Belief Updates through Dynamic Modal Logic:
Unveiling the Structure of Belief

Daniel Díaz Quílez

May 25$^{\text{th}}$, 2025

Abstract

Rational agents are constantly updating their beliefs. Researchers need a framework for reasoning about belief updates to improve the decision-making of rational agents. Classical modal logic captures static belief states well, but does not account for the dynamics of belief updates. To address this gap, we have formalized belief updates with dynamic modal logic, providing a formal language to model chains of beliefs. We illustrate our approach by formalizing expansions, contractions, and revisions, and by proving the equivalence of two chains of beliefs that connect seemingly unrelated formulas.

1 Introduction

Beliefs are dynamic. They change every day. Throughout our lives, we talk, see, think — and our beliefs change. Understanding how rational agents update their beliefs is crucial to fields like Theory Change, AI, and Game Theory [1]. To formalize belief updates, Carlos Alchourrón, Peter Gärdenfors, and David Makinson developed AGM [2] — a formal framework that describes how rational agents update their beliefs when presented with new information.

AGM classifies belief updates into expansions, contractions, and revisions. Expansions are straightforward: an agent adds a new belief φ to its set of beliefs. In an expansion, φ does not conflict with any of the agent's previous beliefs, so the agent can believe φ. Contractions and revisions, however, are more subtle. In a contraction, an agent removes a belief φ from its set of beliefs, so the agent updates its set of beliefs so that they do not logically imply φ. In a revision, an agent adds a belief φ that is inconsistent with its set of beliefs, so the agent discards the beliefs inconsistent with φ, and only then believes φ.

To model these operations in a way interpretable by a computer, we use dynamic modal logic: an extension of classical logic that formalizes statements like "after an action α takes place, then a formula φ must hold" (see Remark 4.3 for why classical logic cannot express these statements). The formalization of these statements enables us to prove chains of beliefs, or meta-beliefs — statements like "if after learning φ an agent believes A, then after learning φ the agent must believe B". These meta-beliefs capture how beliefs relate to one another and evolve through updates, encapsulating the structure of belief.

With this goal in mind, let us now formally define belief states and how they evolve.

2 Belief States and their Dynamics

Let $\mathcal{L}$ be the set of all propositional formulas. A belief base is a consistent subset $B \subset \mathcal{L}$ that represents the explicit beliefs of a rational agent. To capture the idea that a rational agent believes all the logical consequences of its beliefs, we introduce the notion of an epistemic state.

Definition 2.1. An epistemic state K is the closure of a belief base B under logical consequence, denoted $Cn(B)$. In the modal logic setting, $Cn(B)$ also contains all modal consequences of B.

Example 2.2. Let $B = \{\varphi, \varphi \to \psi\}$ be the belief base of an epistemic state K. Then

$$K = Cn(B) = Cn(\{\varphi, \varphi \to \psi\}) = Cn(\{\varphi, \varphi \to \psi, \psi\}).$$

An epistemic state $K = Cn(B)$ represents all the beliefs of a rational agent. In AGM, a rational agent updates its beliefs through expansions, contractions, and revisions.

- An **expansion** adds to the epistemic state K a formula φ that is consistent with the belief base B. We denote an expansion by $K + \varphi$, and the resulting epistemic state is $K + \varphi = Cn(B \cup \{\varphi\})$.

- A **contraction** removes from the epistemic state K a formula φ while preserving as much of K as possible. We denote a contraction by $K - \varphi$, and the resulting epistemic state is $K - \varphi = Cn(B_{\text{new}})$ where B_{new} is a maximal subset of B such that $\varphi \notin Cn(B_{\text{new}})$.

- A **revision** adds to the epistemic state K a formula φ that is inconsistent with the belief base B. We denote a revision by $K * \varphi$, and the resulting epistemic state is $K * \varphi = Cn(B_{\text{new}} \cup \{\varphi\})$ where B_{new} is a maximal subset of B consistent with φ.

Definition 2.3. Let K be an epistemic state and φ a propositional formula. In AGM,

$$K * \varphi = (K - \neg\varphi) + \varphi.$$

Example 2.4. Let $K = \{\varphi \to \psi\}$ be an epistemic state.

- To **expand** K by φ, we compute

$$K + \varphi = Cn(\{\varphi \to \psi, \varphi\}) = Cn(\{\varphi \to \psi, \varphi, \psi\}).$$

- To **contract** $K + \varphi$ by ψ, notice that $K + \varphi = Cn(\{\varphi \to \psi, \varphi, \psi\})$ as above, and that removing ψ from $K + \varphi$ is not enough since ψ is a logical consequence of φ and $\varphi \to \psi$. There are therefore two possible results:

$$(K + \varphi) - \psi = \begin{cases} Cn(\{\varphi \to \psi\}) \\ Cn(\{\varphi\}) \end{cases}$$

  AGM uses a selection function to decide between the two possible results, but since we do not need it for our purposes, we avoid introducing it and consider both.

- To **revise** $K + \varphi$ by $\neg\psi$, we apply Definition 2.3 and the results above,

$$(K + \varphi) * \neg\psi = ((K + \varphi) - \psi) + \neg\psi = \begin{cases} Cn(\{\varphi \to \psi, \neg\psi\}) \\ Cn(\{\varphi, \neg\psi\}) \end{cases}$$

We have now defined the basic operations of AGM. To reason about chains of beliefs and uncover the structure of belief, we must be able to model these operations in a computer. AGM defines expansions, contractions, and revisions using set theory — an elegant choice in pure mathematics, but difficult to implement in practice.

Propositional logic is a more practical starting point. It is simple, well-understood, and easy to model in a computer. But it is too limited for our purposes. Its statements are static. They are either true or false, without reference to change or context. To overcome this limitation, we turn to modal logic.

3 Modal Logic

Modal logic extends propositional logic by introducing the notion of possible worlds [3]. In modal logic, the truth of a statement depends on the context — or world — in which it is evaluated. This characteristic makes modal logic ideal to model belief dynamics, since each possible world can represent an epistemic state. Let us define the notion of a possible world more formally.

Definition 3.1. A frame $\mathcal{F}$ is a tuple (W, R) where W is a non-empty set representing the possible worlds and $R \subset W \times W$ is a relation representing the accessibility between worlds. That is, if w and v are two possible worlds in W, then wRv means that the world v is accessible from the world w.

Example 3.2. Let $\mathcal{F} = (W, R)$ be a frame with three possible worlds, $W = \{w_0, w_1, w_2\}$, and an accessibility relation $R \subset W \times W$ such that $w_0 R w_1$ and $w_0 R w_2$. An interpretation of this frame could go as follows. The world w_0 is the present, the world w_1 is a version of tomorrow in which I have a cup of hot chocolate, and the world w_2 is a version of tomorrow in which I do not have a cup of hot chocolate. The accessibility relation R illustrates that from the present world w_0, both worlds w_1 and w_2 are accessible, since tomorrow I may or I may not have a cup of hot chocolate.

Now that we have a precise setting, we need to define how formulas are verified, and as we said before, we want the verification process to take context into account.

Definition 3.3. A Kripke model M is a triple (W, R, V) where (W, R) constitutes a frame and $V : \mathcal{L} \times W \to \{0, 1\}$ is a function that assigns a truth value to each formula in every world. We write $M, w \models \varphi$, read M models φ at w, when $V(\varphi, w) = 1$.

Example 3.4. Building on Example 3.2, let HC be the proposition "I have a cup of hot chocolate." Then:

- $V(HC, w_1) = 1$, since I have a cup of hot chocolate in w_1. Thus, $M, w_1 \models HC$.
- $V(HC, w_2) = 0$, since I do not have a cup of hot chocolate in w_2. Thus, $M, w_2 \not\models HC$.
- $V(HC, w_0)$ is either 0 or 1, depending on whether I currently have a cup of hot chocolate.

Modal logic also extends propositional logic with two new operators: necessity ($\Box\varphi$) and possibility ($\Diamond\varphi$) [3]. A formula is necessary when it holds in every possible world accessible from the current world, and a formula is possible when it holds in some possible world accessible from the current world.

Definition 3.5. Let $M = (W, R, V)$ be a Kripke model and φ a propositional formula. Then

$$M, w \models \Box\varphi \iff M, v \models \varphi \text{ for every } v \in W \text{ such that } wRv$$
$$M, w \models \Diamond\varphi \iff M, v \models \varphi \text{ for some } v \in W \text{ such that } wRv$$

The following theorem captures the relation between necessity and possibility.

Theorem 3.6. Let M be a Kripke model and φ be a modal formula. Then

$$M, w \models \Box\varphi \iff M, w \models \neg\Diamond\neg\varphi.$$

Proof.

$$\begin{aligned} M, w \models \Box\varphi &\iff M, v \models \varphi \text{ for every } v \in W \text{ such that } wRv \\ &\iff \text{There is no } v \in W \text{ such that } wRv \text{ and } M, v \models \neg\varphi \\ &\iff M, w \not\models \Diamond\neg\varphi \\ &\iff M, w \models \neg\Diamond\neg\varphi \end{aligned}$$

□

To model belief updates within modal logic, we must relate possible worlds to epistemic states and introduce actions that mirror expansions, contractions, and revisions.

4 Dynamic Modal Logic

Dynamic modal logic extends modal logic by allowing actions to update possible worlds. In a world w, the formula $[\alpha]\varphi$ means that φ holds in every possible world resulting from applying the action α to w, while $\langle\alpha\rangle\varphi$ means that φ holds in at least one such possible world.

In dynamic modal logic, possible worlds and accessibility relations have a clear interpretation. Possible worlds represent different states of a system, while accessibility relations represent how actions transform a state into another. For example, the world v representing the present with a broken pen is accessible from the world w representing the present with an unbroken pen by applying an action α — such as breaking the pen — to w, denoted $wR_\alpha v$.

Definition 4.1. Let $M = (W, R, V)$ be a Kripke model and φ a modal formula. Then

$$\begin{aligned} M, w \models [\alpha]\varphi &\iff M, v \models \varphi \text{ for every } v \in W \text{ such that } wR_\alpha v \\ M, w \models \langle\alpha\rangle\varphi &\iff M, v \models \varphi \text{ for some } v \in W \text{ such that } wR_\alpha v \end{aligned}$$

Example 4.2. Elaborating on Example 3.2, the world w_1 (where I have a cup of hot chocolate) is accessible from the present world w_0 via at least two different accessibility relations: the relation R_α, where α represents making a cup of hot chocolate, and the relation R_β, where β represents buying a cup of hot chocolate.

Remark 4.3. The statement $[\alpha]\varphi$ is not equivalent to any formula ψ. Formulas describe truths in a model, while actions change the model. A formula cannot change the valuation of a model, so no formula ψ can be equivalent to the action "make φ true".

Since dynamic modal logic introduces a necessity and a possibility operator for each possible action α, it is sometimes insightful to write $[\alpha]$ as $\Box_\alpha$ and $\langle\alpha\rangle$ as $\Diamond_\alpha$. The same interpretations and equivalences apply as in standard modal logic. In particular, consider the following theorem.

Theorem 4.4. Let M be a Kripke model, α an action, and φ a modal formula. Then

$$M, w \models [\alpha]\varphi \iff M, w \models \neg\langle\alpha\rangle\neg\varphi.$$

holds.

Proof. By writing $[\alpha]$ as $\Box_\alpha$ and $\langle\alpha\rangle$ as $\Diamond_\alpha$, the equivalence follows from Theorem 3.6:

$$\begin{aligned} M, w \models [\alpha]\varphi &\iff M, w \models \Box_\alpha\varphi \\ &\iff M, w \models \neg\Diamond_\alpha\neg\varphi \\ &\iff M, w \models \neg\langle\alpha\rangle\neg\varphi. \end{aligned}$$

$\square$

5 Belief Updates with Dynamic Modal Logic

There are many different approaches to modeling belief updates with dynamic modal logic. Most papers formalize the AGM postulates as axioms and then translate them into dynamic modal logic [4]. This approach is mathematically precise, but too rigorous for our purposes, so instead we follow the Kripke–Lewis approach and model each AGM operation with a dynamic action [5].

It is natural to represent epistemic states as possible worlds and the three operations of AGM as distinct sets of actions. We denote expansions by φ as $[+\varphi]$ and $\langle+\varphi\rangle$, contractions by φ as $[-\varphi]$ and $\langle-\varphi\rangle$, and revisions by φ as $[*\varphi]$ and $\langle*\varphi\rangle$ [6]. The formula $[-\varphi]\psi$, for example, means that ψ holds in every epistemic state $K - \varphi$.

Definition 5.1. Let $M = (W, R, V)$ be a Kripke model where the world $w \in W$ represents the epistemic state K. Then

$$M, w \models [\times\varphi]\psi \iff M, v \models \psi \text{ for every } v \in W \text{ such that } wR_{\times\varphi}v$$
$$M, w \models \langle\times\varphi\rangle\psi \iff M, v \models \psi \text{ for some } v \in W \text{ such that } wR_{\times\varphi}v$$

where $R_{\times\varphi}$ updates $w = K$ to $v = K \times \varphi$ and $\times$ is either $+$, $-$, or $*$.

Example 5.2. Let M be the Kripke model where the world w represents the epistemic state $K = Cn(\{\varphi \to \psi\})$. The expansion of K by φ is

$$K + \varphi = Cn(K \cup \{\varphi\}) = Cn(\{\varphi \to \psi, \varphi\}) = Cn(\{\varphi \to \psi, \varphi, \psi\}),$$

so every world v such that $wR_{+\varphi}v$ represents the epistemic state $K + \varphi$. In particular, since $\psi \in K + \varphi$, we have that $M, v \models \psi$ in every v such that $wR_{+\varphi}v$, so $M, w \models [+\varphi]\psi$.

With the setting in place, we are now ready to model belief updates with dynamic modal logic. Our approach is particularly useful to state and prove meta-beliefs. As in the introduction, a meta-belief is a chain of beliefs. We call them meta-beliefs because they are claims about the relationships between beliefs, and therefore about the structure of belief itself.

Remark 5.3. Recall that $\varphi \to \psi$ is equivalent to $\neg\varphi \vee \psi$.

Example 5.4. $\vdash [+\varphi]\psi \to [+\varphi](\varphi \to \psi)$.

Proof. Although there is a simpler proof, we present the following to illustrate how to apply the machinery we have developed so far.

Suppose towards contradiction that there is a Kripke model $M = (W, R, V)$ in which the claim does not hold at some world $w_0 \in W$. Then

1. $M, w_0 \models \neg([+\varphi]\psi \to [+\varphi](\varphi \to \psi))$ — Assumption
2. $M, w_0 \models \neg(\neg[+\varphi]\psi \vee [+\varphi](\varphi \to \psi))$ — Remark 5.3 to (1)
3. $M, w_0 \models \neg\neg[+\varphi]\psi \wedge \neg[+\varphi](\varphi \to \psi)$ — $\neg(A \vee B) \to \neg A \wedge \neg B$ to (2)
4. $M, w_0 \models \neg\neg[+\varphi]\psi$ — $A \wedge B \to A$ to (3)
5. $M, w_0 \models [+\varphi]\psi$ — $\neg\neg A \to A$ to (4)
6. $M, w_0 \models \neg[+\varphi](\varphi \to \psi)$ — $A \wedge B \to B$ to (3)
7. $M, w_0 \models \neg\neg\langle+\varphi\rangle\neg(\varphi \to \psi)$ — Theorem 4.4 to (6)

8. $M, w_0 \models \langle +\varphi \rangle \neg(\varphi \to \psi)$ $\qquad$ $\neg\neg A \to A$ to (7)

9. $M, w_1 \models \neg(\varphi \to \psi)$ $\qquad$ Definition 4.1 to (8) where $w_0 R_{+\varphi} w_1$

10. $M, w_1 \models \neg(\neg\varphi \vee \psi)$ $\qquad$ Remark 5.3 to (9)

11. $M, w_1 \models \neg\neg\varphi \wedge \neg\psi$ $\qquad$ $\neg(A \vee B) \to \neg A \wedge \neg B$ to (10)

12. $M, w_1 \models \neg\psi$ $\qquad$ $A \wedge B \to B$ to (11)

13. $M, w_1 \models \psi$ $\qquad$ Definition 4.1 to (5) since $w_0 R_{+\varphi} w_1$

— a contradiction since $M, w_1 \models \psi$ and $M, w_1 \models \neg\psi$. Thus, M cannot exist and every model verifies the claim. □

Another advantage of our approach is that it handles naturally iterated belief updates [7]. An iterated belief update is a formula such as $[+\psi][-\varphi]\chi$. In the epistemic state K, this formula means that χ holds in every epistemic state resulting from expanding K by ψ and then contracting $K + \psi$ by φ. To reason about iterated belief updates, we first need to prove two intermediary results.

Proposition 5.5. The equivalence relations induced by expansions are transitive. That is,

$$wR_{+\varphi}v \text{ and } vR_{+\psi}u \implies wR_{+\varphi\wedge\psi}u.$$

Proof. Let w represent an epistemic state K, v represent $K + \varphi$, and u represent $(K + \varphi) + \psi$. Then $wR_{+\varphi}v$ and $vR_{+\psi}u$. Now, notice that $(K + \varphi) + \psi = K + (\varphi \wedge \psi)$, so $wR_{+\varphi\wedge\psi}u$. □

Theorem 5.6. Let $M = (W, R, V)$ be a Kripke model. Let $w \in W$ represent K and $v \in W$ represent $K + \psi$. Then:

$$M, w \models [+\varphi]\chi \implies M, v \models [+\varphi]\chi.$$

Proof. Suppose that $M, w \models [+\varphi]\chi$, and let $u \in W$ be any possible world accessible by expanding v by φ. That is, $vR_{+\varphi}u$. We want to show that χ holds in u. By assumption, w represents K and v represents $K + \psi$, so $wR_{+\psi}v$. By Proposition 5.5, the equivalence relations induced by expansions are transitive, so $wR_{+\psi}v$ and $vR_{+\varphi}u$ imply that $wR_{+\psi\wedge\varphi}u$. Since $M, w \models [+\varphi]\chi$, then χ holds in every possible world accessible from w, and since u is accessible from w by $wR_{+\psi\wedge\varphi}u$, then χ holds in u. □

We can now prove a meta-belief involving iterated belief updates.

Example 5.7. $\vdash [+\varphi]\chi \to [+\psi][+\varphi]\chi$.

Proof. Suppose towards contradiction that there exists a Kripke model $M = (W, R, V)$ in which the claim does not hold at some world $w_0 \in W$. Then

1. $M, w_0 \models \neg([+\varphi]\chi \to [+\psi][+\varphi]\chi)$ — Assumption
2. $M, w_0 \models \neg(\neg[+\varphi]\chi \vee [+\psi][+\varphi]\chi)$ — Remark 5.3 to (1)
3. $M, w_0 \models \neg\neg[+\varphi]\chi \wedge \neg[+\psi][+\varphi]\chi$ — $\neg(A \vee B) = \neg A \wedge \neg B$ to (2)
4. $M, w_0 \models \neg\neg[+\varphi]\chi$ — $A \wedge B \to A$ to (3)
5. $M, w_0 \models [+\varphi]\chi$ — $\neg\neg A \to A$ to (4)
6. $M, w_0 \models \neg[+\psi][+\varphi]\chi$ — $A \wedge B \to B$ to (3)
7. $M, w_0 \models \neg\neg\langle+\psi\rangle\neg[+\varphi]\chi$ — Theorem 4.4 to (6)
8. $M, w_0 \models \langle+\psi\rangle\neg[+\varphi]\chi$ — $\neg\neg A \to A$ to (7)
9. $M, w_1 \models \neg[+\varphi]\chi$ — Definition 4.1 to (8) where $w_0 R_{+\psi} w_1$
10. $M, w_1 \models \neg\neg\langle+\varphi\rangle\neg\chi$ — Theorem 4.4 to (9)
11. $M, w_1 \models \langle+\varphi\rangle\neg\chi$ — $\neg\neg A \to A$ to (10)
12. $M, w_2 \models \neg\chi$ — Definition 4.1 to (11) where $w_1 R_{+\varphi} w_2$
13. $M, w_1 \models [+\varphi]\chi$ — Theorem 5.6 to (5) since $w_0 R_{+\psi} w_1$ and $w_1 R_{+\varphi} w_2$
14. $M, w_2 \models \chi$ — Definition 4.1 to (13) since $w_1 R_{+\varphi} w_2$

— a contradiction since $M, w_2 \models \chi$ and $M, w_2 \models \neg\chi$. Thus, M cannot exist and every model verifies the claim. □

Our framework is now complete. We are now able to model belief updates within dynamic modal logic and to state and prove statements about meta-beliefs.

6 Conclusion

In summary, we have illustrated how to model belief updates in a way that is interpretable by a computer. We achieved this by representing expansions, contractions, and revisions as the indexed modalities $[+\varphi]$, $[-\varphi]$, and $[*]\varphi$ in dynamic modal logic. This framework allows us to state and prove statements about meta-beliefs — statements like "if a rational agent believes that learning α implies A, then it must believe that learning β implies B". These meta-beliefs capture how individual beliefs relate to one another — revealing the structure of belief itself.

In future work, we could apply our framework to model the epistemic state of a Machine Learning model in training, and automate our framework to map beliefs that are related. We could also expand our framework to model scenarios with probabilistic certainty.

References

[1] Peter Gärdenfors. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. Cambridge, MA: MIT Press, 1988.

[2] Peter Gärdenfors Carlos Alchourrón and David Makinson. "On the Logic of Theory Change: Partial meet contraction and revision functions". In: *Journal of Symbolic Logic* 50.2 (1985), pp. 510–530. DOI: 10.2307/2274239.

[3] Graham Priest. *An Introduction to Non-Classical Logic: From If to Is*. 2nd. Cambridge: Cambridge University Press, 2008.

[4] André Fuhrmann. "On the Modal Logic of Theory Change". In: (1991). Ed. by André Fuhrmann and Michael Morreau, pp. 259–281.

[5] Giacomo Bonanno. "A Kripke-Lewis Semantics for Belief Update and Belief Revision". In: (2023). arXiv: 2310.11506 [cs.LO]. URL: https://arxiv.org/abs/2310.11506.

[6] Giacomo Bonanno. "A Modal Logic Translation of the AGM Axioms for Belief Revision". In: (2025). arXiv: 2502.14176 [cs.LO]. URL: https://arxiv.org/abs/2502.14176.

[7] Adnan Darwiche and Judea Pearl. "On the Logic of Iterated Belief Revision". In: *Artificial Intelligence* 89.1 (1997), pp. 1–29. ISSN: 0004-3702. DOI: https://doi.org/10.1016/S0004-3702(96)00038-0. URL: https://www.sciencedirect.com/science/article/pii/S0004370296000380.