Notre Dame Journal of Formal Logic
Volume 37, Number 2, Spring 1996

Minimal Temporal Epistemic Logic

JOERI ENGELFRIET

Abstract In the study of nonmonotonic reasoning the main emphasis has been on static (declarative) aspects. Only recently has there been interest in the dynamic aspects of reasoning processes, particularly in artificial intelligence. We study the dynamics of reasoning processes by using a temporal logic to specify them and to reason about their properties, just as is common in theoretical computer science. This logic is composed of a base temporal epistemic logic with a preference relation on models, and an associated nonmonotonic inference relation, in the style of Shoham, to account for the nonmonotonicity. We present an axiomatic proof system for the base logic and study decidability and complexity for both the base logic and the nonmonotonic inference relation based on it. Then we look at an interesting class of formulas, prove a representation result for it, and provide a link with the rule of monotonicity.

1 Introduction In theoretical computer science, temporal logic has been widely recognized as a valuable tool for specifying processes and reasoning about their properties. In the study of nonmonotonic reasoning the temporal view is not very common, partly because (nonmonotonic) logic is usually thought of as a purely static notion. However, in nonmonotonic reasoning dynamic aspects of reasoning processes can be interesting to study and often influence the static aspects, just as is common in computer science, where we often have declarative semantics next to procedural semantics of processes. There are also differences between the notion of process in computer science and a reasoning process, for instance in the nature of a state: in a computer it is composed of the values of the variables, whereas in a reasoning process it consists of the facts which are believed (or derived) at that time.

A number of examples in which a temporal logic is used to specify reasoning processes can be found in Engelfriet and Treur [6], where such specifications are introduced for default logic (see Reiter [18]), classical inference systems, and meta-level architectures. Also, in Engelfriet and Treur [7] it is shown that there exists a large class of reasoning processes that can be specified in this temporal logic. Therefore it seems justified to study this temporal logic formalism in more detail, which will be done in the present paper.

In Section 2 we introduce the temporal logic which is the basis of the framework, and in Section 3 an extra restriction is imposed upon this logic. Section 4 describes the notions of minimal models and minimal entailment which will be studied in the rest of the paper. In Section 5 decidability of this notion is established, and Section 6 gives complexity results for both the base logic and minimal entailment. In Section 7 we look at a special class of formulas and prove a link with the rule of monotonicity. Section 8 gives conclusions and suggestions for further research.

2 Temporal epistemic logic When designing a logic capable of describing the behavior of reasoning processes over time, two important decisions have to be made: which temporal ontology is suited best for the purpose, and what is a state in a reasoning process? We view a reasoning process, performed by an agent for instance, as a stepwise process: the agent starts out with some initial facts (possibly none) and attempts to derive consequences by applying rules; a new state in which the agent has more knowledge results. The agent will then try again to derive new facts resulting in a next state, et cetera, possibly *ad infinitum*. This suggests a temporal ontology which is discrete and has a starting point (the natural numbers seem most suited). In theoretical computer science there has been much debate about whether time should be linear or branching (towards the future) (see de Bakker, de Roever, and Rozenberg [3]). The most important differences between these two approaches are that linear time logics have in general a lower complexity but also less expressivity than the corresponding branching time logics. Although some results in [6] on specifying proof systems in temporal logic seem to suggest that sometimes the higher expressivity of branching time logic is needed, we will confine ourselves here to using linear time.

As suggested above, the important thing about the state of a reasoning agent at a particular moment is the knowledge he has derived. Kripke semantics can be used to formalize such an *information state*. We will take propositional logic as the basic language in which the agent can describe his knowledge. A modal operator K will be used to denote the agent's knowledge. In principle the agent may perform (positive and negative) introspection, which suggests an **S5** logic to describe knowledge.

Definition 2.1 (Epistemic language) Let P be a (finite or countably infinite) set of propositional atoms. The language $\mathcal{L}_{\mathbf{S5}}$ is the smallest set closed under:

1. if $p \in P$ then $p \in \mathcal{L}_{\mathbf{S5}}$;
2. if $\varphi, \psi \in \mathcal{L}_{\mathbf{S5}}$ then $\mathsf{K}\varphi,\ \varphi \wedge \psi,\ \neg\varphi \in \mathcal{L}_{\mathbf{S5}}$.

Furthermore, we introduce the following abbreviations:

$$\varphi \vee \psi \equiv \neg(\neg\varphi \wedge \neg\psi),$$
$$\varphi \to \psi \equiv \neg\varphi \vee \psi,$$
$$\mathsf{M}\varphi \equiv \neg\mathsf{K}\neg\varphi,$$
$$\top \equiv p \vee \neg p,$$
$$\bot \equiv \neg\top.$$

If every atom occurring in a formula φ is in the scope of a K operator, we call φ *subjective*.

An example of a subjective formula is $\neg\mathsf{K}p \wedge \mathsf{K}(q \to p)$, whereas $\mathsf{K}(p \wedge q) \vee s$ is not subjective. In the rest of this paper we will be especially interested in subjective formulas since they describe (only) the knowledge and ignorance of the agent. As we want to talk about the knowledge of the agent changing over time, the epistemic language will be temporalized below.

In the usual **S5** semantics a model is a triple (W, R, π) where W is a set of worlds, R is an equivalence relation on W, and π is a function that assigns a propositional valuation to each world in W. We may however (see e.g. Meyer and van der Hoek [16]), in the case of one agent, restrict ourselves to *normal* **S5**-models, in which the relation is universal (each world is accessible from every other world) and worlds are identified with propositional valuations.

Definition 2.2 (**S5** semantics) A propositional valuation of signature P is a function from P into $\{0, 1\}$, where 0 stands for false and 1 for true. The set of such valuations will be denoted by $\mathrm{Mod}(P)$. A normal **S5**-model M is a nonempty set of valuations. The truth of an **S5**-formula φ in such a model, evaluated in a world $m \in M$, denoted $(M, m) \models_{\mathbf{S5}} \varphi$, is defined inductively:

$$(M, m) \models_{\mathbf{S5}} p \iff m(p) = 1, \text{ for } p \in P \tag{1}$$

$$(M, m) \models_{\mathbf{S5}} \varphi \wedge \psi \iff (M, m) \models_{\mathbf{S5}} \varphi \text{ and } (M, m) \models_{\mathbf{S5}} \psi \tag{2}$$

$$(M, m) \models_{\mathbf{S5}} \neg\varphi \iff \text{it is not the case that } (M, m) \models_{\mathbf{S5}} \varphi \tag{3}$$

$$(M, m) \models_{\mathbf{S5}} \mathsf{K}\varphi \iff (M, m') \models_{\mathbf{S5}} \varphi \text{ for every } m' \in M. \tag{4}$$

A pair (M, m) where M is a normal **S5**-model and $m \in M$ (the *current world*) is called an *epistemic state*, and the set of such pairs is denoted by ES(P), or simply ES.

It is easy to see that the truth of a subjective **S5**-formula in a model is independent of the world in which it is evaluated, so if we restrict ourselves to subjective formulas, the world m in which it is evaluated is often left out.

Remark 2.3 Note that an **S5**-formula is subjective if and only if it is equivalent to a formula of the form $\mathsf{K}\varphi$ with $\varphi \in \mathcal{L}_{\mathbf{S5}}$.

Axiomatizations for **S5** are known from the literature (e.g. Halpern and Moses [13]).

Definition 2.4 (Axiom system for **S5**) The axiom system of **S5** consists of:

1. All instances of propositional tautologies
2. $\mathsf{K}(\varphi \to \psi) \to (\mathsf{K}\varphi \to \mathsf{K}\psi)$ (K)
3. $\mathsf{K}\varphi \to \varphi$ (T)
4. $\mathsf{K}\varphi \to \mathsf{K}\mathsf{K}\varphi$ (*Positive Introspection*)
5. $\neg\mathsf{K}\varphi \to \mathsf{K}\neg\mathsf{K}\varphi$ (*Negative Introspection*)

and the following two rules:

1. $\dfrac{\varphi \quad \varphi \to \psi}{\psi}$ (*Modus Ponens*)
2. $\dfrac{\varphi}{\mathsf{K}\varphi}$ (*Necessitation*)

If there is a proof for φ using this system, we will denote this by $\vdash_{\mathbf{S5}} \varphi$.

It is well known that this system is sound and complete with respect to the class of normal **S5**-models.

In order to describe past and future we will introduce temporal operators P, H, F, G, and $\Box$, standing for “sometimes in the past,” “always in the past,” “sometimes in the future," "always in the future," and "always" respectively. Note that we do not want to talk about the agent's knowledge of the future and past, but about the future and past of the agent's knowledge. Therefore temporal operators need never occur within the scope of the epistemic K operator. This is reflected in the definition of the temporal epistemic language.

Definition 2.5 (Temporal epistemic language) The language $\mathcal{L}_{\mathbf{TEL}}$ is the smallest set closed under:

1. if $\varphi \in \mathcal{L}_{\mathbf{S5}}$ then $\varphi \in \mathcal{L}_{\mathbf{TEL}}$;
2. if $\alpha, \beta \in \mathcal{L}_{\mathbf{TEL}}$ then $\alpha \wedge \beta, \neg\alpha, P\alpha, F\alpha \in \mathcal{L}_{\mathbf{TEL}}$.

Again the abbreviations for $\vee$, $\rightarrow$, $\top$, and $\perp$ are introduced, as well as:

$G\alpha \equiv \neg F \neg \alpha$,
$H\alpha \equiv \neg P \neg \alpha$, and
$\Box\alpha \equiv H\alpha \wedge \alpha \wedge G\alpha$.

If in the first clause we restrict ourselves to subjective **S5**-formulas, we get the set of *subjective* **TEL**-*formulas*.

In the rest of this paper we will be interested in subjective **TEL**-formulas since they describe how the knowledge of the agent changes over time. Based on the set of natural numbers (starting at zero) as flow of time and the notion of epistemic state as formalization of a state, the following semantics is introduced for temporal epistemic logic (**TEL**).

Definition 2.6 (Semantics of **TEL**) A **TEL**-*model* is a function $\mathcal{M} : \mathbb{N} \to \mathrm{ES}$. The truth of a formula $\varphi \in \mathcal{L}_{\mathbf{TEL}}$ in $\mathcal{M}$ at time point $t \in \mathbb{N}$, denoted $(\mathcal{M}, t) \models \varphi$, is defined inductively as follows:

$$(\mathcal{M}, t) \models \varphi \iff \mathcal{M}(t) \models_{\mathbf{S5}} \varphi, \text{ if } \varphi \in \mathcal{L}_{\mathbf{S5}} \tag{1}$$

$$(\mathcal{M}, t) \models \varphi \wedge \psi \iff (\mathcal{M}, t) \models \varphi \text{ and } (\mathcal{M}, t) \models \psi \tag{2}$$

$$(\mathcal{M}, t) \models \neg\varphi \iff \text{it is not the case that } (\mathcal{M}, t) \models \varphi \tag{3}$$

$$(\mathcal{M}, t) \models P\varphi \iff \exists s \in \mathbb{N} \text{ such that } s < t \text{ and } (\mathcal{M}, s) \models \varphi \tag{4}$$

$$(\mathcal{M}, t) \models F\varphi \iff \exists s \in \mathbb{N} \text{ such that } t < s \text{ and } (\mathcal{M}, s) \models \varphi \tag{5}$$

A formula φ is true in a model $\mathcal{M}$, denoted $\mathcal{M} \models \varphi$, if for all $t \in \mathbb{N}$, $(\mathcal{M}, t) \models \varphi$. If φ is true in all models we write $\models \varphi$ (φ is *valid*), and we write $\psi \models \varphi$ (φ is a *semantical consequence* of ψ) if for all models $\mathcal{M}$ and $t \in \mathbb{N}$, $(\mathcal{M}, t) \models \psi$ implies $(\mathcal{M}, t) \models \varphi$. We will often write $\mathcal{M}_t$ for $\mathcal{M}(t)$.

For future use we give the following definition. (Here O^i stands for a sequence of O operators of length i, where $O \in \{P, H, F, G, \Box\}$. Furthermore $O^0\alpha$ stands for α.)

Definition 2.7 For $i \in \mathbb{N}$ define $at_i := P^i\top \wedge H^{i+1}\perp$.

It is easy to see that $(\mathcal{M}, j) \models at_i$ if and only if $i = j$.

We would like to find an axiom system for **TEL**. The idea is to use the axioms of an **S5**-system together with axioms for tense logic over the natural numbers. Instead of proving soundness and completeness for the resulting system from scratch, we will use results from Finger and Gabbay [8] where a general method for temporalizing a given logic system is presented. In their notation, **TEL** would be T(**S5**). We cannot directly apply their results since they use the temporal operators Since and Until, but adaptation to our situation is easy. Our class of flows of time contains only the set of natural numbers. First we will give an axiomatic system for propositional tense logic over the natural numbers (from Goldblatt [10]), which is sound and complete with respect to $\mathbb{N}$.

Definition 2.8 (Tense logic over the natural numbers) The axiom system for tense logic over $\mathbb{N}$ consists of:

1. All instances of propositional tautologies
2. $G(\varphi \to \psi) \to (G\varphi \to G\psi)$
3. $H(\varphi \to \psi) \to (H\varphi \to H\psi)$
4. $\varphi \to HF\varphi$ (C_P)
5. $\varphi \to GP\varphi$ (C_F)
6. $H\varphi \to HH\varphi$ (4_P)
7. $G\varphi \to GG\varphi$ (4_F)
8. $F(\top)$ (D_F)
9. $G(G\varphi \to \varphi) \to (FG\varphi \to G\varphi)$ (Z_F)
10. $H(H\varphi \to \varphi) \to H\varphi$ (W_P)

and the following rules:

1. $\dfrac{\varphi \quad \varphi \to \psi}{\psi}$ (*Modus Ponens*)
2. $\dfrac{\varphi}{G\varphi} \qquad \dfrac{\varphi}{H\varphi}$ (*Necessitation*)

Using the axiom systems for **S5** and tense logic, Definition 2.6 of [8] allows us to give an axiomatization for **TEL**.

Definition 2.9 (Axiomatization for **TEL**) The axiom system of **TEL** consists of:

1. The axioms 1–10 of Definition 2.8
2. The inference rules 1 and 2 of Definition 2.8
3. For every formula $\alpha \in \mathcal{L}_{\mathbf{S5}}$, if $\vdash_{\mathbf{S5}} \alpha$ then $\vdash_{\mathbf{TEL}} \alpha$ *(Preserve)*.

Using Theorem 2.2 of [8], soundness of **S5** and soundness of the axiom system for tense logic over $\mathbb{N}$, we immediately have the following theorem.

Theorem 2.10 (Soundness of **TEL**) *The axiom system* **TEL** *is sound.*

Theorem 2.3 of [8] states that if the system to be temporalized is complete and the axiomatization of the logic with Since and Until is complete over a class of *linear* flows of time, then the "merged" axiomatization is complete for the temporalized logic. Our class of flows of time (consisting only of the natural numbers) is a subclass of the linear flows of time. A slight adaptation of their proof yields the same result for temporalizing over the temporal operators used in **TEL**. Therefore we have the following.

Theorem 2.11 (Completeness of **TEL**) *The axiom system* **TEL** *is complete.*

Again borrowing from [8], Theorem 3.1, and using the fact that both **S5** ([16]) and tense logic over the natural numbers (Sistla and Clarke [19]) are decidable, we have the following theorem.

Theorem 2.12 (Decidability of TEL) *The logic* TEL *is decidable.*

In the next section we will impose an extra restriction on our models.

3 Conservativity We want to use subjective temporal formulas for describing the behavior of a reasoning agent. The reasoning will be assumed to be *conservative*, that is, the agent's knowledge will increase as he is reasoning. Although the actual implementation of the reasoning behavior may involve backtracking or the addition of extra assumptions which may later be retracted, we are interested only in the increase of knowledge over time: adding assumptions and later retracting them is assumed to be done in one step. This presupposes a world which does not change. We will restrict ourselves to conservative behavior here, though we agree that it may be worthwhile to investigate nonconservative behavior as well.

In the following we are interested only in subjective formulas, so we delete the world from the epistemic state. Thus in the following, we consider ES to be the set of all normal **S5**-models, i.e., the powerset of $\mathrm{Mod}(P)$ without the empty set. We will study consequence relations between formulas, and it will turn out that these notions are independent of the propositional signature. Therefore the propositional signature can and will be assumed to be finite.

Definition 3.1 (Conservative models)

1. We define the *degree-of-information ordering* $\leq$ on information states as follows:
$$\text{for } M_1, M_2 \in \mathrm{ES}, M_1 \leq M_2 \iff M_2 \subseteq M_1$$
We write $M_1 < M_2$ if $M_1 \leq M_2$ and $M_1 \neq M_2$.
2. A TEL-model $\mathcal{M}$ is called *conservative* if for all $s \in \mathbb{N}$:
$$\mathcal{M}_s \leq \mathcal{M}_{s+1}$$
3. Validity and semantical consequence restricted to the class of conservative models (TELC-*models*) will be denoted by $\models^c$.

The definition of the degree-of-information ordering is based on the observation that the more valuations one considers to be possible, the less knowledge (or information) one has. Note that for any conservative model $\mathcal{M}$, time point $s \in \mathbb{N}$, and propositional formula φ: if $(\mathcal{M}, s) \models \mathsf{K}\varphi$, then for $t > s$ also $(\mathcal{M}, t) \models \mathsf{K}\varphi$. This means that whenever a propositional formula is known, it will remain known in the future.

The notions $\models$ and $\models^c$ are not compact: the set $\{P^i(\top) | i \in \mathbb{N}\}$ is not satisfiable, whereas each finite subset is (for both notions).

Proposition 3.2 (Axiomatization) *Let* $C = \{\Box(\mathsf{K}\alpha \to G(\mathsf{K}\alpha)) | \alpha$ *a propositional formula*$\}$. *For each* TEL-*model* $\mathcal{M}$ *the following are equivalent:*

1. $\mathcal{M}$ *is conservative*
2. $\mathcal{M} \models C$

3. $(\mathcal{M}, t) \models C$ *for some* $t \in \mathbb{N}$.

Furthermore, the axiom system **TELC***, consisting of* **TEL** *plus the axioms of* C*, is sound and complete with respect to the class of* **TELC***-models.*

Proof: Let $\mathcal{M}$ be conservative and let $t \in \mathbb{N}$. Suppose $(\mathcal{M}, t) \models \mathsf{K}\alpha$ and take $s > t$ arbitrary. Then for all $m \in \mathcal{M}_t$, $m \models \alpha$. Take $m \in \mathcal{M}_s$, then since $\mathcal{M}$ is conservative we have $\mathcal{M}_s \subseteq \mathcal{M}_t$, so $m \in \mathcal{M}_t$ and $m \models \alpha$. Therefore $(\mathcal{M}, s) \models \mathsf{K}\alpha$, and since s was arbitrary we have $(\mathcal{M}, t) \models G(\mathsf{K}\alpha)$, so $(\mathcal{M}, t) \models \mathsf{K}\alpha \to G(\mathsf{K}\alpha)$. We have $(\mathcal{M}, 0) \models \Box(\mathsf{K}\alpha \to G(\mathsf{K}\alpha))$.

Suppose on the other hand that $(\mathcal{M}, t) \models C$ for some $t \in \mathbb{N}$, but $\mathcal{M}$ is not conservative. Then there exists $s \in \mathbb{N}$ and $m \in \mathcal{M}_{s+1}$ with $m \notin \mathcal{M}_s$. Let φ_m be the conjunction of the literals that are true in m (i.e., $\varphi_m = \bigwedge\{p \in P | m \models p\} \wedge \bigwedge\{\neg p | p \in P, m \not\models p\}$; this is a propositional formula since P was assumed finite in the remark above Definition 3.1). Then since $m \notin \mathcal{M}_s$ and for all $m' \neq m$, $m' \models \neg\varphi_m$, we have $(\mathcal{M}, s) \models \mathsf{K}(\neg\varphi_m)$, but as $m \in \mathcal{M}_{s+1}$ and $m \not\models \neg\varphi_m$, $(\mathcal{M}, s+1) \not\models \mathsf{K}(\neg\varphi_m)$, so $(\mathcal{M}, s) \not\models G(\mathsf{K}(\neg\varphi_m))$. Thus $(\mathcal{M}, t) \not\models \Box(\mathsf{K}(\neg\varphi_m) \to G(\mathsf{K}(\neg\varphi_m)))$, a contradiction.

The above shows that the axioms of C are sound. Now suppose $\models^c \varphi$, then we have for all **TEL**-models $\mathcal{M}$: if $\mathcal{M}$ is conservative then $\mathcal{M} \models \varphi$. Since there are only a finite number of nonequivalent propositional formulas over P, C can be taken to be finite, and therefore we can take the conjunction of its elements. So if $(\mathcal{M}, s) \models \bigwedge C$ then $\mathcal{M}$ is conservative, so $\mathcal{M} \models \varphi$, and therefore $(\mathcal{M}, s) \models \varphi$. Thus we have $\bigwedge C \models \varphi$, and using the deduction lemma for **TEL** (which can be easily verified), $\models \bigwedge C \to \varphi$, from which by the completeness of **TEL** it follows that $\vdash_{\mathbf{TEL}} \bigwedge C \to \varphi$. Since **TELC** contains **TEL** and the axioms of C and has Modus Ponens as inference rule, we conclude $\vdash_{\mathbf{TELC}} \varphi$. □

We also have that **TELC** is decidable.

Proposition 3.3 (Decidability of **TELC**) *The logic* **TELC** *is decidable.*

Proof: Checking whether $\vdash_{\mathbf{TELC}} \varphi$ reduces to checking $\vdash_{\mathbf{TEL}} \bigwedge C \to \varphi$, where C is the set of rules $\Box(\mathsf{K}\alpha \to G(\mathsf{K}\alpha))$ for all nonequivalent propositional formulas α in the proposition letters of φ. This is decidable by Theorem 2.12. □

Using **TELC** as our base logic we will now consider minimal conservative models and minimal entailment.

4 Minimal models and minimal entailment To describe the behavior of a reasoning agent over time, we assume we have a finite number of subjective **TEL**-formulas (or just a single one, the finite conjunction of these formulas). We are interested in the consequences of this description. It is for instance possible to describe the behavior of an agent performing default reasoning by translating a default rule $(\alpha : \beta)/\gamma$ into the **TEL**-rule $\mathsf{K}\alpha \wedge G(\neg\mathsf{K}\neg\beta) \to G(\mathsf{K}\gamma)$, as described in Engelfriet and Treur [5]. This description forces conclusions to be added in certain circumstances. However, we want the knowledge of the agent to be minimal: only those facts should be known which are prescribed by the description to be known, and no other facts. So we make the explicit assumption that “all the agent knows” is what is dictated by the description. Apart from the temporal aspect, this is similar in spirit to the theory of epistemic states of Halpern and Moses [12], introduced to formalize the notion of “knowing only φ.” For a broader discussion of minimalization of models, see for instance van Benthem [21].

We will formalize this minimality by introducing a preference relation over **TELC**-models which favors models with as little propositional knowledge as possible. Formulas are assumed to be subjective.

Definition 4.1 (Minimal models and entailment)

1. We extend the degree of information ordering to **TELC**-models $\mathcal{M}, \mathcal{N}$:

$$\mathcal{M} \leq \mathcal{N} \iff \text{for all } s \in \mathbb{N} : \mathcal{M}_s \leq \mathcal{N}_s.$$

   We write $\mathcal{M} < \mathcal{N}$ if $\mathcal{M} \leq \mathcal{N}$ and $\mathcal{M} \neq \mathcal{N}$.
2. A **TELC**-model $\mathcal{M}$ is a *minimal conservative model* of φ, denoted $\mathcal{M} \models_{\min} \varphi$, if $\mathcal{M} \models \varphi$ and for all conservative models $\mathcal{N}$, if $\mathcal{N} \models \varphi$ and $\mathcal{N} \leq \mathcal{M}$ then $\mathcal{N} = \mathcal{M}$.
3. For **TEL**-formulas φ, ψ, we say φ is a *minimal conservative consequence* of ψ or ψ *minimally entails* φ, denoted $\psi \models^c_{\min} \varphi$, if for all minimal conservative models $\mathcal{M}$ of ψ, $\mathcal{M} \models \varphi$ holds.

For a subjective formula φ (which describes the reasoning of an agent), its minimal models represent the process of the agent's reasoning in time. We can then use minimal consequence to infer properties of this reasoning process.

Note that the notion of minimal entailment strengthens the notion of conservative entailment in the sense that $\varphi \models^c \psi$ implies $\varphi \models^c_{\min} \psi$. An easy example, even without temporal operators, shows that it is a proper extension: although $\mathsf{K}\, p \not\models^c \neg \mathsf{K}\, q$, we do have $\mathsf{K}\, p \models^c_{\min} \neg \mathsf{K}\, q$.

The minimal consequence relation defined here on **TEL**-formulas can be seen as a temporalization of Ground **S5** (or Minimal **S5**) as studied in for example Donini, Nardi, and Rosati [4], which in turn is a generalization of the entailment relation of [12] mentioned before. Semantically, Minimal **S5** can be defined in a way similar to minimal conservative consequence: a normal **S5**-model M is a minimal model of an **S5**-formula α if $M \models_{\mathbf{S5}} \mathsf{K}\,\alpha$ and for all **S5**-models N, if $N \models_{\mathbf{S5}} \mathsf{K}\,\alpha$ and $N \leq M$ (where $\leq$ is the degree-of-information ordering on **S5**-models of Definition 3.1), then $N = M$. For **S5**-formulas α, β, we define $\alpha \models^{\mathbf{S5}}_{\min} \beta$ if $\mathsf{K}\,\beta$ is true in all minimal models of α. The following is easy to prove.

Proposition 4.2 *Let α, β be* **S5**-*formulas, then:*

$$\alpha \models^{\mathbf{S5}}_{\min} \beta \iff \mathsf{K}\,\alpha \models^c_{\min} \mathsf{K}\,\beta.$$

So there is an almost trivial reduction of Minimal **S5** to our minimal conservative consequence. We will use this fact later on when we discuss complexity. If $\mathsf{K}\,\alpha$ has only one minimal model, then α is called *honest*, and when we restrict the premises to honest formulas, we get the entailment relation of [12].

Since we are working with a fixed propositional signature P, the above definition of minimal entailment seems to depend on P, but this is not actually the case.

Proposition 4.3 *The notion $\models^c_{\min}$ is independent of the propositional signature.*

Proof: For a propositional signature P we write $\mathcal{L}_P$ to denote the temporal epistemic language based on P and P-$\models^c_{\min}$ to denote the associated notion of minimal conservative consequence. It is sufficient to show that for two signatures P, Q with $P \subseteq Q$ we have that for all formulas φ, ψ in $\mathcal{L}_P$: φ P-$\models^c_{\min}$ ψ if and only if φ Q-$\models^c_{\min}$ ψ.

Let P, Q be two propositional signatures with $P \subseteq Q$. For a propositional valuation m of signature Q, $m|_P$ denotes the restriction of m to atoms of P. Consider the following constructions:

- For a **TEL**-model $\mathcal{M}$ based on Q, we define its restriction to P, $\mathcal{M}|_P$ by:
$$(\mathcal{M}|_P)_s = \{m|_P : m \in \mathcal{M}_s\}.$$
- For a **TEL**-model $\mathcal{M}$ based on P, we define its extension to Q, $\mathcal{M}|^Q$ by:
$$(\mathcal{M}|^Q)_s = \{m \in \mathrm{Mod}(Q) : m|_P \in \mathcal{M}_s\}.$$

By induction on $\varphi \in \mathcal{L}_P$ it is easy to see that truth of φ at a point in time is preserved under these constructions.

Now suppose that $\mathcal{M}$ is a conservative **TEL**-model based on Q and $\mathcal{M} \models_{\min} \varphi$ (with the notion of $\models_{\min}$ based on Q). Then $\mathcal{M}|_P \models_{\min} \varphi$ (with the notion of $\models_{\min}$ based on P): for suppose $\mathcal{N}$ is a conservative **TEL**-model based on P with $\mathcal{N} < \mathcal{M}|_P$ and $\mathcal{N} \models \varphi$, then (!) $\mathcal{N}|^Q < \mathcal{M}$ and $\mathcal{N}|^Q \models \varphi$.

Conversely, suppose that $\mathcal{M}$ is a conservative **TEL**-model based on P and $\mathcal{M} \models_{\min} \varphi$. Then $\mathcal{M}|^Q \models_{\min} \varphi$: for suppose $\mathcal{N}$ is a conservative **TEL**-model based on Q with $\mathcal{N} < \mathcal{M}|^Q$ and $\mathcal{N} \models \varphi$, then (!) $\mathcal{N}|_P < \mathcal{M}$ and $\mathcal{N}|_P \models \varphi$.

It is now easy to see that φ P-$\models^c_{\min}$ ψ if and only if φ Q-$\models^c_{\min}$ ψ. □

As an example of the use of these notions, it has been shown in [5] that minimal entailment can capture skeptical consequence in default logic (see [18]). A default theory consists of a set of propositional formulas, called the *axioms* and denoted by W, and a set D of defaults of the form $(\alpha : \beta)/\gamma$, where α, β, and γ are propositional formulas. Such a default has the intended meaning: if you believe α and β is consistent with your beliefs, then you should also believe γ. The theory of Reiter ([18]) then prescribes how, using the default rules, you can extend W to a set of formulas, called an *extension*.

Definition 4.4 (Reiter extension) Let $\langle W, D\rangle$ be a default theory. A set of propositional sentences E is a *Reiter extension* of $\langle W, D\rangle$ if and only if:

$$\begin{aligned} E &= \bigcup_{i=0}^{\infty} E_i \qquad \text{with} \\ E_0 &= Cn(W) \qquad \text{and for } i \geq 0 : \\ E_{i+1} &= Cn(E_i \cup \{\gamma | (\alpha : \beta)/\gamma \in D, \alpha \in E_i \text{ and } \neg\beta \notin E\}), \end{aligned}$$

where $Cn(A)$ denotes the set of all propositional consequences of A.

Note that in the definition, the sets E_i depend on E, making the definition nonconstructive. In general for a default theory there may be multiple extensions. If a formula φ is in all of these extensions, we call φ a *skeptical consequence* of the default theory.

Example 4.5 (Default logic) Let a finite default theory $\Delta = \langle W, D\rangle$ be given and let $\psi = \bigwedge\{\mathsf{K}\alpha \wedge G(\neg\mathsf{K}\neg\beta) \rightarrow G(\mathsf{K}\gamma) | (\alpha : \beta)/\gamma \in D\} \wedge \bigwedge\{\mathsf{K}\alpha | \alpha \in W\}$. Then φ is a skeptical consequence of Δ if and only if $\psi \models^c_{\min} F(\mathsf{K}\varphi)$ (see [5]).

We are interested in the complexity of minimal entailment; we will first concentrate on the decidability.

5 Decidability of minimal entailment The first question to be asked when investigating the complexity of a notion is whether it is decidable. The notion of minimal entailment will turn out to be decidable, but in order to prove that we will first need some lemmas.

Observation 5.1 A conservative **TEL**-model $\mathcal{M}$ consists of a sequence of normal **S5**-models. These models consist of a finite number of propositional valuations, since P is assumed to be finite. Furthermore the sequence is (not necessarily strictly) decreasing. Therefore there must exist a time point $s \in \mathbb{N}$ such that for all $t > s : \mathcal{M}_t = \mathcal{M}_s$. If s_0 is the smallest point for which this is true, we say that $\mathcal{M}$ *stabilizes* at s_0.

Since all **TELC**-models stabilize, it is possible to store them in finite space.

The idea in the proof of decidability is that for each formula ψ there is a number n_ψ such that a minimal model of ψ must stabilize before n_ψ. Then there is only a finite number of models to be checked, and since they stabilize, it is always possible to check whether a temporal formula holds in them. To obtain the upper bound n_ψ one reasons that if there exists a long enough sequence of identical states in a model before it stabilizes, then it is possible to insert an extra (identical) state into this sequence without disturbing the truth of ψ. Since this enlarged model is smaller (with respect to $\leq$) than the original, the original model could not have been a minimal model of ψ. The length of such a sequence depends on the depth of nesting of temporal operators in ψ. We will now formalize these ideas.

Definition 5.2 (Depth) The *depth of nesting* of temporal operators in a formula φ, $\text{depth}(\varphi)$, is defined inductively as follows:

- $\text{depth}(\varphi) = 0$, if $\varphi \in \mathcal{L}_{\mathbf{S5}}$
- $\text{depth}(\alpha \wedge \beta) = \max\{\text{depth}(\alpha), \text{depth}(\beta)\}$
- $\text{depth}(\neg\alpha) = \text{depth}(\alpha)$
- $\text{depth}(P\alpha) = \text{depth}(F\alpha) = \text{depth}(\alpha) + 1$

The first lemma states that in a sequence of identical states, formulas with small enough depth cannot discriminate between states in the middle of the sequence. Lemmas 5.3, 5.4, and Fact 5.5 are also valid for nonsubjective formulas.

Lemma 5.3 *If $\mathcal{M}$ is a* **TEL**-*model such that for some $N \geq 1,\ s \geq N$:*

$$\mathcal{M}_s = \mathcal{M}_{s+i} = \mathcal{M}_{s-i} \quad \textit{for all } 1 \leq i \leq N,$$

then for all φ with $\text{depth}(\varphi) < N$ *and* $1 \leq j \leq N - \text{depth}(\varphi)$*:*

$$(\mathcal{M}, s - j) \models \varphi \Leftrightarrow (\mathcal{M}, s) \models \varphi \Leftrightarrow (\mathcal{M}, s + j) \models \varphi.$$

Proof: By induction on φ, where the only interesting cases are the temporal operators (the abbreviation "i.h." stands for induction hypothesis).

$F\alpha$: Let $1 \leq j \leq N - \text{depth}(F\alpha)$. The implications from right to left are trivial, so we will prove only $(\mathcal{M}, s - j) \models F\alpha \Rightarrow (\mathcal{M}, s + j) \models F\alpha$. Suppose $(\mathcal{M}, s - j) \models F\alpha$. There exists $n \in \mathbb{N},\ n > s - j$ with $(\mathcal{M}, n) \models \alpha$. If $n > s + j$ then $(\mathcal{M}, s + j) \models F\alpha$, so suppose $s - j < n \leq s + j$.

1. If $n = s - k$ with $1 \leq k < j$ then $1 \leq k < j \leq N - \text{depth}(F\alpha) < N - \text{depth}(\alpha)$ and by the i.h. we get $(\mathcal{M}, s) \models \alpha$.
2. If $n = s$ then $(\mathcal{M}, s) \models \alpha$.
3. If $n = s + k$ with $1 \leq k \leq j$ then $1 \leq k \leq j \leq N - \text{depth}(F\alpha) < N - \text{depth}(\alpha)$, so by the i.h. $(\mathcal{M}, s) \models \alpha$.

So we have $(\mathcal{M}, s) \models \alpha$ and $1 \leq j + 1 \leq N - (\text{depth}(F\alpha) - 1) = N - \text{depth}(\alpha)$, so by the i.h. we have $(\mathcal{M}, s + (j + 1)) \models \alpha$, and so $(\mathcal{M}, s + j) \models F\alpha$.

$P\alpha$: Analogous to $F\alpha$. □

We will often use this lemma with $j = 1$ and $N = \text{depth}(\varphi) + 1$. The following example shows that we really need that many identical states.

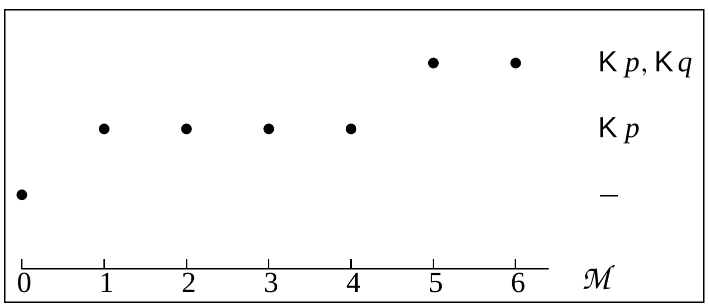

This picture represents the model in which nothing is known at time point 0, p is known from time point 1 onwards, and q is known from time point 5. We have $(\mathcal{M}, 3) \not\models G(\mathsf{K}\,q)$ but $(\mathcal{M}, 3 + 1) \models G(\mathsf{K}\,q)$ (we need an extra $\mathsf{K}\,p$ state between 4 and 5); also $(\mathcal{M}, 2 - 1) \models H(\neg\mathsf{K}\,p)$ but $(\mathcal{M}, 2) \not\models H(\neg\mathsf{K}\,p)$ (we need an extra $\mathsf{K}\,p$ state between 0 and 1).

The next lemma shows that if we have a sequence of identical states, a middle state can be duplicated or removed without changing the truth of formulas with sufficiently small depth of operator-nesting.

Lemma 5.4 *Let $\mathcal{M}$ be a model as in Lemma 5.3. Define $f : \mathbb{N} \to \mathbb{N}$ as follows:*

$$f(n) = \begin{cases} n & \text{if } n \leq s \\ n - 1 & \text{if } n > s. \end{cases}$$

and let $\mathcal{N}$ be a model satisfying $\mathcal{N}_i = \mathcal{M}_{f(i)}$ for all $i \in \mathbb{N}$. Then for all formulas φ with $\text{depth}(\varphi) \leq N$ *we have:*

$$(\mathcal{N}, i) \models \varphi \iff (\mathcal{M},\ f(i)) \models \varphi \text{ for all } i \in \mathbb{N}.$$

Proof: By induction on φ, where the only nontrivial cases are the operators (for which we will take H and G).

$H\varphi$: Suppose $(\mathcal{N}, i) \models H\varphi$. Take $k < f(i)$. Then there exists $t < i$ such that $f(t) = k$ and then $(\mathcal{N}, t) \models \varphi$, so by the i.h. $(\mathcal{M}, k) \models \varphi$. Thus $(\mathcal{M}, f(i)) \models H\varphi$.

Suppose $(\mathcal{M}, f(i)) \models H\varphi$.

- If $i \leq s$, take $k < i$ then $f(k) < f(i)$, so $(\mathcal{M}, f(k)) \models \varphi$ and by the i.h. $(\mathcal{N}, k) \models \varphi$. We have $(\mathcal{N}, i) \models H\varphi$.
- If $i \geq s + 1$, take $k < i$;
  - If $k \neq s$ then $f(k) < f(i)$, so $(\mathcal{M}, f(k)) \models \varphi$ and by the i.h. $(\mathcal{N}, k) \models \varphi$.
  - If $k = s$ then $s - 1 < f(i)$, so $(\mathcal{M}, s - 1) \models \varphi$. As depth$(H\varphi) \leq N$ we have $1 \leq 1 \leq N -$ depth(φ), and by Lemma 5.3 we have $(\mathcal{M}, s) \models \varphi$. By the i.h. $(\mathcal{N}, s) \models \varphi$, or $(\mathcal{N}, k) \models \varphi$.

So we have $(\mathcal{N}, i) \models H\varphi$.

$G\varphi$: Analogous. □

The following picture sketches the situation with $N = 2$.

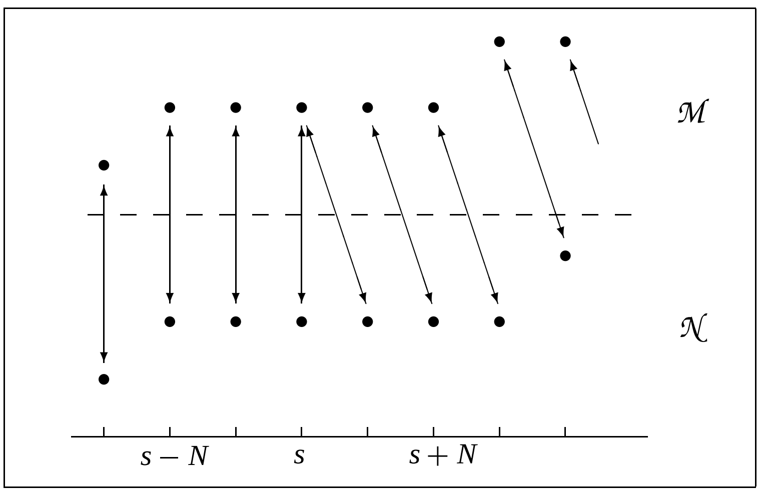

Another way of proving this lemma is to show that there exist bisimulations up to N between these two models. The main use of the lemma lies in the possibility of enlarging or reducing sequences of identical states in a model without disturbing truth of formulas with sufficiently small depth of nesting.

Fact 5.5 *For the models $\mathcal{M}$, $\mathcal{N}$ of Lemma 5.4 the following holds: if $\mathcal{M}$ is conservative then $\mathcal{N}$ is conservative and vice versa, $\mathcal{N} \leq \mathcal{M}$, and if there exists $t \geq s + N$ such that $\mathcal{M}_t < \mathcal{M}_{t+1}$ then $\mathcal{N} < \mathcal{M}$.*

Proof: Take $s \in \mathbb{N}$, then $\mathcal{N}_s = \mathcal{M}_{f(s)}$. Since $f(s) \leq s$ and $\mathcal{M}$ is conservative we have $\mathcal{M}_{f(s)} \leq \mathcal{M}_s$ so $\mathcal{N}_s \leq \mathcal{M}_s$. If there exists $t \geq s + N$ such that $\mathcal{M}_t < \mathcal{M}_{t+1}$ then $\mathcal{N}_{t+1} = \mathcal{M}_{f(t+1)} = \mathcal{M}_t < \mathcal{M}_{t+1}$. □

This fact and the previous lemma allow us to conclude that for each formula there is a time point such that the minimal models of the formula must stabilize before this point. From now on we will again restrict ourselves to subjective formulas.

Lemma 5.6 *Suppose the propositional signature P consists of n atoms. If a conservative model $\mathcal{M}$ of signature P is a minimal model of a subjective formula φ then it stabilizes on or before time point $(2^n - 1) \cdot 2 \cdot \text{depth}(\varphi)$.*

Proof: First we will show that a minimal model $\mathcal{M}$ of φ cannot have more than $2 \cdot \text{depth}(\varphi)$ successive identical states before it stabilizes. Suppose $\mathcal{M} \models_{\min} \varphi$ and it has at least $2 \cdot \text{depth}(\varphi) + 1$ successive identical states before it stabilizes. So there exists $s \geq \text{depth}(\varphi)$ such that $\mathcal{M}_s = \mathcal{M}_{s+i} = \mathcal{M}_{s-i}$ for all $1 \leq i \leq \text{depth}(\varphi)$, and $t \geq s + \text{depth}(\varphi)$ such that $\mathcal{M}_t < \mathcal{M}_{t+1}$. Now consider the model $\mathcal{N}$ as described in Lemma 5.4. Since $\mathcal{M} \models \varphi$ we have $\mathcal{N} \models \varphi$, and by Fact 5.5 we have $\mathcal{N} < \mathcal{M}$. Therefore $\mathcal{M}$ cannot be a minimal model of φ.

As P has n atoms, there exist 2^n different propositional models. Since a conservative model $\mathcal{M}$ consists of a decreasing sequence of (nonempty) sets of propositional models, there are at most $2^n - 1$ points s such that $\mathcal{M}_s < \mathcal{M}_{s+1}$. If $\mathcal{M}$ is a minimal model of φ then there can be at most $2 \cdot \text{depth}(\varphi)$ successive identical states before it stabilizes, and therefore $\mathcal{M}$ must stabilize on or before time point $(2^n - 1) \cdot 2 \cdot \text{depth}(\varphi)$. □

Lemma 5.7 *For a conservative model $\mathcal{M}$, $s \in \mathbb{N}$ and a formula φ it is decidable whether $(\mathcal{M}, s) \models \varphi$.*

Proof: Suppose we have a conservative model $\mathcal{M}$ and $s \in \mathbb{N}$. By Observation 5.1, $\mathcal{M}$ stabilizes at some point s_0. It is easily seen from Lemma 5.3 that for a formula φ we have $(\mathcal{M}, t) \models \varphi \iff (\mathcal{M}, u) \models \varphi$ for all $t, u \geq s_0 + \text{depth}(\varphi)$. Then use induction on φ. □

Most importantly, it is decidable if a model is a minimal model of a subjective formula.

Lemma 5.8 *For a conservative model $\mathcal{M}$ and a subjective formula φ it is decidable whether $\mathcal{M} \models_{\min} \varphi$.*

Proof: First, we need to check whether $\mathcal{M} \models \varphi$, which is equivalent to checking $(\mathcal{M}, 0) \models \Box\varphi$, which is decidable by Lemma 5.7. Suppose P has n atoms. If $\mathcal{M}$ stabilizes after time point $(2^n - 1) \cdot 2 \cdot \text{depth}(\varphi)$ it is not a minimal model of φ by Lemma 5.6. So suppose $\mathcal{M} \models \varphi$ and $\mathcal{M}$ stabilizes on or before time point $(2^n - 1) \cdot 2 \cdot \text{depth}(\varphi)$.

In order to check whether $\mathcal{M} \models_{\min} \varphi$ we have to see if there exists a conservative model smaller than $\mathcal{M}$ which satisfies φ. Of course in general there are an infinite number of conservative models smaller than $\mathcal{M}$, but we will show that we have only to consider models which stabilize not later than time point $(2^n - 1) \cdot (4 \cdot \text{depth}(\varphi) + 1)$. In other words, we will show that if there exists a conservative model smaller than $\mathcal{M}$ satisfying φ, there also exists such a model which stabilizes on or before point $(2^n - 1) \cdot (4 \cdot \text{depth}(\varphi) + 1)$. The converse of this statement is of course trivial.

Suppose we have a conservative model $\mathcal{N}$ with $\mathcal{N} < \mathcal{M}$ and $\mathcal{N} \models \varphi$, and let s be the stabilizing point of $\mathcal{N}$. If $s \leq (2^n - 1) \cdot (4 \cdot \text{depth}(\varphi) + 1)$ then we are done, so suppose not. Now consider the following procedure for constructing a model $\mathcal{N}'$: if there exists a sequence of more than $2 \cdot \text{depth}(\varphi) + 1$ successive identical states in $\mathcal{N}$ between time points $(2^n - 1) \cdot 2 \cdot \text{depth}(\varphi)$ and s, then we delete points from this sequence until it has length $2 \cdot \text{depth}(\varphi) + 1$. Lemma 5.4 ensures that we can do this without disturbing the truth of φ. It is also easy to see that the result is conservative and still (strictly) smaller than $\mathcal{M}$. Let $\mathcal{N}'$ be the model which results from applying this procedure for every such sequence. Then $\mathcal{N}' \models \varphi$ and $\mathcal{N}' < \mathcal{M}$. Let s' be the stabilizing point of $\mathcal{N}'$. Then in $\mathcal{N}'$ there are at most $2^n - 1$ points t with $(2^n - 1) \cdot 2 \cdot \mathrm{depth}(\varphi) \leq t < s$ and $\mathcal{N}'_t < \mathcal{N}'_{t+1}$. Between such points there are at most $2 \cdot \mathrm{depth}(\varphi) + 1$ identical states and therefore $s \leq (2^n - 1) \cdot 2 \cdot \mathrm{depth}(\varphi) + (2^n - 1) \cdot (2 \cdot \mathrm{depth}(\varphi) + 1) = (2^n - 1) \cdot (4 \cdot \mathrm{depth}(\varphi) + 1)$.

It is easy to see that, given the finite signature, there are only a finite number of conservative models which stabilize not later than time point $(2^n - 1) \cdot (4 \cdot \mathrm{depth}(\varphi) + 1)$. For each such model $\mathcal{N}$ we can check whether $\mathcal{N} < \mathcal{M}$ (only the first $(2^n - 1) \cdot (4 \cdot \mathrm{depth}(\varphi) + 1)$ time points have to be considered), and we can check whether $\mathcal{N} \models \varphi$ (again decidable). If we find such a model then $\mathcal{M} \not\models_{\min} \varphi$, otherwise $\mathcal{M} \models_{\min} \varphi$. □

Now we are ready to prove decidability of minimal entailment.

Theorem 5.9 (Decidability of minimal entailment) *For two subjective formulas φ, ψ it is decidable whether $\varphi \models^c_{\min} \psi$.*

Proof: We can take the signature P to consist of the atoms occurring in φ and ψ. Suppose there are n such atoms. Then Lemma 5.6 states that we have only to consider models which stabilize not later than time point $(2^n - 1) \cdot 2 \cdot \mathrm{depth}(\varphi)$, and since the signature is finite, there are only finitely many such models. For each such model $\mathcal{M}$ it is decidable by Lemma 5.8 whether $\mathcal{M} \models_{\min} \varphi$. Now we have only to check for each of these (finitely many) minimal models $\mathcal{M}$ of φ whether $\mathcal{M} \models \psi$, which is decidable by Lemma 5.7. □

Of course the procedure given in the proof will be very inefficient.

Having established that both **TELC** and minimal entailment are decidable, in the next section we will look at the complexity of these notions, and in particular whether the minimalization process has a structural impact on complexity.

6 Complexity We will first give a brief overview of the relevant concepts of complexity theory needed in the rest of this chapter. This is meant as a reminder for the reader, not as an introduction to this field (see Johnson [14] for a good introduction). Especially the Polynomial Hierarchy (PH) will concern us here. The Polynomial Hierarchy is a hierarchy of classes of problems of increasing complexity. The two best known complexity classes in PH are P and NP. The basic notion in defining complexity classes is the Turing Machine (TM). The class P consists of all problems solvable by a deterministic TM running in time polynomial in the length of the input. Problems solvable by a nondeterministic TM running in polynomial time form the class NP. For any complexity class C, the class co-C consists of the problems whose complement is in C. In order to define the other classes in PH, we need the notion of an oracle TM, which is a TM that has access to an oracle for a particular decision problem: all instances of that problem can be solved in one time step by consulting the oracle. Formally, if C is a complexity class then the class NP^C consists of those problems solvable by a nondeterministic TM with access to an oracle for a problem in C, running in time polynomial in the input size. Now set:

$$\Sigma_0^P = \Pi_0^P = P, \qquad \text{and for } k \geq 0:$$

$$\Sigma^P_{k+1} = \mathrm{NP}^{\Sigma^P_k} \text{ and } \Pi^P_{k+1} = \text{co-}\Sigma^P_{k+1}$$

Note that $\Sigma^P_1 = \mathrm{NP}$ and $\Pi^P_1 = \text{co-NP}$. For a problem p, if for any problem in class C there is a polynomial transformation of that problem to p, then p is called C-hard. If p is in C and is C-hard, it is called C-complete. If a C-hard problem can be (polynomially) transformed to p, p is also C-hard.

In order to study its complexity we will first look at satisfiability of **TELC**. Without loss of generality we restrict ourselves to satisfiability of subjective formulas in time point 0.

Definition 6.1 (**TELC**(0)-SAT) A subjective formula φ is in **TELC**(0)-SAT if there exists a **TELC**-model $\mathcal{M}$ such that $(\mathcal{M}, 0) \models \varphi$.

Remark 6.2 It is easy to see that **TELC**(0)-SAT is polynomially reducible (and vice versa) to satisfiability (in any time point): φ is satisfiable if and only if $\varphi \vee F\varphi$ is in **TELC**(0)-SAT, and φ is in **TELC**(0)-SAT if and only if $\Box(at_0 \to \varphi)$ is satisfiable.

Definition 6.3 (Size of a **TELC**-model) For a **TELC**-model $\mathcal{M}$ we call its stabilizing point the *size* of $\mathcal{M}$, denoted $\text{size}(\mathcal{M})$.

Definition 6.4 (Subformula) Let $\text{Subf}(\varphi)$ denote the subformulas of φ, where maximal **S5**-subformulas of φ are not further decomposed, and let $\text{Subf}\,\mathbf{S5}(\varphi)$ denote the set of subformulas of φ which are in $\mathcal{L}_{\mathbf{S5}}$.

We give an example to clarify this definition: $\text{Subf}(G(\mathsf{K}\,p \wedge \mathsf{K}\,q)) = \{G(\mathsf{K}\,p \wedge \mathsf{K}\,q),\ \mathsf{K}\,p \wedge \mathsf{K}\,q\}$ and $\text{Subf}\,\mathbf{S5}(G(\mathsf{K}\,p \wedge \mathsf{K}\,q)) = \{\mathsf{K}\,p \wedge \mathsf{K}\,q, \mathsf{K}\,p, \mathsf{K}\,q, p, q\}$. So $\text{Subf}(\varphi) \cup \text{Subf}\,\mathbf{S5}(\varphi)$ is the set of all subformulas of φ.

First we will prove a small-model theorem for **TELC**. Let $\text{length}(\varphi)$ denote the length of the formula φ as a string.

Lemma 6.5 (Small model theorem) *If a subjective formula φ is in* **TELC***(0)*-SAT *then there exists a* **TELC***-model $\mathcal{M}$ such that $(\mathcal{M}, 0) \models \varphi$, $\text{size}(\mathcal{M}) \leq 4 \cdot (\text{length}(\varphi))^2$, and for all $i \in \mathbb{N}$ the* **S5***-model $\mathcal{M}_i$ contains not more than $2 \cdot \text{length}(\varphi)$ valuations.*

Proof: Suppose for some **TELC**-model $\mathcal{N}$ we have $(\mathcal{N}, 0) \models \varphi$ and let $s_{\mathcal{N}}$ be the stabilizing point of $\mathcal{N}$. Let $\mathcal{L}_0$ denote the propositional language based on P.

Now let $A = \{\psi, \neg\psi | \psi \in \mathcal{L}_0, \psi \in \text{Subf}\,\mathbf{S5}(\varphi)\}$ and for $i \in \mathbb{N}$:

$$B(i) = \{\mathsf{K}\,\psi | \psi \in A, \mathcal{N}_i \models \mathsf{K}\,\psi\} \cup \{\neg\,\mathsf{K}\,\psi | \psi \in A, \mathcal{N}_i \not\models \mathsf{K}\,\psi\}.$$

Based on these sets we will define a **TELC**-model $\mathcal{N}'$.

For each $\neg\,\mathsf{K}\,\psi \in B(s_{\mathcal{N}})$ choose a valuation $m \in \text{Mod}(P)$ such that $m \not\models \psi$ and $m \models \alpha$ for each $\mathsf{K}\,\alpha \in B(s_{\mathcal{N}})$ (such a valuation exists since $(\mathcal{N}, s_{\mathcal{N}}) \not\models \mathsf{K}\,\psi$ and $(\mathcal{N}, s_{\mathcal{N}}) \models \mathsf{K}\,\alpha$ for each $\mathsf{K}\,\alpha \in B(s_{\mathcal{N}})$). Let M be the set of these valuations. We have $M \models B(s_{\mathcal{N}})$. If there are no formulas $\neg\,\mathsf{K}\,\psi \in B(s_{\mathcal{N}})$ then choose any valuation m with $m \models \alpha$ for each $\mathsf{K}\,\alpha \in B(s_{\mathcal{N}})$ (which again exists). Set $\mathcal{N}'_j = M$ for all $j \geq s_{\mathcal{N}}$. It easy to verify that $\mathcal{N}'_j \models B(j)$ for all $j \geq s_{\mathcal{N}}$.

Now using induction on $s_{\mathcal{N}} > j \geq 0$, let $B(j) \backslash B(j+1) = \{\neg\,\mathsf{K}\,\psi_1, \ldots, \neg\,\mathsf{K}\,\psi_n\}$ (because $\mathcal{N}$ is conservative there will be no formulas $\mathsf{K}\,\psi$ in this set). For $k = 1, \ldots, n$ choose a valuation m_k with $m_k \not\models \psi_k$ and $m \models \alpha$ for each $\mathsf{K}\,\alpha \in B(j)$ (again such valuations exist). Let $\mathcal{N}'_j = \mathcal{N}'_{j+1} \cup \{m_1, \ldots, m_k\}$. It is again easy to verify that $\mathcal{N}'_j \models B(j)$.

The resulting model $\mathcal{N}'$ has the following properties:

1. $\mathcal{N}'$ is a **TELC**-model.
2. $\mathcal{N}'_j \models B(j)$ for all $j \in \mathbb{N}$.
3. The number of valuations of $\mathcal{N}'_j$ is smaller than the number of elements in $A(\leq 2 \cdot \text{length}(\varphi))$.
4. $(\mathcal{N}', 0) \models \varphi$: Take $\psi \in \text{Subf}(\varphi) \cap \mathcal{L}_{\mathbf{S5}}$ (which must be subjective). Then using a normal form described in [16] it is easy to see that ψ is equivalent to a formula $\psi' = \delta_1 \vee \cdots \vee \delta_m$ with for $i = 1, \ldots, m$: $\delta_i = \mathsf{K}\varphi_{1,i} \wedge \cdots \wedge \mathsf{K}\varphi_{k(i),i} \wedge \neg\mathsf{K}\psi_{1,i} \wedge \cdots \wedge \neg\mathsf{K}\psi_{\ell(i),i}$, with $\varphi_{jk}, \psi_{jk} \in A$. So using (2) we have:
$$\mathcal{N}'_i \models \mathsf{K}\varphi_{jk} \iff \mathcal{N}_i \models \mathsf{K}\varphi_{jk} \quad \text{and}$$
$$\mathcal{N}'_i \models \neg\mathsf{K}\psi_{jk} \iff \mathcal{N}_i \models \neg\mathsf{K}\psi_{jk},$$
so $\mathcal{N}'_i \models \psi' \iff \mathcal{N}_i \models \psi'$ and thus $\mathcal{N}'_i \models \psi \iff \mathcal{N}_i \models \psi$. An easy induction gives: for all $i \in \mathbb{N}$, for all $\psi \in \text{Subf}(\varphi) : (\mathcal{N}', i) \models \psi \iff (\mathcal{N}, i) \models \psi$ and therefore $(\mathcal{N}', 0) \models \varphi$.
5. The number of i for which $\mathcal{N}'_i < \mathcal{N}'_{i+1}$ is less than $2 \cdot \text{length}(\varphi)$: real changes occur at most once for each $\neg\mathsf{K}\psi$ with $\psi \in A$ and A contains at most $2 \cdot \text{length}(\varphi)$ elements.

Now construct the model $\mathcal{M}$ as follows: for each sequence of more than $2 \cdot \text{depth}(\varphi) + 1$ identical states in $\mathcal{N}'$, before its stabilizing point, delete states from this sequence until it has length $2 \cdot \text{depth}(\varphi) + 1$. Let $\mathcal{M}$ be the resulting model. Now Lemma 5.4 ensures that $(\mathcal{M}, 0) \models \varphi$. Furthermore $2 \cdot \text{depth}(\varphi) + 1 \leq 2 \cdot \text{length}(\varphi)$ so that $\text{size}(\mathcal{M}) \leq (2 \cdot \text{length}(\varphi))^2$. □

With this lemma we can show that **TELC**(0)-SAT is in NP, using methods similar to those in, e.g., [19] and Ladner [15].

Theorem 6.6 **TELC***(0)*-SAT *is in* NP.

Proof: For a subjective formula φ we present the following nondeterministic algorithm to verify if φ is in **TELC**(0)-SAT. A nondeterministic Turing Machine (M) guesses $4 \cdot (\text{length}(\varphi))^2$ Kripke models $\mathcal{M}_i$ with each not more than $2 \cdot \text{length}(\varphi)$ valuations, such that $\mathcal{M}_i \supseteq \mathcal{M}_{i+1}$. $\mathcal{M}$ will be this model, remaining constant after time point $4 \cdot (\text{length}(\varphi))^2$. Then it verifies if $(\mathcal{M}, 0) \models \varphi$ as follows: for each $i \in \{0, \ldots, 4 \cdot (\text{length}(\varphi))^2 + \text{length}(\varphi)\}$, M maintains a set label(i) which is initialized to the empty set and at the end will contain the subformulas of φ true at time point i. Now for each $\psi \in \text{Subf}(\varphi)$ we do the following (starting with the **S5**-subformulas, and treating ψ only if all of its subformulas have already been treated): for each $i \in \{0, \ldots, 4 \cdot (\text{length}(\varphi))^2 + \text{length}(\varphi)\}$ update label(i) as follows:

1. Add $\psi \in \mathcal{L}_{\mathbf{S5}}$ to label(i) if and only if $\mathcal{M}_i \models \psi$ (this can be checked in time polynomial in the number of states in $\mathcal{M}_i$, using a labeling algorithm similar to the one described here, see, e.g., [13]).
2. Add $\neg\psi$ to label(i) if and only if $\psi \notin$ label(i).
3. Add $\alpha \wedge \beta$ to label(i) if and only if $\alpha \in$ label(i) and $\beta \in$ label(i).

4. Add $F\alpha$ to label(i) if and only if $\alpha \in$ label(j) for some $j > i$ (If $i = 4 \cdot (\text{length}(\varphi))^2 + \text{length}(\varphi)$ then add $F\alpha$ to label(i) if and only if $\alpha \in$ label(i)).
5. Add $P\alpha$ to label(i) if and only if $\alpha \in$ label(j) for some $j < i$.

Now we have $(\mathcal{M}, 0) \models \varphi$ if and only if $\varphi \in$ label(0) at the end of this procedure. It is easy to verify that this algorithm works properly in time polynomial in length(φ). Lemma 6.5 ensures that there is a guess for which M halts in an accepting state if and only if φ is in **TELC**(0)-SAT. □

This gives us the following corollary.

Corollary 6.7 **TELC** *satisfiablity is* NP-*complete.*

Proof: The reduction given in Remark 6.2 ensures that **TELC** satisfiablity is in NP, and clearly a propositional formula φ is satisfiable if and only if $\mathsf{M}\,\varphi$ is **TELC** satisfiable. As satisfiability of propositional formulas is NP-complete, **TELC** satisfiability is also NP-complete. □

We would like to show that the minimalization of models makes the consequence relation more complex, and we can do this using the reduction of Minimal **S5** to minimal conservative consequence, as described in Proposition 4.2.

Proposition 6.8 *Minimal conservative consequence is* Π_3^P*-hard.*

Proof: The reduction of Proposition 4.2 is clearly polynomial, and Minimal **S5** is Π_3^P-complete ([4]). □

So minimal consequence is harder than **TELC**-consequence (which is Π_1^P-complete, or co-NP-complete), provided the polynomial hierarchy does not collapse (see [14]).

In [7] a sublanguage of the subjective part of $\mathcal{L}_{\mathbf{TEL}}$ is proposed as a specification language for (conservative) reasoning processes, and it is shown that this language is suited for this task. We will now look at the complexity of minimal entailment restricted to this language. Let $H_0\varphi$ be an abbreviation for $(at_0 \to \varphi)$.

Definition 6.9 The language $\mathcal{L}'$ is the smallest set such that:

1. If $\alpha \in \mathcal{L}_0$ then $\mathsf{K}\,\alpha \in \mathcal{L}'$.
2. If $\alpha, \beta, \gamma, \psi$, and $\varphi \in \mathcal{L}_0$ then $H_0(\mathsf{K}\,\alpha) \wedge H_0(\neg\,\mathsf{K}\,\beta) \wedge \mathsf{K}\,\gamma \wedge G(\neg\,\mathsf{K}(\neg\psi)) \to G(\mathsf{K}\,\varphi) \in \mathcal{L}'$.
3. If $\varphi, \psi \in \mathcal{L}'$ then $\varphi \wedge \psi \in \mathcal{L}'$.

For $\varphi \in \mathcal{L}'$ and $\psi = F(\mathsf{K}\,\alpha)$ with $\alpha \in \mathcal{L}_0$ we define $\varphi \models'^{\,c}_{\min} \psi$ if and only if $\varphi \models^{c}_{\min} \psi$.

The basis of the language is formed by the formulas in “rule format” of item (2) of the definition. It prescribes the inference of a conclusion (φ) if some conditions are met. These conditions may refer to the facts which are (un)known at the start of the reasoning process (the part with the H_0-operators), to facts currently known (γ), and it may contain a “global consistency check” (ψ) in analogy with the translated rules for default logic. If $G(\neg\,\mathsf{K}(\neg\psi))$ is true at some point in time, then $\neg\psi$ is never known in the future, which means that ψ remains consistent with what the agent knows. The formulas of item (1) just prescribe facts which should be known from the start (initial knowledge). Conjunctions are allowed to make a single formula of rules and initial facts. The formula $F(\mathsf{K}\,\alpha)$ expresses that α will be known sometime in the future (and can be regarded as a conclusion of the reasoning process).

Since we can reduce default logic to this fragment (see Example 4.5) and default logic is Π_2^P-complete (Gottlob [11], Stillman [20], see also Papadimitriou and Sideri [17]), $\models'^{c}_{\min}$ is Π_2^P-hard. However, it is no harder than that.

Proposition 6.10 *$\models'^{c}_{\min}$ is Π_2^P-complete.*

Proof: We will describe a nondeterministic Turing Machine M with access to an NP-oracle for determining whether *not* $\varphi \models'^{c}_{\min} \psi$ (similar to the proofs in [20], [17] or [11]). A minimal model of φ can have no identical states before it stabilizes. For each conjunct $H_0(\mathsf{K}\,\alpha) \wedge H_0(\neg\,\mathsf{K}\,\beta) \wedge \mathsf{K}\,\gamma \wedge G(\neg\,\mathsf{K}(\neg\delta)) \to G(\mathsf{K}\,\epsilon)$ in φ, M guesses a time point $i \geq 1$ but not more than n, where n is the number of these conjuncts plus one, from which time onwards ϵ will be assumed to hold (or it guesses that ϵ will never hold). Denote for $i \in \{0, \ldots, n\}$, the set of formulas assumed to hold at i plus the formulas α for which there is a conjunct $\mathsf{K}\,\alpha$ in φ, by $A(i)$. Then M uses the NP-oracle to perform the following:

1. Let $f(\epsilon)$ be the point from which ϵ is assumed to hold (and so $f(\epsilon) \in \{1, \ldots, n, \infty\}$). Now it checks for all $i \in \{1, \ldots, n\}$ if $\{\mathsf{K}\,\epsilon \mid f(\epsilon) \leq i\} \cup \{\neg\,\mathsf{K}\,\epsilon \mid f(\epsilon) > i\}$ is **S5**-satisfiable (using the oracle; note that **S5**-satisfiability is in NP). If not, it halts in a rejecting state (the guess does not induce a TELC-model).
2. For each conjunct $H_0(\mathsf{K}\,\alpha) \wedge H_0(\neg\,\mathsf{K}\,\beta) \wedge \mathsf{K}\,\gamma \wedge G(\neg\,\mathsf{K}(\neg\delta)) \to G(\mathsf{K}\,\epsilon)$ and for each time point $i \in \{0, \ldots, n\}$ it computes whether $A(0) \models \alpha$, whether $A(0) \not\models \beta$, whether $A(i) \models \gamma$ and whether for no $i < j \leq n$, $A(j) \models \delta$, using the NP-oracle. If this is true for no time point then it checks whether ϵ is assumed never to hold; otherwise it takes the first such point and checks whether ϵ is assumed to hold from the next time point on. If these conditions are violated then M halts in a rejecting state (the guess does not induce a minimal model of φ).
3. It checks if $A(n) \models \chi$ (when $\psi = F(\mathsf{K}\,\chi)$). If this is the case then in this minimal model of φ, ψ holds, so M halts in a rejecting state (the guess does not induce a minimal model of φ in which ψ fails). Otherwise it halts in an accepting state (the guess induces a minimal model of φ in which ψ does not hold).

This nondeterministic algorithm is polynomial in φ (using an NP-oracle for propositional consequence and **S5**-satisfiability) so the converse of $\models'^{c}_{\min}$ is in Σ_2^P which implies that $\models'^{c}_{\min}$ is in Π_2^P. Together with Π_2^P-hardness this gives the desired result. □

Apart from default logic, skeptical consequence relations of many other well-known nonmonotonic logics such as McDermott and Doyle's nonmonotonic logic, autoepistemic logic, and nonmonotonic logic **N** are Π_2^P-complete ([19]), which means that we can reduce these relations to minimal consequence (or even $\models'^{c}_{\min}$), using a polynomial reduction. Further research is needed to find these reductions.

We would also like to have an upper bound on the complexity of minimal consequence. In order to get this, we need to sharpen some previous lemmas. Lemma 5.6 gave an upper bound on the size of minimal models of φ, but it is not polynomial in the length of φ. We already know that the length of a sequence of identical states in a minimal model is polynomially bounded, so we will try to find a polynomial bound on the number of transitions between nonidentical states in a minimal model. The key is that in a minimal model of φ, after such a transition occurs, the agent will know (at least) one of the subformulas of φ he did not know before. In fact, a minimal model of φ is uniquely determined by the subformulas of φ which are true at any moment in time. We will now make this formal.

Definition 6.11 For a subjective formula φ, define $A(\varphi) = \{\psi, \neg\psi | \psi \in \mathcal{L}_0 \cap \mathrm{Subf}\,\mathbf{S5}(\varphi)\}$. A TELC-model $\mathcal{M}$ of φ is *based on* φ (abbreviated $\mathrm{bo}(\varphi)$) if there exist sets $A(i)$ for each $i \in \mathbb{N}$ with $A(0) \subseteq A(1) \subseteq \cdots \subseteq A(\varphi)$ and $\mathcal{M}_i = \mathrm{Mod}(A(i)) = \{m \in \mathrm{Mod}(P) | m \models A(i)\}$.

Lemma 6.12 *If* $\mathcal{M} \models_{\min} \varphi$ *then* $\mathcal{M}$ *is* $\mathrm{bo}(\varphi)$ *and* $\mathrm{size}(\mathcal{M}) \leq 4 \cdot (\mathrm{length}(\varphi))^2$.

Proof: Suppose $\mathcal{M}$ is not based on φ. Define $A(i) = \{\alpha, \neg\alpha | \alpha \in A(\varphi) \text{ and } \mathcal{M}_i \models \mathsf{K}\alpha\}$ and let $\mathcal{N}_i = \mathrm{Mod}(A(i))$. Clearly $A(0) \subseteq A(1) \subseteq \cdots \subseteq A(\varphi)$, so $\mathcal{N}$ is a TELC-model and $\mathcal{N} < \mathcal{M}$. Furthermore for all $\alpha \in \mathcal{L}_0 \cap \mathrm{Subf}\,\mathbf{S5}(\varphi)$ we have $\mathcal{M}_i \models \mathsf{K}\alpha \iff \mathcal{N}_i \models \mathsf{K}\alpha$ and $\mathcal{M}_i \models \mathsf{M}\alpha \iff \mathcal{N}_i \models \mathsf{M}\alpha$, so using the same argument as in the proof of Lemma 6.5 we have $\mathcal{N} \models \varphi$. This contradicts the assumption that $\mathcal{M} \models_{\min} \varphi$, so $\mathcal{M}$ is based on φ. But then the number of changes in $\mathcal{M}$ (the points $i \in \mathbb{N}$ where $\mathcal{M}_i < \mathcal{M}_{i+1}$) cannot be larger than the number of elements of $A(\varphi)$ and in between such updates there cannot be sequences of identical states longer than $2 \cdot \mathrm{depth}(\varphi) + 1$ so $\mathrm{size}(\mathcal{M}) \leq 4 \cdot (\mathrm{length}(\varphi))^2$. □

Notice that a model $\mathcal{M}$ based on φ can equivalently be described by giving for each formula in $A(\varphi)$ the time point at which it is known in $\mathcal{M}$, or "infinity" if this is never the case. We have a similar result for models which refute that $\mathcal{M}$ is a minimal model of φ.

Lemma 6.13 *If* $\mathcal{M} \models \varphi$ *but* $\mathcal{M} \not\models_{\min} \varphi$*, then there exists a* TELC*-model* $\mathcal{N}$ *such that* $\mathcal{N} < \mathcal{M}$, $\mathcal{N} \models \varphi$*, and* $\mathcal{N}$ *is based on* φ *with* $\mathrm{size}(\mathcal{N}) \leq \mathrm{size}(\mathcal{M}) + 4 \cdot (\mathrm{length}(\varphi))^2$.

Proof: Suppose $\mathcal{M} \models \varphi$ but $\mathcal{M} \not\models_{\min} \varphi$ then there is a TELC-model $\mathcal{M}'$ with $\mathcal{M}' < \mathcal{M}$ and $\mathcal{M}' \models \varphi$. In the same way as in the proof of Lemma 6.12 we can make a model $\mathcal{M}''$ which is a model of φ based on φ and $\mathcal{M}'' \leq \mathcal{M}'$. Now from any sequence of identical states in $\mathcal{M}''$ after $\mathrm{size}(\mathcal{M})$ but before $\mathrm{size}(\mathcal{M}'')$ with length more than $2 \cdot \mathrm{depth}(\varphi) + 1$ we can delete states until it has length $2 \cdot \mathrm{depth}(\varphi) + 1$. Let $\mathcal{N}$ be the resulting model (this construction is the same as the one used in the proof of Lemma 5.8). So we have $\mathcal{N} < \mathcal{M}$, $\mathcal{N} \models \varphi$, and $\mathcal{N}$ is based on φ. Furthermore, $\mathcal{N}$ has less than $2 \cdot \mathrm{length}(\varphi)$ updates, and sequences between $\mathrm{size}(\mathcal{M})$ and $\mathrm{size}(\mathcal{N})$ have length no greater than $2 \cdot \mathrm{depth}(\varphi) + 1$, so $\mathrm{size}(\mathcal{N}) \leq \mathrm{size}(\mathcal{M}) + 2 \cdot \mathrm{length}(\varphi) \cdot 2 \cdot \mathrm{length}(\varphi) = \mathrm{size}(\mathcal{M}) + 4 \cdot (\mathrm{length}(\varphi))^2$. □

Lemma 6.14 *Deciding for a formula* φ *and a model* $\mathcal{M}$ *based on* φ *whether* $\mathcal{M} \models_{\min} \varphi$ *is in* Π_2^P.

Proof: We assume the model $\mathcal{M}$ encoded as described in the remark after Lemma 6.12: there is a function $f : A(\varphi) \to \mathbb{N} \cup \{\infty\}$ such that $f(\alpha)$ gives the time point from which α is known. We will show that deciding whether $\mathcal{M} \not\models_{\min} \varphi$ is in Σ_2^P by describing a nondeterministic Turing Machine M with access to an NP-oracle. Let $\mathrm{size}(\mathcal{M}) = \max(f[A(\varphi)] \backslash \{\infty\})$ (if $f[A(\varphi)] = \{\infty\}$, then let $\mathrm{size}(\mathcal{M}) = 0$). First we check whether $\text{size}(\mathcal{M}) \leq 4 \cdot (\text{length}(\varphi))^2$; if not we halt in an accepting state. Otherwise we use a labeling algorithm as described earlier to check if $\mathcal{M} \models \varphi$. The range of time points we have to check is from 0 to $\text{size}(\mathcal{M}) + \text{length}(\varphi)$. The subformulas in $\text{Subf}(\varphi) \cap \mathcal{L}_{\mathbf{S5}}$ are treated as follows: for such a formula α and time point i it is checked (using the NP-oracle) if $\{\mathsf{K}\epsilon \mid f(\epsilon) \leq i\} \cup \{\neg\mathsf{K}\epsilon \mid f(\epsilon) > i\} \models_{\mathbf{S5}} \alpha$. If so, α is added to label(i), otherwise not. If $\mathcal{M} \not\models \varphi$, M halts in an accepting state (certainly $\mathcal{M} \not\models_{\min} \varphi$). Otherwise M guesses a **TELC**-model $\mathcal{N}$ by guessing a function $g : A(\varphi) \to \mathbb{N} \cup \{\infty\}$ such that:

1. $f(\epsilon) \leq g(\epsilon)$;
2. either $g(\epsilon) \leq \text{size}(\mathcal{M}) + 4 \cdot (\text{length}(\varphi))^2$ or $g(\epsilon) = \infty$;
3. for at least one $\epsilon \in A(\varphi)$ we have $g(\epsilon) > f(\epsilon)$.

Then it checks for $i \in \{0, \ldots, \text{size}(\mathcal{M}) + 4 \cdot (\text{length}(\varphi))^2\}$ whether $\{\mathsf{K}\epsilon \mid g(\epsilon) \leq i\} \cup \{\neg\mathsf{K}\epsilon \mid g(\epsilon) > i\}$ is **S5**-consistent, using the oracle. If not, we halt in a rejecting state (g does not describe a **TELC**-model). Otherwise we know that g induces a **TELC**-model $\mathcal{N}$ with $\mathcal{N} < \mathcal{M}$ (if such a guess is not possible then we halt in a rejecting state because $\mathcal{M} \models_{\min} \varphi$). Next we use the labeling algorithm to check whether $\mathcal{N} \models \varphi$; if not we halt in a rejecting state, otherwise in an accepting state: $\mathcal{N}$ is a smaller model of φ. It is clear that the algorithm works in polynomial time (using the NP-oracle). Lemma 6.13 ensures that there is a guess for which M halts in an accepting state if and only if $\mathcal{M} \not\models_{\min} \varphi$. Thus deciding if $\mathcal{M} \not\models_{\min} \varphi$ is in Σ_2^P, so the complement is in Π_2^P. □

Theorem 6.15 *Deciding whether $\varphi \models_{\min}^c \psi$ is in Π_3^P.*

Proof: We will show that deciding whether *not* $\varphi \models_{\min}^c \psi$ is in Σ_3^P by giving a nondeterministic Turing Machine M with access to a Π_2^P-oracle. First M guesses a **TELC**-model $\mathcal{M}$ based on φ by guessing a function $f : A(\varphi) \to \mathbb{N} \cup \{\infty\}$ such that for all $\epsilon \in A(\varphi)$ either $f(\epsilon) \leq 4 \cdot (\text{length}(\varphi))^2$ or $f(\epsilon) = \infty$. Then it checks for $i \in \{0, \ldots, 4 \cdot (\text{length}(\varphi))^2\}$ whether $\{\mathsf{K}\epsilon \mid f(\epsilon) \leq i\} \cup \{\neg\mathsf{K}\epsilon \mid f(\epsilon) > i\}$ is **S5**-consistent, using the oracle. If not it halts in a rejecting state (f does not induce a **TELC**-model). Now it uses the Π_2^P-oracle to determine whether $\mathcal{M} \models_{\min} \varphi$. If not it halts in a rejecting state. Otherwise it uses a labeling algorithm to check whether $\mathcal{M} \models \psi$ (as in the proof of the previous lemma, using the Π_2^P-oracle for **S5**-consequence); if this is true M halts in a rejecting state, otherwise in an accepting state. The algorithm works in polynomial time, and Lemma 6.12 ensures there is a guess for which M halts in an accepting state if and only if *not* $\varphi \models_{\min}^c \psi$. So as this is in Σ_3^P, the complement is in Π_3^P. □

Combining this with Proposition 6.8, we immediately get the following.

Corollary 6.16 *Minimal conservative consequence is Π_3^P-complete.*

7 Downward persistence The entailment relation we have defined is a nonmonotonic one, which means that one can have that $\alpha \models_{\min}^c \gamma$ but not $\alpha \wedge \beta \models_{\min}^c \gamma$ for some formulas α, β, and γ (see Gabbay, Hogger and Robinson [9]). We are interested in the class of formulas β which can be added to the premises without disturbing any of the conclusions. It will turn out that this is the class of downward persistent formulas (see also [21]). In the rest of this section we will investigate the class of formulas which are preserved under decreasing or increasing (with respect to $\leq$) the models. Since our logic is essentially a temporalized version of **S5**, we will first look at **S5**-formulas preserved under going to larger and smaller models.

Definition 7.1 (Preservation under supermodels)

1. An **S5**-formula φ is *preserved under supermodels* if for any two **S5**-models M, N such that $N \subseteq M$, and $m \in N$: if $(N, m) \models_{\mathbf{S5}} \varphi$ then $(M, m) \models_{\mathbf{S5}} \varphi$.
2. Define the class of **S5**-formulas DIAM by:

$$\text{DIAM} := p|\neg p|\text{DIAM} \wedge \text{DIAM}|\text{DIAM} \vee \text{DIAM}|\,\mathsf{M}(\text{DIAM})$$

We want to prove that formulas in this class are the only ones (up to equivalence) which are preserved under supermodels.

Theorem 7.2 *An* **S5**-*formula* φ *is preserved under supermodels if and only if it is* **S5**-*equivalent to a formula in* DIAM.

Proof: It is easy to see that a formula equivalent to one in DIAM is preserved under supermodels. Now let φ be preserved under supermodels. Suppose $\text{Mod}(P) = \{m_1, \ldots, m_n\}$. For $i = 1, \ldots, n$ define $A(i) = \min\{N \subseteq \text{Mod}(P) | (N, m_i) \models_{\mathbf{S5}} \varphi\}$, where for a set $\mathcal{B}$ of **S5**-models, $\min \mathcal{B} = \{N \in \mathcal{B}|$ there is no $M \in \mathcal{B}$ such that M is a proper subset of $N\}$. Define for $j = 1, \ldots, n : \alpha_j := \bigwedge\{p | p \in P, m_j \models p\} \wedge \bigwedge\{\neg p | p \in P, m_j \not\models p\}$, and for an **S5**-model M, $\varphi_M = \bigwedge\{\mathsf{M}\alpha_j | j = 1, \ldots, n$ and $m_j \in M\}$. It is easy to see that for an **S5**-model N we have that $M \subseteq N$ if and only if $(N, m) \models_{\mathbf{S5}} \varphi_M$ for some or all $m \in N$. Now define for $j = 1, \ldots, n$:

$$\psi_j = \begin{cases} \alpha_j \wedge \bigvee\{\varphi_M | M \in A(j)\} & \text{if there exists an } \mathbf{S5}\text{-model } N \text{ with } (N, m_j) \models_{\mathbf{S5}} \varphi, \\ \bot & \text{otherwise.} \end{cases}$$

Note that $\bot$ is equivalent to $\mathsf{M}(p \wedge \neg p)$. Now let $\psi = \bigvee\{\psi_j | j = 1, \ldots, n\}$. Then ψ is in DIAM. We will show that ψ is equivalent to φ.

Suppose $(N, m_i) \models_{\mathbf{S5}} \varphi$. Then there exists an $M \in A(i)$ with $M \subseteq N$, so $(N, m_i) \models_{\mathbf{S5}} \varphi_M$ and $(N, m_i) \models_{\mathbf{S5}} \alpha_i$. Hence $(N, m_i) \models_{\mathbf{S5}} \psi_i$ and $(N, m_i) \models_{\mathbf{S5}} \psi$.

Suppose that $(N, m_i) \models_{\mathbf{S5}} \psi$. Then there exists a j such that $(N, m_i) \models_{\mathbf{S5}} \psi_j$, but then $i = j$ and there exists $M \in A(i)$ such that $(N, m_i) \models_{\mathbf{S5}} \varphi_M$. Thus $M \subseteq N$ and $(M, m_i) \models_{\mathbf{S5}} \varphi$, but since φ is preserved under supermodels we have $(N, m_i) \models_{\mathbf{S5}} \varphi$. □

We are also interested in formulas preserved under taking submodels.

Definition 7.3 (Preservation under submodels)

1. An **S5**-formula φ is *preserved under submodels* if for any two **S5**-models $\mathcal{M}, \mathcal{N}$ such that $\mathcal{N} \subseteq \mathcal{M}$, and $m \in \mathcal{N}$: if $(\mathcal{M}, m) \models_{\mathbf{S5}} \varphi$ then $(\mathcal{N}, m) \models_{\mathbf{S5}} \varphi$.
2. Define the class of **S5**-formulas BOX by:

$$\text{BOX} := p|\neg p|\text{BOX} \wedge \text{BOX}|\text{BOX} \vee \text{BOX}|\,\mathsf{K}(\text{BOX}).$$

Proposition 7.4 *An* **S5**-*formula* φ *is preserved under submodels if and only if it is equivalent to a formula in* BOX.

Proof: Easy. □

Now we are ready to use these results to get a preservation result for **TEL**-formulas. As we were interested in downward persistent formulas because of the link with the rule of monotonicity for minimal consequence, the definition of downward persistence should use the corresponding notion of satisfaction of a formula in a model ($\mathcal{M} \models \varphi$). Also the notion of equivalence between formulas should be based on this notion.

Definition 7.5 (Upward and downward persistence)

1. A subjective **TEL**-formula φ is called

   downward persistent (dp) if for all **TELC**-models $\mathcal{M}$, $\mathcal{N}$:

   $$\text{if } \mathcal{M} \leq \mathcal{N} \text{ and } \mathcal{N} \models \varphi \text{ then } \mathcal{M} \models \varphi;$$

   upward persistent (up) if for all **TELC**-models $\mathcal{M}$, $\mathcal{N}$:

   $$\text{if } \mathcal{M} \leq \mathcal{N} \text{ and } \mathcal{M} \models \varphi \text{ then } \mathcal{N} \models \varphi.$$

2. Define $\text{DP} := \text{DIAM}|\text{DP} \wedge \text{DP}|\text{DP} \vee \text{DP}|F(\text{DP})|G(\text{DP})|P(\text{DP})|H(\text{DP})$
   $\text{UP} := \text{BOX}|\text{UP} \wedge \text{UP}|\text{UP} \vee \text{UP}|F(\text{UP})|G(\text{UP})|P(\text{UP})|H(\text{UP})$
3. For two subjective **TEL**-formulas φ, ψ:

   $$\varphi \sim \psi \iff \text{for all } \textbf{TELC}\text{-models } \mathcal{M} : \mathcal{M} \models \varphi \iff \mathcal{M} \models \psi.$$

We can link the notion of $\sim$ with the notion $\models^c$: if we denote $\varphi \models^c \psi$ and $\psi \models^c \varphi$ by $\varphi \equiv^c \psi$ then: $\varphi \sim \psi \iff \Box\varphi \equiv^c \Box\psi$. This implies that $\sim$ is decidable.

Now we are ready to prove the following.

Theorem 7.6 *A subjective* **TEL**-*formula* φ *is downward persistent if and only if it is equivalent (in the sense of* $\sim$*) to a subjective formula in* DP.

Proof: For a subjective (!) formula $\varphi \in \text{DP}$ one can easily prove that for all **TELC**-models $\mathcal{M}$, $\mathcal{N}$ and $i \in \mathbb{N}$: if $\mathcal{M} \leq \mathcal{N}$ and $(\mathcal{N}, i) \models \varphi$ then $(\mathcal{M}, i) \models \varphi$. This implies that a formula equivalent (in the sense of $\sim$) to one in DP is dp.

Suppose φ is a subjective dp formula. We will construct its equivalent in DP. If there is no **TELC**-model $\mathcal{M}$ such that $\mathcal{M} \models \varphi$ then φ is equivalent to $\perp$. Note that $\perp$ is equivalent to $\mathsf{M}(p \wedge \neg p)$, which is a subjective formula in DP. Suppose we have a propositional signature P with m atoms. For a set of **TELC**-models $\mathcal{B}$ define $\max \mathcal{B} = \{\mathcal{M} \in \mathcal{B} | \text{there is no } \mathcal{N} \in \mathcal{B} \text{ with } \mathcal{M} < \mathcal{N}\}$. If there is a **TELC**-model $\mathcal{M}$ such that $\mathcal{M} \models \varphi$, then we define $\mathcal{A} = \max\{\mathcal{M} | \mathcal{M} \models \varphi\}$. Suppose $\mathcal{M} \models \varphi$ and $\mathcal{M}$ stabilizes after time point $(2^m - 1) \cdot (2 \cdot \text{depth}(\varphi) + 1)$. Then we can delete points in sequences of more than $(2 \cdot \text{depth}(\varphi) + 1)$ identical states before the stabilizing point without disturbing the truth of φ. If we do this for each such a sequence we end up with a model of φ which is larger (with respect to $\leq$) than $\mathcal{M}$ and stabilizes not later than $(2^m - 1) \cdot (2 \cdot \text{depth}(\varphi) + 1)$. Thus: $\mathcal{A} = \max\{\mathcal{M} | \mathcal{M} \models \varphi \text{ and } \mathcal{M} \text{ stabilizes not later than } (2m - 1) \cdot (2 \cdot \text{depth}(\varphi) + 1)\}$. As the set we take the maximal elements of is nonempty and finite and the relation $<$ on **TELC**-models is transitive and irreflexive, $\mathcal{A}$ is nonempty and finite. Note that the argument used here (for maximal models) is similar to the one used for minimal models in the proof of Lemma 5.6: there the idea was that a model which is too long can be enlarged (yielding a smaller model with respect to $\leq$), whereas here the idea is that if a model is too long, it can be reduced (yielding a bigger model with respect to $\leq$).

Suppose $\mathrm{Mod}(P) = \{m_1, \ldots, m_n\}$ (with of course $n = 2^m$). Again define for $j = 1, \ldots, n : \alpha_j := \bigwedge\{p | p \in P, m_j \models p\} \wedge \bigwedge\{\neg p | p \in P, m_j \not\models p\}$. Now define for $i = 1, \ldots, n$ and for a **TELC**-model $\mathcal{M} : n(i, \mathcal{M}) = \sup\{j \in \mathbb{N} | m_i \in \mathcal{M}_j\}$, where $\sup \varnothing = -\infty$. Let

$$\psi(i, \mathcal{M}) = \begin{cases} \Box(at_{n(i,\mathcal{M})} \to \mathsf{M}\,\alpha_i) & \text{if } n(i, \mathcal{M}) \in \mathbb{N} \\ \Box(\mathsf{M}\,\alpha_i) & \text{if } n(i, \mathcal{M}) = \infty \\ \top & \text{if } n(i, \mathcal{M}) = -\infty \end{cases}$$

(Note that $\top$ is equivalent to $sf M(p \vee \neg p)$.)

Furthermore, define $\psi_{\mathcal{M}} = \bigwedge\{\psi(i, \mathcal{M}) | i = 1, \ldots, n\}$. Now it can easily be proven that $\mathcal{N} \models \psi_{\mathcal{M}} \iff \mathcal{N} \leq \mathcal{M}$: the formulas $\psi(i, \mathcal{M})$ make sure that the valuation m_i is in $\mathcal{N}_t$ at least until the last time point s for which m_i is in $\mathcal{M}_s$. Finally, define: $\psi = \bigvee\{\psi_{\mathcal{M}} | \mathcal{M} \in \mathcal{A}\}$. Then ψ is in DP and $\varphi \sim \psi$:

- Suppose $\mathcal{M} \models \varphi$. Then there exists $\mathcal{N} \in \mathcal{A}$ with $\mathcal{M} \leq \mathcal{N}$ (!), so $\mathcal{M} \models \psi_{\mathcal{N}}$ and $\mathcal{M} \models \psi$.
- Suppose $\mathcal{M} \models \psi$. Then there exists $\mathcal{N} \in \mathcal{A}$ with $\mathcal{M} \models \psi_{\mathcal{N}}$, so $\mathcal{M} \leq \mathcal{N}$; and as $\mathcal{N} \in \mathcal{A}$ we have $\mathcal{N} \models \varphi$, and φ was dp, so $\mathcal{M} \models \varphi$. □

As in the case of **S5**-formulas we have the following.

Proposition 7.7 *A subjective* **TEL**-*formula φ is upward persistent if and only if it is equivalent (in the sense of $\sim$) to a subjective formula in* UP.

Proof: If φ is up then $\neg\Box\varphi$ is dp so by the previous theorem $\neg\Box\varphi \sim \psi$ for some $\psi \in$ DP. Then $\varphi \sim \neg\Box\psi$ and $\neg\Box\psi$ is equivalent to some formula in UP. □

Furthermore, the property of downward persistence is decidable.

Proposition 7.8 *For a subjective formula φ it is decidable whether φ is dp.*

Proof: Suppose P contains n propositional atoms. We will prove that φ is dp if and only if for all **TELC**-models $\mathcal{M}, \mathcal{N}$ with $\mathrm{size}(\mathcal{M}) \leq (2^n - 1) \cdot (2 \cdot \mathrm{depth}(\varphi) + 1)$, $\mathrm{size}(\mathcal{N}) \leq 2 \cdot (2^n - 1) \cdot (2 \cdot \mathrm{depth}(\varphi) + 1)$: if $\mathcal{N} \leq \mathcal{M}$ and $\mathcal{M} \models \varphi$ then $\mathcal{N} \models \varphi$. This implies the decidability of dp.

Suppose φ is not dp, then there exist **TELC**-models $\mathcal{M}, \mathcal{N}$ with $\mathcal{N} < \mathcal{M}$, $\mathcal{M} \models \varphi$, and $\mathcal{N} \not\models \varphi$. Now we construct a **TELC**-model $\mathcal{M}'$ by deleting points from sequences of more than $2 \cdot \mathrm{depth}(\varphi) + 1$ identical states before the stabilizing point from $\mathcal{M}$ until each such sequence is exactly $2 \cdot \mathrm{depth}(\varphi) + 1$ states long. Then $\mathrm{size}(\mathcal{M}') \leq (2^n - 1) \cdot (2 \cdot \mathrm{depth}(\varphi) + 1)$, $\mathcal{N} < \mathcal{M}'$, and $\mathcal{M}' \models \varphi$ (by Lemma 5.4). Now we construct a model $\mathcal{N}'$ using the following procedure. First we identify all sequences of identical states in $\mathcal{N}$ after time point $(2^n - 1) \cdot (2 \cdot \mathrm{depth}(\varphi) + 1)$ but before the stabilizing point of $\mathcal{N}$ of length more than $(2 \cdot \mathrm{depth}(\varphi) + 1)$ points. From each such sequence we delete points until it has length $(2 \cdot \text{depth}(\varphi) + 1)$. Then $\text{size}(\mathcal{N}') \leq 2 \cdot (2^n - 1) \cdot (2 \cdot \text{depth}(\varphi) + 1)$, $\mathcal{N}' \not\models \varphi$ (Lemma 5.4), and it is easily checked that $\mathcal{N}' < \mathcal{M}'$. □

Similarly it is decidable whether a formula is up, and this gives us another way of verifying **TELC** theorems since $\vdash_{\mathbf{TELC}} \varphi \iff \mathcal{M}^{ti} \models \varphi$ and φ is up, where $\mathcal{M}^{ti}$ is the *totally ignorant* model defined by $\mathcal{M}^{ti}_s = \text{Mod}(P)$ for all s (note that for all **TELC**-models $\mathcal{N}$ we have $\mathcal{M}^{ti} \leq \mathcal{N}$; use soundness and completeness of **TELC**). Since **TELC**-theoremhood is co-NP-complete, we have the following as an immediate consequence.

Corollary 7.9 *Upward persistence for subjective formulas is co*-NP-*hard.*

For a valuation $m \in \text{Mod}(P)$ we can define the **TELC**-model $\mathcal{M}^m$ by $(\mathcal{M}^m)_t = \{m\}$ for all t. It is easy to see that such a model is maximal in the ordering $\leq$, and this gives us another way of checking **TELC** theorems since $\vdash_{\mathbf{TELC}} \varphi \iff \varphi$ is dp and $\mathcal{M}^m \models \varphi$ for all $m \in \text{Mod}(P)$. Furthermore we have: φ up and dp $\iff \vdash_{\mathbf{TELC}} \varphi$ or $\varphi \sim \bot$, which gives us the following.

Corollary 7.10 *Checking whether a subjective formula is downward and upward persistent is co*-NP-*complete.*

One of the reasons we were interested in formulas preserved under shrinking models was the link to monotonicity, which we can now prove with the following proposition.

Proposition 7.11 *If a subjective formula β is downward persistent then for all subjective formulas α, γ: if $\alpha \models^c_{\min} \gamma$ then $\alpha \wedge \beta \models^c_{\min} \gamma$.*

Proof: Suppose β is downward persistent and that for two formulas α, γ we have $\alpha \models^c_{\min} \gamma$. Take a minimal model $\mathcal{M}$ of $\alpha \wedge \beta$, then $\mathcal{M} \models \alpha \wedge \beta$, so $\mathcal{M} \models \alpha$. But $\mathcal{M}$ is also minimal with respect to this property, for suppose $\mathcal{N} \leq \mathcal{M}$ and $\mathcal{N} \models \alpha$, then since β is downward persistent, we also have $\mathcal{N} \models \beta$, so $\mathcal{N} \models \alpha \wedge \beta$. But since $\mathcal{M}$ was a minimal model of $\alpha \wedge \beta$ we must have $\mathcal{N} = \mathcal{M}$. So $\mathcal{M}$ is a minimal model of α so $\mathcal{M} \models \gamma$. We have proved that $\alpha \wedge \beta \models^c_{\min} \gamma$. □

We have given a syntactical characterization of downward persistent formulas and the link with monotonicity, but it is also possible to characterize the downward persistent formulas using monotonicity (referring only to minimal entailment).

Proposition 7.12 *A subjective formula φ is downward persistent if and only if $\forall \alpha, \beta : \alpha \models^c_{\min} \beta \Rightarrow \alpha \wedge \varphi \models^c_{\min} \beta$.*

Proof: The "only if" part is Proposition 7.11.

Suppose φ is not dp, then there exist **TELC**-models $\mathcal{M}, \mathcal{N}$ such that $\mathcal{N} < \mathcal{M}$, $\mathcal{M} \models \varphi$, but $\mathcal{N} \not\models \varphi$. For a **TELC**-model $\mathfrak{L}$, define (using notation from the proof of Theorem 7.6): $m(i, \mathfrak{L}) = \min\{j \in \mathbb{N} | m_i \notin \mathcal{M}_j\}$ where $\min \varnothing = \infty$ and $\psi^{\mathfrak{L}} = \{\Box(at_{m(i,\mathfrak{L})} \to \mathsf{K}(\neg\alpha_i)) | i = 1, \ldots, n, m(i, \mathfrak{L}) < \infty\}$.

It is easy to see that for a **TELC**-model $\mathcal{K}$, $\mathcal{K} \models \psi^{\mathfrak{L}}$ if and only if $\mathcal{K} \geq \mathfrak{L}$. Now take $\alpha = (\psi^{\mathcal{N}} \wedge (\Box\varphi \to \psi^{\mathcal{M}}))$ and $\beta = \Diamond(\neg\varphi)$. Any **TELC**-model $\mathfrak{L}$ of α has to satisfy $\mathfrak{L} \geq \mathcal{N}$, and $\mathcal{N} \models \alpha$ ($\mathcal{N} \models \Box\varphi \to \psi^{\mathcal{M}}$ since $(\mathcal{N}, i) \not\models \Box\varphi$ for all $i \in \mathbb{N}$). Therefore $\mathcal{N} \models_{\min} \alpha$, and it is the only minimal model of α. Since $\mathcal{N} \models \Diamond(\neg\varphi)$ we have $\alpha \models^c_{\min} \beta$. Any **TELC**-model $\mathfrak{L}$ of $\alpha \wedge \varphi$ has $\mathfrak{L} \models \Box\varphi$, so $\mathfrak{L} \models \psi^{\mathcal{M}}$, which implies $\mathfrak{L} \geq \mathcal{M}$. Also $\mathcal{M} \models \alpha \wedge \varphi$ (since $\mathcal{N} < \mathcal{M}$), so $\mathcal{M}$ is the unique minimal model of $\alpha \wedge \varphi$. But $\mathcal{M} \not\models \beta$, and therefore we do not have $\alpha \wedge \varphi \models^c_{\min} \beta$. □

So this proposition says that a formula is downward persistent if and only if you can always be sure that adding this formula to your knowledge does not disturb any consequences.

8 Conclusions and further research The logic **TELC** was proposed to describe the behavior of a conservative reasoning agent. This logic was shown to be decidable, and a sound and complete axiomatization was given. Based on this logic we defined a notion of minimal entailment and studied the decidability and complexity. **TELC** was found to be co-NP-complete and minimal consequence was shown to be Π^P_3-complete. Furthermore, a syntactical characterization of formulas preserved under going to smaller models was presented and a link with monotonicity was given.

The fact that the interaction between the epistemic part and the temporal part of the logic is quite limited (only conservativity gives a link) is important for the results in this paper. No interaction axioms were required for **TEL**, and the soundness and completeness results easily followed from [8]. The syntactical characterization of Section 7 was obtained by first treating **S5** and using this for **TEL**. Compositionality makes things easier.

The translation of default logic into **TEL** is already known ([5]); further work is needed to find the translations for other nonmonotonic logics such as autoepistemic logic.

Although a decision procedure is sketched for minimal entailment, we would also like to have an axiomatization. This might not be easy: it would immediately yield an axiomatization for default logic, which has not been given before.

We have characterized the downward persistent formulas. We would like to find a similar result for the class of formulas which have no minimal models (like $F(\mathsf{K}\, p)$). These are the formulas which are in a sense not “honest” since they do not describe the reasoning behavior of an agent properly.

The use of **S5** as the logic to describe the knowledge of the agent at any point in time (allowing negative introspection) is not always realistic. If we use another modal logic such as **S4**, many results in this paper would have to be re-examined; in particular the complexity might be higher. A number of constructions used in the proofs will no longer work, and we might have to use methods like those in, for instance, Andréka, van Benthem and Németi [2].

It would also be interesting to lift the restriction of conservativity. This plays an important role in many of the proofs in this article but does not allow retraction, which is needed for belief revision (see for instance Alchourrón, Gärdenfors and Makinson [1]). In the nonconservative case, we would also like to extend the language with operators like Next, Since, and Until.

Finally, we would like to extend the framework to the case of many agents, also allowing communication between agents and interaction with the outside world. It is not straightforward how to extend the information ordering to this case. Some ideas on how to do this are given in [12].

Acknowledgments I would like to thank Johan van Benthem and Dick de Jongh for their remarks, suggestions, and corrections and for their enthusiasm. I am grateful to Thomas Eiter and Georg Gottlob for their comments leading to the lower bound for the complexity of minimal consequence. Furthermore I would like to thank Pascal van Eck and Rineke Verbrugge for commenting on earlier versions of this paper.

REFERENCES

[1] Alchourrón, C. E., P. Gärdenfors, and D. Makinson, "On the logic of theory change: partial meet contraction and revision functions," *The Journal of Symbolic Logic*, vol. 50 (1985), pp. 510–530. Zbl 0578.03011 MR 87c:03020 8

[2] Andréka, H., J. van Benthem, and I. Németi, "Back and forth between modal logic and classical logic," Institute for Language, Logic and Information, Report ML-95-04, Amsterdam, 1995. Forthcoming in: *Bulletin of the Interest Group for Pure and Applied Logics*. Zbl 0840.03010 MR 96k:03036 8

[3] de Bakker, J. W., W.-P. de Roever, and G. Rozenberg, *Linear Time, Branching Time and Partial Orders in Logics and Models for Concurrency*, Lecture Notes in Computer Science, vol. 354, Springer-Verlag, Berlin, 1989. Zbl 0683.68001 MR 90i:68009 2

[4] Donini, F. M., D. Nardi, and R. Rosati, "Ground nonmonotonic modal logics for knowledge representation," pp. 133-144 in *Proceedings of the Second World Congress on Fundamentals of Artificial Intelligence,WOCFAI-95*, edited by M. De Glas and Z. Pawlak, Angkor, Paris, 1995. 4, 6

[5] Engelfriet, J., and J. Treur, "A temporal model theory for default logic," pp. 91–96 in *Proceedings of the 2nd European Conference on Symbolic and Quantitative Approaches to Reasoning and Uncertainty, ECSQARU '93*, edited by M. Clarke, R. Kruse, and S. Moral, Lecture Notes in Computer Science, vol. 747, Springer-Verlag, Berlin, 1993. MR 95e:03092 4, 4, 4.5, 8

[6] Engelfriet, J., and J. Treur, "Temporal theories of reasoning," pp. 279–299 in *Logics in Artificial Intelligence, Proceedings of the 4th European Workshop on Logics in Artificial Intelligence, JELIA '94*, edited by C. MacNish, D. Pearce, and L. M. Pereira, Lecture Notes in Artificial Intelligence, vol. 838, Springer-Verlag, Berlin, 1994. MR 1320729 1, 2

[7] Engelfriet, J., and J. Treur, "Specification of nonmonotonic reasoning," pp. 111–125 in *Practical Reasoning, Proceedings of the International Conference on Formal and Applied Practical Reasoning, FAPR'96*, edited by D. M. Gabbay and H. J. Ohlbach, Lecture Notes in Artificial Intelligence, vol. 1085, Springer-Verlag, Berlin, 1996. MR 1457030 1, 6

[8] Finger, M., and D. M. Gabbay, "Adding a temporal dimension to a logic system," *Journal of Logic, Language and Information*, vol. 1 (1992), pp. 203–233. Zbl 0798.03031 MR 95h:03037 2, 2, 2, 2, 2, 8

[9] Gabbay, D. M., C. J. Hogger, and J. A. Robinson, *Handbook of Logic in Artificial Intelligence and Logic Programming, vol. 3: Nonmonotonic Reasoning and Uncertain Reasoning*, Oxford University Press, Oxford, 1994. Zbl 0804.03017 MR 95b:03003 7

[10] Goldblatt, R., *Logics of Time and Computation*, CSLI Lecture Notes, No. 7, Stanford University Press, Palo Alto, 1992. Zbl 0635.03024 MR 88k:03058 2

[11] Gottlob, G., "Complexity results for nonmonotonic logics," *Journal of Logic and Computation*, vol. 2, (1992), pp. 397–425. Zbl 0765.03012 MR 93i:03034 6, 6

[12] Halpern, J. Y., and Y. Moses, "Towards a theory of knowledge and ignorance," pp. 125–143 in *Proceedings of the Workshop on Non-monotonic Reasoning*, AAAI Press, Menlo Park, 1984. 4, 4, 4, 8

[13] Halpern, J. Y., and Y. Moses, "A guide to the modal logics of knowledge and belief," pp. 480–490 in *Proceedings of the 9th IJCAI*, edited by A. Joshi, Kaufmann, Los Altos, 1985. 2, 1

[14] Johnson, D. S., "A catalog of complexity classes," pp. 67–161 in *Handbook of Theoretical Computer Science, vol. A: Algorithms and complexity,* edited by J. van Leeuwen, Elsevier Science Publishers, Amsterdam, 1990. Zbl 0900.68246 MR 1127168 6, 6

[15] Ladner, R. E., "The computational complexity of provability in systems of modal propositional logic," *SIAM Journal on Computing*, vol. 6 (1977), pp. 467–480. Zbl 0373.02025 MR 56:8326 6

[16] Meyer, J.-J. Ch., and W. van der Hoek, *Epistemic Logic for Computer Science and Artificial Intelligence*, Cambridge Tracts in Theoretical Computer Science, vol. 41, Cambridge University Press, Cambridge, 1995. Zbl 0868.03001 MR 96m:68166 2, 2, 4

[17] Papadimitriou, C. H., and M. Sideri, "On finding extensions of default theories," pp. 276–281 in *Proceedings of the 4th International Conference on Database Theory, ICDT-92*, edited by J. Biskup and R. Hull, Lecture Notes in Computer Science, vol. 646, Springer-Verlag, Berlin, 1992. 6, 6

[18] Reiter, R., "A logic for default reasoning," *Artificial Intelligence*, vol. 13 (1980), pp. 81–132. Zbl 0435.68069 MR 83e:68138 1, 4, 4

[19] Sistla, A. P., and E. M. Clarke, "The complexity of propositional linear temporal logics," *Journal of the ACM*, vol. 32 (1985), pp. 733–749. Zbl 0632.68034 MR 87j:68077 2, 6

[20] Stillman, J., "The complexity of propositional default logics," pp. 794 –799 in *Proceedings AAAI-92*, edited by W. Swartout, AAAI Press, Menlo Park, 1992. MR 93k:68089 6, 6

[21] van Benthem, J. F. A. K., "Semantic parallels in natural language and computation," pp. 331–375 in *Logic Colloquium, Granada 1987*, edited by H. D. Ebbinghaus et al., North Holland, Amsterdam, 1989. Zbl 0694.03022 MR 91k:03075 4, 7

Department of Mathematics and Computer Science
Vrije Universiteit Amsterdam
De Boelelaan 1081a
1081HV Amsterdam
The Netherlands
email: joeri@cs.vu.nl
web: http://www.cs.vu.nl/∼joeri