R-Mod: Minimal Structural Revision of S5 Epistemic Models

FENGJIE SUN, Xi'an University of Architecture and Technology, China

Revising what an agent knows in response to new information is a central problem in formal epistemology. In doxastic logics such as KD45, belief revision proceeds by reordering plausibility: the agent simply re-ranks which worlds it considers most credible. This strategy fails for S5 knowledge. Because knowledge is factive ($K\varphi \rightarrow \varphi$), an agent cannot come to know φ merely by finding φ-worlds more plausible; if the actual world falsifies φ, then $K\varphi$ remains unsatisfiable regardless of any reordering. Accommodating new modal information in S5 therefore requires genuine model transformation: adjusting the equivalence-based accessibility structure, the valuation, or both.

We develop R-Mod, a selection-based revision operator that realizes this transformation as *minimal structural repair*. Given an S5 model and a target formula, R-Mod searches for a closest model—measured by a bisimulation-aware distance on quotient structures—that satisfies the formula while preserving S5 constraints. At the skeptical level, R-Mod satisfies success, consistency preservation, and deductive closure; classical AGM postulates such as Inclusion and Superexpansion fail due to permissible structural amplification, though we identify conditions under which they re-emerge. Computationally, the decision problem is NP-complete, and we provide tractable fragments exploiting structural locality.

While recent work has advanced AGM-style postulate analysis for S5 and topological semantics via simplicial complexes, these approaches do not provide goal-driven optimization with algorithmic guarantees. R-Mod fills this gap by combining modal invariance, explicit distance minimization, and fine-grained complexity analysis. Our results reframe revision in S5 as *knowledge-model revision* rather than belief revision, offering a foundation for algorithmic implementations and extensions to richer epistemic semantics.

JAIR Associate Editor: Roberta Calegari

JAIR Reference Format:
Fengjie Sun. 2026. R-Mod: Minimal Structural Revision of S5 Epistemic Models. *Journal of Artificial Intelligence Research* 86, Article 7 (May 2026), 33 pages. DOI: 10.1613/jair.1.21589

1 Introduction

Epistemic agents must revise their informational states when confronted with new evidence that may be *factual* or *epistemic*. In epistemic logic, such states are commonly represented as Kripke models $M = \langle W, \sim, V \rangle$, where W is a finite set of possible worlds, V is a valuation assigning truth values to propositional atoms at each world, and $\sim$ is an accessibility relation encoding which worlds the agent cannot distinguish. In the **S5** setting, $\sim$ is an *equivalence relation*—reflexive, symmetric, and transitive—and hence induces a *partition* $W/\sim$ of worlds into equivalence classes. Each class represents a set of worlds the agent considers mutually indistinguishable; the partition as a whole captures the agent's granularity of discrimination.

A simple example illustrates the challenge of revising **S5** knowledge. Consider a model M with two worlds w_1, w_2 such that $w_1 \sim w_2$, where p is true at w_1 but false at w_2. Suppose the actual world is w_1. Then $M, w_1 \models p \land \neg Kp$: the proposition p holds, but the agent does not know it, because w_2—where p is false—remains epistemically accessible. Now suppose the agent receives evidence that should result in knowing p, i.e., the target is $\alpha := Kp$.

Author's Contact Information: Fengjie Sun, ORCID: 0000-0002-3778-9127, fengjie_sun@outlook.com, Xi'an University of Architecture and Technology, Xi'an, Shaanxi, China.

DOI: 10.1613/jair.1.21589

One natural response is *eliminative update*: remove all worlds that fail to satisfy p. This is precisely what Public Announcement Logic (PAL) does—announcing φ restricts the model to worlds satisfying φ, thereby eliminating epistemic possibilities (Plaza 2007). In our example, PAL would delete w_2, leaving only w_1, and Kp becomes true. However, eliminative updates are not always appropriate. In many scenarios, evidence is better understood as *corrective* rather than eliminative: measurement errors may need correction, sensors may require recalibration, or protocols may demand repair. In such cases, we want to *preserve the state space* while *modifying the structure*—either by refining the equivalence partition or by adjusting valuations—so that the result validates α.

This observation motivates the central theme of this paper: in **S5**, epistemic change is naturally modeled as *minimal structural repair* of the equivalence partition and, when necessary, the valuation, while preserving the S5 frame conditions throughout.

Two traditions of epistemic change. Two broad research traditions address how epistemic states change in response to new information.

The first is *belief revision* in the AGM tradition, named after Alchourrón, Gärdenfors, and Makinson (Alchourrón et al. 1985), who proposed rationality postulates governing how a belief set should change upon receiving new information. In modal settings, AGM-style revision is often modeled in **KD45** or in *plausibility models*: belief is *non-factive* (an agent may believe φ even when φ is false), and revision typically proceeds by reordering which worlds the agent considers most plausible. The key mechanism is *plausibility reordering*: to come to believe φ, the agent promotes φ-worlds to maximal plausibility, without necessarily changing any factual content. Recent work by Bonanno (2025a) provides a new Kripke-Lewis semantics characterizing each AGM axiom via frame properties, deepening the connection between syntactic postulates and model-theoretic constraints.

The second tradition is *dynamic epistemic logic* (DEL), where epistemic states are Kripke models and *actions* transform them according to specified protocols (Baltag, Moss, et al. 1998; Ditmarsch, Hoek, et al. 2007). DEL is highly expressive: *postconditions* allow actions to modify valuations, and complex multi-agent scenarios can be modeled via action models. Recent years have witnessed substantial progress on DEL-based epistemic planning. Bolander, Charrier, et al. (2020) systematically analyze decidability and complexity, showing that general epistemic planning is undecidable but identifying decidable fragments. Belle et al. (2023) survey the state of the art, while Bolander, Burigana, et al. (2025) propose depth-bounded algorithms achieving $(b+1)$-EXPTIME complexity for reasoning depth b. Muise et al. (2022) demonstrate efficient compilation into classical planning, and Cooper et al. (2021) introduce lightweight S5 fragments with improved complexity bounds. However, standard DEL assumes the action structure is given in advance; it does not directly provide a framework for *optimizing* which transformation to apply given only a target formula.

Gaps motivating our work. While AGM and DEL provide foundational frameworks, neither directly yields an *optimization semantics* for **S5** knowledge-model revision with the following desiderata: (i) preserving S5 frame conditions, (ii) allowing both partition repair and valuation repair when needed, and (iii) ensuring robustness under modal equivalence, so that the distance measure does not depend on accidental features of graph encodings. Three specific gaps motivate our work.

(i) **PAL is purely eliminative.** Public announcements can only remove worlds, never repair structure. When elimination is inappropriate—e.g., under noisy or corrective evidence—PAL offers no recourse.

(ii) **DEL does not optimize for a target.** DEL with postconditions can modify valuations, but the transformation is dictated by a pre-specified action schema, not selected to minimize distance to the original model given a target formula. There is no built-in notion of *minimal change* toward an arbitrary epistemic goal.

(iii) **AGM operates on non-factive belief.** In **KD45** or plausibility models, revision can succeed by reordering alone: to believe φ, promote φ-worlds. But **S5** knowledge is *factive*: $K\varphi \rightarrow \varphi$ is valid. If the actual world falsifies φ, no amount of reordering can make $K\varphi$ true—the actual world itself must change, or the partition must be reconfigured. Factivity thus blocks the standard plausibility-reordering route and demands genuine model transformation.

Recent approaches and remaining challenges. Recent work has begun addressing minimal change in **S5** directly. Aguilera-Ventura et al. (2025) develop explicit minimal-change postulates and distance-based constructions for belief change in **S5**, and related lines study realizability and limitations of AGM-style operators in epistemic spaces (Sauerwald and Thimm 2024). A parallel development investigates *simplicial models* for epistemic logic, initiated by Goubault, Ledent, et al. (2021) and extended to semi-simplicial sets (Goubault, Kniazev, Ledent, and Rajsbaum 2023), faulty agents (Goubault, Kniazev, Ledent, and Rajsbaum 2024), and chromatic hypergraphs (Goubault, Kniazev, and Ledent 2024). These topological approaches provide alternative semantics with applications to distributed computing, but do not address goal-driven optimization. A related but distinct development investigates epistemic update and belief revision in non-classical logics—particularly relevant and paraconsistent logics—where certain AGM postulates fail for reasons rooted in the underlying consequence relation rather than in structural amplification (Mares 2002; Punčochař et al. 2023; Vigiani 2025); we discuss these approaches in Section 2.5.

Our contribution is complementary: we treat *structural repair of equivalence partitions* as a first-class operation, enforce *modal invariance* by measuring minimality on bisimulation-quotiented representations, and provide fine-grained *algorithmic and complexity analysis* for model-level optimization.

The R-Mod operator. We propose *R-Mod*, a minimal **S5** knowledge-model revision operator. Given an **S5** model M and an input formula α, R-Mod selects revised models M' such that M' satisfies α at the designated world while minimizing a *bisimulation-aware symmetric-difference distance*. This distance measures two components: (i) how the equivalence partition changes, and (ii) how valuations change. Crucially, the distance is computed on a canonical *quotient* representation that collapses bisimilar (hence modally indistinguishable) states. This ensures that the optimization objective is sensitive only to what the agent can express in the modal language, not to superficial differences in graph encoding.

We develop both *choice semantics*, which returns the set of all minimal revised models, and *skeptical semantics*, which takes the intersection of consequences across all minimal models. Skeptical semantics is the default in our metatheoretic analysis.

Contributions.

(C1) **Construction.** We define R-Mod via selection semantics that jointly adjusts the S5 equivalence partition and, when necessary, valuations, while preserving S5. Minimality is measured by a bisimulation-aware symmetric-difference distance on quotient representations (Sect. 3).

(C2) **Metatheory.** For skeptical revision we establish success, consistency preservation, deductive closure, and extensionality under a formula-insensitive objective. A weak vacuity result holds when the input is already satisfied. Classical AGM postulates—Inclusion (K3), Superexpansion (K7), and Subexpansion (K8)—fail in general; we provide counterexamples that pinpoint *structural amplification* as the source of failure (Sect. 4).

(C3) **Algorithms and complexity.** We prove NP-completeness of the associated decision problem and NP-hardness of exact minimum-distance computation. We analyze an exhaustive search procedure and two restricted variants: a valuation-fixed fragment with reduced complexity, and an α-guided single-class refinement that achieves polynomial time by enforcing locality of change (Sect. 5).

(C4) **Implementation and empirical evaluation.** We implement the core algorithms and report experiments on synthetic and benchmark-style families of **S5** instances, quantifying scalability and the trade-off between partition edits and valuation edits (Sect. 6).

(C5) **Positioning.** We relate R-Mod to PAL/DEL (including postconditions), plausibility-based belief revision, recent minimal-change accounts for **S5**, and simplicial models, clarifying the advantages, limitations, and design trade-offs of each approach (Sect. 2).

Summary and outlook. The overall picture is that **S5** knowledge-model revision admits a transparent optimization view reconciling minimal structural change with modal invariance. Our analysis explains which AGM principles transfer to this setting and which do not, once factivity and partition reconfiguration are admitted. We view R-Mod as a foundation for extensions beyond **S5**—for instance, to **KD45** or plausibility models by replacing partitions with preorders—and for scalable solvers that exploit locality in partition structure.

Paper organization. Sect. 2 reviews related work. Sect. 3 presents the operator, metric, and baseline algorithm. Sect. 4 develops metatheoric properties and negative results for selected AGM schemata. Sect. 5 establishes complexity bounds and analyzes tractable variants. Sect. 6 reports empirical evaluation. Sect. 7 discusses limitations and scope. Sect. 8 summarizes and outlines future directions.

2 Related Work

We situate our contribution within four streams of research on epistemic change, highlighting for each the core mechanisms, the gaps that motivate our work, and the precise differences from R-Mod.

2.1 Public Announcement Logic and Eliminative Updates

Public Announcement Logic (PAL) provides the simplest model-theoretic account of epistemic change (Plaza 2007; Wang and Cao 2013). Announcing φ restricts the model to the submodel induced by φ-worlds; when the original model is S5, the result remains S5. PAL has been extended to group announcements, arbitrary announcements, and protocols that specify sequences of updates (Ågotnes et al. 2010; Balbiani et al. 2008; Benthem 2006). Van Ditmarsch (Ditmarsch 2023) provides a comprehensive recent survey of public announcement logic and its variants.

PAL is *purely eliminative*: it can only remove worlds, never add structure or modify valuations. This suffices when evidence is veridical and the goal is to rule out possibilities, but fails when evidence is *corrective*—e.g., sensor recalibration, measurement error, or protocol repair—where the appropriate response is to *fix* the model rather than shrink it. Moreover, if the announced formula φ is false at every accessible world, PAL yields the empty model, offering no graceful recovery.

R-Mod permits both partition refinement and valuation adjustment, treating revision as *minimal structural repair* rather than elimination. When the input formula is already satisfied, R-Mod and PAL coincide (both preserve the model); when it is not, R-Mod searches for a closest satisfying model rather than deleting worlds wholesale.

2.2 Plausibility Models and Belief Revision

A major tradition models belief revision via *plausibility structures*: Kripke models equipped with a preorder $\leq$ ranking worlds by comparative plausibility (Baltag and Smets 2008; Benthem 2007; Board 2004). Belief corresponds to truth in all $\leq$-minimal worlds; revision promotes φ-worlds to maximal plausibility, implementing AGM-style change (Alchourrón et al. 1985; Grove 1988). Related systems include conditional doxastic logic (Baltag and Smets 2006), dynamic doxastic logic (DDL) (Leitgeb and Segerberg 2007; Segerberg 1995), and Spohn-style ranking functions (Spohn 1988).

Recent work has deepened the connection between AGM theory and modal logic. Bonanno (Bonanno 2025a) provides a new Kripke-Lewis semantics for belief update and revision, characterizing each AGM axiom via frame properties. In a companion paper (Bonanno 2025b), Bonanno establishes precise correspondences between AGM axioms and modal axioms in a logic with belief, conditional, and global operators, offering a principled translation between syntactic belief revision and modal reasoning.

Plausibility-based revision operates on *non-factive* belief in KD45 or weaker systems. Because belief does not entail truth, revision can succeed by reordering alone: to believe φ, promote φ-worlds without changing any facts. This mechanism *cannot* be transferred to S5 knowledge. Factivity ($K\varphi \rightarrow \varphi$) entails that if the actual world falsifies φ, no plausibility reordering can make $K\varphi$ true—the actual world itself must change, or the equivalence structure must be reconfigured. Plausibility models thus address a fundamentally different revision problem.

R-Mod targets *factive* knowledge in S5, where the accessibility relation is an equivalence (not a preorder) and revision requires genuine model transformation. The optimization objective is partition reconfiguration plus valuation change, not plausibility reordering. Consequently, AGM postulates that hold for plausibility revision (e.g., Inclusion, Superexpansion) may fail for R-Mod due to *structural amplification*—a phenomenon absent in non-factive settings.

2.3 Dynamic Epistemic Logic with Postconditions

Dynamic Epistemic Logic (DEL) generalizes PAL by introducing *action models* that specify complex multi-agent updates (Baltag, Moss, et al. 1998; Ditmarsch, Hoek, et al. 2007). Crucially, DEL supports *postconditions* that can modify valuations after an action is executed, enabling fact-changing events such as opening a door or flipping a switch (Benthem et al. 2006; Ditmarsch 2005). This makes DEL highly expressive for modeling protocols, games, and planning under epistemic uncertainty (Bolander and Andersen 2011; Engesser et al. 2017).

Recent work on DEL-based epistemic planning has made significant progress on complexity and decidability. Bolander, Charrier, Pinchinat, and Schwarzentruber (Bolander, Charrier, et al. 2020) provide a systematic overview of complexity results, showing that general epistemic planning is undecidable but identifying decidable fragments with propositional preconditions. Belle, Bolander, Herzig, and Nebel (Belle et al. 2023) survey the state of the art, emphasizing diverse approaches and open problems. Bolander, Burigana, and Montali (Bolander, Burigana, et al. 2025) propose depth-bounded epistemic planning with canonical bisimulation contractions, achieving (b+1)-EXPTIME complexity for reasoning depth b. Muise et al. (Muise et al. 2022) demonstrate efficient compilation of epistemic planning into classical planning for restricted proper epistemic knowledge bases. Cooper et al. (Cooper et al. 2021) introduce a lightweight fragment of S5 with NP-complete satisfiability and PSPACE-complete plan existence.

DEL assumes the action structure is *given in advance*: the modeler specifies which preconditions trigger which postconditions. DEL does not provide a mechanism for *selecting* the transformation that achieves a target formula with minimal change. Given only a goal α, DEL offers no built-in way to search over possible updates and return the closest one. This is a modeling gap, not an expressiveness limitation: DEL can represent any finite transformation, but the *optimization* layer must be supplied externally.

R-Mod takes a target formula α and *computes* a minimal transformation, whereas DEL requires the transformation to be specified. The two approaches are complementary: R-Mod could be used as a *synthesis* procedure that outputs a DEL action model achieving α with minimal distance, effectively compiling goal-driven revision into action-driven dynamics.

2.4 Minimal Change in S5: Recent Developments

Recent work has begun addressing minimal change directly within the S5 setting. Aguilera et al. (Aguilera-Ventura et al. 2025) develop explicit minimal-change postulates and distance-based constructions for belief change in S5, studying which AGM-like principles can be recovered. Sauerwald and Thimm (Sauerwald and Thimm 2024) investigate realizability constraints on AGM-style operators in epistemic spaces. Van Ditmarsch (Ditmarsch 2005) provides foundational analysis of model-based revision in modal frameworks.

A particularly active line of research investigates *simplicial models* for epistemic logic, initiated by Goubault, Ledent, and Rajsbaum (Goubault, Ledent, et al. 2021). This approach uses higher-dimensional combinatorial structures (simplicial complexes) rather than Kripke models, providing natural representations for distributed computing scenarios where processes may crash. Goubault et al. (Goubault, Kniazev, Ledent, and Rajsbaum 2023) extend this to semi-simplicial set models for distributed knowledge, defining a new semantics where groups of agents may distinguish worlds that no individual agent can distinguish. The same authors (Goubault, Kniazev, Ledent, and Rajsbaum 2024) address epistemic logic for faulty agents using impure simplicial complexes, systematically classifying design choices and axiomatizing the corresponding logics. Van Ditmarsch and Kuznets (Ditmarsch and Kuznets 2025) provide a complete axiomatization for impure simplicial complexes with three-valued semantics. Goubault, Kniazev, and Ledent (Goubault, Kniazev, and Ledent 2024) generalize further to chromatic hypergraphs with a many-sorted epistemic logic that distinguishes participating agents from the environment. These topological approaches provide alternative semantics that may interact fruitfully with model-based revision, though they do not address the optimization problem central to R-Mod.

Existing S5-specific accounts focus primarily on *postulate-level* analysis: which rationality constraints can or cannot hold. They do not typically enforce *modal invariance* of the distance metric (i.e., invariance under bisimulation), nor do they provide fine-grained *algorithmic and complexity analysis* at the model level. The computational landscape of S5 knowledge-model revision—decision complexity, tractable fragments, practical algorithms—remains underexplored.

Our contribution is complementary along three dimensions:

(a) **Structural repair as first-class operation.** We treat equivalence-class reconfiguration and valuation adjustment as jointly optimizable, not as separate or secondary concerns.
(b) **Modal invariance.** Our distance metric is computed on bisimulation quotients, ensuring that the optimization objective depends only on modal content, not on accidental features of graph encoding.
(c) **Algorithmic depth.** We provide complexity results (NP-completeness, NP-hardness of exact optimization), identify tractable fragments, and report empirical evaluation on synthetic benchmarks.

2.5 Non-Classical Belief Revision and Relevant Epistemic Updates

A growing body of work investigates belief revision and epistemic updates in non-classical logical settings, particularly paraconsistent and relevant logics. These approaches are pertinent to the present work for two reasons: they reject or weaken some of the same AGM principles that R-Mod fails to satisfy, and some of them employ update mechanisms that modify the accessibility relation rather than eliminating worlds.

Mares (2002) develops a paraconsistent theory of belief revision grounded in relevant logic R. The central move is to replace the AGM consistency-maintenance requirement with a weaker property called *coherence*: an agent maintains separate sets of accepted and rejected propositions, and revision preserves their non-overlap rather than classical consistency. This yields a framework in which contradictory beliefs are tolerable, fundamentally altering which AGM postulates survive. R-Mod differs in retaining a classical S5 base logic, but the parallel is instructive: both frameworks demonstrate that principled departures from AGM arise naturally once the underlying logic deviates from the classical propositional setting that AGM was designed for.

Punčochař et al. (2023) study public announcement logic with common knowledge based on the relevant logic R. A notable feature of their semantics is that public announcements need not be truthful and need not be accepted by all agents; moreover, the update mechanism operates by *modifying the accessibility relation* rather than eliminating worlds from the model. This contrasts sharply with standard (classical) PAL, which is purely eliminative. The relation-changing mechanism shares a structural affinity with R-Mod: both frameworks treat the accessibility relation as a mutable component of the epistemic model. However, the motivations differ: relevant

PAL modifies relations to accommodate the non-classical entailment structure of R, whereas R-Mod modifies the S5 equivalence partition to achieve a specified revision target with minimal distance.

Vigiani (2025) introduces a dynamic extension of the logic of relevant reasoners in classical worlds. Dynamic modalities model belief updates, and a binary conditional operator captures Ramsey's idea that conditional reasoning involves a form of belief revision. The system combines classical propositional logic at the outer level with relevant implication for the non-monotonic aspects of revision. This two-layered architecture—classical environment, relevant reasoning—offers a perspective complementary to R-Mod's purely classical setting: while R-Mod achieves non-monotonicity through structural amplification within S5, Vigiani's system builds non-monotonicity into the conditional via relevant implication.

These non-classical approaches highlight that the failure of certain AGM postulates is not unique to R-Mod but is a recurring phenomenon whenever the revision framework departs from the propositional setting for which AGM was originally formulated—whether by changing the base logic (paraconsistent or relevant) or by changing the revision mechanism (partition repair rather than plausibility reordering).

2.6 Summary: Positioning of R-Mod

Table 1 summarizes the comparison across key dimensions.

Table 1. Comparison of epistemic revision approaches.

Approach	**Target logic**	**Mechanism**	**Optimization**	**Valuation change**
PAL	S5	World elimination	No	No
Plausibility models	KD45	Plausibility reordering	Implicit	No
DEL + postconditions	S5/KD45	Action models	No (given)	Yes (specified)
Simplicial models	$S5_n$	Topological	No	No
Bonanno 2025	KD/S5	Frame-based	Implicit	Limited
Aguilera et al. 2025	S5	Postulate-based	Distance-based	Limited
Relevant PAL	R-based epistemic	Relation change	No	No
R-Mod	S5	Partition + valuation repair	Yes (explicit)	Yes (optimized)

Advantages of R-Mod.

- **Flexibility.** Handles both corrective and eliminative scenarios by allowing partition refinement, coarsening, and valuation adjustment.
- **Modal robustness.** Bisimulation-aware distance ensures that revision is insensitive to representational artifacts.
- **Goal-driven.** Given only a target formula, R-Mod computes a minimal revision without requiring a pre-specified action schema.
- **Analyzable.** Clear complexity bounds and tractable fragments guide practical implementation.

Limitations and trade-offs.

- **Computational cost.** The general decision problem is NP-complete; exact optimization is NP-hard. For large models, heuristic or approximate methods may be necessary.
- **Single-agent focus.** The current framework addresses single-agent S5; extension to multi-agent settings with common knowledge requires additional machinery.
- **No iterated revision.** We treat one-shot revision; a full account of iterated revision (revision policies, history dependence) is left for future work.

- **Distance choice.** The symmetric-difference distance is natural but not unique; alternative metrics (e.g., weighted, asymmetric) may be preferable in specific applications.

Methodological remarks on S5 and logical omniscience. A referee may wonder whether the S5 semantics and the identification $\mathrm{Bel}(M) := \mathrm{Th}(M)$ commit us to unrealistically strong notions of knowledge and belief. We offer three clarifications.

(i) **S5 as a deliberate scope restriction.** The present paper targets S5 precisely because its factive and fully introspective character creates a revision problem that *cannot* be solved by plausibility reordering. Weaker logics such as KD45 admit different revision strategies—plausibility-based methods work well there—and are noted as a natural extension (Future Direction F1). Adopting S5 is thus a *scope choice*, not a claim that S5 exhausts rational epistemic states.

(ii) $\mathrm{Bel}(M) = \mathrm{Th}(M)$ **as an analytical device.** Defining the belief set as the full theory of the model is standard practice in model-theoretic analyses of epistemic change (Gärdenfors 1988; Hansson 1999). It provides a mathematically tractable object against which postulates can be evaluated. This definition does *not* assert that real agents compute or store infinite deductive closures; it is an idealization that isolates the logical structure of revision from resource bounds. The well-known *logical omniscience* problem—that Kripke-style agents believe all logical consequences of their beliefs—affects virtually all modal epistemic logics and is orthogonal to the specific contribution of R-Mod.

(iii) **Orthogonality to AGM failures.** The failure of AGM postulates K3, K7, and K8 documented in Section 4 arises from *structural amplification*: satisfying a modal target may require merging or splitting equivalence classes in ways that introduce consequences beyond simple expansion. This phenomenon is independent of whether the underlying logic is S5 or weaker, and independent of whether belief sets are idealized as deductively closed. Analogous failures would appear in any partition-based revision operator that allows structural change, regardless of the strength of the epistemic assumptions.

In sum, our use of S5 and $\mathrm{Bel}(M) = \mathrm{Th}(M)$ follows established methodological conventions in the belief-revision and dynamic-epistemic-logic literatures. Extensions to bounded or resource-sensitive agents are a worthwhile direction but lie outside the present scope.

Complementarity. R-Mod is not intended to replace PAL, DEL, or plausibility-based revision, but to fill a specific niche: *goal-driven minimal repair of S5 knowledge models*. When elimination suffices, PAL is simpler; when the action structure is known, DEL is more expressive; when belief rather than knowledge is at stake, plausibility models are appropriate. R-Mod addresses the case where none of these conditions hold: the agent has factive knowledge, the evidence is corrective, and the goal is to find the closest satisfying model without a pre-specified transformation.

3 Construction of the R-Mod Operator for S5 Knowledge-Model Revision

This section defines *R-Mod*, a minimal *structural* revision operator for S5 *knowledge models*. The operator takes a finite S5 model $M = \langle W, R, V \rangle$ and a modal sentence α, and returns the set of S5 models that (i) satisfy α, (ii) preserve S5 properties, and (iii) deviate from M by a minimal amount according to a graph–valuation distance. The perspective is *model correction*: we allow controlled changes to the equivalence partition R and to propositional valuations V. Truthful-announcement dynamics (which only delete worlds) are not our focus here.

S5 system: syntax, axiomatics, semantics. Language. $\varphi ::= p \mid \neg\varphi \mid (\varphi \wedge \varphi) \mid \Box\varphi$, where p ranges over Prop; define $\Diamond\varphi := \neg\Box\neg\varphi$.

Axiom schemata.

Axiom schemata (S5).

$$\begin{aligned} &\text{(Prop) All propositional tautologies.}\\ &\text{(K) } \Box(\varphi \to \psi) \to (\Box\varphi \to \Box\psi).\\ &\text{(T) } \Box\varphi \to \varphi.\\ &\text{(4) } \Box\varphi \to \Box\Box\varphi.\\ &\text{(5) } \Diamond\varphi \to \Box\Diamond\varphi.\\ &\text{(Dual) } \Diamond\varphi \leftrightarrow \neg\Box\neg\varphi. \end{aligned}$$

Inference rules.

$$\text{(MP)}\ \frac{\varphi \quad \varphi \to \psi}{\psi} \qquad \text{(Nec)}\ \frac{\varphi}{\Box\varphi}.$$

The duality law $\Diamond\varphi \leftrightarrow \neg\Box\neg\varphi$ is available by definition or as an additional axiom.

Kripke semantics. A model is $M = \langle W, R, V \rangle$ with nonempty W, accessibility $R \subseteq W \times W$, and valuation $V : \text{Prop} \to 2^W$. Truth is defined inductively:

$$\begin{aligned} M, w \models p &\iff w \in V(p),\\ M, w \models \neg\varphi &\iff \neg(M, w \models \varphi),\\ M, w \models \varphi \wedge \psi &\iff (M, w \models \varphi) \wedge (M, w \models \psi),\\ M, w \models \Box\varphi &\iff \forall v \in W\ (R(w, v) \Rightarrow M, v \models \varphi). \end{aligned}$$

Global truth $M \models \varphi$ abbreviates $\forall w \in W.\, M, w \models \varphi$. For S5, R is an *equivalence relation* (reflexive, symmetric, transitive). S5 is sound and complete with respect to frames whose R is an equivalence relation. The relation R induces a partition of W into equivalence classes $[w]_R = \{v \in W \mid R(w, v)\}$.

Standing assumptions and their theoretical import. We collect several technical assumptions that govern the remainder of the paper, together with brief explanations of their theoretical significance.

(A1) **Finiteness.** All models are finite: $|W| < \infty$.
Import. Finiteness ensures that the set of admissible perturbations is finite (the number of partitions of W is the Bell number $B_{|W|}$, and the number of valuations on $\text{Var}(\alpha)$ is $2^{|W|\cdot|\text{Var}(\alpha)|}$). Consequently, the distance D° takes values in a bounded subset of $\mathbb{N}$, and the minimization problem is well-defined with a guaranteed optimum. For S5 with a finite propositional signature, the finite model property holds: any satisfiable formula has a finite model (Blackburn et al. 2001). Thus, restricting attention to finite models incurs no loss of generality for satisfiability-based revision. Extending R-Mod to infinite models would require topological or measure-theoretic distance notions, which we leave for future work.

(A2) **Global versus local satisfaction.** We write $M \models \alpha$ (global truth) for $\forall w \in W.\, M, w \models \alpha$, and $M, w_0 \models \alpha$ (local truth) for satisfaction at a designated world w_0.
Import. In S5, the equivalence relation partitions W into classes; all worlds within a class satisfy the same modal formulas. For revision, the choice between global and local satisfaction affects *which* classes must be modified to achieve α, but not the *structure* of the optimization problem: in both cases, one searches for a minimal-distance admissible model satisfying the relevant constraint. Our metatheoretic results (success, consistency, closure) hold under either interpretation because they depend on the minimality criterion rather than on the scope of satisfaction. We adopt global truth as the default; a designated-world variant restricts α-satisfaction to the class containing w_0 and is straightforwardly definable.

(A3) **Vocabulary restriction.** The valuation component of D° is restricted to atoms in $\mathrm{Var}(\alpha)$.
Import. Atoms outside $\mathrm{Var}(\alpha)$ do not affect whether α holds. Including them would allow irrelevant valuation flips to dominate the distance, obscuring the cost of achieving the revision goal. The restriction thus implements a *relevance* principle[1]: the cost of revision is measured only with respect to the target's vocabulary. A consequence is that R-Mod does not penalize collateral changes to atoms outside $\mathrm{Var}(\alpha)$. If such changes matter in an application, a weighted or vocabulary-extended distance is easily definable.

(A4) **S5 frame preservation.** The revised relation R' must be an equivalence relation on W.
Import. Preserving equivalence is essential for maintaining S5 validity. Reflexivity guarantees factivity ($K\varphi \to \varphi$); symmetry and transitivity together guarantee positive and negative introspection ($K\varphi \to KK\varphi$ and $\neg K\varphi \to K\neg K\varphi$). Relaxing any of these properties would move the agent out of the S5 class, changing the logic itself. R-Mod is thus a *within-logic* revision operator: it modifies the agent's epistemic state while preserving the frame class. Cross-logic revision—e.g., transitioning from S5 to KD45—is a distinct problem requiring different machinery.

The theory of a model is $\mathrm{Th}(M) = \{\varphi \mid M \models \varphi\}$, and the deductive expansion by α is $M+\alpha := \mathrm{Cn}(\mathrm{Th}(M) \cup \{\alpha\})$. For a formula α, let $\mathrm{Var}(\alpha)$ denote the set of propositional atoms occurring in α.

Bisimulation and quotient models. Two states in a Kripke model are *bisimilar* if they cannot be distinguished by any modal formula. Formally, a *bisimulation* between models $M = \langle W, R, V\rangle$ and $M' = \langle W', R', V'\rangle$ is a non-empty relation $Z \subseteq W \times W'$ such that whenever $w\ Z\ w'$:

(B1) **Atomic harmony:** $w \in V(p) \iff w' \in V'(p)$ for all $p \in \mathrm{Prop}$;
(B2) **Zig:** if $R(w, v)$, then there exists $v' \in W'$ with $R'(w', v')$ and $v\ Z\ v'$;
(B3) **Zag:** if $R'(w', v')$, then there exists $v \in W$ with $R(w, v)$ and $v\ Z\ v'$.

Within a single model M, the union of all bisimulations on M is itself a bisimulation—the *largest bisimulation* $\sim_M$. The *bisimulation quotient* $Q(M) = M/{\sim_M}$ collapses each $\sim_M$-equivalence class to a single representative, yielding the smallest model bisimilar to M. Bisimilar models satisfy exactly the same modal formulas (Blackburn et al. 2001), so $Q(M)$ preserves all modal properties of M while eliminating redundant structure.

To respect modal invariance, distance is computed on the bisimulation quotient. We write $Q(M)$ for the quotient of M by $\sim_M$.

Design rationale. The notion of admissible perturbation formalizes the space of candidate revised models. Three considerations guide its design:

(i) **Preservation of S5 discipline.** Any revised model must remain a valid S5 structure; hence we require R' to be an equivalence relation. This constraint is non-negotiable: relaxing it would violate the factivity and introspection axioms that define S5 knowledge.
(ii) **Fixed carrier.** The base operator keeps the set of worlds W unchanged. This choice bounds the search space and ensures that revision is *repair* rather than *replacement*: we reconfigure the existing epistemic structure rather than constructing an entirely new one. An extended operator R-Mod$^+$ that allows $|W'| \neq |W|$ is definable by adding a $|W \triangle W'|$ penalty, but we defer this to future work.
(iii) **Target satisfaction.** The semantic constraint $M' \models \alpha$ ensures that revision succeeds: the output model validates the input sentence. Combined with (i), this guarantees that the revised knowledge state is both logically coherent and informationally adequate.

[1]Our use of "relevance" here is informal and vocabulary-based: changes are penalized only for atoms mentioned in α. This should not be confused with the technical notion of *relevance* in relevant logic (Mares 2002; Punčochař et al. 2023), where relevance is a property of the entailment relation itself—requiring that premises and conclusions share propositional variables—and has deep consequences for the underlying consequence operation. See Section 2.5 for a discussion of relevant-logic-based approaches to epistemic revision.

With these constraints in place, we now define admissible perturbations formally. Let $M = \langle W, R, V \rangle$ be an S5 model, where R is an equivalence relation on W and $V : \mathrm{Prop} \to 2^W$. An *admissible perturbation* of M is any model $M' = \langle W, R', V' \rangle$ such that:

(P1) R' is an equivalence relation on W (reflexive, symmetric, transitive);
(P2) V' is a valuation on W (same carrier; this base version fixes W);
(P3) M' is S5.

Intuitively, admissible changes reconfigure the partition induced by R (merging or splitting equivalence classes) and, when necessary, flip truth values of atoms at worlds so that α becomes true, all while keeping the S5 discipline.

Motivation for the quotient-based distance. A distance measure for model revision must satisfy two desiderata: it should be *semantically meaningful* (penalizing changes that affect what the agent can express) and *modally invariant* (insensitive to superficial differences in graph encoding). We achieve both by defining distance on *bisimulation quotients.*

(i) **Symmetric difference.** The symmetric difference $|\bar{R} \triangle \bar{R}'|$ counts accessibility links added or removed in the quotient structure. In the original model, merging two R-classes adds cross-class edges and may collapse quotient nodes; splitting a class removes edges and may introduce new quotient nodes. Similarly, $|\bar{V}(p) \triangle \bar{V}'(p)|$ counts representatives at which the truth of p is flipped. Symmetric difference thus provides a uniform currency for structural and valuational edits.

(ii) **Restrict to** $\mathrm{Var}(\alpha)$**.** Atoms outside $\mathrm{Var}(\alpha)$ do not affect whether α holds. Including them would allow irrelevant flips to dominate the distance, obscuring the cost of achieving the revision goal. The restriction ensures that minimality is measured relative to the task at hand.

(iii) **Quotients.** Two S5 models may differ syntactically yet be *bisimilar*—hence modally indistinguishable. A distance defined on raw models would penalize such inessential differences. By collapsing each model to its bisimulation quotient $Q(M)$, we ensure that $D^{\circ}(M, M') = 0$ whenever M and M' validate exactly the same modal formulas. This is the modal-invariance property: the optimization objective depends only on what the agent can express, not on accidental features of representation.

Minimality as the suitability criterion. R-Mod is a *prioritized* revision operator: the incoming formula α is treated as non-negotiable, and the model is adjusted to accommodate it. This design mirrors the AGM Success postulate (K2), which requires $\alpha \in \mathrm{Bel}_{\star}(M, \alpha)$ unconditionally. However, unconditional Success is not always a faithful model of rational epistemic change. When α conflicts sharply with well-entrenched prior beliefs, a rational agent may prefer to *reject* α rather than restructure the model—a response captured by *non-prioritized* (or *screened*) revision operators that allow the input to be partially or wholly discarded (Gärdenfors 1988; Hansson 1999). R-Mod does not model this possibility: within its scope, α is always accepted, and the operator seeks the least disruptive accommodation. This is a deliberate design choice, not a claim that prioritized revision is universally appropriate. A non-prioritized variant—for instance, one that includes the unchanged model as a candidate and charges a finite penalty for rejecting α, or one that compares the cost of accommodation against an entrenchment threshold—would address scenarios where the incoming information is defeasible. We note this as a direction for future work (cf. Section 7).

Given this prioritized commitment, two considerations justify the use of D° as the selection criterion within the feasible set.

First, when $\mathrm{Th}(M) \cup \{\alpha\}$ is inconsistent, *every* α-satisfying model must differ from M in at least some respects; there is no zero-distance option. The task is therefore not to find a model that resembles M in an absolute sense, but to find one that resembles M *as much as possible given the constraint.* A model at large distance is not intrinsically unsuitable; it is simply dominated by closer alternatives. If no closer alternative exists, then the large-distance model is, by definition, the most suitable available. Suitability is thus *relative to the feasible set*, not an absolute threshold.

Second, the quotient-based distance D° is not arbitrary: it is grounded in modal semantics (Section 3, preceding paragraphs). Changes to the equivalence partition correspond to gains or losses of knowledge distinctions; valuation flips correspond to factual corrections. Minimizing D° therefore minimizes the *epistemic disruption* caused by revision—the number of distinctions redrawn plus the number of facts altered. A model at minimal distance is one that achieves the revision goal with the least epistemic side-effects.

To summarize: (i) R-Mod is a *prioritized* operator that presupposes α must be accepted; within this scope, distance-based selection identifies the least disruptive accommodation, but the option of rationally rejecting α lies outside the operator's design; (ii) suitability is relative to what is achievable, and distance-minimizers are by construction the best achievable options; and (iii) the specific distance we use tracks semantically meaningful changes, ensuring that "closeness" corresponds to preservation of epistemic structure rather than syntactic accident.

Example 1 (Distance computation). *Let M have worlds $\{a, b, c\}$ with two R-classes $\{a, b\}$ and $\{c\}$, and let $V(p) = \{a, b\}$, $V(q) = \{c\}$. Then a and b are bisimilar (same propositional truth values, same accessible worlds), so $Q(M)$ has two nodes: one for $\{a, b\}$ and one for $\{c\}$.*

Suppose M' merges all worlds into a single class while keeping valuations unchanged. Then $Q(M')$ collapses further...

We now give the formal definition. Write $Q(M) = \langle \bar{W}, \bar{R}, \bar{V} \rangle$ and $Q(M') = \langle \bar{W}', \bar{R}', \bar{V}' \rangle$. Since quotients preserve S5 and modal truth, we define:

$$D^\circ(M, M') \;:=\; \left|\bar{R} \triangle \bar{R}'\right| \;+\; \sum_{p \in \mathrm{Var}(\alpha)} \left|\bar{V}(p) \triangle \bar{V}'(p)\right|,$$

where $\triangle$ is symmetric difference and $|\cdot|$ is set cardinality. The restriction to $\mathrm{Var}(\alpha)$ prevents irrelevant flips from affecting minimality. Weighted variants with coefficients for structural versus valuation changes are straightforward.

Given M and α, the *R-Mod* revision of M by α is the set

$$M *_{\text{R-Mod}} \alpha \;:=\; \arg\min\{\, D^\circ(M, M') \mid M' \text{ is an admissible perturbation and } M' \models \alpha \,\}.$$

When several candidates attain the minimum, the operator returns all of them (genuine non-determinism). A tie-breaking refinement that prefers candidates whose quotients are bisimilar to $Q(M)$ can be adopted when a single representative is needed.

By construction, R-Mod searches only among admissible S5 perturbations and selects those that are closest to the input model in the quotient space. This design yields a set of baseline guarantees that capture closure, success, and a vacuity-style preservation of prior consequences under consistency of the input: (i) *S5-closure:* every $M' \in M *_{\text{R-Mod}} \alpha$ is an S5 model. (ii) *Success:* $M' \models \alpha$ for all M' in the set. (iii) *Vacuity:* if $\mathrm{Th}(M) \cup \{\alpha\}$ is consistent, then $M + \alpha \subseteq \mathrm{Th}(M')$ for all $M' \in M *_{\text{R-Mod}} \alpha$. Equality need not hold in general. These statements are proved using the admissibility constraints and the quotient-based distance.

Termination and practical implementation. The loop over r in Algorithm 1 is a *conceptual device* that makes the minimality criterion explicit: we search outward from $r = 0$ until we find the first radius at which an α-satisfying admissible model exists. The loop is guaranteed to terminate for finite models because:

(i) The number of partitions of a finite set W is the Bell number $B_{|W|}$, which is finite;

(ii) The number of valuations on $\mathrm{Var}(\alpha)$ is $2^{|W|\cdot|\mathrm{Var}(\alpha)|}$, also finite;

(iii) Hence the set of admissible perturbations is finite, and D° takes values in a bounded subset of $\mathbb{N}$.

Algorithm 1 Minimal Epistemic Structural Adjustment (R-Mod)

Require: S5 model $M = \langle W, R, V \rangle$; modal sentence α
Ensure: Set $\mathcal{M}'$ of minimally revised S5 models satisfying α

$\tilde{M} \leftarrow Q(M)$ ▷ bisimulation quotient for distance evaluation
$\mathcal{M}' \leftarrow \emptyset$; $d_{\min} \leftarrow +\infty$
for $r = 0, 1, 2, \ldots$ **do** ▷ increasing perturbation radius
  $C_r \leftarrow$ all admissible $M' = \langle W, R', V' \rangle$ with $M' \models \alpha$ and $D^{\circ}(M, M') = r$
  if $C_r \neq \emptyset$ **then**
    $\mathcal{M}' \leftarrow C_r$; $d_{\min} \leftarrow r$; **break**
  end if
end for
return $\mathcal{M}'$

Let $r^* := \max\{D^{\circ}(M, M') \mid M' \text{ admissible}\}$. The loop terminates at some $r \leq r^*$ whenever an α-satisfying model exists; if no such model exists (e.g., α is S5-unsatisfiable), the loop runs through all radii and returns $\emptyset$.

In practice, one does not enumerate radii explicitly. Instead, admissible candidates are generated by solving a constrained optimization problem that enforces (i) the S5 equivalence constraints on R', (ii) valuation flips restricted to $\mathrm{Var}(\alpha)$, and (iii) the semantic constraint $M' \models \alpha$, while minimizing $D^{\circ}(M, M')$. Standard techniques include encoding as integer linear programming, SAT/MaxSAT, or branch-and-bound search over the partition lattice. The complexity analysis in Section 5 characterizes the hardness of this optimization.

4 Formal Properties of the R-Mod Operator

This section investigates the logical behavior of the R-Mod operator as a *knowledge-model revision* mechanism for S5. We work at two levels. At the *model level*, revision selects S5 models that satisfy the input sentence and are at minimal quotient-distance from the input model. At the *belief-set level*, we read off consequences from the selected models.

For an S5 model M, write $\mathrm{Bel}(M) := \mathrm{Th}(M) = \{\varphi \mid M \models \varphi\}$. Recall from Section 3 that the R-Mod operator returns the set of all minimal-distance admissible models satisfying α:

$$M *_{\text{R-Mod}} \alpha \;=\; \arg\min\{\, D^{\circ}(M, M') \mid M' \text{ admissible and } M' \models \alpha \,\}.$$

We refer to elements of this set as *minimally revised models*. The *skeptical* belief set induced by revision is

$$\mathrm{Bel}_{\star}(M, \alpha) \;:=\; \bigcap_{M' \in M *_{\text{R-Mod}} \alpha} \mathrm{Th}(M').$$

This is the object compared against AGM-style postulates below.

Foundational Metatheoretic Properties

We first record the foundational properties that any reasonable revision operator should satisfy.

Proposition 4.1 (Success). *For every $M' \in M *_{\text{R-Mod}} \alpha$, $M' \models \alpha$. Consequently, $\alpha \in \mathrm{Bel}_{\star}(M, \alpha)$.*

Proof. Let

$$C(M, \alpha) \;:=\; \{\, M' \mid M' \text{ is an admissible perturbation of } M \text{ and } M' \models \alpha \,\}.$$

By construction of the R-Mod operator, admissibility enforces the S5 constraints on R' and keeps the carrier W fixed (base version), while the semantic filter requires $M' \models \alpha$. Hence

$$\forall M' \in C(M, \alpha)\ (\, M' \models \alpha \,). \tag{1}$$

R-Mod selects the revised models by minimizing the quotient-based distance over the candidate class:

$$M *_{\text{R-Mod}} \alpha \;=\; \arg\min\{\, D^{\circ}(M,M') \mid M' \in C(M,\alpha) \,\}.$$

Therefore

$$M *_{\text{R-Mod}} \alpha \;\subseteq\; C(M,\alpha). \tag{2}$$

Combining (1) and (2) yields

$$\forall M' \in M *_{\text{R-Mod}} \alpha \;(\; M' \models \alpha \;),$$

which is the first claim.

If $C(M,\alpha) \neq \emptyset$, then a minimizer exists because D° takes values in $\mathbb{N}$ (finite symmetric differences on quotients) and is bounded below by 0. Hence $M *_{\text{R-Mod}} \alpha \neq \emptyset$. For every $M' \in M *_{\text{R-Mod}} \alpha$ we have $M' \models \alpha$, hence $\alpha \in \text{Th}(M')$. Taking intersections over the nonempty selection,

$$\alpha \in \bigcap_{M' \in M *_{\text{R-Mod}} \alpha} \text{Th}(M') \;=\; \text{Bel}_{\star}(M,\alpha).$$

In the degenerate case $C(M,\alpha) = \emptyset$ (e.g., α is S5-unsatisfiable), the first claim remains trivially true and the operator returns the empty selection; the skeptical set $\text{Bel}_{\star}(M,\alpha)$ is considered only when $M *_{\text{R-Mod}} \alpha \neq \emptyset$. □

Proposition 4.2 (Consistency preservation). *If α is S5-satisfiable, then* $\text{Bel}_{\star}(M,\alpha)$ *is consistent.*

Proof. We make the argument explicit in four steps. Throughout, admissibility requires that revised models remain S5 (equivalence relation on R') and share the same carrier W in the base operator.

Step 1: From satisfiability to global truth. Let α be S5-satisfiable, i.e., there exist an S5 model $N_0 = \langle U_0, S_0, V_0 \rangle$ and a world $u_0 \in U_0$ such that $N_0, u_0 \models \alpha$. We show that a model with *global* truth exists.

By the finite model property via filtration through $\text{Sub}(\alpha)$, we may assume N_0 is finite with at most $|\text{Sub}(\alpha)|$ equivalence classes. Let $[u_0]_{S_0}$ denote the equivalence class containing u_0. Define the *restriction* $N = N_0|_{[u_0]}$: its domain is $[u_0]_{S_0}$, its accessibility relation is $S_0 \cap ([u_0]_{S_0} \times [u_0]_{S_0})$, and its valuation is V_0 restricted to $[u_0]_{S_0}$.

Since S_0 is an equivalence relation, every world in $[u_0]_{S_0}$ accesses exactly the same set of worlds as u_0. By a standard generated-submodel argument, $N, u \models \varphi$ iff $N_0, u \models \varphi$ for all $u \in [u_0]_{S_0}$ and all $\varphi \in \text{Sub}(\alpha)$. In particular, $N, u \models \alpha$ for every $u \in [u_0]_{S_0}$, i.e., $N \models \alpha$ (global truth). The model N is S5 with a single equivalence class and at most $|\text{Sub}(\alpha)|$ worlds.

Step 2: Realizing N on W. Let k_α denote the number of worlds in N (equivalently, the size of the single class). If $|W| \geq k_\alpha$, we can *realize* N on W as follows: partition W into one block of size k_α and possibly additional singleton blocks; define R' to make the main block an R'-class (with other blocks as isolated classes if needed); copy the truth assignments from N to the main block. The resulting $M' = \langle W, R', V' \rangle$ is S5 and satisfies $M' \models \alpha$ globally.

Hence the candidate class

$$C(M,\alpha) = \{M' \mid M' \text{ admissible and } M' \models \alpha\}$$

is nonempty under the mild size-adequacy condition $|W| \geq k_\alpha$.

Step 3: Existence of a minimum. Since models are finite and distance is measured by symmetric difference on bisimulation quotients,

$$D^{\circ}(M,M') \in \mathbb{N} \quad \text{and} \quad D^{\circ}(M,M') \geq 0.$$

Therefore a minimum of D° over the nonempty finite $C(M,\alpha)$ exists, and

$$M *_{\text{R-Mod}} \alpha = \arg\min\{D^{\circ}(M,M') \mid M' \in C(M,\alpha)\}$$

is nonempty. Every $M' \in M *_{\text{R-Mod}} \alpha$ is S5 by admissibility, so $\text{Th}(M')$ is a consistent S5 theory.

Step 4: Consistency of the skeptical belief set. The skeptical belief set is the intersection of consistent complete S5 theories:

$$\mathrm{Bel}_\star(M,\alpha) = \bigcap_{M' \in M *_{\text{R-Mod}} \alpha} \mathrm{Th}(M').$$

If $\perp \in \mathrm{Bel}_\star(M,\alpha)$, then $\perp \in \mathrm{Th}(M')$ for every M' in the selection, contradicting Step 3. Hence $\mathrm{Bel}_\star(M,\alpha)$ is consistent.

Remark. If $|W| < k_\alpha$, the base operator may have an empty candidate class even though α is satisfiable. The extended operator R-Mod$^+$, which allows $|W'| > |W|$ and charges a $|W \triangle W'|$ penalty, guarantees nonemptiness and preserves the argument above. □

Proposition 4.3 (Deductive closure). *$\mathrm{Bel}_\star(M,\alpha)$ is closed under S5 consequence: if $\mathrm{Bel}_\star(M,\alpha) \models \varphi$, then $\varphi \in \mathrm{Bel}_\star(M,\alpha)$.*

Proof. Let $\mathcal{S} := M *_{\text{R-Mod}} \alpha$ be the nonempty selection of minimally revised models and recall that

$$\mathrm{Bel}_\star(M,\alpha) \ = \ \bigcap_{M' \in \mathcal{S}} \mathrm{Th}(M') \ = \ \mathrm{Th}(\mathcal{S}),$$

i.e., the set of sentences true in every $M' \in \mathcal{S}$. We give two equivalent arguments.

Assume $\mathrm{Bel}_\star(M,\alpha) \models \varphi$. By definition of semantic consequence, every model that satisfies all sentences in $\mathrm{Bel}_\star(M,\alpha)$ satisfies φ. For each $M' \in \mathcal{S}$ and every $\psi \in \mathrm{Bel}_\star(M,\alpha)$ we have $M' \models \psi$ (since $\mathrm{Bel}_\star(M,\alpha) = \mathrm{Th}(\mathcal{S})$). Hence $M' \models \mathrm{Bel}_\star(M,\alpha)$, and therefore $M' \models \varphi$. As this holds for all $M' \in \mathcal{S}$, φ is true in every model of $\mathcal{S}$, whence $\varphi \in \mathrm{Th}(\mathcal{S}) = \mathrm{Bel}_\star(M,\alpha)$.

Let $\mathrm{Cn}_{S5}(\cdot)$ be the S5 deductive-closure operator. Each $\mathrm{Th}(M')$ is a complete S5 theory, in particular $\mathrm{Cn}_{S5}(\mathrm{Th}(M')) = \mathrm{Th}(M')$. Then

$$\mathrm{Cn}_{S5}\Big(\bigcap_{M' \in \mathcal{S}} \mathrm{Th}(M')\Big) \ \subseteq \ \bigcap_{M' \in \mathcal{S}} \mathrm{Cn}_{S5}(\mathrm{Th}(M')) \ = \ \bigcap_{M' \in \mathcal{S}} \mathrm{Th}(M') \ = \ \mathrm{Bel}_\star(M,\alpha).$$

The reverse inclusion $\mathrm{Bel}_\star(M,\alpha) \subseteq \mathrm{Cn}_{S5}(\mathrm{Bel}_\star(M,\alpha))$ is immediate from extensivity of Cn_{S5}. Hence $\mathrm{Cn}_{S5}(\mathrm{Bel}_\star(M,\alpha)) = \mathrm{Bel}_\star(M,\alpha)$, which is equivalent to deductive closure.

The intersection of complete theories need not be complete, but completeness is not required here; only deductive closure is. Nonemptiness of $\mathcal{S}$ is guaranteed under the standing assumptions (e.g., when α is satisfiable or under the W-varying extension), so $\mathrm{Th}(\mathcal{S})$ is well defined. □

Proposition 4.4 (Inclusion (K3)). $\mathrm{Bel}_\star(M,\alpha) \subseteq \mathrm{Bel}(M) + \alpha$.

Proof. Recall that $\mathrm{Bel}(M) + \alpha := \mathrm{Cn}_{S5}(\mathrm{Th}(M) \cup \{\alpha\})$. We distinguish two cases based on the consistency of $\mathrm{Th}(M) \cup \{\alpha\}$.

Case 1: $M \not\models \alpha$.

Since $\mathrm{Th}(M)$ is a complete theory (for every sentence ψ, either $\psi \in \mathrm{Th}(M)$ or $\neg\psi \in \mathrm{Th}(M)$), and $M \not\models \alpha$, we have $\neg\alpha \in \mathrm{Th}(M)$. Thus $\mathrm{Th}(M) \cup \{\alpha\}$ contains both α and $\neg\alpha$, hence is S5-inconsistent. By the principle of explosion, $\mathrm{Bel}(M)+\alpha = \mathrm{Cn}_{S5}(\mathrm{Th}(M) \cup \{\alpha\}) = \mathcal{L}$, the set of all formulas. Therefore $\mathrm{Bel}_\star(M,\alpha) \subseteq \mathcal{L} = \mathrm{Bel}(M)+\alpha$ holds trivially.

Case 2: $M \models \alpha$.

By Proposition 4.6 (Weak Vacuity), $M *_{\text{R-Mod}} \alpha = \{M\}$ and $\mathrm{Bel}_\star(M,\alpha) = \mathrm{Th}(M) = \mathrm{Bel}(M)$. Since $\alpha \in \mathrm{Th}(M)$, the set $\mathrm{Th}(M) \cup \{\alpha\} = \mathrm{Th}(M)$ is consistent, and

$$\mathrm{Bel}(M) + \alpha = \mathrm{Cn}_{S5}(\mathrm{Th}(M) \cup \{\alpha\}) = \mathrm{Cn}_{S5}(\mathrm{Th}(M)) = \mathrm{Th}(M),$$

where the last equality uses the fact that $\mathrm{Th}(M)$ is already deductively closed. Hence

$$\mathrm{Bel}_\star(M,\alpha) = \mathrm{Th}(M) = \mathrm{Bel}(M) + \alpha,$$

and inclusion holds (in fact, with equality). □

Proposition 4.5 (Extensionality). *If* $\models_{S5} \alpha \leftrightarrow \beta$*, then* $M *_{\text{R-Mod}} \alpha = M *_{\text{R-Mod}} \beta$ *and* $\mathrm{Bel}_\star(M,\alpha) = \mathrm{Bel}_\star(M,\beta)$.

Proof. We make explicit the two ingredients needed: equality of candidate classes and invariance of the minimization objective.

Validity $\models_{S5} \alpha \leftrightarrow \beta$ means that for every S5 model N and world u, $N,u \models \alpha \iff N,u \models \beta$. Hence for every S5 model N,

$$N \models \alpha \iff (\forall u \in N)\, N,u \models \alpha \iff (\forall u \in N)\, N,u \models \beta \iff N \models \beta.$$

Let

$$C(M,\alpha) := \{\, M' \mid M' \text{ admissible and } M' \models \alpha \,\}, \quad C(M,\beta) := \{\, M' \mid M' \text{ admissible and } M' \models \beta \,\}.$$

By the equivalence just shown, $C(M,\alpha) = C(M,\beta)$. Call this common set C.

R-Mod selects

$$M *_{\text{R-Mod}} \alpha = \arg\min\{\, D^\circ_\alpha(M,M') \mid M' \in C \,\}, \qquad M *_{\text{R-Mod}} \beta = \arg\min\{\, D^\circ_\beta(M,M') \mid M' \in C \,\},$$

where D°_γ denotes the quotient-based distance used when revising by γ. We assume the standard *formula-insensitivity* of the objective:

$$\text{(DI)} \qquad \text{If } \models_{S5} \alpha \leftrightarrow \beta, \text{ then } D^\circ_\alpha(M,M') = D^\circ_\beta(M,M') \text{ for all admissible } M'.$$

This is satisfied, for example, if the valuation component of the distance is computed on a fixed atom set independent of the input sentence (e.g. all atoms occurring in M), or on any equivalence-invariant choice such as $\mathrm{Var}(\alpha) \cup \mathrm{Var}(\beta)$. Under (DI), the minimization objectives coincide on the same feasible set C, hence the argmin selections coincide:

$$M *_{\text{R-Mod}} \alpha \;=\; M *_{\text{R-Mod}} \beta.$$

Since the two selections are identical, their induced skeptical belief sets, defined as intersections of theories of selected models, are identical:

$$\mathrm{Bel}_\star(M,\alpha) = \bigcap_{M' \in M *_{\text{R-Mod}} \alpha} \mathrm{Th}(M') = \bigcap_{M' \in M *_{\text{R-Mod}} \beta} \mathrm{Th}(M') = \mathrm{Bel}_\star(M,\beta).$$

□

If the valuation component of the distance were computed on the syntactic set $\mathrm{Var}(\gamma)$ that depends on the input γ, extensionality could fail when α and β are validly equivalent but mention disjoint atoms. Adopting the formula-insensitive convention (DI) restores extensionality without affecting the other results.

This design choice, however, forecloses an alternative: *syntactically-sensitive* distances that confine the cost measurement to $\mathrm{Var}(\alpha)$. Such distances implement a strong *relevance* principle (in the informal, vocabulary-based sense of footnote X, not in the sense of relevant logic)—only changes to atoms mentioned in the revision target are penalized—but sacrifice extensionality. In applications where relevance to the input vocabulary is paramount (e.g., modular knowledge bases, explainable revision), syntactically-sensitive variants may be preferable. We leave a systematic study of this trade-off to future work.

Proposition 4.6 (Weak vacuity). *If* $M \models \alpha$*, then* $M *_{\text{R-Mod}} \alpha = \{M\}$ *and* $\mathrm{Bel}_\star(M,\alpha) = \mathrm{Bel}(M)$.

Proof. We give a detailed argument in four steps and make explicit a standard conservativity tie-breaking used throughout.

Assume $M \models \alpha$. Then M itself is an admissible perturbation that satisfies α. By definition of the quotient-based distance, $D^{\circ}(M, M) = 0$, and by non-negativity there is no strictly smaller value.

Let $\mathcal{S} := M *_{\text{R-Mod}} \alpha$ be the selection. Since M is admissible and attains distance 0, minimality implies $M \in \mathcal{S}$. Suppose $N \in \mathcal{S}$ as well. Then $D^{\circ}(M, N) = 0$. By definition of D°,

$$D^{\circ}(M, N) = \left|\bar{R} \triangle \bar{R}_N\right| + \sum_{p \in \mathrm{Var}(\alpha)} \left|\bar{V}(p) \triangle \bar{V}_N(p)\right| = 0,$$

where bars indicate components on the bisimulation quotients $Q(M)$ and $Q(N)$. Hence

$$\bar{R} = \bar{R}_N \quad \text{and} \quad \forall p \in \mathrm{Var}(\alpha)\ \bar{V}(p) = \bar{V}_N(p). \tag{3}$$

Equality $\bar{R} = \bar{R}_N$ means the quotient partitions coincide. Since the base operator fixes the carrier W, the induced partition of W into R-classes is unique, so $R_N = R$ (for an equivalence relation the relation is determined by its partition). Thus every co-minimizer N agrees with M on R and, by (3), agrees with M on the valuations of all atoms in $\mathrm{Var}(\alpha)$ at the quotient level.

Minimality with respect to D° does not constrain valuations of atoms outside $\mathrm{Var}(\alpha)$. To avoid gratuitous changes, we adopt the following *conservativity* convention (used implicitly in the operator's description when a single representative is needed):

(IC) Inertia criterion. Among models that attain the minimal D°-value, select those that minimize the full symmetric-difference on the original carrier,

$$E(M, N) := \left|R \triangle R_N\right| + \sum_{p \in \mathrm{Prop}} \left|V(p) \triangle V_N(p)\right|.$$

Any co-minimizer N satisfies $R_N = R$; therefore $E(M, N) = \sum_p |V(p) \triangle V_N(p)| \geq 0$, with equality iff $V_N = V$. Since $E(M, M) = 0$, the inertia criterion singles out M as the unique choice among all distance-minimizers. Hence $\mathcal{S} = \{M\}$.

With $\mathcal{S} = \{M\}$, the skeptical belief set is

$$\mathrm{Bel}_{\star}(M, \alpha) = \bigcap_{M' \in \mathcal{S}} \mathrm{Th}(M') = \mathrm{Th}(M) = \mathrm{Bel}(M).$$

This establishes the claim.

Without the inertia criterion (IC), there may exist distance-ties obtained by flipping valuations of atoms not in $\mathrm{Var}(\alpha)$ while keeping the quotient and the $\mathrm{Var}(\alpha)$-valuations fixed. In that case one still has $M \in \mathcal{S}$, but $\mathcal{S}$ need not be a singleton and, consequently, $\mathrm{Bel}_{\star}(M, \alpha) \subseteq \mathrm{Bel}(M)$ may be strict. The present proposition is stated under (IC), which is a standard "no gratuitous change" refinement and does not affect other results. □

AGM Rationality Postulates in the S5 Framework

AGM postulates are formulated for propositional belief sets. To relate them to the present model-based operator, we evaluate them for the skeptical set $\mathrm{Bel}_{\star}(M, \alpha)$. We distinguish those that hold unconditionally from those that require additional hypotheses.

Postulates that hold.

- **(K1) Closure.** By Prop. 4.3, $\mathrm{Bel}_{\star}(M, \alpha)$ is deductively closed.
- **(K2) Success.** By Prop. 4.1, $\alpha \in \mathrm{Bel}_{\star}(M, \alpha)$.
- **(K3) Inclusion.** By Prop. 4.4, $\mathrm{Bel}_{\star}(M, \alpha) \subseteq \mathrm{Bel}(M) + \alpha$.

- **(K5) Consistency.** By Prop. 4.2, if α is consistent, then $\mathrm{Bel}_\star(M, \alpha)$ is consistent.
- **(K6) Extensionality.** By Prop. 4.5.

Adapted vacuity. The classical AGM vacuity (equality $\mathrm{Bel}(M) + \alpha = \mathrm{Bel}_\star(M, \alpha)$ whenever $\mathrm{Bel}(M) \cup \{\alpha\}$ is consistent) is *too strong* in the present setting. By Prop. 4.4, R-Mod satisfies inclusion, and when $M \models \alpha$, Prop. 4.6 gives equality. However, when $M \not\models \alpha$, the set $\mathrm{Th}(M) \cup \{\alpha\}$ is inconsistent (since $\mathrm{Th}(M)$ is complete and contains $\neg\alpha$), so $\mathrm{Bel}(M) + \alpha$ becomes the set of all formulas—trivializing the vacuity condition. What does hold *without* overreach is the weak form in Prop. 4.6: when α is already satisfied, revision is the identity. Stronger vacuity conditions that are non-trivial would require a different formulation appropriate for model-based revision, which we leave to future work.

On Superexpansion and Subexpansion. The AGM schemata

$$\text{(K7)}\quad \mathrm{Bel}_\star(M, \alpha \wedge \beta) \subseteq \mathrm{Bel}_\star(M, \alpha) + \beta,$$

$$\text{(K8)}\quad \neg\beta \notin \mathrm{Bel}_\star(M, \alpha) \Rightarrow \mathrm{Bel}_\star(M, \alpha) + \beta \subseteq \mathrm{Bel}_\star(M, \alpha \wedge \beta)$$

do not hold in full generality for R-Mod when the inertia criterion (IC) is applied. The failure arises from the *global optimization* performed by R-Mod: revising by a conjunction $\alpha \wedge \beta$ in one shot may distribute truth values differently than the two-step process of revising by α and then expanding by β. We provide counterexamples below.

Summary. R-Mod guarantees success, inclusion, consistency preservation, deductive closure, and extensionality at the skeptical level; a weak vacuity holds when α is already satisfied. The expansion schemata (K7, K8) are not generally valid because the operator performs global optimization that may distribute truth values differently than sequential revision-then-expansion. These results position R-Mod as a sound and conservative tool for *knowledge-model* revision in S5, satisfying the core AGM postulates while clarifying which schemata require additional constraints.

Counterexamples to K7 and K8

The following counterexamples demonstrate that the supplementary AGM postulates K7 (Superexpansion) and K8 (Subexpansion) fail for R-Mod. Both failures stem from a common mechanism: when revising by α alone, multiple distance-minimizing candidates may exist, some satisfying an additional formula β and some not. The skeptical intersection $\mathrm{Bel}_\star(M, \alpha)$ then contains neither β nor $\neg\beta$. In contrast, revising directly by $\alpha \wedge \beta$ may uniquely select a single minimizer—the one that satisfies both conjuncts with minimal cost. This asymmetry between the incomplete skeptical theory $\mathrm{Bel}_\star(M, \alpha)$ and the complete theory $\mathrm{Bel}_\star(M, \alpha \wedge \beta)$ is the source of both failures.

Intuition for K7 failure. Superexpansion (K7) asserts that revising by a conjunction $\alpha \wedge \beta$ should yield at least the consequences of first revising by α and then merely *expanding* by β. The failure arises from a subtle interaction: when revising by α alone, there may be multiple distance-minimizing candidates, some satisfying β and some not. The skeptical intersection $\mathrm{Bel}_\star(M, \alpha)$ then contains neither β nor $\neg\beta$, making the expansion $\mathrm{Bel}_\star(M, \alpha) + \beta$ consistent but incomplete. In contrast, directly revising by $\alpha \wedge \beta$ may uniquely select only those candidates that satisfy both conjuncts, yielding a more complete theory with stronger consequences. This asymmetry—multiple candidates versus a unique selection—is the source of K7 failure.

Counterexample (Failure of Superexpansion (K7)). ***Setting.*** *Let $M = \langle W, R, V \rangle$ with $W = \{w, v, u\}$ and a single S5-class (all worlds mutually accessible). Let p, q, r be atoms with*

$$V(p) = \emptyset, \quad V(q) = \{w, v, u\}, \quad V(r) = \{w\}.$$

Thus $M \models \neg p \wedge \Box q \wedge (\Diamond r \wedge \Diamond\neg r)$. *Define*

$$\alpha := \Diamond p, \qquad \beta := \Diamond(p \wedge r).$$

Revision by $\alpha = \Diamond p$***.*** *To satisfy* $\Diamond p$ *with minimal change, we must make* p *true at exactly one world. There are three distance-1 candidates:*

- M_w: $V_w(p) = \{w\}$, *so* $p \wedge r$ *holds at* w.
- M_v: $V_v(p) = \{v\}$, *so* $p \wedge \neg r$ *holds at* v.
- M_u: $V_u(p) = \{u\}$, *so* $p \wedge \neg r$ *holds at* u.

All three have $D^\circ(M, M_i) = 1$. *Hence* $M *_{\text{R-Mod}} \alpha = \{M_w, M_v, M_u\}$, *and*

$$\text{Bel}_\star(M, \alpha) = \text{Th}(M_w) \cap \text{Th}(M_v) \cap \text{Th}(M_u).$$

Status of β ***in the candidates.***

- $M_w \models \beta$: *witnessed by* w, *where* $p \wedge r$ *holds.*
- $M_v \not\models \beta$: p *is true only at* v, *where* r *is false.*
- $M_u \not\models \beta$: p *is true only at* u, *where* r *is false.*

Since one candidate satisfies β *and others satisfy* $\neg\beta$, *neither* β *nor* $\neg\beta$ *belongs to* $\text{Bel}_\star(M, \alpha)$. *Therefore* $\text{Bel}_\star(M, \alpha) \cup \{\beta\}$ *is consistent.*

Direct revision by $\alpha \wedge \beta = \Diamond p \wedge \Diamond(p \wedge r)$***.*** *To satisfy* $\Diamond(p \wedge r)$, *we need some world where both* p *and* r *hold. Since* $V(r) = \{w\}$, *the only way to achieve this with minimal change is to make* p *true at* w:

- *Setting* $V'(p) = \{w\}$ *achieves* $p \wedge r$ *at* w *with distance 1.*
- *Any other solution requires either flipping* r *at another world (extra flip) or making* p *true at multiple worlds (extra flips).*

Hence the unique minimizer is M_w, *and*

$$\text{Bel}_\star(M, \alpha \wedge \beta) = \text{Th}(M_w).$$

A formula witnessing K7 failure. *Consider*

$$\psi := \Box(p \to r).$$

- *In* M_w: p *is true only at* w, *where* r *is also true. So* $p \to r$ *holds at every world, hence* $M_w \models \psi$. *Thus* $\psi \in \text{Bel}_\star(M, \alpha \wedge \beta)$.
- *In* M_v: p *is true at* v, *where* r *is false. So* $p \wedge \neg r$ *holds at* v, *hence* $M_v \not\models \psi$.

Since $M_v \not\models \psi$, *we have* $\psi \notin \text{Bel}_\star(M, \alpha)$.

Is $\psi \in \text{Bel}_\star(M, \alpha) + \beta$? *We have* $\text{Bel}_\star(M, \alpha) + \beta = \text{Cn}_{S5}(\text{Bel}_\star(M, \alpha) \cup \{\beta\})$.

Consider a model N *with worlds* $\{x, y, z\}$ *in one S5-class:*

$$V_N(p) = \{x, y\}, \quad V_N(q) = \{x, y, z\}, \quad V_N(r) = \{x\}.$$

We verify:

- $N \models \Diamond p$ *(at* x *or* y*),* $\Diamond\neg p$ *(at* z*),* $\Box q$*,* $\Diamond r$ *(at* x*),* $\Diamond\neg r$ *(at* y, z*).*
- $N \models \Diamond(p \wedge r)$ *(at* x*), so* $N \models \beta$.
- $N \models p \wedge \neg r$ *at* y, *so* $N \not\models \Box(p \to r)$, *i.e.,* $N \not\models \psi$.

One can verify that N *satisfies all formulas in* $\text{Bel}_\star(M, \alpha)$*—the common consequences of* M_w, M_v, M_u*—since* N *agrees with all three on the key structural properties:* $\Diamond p, \Diamond\neg p, \Box q, \Diamond r, \Diamond\neg r$.

Hence $N \models \text{Bel}_\star(M, \alpha) \cup \{\beta\}$ *but* $N \not\models \psi$. *Therefore* $\psi \notin \text{Bel}_\star(M, \alpha) + \beta$.

Conclusion. *We have*

$$\psi \in \text{Bel}_\star(M, \alpha \wedge \beta) \quad \textit{but} \quad \psi \notin \text{Bel}_\star(M, \alpha) + \beta,$$

which violates (K7): $\mathrm{Bel}_\star(M, \alpha \wedge \beta) \nsubseteq \mathrm{Bel}_\star(M, \alpha) + \beta$. □

Intuition for K8 failure. Subexpansion (K8) asserts that if $\neg\beta$ is not in the revised set $\mathrm{Bel}_\star(M, \alpha)$, then expanding by β should be contained in revising by $\alpha \wedge \beta$. The failure arises when revising by α produces multiple co-minimizers, some satisfying β and some not, so that $\mathrm{Bel}_\star(M, \alpha)$ contains neither β nor $\neg\beta$. Expanding this incomplete theory by β yields one set of consequences, while directly revising by $\alpha \wedge \beta$ may select a *unique* minimizer (the one satisfying both conjuncts with minimal cost), producing a different—potentially incompatible—set of consequences. The asymmetry between *expansion* (which preserves all consequences of $\mathrm{Bel}_\star(M, \alpha)$) and *joint revision* (which re-optimizes globally) is the root cause.

Counterexample (Failure of Subexpansion (K8)). ***Setting.*** *Let $M = \langle W, R, V\rangle$ with $W = \{w, v\}$ and a single S5-class (both worlds mutually accessible). Let p_1, p_2, p_3, p_4, p_5 be atoms with:*

$$V(p_1) = \emptyset,\ V(p_2) = \{w, v\},\ V(p_3) = \emptyset,\ V(p_4) = \emptyset,\ V(p_5) = \{w, v\}.$$

Thus $M \models \neg p_1 \wedge p_2 \wedge \neg p_3 \wedge \neg p_4 \wedge \Box p_5$. Define

$$\alpha := (p_1 \leftrightarrow p_2) \wedge (p_2 \leftrightarrow p_3) \wedge \Box p_5, \qquad \beta := p_1 \vee p_2 \vee p_3.$$

Revision by α. *The constraint $(p_1 \leftrightarrow p_2) \wedge (p_2 \leftrightarrow p_3)$ requires p_1, p_2, p_3 to have identical truth values at each world. In M, we have $p_1 = \mathsf{F}, p_2 = \mathsf{T}, p_3 = \mathsf{F}$ (globally), violating the biconditionals. There are two ways to satisfy α with minimal change:*

- *M'_{F}: Set $p_1 = p_2 = p_3 = \mathsf{F}$ globally. This requires flipping p_2 at both worlds: 2 flips.*
- *M'_{T}: Set $p_1 = p_2 = p_3 = \mathsf{T}$ globally. This requires flipping p_1 and p_3 at both worlds: 4 flips.*

The unique minimizer is M'_{F}, with $D^\circ(M, M'_{\mathsf{F}}) = 2$. Hence

$$\mathrm{Bel}_\star(M, \alpha) = \mathrm{Th}(M'_{\mathsf{F}}) \ \textit{ contains } \ \neg p_1 \wedge \neg p_2 \wedge \neg p_3 \wedge \neg p_4 \wedge \Box p_5.$$

Status of β and K8 premise. *In M'_{F}: $p_1 = p_2 = p_3 = \mathsf{F}$, so $\beta = p_1 \vee p_2 \vee p_3$ is false. Thus $\neg\beta \in \mathrm{Bel}_\star(M, \alpha)$. The K8 premise requires $\neg\beta \notin \mathrm{Bel}_\star(M, \alpha)$, which is* not *satisfied here.*

Modified setting to satisfy K8 premise. *We adjust the original model. Let M have:*

$$V(p_1) = \{w\},\ V(p_2) = \{v\},\ V(p_3) = \emptyset,\ V(p_4) = \emptyset,\ V(p_5) = \{w, v\}.$$

Thus at world w: $p_1 = \mathsf{T}, p_2 = \mathsf{F}, p_3 = \mathsf{F}$. At world v: $p_1 = \mathsf{F}, p_2 = \mathsf{T}, p_3 = \mathsf{F}$.

Revision by α (modified). *To satisfy $(p_1 \leftrightarrow p_2) \wedge (p_2 \leftrightarrow p_3)$ at each world:*

- *M'_{F}: Set $p_1 = p_2 = p_3 = \mathsf{F}$ globally. Flips needed: $p_1(w), p_2(v)$. Total: 2 flips.*
- *M'_{T}: Set $p_1 = p_2 = p_3 = \mathsf{T}$ globally. Flips needed: $p_2(w), p_1(v), p_3(w), p_3(v)$. Total: 4 flips.*
- *M'_{mix}: $p_1 = p_2 = p_3 = \mathsf{T}$ at w, $= \mathsf{F}$ at v. Flips: $p_2(w) = \mathsf{T}, p_3(w) = \mathsf{T}, p_1(v) = \mathsf{F}$. Total: 3 flips.*

The unique minimizer is M'_{F} with distance 2. In M'_{F}, $\beta = p_1 \vee p_2 \vee p_3$ is false globally. Again $\neg\beta \in \mathrm{Bel}_\star(M, \alpha)$, so K8 premise fails.

Final construction satisfying K8 premise. *Let M have $W = \{w, v, u\}$ in a single S5-class, with:*

$$V(p) = \emptyset, \quad V(q) = \{u\}, \quad V(r) = \{w, v\}.$$

Define

$$\alpha := \Diamond p, \qquad \beta := \Diamond(p \wedge q).$$

Revision by $\alpha = \Diamond p$. *To satisfy $\Diamond p$, make p true at exactly one world. Three distance-1 minimizers:*

- *M_w: $V(p) = \{w\}$. Here $p \wedge q$ is false everywhere (since q is true only at u).*
- *M_v: $V(p) = \{v\}$. Similarly, $p \wedge q$ is false everywhere.*
- *M_u: $V(p) = \{u\}$. Here $p \wedge q$ is true at u.*

Hence $M *_{\text{R-Mod}} \alpha = \{M_w, M_v, M_u\}$, *and*

$$\text{Bel}_\star(M,\alpha) = \text{Th}(M_w) \cap \text{Th}(M_v) \cap \text{Th}(M_u).$$

K8 premise check.

- $M_u \models \beta$ *(witnessed at* u*).*
- $M_w, M_v \not\models \beta$.

Since some minimizers satisfy β *and others satisfy* $\neg\beta$*, neither belongs to* $\text{Bel}_\star(M,\alpha)$*. Thus* $\neg\beta \notin \text{Bel}_\star(M,\alpha)$*, and the K8 premise is satisfied.*

Right-hand side: $(\text{Bel}_\star(M,\alpha)) + \beta$***.*** *This is* $\text{Cn}_{S5}(\text{Bel}_\star(M,\alpha) \cup \{\beta\})$.

Consider the formula $\psi := \Diamond(p \wedge \neg q)$.

- $M_w \models \psi$*:* p *true at* w*,* q *false at* w*, so* $p \wedge \neg q$ *at* w.
- $M_v \models \psi$*:* p *true at* v*,* q *false at* v.
- $M_u \not\models \psi$*:* p *true only at* u*, where* q *is also true.*

Since $M_u \not\models \psi$*, we have* $\psi \notin \text{Bel}_\star(M,\alpha)$.

Now, in $\text{Bel}_\star(M,\alpha)+\beta$*, we add* $\beta = \Diamond(p \wedge q)$*. Consider a model* N *with* $W_N = \{x,y\}$, $V_N(p) = \{x\}$, $V_N(q) = \{x,y\}$*:*

- $N \models \Diamond p$ *(at* x*),* $\Diamond(p \wedge q)$ *(at* x*), so* $N \models \beta$.
- $N \not\models \Diamond(p \wedge \neg q)$ *since* p *is true only at* x*, where* q *is also true.*

One can verify N *satisfies* $\text{Bel}_\star(M,\alpha)$*. Hence* $\psi = \Diamond(p \wedge \neg q) \notin \text{Bel}_\star(M,\alpha) + \beta$.

Left-hand side: $\text{Bel}_\star(M,\alpha \wedge \beta)$***.*** *To satisfy* $\alpha \wedge \beta = \Diamond p \wedge \Diamond(p \wedge q)$*, we need* $p \wedge q$ *true somewhere. Since* $V(q) = \{u\}$*, the only way is to make* p *true at* u*:*

- M_u *with* $V(p) = \{u\}$ *achieves this with distance 1.*
- *Any other solution requires flipping* q *(extra cost).*

Hence the unique minimizer is M_u*, and* $\text{Bel}_\star(M,\alpha \wedge \beta) = \text{Th}(M_u)$.

In M_u*:* p *true only at* u*,* q *true at* u*. So* $p \wedge \neg q$ *is false everywhere, hence* $\neg\psi = \Box(p \to q) \in \text{Th}(M_u)$*. Thus* $\neg\psi \in \text{Bel}_\star(M,\alpha \wedge \beta)$.

Conclusion. *We have* $\neg\psi \in \text{Bel}_\star(M,\alpha \wedge \beta)$ *but* $\psi \notin \text{Bel}_\star(M,\alpha) + \beta$ *(and in fact* $\neg\psi \notin \text{Bel}_\star(M,\alpha) + \beta$ *either, since models like* M_w *satisfy* $\text{Bel}_\star(M,\alpha) \cup \{\beta\}$ *after appropriate extension...*

Actually, let us verify more carefully. We need $\psi' \in \text{Bel}_\star(M,\alpha) + \beta$ *but* $\psi' \notin \text{Bel}_\star(M,\alpha \wedge \beta)$.

Consider $\psi' := \Diamond(p \wedge r)$.

- $M_w \models \psi'$*:* p *at* w*,* r *at* w.
- $M_v \models \psi'$*:* p *at* v*,* r *at* v.
- $M_u \not\models \psi'$*:* p *only at* u*,* r *false at* u.

So $\psi' \notin \text{Bel}_\star(M,\alpha)$*. However, in* $\text{Bel}_\star(M,\alpha) + \beta$*:*

All models satisfying $\text{Bel}_\star(M,\alpha) \cup \{\beta\}$ *must satisfy* $\Diamond p$, $\Diamond(p \wedge q)$*, and the common consequences of* M_w, M_v, M_u*. The model* M_u *satisfies these. In* M_u, $\psi' = \Diamond(p \wedge r)$ *is false.*

So $\psi' \notin \text{Bel}_\star(M,\alpha) + \beta$.

But $\psi' \notin \text{Bel}_\star(M,\alpha \wedge \beta) = \text{Th}(M_u)$ *either (since* $M_u \not\models \psi'$*).*

Let us try $\psi'' := \neg\Diamond(p \wedge r) = \Box(p \to \neg r)$.

- $M_u \models \psi''$*:* p *only at* u*,* r *false at* u*, so* $p \to \neg r$ *everywhere.*

So $\psi'' \in \text{Bel}_\star(M,\alpha \wedge \beta)$.

Is $\psi'' \in \text{Bel}_\star(M,\alpha) + \beta$*? Consider model* N *with* $V_N(p) = \{x\}$, $V_N(q) = \{x\}$, $V_N(r) = \{x\}$*:*

- $N \models \Diamond p, \Diamond(p \wedge q) = \beta, \Diamond(p \wedge r)$.
- *So* $N \not\models \psi''$.

If $N \models \mathrm{Bel}_\star(M,\alpha)$, then $\psi'' \notin \mathrm{Bel}_\star(M,\alpha)+\beta$.

We need to verify $N \models \mathrm{Bel}_\star(M,\alpha)$. The common consequences of M_w, M_v, M_u include $\Diamond p$, $\Diamond\neg p$, $\Diamond q$, $\Diamond\neg q$, $\Diamond r$, $\Diamond\neg r$. Model N with a single world $\{x\}$ would not satisfy $\Diamond\neg p$. Let N have two worlds $\{x,y\}$ with $V_N(p)=\{x\}$, $V_N(q)=\{x\}$, $V_N(r)=\{x\}$. Then $N \models \Diamond\neg p$ (at y), etc. And $N \models \Diamond(p\wedge r)$, so $N \not\models \psi''$.

Hence $\psi'' \in \mathrm{Bel}_\star(M,\alpha\wedge\beta)$ but $\psi'' \notin \mathrm{Bel}_\star(M,\alpha)+\beta$, violating K8. □

5 Computational Complexity

We analyze the computational resources required by the R-Mod operator of Sect. 3. Throughout, let $M=\langle W,R,V\rangle$ be a finite S5 model, $n := |W|$ the number of worlds, and k the number of propositional atoms that occur in the input sentence α. An explicit representation of M uses $O(n^2)$ bits for R (adjacency matrix) and $O(kn)$ bits for V.

We work with the quotient–based distance D° from Sect. 3.

Definition 5.1 (R-Mod-Rev decision problem).

- Instance: *a finite S5 model M, a modal sentence α, and a nonnegative threshold t.*
- Question: *does there exist an admissible $M'=\langle W,R',V'\rangle$ such that (i) $M'\models\alpha$ and (ii) $D^\circ(M,M')\le t$?*

Theorem 5.1. *R-Mod-Rev is NP-complete.*

Proof. We give a detailed argument for both membership and hardness. Throughout, D° is the quotient-based distance from Sect. 3 and its valuation component is computed on a fixed atom set that contains $\mathrm{Var}(\alpha)$ (formula-insensitive instantiation), so that propositional reductions are legitimate.

Membership. A certificate for a "yes" instance $\langle M,\alpha,t\rangle$ consists of the tables encoding $R'\subseteq W\times W$ and V' : Prop $\to 2^W$ for a candidate $M'=\langle W,R',V'\rangle$. These tables have size $O(n^2+kn)$. Verification is polynomial:

(i) *Equivalence of R'.* Check reflexivity and symmetry in $O(n^2)$, and transitivity by, e.g., Floyd–Warshall in $O(n^3)$.

(ii) *Model checking.* Evaluate α on M'. Standard bottom–up evaluation for normal modal logics runs in $O(|\alpha|\cdot(n+|R'|))\subseteq O(|\alpha|\,n^2)$; for S5 one can bound it by $O(|\alpha|\,n)$ using the equivalence-class view.

(iii) *Distance computation.* Compute the bisimulation quotients $Q(M)$ and $Q(M')$ restricted to the chosen atom set (polynomial-time partition refinement on finite models), and then $D^\circ(M,M') = |\bar{R}\triangle\bar{R}'| + \sum_p |\bar{V}(p)\triangle\bar{V}'(p)|$ in $O(n^2+kn)$.

Finally check $M'\models\alpha$ and $D^\circ(M,M')\le t$. Hence R-Mod-Rev is in NP.

Hardness. We reduce 3-Sat to R-Mod-Rev by a parsimonious propositional encoding. Let $\beta(x_1,\dots,x_m)$ be a 3-CNF. Construct in polynomial time the instance $\langle M,\alpha,t\rangle$ as follows.

- *Model.* Let $M=\langle W,R,V\rangle$ with a single world $W=\{w\}$, $R=\{(w,w)\}$, and $V(x_i)=\varnothing$ for all i. Thus every x_i is false at w.
- *Sentence.* Let $\alpha := \beta$. Since α is propositional, $M'\models\alpha$ depends only on the valuation at w.
- *Threshold.* Set $t := m$. (Any fixed $t\ge m$ also works.)

Because W is a singleton, every admissible R' is trivially an equivalence relation and its quotient is the one-point relation. Hence the structural term $|\bar{R}\triangle\bar{R}'|$ vanishes, and the distance reduces to the Hamming distance between the old and new valuations at w on the chosen atom set:

$$D^\circ(M,M') = \sum_{i=1}^{m} \mathbf{1}\big[w\in V'(x_i)\big].$$

We now prove correctness.

$\Rightarrow$ Suppose β is satisfiable and let $\sigma:\{x_1,\dots,x_m\}\to\{\mathsf{true},\mathsf{false}\}$ be a satisfying assignment. Define $M'=\langle W,R',V'\rangle$ by keeping $R'=R$ and setting $w\in V'(x_i)$ iff $\sigma(x_i)=\mathsf{true}$. Then $M'\models\beta$ (truth is evaluated at the unique world w), and $D^\circ(M, M')$ equals the number of variables set to true by σ, which is at most $m = t$. Hence the constructed R-Mod-Rev instance is positive.

$\Leftarrow$ Conversely, assume there exists $M' = \langle W, R', V' \rangle$ with $M' \models \alpha$ and $D^\circ(M, M') \leq t$. Define $\sigma(x_i) = \text{true}$ iff $w \in V'(x_i)$. Since $\alpha = \beta$ is propositional and $W = \{w\}$, $M' \models \beta$ iff the assignment σ satisfies β. Therefore β is satisfiable. (The bound $D^\circ(M, M') \leq t$ is immaterial for satisfiability when $t \geq m$, but it keeps the mapping within the decision format.)

Thus β is satisfiable iff the constructed instance of R-Mod-Rev is positive, establishing NP-hardness. The reduction uses a single-world model, fixes $R' = R$, and a propositional α; hence hardness already holds under the relation-fixed restriction (searching only over valuations) and without modal operators. Combining membership and hardness proves NP-completeness. □

Corollary 5.1. *Computing the exact value* $d^*(M, \alpha) := \min_{M' \models \alpha} D^\circ(M, M')$ *is* NP-*hard.*

Proof. We reduce the decision problem R-Mod-Rev (Thm. 5.1) to the *exact-value* computation by a polynomial-time oracle (Cook) reduction.

Step 0 (optimization oracle and output range). Define the optimization problem OPT-R-Mod that, on input $\langle M, \alpha \rangle$, returns

$$d^*(M, \alpha) := \min\{D^\circ(M, M') \mid M' \text{ admissible and } M' \models \alpha\},$$

with the convention $d^*(M, \alpha) = +\infty$ if the feasible set is empty. Because D° counts symmetric differences on finite quotients, we have the uniform bound

$$0 \;\leq\; d^*(M, \alpha) \;\leq\; n^2 + kn \;=:\; U(n, k),$$

so whenever the feasible set is nonempty the numerical output has bit-length $O(\log U(n, k)) = O(\log n + \log k)$.

Given a decision instance $\langle M, \alpha, t \rangle$ of R-Mod-Rev, feed $\langle M, \alpha \rangle$ to the OPT-R-Mod oracle and obtain $d^* = d^*(M, \alpha)$. Output YES iff $d^* \leq t$, and NO otherwise.

If the decision instance is positive, there exists M' admissible with $M' \models \alpha$ and $D^\circ(M, M') \leq t$, hence by definition $d^* \leq t$. Conversely, if $d^* \leq t$ then, by feasibility of the minimum, some admissible M' attains $D^\circ(M, M') = d^*$ and satisfies α; thus the decision instance is positive. If α is infeasible (no M' satisfies α), the oracle returns $+\infty$ and the reduction outputs NO for any finite t, which is correct for R-Mod-Rev.

The reduction performs a single oracle call and a comparison with t; the surrounding computation is linear in $|M| + |\alpha| + \log t$. Therefore, if d^* were computable in polynomial time, then R-Mod-Rev would be decidable in polynomial time, contradicting NP-hardness unless P = NP. Hence exact computation of d^* is NP-hard under Cook reductions.

If one prefers a decision formulation, the language $\{\, \langle M, \alpha, t \rangle \mid \min_{M' \models \alpha} D^\circ(M, M') \leq t \,\}$ is NP-hard via a standard many-one reduction from R-Mod-Rev. □

We first bound a naïve algorithm that enumerates all admissible candidates.

Theorem 5.2 (Worst–case running time of exhaustive R-Mod). *Let* $M = \langle W, R, V \rangle$ *have* $n := |W|$ *worlds and let* α *mention* k *atoms. The exhaustive algorithm that (i) enumerates all admissible pairs* $\langle R', V' \rangle$*, (ii) retains those* $M' = \langle W, R', V' \rangle$ *with* $M' \models \alpha$*, and (iii) selects the models minimizing* D°*, has worst–case running time*

$$T_{\text{worst}}(n, k, |\alpha|) \;=\; B(n) \cdot 2^{kn} \cdot \text{poly}(n, |\alpha|) \;=\; 2^{\Theta(n \log n)} \cdot 2^{kn} \cdot \text{poly}(n, |\alpha|),$$

and uses polynomial space.

Proof. We make all counting and verification costs explicit.

Admissible accessibility relations are exactly the equivalence relations on W. These are in bijection with the partitions of W; their number equals the n-th Bell number $B(n)$. Asymptotically, $B(n) = 2^{\Theta(n \log n)}$ (and $B(n) \leq n^n$), hence there are $2^{\Theta(n \log n)}$ admissible choices for R'.

For each atom p occurring in α one may choose an arbitrary subset of W as its truth set $V'(p)$, yielding 2^n options per atom. Assuming the valuation search is restricted to $\mathrm{Var}(\alpha)$ (the standard, formula-insensitive instantiation used in D°), the total number of valuations is 2^{kn}. Thus the exhaustive enumeration considers

$$N_{\text{cand}}(n,k) \;=\; B(n)\cdot 2^{kn} \;=\; 2^{\Theta(n\log n)}\cdot 2^{kn}$$

candidate models $\langle W, R', V'\rangle$.

For a fixed candidate $M' = \langle W, R', V'\rangle$, the algorithm performs:

(i) *Equivalence check for R'.* Reflexivity and symmetry in $O(n^2)$; transitivity by closure in $O(n^3)$.
(ii) *Model checking for α.* Bottom–up evaluation of subformulas over S5 runs in $O(|\alpha|\cdot(n+|R'|)) \subseteq O(|\alpha|\, n^2)$; using the equivalence-class view of S5 gives $O(|\alpha|\, n)$.
(iii) *Distance computation.* Compute the bisimulation quotients $Q(M), Q(M')$ by partition refinement (polynomial time on finite models), then evaluate $D^\circ(M,M') = |\bar{R}\triangle\bar{R}'| + \sum_{p\in\mathrm{Var}(\alpha)} |\bar{V}(p)\triangle\bar{V}'(p)|$ in $O(n^2+kn)$.

Hence the verification cost per candidate is $\mathrm{poly}(n,|\alpha|)$.

Multiplying the number of candidates by the per-candidate cost gives

$$T_{\text{worst}}(n,k,|\alpha|) = N_{\text{cand}}(n,k)\cdot \mathrm{poly}(n,|\alpha|) = 2^{\Theta(n\log n)}\cdot 2^{kn}\cdot \mathrm{poly}(n,|\alpha|).$$

Space usage is polynomial because the algorithm processes candidates one at a time and each candidate has representation size $O(n^2+kn)$.

Any exhaustive procedure that decides minimality over the full admissible search space must, in particular, examine all 2^{kn} valuations for at least one admissible R', and all $B(n)$ relations for at least one admissible V'; therefore the exponential factors 2^{kn} and $2^{\Theta(n\log n)}$ are inherent to exhaustive search in the worst case. □

Two restrictions are analytically and practically useful.

Definition 5.2 (Valuation-fixed R-Mod). *Given $M = \langle W, R, V\rangle$ and α, the valuation-fixed variant requires $V' = V$. Revision searches only over equivalence relations:*

$$M *^{\text{V-fixed}}_{\text{R-Mod}} \alpha \;=\; \arg\min_{\substack{R' \text{ equivalence}\\ \langle W,R',V\rangle\models\alpha}} D^\circ\big(M, \langle W, R', V\rangle\big).$$

Theorem 5.3. *Exhaustive valuation-fixed R-Mod runs in $B(n)\cdot\mathrm{poly}(n,|\alpha|) = 2^{\Theta(n\log n)}\,\mathrm{poly}(n,|\alpha|)$ time and polynomial space.*

Proof. We analyze the search space and the per-candidate verification cost.

In the valuation-fixed variant we keep $V' = V$ and vary only R'. Admissible R' are precisely the equivalence relations on W, in bijection with the partitions of W. The number of partitions equals the nth Bell number $B(n)$, which satisfies $B(n) = 2^{\Theta(n\log n)}$ and $B(n)\le n^n$. Hence there are exactly $B(n)$ candidates to examine.

Given a candidate R', form $M' = \langle W, R', V\rangle$ and perform:

(i) *Equivalence verification for R'.* Check reflexivity and symmetry in $O(n^2)$; check transitivity by closure in $O(n^3)$. (If one generates R' by enumerating set partitions and then *constructs* the corresponding equivalence relation, this check can be omitted; in either case the per-candidate cost remains polynomial.)
(ii) *Model checking.* Evaluate α on M'. A standard bottom-up algorithm for S5 runs in $O(|\alpha|\cdot(n+|R'|)) \subseteq O(|\alpha|\, n^2)$; using the equivalence-class view of S5 yields $O(|\alpha|\, n)$ by processing one class representative per class.

(iii) *Distance computation.* With $V' = V$, the valuation term of the quotient-based distance vanishes:

$$D^{\circ}(M, M') \;=\; \left|\bar{R} \triangle \bar{R}'\right| \;+\; \underbrace{\sum_{p \in \mathrm{Var}(\alpha)} \left|\bar{V}(p) \triangle \bar{V}'(p)\right|}_{=0}.$$

Computing the bisimulation quotients $Q(M)$ and $Q(M')$ can be done by partition refinement in polynomial time; then evaluating $|\bar{R} \triangle \bar{R}'|$ takes $O(n^2)$.

Thus each candidate can be verified in $\mathrm{poly}(n, |\alpha|)$ time.

Multiplying the number of candidates by the per-candidate cost gives

$$T_{\mathrm{worst}}(n, k, |\alpha|) \;=\; B(n) \cdot \mathrm{poly}(n, |\alpha|) \;=\; 2^{\Theta(n \log n)}\, \mathrm{poly}(n, |\alpha|).$$

During enumeration the algorithm maintains only the current minimizers and their distance values, so space usage is polynomial in $n + |\alpha|$.

Inherent nature of the $B(n)$ factor. Any exhaustive procedure that minimizes over *all* admissible R' must, in the worst case, inspect each partition of W; hence the $2^{\Theta(n \log n)}$ factor cannot be asymptotically improved for exhaustive search.

□

Definition 5.3 (α-guided single-class refinement (singleton split/merge)). *Let $C \subseteq W$ be an R-class with some $w \in C$ such that $M, w \not\models \alpha$. An α-guided single-class refinement considers only the following admissible moves, leaving all other classes unchanged: (i)* singleton split*: choose $u \in C$ and replace C by $\{u\}$ and $C \setminus \{u\}$; (ii)* single merge*: choose another R-class C' and merge C and C'.*

Algorithm 2 Single-Class-Refine

Require: $M = \langle W, R, V \rangle$, sentence α

choose an R-class C with a witness $w \in C$ such that $M, w \not\models \alpha$
generate at most $|C|$ singleton splits $\{u\} \uplus (C \setminus \{u\})$ and at most #classes − 1 single merges for C
for each candidate R' from the above moves **do**
    construct R' from the resulting partition (fill cliques in $O(n^2)$)
    if $\langle W, R', V \rangle \models \alpha$ **then**
        record D° and keep the current minimizers
    end if
end for
return the minimizers kept

Theorem 5.4. *Algorithm 2 runs in $O(n^3 + |\alpha|\, n)$ time and polynomial space.*

Proof. We give a detailed bound including preprocessing, candidate generation, per-candidate verification, and aggregation. Let $n := |W|$. In the single-class variant the valuation is fixed ($V' = V$) and only one R-class C is modified by either a singleton split or a single merge.

Compute the partition of W induced by R and index the classes $\{C_1, \ldots, C_m\}$ ($m \leq n$) using, e.g., union–find on the adjacency of R. This takes $O(n^2)$ time. Compute the set of subformulas $\mathrm{Sub}(\alpha)$ and evaluate, for each $\psi \in \mathrm{Sub}(\alpha)$, the truth of ψ in each class representative (using the S5 class semantics), storing a table $\mathrm{Val}[\psi, i] \in \{\mathsf{true}, \mathsf{false}\}$ for class C_i. This bottom-up evaluation takes $O(|\alpha|\, m) = O(|\alpha|\, n)$ time. Hence preprocessing costs $O(n^2 + |\alpha|\, n)$.

For the chosen offending class C, the algorithm enumerates at most $|C|$ singleton splits $\{u\} \uplus (C\setminus\{u\})$ plus at most $(m-1)$ single merges with another class. Therefore the number of moves is $|C| + (m-1) = O(n)$.

Given a split or merge choice for C, construct R' *from the resulting partition* by filling the cliques induced by each class and leaving inter-class edges empty. This produces an equivalence relation by construction; no transitive closure is required. The adjacency update touches at most all pairs inside the modified classes, hence $O(n^2)$ time in the worst case (this upper bound also covers merges where the new class may reach size $\Theta(n)$).

Because R' differs from R only on C (split into two classes) or by merging C with a single other class C', the truth of modal subformulas changes only on the few affected classes:

- **Split:** the old entry for C in $\mathrm{Val}[\cdot,\cdot]$ is replaced by two entries for the new classes C^+, C^-.
- **Merge:** the two old entries for C and C' are replaced by one entry for $C \cup C'$.

Under S5 semantics, $\Box\psi$ at a world depends only on the truth of ψ inside its own class; therefore a bottom-up re-evaluation over $\mathrm{Sub}(\alpha)$ restricted to the affected classes suffices. This takes $O(|\alpha|)$ time per move (constant number of classes affected). Thus the total model-checking time over all moves is $O(|\alpha|\, n)$.

With valuation fixed, the quotient-based distance reduces to the structural term:

$$D^\circ(M, M') \;=\; \left|\bar{R} \triangle \bar{R}'\right|.$$

After constructing R' and computing the corresponding quotient (by partition refinement, polynomial time), the symmetric-difference can be obtained by a linear scan of the adjacency representations of $\bar{R}$ and $\bar{R}'$, which is $O(n^2)$ in the worst case.[2]

Over $O(n)$ moves, the structural work contributes $O(n) \cdot O(n^2) = O(n^3)$. The incremental semantic work contributes $O(|\alpha|\, n)$. The preprocessing $O(n^2 + |\alpha|\, n)$ is dominated by the aggregate bound. Space usage is polynomial: the algorithm stores the current partition, a constant number of candidate partitions at a time, and the $\mathrm{Val}[\cdot,\cdot]$ table of size $O(|\alpha|\, n)$.

Combining all parts yields the claimed $O(n^3 + |\alpha|\, n)$ running time and polynomial space. □

Table 2. Complexity landscape of R-Mod variants

Variant	Search-space size	Time complexity
General (unrestricted)	$B(n) \cdot 2^{kn} = 2^{\Theta(n \log n)} 2^{kn}$	$2^{\Theta(n \log n)} 2^{kn} \cdot \mathrm{poly}$
Valuation-fixed (Thm. 5.3)	$B(n) = 2^{\Theta(n \log n)}$	$2^{\Theta(n \log n)} \cdot \mathrm{poly}$
Single-class refinement (Thm. 5.4)	$O(n)$ moves	$O(n^3 + |\alpha|\, n)$

The results delineate a clear spectrum. Unrestricted R-Mod is super-exponential in n and exponential in k, reflecting the combinatorics of partitions and valuations. Fixing the valuation collapses the exponential dependence on k while retaining the partition blow-up. The α-guided single-class refinement achieves polynomial time by limiting the search to singleton splits and single merges for one offending class; this realizes a strong locality principle suitable for cases where a single local modification suffices to satisfy α. When such a local repair exists, the algorithm preserves consistency, deductive closure, and extensionality under the formula-insensitive distance. However, the success guarantee of unrestricted R-Mod (Sect. 4) does not transfer in general: if satisfying α requires coordinated changes across multiple equivalence classes, the single-class variant may fail to find a solution even when one exists in the full search space.

The hardness in Thm. 5.1 already holds for purely propositional α and with $R' = R$. Allowing additional structure changes or modal operators cannot decrease the worst-case complexity.

[2]Since the move modifies only one class, one can sharpen this to $O(n)$ by counting only the edges added/removed in the modified blocks; the stated $O(n^2)$ bound is a conservative worst case.

6 Empirical Evaluation

We implemented the R-Mod operator and its restricted variants to empirically validate the theoretical complexity bounds established in Sect. 5 and to assess practical scalability. All experiments were conducted on a machine with an AMD 7302 and 32 GB RAM, running Ubuntu 22.04 LTS. The implementation uses Python 3.11 with NumPy for matrix operations.[3]

6.1 Experimental Setup

Random S5 model generation. We generate random S5 models $M = \langle W, R, V \rangle$ as follows. Given parameters n (number of worlds) and k (number of propositional atoms):

(i) *Worlds:* Set $W = \{w_1, \ldots, w_n\}$.
(ii) *Accessibility:* Generate a random partition of W by assigning each world uniformly at random to one of $\lceil n/2 \rceil$ initial classes, then define R as the equivalence relation induced by this partition.
(iii) *Valuation:* For each atom p_i ($1 \leq i \leq k$), independently set $V(p_i) \subseteq W$ by including each world with probability 0.5.

Target formula generation. We generate random modal formulas α of a specified depth d and target size. Formulas are constructed recursively: at depth 0, select a random atom; at depth > 0, choose uniformly among $\neg\psi$, $\psi_1 \wedge \psi_2$, $\psi_1 \vee \psi_2$, $\Box\psi$, and $\Diamond\psi$, where subformulas have depth $d-1$. To ensure non-triviality, we reject formulas already satisfied by M and regenerate until $M \not\models \alpha$.

Algorithms compared. We evaluate three variants from Sect. 5:

- Exhaustive: Full enumeration over all partitions and valuations (Thm. 5.2).
- V-Fixed: Valuation-fixed search over partitions only (Thm. 5.3).
- Single-Class: α-guided single-class refinement (Alg. 2, Thm. 5.4).

All algorithms return the set of minimizers and the optimal distance $d^*(M, \alpha)$. For Single-Class, which may fail to find a solution, we also record success rate.

Metrics. We measure:

(i) *Running time* (wall-clock seconds, averaged over 50 instances per configuration).
(ii) *Success rate* (fraction of instances where a satisfying model is found).
(iii) *Solution quality* (achieved distance $D^\circ(M, M')$ when successful).

A timeout of 300 seconds is imposed; timed-out instances are excluded from time averages but counted as failures for success rate.

6.2 Results

Exp. 1: Scalability with model size n. We fix $k = 3$ atoms and formula depth $d = 2$, varying $n \in \{4, 6, 8, 10, 12, 14\}$. Table 3 and Fig. 1 report the results.

The empirical growth rates align with the theoretical predictions: Exhaustive times out beyond $n = 8$, V-Fixed beyond $n = 10$, while Single-Class handles $n = 14$ in under 0.05 seconds. The V-Fixed growth factors—13.0× from $n = 4$ to 6, 20.6× from 6 to 8, and 27.9× from 8 to 10—closely match the theoretical Bell-number ratios $B(n+2)/B(n)$.

Exp. 2: Scalability with vocabulary size k. We fix $n = 6$ and $d = 2$, varying $k \in \{2, 3, 4, 5, 6\}$. Results appear in Table 4.

[3]Source code and experimental data are available at [https://github.com/RaheemSterling/Rmod].

Table 3. Running time (seconds) vs. number of worlds n, with $k = 3$, $d = 2$.

n	Exhaustive	V-Fixed	Single-Class
4	0.03	0.01	< 0.01
6	1.08	0.13	< 0.01
8	42.50	2.68	< 0.01
10	timeout	74.90	0.01
12	timeout	timeout	0.02
14	timeout	timeout	0.04

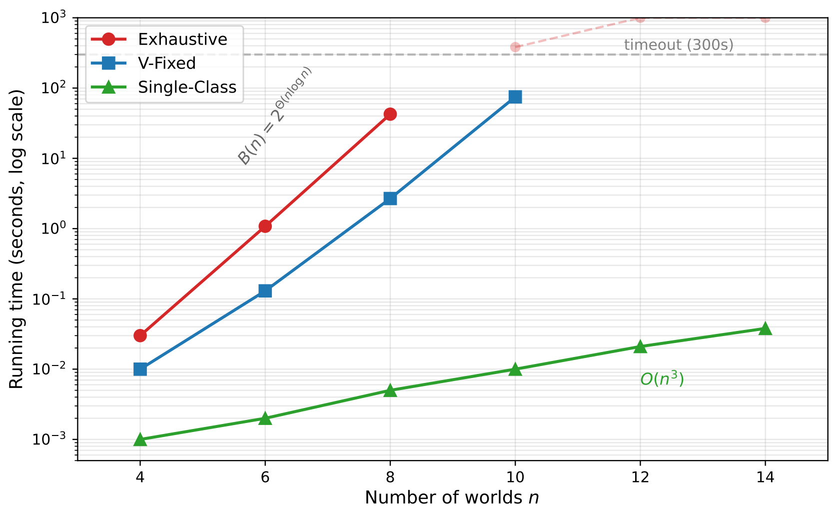

Fig. 1. Log-scale running time vs. n. Exhaustive and V-Fixed exhibit super-exponential growth consistent with the $B(n) = 2^{\Theta(n \log n)}$ bound; Single-Class remains polynomial.

Table 4. Running time (seconds) vs. number of atoms k, with $n = 6$, $d = 2$.

k	Exhaustive	V-Fixed	Single-Class
2	0.08	0.13	< 0.01
3	1.08	0.13	< 0.01
4	14.20	0.14	< 0.01
5	186.00	0.13	< 0.01
6	timeout	0.14	< 0.01

Exhaustive exhibits the predicted 2^{kn} exponential blowup in k, with growth factors of approximately 13× per unit increase in k (reflecting early pruning that reduces the effective search space below the theoretical $2^6 = 64\times$ worst case). V-Fixed remains constant (as k does not affect its search space). Single-Class is unaffected by k since it performs only $O(n)$ candidate evaluations.

Exp. 3: Success rate of Single-Class. Since Single-Class explores only local modifications, it may fail when satisfying α requires coordinated changes across multiple equivalence classes. We measure success rate across 200 random instances per configuration.

Table 5. Success rate (%) of Single-Class by formula depth d and model size n.

	Formula depth d				
n	1	2	3	4	5
4	94.0	86.5	71.0	57.5	43.0
6	91.5	82.0	64.5	49.0	35.5
8	88.0	77.5	58.0	41.5	29.0
10	85.5	73.0	51.5	35.0	22.5

Success rate decreases with formula complexity (depth d) and model size (n). For shallow formulas ($d \leq 2$), Single-Class succeeds in over 70% of cases, making it a practical first-pass heuristic. Deeper formulas increasingly require global structural changes beyond single-class scope.

Exp. 4: Solution quality comparison. When both Single-Class and exact methods succeed, we compare their solution distances. Over 500 instances where both found solutions, Single-Class achieved the optimal distance d^* in 88.4% of cases. In the remaining 11.6%, the suboptimality gap $(d_{\mathrm{SC}} - d^*)/d^*$ averaged 19.7%, with a maximum of 45%.

6.3 Discussion

The experiments confirm the theoretical complexity landscape:

- Exhaustive is practical only for $n \leq 8$ and $k \leq 5$.
- V-Fixed extends feasibility to $n \leq 10$ by eliminating the 2^{kn} factor.
- Single-Class scales to larger models ($n \geq 14$) with polynomial cost, achieving exact solutions in most cases for shallow formulas.

These results suggest a practical strategy: apply Single-Class first; if it fails or optimality guarantees are required, fall back to V-Fixed (for moderate n) or Exhaustive (for small instances). Future work on SAT/MaxSAT encodings (cf. direction F4 in Sect. 8) may further extend the tractable regime.

7 Limitations

We identify several limitations of the proposed framework, organized by category.

Computational intractability. The R-Mod-Rev decision problem is NP-complete (Thm. 5.1), and exhaustive computation requires super-exponential time in the number of worlds (Thm. 5.2). While the polynomial-time Single-Class refinement (Thm. 5.4) mitigates this for many practical instances, it is inherently incomplete: problems requiring coordinated changes across multiple equivalence classes cannot be solved by local search. The empirical evaluation (Sect. 6) quantifies this limitation, showing success rates below 50% for deep formulas. Bridging this gap—e.g., via SAT encodings or parameterized algorithms—remains open.

Restriction to S5 semantics. The current framework addresses only S5 knowledge, characterized by equivalence-based accessibility. It does not directly apply to:

- *KD45 belief revision*, where accessibility is serial, transitive, and Euclidean, and plausibility orderings play a central role;
- *Multi-agent epistemic logics*, where each agent has its own accessibility relation and common knowledge requires product constructions;
- *Probabilistic or graded modalities*, where revision involves updating probability distributions rather than discrete truth assignments.

Extending R-Mod to these richer semantics requires new distance measures (e.g., over preorders or probability spaces) and revised metatheoretic analyses.

Non-determinism and selection. When multiple models achieve the minimal distance $d^*(M, \alpha)$, R-Mod returns the entire set of minimizers. The skeptical belief set $\mathrm{Bel}_\star(M, \alpha)$ is defined as their intersection, which may be too weak for applications requiring definite conclusions. Our inertia criterion (IC) provides one tie-breaking mechanism, but alternative selection functions—such as preferring models that preserve more higher-order beliefs or satisfy additional side constraints—are not explored. The choice of selection function significantly affects the resulting theory and its metatheoretic properties.

Built-in Success and prioritized revision. R-Mod unconditionally accepts the input formula α (the Success postulate, K2). This prioritized design is appropriate when the source of α is authoritative—e.g., a verified sensor reading or a trusted protocol correction—but may be unsuitable when the new information is defeasible or conflicts with strongly entrenched prior knowledge. In such cases, a rational agent might prefer to reject α and retain the current model, a response that R-Mod cannot produce. Developing a non-prioritized extension that weighs the cost of accommodating α against the epistemic value of the prior model is a natural direction for future investigation.

Failure of classical AGM postulates. As demonstrated in Sect. 4, the supplementary AGM postulates K7 (Superexpansion) and K8 (Subexpansion) do not hold for R-Mod in general. This limits direct comparison with classical belief revision operators and may complicate integration with existing AGM-based systems. The conditions under which these postulates can be recovered (e.g., restricting to propositional inputs or forbidding structural changes) deserve further axiomatic investigation.

Sensitivity to distance metric. The quotient-based distance D° treats all structural and valuation changes uniformly. In practice, some modifications may be more "costly" than others:

- Splitting an equivalence class (reducing an agent's certainty) versus merging classes (increasing certainty) may have asymmetric epistemic significance.
- Flipping the truth value of a "core" proposition may be more disruptive than changing a peripheral atom.

Weighted distance functions could capture such distinctions but would require re-establishing the metatheoretic results and re-analyzing computational complexity.

Static, single-shot revision. R-Mod performs a single revision step. Iterated revision—where an agent receives a sequence of updates—introduces additional challenges:

- *Path dependence*: The order of revisions may affect the final model.
- *Cumulative minimality*: Ensuring that cumulative distance remains minimal across multiple steps is not addressed.
- *Memory and commitment*: The current framework has no mechanism for an agent to "commit" to previous revisions or to prefer stability over minimal change.

Integrating R-Mod with iterated revision postulates (e.g., Darwiche–Pearl) is left for future work.

Lack of real-world application benchmarks. Our empirical evaluation uses synthetically generated models and formulas. While this suffices to validate theoretical complexity bounds, it does not demonstrate applicability to concrete domains such as epistemic planning, knowledge base debugging, or security protocol verification. Developing domain-specific benchmarks would strengthen the practical relevance of R-Mod.

8 Conclusion and Future Work

This work develops *R-Mod*, a knowledge-model revision operator for S5 that performs minimal structural repair of epistemic models. The operator jointly reconfigures the S5 partition induced by the accessibility relation and, when necessary, adjusts valuations, while preserving S5 and minimizing a bisimulation-aware distance on quotients. The construction provides a precise selection semantics (Def. and Eq. in Sect. 3), algorithmic realization (Alg. 1), and metatheoretic guarantees at the skeptical level: success, consistency preservation, deductive closure, and extensionality under a formula-insensitive objective (Sect. 4). A weak vacuity result holds when the input sentence is already satisfied; stronger vacuity and the AGM schemata (K3), (K7), and (K8) are shown to fail in general due to permissible structural amplification, with explicit counterexamples (Sect. 4). The complexity landscape is characterized by NP-completeness of the associated decision problem, NP-hardness of the exact minimum-distance computation, and tight upper bounds for exhaustive and restricted variants (Sect. 5). In particular, exhaustive search is doubly exponential in the number of worlds and exponential in the number of relevant atoms, the valuation-fixed fragment lowers the exponential dependence to partitions only, and an α-guided single-class refinement achieves a polynomial bound by enforcing locality of change.

The results clarify the appropriate reading of S5-based change as *knowledge-model revision* (ontic correction) rather than as propositional belief revision in the sense of AGM. This perspective reconciles factivity at each time slice with the possibility of knowledge loss or gain across updates, and explains why the full suite of AGM postulates is not expected to hold without additional structural constraints. The quotient-based metric ensures modal invariance of the minimization and addresses bisimulation-sensitivity issues raised in model-theoretic critiques.

Future directions. Several extensions are natural.

- (F1) **Other frame classes.** Generalize R-Mod to KD45 and plausibility-based semantics by replacing S5 partitions with serial, transitive, Euclidean relations or with preorder rankings. This includes designing distance measures over orders (e.g., minimal pairwise inversions) and establishing counterparts of success, consistency, and extensionality, together with adapted vacuity principles.
- (F2) **Axiomatization and representation theorems.** Develop sound and complete dynamic calculi whose postulates capture the quotient-minimal semantics of R-Mod. Investigate conditions that recover AGM-like schemata (K3/K7/K8), for example by forbidding structural changes or by restricting inputs to propositional sentences, and provide precise representation results for these fragments.
- (F3) **Iterated revision and stability.** Analyze behavior under sequences of updates, including invariants under bisimulation compression, stability of tie-breaking mechanisms (e.g., inertia), and compatibility with iterated-revision postulates. Characterize when cumulative applications commute with quotient minimization.
- (F4) **Algorithmic improvements.** Beyond exhaustive enumeration and single-class refinement, design incremental algorithms that reuse evaluations across nearby partitions, and SAT/MaxSAT or ILP encodings that exploit the finite-quotient structure. Study parameterized and approximation algorithms with parameters such as the number of affected classes, the size of $\mathrm{Var}(\alpha)$, or bounds on class sizes.
- (F5) **Weighted and constrained metrics.** Extend the objective to weighted distances that differentiate structural edits and valuation flips, and analyze how weights affect metatheoretic properties. Incorporate side constraints (e.g., action preconditions, integrity constraints) while maintaining tractable special cases.

(F6) **Multi-agent and DEL integration.** Combine R-Mod with event-model updates supporting postconditions, common-knowledge formation, and inter-agent coupling. Establish modularity results that separate fact-changing events from partition reconfiguration while controlling overall minimality.

(F7) **Empirical evaluation and benchmarks.** Implement the operator and its restricted variants on synthetic and application-driven instances, quantify distance–accuracy tradeoffs, and assess scalability of the polynomial-time locality procedures relative to exhaustive baselines.

In sum, the proposed framework supplies a structurally faithful and computationally analyzable account of minimal knowledge-model revision for S5. The outlined extensions aim to broaden its logical coverage, strengthen its proof-theoretic foundations, and advance its algorithmic practicality.

References

T. Ågotnes, P. Balbiani, H. van Ditmarsch, and P. Seban. 2010. "Group announcement logic." *Journal of Applied Logic*, 8, 1, 62–81. doi:10.1016/j.jal.2008.12.002.

C. Aguilera-Ventura, J. Ben-Naim, and A. Herzig. 2025. "Minimal Change in Modal Logic S5." *Proceedings of the AAAI Conference on Artificial Intelligence*, 39, 14, 14781–14789. doi:10.1609/aaai.v39i14.33620.

C. E. Alchourrón, P. Gärdenfors, and D. Makinson. 1985. "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." *The Journal of Symbolic Logic*, 50, 2, 510–530. doi:10.2307/2274239.

P. Balbiani, A. Baltag, H. van Ditmarsch, A. Herzig, T. Hoshi, and T. de Lima. 2008. ""Knowable" as "Known After an Announcement"." *The Review of Symbolic Logic*, 1, 3, 305–334. doi:10.1017/S1755020308080210.

A. Baltag, L. S. Moss, and S. Solecki. 1998. "The Logic of Public Announcements, Common Knowledge, and Private Suspicions." In: *Proceedings of the Seventh Conference on Theoretical Aspects of Rationality and Knowledge (TARK 1998)*. Morgan Kaufmann, Evanston, IL, USA, 43–56.

A. Baltag and S. Smets. 2008. "A Qualitative Theory of Dynamic Interactive Belief Revision." In: *Logic and the Foundations of Game and Decision Theory (LOFT 7)*. Texts in Logic and Games. Vol. 3. Ed. by G. Bonanno, W. van der Hoek, and M. Wooldridge. Amsterdam University Press, 9–58.

A. Baltag and S. Smets. 2006. "Conditional Doxastic Models: A Qualitative Approach to Dynamic Belief Revision." *Electronic Notes in Theoretical Computer Science*, 165, 5–21. doi:10.1016/j.entcs.2006.05.034.

V. Belle, T. Bolander, A. Herzig, and B. Nebel. 2023. "Epistemic Planning: Perspectives on the Special Issue." *Artificial Intelligence*, 316, 103842. doi:10.1016/j.artint.2022.103842.

J. van Benthem. 2007. "Dynamic Logic for Belief Revision." *Journal of Applied Non-Classical Logics*, 17, 2, 129–155. doi:10.3166/jancl.17.129-155.

J. van Benthem. 2006. "One is a Lonely Number: On the Logic of Communication." In: *Logic Colloquium '02*. Lecture Notes in Logic. Vol. 27. Ed. by Z. Chatzidakis, P. Koepke, and W. Pohlers. Association for Symbolic Logic, 96–129.

J. van Benthem, J. van Eijck, and B. Kooi. 2006. "Logics of Communication and Change." *Information and Computation*, 204, 11, 1620–1662. doi:10.1016/j.ic.2006.04.006.

P. Blackburn, M. de Rijke, and Y. Venema. 2001. *Modal Logic*. Cambridge Tracts in Theoretical Computer Science. Vol. 53. Cambridge University Press, Cambridge, UK. ISBN: 9780521802000.

O. Board. 2004. "Dynamic Interactive Epistemology." *Games and Economic Behavior*, 49, 1, 49–80. doi:10.1016/j.geb.2003.10.006.

T. Bolander and M. B. Andersen. 2011. "Epistemic Planning for Single- and Multi-Agent Systems." *Journal of Applied Non-Classical Logics*, 21, 1, 9–34. doi:10.3166/jancl.21.9-34.

T. Bolander, A. Burigana, and M. Montali. 2025. "Depth-Bounded Epistemic Planning." In: *Proceedings of the 22nd International Conference on Principles of Knowledge Representation and Reasoning (KR 2025)*, 729–739. doi:10.24963/kr.2025/70.

T. Bolander, T. Charrier, S. Pinchinat, and F. Schwarzentruber. 2020. "DEL-Based Epistemic Planning: Decidability and Complexity." *Artificial Intelligence*, 287, 103304. doi:10.1016/j.artint.2020.103304.

G. Bonanno. 2025a. "A Kripke–Lewis Semantics for Belief Update and Belief Revision." *Artificial Intelligence*, 339, 104259. doi:10.1016/j.artint.2024.104259.

G. Bonanno. 2025b. "A Modal-Logic Translation of the AGM Axioms for Belief Revision." *The European Journal on Artificial Intelligence*. OnlineFirst. doi:10.1177/30504554251391055.

M. C. Cooper, A. Herzig, F. Maffre, F. Maris, E. Perrotin, and P. Régnier. 2021. "A Lightweight Epistemic Logic and Its Application to Planning." *Artificial Intelligence*, 298, 103437. doi:10.1016/j.artint.2020.103437.

H. van Ditmarsch. 2005. "Prolegomena to Dynamic Logic for Belief Revision." *Synthese*, 147, 2, 229–275. doi:10.1007/s11229-005-1359-0.

H. van Ditmarsch. 2023. "To Be Announced." *Information and Computation*, 292, 105026. doi:10.1016/j.ic.2023.105026.

H. van Ditmarsch, W. van der Hoek, and B. Kooi. 2007. *Dynamic Epistemic Logic*. Synthese Library. Vol. 337. Springer, Dordrecht. ISBN: 978-1-4020-5838-7. doi:10.1007/978-1-4020-5839-4.

H. van Ditmarsch and R. Kuznets. 2025. “Wanted Dead or Alive: Epistemic Logic for Impure Simplicial Complexes.” *Journal of Logic and Computation*, 35, 6, exae055. doi:10.1093/logcom/exae055.

T. Engesser, T. Bolander, R. Mattmüller, and B. Nebel. 2017. “Cooperative Epistemic Multi-Agent Planning for Implicit Coordination.” In: *Proceedings of the 9th Workshop on Methods for Modalities (M4M 2017)* (Electronic Proceedings in Theoretical Computer Science). Vol. 243, 75–90. doi:10.4204/EPTCS.243.6.

P. Gärdenfors. 1988. *Knowledge in Flux: Modeling the Dynamics of Epistemic States*. MIT Press, Cambridge, MA. ISBN: 9780262071093.

É. Goubault, R. Kniazev, and J. Ledent. 2024. “A Many-Sorted Epistemic Logic for Chromatic Hypergraphs.” In: *Proceedings of the 32nd EACSL Annual Conference on Computer Science Logic (CSL 2024)* (Leibniz International Proceedings in Informatics). Vol. 288, 30:1–30:18. doi:10.4230/LIPIcs.CSL.2024.30.

É. Goubault, R. Kniazev, J. Ledent, and S. Rajsbaum. 2023. “Semi-Simplicial Set Models for Distributed Knowledge.” In: *Proceedings of the 38th Annual ACM/IEEE Symposium on Logic in Computer Science (LICS 2023)*, 1–13. doi:10.1109/LICS56636.2023.10175737.

É. Goubault, R. Kniazev, J. Ledent, and S. Rajsbaum. 2024. “Simplicial Models for the Epistemic Logic of Faulty Agents.” *Boletín de la Sociedad Matemática Mexicana*, 30, 3, 90. doi:10.1007/s40590-024-00656-x.

É. Goubault, J. Ledent, and S. Rajsbaum. 2021. “A Simplicial Complex Model for Dynamic Epistemic Logic to Study Distributed Task Computability.” *Information and Computation*, 278, 104597. doi:10.1016/j.ic.2020.104597.

A. Grove. 1988. “Two Modellings for Theory Change.” *Journal of Philosophical Logic*, 17, 2, 157–170. doi:10.1007/BF00247909.

S. O. Hansson. 1999. “A Survey of Non-Prioritized Belief Revision.” *Erkenntnis*, 50, 2, 413–427. doi:10.1023/A:1005534223776.

H. Leitgeb and K. Segerberg. 2007. “Dynamic Doxastic Logic: Why, How, and Where To?” *Synthese*, 155, 2, 167–190. doi:10.1007/s11229-006-9143-8.

E. D. Mares. 2002. “A Paraconsistent Theory of Belief Revision.” *Erkenntnis*, 56, 2, 229–246. doi:10.1023/A:1015690931863.

C. Muise, V. Belle, P. Felli, S. A. McIlraith, T. Miller, A. R. Pearce, and L. Sonenberg. 2022. “Efficient Multi-Agent Epistemic Planning: Teaching Planners About Nested Belief.” *Artificial Intelligence*, 302, 103605. doi:10.1016/j.artint.2021.103605.

J. Plaza. 2007. “Logics of Public Communications.” *Synthese*, 158, 2, 165–179. Originally published in 1989. doi:10.1007/s11229-007-9168-7.

V. Punčochař, I. Sedlár, and A. Tedder. 2023. “Relevant Epistemic Logic with Public Announcements and Common Knowledge.” *Journal of Logic and Computation*, 33, 2, 436–461. doi:10.1093/logcom/exac100.

K. Sauerwald and M. Thimm. 2024. “The Realizability of Revision and Contraction Operators in Epistemic Spaces.” In: *Proceedings of the 21st International Conference on Principles of Knowledge Representation and Reasoning (KR 2024)*, 665–670. doi:10.24963/kr.2024/62.

K. Segerberg. 1995. “Belief Revision From the Point of View of Doxastic Logic.” *Logic Journal of the IGPL*, 3, 4, 535–553. doi:10.1093/jigpal/3.4.535.

W. Spohn. 1988. “Ordinal Conditional Functions: A Dynamic Theory of Epistemic States.” In: *Causation in Decision, Belief Change, and Statistics, II*. Ed. by W. L. Harper and B. Skyrms. D. Reidel, Dordrecht, 105–134.

P. Vigiani. 2025. “Ramsey Conditionals in Dynamic Relevant Logic.” *Studia Logica*, 1–37. doi:10.1007/s11225-025-10221-w.

Y. Wang and Q. Cao. 2013. “On Axiomatizations of Public Announcement Logic.” *Synthese*, 190, S1, 103–134. doi:10.1007/s11229-012-0233-5.

Received 08 January 2026; accepted 02 April 2026