitions.

Faithfulness and logicity are easy to verify: by the first clause of the definition of the assignment, we have that the least elements of $\preceq_I$ are precisely those $(\mathbb{G}, \sigma)$ with $\mathbb{G}, \sigma \models_w^n I$, and the definition of the assignment is completely syntax-invariant and based on satisfying formulas.

So applying our representation theorem, we get an intention system revision operator compatible with the assignment, and satisfying all our postulates. By compatibility, this operator computes the revision $I \bullet_\varphi J$ of basic intention systems (with other revision reducing to these cases by disjunctive factoring) by those profiles that weakly satisfy J and are minimal in the ordering of I . Conceptually, this says that the revision is obtained by taking J and maintaining I as long as possible, for as many coalitions as possible.

6 Conclusion

This work provided the first general framework for revising interconnected systems of multiagent intentions. Intentions for every coalition are expressed in an appropriate logical language, coherence is analysed via feasibility and enforceability criteria, and rational revision operators are axiomatised and semantically characterised by a Katsuno–Mendelzon-style theorem that allows non-total preorders. The framework provides a unifying theoretical foundation for maintaining joint commitments in complex autonomous systems. Future work includes the development of more concrete revision operators instantiating the framework, integration of stochastic environments, and unifying our framework with methods for individual and/or coalition-level intention revision.

References

- Bratman, M. 1987. *Intention, Plans, and Practical Reason*. Cambridge: Cambridge, MA: Harvard University Press.
- Brihaye, T.; Da Costa, A.; Laroussinie, F.; and Markey, N. 2009. ATL with Strategy Contexts and Bounded Memory. In Artemov, S.; and Nerode, A., eds., *Logical Foundations of Computer Science*, 92–106. Berlin, Heidelberg: Springer Berlin Heidelberg. ISBN 978-3-540-92687-0.
- Cawsey, A.; Galliers, J.; Logan, B.; Reece, S.; and Sparck Jones, K. 1993. Revising Beliefs and Intentions: A Unified Framework for Agent Interaction. In Sloman, A.; Hogg, D.; Humphreys, G.; Ramsay, A.; and Partridge, D., eds., *Prospects for Artificial Intelligence: Proceedings of AISB'93, The Ninth Biennial Conference of the Society for the Study of Artificial Intelligence and the Simulation of Behaviour*, Frontiers in Artificial Intelligence Applications, 130–139. IOS Press.
- Dann, M.; Yao, Y.; Alechina, N.; Logan, B.; Meneguzzi, F.; and Thangarajah, J. 2023. Multi-Agent Intention Recognition and Progression. In Elkind, E., ed., *Proceedings of the Thirty-Second International Joint Conference on Artificial Intelligence, IJCAI-23*, 91–99. International Joint Conferences on Artificial Intelligence Organization.
- Delgrande, J. P.; Peppas, P.; and Woltran, S. 2018. General Belief Revision. *J. ACM*, 65(5).

- Dunin-Keplicz, B.; and Verbrugge, R. 2002. Evolution of collective commitment during reconfiguration. In *The First International Joint Conference on Autonomous Agents & Multiagent Systems, AAMAS 2002, July 15-19, 2002, Bologna, Italy, Proceedings*, 1141–1142. ACM.
- Falakh, F. M.; Rudolph, S.; and Sauerwald, K. 2023. AGM Belief Revision, Semantically. arXiv:2112.13557.
- Gladyshev, M.; Alechina, N.; Dastani, M.; and Doder, D. 2025. Reasoning about group responsibility for exceeding risk threshold in one-shot games. *Information and Computation*, 303: 105257.
- Goranko, V.; and Otto, M. 2007. 5 Model theory of modal logic. In Blackburn, P.; van, J.; and Wolter, F., eds., *Handbook of Modal Logic*, volume 3 of *Studies in Logic and Practical Reasoning*, 249–329. Elsevier.
- Herzig, A.; Lorini, E.; and Walther, D. 2013. Reasoning about Actions Meets Strategic Logics. In *Proceedings of the 4th International Workshop on Logic, Rationality, and Interaction - Volume 8196, LORI 2013*, 162–175. Berlin, Heidelberg: Springer-Verlag. ISBN 9783642409479.
- Icard, T.; Pacuit, E.; and Shoham, Y. 2010. Joint revision of belief and intention. In *Proceedings of the Twelfth International Conference on Principles of Knowledge Representation and Reasoning, KR'10*, 572–574. AAAI Press. ISBN 1577354516.
- Katsuno, H.; and Mendelzon, A. O. 1991. Propositional knowledge base revision and minimal change. *Artificial Intelligence*, 52(3): 263–294.
- Lindström, S.; and Rabinowicz, W. 1989. On Probabilistic Representation of Non-Probabilistic Belief Revision. *Journal of Philosophical Logic*, 18(1): 69–101.
- Motamed, N.; Alechina, N.; Dastani, M.; and Doder, D. 2024. Revising Beliefs and Intentions in Stochastic Environments. In Larson, K., ed., *Proceedings of the Thirty-Third International Joint Conference on Artificial Intelligence, IJCAI-24*, 3513–3522. International Joint Conferences on Artificial Intelligence Organization. Main Track.
- Motamed, N.; Alechina, N.; Dastani, M.; Doder, D.; and Logan, B. 2023. Probabilistic Temporal Logic for Reasoning about Bounded Policies. In Elkind, E., ed., *Proceedings of the 32nd International Joint Conference on Artificial Intelligence (IJCAI-23)*, 3296–3303. Main Track.
- de Rijke, M. 1994. Meeting Some Neighbours. In *Logic and Information Flow*. The MIT Press. ISBN 9780262285421.
- Shoham, Y. 2009. Logical Theories of Intention and the Database Perspective. *Journal of Philosophical Logic*, 38(6): 633–647.
- Yazdanpanah, V.; Dastani, M.; Jamroga, W.; Alechina, N.; and Logan, B. 2019. Strategic Responsibility Under Imperfect Information. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS '19, Montreal, QC, Canada, May 13-17, 2019*, 592–600.
- van Zee, M.; Doder, D.; van der Torre, L.; Dastani, M.; Icard, T.; and Pacuit, E. 2020. Intention as commitment toward time. *Artificial Intelligence*, 283: 103270.
- Zhang, J.; Nguyen, X.; and Kowalczyk, R. 2007. Graph-based multi-agent replanning algorithm. 122.
- Zhang, Z.; Zeng, Y.; Jiang, W.; Pan, Y.; and Tang, J. 2023. Intention recognition for multiple agents. *Inf. Sci.*, 628: 360–376.