Constructive Preference Relations: Navigating Undecidability in Rational LTL Contraction

Hannes Gaißer[1], Dominik Klumpp[2] and Jandson S Ribeiro[3]

[1]University of Tübingen, Germany
[2]LIX – CNRS – École Polytechnique, France
[3]Cardiff University, United Kingdom

Abstract

We study the computational aspects of *epistemic preference relations* in non-classical logics, particularly *linear temporal logic* (LTL). Epistemic preferences form the backbone of *belief contraction operators*, which describe how to rationally relinquish obsolete beliefs. These preference relations have to satisfy certain innocuous conditions; and constructing such relations is usually assumed to be a trivial process. However, in the case of LTL, where relations are represented with Büchi automata, we show that this is a challenging task: the core condition, which guarantees the success of contraction, is in fact *undecidable.* Towards achieving effective LTL belief contraction, we then propose several concrete constructions of novel preference relations that satisfy the required conditions *by design.* These constructions include, among others, (1) generalisations of distance measures (e.g. Dalal) beyond the classical setting, as well as (2) the ability to hierarchically compose different preference relations. Our results not only provide rich families of preference relations for LTL, but also generalise the limited pool of concrete preference relations for the classical cases, allowing us to go beyond Dalal to achieve full rationality.

1. Introduction

In *belief change* [1, 2], the process of relinquishing an obsolete belief is studied under the name of *contraction.* Ideally, only beliefs associated with the entailment of the obsolete belief should be discarded, so the incurred changes are minimised. This principle of minimal change is conceptualised via sets of rationality postulates that prescribe what a rational change is. Families of constructive operators that abide by such postulates were identified, called rational operators. Contraction is central in belief change, as several other change processes internalise contraction. For example, *revision*, the process of accommodating a new belief while maintaining consistency, consists of contracting potential conflicts with the new belief and then incorporating it. This makes contraction an ideal component for identifying and understanding how modifications are conducted and, therefore, to properly address minimal change. We consider belief contraction in the setting of reasoning with temporal information, specifically *linear temporal logic* (LTL, for short) [3].

Preference relations are the *de facto* strategy to construct rational change operators. For instance, in *iterated belief revision* [4, 5] whose focus is on the impact of the order in which new beliefs arrive, total relations and rankings[1] are the canonical way for representing an agent's beliefs, while the constructive operators essentially bend a preference relation into a new one. Although most of the research in belief change concentrates in the classical settings [6, 7], such as classical propositional logic, first order logic and some sub-classical cases, such as Horn, recently efforts have been put to extend belief change to more expressive logics capable of expressing intricate knowledge domains such as description logics [8, 9], CTL [10, 11] and LTL [3]. Due to the higher expressive power of these logics, the constructive contraction operators in these settings differ from the classical ones, though preferences relations still remain the main component to obtain such rational constructive operators [12, 13, 14]. In this sense, preference relations form the backbone of rational operators, even beyond the classical cases. Such preference relations must satisfy some innocuous conditions to guarantee rationality of the contraction operators.

24th International Workshop on Nonmonotonic Reasoning, July 17–19, 2026, Lisbon, Portugal

[1]A ranking is a special kind of preference relation where each object is assigned a confidence/preference number.

Belief change investigates and provides solutions for problems intrinsically connected with the flux of knowledge such as *ontology repair, reasoning in the presence of inconsistency, data integration* and *non-monotonic reasoning*, to cite a few. This makes belief change a central component to successfully support and implement various forms of knowledge dynamics and their respective applications. Therefore, it is paramount to understand the computational powers and limitations of belief change. Despite this central role of belief change in handling knowledge dynamics, the computational aspects of belief change are almost an untouched subject. The few works that look at it concentrate on specific change operators within classical propositional logics or Horn [15, 16, 17, 18]. Only recently, this question started to receive attention for logics beyond the classical spectrum, and revealed dire disruptive negative results. Guerra and Wassermann [19] have investigated some decidability problems regarding the traditional partial meet contraction operators in LTL, and have shown that deciding whether a given contraction operator is a partial meet operator is undecidable. Souza [20] investigated the problem of reasoning about belief change when the underlying knowledge is specified in CTL, and revision operators are encapsulated in the object language of preference logics. Klumpp and Ribeiro [14] have shown that there are uncountably many uncomputable contraction operators in LTL and other non-classical logics. They framed the precise class of computable contraction operators in LTL, using *Büchi automata*, and *Büchi-Mealy automata* to specify preference relations.

Given the central role of contraction in supporting and obtaining other forms of belief change operators, contraction is the ideal candidate for studying the computational aspects of belief change, as by effectively computing contraction, we lay the foundation to construct other forms of change operators. Following the results by Klumpp and Ribeiro [14], using Büchi-Mealy automata to specify preference relation in LTL, we consider the problem of verifying whether a preference relation specified as a given Büchi-Mealy automaton satisfies the conditions required to secure rationality of contraction. We show both positive and negative results, and discuss how to address the negative results.

Contributions and roadmap. Section 2 presents the basic notation and definitions used throughout the paper, and briefly reviews LTL and Büchi automata, while Section 3 reviews the rationality postulates of belief contraction, and the respective preference relations and rational constructions for LTL. The subsequent two sections contain our main contributions:

- Section 4 investigates the decidability of the two central conditions for rationality. One of them, *mirroring*, is decidable (Theorem 17), while the other condition, *maximal cut*, which ensures success of the contraction, is undecidable (Theorem 27). The same applies to a stricter form of maximal cut, *well-foundedness*. We explain why this rules out the possibility of certain representation results.

- In Section 5, we show how to achieve effective LTL contraction despite these issues, by presenting constructive families of preference relations that guarantee *mirroring* and *maximal cut* by design. Such families of relations not only provide a way to compute contraction for LTL, but also provide richer alternatives of preference relations for the classical settings.

Finally, Section 6 discusses our results, their impact on the field and potential future research paths.

An extended version of the paper including proofs of our results is available online [21].

2. LTL and Büchi Automata

We introduce some notations and definitions used throughout the paper, and briefly review *linear temporal logic*, LTL for short. The power set of a set X is denoted $\mathcal{P}(X)$. For the remainder of the paper, we fix a finite, non-empty set of atomic propositions AP.

Definition 1 (LTL Formulae). The formulae of LTL *are generated by the grammar*

$$\varphi ::= \bot \mid p \mid \neg\varphi \mid \varphi \vee \varphi \mid \mathbf{X}\,\varphi \mid \varphi\,\mathbf{U}\,\varphi\ ,$$

where p *ranges over* AP*.* Fm_{LTL} *denotes the set of all LTL formulae.*

In LTL, time is interpreted as a single timeline that unfolds to the future. Formally, timelines are *traces*, infinite sequences $\tau = a_0 a_1 \cdots$ where each $a_i \in \mathcal{P}(AP)$ is the set of atomic propositions that holds at the time step i. We write τ^i for the trace $a_i a_{i+1} \cdots$. The semantic of LTL is given in terms of Kripke structures and their traces.

Definition 2 (Kripke Structure). *A* Kripke structure $M = (S, I, T, \lambda)$ *consists of a finite set of states S; a non-empty set of initial states $I \subseteq S$; a left-total transition relation $T \subseteq S \times S$, i.e., for all $s \in S$ there exists $s' \in S$ with $(s, s') \in T$; and a labeling $\lambda : S \to \mathcal{P}(AP)$ of states with sets of atomic propositions.*

A trace of a Kripke structure M is a sequence $\tau = \lambda(s_0)\lambda(s_1)\lambda(s_2)\cdots$ with $s_0 \in I$, and for all $i \geq 0$, $s_i \in S$ and $(s_i, s_{i+1}) \in T$. The set of all traces of a Kripke structure M is given by $\mathit{Traces}(M)$.

The satisfaction relation between Kripke structures and LTL formulae is defined in terms of the satisfaction between the Kripke structure's traces and LTL formulae.

Definition 3 (Satisfaction). *The* satisfaction relation *is the least relation* $\models\ \subseteq \mathcal{P}(AP)^\omega \times \mathit{Fm}_{LTL}$ *between traces and LTL formulae such that, for all $\tau = a_0 a_1 \cdots \in \mathcal{P}(AP)^\omega$:*

$$\begin{array}{lll@{\qquad}lll}
\tau \not\models \bot & & & \tau \models \varphi_1 \vee \varphi_2 & \textit{iff} & \tau \models \varphi_1 \textit{ or } \tau \models \varphi_2 \\
\tau \models p & \textit{iff} & p \in a_0 & \tau \models \mathbf{X}\,\varphi & \textit{iff} & \tau^1 \models \varphi \\
\tau \models \neg\varphi & \textit{iff} & \pi \not\models \varphi & \tau \models \varphi_1\,\mathbf{U}\,\varphi_2 & \textit{iff} & \textit{there exists } i \geq 0 \textit{ s.t. } \tau^i \models \varphi_2 \\
& & & & & \textit{and for all } j < i, \tau^j \models \varphi_1
\end{array}$$

The set of all traces satisfying an LTL formula φ is denoted $\|\varphi\|$. Traces that do not satisfy φ are called *contra-traces* of φ. As a trace does not satisfy a formula φ iff it satisfies its negation $\neg\varphi$, the set of all contra-traces of φ is $\|\neg\varphi\|$. A Kripke structure M satisfies a formula φ, denoted $M \models \varphi$, iff all traces of M satisfy φ. M satisfies a set X of formulae, $M \models X$, iff $M \models \varphi$ for all $\varphi \in X$.

From the satisfaction relation, we define a consequence operator $\mathit{Cn}_{LTL} : \mathcal{P}(\mathit{Fm}_{LTL}) \to \mathcal{P}(\mathit{Fm}_{LTL})$ which maps each set of formulae X to all its consequences.

Definition 4 (Consequence Operator). *The* consequence operator Cn_{LTL} *maps each set X of LTL formulae to the set of all formulae ψ, such that for all Kripke structures M, if $M \models X$ then also $M \models \psi$.*

A *theory* is a logically closed set of formulae $\mathcal{K}$, that is, $\mathit{Cn}_{LTL}(\mathcal{K}) = \mathcal{K}$. The expansion of a theory $\mathcal{K}$ by a formula φ is the theory $\mathcal{K} + \varphi := \mathit{Cn}_{LTL}(\mathcal{K} \cup \{\varphi\})$. A theory $\mathcal{K}$ is *consistent* if $\mathcal{K} \neq \mathit{Fm}_{LTL}$, and it is *complete* if for all formulae $\varphi \notin \mathcal{K}$, we have $\mathcal{K} + \varphi = \mathit{Fm}_{LTL}$. The set of all complete consistent theories is denoted as CCT.

Büchi automata are finite automata widely used in formal specification and verification of systems, especially in LTL model checking [11], as well as in planning [22, 23] when goals are specified in LTL.

Definition 5 (Büchi Automata). *A* Büchi automaton $A = (Q, \Sigma, \Delta, Q_0, R)$ *consists of a finite set of states Q; a finite, nonempty alphabet Σ (whose elements are called* letters*); a transition relation $\Delta \subseteq Q \times \Sigma \times Q$; a set of initial states $Q_0 \subseteq Q$; and a set of* recurrence states $R \subseteq Q$.

A Büchi automaton accepts an infinite word over the alphabet Σ if the automaton visits a recurrence state infinitely often while reading the word. Formally, an infinite word is a sequence $a_0 a_1 \ldots$ with $a_i \in \Sigma$ for all i. For a finite word $\rho = a_0 \ldots a_n$, with $n \geq 0$, let ρ^ω denote the infinite word derived by infinite repetition of ρ (we assume $\cdot^\omega$ binds stronger than concatenation). The set of all infinite words is denoted by Σ^ω. An infinite word $a_0 a_1 a_2 \ldots \in \Sigma^\omega$ is *accepted* by a Büchi automaton $A = (Q, \Sigma, \Delta, Q_0, R)$ if there exists a sequence $q_0, q_1, q_2, \ldots$ of states $q_i \in Q$ such that $q_0 \in Q_0$ is an initial state, for all i we have that $(q_i, a_i, q_{i+1}) \in \Delta$ and there are infinitely many $i \in \mathbb{N}$ with $q_i \in R$. The set $\mathcal{L}(A)$ of all accepted words is the *language* of A.

Emptiness of a Büchi automaton's language is decidable. Further, Büchi automata for the union, intersection and complement of the languages of given Büchi automata can be effectively constructed [24]. Unless otherwise noted, we always consider Büchi automata over the alphabet $\Sigma = \mathcal{P}(AP)$, where letters are sets of atomic propositions and infinite words are traces. The automata-theoretic treatment of LTL is based on the following result:

Proposition 6 (Clarke et al. [11])**.** *For each LTL formula φ and Kripke structure M, there exist Büchi automata A_φ and A_M that accept precisely $\mathcal{L}(A_\varphi) = \|\varphi\|$ resp. $\mathcal{L}(A_M) = \mathit{Traces}(M)$.*

We use Büchi automata as structures to represent an agent's beliefs. In the original AGM paradigm of belief change [1, 25], an agent's corpus of beliefs is represented as a theory. However, Büchi automata accept traces, rather than formulae. This bridge is accomplished by the notion of support, which takes all formulae satisfied by all traces accepted by a Büchi automaton. Klumpp and Ribeiro [14] have shown that the support is indeed a theory.

Definition 7 (Support [14])**.** *The* support *of a Büchi automaton A is the set $\mathcal{S}(A) := \{\, \varphi \in Fm_{LTL} \mid \forall \pi \in \mathcal{L}(A) \,.\, \pi \models \varphi \,\}$. If $\varphi \in \mathcal{S}(A)$, we say that A* supports *φ.*

The space of all theories represented by Büchi automata is called Büchi excerpt, denoted $\mathbb{E}_{\text{Büchi}}$. Notably, some theories cannot be represented by Büchi automata. In non-finitary logics like LTL, this is unavoidable for any method of finite representation [14].

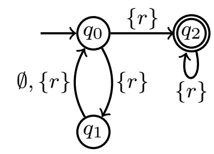

Figure 1: Büchi automaton with initial state q_0 and recurrence state q_2.

Example 8. *Let the atomic proposition r denote* "it rains"*. Figure 1 shows a Büchi automaton A, which supports the formulae r (*"it rains today"*), $\mathbf{X}^2\, r$ (*"it will rain in two days"*), $\mathbf{X}^4\, r$ (*"it will rain in 4 days"*), etc. However, A does not support $\mathbf{G}\, r$ (*"it rains every day"*) because it accepts the trace $\{r\}\,\emptyset\,\{r\}^\omega$.*

3. Belief Contraction in LTL

The field of belief change is founded on the AGM paradigm [25, 1], originally devised for classical logics. In the original, a logic is treated as a pair (Fm, Cn) where Fm denotes the language of the logic, and Cn is a consequence operator mapping each set of formulae to the set of all its entailments. The epistemic state of an agent is represented as a theory. A contraction function on a theory $\mathcal{K}$[2] is a function $\dot{-} : Fm \to \mathcal{P}(Fm)$ that, given an obsolete piece of information φ, produces a theory which does not entail φ. Contraction functions are subject to the following rationality postulates [1]:

$(\mathbf{K_1^-})$ $\mathcal{K} \dot{-} \varphi = Cn(\mathcal{K} \dot{-} \varphi)$	(closure)	$(\mathbf{K_2^-})$ $\mathcal{K} \dot{-} \varphi \subseteq \mathcal{K}$	(inclusion)
$(\mathbf{K_3^-})$ If $\varphi \notin \mathcal{K}$ then $\mathcal{K} \dot{-} \varphi = \mathcal{K}$	(vacuity)	$(\mathbf{K_4^-})$ If $\varphi \notin Cn(\emptyset)$ then $\varphi \notin \mathcal{K} \dot{-} \varphi$	(success)
$(\mathbf{K_5^-})$ $\mathcal{K} \subseteq (\mathcal{K} \dot{-} \varphi) + \varphi$	(recovery)	$(\mathbf{K_6^-})$ If $\varphi \equiv \psi$ then $\mathcal{K} \dot{-} \varphi = \mathcal{K} \dot{-} \psi$	(extensionality)

$(\mathbf{K_7^-})$ $(\mathcal{K} \dot{-} \varphi) \cap (\mathcal{K} \dot{-} \psi) \subseteq \mathcal{K} \dot{-} (\varphi \wedge \psi)$ $\qquad$ $(\mathbf{K_8^-})$ If $\varphi \notin \mathcal{K} \dot{-} (\varphi \wedge \psi)$ then $\mathcal{K} \dot{-} (\varphi \wedge \psi) \subseteq \mathcal{K} \dot{-} \varphi$

For a detailed discussion on the rationale of these postulates, see [25, 1, 2]. A contraction function that satisfies $(\mathbf{K_1^-})$ to $(\mathbf{K_6^-})$ is called a *rational* contraction function. If a contraction function satisfies all the eight rationality postulates, we say that it is *fully rational.*

Several classes of (fully) rational AGM contraction functions were proposed on classical logics [2]. In more expressive, non-classical logics such as LTL, however, such classical classes of functions are too weak to capture recovery. To embrace more expressive logics, Ribeiro et al. [12] have proposed new classes of (fully) rational contraction functions. Upon these, Klumpp and Ribeiro [14] characterised precisely the class of all rational AGM contraction functions on LTL theories in the Büchi excerpt, called Büchi contraction functions.

Büchi contraction functions work by selecting the most plausible contra-traces of the LTL formula to be contracted, and incorporating these traces. The selection is realised by a *Büchi choice function*, guaranteeing that the set of selected contra-traces is represented as a Büchi automaton.

Definition 9 (Büchi Choice Functions)**.** *A* Büchi choice function *γ maps each LTL formula to a single Büchi automaton, such that for all LTL formulae φ and ψ,*

[2]Note that we are following the traditional local definition of contraction function rather than the global one. As thoughtfully studied by Hansson [2], global functions do not bring any benefit on one-shot contractions.

(BF1) *the language accepted by $\gamma(\varphi)$ is non-empty;*

(BF2) *$\gamma(\varphi)$ supports $\neg\varphi$, if φ is not a tautology; and*

(BF3) *$\gamma(\varphi)$ and $\gamma(\psi)$ accept the same language, if $\varphi \equiv \psi$.*

As the purpose is to relinquish a formula φ, the choice function must select only among the contra-traces of φ **(BF2)**, except for tautologies, which cannot be removed. To be successful, at least one contra-trace of φ must be selected **(BF1)**. The choice is not syntax-sensitive **(BF3)**. The contraction corresponds to assimilating all selected traces into the current Büchi automaton, as long as the formula is not a tautology, in which case the theory remains untouched.

Definition 10 (Büchi Contraction Functions)**.** *Let $\mathcal{K}$ be a theory in the Büchi excerpt and let γ be a Büchi choice function. The* Büchi Contraction Function *on $\mathcal{K}$ induced by γ is the function $\dot{-}_\gamma$ such that $\mathcal{K} \dot{-}_\gamma \varphi = \mathcal{K} \cap \mathcal{S}(\gamma(\varphi))$ if $\varphi \notin Cn(\emptyset)$ and $\varphi \in \mathcal{K}$, or $\mathcal{K} \dot{-}_\gamma \varphi = \mathcal{K}$ otherwise.*

Theorem 11 (Klumpp and Ribeiro [14])**.** *A contraction function $\dot{-}$ in the* Büchi *excerpt is rational if and only if $\dot{-}$ is a* Büchi *contraction function.*

For the supplementary postulates ($\mathbf{K_7^-}$) and ($\mathbf{K_8^-}$), the choice function has to be constrained further. The main idea is that the choices must be based on epistemic preferences an agent has about its beliefs, with some beliefs being more plausible than others. Such preferences are realised by a binary relation $\prec$ between the traces, where $\tau \prec \tau'$ indicates that τ' is more plausible than τ[3]The choice then boils down to choosing the most plausible contra-traces of the formula φ being relinquished, that is, $\mathcal{L}(\gamma_\prec(\varphi)) = \max_\prec \|\neg\varphi\|$. The preference relation is subject to two central conditions:

(Maximal Cut) $\max_\prec \|\neg\varphi\| \neq \emptyset$, if φ is not a tautology.

(Mirroring) If $\tau_1 \not\prec \tau_2$ and $\tau_2 \not\prec \tau_1$ but $\tau_1 \prec \tau_3$ then $\tau_2 \prec \tau_3$.

(Maximal Cut) guarantees the success of the contraction, by ensuring the existence of maximally plausible contra-traces for each formula. The second property **(Mirroring)** enforces that if two traces are incomparable, then the preferences over them must mimic one another. To obtain a Büchi choice function, it still necessary to guarantee that $\max_\prec \|\neg\varphi\|$ is recognised by a Büchi automaton. To this end, Klumpp and Ribeiro [14] proposed the use of *Büchi-Mealy automata* to represent such relations.

Definition 12 (Büchi-Mealy Automata)**.** *A* Büchi-Mealy automaton *(BMA, for short) is a Büchi automaton over the alphabet $\Sigma_{\text{BM}} = \mathcal{P}(AP) \times \mathcal{P}(AP)$.*

A Büchi-Mealy automaton B accepts infinite sequences of pairs $(a_1, b_1) \cdots (a_i, b_i) \cdots$. Such an infinite sequence encodes a pair of traces (τ_a, τ_b), where $\tau_a = a_1 \cdots a_i \cdots$ and $\tau_b = b_1 \cdots b_i \cdots$. Hence, a BMA B recognises the binary relation $\mathcal{R}(B) = \{(a_1 \ldots, b_1 \ldots) \in \Sigma^\omega \times \Sigma^\omega \mid (a_1, b_1) \cdots \in \mathcal{L}(B)\}$.

Klumpp and Ribeiro [14] have shown that the relation $\mathcal{R}(B)$ enjoys a key property: for each LTL formula φ, the most plausible traces of a formula φ wrt. $\mathcal{R}(B)$ are recognised by a Büchi automaton, which can be effectively computed. In fact, this property, together with **(Maximal Cut)**, guarantees that the preference relation $\mathcal{R}(B)$ of a BMA B induces a Büchi choice function γ_B with $\mathcal{L}\big(\gamma_B(\varphi)\big) = \max_{\mathcal{R}(B)} \|\neg\varphi\|$. The respective Büchi contraction function, denoted by $\dot{-}_{\mathcal{R}(B)}$, is guaranteed to be computable. If $\mathcal{R}(B)$ additionally satisfies **(Mirroring)**, the contraction function $\dot{-}_{\mathcal{R}(B)}$ is fully rational.

Theorem 13 (Klumpp and Ribeiro [14])**.** *If the preference relation $\mathcal{R}(B)$ of a BMA satisfies both* **(Maximal Cut)** *and* **(Mirroring)***, the induced* Büchi *contraction function is fully rational and computable.*

Example 14. *Let the proposition r once again denote that* "it rains"*. Figure 2 shows an example of a Büchi-Mealy automaton B. By convention, we write x/y on the edges in place of the pair (x, y).*

As shown to the right of the figure, the relation $\mathcal{R}(B)$ considers the trace $\emptyset\{r\}\emptyset^\omega$ (in which it does

[3]As most of our results are founded on Ribeiro et al. [12], we follow their interpretation on preferences as "the higher, the better", unlike the classical interpretation founded on the entrenchment notion "the more grounded, the better".

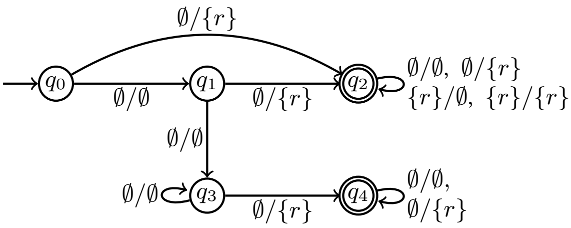

$\emptyset\,\{r\}\,\emptyset\,\emptyset\,\emptyset\cdots$	$\prec$	$\{r\}\,\emptyset\,\emptyset\,\emptyset\,\emptyset\cdots$
$\emptyset\,\emptyset\,\{r\}\,\emptyset\,\emptyset\cdots$	$\prec$	$\emptyset\,\{r\}\,\emptyset\,\emptyset\,\emptyset\cdots$
$\emptyset\,\emptyset\,\emptyset\,\{r\}\,\emptyset\cdots$	$\#$	$\emptyset\,\emptyset\,\{r\}\,\emptyset\,\emptyset\cdots$
$\emptyset\,\emptyset\,\emptyset\,\emptyset\,\emptyset\cdots$	$\prec$	$\emptyset\,\emptyset\,\emptyset\,\{r\}\,\emptyset\cdots$

Figure 2: A Büchi-Mealy automaton B (on the left), along with some examples of preferences between traces, as given by the relation $\prec\, = \mathcal{R}(B)$. The notation $\tau_1 \# \tau_2$ indicates that both $\tau_1 \not\prec \tau_2$ and $\tau_2 \not\prec \tau_1$ hold.

not rain today, but tomorrow) less plausible than the trace $\{r\}\emptyset^\omega$ (in which it rains today), due to the corresponding run $q_0\, q_2^\omega$. Similarly, the preference on the second line (preferring a trace in which it rains tomorrow over one in which it rains only in two days) is witnessed by the run $q_0\, q_1\, q_2^\omega$. There is no preference between the traces in the third line (rain in two vs. in three days), as there is no corresponding run (after the partial run $q_0\, q_1\, q_3\, q_4$, the automaton B is "stuck"). Finally, B considers the trace $\emptyset^\omega$ (in which it never rains) less plausible than $\emptyset\emptyset\emptyset\{r\}\emptyset^\omega$, by the run $q_0\, q_1\, q_3\, q_3\, q_4^\omega$. It can be shown that $\mathcal{R}(B)$ satisfies **(Mirroring)** *and* **(Maximal Cut)**. *We return to this point in Example 34.*

For an effective realization of fully rational LTL belief contraction based on epistemic preferences represented as BMA, it is thus necessary to ensure that the particular BMA relation satisfies **(Mirroring)** and **(Maximal Cut)**. While it is possible to manually verify these properties for individual BMA relations, allowing for the construction of bespoke contraction operators, a general framework of belief contraction requires a general solution to this problem. Two main alternatives are: (i) identify general constructions guaranteeing that the underlying relation can be encoded as a BMA and satisfies **(Maximal Cut)** and **(Mirroring)**, or (ii) provide devices capable of deciding whether the relation of a given BMA satisfies **(Maximal Cut)** and **(Mirroring)**. In the next section, we show some fundamental limitations of such a general solution.

4. On Deciding Mirroring and Maximal Cut

We investigate the *decidability* of the two fundamental properties of BMA relations that guarantee full rationality: **(Mirroring)** and **(Maximal Cut)**. We show that **(Mirroring)** is decidable, but **(Maximal Cut)** is undecidable. As we see later on, this result precludes the possibility of a constructive representation result for fully rational BMA relations. In addition to **(Maximal Cut)**, we also investigate the conceptually simpler but strictly stronger property *well-foundedness*. While this shift does not improve the decidability situation, well-foundedness forms the basis for a set of pragmatic constructions that guarantee full rationality (Section 5).

4.1. Mirroring

We begin our investigation with the property **(Mirroring)**. While the definition of this property in Section 3 is given in a *point-wise fashion* (by quantifying over individual points, i.e., traces), the property can be equivalently stated *algebraically*, using standard operators from relational algebra.

Lemma 15. *A relation R satisfies* **(Mirroring)** *if and only if the inclusion $(\overline{R} \cap (\overline{R})^{-1}) \circ R \;\subseteq\; R$ holds, where $\overline{\cdot}$ denotes the complement, $(\cdot)^{-1}$ denotes the inverse relation, and $\circ$ denotes relational composition.*

This algebraic characterisation is useful to establish decidability, due to the following observation:

Lemma 16. *Büchi-Mealy automata are effectively closed under union, intersection, complement, inverse relation and relational composition.*

Moreover, the inclusion between Büchi-Mealy automata is decidable, as a special case of Büchi automata inclusion [24]. We conclude:

Theorem 17. *The BMA mirroring problem, as given below, is decidable in* ExpTime.

> ***Input:*** *A Büchi-Mealy automaton B.*
> ***Output:*** yes *if $\mathcal{R}(B)$ satisfies* **(Mirroring)**, no *otherwise.*

Proof. It suffices to construct a BMA for the left-hand side of the inclusion in Lemma 15, where $R = \mathcal{R}(B)$, and decide (in ExpTime) the inclusion between this BMA and B. The dominating factor is the complementation of B, which takes exponential time and produces a complement automaton that may be exponentially larger than B [26]. The other operations are at worst quadratic. □

4.2. On Maximal Cut and Well-Foundedness

Before we show undecidability of **(Maximal Cut)** in Section 4.3, let us spend a moment studying the connection (and distinctions) between **(Maximal Cut)** and another, more classical property of relations: *well-foundedness.* We begin our comparison with an equivalent reformulation of **(Maximal Cut)**, which serves to highlight the connection between the properties:

(Maximal Cut) For every nonempty set $X \subseteq \Sigma^\omega$, if there exists $\varphi \in Fm_{LTL}$ with $X = \|\varphi\|$, then $\max_{\prec}(X) \neq \emptyset$.

As highlighted, this property requires a careful determination which sets X of traces can be represented by an LTL formula φ. By contrast, *well-foundedness* omits this restriction of the set X:

(Well-Founded) For every nonempty set $X \subseteq \Sigma^\omega$, it holds that $\max_{\prec}(X) \neq \emptyset$.
Or equivalently: There exists no infinite sequence $\tau_0 \prec \tau_1 \prec \ldots$, with $\tau_i \in \Sigma^\omega$ for all $i \in \mathbb{N}$.[4]

Clearly, **(Well-Founded)** implies **(Maximal Cut)**. The reverse implication depends on the logic. For instance, in classical propositional logic (over a finite signature), we have:

Observation 18. **(Maximal Cut)** *implies* **(Well-Founded)** *in propositional logic, but not in Horn logic.*

Generally, whether or not **(Maximal Cut)** and **(Well-Founded)** coincide depends on two parameters. The first parameter is the expressivity of single formulae in the logic: as we have seen, limiting propositional logic to Horn prevents certain sets from being expressed as formulae, leading to a rupture between the two properties. In non-finitary logics such as LTL, this rupture is in fact unavoidable: the countably many LTL formulae cannot capture the uncountably many sets of traces [14].

Proposition 19. *In LTL,* **(Maximal Cut)** *and* **(Well-Founded)** *do* not *coincide in general.*

Proof. Proposition 20 below proves a stronger result. □

The second parameter is the class of relations we consider. The observations above refer to the class of all preference relations. Yet, in non-finitary logics like LTL, most relations between the infinitely many traces cannot be finitely represented, and are thus not relevant when it comes to *effective* contraction. In the particular case of LTL, we therefore ask ourselves whether the rupture between the two properties persists when we restrict to (finitely-representable) BMA relations. It turns out that this is the case:

Proposition 20. *There exist BMA relations that satisfy* **(Mirroring)** *and* **(Maximal Cut)** *but do not satisfy* **(Well-Founded)**.

Proof sketch. For every $i \in \mathbb{N}$, let $\sigma_i := (\{p\}\{p\})^i\,\emptyset\,\{p\}^\omega$. Let $\tau \prec \tau'$ hold if and only if (1) $\tau = \sigma_i$ and $\tau' = \sigma_j$ for $i, j \in \mathbb{N}$ with $i < j$; or (2) $\tau = \sigma_i$ for $i \in \mathbb{N}$, but $\tau' \neq \sigma_j$ for all $j \in \mathbb{N}$. This relation does not satisfy **(Well-Founded)**, as the σ_i form an infinite ascending sequence. But $\prec$ satisfies **(Mirroring)** and **(Maximal Cut)**. For **(Maximal Cut)**, we appeal to a result by Wolper [27] showing that for every LTL formula φ, either some trace τ with $\tau \neq \sigma_j$ for all $j \in \mathbb{N}$ satisfies $\neg\varphi$ (and hence, τ is maximal in $\|\neg\varphi\|$), or there exists a (finite) maximal i such that σ_i satisfies $\neg\varphi$ (and hence, σ_i is maximal in $\|\neg\varphi\|$). It remains only to note that there exists a BMA recognizing the relation $\prec$. □

[4] Due to this formulation, **(Well-Founded)** is sometimes called the *ascending chain condition* (ACC).

As before, the rupture stems from a limitation of the expressivity of single LTL formulae. In this case, it is particularly the disconnect between the expressive powers of LTL resp. Büchi automata.

As we turn below to the decidability question for **(Well-Founded)** resp. **(Maximal Cut)** on BMA relations, a restricted class of BMA relations in which the two properties coincide plays a key role.

4.3. Undecidability of Maximal Cut

We show that **(Maximal Cut)** is undecidable for BMA relations via a reduction from a known undecidable problem. Specifically, we reduce the *immortality* of Turing machines (explained further below) to *well-foundedness* of BMA relations.

Central to our proof is the construction of a Büchi-Mealy automaton that simulates a single step of a given Turing machine. To extend the result to **(Maximal Cut)**, we formulate our construction using *linear bounded automata* (LBA), a special type of Turing machines.

Linear bounded automata function like regular Turing machines, consisting of a *read-write head* and a *tape*, a string of symbols. In each step, the head reads a symbol a on the tape, and depending on the state q of the automaton, replaces a with a new symbol b, moves left or right, and transitions into a new state q'. This behaviour is specified by a transition function δ. What distinguishes LBA from general Turing machines is that they may only use the space on the tape that is initially occupied by the input.

Definition 21. *A* linear bounded automaton $T = (Q, \Sigma, \Gamma, q_0, \delta)$ *consists of a finite set of states* Q*, a finite non-empty input alphabet* Σ*, a finite tape alphabet* Γ *which includes* Σ *and additional symbols* $<$ *and* $>$ *marking the left and right end of the tape, a start state* $q_0 \in Q$*, and a partial transition function* $\delta : Q \times \Gamma \to Q \times \Gamma \times \{L, R\}$*. The function* δ *must not include transitions that overwrite* $<$ *or* $>$ *with other symbols, move to the left of* $<$ *resp. the right of* $>$*, or overwrite other symbols with* $<$ *or* $>$*.*

Every combination of state $q \in Q$, tape content $w \in \Gamma^*$ and head position, denoted by the •-symbol above the corresponding symbol in w, form a *configuration* of a given linear bounded automaton. $qu\dot{a}v$ denotes the configuration where $w = uav$ for $u, v \in \Gamma^*$ and the automaton is in state q with the head above the symbol a. The function δ defines a successor relation $\to$ on configurations of T in the expected way. For example: $q{<}\dot{a}{>} \to q'{<}b\dot{>}$ if $\delta(q, a) = (q', b, R)$. A configuration c is *halting* if it has no successor, *mortal* if there is a halting configuration c' such that $c \to^* c'$, and *immortal* otherwise.

The following well-known property of linear bounded automata is central to our proof.

Lemma 22 (Chowdhary [28])**.** *Any infinite run of successive configurations* $c_1 \to c_2 \to \dots$ *of an LBA* T *only consists of finitely many unique configurations.*

Proof sketch. The tape never changes length, so all successive configurations must have the same tape length n. There are only finitely many configurations of a given length. ☐

For our proof, we use a variant of the halting problem, which is undecidable for Turing machines as well as linear bounded automata. This *immortality problem* [29] asks whether or not a given machine has any immortal configuration. No input is specified, in fact the immortal configuration does not even have to be reachable from any initial configuration.

Proposition 23 (Hooper [29])**.** *The* LBA Immortality Problem*, as given below, is undecidable.*

Input: *A linear bounded automaton* T*.*
Output: yes *if* T *has an immortal configuration,* no *otherwise.*

Our reduction proof rests on the following construction.

Lemma 24. *For a given LBA* T*, we can construct a Büchi-Mealy automaton* B_T *such that* $\mathcal{R}(B_T)$ *is* **(Well-Founded)** *if and only if* T *has no immortal configuration.*

Proof sketch. We construct a Büchi-Mealy automaton that accepts only trace-pairs that describe successive configurations $c_1 \to c_2$ by non-deterministically guessing the state q and the current symbol a in c_1, and then checking if the two configurations match according to the rule $\delta(q, a)$. For instance, if $\delta(q, a) = (q', a', R)$, the automaton expects states q and q' in the first and second trace respectively, then it expects the traces to match symbol-by-symbol until it finds $\dot{a}$ and a' in the first and second trace respectively. I.e., the head was above symbol a on the tape, which is now replaced by a' and the head is positioned elsewhere. The head moves to the right according to the rule, so the automaton expects it on the next symbol in the second trace. Afterwards, the traces must match again until the symbol $>$ appears on both traces, marking the end of the configuration.

Every immortal configuration c corresponds to an infinite run $c \to c_1 \to c_2 \to \dots$ of T and thus to to an infinite ascending chain in $\mathcal{R}(B_T)$, and vice versa. □

This construction allows us to reduce LBA immortality to BMA well-foundedness.

Theorem 25. *The Well-Foundedness problem for Büchi-Mealy automata, as given below, is undecidable.*

> ***Input:*** *A Büchi-Mealy automaton B.*
> ***Output:*** $\texttt{yes}$ *if $\mathcal{R}(B)$ satisfies* **(Well-Founded)***,* $\texttt{no}$ *otherwise.*

Let us shift our focus back to **(Maximal Cut)**, our original property of interest. We show its undecidability by linking it once again to well-foundedness – indeed, on relations corresponding to LBA, these two properties coincide:

Lemma 26. *The relation $\mathcal{R}(B_T)$, for an LBA T, is well-founded if and only if it satisfies* **(Maximal Cut)***.*

Proof sketch. We need only show that **(Maximal Cut)** implies **(Well-Founded)**; the opposite direction always holds. Therefore, suppose contrapositively that $\mathcal{R}(B_T)$ were not well-founded. Then by Lemma 24, T has some immortal configuration from which there exists an infinite run. By Lemma 22, this run consists of only finitely many configurations. We construct an LTL formula φ that describes exactly all the configurations of this run. Then, $\|\varphi\|$ cannot have a maximal element (it would correspond to a halting configuration in the infinite run). Hence, $\mathcal{R}(B_T)$ violates **(Maximal Cut)**; the formula $\neg\varphi$ serves as the counter-example. □

Hence, by the same construction as before, we can also reduce the LBA immortality problem to **(Maximal Cut)**. This is the last ingredient for our central theorem:

Theorem 27. *The following problem is undecidable:*

> ***Input:*** *A Büchi-Mealy automaton B.*
> ***Output:*** $\texttt{yes}$ *if $\mathcal{R}(B)$ satisfies* **(Maximal Cut)***,* $\texttt{no}$ *otherwise.*

Interestingly, our proof of undecidability does not even require the full expressive power of LTL. In fact, the only temporal operator used to construct the formula φ in Lemma 26 is $\mathbf{X}$ ("next"). The fixpoint operators ($\mathbf{U}$, or the derived operators $\mathbf{F}$ resp. $\mathbf{G}$) that allow reasoning over an unbounded temporal horizon are not needed for this result.

We conclude this section with a discussion on the impact of our undecidability result, namely that it implies the impossibility of a *constructive representation result.* To that end, we note:

Observation 28. *The* **(Maximal Cut)** *problem on BMA, as given in Theorem 27, is co-semi-decidable.*

The source of undecidability for **(Maximal Cut)** are thus the cases where the property indeed holds. It is not semi-decidable, i.e. it is impossible to *certify* the property in finite time. Suppose now a (finite) system of effective constructions of BMA, such that the resulting BMA are guaranteed to satisfy **(Maximal Cut)** by design. Indeed, we give just such a system in the next section. Our observation above implies that such a system *cannot possibly be complete*, i.e., there can be no representation theorem for **(Maximal Cut)**. There always must exist some BMA relations satisfying **(Maximal Cut)** that cannot be constructed with this system. Contradicting our observation, such a complete system would provide a semi-decision procedure for **(Maximal Cut)**: Given a BMA B, enumerate all BMA B' constructible with the system and decide one-by-one whether B and B' are equivalent.

Similarly, this observation precludes the possibility of a stronger, decidable (perhaps even syntactic) condition **(MC+)** on BMA that implies **(Maximal Cut)** and is "quasi-complete", in the sense that every BMA satisfying **(Maximal Cut)** is *equivalent* to one that satisfies **(MC+)**. This also yields a semi-decision procedure: to verify **(Maximal Cut)** for a given B, enumerate all BMA B', and decide for each of them whether B' satisfies **(MC+)**, and if so, whether B and B' are equivalent.

Although we cannot overcome undecidability with easy alternatives, the goal of effective LTL contraction is not invalidated. Rather this result motivates us to study epistemic preference relations suited for particular applications of contraction. The next section presents a first step in this direction.

5. Constructions for Rational Preferences

As maximal cut is undecidable, we cannot in general take a proposed BMA B and automatically check whether its relation satisfies **(Maximal Cut)**, in addition to **(Mirroring)**, thus certifying that B may be used to perform fully rational contractions. Instead, we propose in the following a number of constructions that guarantee *by design* that the resulting BMA relations satisfy both conditions.

As a first step, we analyze more closely the class of relations that satisfy **(Mirroring)** and **(Maximal Cut)**. To this end, we introduce a subclass of preference relations, called *rankings*.

Definition 29. *A relation $\prec \subseteq \Sigma^\omega \times \Sigma^\omega$ is a* ranking *over a totally-ordered set $(Y, <_Y)$ if there exists a function $f : \Sigma^\omega \to Y$ such that $\tau_1 \prec \tau_2$ if and only if $f(\tau_1) <_Y f(\tau_2)$.*

We say that the *ranking function f induces the ranking $\prec$*. Every ranking satisfies **(Mirroring)**, but not necessarily **(Maximal Cut)**. Though focusing on rankings might seem like a restriction, it is not:

Theorem 30. *For a BMA B, if $\mathcal{R}(B)$ satisfies* **(Mirroring)** *and* **(Maximal Cut)***, then $\mathcal{R}(B)$ is a ranking.*

Proof sketch. Let $\mathcal{R}(B)$ satisfy **(Mirroring)** and **(Maximal Cut)**. The relation $\mathcal{R}(B)$ must be acyclic, as from any cycle we can construct a formula violating **(Maximal Cut)**. By **(Mirroring)**, the elements incomparable wrt. $\mathcal{R}(B)$ form equivalence classes, and $\mathcal{R}(B)$ corresponds to a total order over these equivalence classes. Mapping each element to its equivalence class, we derive the ranking function f that induces $\mathcal{R}(B)$. □

This result states that, fundamentally, due to the expressiveness of LTL, *the only BMA relations guaranteeing full rationality* are rankings: traces are partitioned into subsets of equi-preferable traces, and these subsets must be *totally* ordered.

While we here state the theorem specifically for LTL and for BMA relations over traces, the underlying insight is broader. Santos [30] shows a representation result between fully-rational AGM contraction operators and *Blade Contraction Functions* induced by relations over complete consistent theories (CCTs), for all boolean Tarskian logics. In this general setting, Theorem 30 carries over to a wide class of expressive logics, the *compendious logics* [14]. We derive the following representation result:

Theorem 31. *Every fully rational AGM contraction operator over a compendious logic [14] is a Blade Contraction Function [30] induced by a* ranking *over CCTs that satisfies* maximal cut *[30], and vice versa.*

Consequently, we confine our study of BMA relations to rankings. We call such rankings *BMA rankings*. To ensure **(Maximal Cut)**, we focus on rankings over well-founded sets.

Observation 32. *If $\prec$ is a ranking over $(Y, <_Y)$ and $<_Y$ is well-founded, then $\prec$ is well-founded.*

In the remainder of this section, we provide three constructions of BMA relations that represent well-founded rankings, and one construction to *hierarchically compose* multiple BMA rankings while preserving well-foundedness. Together, these constructions induce a large class of BMA rankings and corresponding fully-rational contractions.

(1) Lists of Ranks. The most straightforward way to ensure **(Well-Founded)** is to consider finite rankings. Let therefore $M_1, \ldots, M_n$ be a partition of the set Σ^ω (i.e., $\bigcup_{i=1}^n M_i = \Sigma^\omega$ and $M_i \cap M_j = \emptyset$ for $i \neq j$). Assume that the traces in M_1 are considered less plausible than those in M_2, those in $M_1 \cup M_2$ are less plausible than those in M_3, etc. The corresponding ranking function f maps traces to the well-ordered set $(\{1, \ldots, n\}, <_\mathbb{N})$, and satisfies $f(\tau) = i$ if and only if $\tau \in M_i$. If each rank M_i is described as a Büchi automaton, the induced ranking can be represented as a BMA.

Proposition 33. *Let $A_1, \ldots, A_n$ be Büchi automata with disjoint languages, such that $\bigcup_{i=1}^n \mathcal{L}(A_i) = \Sigma^\omega$. The ranking function f with $f(\tau) = i$ if and only if $\tau \in \mathcal{L}(A_i)$ induces a well-founded BMA ranking.*

It is natural to define the preference relation in terms of the formulae satisfied by the traces. In particular, we can choose automata A_i corresponding to LTL formulae.

Example 34. *Let the atomic proposition r denote that* "it rains". *We may consider it more plausible that* "it will rain tomorrow" *(in LTL, $\varphi :\equiv \mathbf{X}\, r$) than that* "it will not rain tomorrow" *($\neg\varphi$). This preference corresponds to a BMA relation, as we can see by setting $A_1 = A_{\neg\varphi}$, $A_2 = A_\varphi$ and $n = 2$ in Proposition 33.*

LTL also allows us to express more fine-grained temporal preferences. We can express the attitude that most plausibly, "it will rain today", *somewhat less plausibly* "it will rain tomorrow (but not today)", *and even less plausibly* "it will eventually rain (but not today or tomorrow)". *The alternative* "it will never rain again" *is most implausible. To express this, we set n to 4 and $A_i := A_{\varphi_i}$ (for $i = 1, \ldots, n$), with*

$$\varphi_1 \equiv \mathbf{G}\, \neg r \qquad \varphi_2 \equiv \mathbf{F}\, r \wedge (\neg \mathbf{X}\, r) \wedge \neg r \qquad \varphi_3 \equiv (\mathbf{X}\, r) \wedge \neg r \qquad \varphi_4 \equiv r$$

The BMA recognizing the corresponding ranking is shown in Fig. 2 and discussed in Example 14.

While often useful, this approach allows us to distinguish only finitely many ranks of equi-preferable traces. Consider again the example above, but with the preference that *"it is plausible that it rains* soon*"*. Or more precisely, *"the sooner it rains (according to some trace), the more plausible"*. This requires us to distinguish infinitely many ranks (*"it first rains in n days"*, for each $n \in \mathbb{N}$). As we show next, BMA are expressive enough to represent such preferences.

(2) The sooner, the better. A rich class of infinite rankings can be described by ordering traces according to the earliest time point from which on some plausible LTL formula φ holds. I.e., we prefer traces satisfying φ over traces satisfying only $\mathbf{X}\, \varphi$, over traces satisfying only $\mathbf{X}^2\, \varphi$, etc. Given a trace τ, we write $\mathsf{fst}_\varphi(\tau) = i$ for the minimal $i \in \mathbb{N}$ such that $\tau \models \mathbf{X}^i\, \varphi$, and $\mathsf{fst}_\varphi(\tau) = \infty$ if no such i exists.

Proposition 35. *The ranking function $f : \Sigma^\omega \to \mathbb{Z}_{\leq 0} \cup \{-\infty\}, \tau \mapsto -\mathsf{fst}_\varphi(\tau)$ induces a well-founded BMA ranking.*

The induced BMA ranking is denoted as $\prec_\varphi$. We invert the sign, as traces where the desired property holds at a *smaller* time index are considered *more* preferable. Well-foundedness arises from the fact that there does not exist an infinite descending chain of trace indices $i \in \mathbb{N}$.

Example 36. *The BMA ranking $\prec_\varphi$, for the LTL formula $\varphi \equiv r$, expresses the preference* "it is plausible that it rains soon" *discussed above. Our result also applies to more complex temporal statements. For instance,* "it is plausible that soon, it will never snow again" *is expressed by the BMA ranking $\prec_\psi$, for the LTL formula $\psi \equiv \mathbf{G}\, \neg s$, where the proposition s indicates that* "it snows" *(see Fig. 3).*

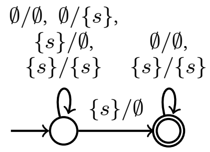

Figure 3: BMA for the ranking $\prec_{\mathbf{G}\,\neg s}$.

(3) Rankings from Distance Measures. A well-known approach to construct epistemic preference relations in the classical setting is based on defining a *distance* between models (propositional valuations). The most preferred models are those with the least distance to the models of the knowledge base $\mathcal{K}$ from which a formula is contracted. I.e., models are *ranked* by their distance from $\mathcal{K}$, and the closest models are most preferred. The most prominent representative of this approach is the distance defined by Dalal [31, 32]. Many other distance measures fail to ensure rationality [33] in the classical setting.

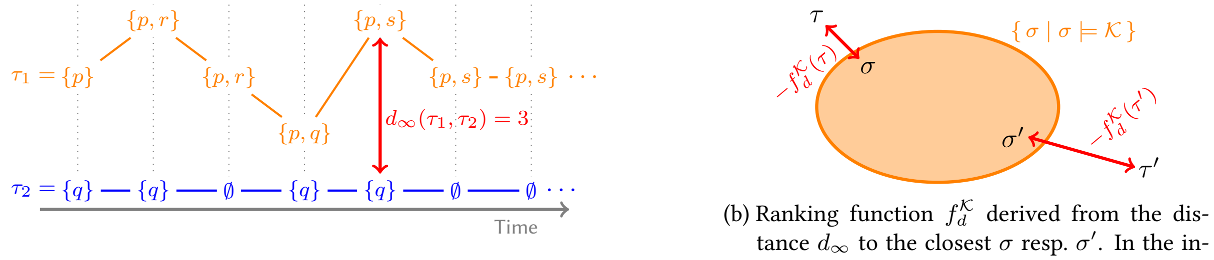

(a) Trace distance d_∞, where d is the Dalal distance.

(b) Ranking function $f_d^{\mathcal{K}}$ derived from the distance d_∞ to the closest σ resp. σ'. In the induced ranking, τ is preferred over τ'.

Figure 4: Construction of a ranking from a notion of distance d on propositional valuations.

We define a well-founded BMA ranking based on a notion of distance between the traces and the given knowledge base $\mathcal{K}$. The construction is parametric in a distance function $d : \Sigma \times \Sigma \to \mathbb{R}_{\geq 0} \cup \{+\infty\}$ on propositional valuations in $\Sigma = \mathcal{P}(AP)$. It may be useful to think of d as a distance in the mathematical sense that it satisfies the *metric axioms* (symmetry, triangle inequality, and $d(x,y) = 0$ iff $x = y$). For instance, the distance of Dalal [31, 32] (also known as *Hamming distance*) satisfies these axioms. Yet, neither rationality nor the constructive realization with BMA depend on such assumptions, thus allowing for the use of many broader notions of distance that fail rationality in the classical usage.

Observe that the distance $d(x,y)$ can only take finitely many values, as $\Sigma \times \Sigma$ is finite; let the range of d be $\{r_0, r_1, \dots, r_n\}$ with $r_0 < r_1 < \dots < r_n$. The function d gives rise to a distance on traces:

$$d_\infty(A_0A_1\dots, B_0B_1\dots) := \max_{i\in\mathbb{N}} d(A_i, B_i)$$

Figure 4a illustrates this distance measure on traces. If d is a metric, the distance d_∞ is the corresponding *Chebyshev metric*. For a consistent theory $\mathcal{K}$, we define the rank $f_d^{\mathcal{K}}(\tau)$ of a trace τ via the distance between τ and the closest trace σ such that $\sigma \models \mathcal{K}$ holds:

$$f_d^{\mathcal{K}}(\tau) := -\min\{\, d_\infty(\tau,\sigma) \mid \sigma \models \mathcal{K} \,\}$$

We invert the sign such that traces τ close to $\mathcal{K}$ are preferred. The minima and maxima in these equations are well-defined, as they operate on nonempty subsets of the finite set $\{r_0, \dots, r_n\}$. Figure 4b illustrates the ranking function. And indeed, this ranking (denoted $\prec_d^{\mathcal{K}}$) can be represented as BMA:

Proposition 37. *If $\mathcal{K} \in \mathbb{E}_{\text{Büchi}}$ is consistent, the ranking function $f_d^{\mathcal{K}}$ induces a well-founded BMA ranking.*

(4) Hierarchical Preferences. A powerful tool when describing (ranked) preferences is the ability to express *hierarchical* preferences. For instance, when discussing the weather, we might prefer (resp. consider as more plausible) traces in which *"soon, it will never snow again"*, but among two traces that agree on the last snow day, we might prefer the trace that is closer to our belief base $\mathcal{K}$.

Given ranking functions $f_i : \Sigma^\omega \to Y_i$ over totally-ordered sets $(Y_i, <_i)$, for $i \in \{1,2\}$, a hierarchical preference can be expressed as the ranking function $f_{12} : \Sigma^\omega \to Y_1 \times Y_2$ with $f(\tau) = (f_1(\tau), f_2(\tau))$, where $Y_1 \times Y_2$ is totally ordered by the lexicographical combination of $<_1, <_2$, i.e., $(y_1, y_2) <_{12} (y_1', y_2')$ if $y_1 <_1 y_1'$, or $y_1 = y_1'$ and $y_2 < y_2'$. BMA rankings are closed under this hierarchical composition.

Proposition 38. *Let $\prec_1, \prec_2$ be BMA rankings induced by ranking functions f_1 resp. f_2. The ranking function f_{12} induces a BMA ranking $\prec_{12}$. If $\prec_1$ and $\prec_2$ are well-founded, so is the ranking $\prec_{12}$.*

We denote the composed ranking as $\prec_1 \triangleright \prec_2$. In combination with the other constructions presented above, hierarchical preference composition allow us to express complex preference relations.

Example 39. *Consider the BMA ranking $\prec_\psi$, for $\psi = \mathbf{G}\,\neg s$ (*"it is plausible that soon, it will never snow again"*), and the BMA ranking $\prec_{Dalal}^{\mathcal{K}}$ induced by the Dalal distance measure. By composing the two rankings, we arrive at an epistemic preference relation $\prec_\psi \triangleright \prec_{Dalal}^{\mathcal{K}}$ that prefers traces in which the last snow day occurs sooner rather than later; and among traces that agree on the last snow day, it prefers traces that are closer to $\mathcal{K}$ (as defined by the Dalal-Chebyshev distance). Through this composition, each of the infinitely many ranks of $\prec_\psi$ is sub-divided into a finite number of sub-ranks based on the Dalal distance.*

6. Conclusion

Epistemic preference relations form the backbone of rational belief change. We have studied the particular case of effective AGM contraction in an expressive, non-finitary logic, namely linear temporal logic (LTL). In this setting, preference relations are represented via so-called *Büchi-Mealy automata* and must satisfy two core properties to ensure rationality of the corresponding contraction. We have shown that while one of the properties, **(Mirroring)**, is decidable, the key property ensuring success of the contraction, **(Maximal Cut)**, is undecidable. This has far-reaching consequences, as it rules out the possibility of a representation result for effective LTL contraction. Consequently, effective LTL contraction requires custom design of suitable preference relations for particular applications.

We have proposed a number of constructions for preference relations that guarantee the required properties *by design*; based on the observation that such relations must necessarily be *rankings*, potentially with an infinite number of ranks. One of our constructions generalizes an idea from classical belief change, namely the construction of preferences based on the distance to the models of the current epistemic state. Constructing preference relations for classical propositional logics is also a challenge. Although some approaches were proposed [31, 34], very few achieve full rationality, being Dalal and some *culpability measures* [33] the few known ones. Dalal can only be used on consistent beliefs, while *culpability measures* work only for inconsistent beliefs. Our approaches easily apply for propositional logics, and provide novel and richer alternatives to obtain preference relations.

The attentive reader may have observed that we have not shown the undecidability of **(Maximal Cut)** *assuming that* **(Mirroring)** *holds*, or equivalently, the undecidability of the conjunction of the two properties. We leave the full resolution of this question to future work. However, let us remark that there are strong indications that undecidability indeed persists under such a restriction. In particular, our proof of undecidability for **(Maximal Cut)** relies on a reduction from a variation of the uniform halting problem, i.e., the question whether a given Turing machine *always terminates*. In the study of *program verification*, a link has long been made between termination and the existence of a (well-founded) *ranking function* (sometimes also called *variant* or *termination measure*). We suspect that closing the loop to our previous observation that relations satisfying both **(Mirroring)** and **(Maximal Cut)** must indeed be rankings would allow for such a proof (though the technical details remain challenging).

Declaration on Generative AI

The author(s) have not employed any Generative AI tools.

References

[1] P. Gärdenfors, Knowledge in flux: Modeling the dynamics of epistemic states., The MIT press, 1988.

[2] S. O. Hansson, A textbook of belief dynamics - theory change and database updating, volume 11 of *Applied logic series*, Kluwer, 1999.

[3] A. Pnueli, The temporal logic of programs, in: FOCS, IEEE Computer Society, 1977, pp. 46–57.

[4] A. Darwiche, J. Pearl, On the logic of iterated belief revision, Artif. Intell. 89 (1997) 1–29.

[5] N. Schwind, S. Konieczny, R. P. Pérez, On the representation of Darwiche and Pearl's epistemic states for iterated belief revision, in: KR, 2022, pp. 320–330.

[6] J. P. Delgrande, P. Peppas, S. Woltran, General belief revision, J. ACM 65 (2018) 29:1–29:34.

[7] R. Booth, T. Meyer, I. Varzinczak, R. Wassermann, On the link between partial meet, kernel, and infra contraction and its application to Horn logic, CoRR abs/1401.3902 (2014). `arXiv:1401.3902`.

[8] F. Baader, I. Horrocks, C. Lutz, U. Sattler, An Introduction to Description Logic, Cambridge University Press, 2017.

[9] F. Baader, R. Wassermann, Contractions based on optimal repairs, in: KR, 2024, pp. 94–105.

[10] E. M. Clarke, E. A. Emerson, Design and synthesis of synchronization skeletons using branching time temporal logic, in: 25 Years of Model Checking, volume 5000 of *Lecture Notes in Computer Science*, Springer, 2008, pp. 196–215.
[11] E. M. Clarke, O. Grumberg, D. Kroening, D. A. Peled, H. Veith, Model checking, 2nd Edition, MIT Press, 2018.
[12] J. S. Ribeiro, A. Nayak, R. Wassermann, Towards belief contraction without compactness, in: KR, AAAI Press, 2018, pp. 287–296.
[13] J. S. Ribeiro, A. Nayak, R. Wassermann, Belief change and non-monotonic reasoning sans compactness, in: AAAI, AAAI Press, 2019, pp. 3019–3026.
[14] D. Klumpp, J. S. Ribeiro, Effective AGM belief contraction: A journey beyond the finitary realm, in: KR, 2025, pp. 429–439.
[15] B. Nebel, How hard is it to revise a belief base?, in: D. Dubois, H. Prade (Eds.), Belief Change, Springer Netherlands, 1998, pp. 77–145.
[16] T. Eiter, G. Gottlob, On the complexity of propositional knowledge base revision, updates, and counterfactuals, Artif. Intell. 57 (1992) 227–270.
[17] N. Schwind, S. Konieczny, J. Lagniez, P. Marquis, On computational aspects of iterated belief change, in: IJCAI, ijcai.org, 2020, pp. 1770–1776.
[18] K. Sauerwald, C. Beierle, Iterated belief change, computationally, in: KR, 2022, pp. 410–414.
[19] P. T. Guerra, R. Wassermann, On the uncomputability of partial meet contraction for linear-time temporal logic, South American Journal of Logic 4 (2018) 385–406.
[20] M. Souza, On AGM belief revision for computational tree logic, in: BRACIS (1), volume 13653 of *Lecture Notes in Computer Science*, Springer, 2022, pp. 413–428.
[21] H. Gaißer, D. Klumpp, J. S. Ribeiro, Constructive preference relations: Navigating undecidability in rational LTL contraction, CoRR abs/2606.16957 (2026). arXiv:2606.16957.
[22] G. D. Giacomo, M. Y. Vardi, Automata-theoretic approach to planning for temporally extended goals, in: ECP, Lecture Notes in Computer Science, Springer, 1999, pp. 226–238.
[23] F. Patrizi, N. Lipovetzky, G. D. Giacomo, H. Geffner, Computing infinite plans for LTL goals using a classical planner, in: IJCAI, IJCAI/AAAI, 2011, pp. 2003–2008.
[24] J. Richard Büchi, Symposium on Decision Problems: On a Decision Method in Restricted Second Order Arithmetic, Studies in Logic and the Foundations of Mathematics, Elsevier, 1966, pp. 1–11.
[25] C. E. Alchourrón, P. Gärdenfors, D. Makinson, On the logic of theory change: Partial meet contraction and revision functions, J. Symb. Log. 50 (1985) 510–530.
[26] S. Schewe, Büchi complementation made tight, in: STACS, LIPIcs, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, Germany, 2009, pp. 661–672.
[27] P. Wolper, Temporal logic can be more expressive, Inf. Control. 56 (1983) 72–99.
[28] K. R. Chowdhary, Linear-Bounded Automata and Context-Sensitive Languages, Springer Nature Singapore, 2025, pp. 431–468.
[29] P. K. Hooper, The undecidability of the Turing machine immortality problem, J. Symb. Log. 31 (1966) 219–234.
[30] J. S. R. Santos, Belief Change without Compactness, Ph.D. thesis, Macquarie University, 2019.
[31] M. Dalal, Investigations into a theory of knowledge base revision, in: AAAI, AAAI Press / The MIT Press, 1988, pp. 475–479.
[32] M. Dalal, Updates in propositional databases, Tech. Report DCS-TR-222, Department of Computer Science, Rutgers University, New Brunswick, NJ, 1988.
[33] J. S. Ribeiro, M. Thimm, Consolidation via tacit culpability measures: Between explicit and implicit degrees of culpability, in: KR, 2021, pp. 529–538.
[34] H. Katsuno, A. O. Mendelzon, Propositional knowledge base revision and minimal change, Artif. Intell. 52 (1992) 263–294.