

A Logic for Temporal Conditionals and a Solution to the Sea Battle Puzzle

Fengkui Ju¹

School of Philosophy, Beijing Normal University

Gianluca Grilletti²

Institute for Logic, Language and Computation, University of Amsterdam

Valentin Goranko³

Department of Philosophy, Stockholm University

and Department of Mathematics, University of Johannesburg (visiting professorship)

Abstract

Temporal reasoning with conditionals is more complex than both classical temporal reasoning and reasoning with timeless conditionals, and can lead to some rather counter-intuitive conclusions. For instance, Aristotle’s famous “Sea Battle Tomorrow” puzzle leads to a fatalistic conclusion: whether there will be a sea battle tomorrow or not, but that is necessarily the case now. We propose a branching-time logic LTC to formalise reasoning about temporal conditionals and provide that logic with adequate formal semantics. The logic LTC extends the Nexttime fragment of CTL*, with operators for model updates, restricting the domain to only future moments where antecedent is still possible to satisfy. We provide formal semantics for these operators that implements the restrictor interpretation of antecedents of temporalized conditionals, by suitably restricting the domain of discourse. As a motivating example, we demonstrate that a naturally formalised in our logic version of the ‘Sea Battle’ argument renders it unsound, thereby providing a solution to the problem with fatalist conclusion that it entails, because its underlying reasoning per cases argument no longer applies when these cases are treated not as material implications but as temporal conditionals. On the technical side, we analyze the semantics of LTC and provide a series of reductions of LTC-formulae, first recursively eliminating the dynamic update operators and then the path quantifiers in such formulae. Using these reductions we obtain a sound and complete axiomatization for LTC, and reduce its decision problem to that of the modal logic KD.

Keywords: conditionals, temporal settings, restrictors, the Sea Battle Puzzle

¹ Email: fengkui.ju@bnu.edu.cn. Supported by the National Social Science Foundation of China (No. 12CZX053) and the Major Program of the National Social Science Foundation of China (NO. 17ZDA026).

² Email: grilletti.gianluca@gmail.com. Supported by the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (No. 680220).

³ Email: valentin.goranko@philosophy.su.se. Partly supported by research grant 2015-04388 of the Swedish Research Council.

1 Introduction

Temporal conditionals The original philosophical motivation for the present work is rooted in the concept of *conditional* (see e.g. [4], [1]). Classical logic only deals with the simplest kind of conditionals, viz. material implications. However, the logical laws for material implications sometimes lead to counter-intuitive, or just plainly inconsistent conclusions when applied to other types of conditionals. Here is a simple example from [19]. There are two marbles in a black box and we know that one of them is blue and the other is red. Someone picks a marble from the box but we do not see which one. Intuitively the sentence “*the picked marble must be blue or the picked marble must be red*” is false. However, each of the following sentences are true: “*the picked marble is not blue or is not red. if the picked marble is not blue then it must be red. if the picked marble is not red then it must be blue.*” These sentences seem to imply that the picked marble must be blue or the picked marble must be red.

Some conditionals involve temporality. Here is an example from chess: “*if you attack the opponent’s queen with your knight in the next move, you will eventually lose the game.*” We call such conditionals *temporal conditionals*. Various logical problems are concerned with such conditionals and a typical one is the *Sea Battle Puzzle*. This puzzle goes back to Aristotle and is also closely related to the Master Argument of Diodorus Cronus. Briefly, it goes as follows⁴: either there will be a sea battle tomorrow or not; if there is a sea battle tomorrow, it is necessarily so; if there is no sea battle tomorrow, it is necessarily so. Thus, either there will necessarily be a sea battle tomorrow or there will necessarily be no sea battle tomorrow. The conclusion of this argument seems fatalistic and unacceptable for many, but its premises seem mostly fine.

Restrictors Several works, including [8], [18], [7] and [20] have proposed treatment of the semantics of conditionals based on the concept of *restrictor*, which can be traced back to Ramsey’s Test [16]. The restrictor-based view interprets conditionals differently from material implication. By this view, the conditional ‘if ϕ then ψ ’ is not a connective relating two sentences. Instead, the if-clause, ‘if ϕ ’, is a device for restricting discourse domains that are collections of possibilities. Thus, ‘if ϕ then ψ ’ is true with respect to a domain iff ψ is true with respect to the domain restricted by ‘if ϕ ’.

Reading ‘if ϕ ’ as a restrictor can resolve some logical problems involving conditionals. The two puzzles mentioned above have the same inference pattern, that is, reasoning per cases: “ ϕ_1 or ϕ_2 ; if ϕ_1 then ψ_1 ; if ϕ_2 then ψ_2 ; therefore, ψ_1 or ψ_2 ”. As pointed out in [2], [6] and [10], reasoning per cases is not generally sound under the restrictor reading of conditionals. Here is why. Assume that ‘ ϕ_1 or ϕ_2 ’, ‘if ϕ_1 then ψ_1 ’ and ‘if ϕ_2 then ψ_2 ’ are true with respect to a discourse

⁴ Aristotle does not present the puzzle clearly and there are different versions of it in the literature. We use the one from [5].

domain Δ . Let Δ^{ϕ_1} and Δ^{ϕ_2} be the respective results of restricting Δ with ‘if ϕ_1 ’ and ‘if ϕ_2 ’. What ‘if ϕ_1 then ψ_1 ’ and ‘if ϕ_2 then ψ_2 ’ say is just that ψ_1 is true with respect to Δ^{ϕ_1} and ψ_2 is true with respect to Δ^{ϕ_2} . But if neither the truth of ψ_1 nor the truth of ψ_2 is upward monotonic relative to discourse domains, then it is possible that neither ψ_1 nor ψ_2 is true with respect to Δ .

Under the restrictor view of conditionals, the two puzzles are not puzzling any more. The discourse domain of the marble puzzle consists of epistemic possibilities concerning the color of the picked marble. *Must* is not an upward monotonic notion relative to classes of epistemic possibilities. So the argument in this puzzle is not sound. The discourse domain of the Sea Battle Puzzle consists of possible futures and *necessity* is not an upward monotonic notion relative to classes of possible futures. Then the argument in this puzzle is not sound either.

Absolute and relative necessity Temporal restrictors shrink discourse domains which consist of *possible* futures. We argue that there are two senses of possibility: *ontological* and *genuine* possibility.

In reality we make decisions to do something or not to do something. So we will not do whatever we are able to do. Call the class of things we are able to do *the ability domain*. Call the ability domain excluding the things we decide to refrain from doing *the intension domain*. A future is ontologically possible if we can realize it by doing things in the ability domain. A future is genuinely possible if we can realize it by doing things in the intension domain.

Here is an example. There will be an exam tomorrow morning and there are two things for a student to do this evening: preparing for the exam and watching a football match. Doing the former will enable the student to pass the exam but doing the latter will not. The student decides to watch the match. In this case, the future where the student passes the exam is just ontologically possible but not genuinely possible.

Respectively, we can distinguish two senses of necessity: *absolute* and *relative* necessity. A proposition is absolutely necessary if it is the case for every ontologically possible future. A proposition is relatively necessary if it is the case for every genuinely possible future.

The second/third premise of the Sea Battle Puzzle, “*if there is a/no sea battle tomorrow, it is necessarily so*”, can be called *the principle of necessity of truth*. As there are two senses of necessity, there are two principles of necessity of truth.

We argue that the principle of absolute necessity of truth does not hold. Assume that a fleet admiral is able to do two things: *a* and *b*. Doing *a* will cause a sea battle tomorrow but doing *b* will not. He decides to do *a*. In this case, it is plausible to say that there will be a sea battle tomorrow. But it is wrong to say that there will be a sea battle tomorrow no matter what the admiral will do in the absolute sense.

We think, however, that the principle of relative necessity of truth does hold. Let Δ be the class of all the choices for which the agent is open. Assume

that it is not relatively necessary that ϕ will be the case. Then there is a b in Δ such that doing b will make ϕ false. In this situation, it is strange to say that ϕ will be the case, as the agent might do b . So if ϕ will be the case, it is relatively necessary that ϕ will be the case.

The first premise of the Sea Battle Puzzle, “*either there will be a sea battle tomorrow or not*”, can be called *the principle of excluded future middle*. We think that this principle is commonly accepted, at least in classical (non-intuitionistic) reasoning. So all the three premises of the puzzle hold if necessity is understood in the relative sense.

However, the conclusion of this puzzle is problematic for both senses of necessity. Again, suppose that doing a will cause a sea battle tomorrow but doing b will not. Assume that the admiral has not made the decision to do a or b yet. Then it is wrong to say that there will necessarily be a sea battle tomorrow, and also wrong to say that there will necessarily be no sea battle tomorrow. So the conclusion of the puzzle is false.

Thus, the Sea Battle Puzzle is only logically problematic for the case of relative necessity, but not (necessarily) for the case of absolute necessity.

Our technical proposal and contributions In this paper we introduce a new system of extended temporal logic LTC to formalise temporal conditionals. This logic can express relative necessity. We will show that the three premises of the Sea Battle Puzzle are valid in this logic but its conclusion is not, thus indicating that the reasoning behind it is not sound with respect to this logic. The logic LTC extends the Nexttime fragment of CTL* with *operators for model updates*, restricting the temporal domain to those future moments where antecedent is still possible to satisfy. As an illustrating example, we demonstrate that a suitably formalised in our logic version of the ‘Sea Battle’ argument is not sound, thereby providing a solution to the problem with the fatalist conclusion that it entails. As discussed earlier, the technical core of the proposed solution is that the underlying reasoning per cases argument no longer applies when these cases are treated not as material implications but as temporal conditionals. On the technical side, we analyze the semantics of LTC and provide a series of reductions of LTC-formulae, first recursively eliminating the dynamic operators and then the path quantifiers in such formulae. Using these reductions we obtain a sound and complete axiomatization for LTC and also eventually reduce its decision problem to that of the modal logic KD.

Thus, we consider the contributions of this work to be two-fold: to develop a logical system for formalization and analysis of temporal conditionals, and to apply it to the clarification and solution of various philosophical logical problems that they present.

Structure of the paper In Section 2 we introduce the logic of temporal conditionals LTC and provide and discuss its formal semantics. In Section 3 we show that the update operator in LTC properly captures temporal conditionals. Section 4 proposes and discusses our formal solution to the Sea Battle Puzzle, based on its formalisation in LTC. In Section 5 we prove that, by means of

effective translations, the model update operator can be eliminated and then the path quantifier and temporal operator in *state formulas* can be replaced by the classical box modality. Using these results, in Section 6 we establish a sound and complete axiomatization of LTC and show its decidability. We conclude with brief remarks in Section 7.

2 The logic for temporal conditionals LTC

We assume that the reader has basic familiarity with modal and branching-time temporal logics, e.g. within [3, Chapters 4, 5, 7.1].

Let Φ_0 be a fixed countable set of atomic propositions and let p range over it. The language Φ_{LTC} of LTC involves Φ_0 , the propositional connectives \top, \neg, \wedge and the modalities \mathbf{A}, \mathbf{X} and $[\cdot]$, where \cdot stands for a formula (see further). Φ_{LTC} and $\Phi_{\mathbf{X}[\cdot]}$, a fragment of Φ_{LTC} , are defined as follows:

$$\begin{aligned}\Phi_{\mathbf{X}[\cdot]} : \psi &::= p \mid \top \mid \neg\psi \mid (\psi \wedge \psi) \mid \mathbf{X}\psi \mid [\psi]\psi \\ \Phi_{\text{LTC}} : \phi &::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \mathbf{X}\phi \mid \mathbf{A}\phi \mid [\psi]\phi\end{aligned}$$

The other propositional connectives: \perp (falsum), $\vee, \rightarrow, \leftrightarrow$ are defined in the usual way. We also define $\mathbf{E}\phi := \neg\mathbf{A}\neg\phi$.

We can naturally distinguish state and path formulae of LTC (just like in the logic CTL*), defined by mutual induction as follows:

$$\begin{aligned}\text{State formulae } \Phi_{\text{LTC}}^s : \quad & \phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid \mathbf{A}\theta \mid [\psi]\phi, \text{ where } \psi \in \Phi_{\text{PC}} \\ \text{Path formulae } \Phi_{\text{LTC}}^p : \quad & \theta ::= \phi \mid \neg\theta \mid (\theta \wedge \theta) \mid \mathbf{X}\theta \mid [\psi]\theta, \text{ where } \psi \in \Phi_{\mathbf{X}[\cdot]}\end{aligned}$$

We define the abbreviation $\Box := \mathbf{A}\mathbf{X}$. What follow are four fragments of Φ_{LTC} which will be used for technical purposes later.

$$\begin{aligned}\Phi_{\text{PC}} : \psi &::= p \mid \top \mid \neg\psi \mid (\psi \wedge \psi) \\ \Phi_{\mathbf{X}} : \psi &::= p \mid \top \mid \neg\psi \mid (\psi \wedge \psi) \mid \mathbf{X}\psi \\ \Phi_{\mathbf{A}\mathbf{X}} : \psi &::= p \mid \top \mid \neg\psi \mid (\psi \wedge \psi) \mid \mathbf{X}\psi \mid \mathbf{A}\phi \\ \Phi_{\Box} : \psi &::= p \mid \top \mid \neg\psi \mid (\psi \wedge \psi) \mid \Box\psi\end{aligned}$$

Let us give some intuition on the logical operators in LTC. Informally, $\mathbf{X}\phi$ means that ϕ will be the case in the next moment. $\mathbf{A}\phi$ means that no matter how the agent will act in the future, ϕ is the case now, that is, ϕ is *necessary*. \mathbf{A} can be viewed as a universal quantifier over possible futures. Later we will see that possible futures in the semantic setting of LTC are *genuinely possible* futures in the intuitive sense. So $\mathbf{A}\phi$ indicates the relative necessity. Respectively, $\mathbf{E}\phi$ states that the agent has a way to act in the future so that ϕ is the case now, that is, ϕ is *possible*⁵.

⁵ It seems strange to say that the agent has a way to act in the future so that ϕ is the case now. Actually this is fine, as whether a sentence involving future is true or not now might be dependent on how the agent will act in the future. Indeed, the truth of “there will be a sea battle tomorrow” depends on what the admirals of the fighting fleets will decide tonight.

Intuitively, $[\psi]\phi$ indicates that given ψ , ϕ is the case. Assume that one's decision to make ψ true will always make ψ true⁶. Then another intended understanding of $[\psi]\phi$ is as follows: the occurrence of ψ in $[\psi]$ represents the action of deciding to make ψ true and $[\psi]\phi$ is read as that ϕ is the case after the agent decides to make ψ true. This will be seen more clearly from the formal semantics.

Now, some technical terminology and notation. Let W be a nonempty set of states and let R be a binary relation on it. A (finite or infinite) sequence $w_0 \dots w_n(\dots)$ of states is called an R -sequence if $w_0 R \dots R w_n(\dots)$. Note that w is an R -sequence, for any $w \in W$. Next, (W, R) is a *tree* if there is a state $r \in W$, called the *root*, such that for any w , then there is a unique R -sequence starting with r and ending with w . It is easy to see that if there is a root, then it is unique and R is irreflexive. Further, R is *serial* if, for any $w \in W$ there is a $u \in W$ such that Rwu . We say that a tree (W, R) is serial if R is.

A serial tree can be understood as a time structure encoding an agent's actions (the transitions) and states in time (the nodes). A branching in the tree is interpreted as a situation where the agent can choose between different possible actions. The seriality corresponds to the fact that the agent can always perform an action at any given time.

Fix a serial tree (W, R) . A finite R -sequence $w_0 \dots w_n$ starting at the root is called a *history* of w_n . For any states w and u , u is a *historical* state of w if there is a R -sequence $u_0 \dots u_n$ such that $0 < n$, $u_0 = u$ and $u_n = w$. If u is a historical state of w , we also say that w is a *future* state of u . Note that no state in a tree can be a historical or future state of itself.

An infinite R -sequence is called a *path*. A path starting at the root is a *timeline*. A path $\pi = w_0 w_1 \dots$ passes through a state x if $x = w_i$ for some i . For any path π , we use $\pi(i)$ to denote the $i + 1$ -th element of π , ${}^i\pi$ the prefix of π to the $i + 1$ -th element, and π^i the suffix of π from the $i + 1$ -th element. For example, if $\pi = w_0 w_1 \dots$, then $\pi(2) = w_2$, ${}^2\pi = w_0 w_1 w_2$ and $\pi^2 = w_2 w_3 \dots$. For any history $w_0 \dots w_n$ and path $u_0 u_1 \dots$, if $w_n = u_0$, let $w_0 \dots w_n \otimes u_0 \dots$ denote the timeline $w_0 \dots w_n u_1 \dots$.

$\mathfrak{M} = (W, R, r, V)$ is a (LTC-)model if (W, R) is a serial tree with r as the root and V is a valuation from Φ_0 to 2^W . Figure 1 illustrates a model.

Now we will define by mutual induction two important semantic concepts: " ϕ being true at a state $\pi(i)$ relative to a timeline π in a model \mathfrak{M} ", denoted $\mathfrak{M}, \pi, i \Vdash \phi$ (Definition 2.1), and "*the result of updating \mathfrak{M} at w with ϕ* ", denoted \mathfrak{M}_w^ϕ (Definition 2.2). The mutual induction is needed because $\mathfrak{M}, \pi, i \Vdash [\phi]\psi$ is defined in terms of $\mathfrak{M}_{\pi(i)}^\phi$ which, in turn, is defined in terms of $\mathfrak{M}, \pi, i \Vdash \phi$.

⁶ In this paper we maintain the assumption that making a decision to do something will always result in this thing being done.




Fig. 1. A fragment of a LTC-model.

Definition 2.1 [Semantics]

- $\mathfrak{M}, \pi, i \Vdash p \iff \pi(i) \in V(p)$
 - $\mathfrak{M}, \pi, i \Vdash \top$
 - $\mathfrak{M}, \pi, i \Vdash \neg\phi \iff \text{not } \mathfrak{M}, \pi, i \Vdash \phi$
 - $\mathfrak{M}, \pi, i \Vdash \phi \wedge \psi \iff \mathfrak{M}, \pi, i \Vdash \phi \text{ and } \mathfrak{M}, \pi, i \Vdash \psi$
 - $\mathfrak{M}, \pi, i \Vdash X\phi \iff \mathfrak{M}, \pi, i + 1 \Vdash \phi$
 - $\mathfrak{M}, \pi, i \Vdash \mathbf{A}\phi \iff \text{for any path } \rho \text{ starting at } \pi(i), \mathfrak{M}, i\pi \otimes \rho, i \Vdash \phi$
 - $\mathfrak{M}, \pi, i \Vdash [\phi]\psi \iff \mathfrak{M}_{\pi(i)}^\phi, \pi, i \Vdash \psi \text{ if } (\mathfrak{M}_{\pi(i)}^\phi, \pi) \text{ is well-defined (see below)}$
- (Note that $\mathfrak{M}, \pi, i \Vdash [\phi]\psi$ holds vacuously if $(\mathfrak{M}_{\pi(i)}^\phi, \pi)$ is not well-defined.)

The truth condition of $\mathbf{A}\phi$ at \mathfrak{M}, π, i can also be equivalently stated as follows: $\mathfrak{M}, \tau, i \Vdash \phi$ for any timeline τ passing through $\pi(i)$. If π is a timeline in \mathfrak{M} and $\pi(i) = w$, we will sometimes write $\mathfrak{M}, \pi, w \Vdash \phi$ instead of $\mathfrak{M}, \pi, i \Vdash \phi$. We say that ϕ is *achievable* at w in \mathfrak{M} if $\mathfrak{M}, \pi, w \Vdash \mathbf{E}\phi$, i.e. $\mathfrak{M}, \pi, w \Vdash \phi$ for some timeline π containing w .

Note that the truth of a state formula at a state, relative to a timeline, is not dependent on the timeline, but this is not so for path formulae. Sometimes, if ϕ is a state formula, we write $\mathfrak{M}, w \Vdash \phi$ without specifying a timeline.

Definition 2.2 [Model updates and well-definedness] Suppose ϕ is achievable at w in \mathfrak{M} . We define the set $X_w^\phi \subseteq W$ as follows: for any $x \in W$, $x \in X_w^\phi$ iff

- (i) x is a future state of w , and
- (ii) there is no timeline ρ passing through w and x such that $\mathfrak{M}, \rho, w \Vdash \phi$.

Now, we define $\mathfrak{M}_w^\phi = (W - X_w^\phi, R', r, V')$ as the *restriction* of \mathfrak{M} to $W - X_w^\phi$, that is, $R' = R \cap (W - X_w^\phi)^2$ and $V'(p) = V(p) \cap (W - X_w^\phi)$ for each $p \in \Phi_0$.

If ϕ is not achievable at w , we declare \mathfrak{M}_w^ϕ undefined because in this case all successors of w are in X_w^ϕ and therefore R' is not serial.

If π is a timeline in \mathfrak{M} , such that $\mathfrak{M}_{\pi(i)}^\phi$ is defined, and π belongs to $\mathfrak{M}_{\pi(i)}^\phi$, then we say that $(\mathfrak{M}_{\pi(i)}^\phi, \pi)$ is *well-defined*.

Suppose that ϕ is achievable at w . It can be easily verified that $r \notin X_w^\phi$ and X_w^ϕ is closed under R . So $(W - X_w^\phi, R')$ is still a tree with r as the root. It can

also be verified that R' is serial when \mathfrak{M}_w^ϕ is defined. Then \mathfrak{M}_w^ϕ is a LTC-model. Later, we will see that a timeline π in \mathfrak{M} is a timeline in $\mathfrak{M}_{\pi(i)}^\phi$ iff $\mathfrak{M}, \pi, i \models \phi$.

Updating \mathfrak{M} with ϕ at w means to shrink \mathfrak{M} by removing the states in X_w^ϕ . The set X_w^ϕ can be understood as follows. Assume that the agent is at w and decides to make ϕ true. After the decision is made, some future states are not possible anymore. A state becomes impossible if, when the agent travels in time to it, it would no longer be possible to make ϕ true at w , no matter what happens afterwards. X_w^ϕ is the collection of these states. Figure 2 illustrates how a formula updates a model.




Fig. 2. This figure illustrates how a model is updated and reduced by a formula. The model on the right is the result of updating the one on the left at w_0 with $Xq \wedge XX-s$.

A formula ϕ of LTC is *valid* (resp., *satisfiable*) if $\mathfrak{M}, \pi, i \models \phi$ for any (resp., for some) \mathfrak{M}, π and i . Note that a path formula ϕ is valid (resp., satisfiable) iff the state formula $\mathbf{A}\phi$ is valid (resp., $\mathbf{E}\phi$ is satisfiable). A set Γ of formulae *entails* a formula ϕ , denoted $\Gamma \models \phi$, if for any \mathfrak{M}, π and i , if $\mathfrak{M}, \pi, i \models \psi$ for each $\psi \in \Gamma$, then $\mathfrak{M}, \pi, i \models \phi$. We write $\psi_1, \dots, \psi_n \models \phi$ for $\{\psi_1, \dots, \psi_n\} \models \phi$. As usual, we also say that ϕ and ψ are *equivalent*, denoted $\psi \equiv \phi$, if $\psi \models \phi$ and $\phi \models \psi$. Note that, as expected, \equiv is a congruence with respect to all operators in the language of LTC, incl. $[\cdot]$: if $\phi, \psi \in \Phi_{X[\cdot]}$ and $\phi \equiv \psi$, then $[\phi]\chi \equiv [\psi]\chi$ for any $\chi \in \Phi_{\text{LTC}}$, and if $\chi, \theta \in \Phi_{\text{LTC}}$ and $\chi \equiv \theta$, then $[\phi]\chi \equiv [\phi]\theta$ for any $\phi \in \Phi_{X[\cdot]}$.

3 Restrictors work as intended

The update with ϕ at a state in a model restricts the class of possible futures starting at that state. The following theorem indicates that it restricts the class of possible futures exactly as we wish: it excludes those possible futures which do not satisfy ϕ .

For any ϕ in Φ_{LTC} , we define *the temporal depth* ϕ^{td} (intuitively, ϕ^{td} indicates how far in the future ϕ can see) and *the combined depth* ϕ^{cd} of ϕ , as follows:

$$\begin{array}{ll}
 p^{td} = 0 & p^{cd} = 0 \\
 \top^{td} = 0 & \top^{cd} = 0 \\
 (\neg\phi)^{td} = \phi^{td} & (\neg\phi)^{cd} = \phi^{cd} \\
 (\phi \wedge \psi)^{td} = \max\{\phi^{td}, \psi^{td}\} & (\phi \wedge \psi)^{cd} = \max\{\phi^{cd}, \psi^{cd}\} \\
 (X\phi)^{td} = \phi^{td} + 1 & (X\phi)^{cd} = \phi^{cd} + 1 \\
 (\mathbf{A}\phi)^{td} = \phi^{td} & (\mathbf{A}\phi)^{cd} = \phi^{cd} + 1 \\
 ([\phi]\psi)^{td} = \max\{\phi^{td}, \psi^{td}\} & ([\phi]\psi)^{cd} = \max\{\phi^{cd}, \psi^{cd}\}
 \end{array}$$

Lemma 3.1 *Let ϕ be a formula in $\Phi_{X[\cdot]}$, \mathfrak{M} be a model and w be a state in it.*

- (a) *For any natural number $n \geq \phi^{td}$ and any pair of timelines π and τ of \mathfrak{M} passing through w and sharing the same n states after w , it holds that $\mathfrak{M}, \pi, w \Vdash \phi$ iff $\mathfrak{M}, \tau, w \Vdash \phi$.*
- (b) *For any timeline π of \mathfrak{M} passing through w , $(\mathfrak{M}_w^\phi, \pi)$ is well-defined iff $\mathfrak{M}, \pi, w \Vdash \phi$.*

Proof. We will prove both claims simultaneously by induction on ϕ . The base cases $\phi = p$ and $\phi = \top$ are straightforward. Suppose that both claims hold for all subformulae of ϕ . We spell out here only the two non-trivial cases.

Case $\phi = X\psi$.

(a) Let $n \geq (X\psi)^{td}$ and π and τ be two timelines of \mathfrak{M} passing through w and sharing the same n states after w . Let $\pi(i) = \tau(i) = w$. We have the following equivalences:

$$\begin{array}{l}
 \mathfrak{M}, \pi, i \Vdash X\psi \\
 \Leftrightarrow \mathfrak{M}, \pi, i + 1 \Vdash \psi \\
 * \Leftrightarrow \mathfrak{M}, \tau, i + 1 \Vdash \psi \\
 \Leftrightarrow \mathfrak{M}, \tau, i \Vdash X\psi
 \end{array}$$

The equivalence marked by $*$ holds by the inductive hypothesis for (a) applied to ψ , as π and τ share the same $n - 1$ states after $\pi(i + 1)$ and $\psi^{td} \leq n - 1$.

(b) Suppose π is a timeline in \mathfrak{M} passing through w and $\pi(i) = w$. Note that $(\mathfrak{M}_w^{X\psi}, \pi)$ is well-defined iff $\mathfrak{M}_w^{X\psi}$ is defined and π is a path in it. Clearly if $\mathfrak{M}, \pi, w \Vdash X\psi$ then $\mathfrak{M}_w^{X\psi}$ is defined. To show that $(\mathfrak{M}_w^{X\psi}, \pi)$ is well-defined iff $\mathfrak{M}, \pi, w \Vdash X\psi$, it now suffices to show that, assuming $\mathfrak{M}_w^{X\psi}$ is defined, π is a path in $\mathfrak{M}_w^{X\psi}$ iff $\mathfrak{M}, \pi, w \Vdash X\psi$. Assume $\mathfrak{M}_w^{X\psi}$ is defined. Then:

$$\begin{array}{l}
 \pi \text{ is a path in } \mathfrak{M}_w^{X\psi} \\
 \Leftrightarrow \text{for any } u \text{ of } \pi, \text{ if } u \text{ is a future state of } w, \text{ then there is a timeline } \\
 \tau \text{ passing through } u \text{ such that } \mathfrak{M}, \tau, i \Vdash X\psi, \text{ i.e. } \mathfrak{M}, \tau, i + 1 \Vdash \psi \\
 * \Leftrightarrow \mathfrak{M}, \pi, i + 1 \Vdash \psi \\
 \Leftrightarrow \mathfrak{M}, \pi, i \Vdash X\psi
 \end{array}$$

Here is why the marked equivalence holds. The direction from the right to the left is clear. Assume the left. Let n be a natural number such that $\psi^{td} \leq n$. Let $u = \pi(i + 1 + n)$. Then there is a timeline τ passing through u such that $\mathfrak{M}, \tau, i + 1 \Vdash \psi$. Note that π and τ share the same n states after $\pi(i + 1)$. By the inductive hypothesis for (a) applied to ψ , we have $\mathfrak{M}, \pi, i + 1 \Vdash \psi$.

Case $\phi = [\psi]\chi$.

(a) Let $n \geq ([\psi]\chi)^{td}$ and π and τ be two timelines of \mathfrak{M} passing through w and sharing the same n states after w . We have the following equivalences:

$$\begin{aligned} & \mathfrak{M}, \pi, w \Vdash [\psi]\chi \\ \Leftrightarrow & (\mathfrak{M}_w^\psi, \pi) \text{ is not well-defined or } \mathfrak{M}_w^\psi, \pi, w \Vdash \chi \\ * \Leftrightarrow & \mathfrak{M}, \pi, w \not\Vdash \psi \text{ or } \mathfrak{M}_w^\psi, \pi, w \Vdash \chi \\ * \Leftrightarrow & \mathfrak{M}, \tau, w \not\Vdash \psi \text{ or } \mathfrak{M}_w^\psi, \tau, w \Vdash \chi \\ * \Leftrightarrow & (\mathfrak{M}_w^\psi, \tau) \text{ is not well-defined or } \mathfrak{M}_w^\psi, \tau, w \Vdash \chi \\ \Leftrightarrow & \mathfrak{M}, \tau, w \Vdash [\psi]\chi \end{aligned}$$

The first and third equivalence marked by * holds by the inductive hypothesis for (b) applied to ψ . The second equivalence marked by * holds by the inductive hypothesis for (a) applied to ψ , as $\psi^{td} \leq n$.

(b) Let π be a timeline in \mathfrak{M} passing through w and $\pi(i) = w$. Again, note that $(\mathfrak{M}_w^{[\psi]\chi}, \pi)$ is well-defined iff $\mathfrak{M}_w^{[\psi]\chi}$ is defined and π is a path in it. Also note that if $\mathfrak{M}, \pi, w \Vdash [\psi]\chi$, then $\mathfrak{M}_w^{[\psi]\chi}$ is defined. To show that $(\mathfrak{M}_w^{[\psi]\chi}, \pi)$ is well-defined iff $\mathfrak{M}, \pi, w \Vdash [\psi]\chi$, it now suffices to show that, assuming $\mathfrak{M}_w^{[\psi]\chi}$ is defined, π is a path in $\mathfrak{M}_w^{[\psi]\chi}$ iff $\mathfrak{M}, \pi, w \Vdash [\psi]\chi$. Assume that $\mathfrak{M}_w^{[\psi]\chi}$ is defined. Then:

$$\begin{aligned} & \pi \text{ is a path in } \mathfrak{M}_w^{[\psi]\chi} \\ \Leftrightarrow & \text{for any } u \text{ of } \pi, \text{ if } u \text{ is a future state of } w, \text{ then there is a timeline} \\ & \tau \text{ passing through } u \text{ such that } \mathfrak{M}, \tau, w \Vdash [\psi]\chi \\ \Leftrightarrow & \text{for any } u \text{ of } \pi, \text{ if } u \text{ is a future state of } w, \text{ then there is a timeline} \\ & \tau \text{ passing through } u \text{ such that if } (\mathfrak{M}_w^\psi, \tau) \text{ is well-defined, then} \\ & \mathfrak{M}_w^\psi, \tau, w \Vdash \chi \\ * \Leftrightarrow & \text{for any } u \text{ of } \pi, \text{ if } u \text{ is a future state of } w, \text{ then there is a timeline} \\ & \tau \text{ passing through } u \text{ such that if } \mathfrak{M}, \tau, w \Vdash \psi, \text{ then } \mathfrak{M}_w^\psi, \tau, w \Vdash \chi \\ * \Leftrightarrow & \text{if } \mathfrak{M}, \pi, w \Vdash \psi, \text{ then } \mathfrak{M}_w^\psi, \pi, w \Vdash \chi \\ \Leftrightarrow & \text{if } (\mathfrak{M}_w^\psi, \pi) \text{ is well-defined, then } \mathfrak{M}_w^\psi, \pi, w \Vdash \chi \\ \Leftrightarrow & \mathfrak{M}, \pi, w \Vdash \chi \end{aligned}$$

The first equivalence marked by * holds by the inductive hypothesis for (b) applied to ψ . Here is why the second *-marked equivalence holds. The direction from the right to the left is clear. Assume the left. Let n be a natural number such that $\psi^{td}, \chi^{td} \leq n$. Let $u = \pi(i+n)$. Then there is a timeline τ passing through u such that if $\mathfrak{M}, \tau, w \Vdash \psi$, then $\mathfrak{M}_w^\psi, \tau, w \Vdash \chi$. Note that π and τ share the same n elements after w . By the inductive hypothesis for (a) applied to ψ and χ , we have that if $\mathfrak{M}, \pi, w \Vdash \psi$, then $\mathfrak{M}_w^\psi, \pi, w \Vdash \chi$. \square

The next corollary follows immediately from Lemma 3.1.

Corollary 3.2 *Let \mathfrak{M} be a model, w a state and ϕ a formula of $\Phi_{X[\cdot]}$ achievable at w (and therefore \mathfrak{M}_w^ϕ is defined). Then any timeline π in \mathfrak{M} passing through w is also a timeline in \mathfrak{M}_w^ϕ iff $\mathfrak{M}, \pi, w \Vdash \phi$.*

The following lemma states that $\Phi_{X[\cdot]}$ can be effectively reduced to Φ_X .

Lemma 3.3 *For any ϕ in $\Phi_{X[\cdot]}$, there is a ψ in Φ_X equivalent to ϕ .*

Proof. We first show that for any $\psi \in \Phi_X$, $[\phi]\psi$ is equivalent to $\phi \rightarrow \psi$. Note that the truth of the formulae in Φ_X at a state relative to a path in a model is completely determined by the information of the state and the path. So for any $\psi \in \Phi_X$, $\mathfrak{M}, \pi, w \Vdash \psi$ iff $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi$. Let ψ be in Φ_X . Then: $\mathfrak{M}, \pi, w \Vdash [\phi]\psi \Leftrightarrow$ if $\mathfrak{M}, \pi, w \Vdash \phi$ then $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi \Leftrightarrow$ if $\mathfrak{M}, \pi, w \Vdash \phi$ then $\mathfrak{M}, \pi, w \Vdash \psi \Leftrightarrow \mathfrak{M}, \pi, w \Vdash \phi \rightarrow \psi$. For any χ in $\Phi_{X[\cdot]}$, by use of the previous result, we can eliminate step-by-step the update operator $[\cdot]$ in χ . \square

Next theorem points out some connections between $[\phi]\psi$ and $\phi \rightarrow \psi$.

Theorem 3.4 $[\phi]\psi$ collapses to $\phi \rightarrow \psi$ if either of the following holds:

- (a) ϕ is a state formula;
- (b) ψ is in $\Phi_{X[\cdot]}$.

Proof. (a) Suppose ϕ is a state formula. Let \mathfrak{M} be a model and w a state in it. Then ϕ is either true or false at w . Suppose first that ϕ is true at w . Then the update with ϕ at w does not change \mathfrak{M} . Then for any timeline π passing through w , $[\phi]\psi$ is true at w relative to π iff ψ – which is now equivalent to $\phi \rightarrow \psi$ – is true at w relative to π . Now, suppose ϕ is false at w . Then the update with ϕ at w fails, so both $[\phi]\psi$ and $\phi \rightarrow \psi$ are trivially true at w relative to any timeline.

(b) Suppose ψ is in $\Phi_{X[\cdot]}$. By Lemma 3.3, there is a $\psi' \in \Phi_X$ equivalent to ψ . Then $[\phi]\psi$ is equivalent to $[\phi]\psi'$. By the proof for Lemma 3.3, $[\phi]\psi'$ is equivalent to $\phi \rightarrow \psi'$. \square

Note that $[\phi]\psi$ does not generally collapse to $\phi \rightarrow \psi$ if ϕ is not a state formula and ψ contains the operator **A**. A typical example is $[Xs]AXs$: e.g., if there is a sea battle tomorrow, it is necessarily so.

4 Our solution to the Sea Battle Puzzle

We first show that the principle of relative necessity of truth, formalised in the logic as $[\phi]A\phi$, is valid.

Lemma 4.1 Let $(\mathfrak{M}_w^\phi, \pi)$ be well-defined. For any ψ in $\Phi_{X[\cdot]}$, $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi$ iff $\mathfrak{M}, \pi, w \Vdash \psi$.

Proof. By Lemma 3.3, $\Phi_{X[\cdot]}$ can be reduced to Φ_X . As mentioned in the proof for Lemma 3.3, the truth of the formulae in Φ_X at a state relative to a timeline in a model is completely determined by the information of the state and the timeline. Then the claim follows. \square

Theorem 4.2 $[\phi]A\phi$ is valid, where $\phi \in \Phi_{X[\cdot]}$.

Proof. Suppose $\mathfrak{M}, \pi, w \not\Vdash [\phi]A\phi$. This means that $(\mathfrak{M}_w^\phi, \pi)$ is well-defined but $\mathfrak{M}_w^\phi, \pi, w \not\Vdash A\phi$. The latter condition entails the existence of a timeline τ in \mathfrak{M}_w^ϕ passing through w such that $\mathfrak{M}_w^\phi, \tau, w \not\Vdash \phi$. By Lemma 4.1, $\mathfrak{M}, \tau, w \not\Vdash \phi$. By Corollary 3.2, τ is not a timeline in \mathfrak{M}_w^ϕ – a contradiction. \square

Let s denote the proposition “there is a sea battle”. The Sea Battle Puzzle

can now be formalised as

$$Xs \vee X\neg s, [Xs]AXs, [X\neg s]AX\neg s \models AXs \vee AX\neg s.$$

It is easy to show that $X\phi \vee X\neg\phi$ is valid. The validity of the other two premises follows from Theorem 4.2. On the other hand, it is also easy to see that $AXs \vee AX\neg s$ is not valid. Therefore, the logical argument behind the Puzzle of Sea Battle is not sound.

Remark 4.3 There have been various proposed solutions to the Sea Battle Puzzle in the literature, including Łukasiewicz’s three-valued logic [12], Prior’s Peircean temporal logic [15], Prior’s Ockhamist temporal logic [15], the true futurist theory [13], the supervaluationist theory [17] and the relativist theory [9]. These solutions accept the validity of the argument and argue that not all its premises are true. They focus on the following issue: how do we ascribe truth values to statements such as “there will be a sea battle tomorrow”? These statements are called in the literature *future contingents*, as they are about the future but do not have an absolute sense. In the first two solutions mentioned above, the principle of excluded future middle fails. In others, the principle of necessity of truth fails. Except for Łukasiewicz’s three-valued logic, the other solutions use branching time models. We refer to [14] for detailed comparison between these solutions.

5 Expressivity of LTC

In this section, we show two results. Firstly, the update operator $[\phi]$ can be eliminated from any formula in Φ_{LTC} . So Φ_{LTC} can be reduced to Φ_{XA} . Secondly, when restricted to state formulas, Φ_{XA} are equally expressive with Φ_{\square} . Later we will see that the two results enable us to obtain a complete axiomatization for LTC.

To eliminate $[\phi]$ from $[\phi]\psi$, the strategy is to ‘massage’ $[\phi]$ into ψ , deeper and deeper, by applying some valid reduction principles, until it meets atomic propositions and dissolves. The only difficulty arises when $[\phi]$ meets the operator X . To handle this, some preliminary treatment of ϕ is needed.

Recall that Φ_{PC} is the set of purely propositional formulae.

Lemma 5.1 *For any ϕ in Φ_X , there are ψ_1, \dots, ψ_n in Φ_{PC} and χ_1, \dots, χ_n in Φ_X such that ϕ is equivalent to $(\psi_1 \vee X\chi_1) \wedge \dots \wedge (\psi_n \vee X\chi_n)$.*

Proof. The claim follows easily by firstly transforming ϕ to a CNF, where all subformulae beginning with X are treated as literals, and then applying the valid equivalences $\neg X\phi \equiv X\neg\phi$ and $X\phi \vee X\psi \equiv X(\phi \vee \psi)$. \square

Lemma 5.2 *For all $\phi, \psi \in \Phi_{X[\]}$ and $\chi \in \Phi_{\text{LTC}}$, $[\phi \wedge \psi]\chi \leftrightarrow [\phi][\psi]\chi$ is valid.*

Proof. Firstly, we show that $\mathfrak{M}_w^{\phi \wedge \psi}$ is defined iff $(\mathfrak{M}_w^\phi)^\psi$ is defined.

Assume that $\mathfrak{M}_w^{\phi \wedge \psi}$ is defined, which implies $\mathfrak{M}, \pi, w \Vdash \phi \wedge \psi$ for some timeline π passing through w . Then $\mathfrak{M}, \pi, w \Vdash \phi$ and $\mathfrak{M}, \pi, w \Vdash \psi$. By Lemma 3.1, \mathfrak{M}_w^ϕ is defined and π is a timeline in \mathfrak{M}_w^ϕ . By Lemma 4.1, $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi$ and thus also $(\mathfrak{M}_w^\phi)^\psi$ is defined.

Now, assume that $(\mathfrak{M}_w^\phi)^\psi$ is defined, and thus $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi$ for some timeline π in \mathfrak{M}_w^ϕ passing through w . Moreover, by Theorem 4.2, $\mathfrak{M}, \pi, w \Vdash [\phi]\mathbf{A}\phi$. As $(\mathfrak{M}_w^\phi, \pi)$ is well-defined, $\mathfrak{M}_w^\phi, \pi, w \Vdash \mathbf{A}\phi$, hence $\mathfrak{M}_w^\phi, \pi, w \Vdash \phi$. Thus, $\mathfrak{M}_w^\phi, \pi, w \Vdash \phi \wedge \psi$ and consequently, by Lemma 4.1, $\mathfrak{M}, \pi, w \Vdash \phi \wedge \psi$. So, $\mathfrak{M}_w^{\phi \wedge \psi}$ is defined, as wanted.

Secondly, we show $\mathfrak{M}_w^{\phi \wedge \psi} = (\mathfrak{M}_w^\phi)^\psi$. It suffices to show that $\mathfrak{M}_w^{\phi \wedge \psi}$ and $(\mathfrak{M}_w^\phi)^\psi$ have the same timelines passing through w . By Lemma 3.1 and Lemma 4.1, the following holds:

- π is a timeline in $\mathfrak{M}_w^{\phi \wedge \psi}$ through w
- $\Leftrightarrow \mathfrak{M}, \pi, w \Vdash \phi \wedge \psi$
- $\Leftrightarrow \pi$ is a timeline in \mathfrak{M}_w^ϕ through w and $\mathfrak{M}, \pi, w \Vdash \psi$
- $\Leftrightarrow \pi$ is a timeline in \mathfrak{M}_w^ϕ through w and $\mathfrak{M}_w^\phi, \pi, w \Vdash \psi$
- $\Leftrightarrow \pi$ is a timeline in $(\mathfrak{M}_w^\phi)^\psi$ through w

Consequently,

- $\mathfrak{M}, \pi, w \Vdash [\phi \wedge \psi]\chi$
- \Leftrightarrow if $(\mathfrak{M}_w^{\phi \wedge \psi}, \pi)$ is well-defined, then $\mathfrak{M}_w^{\phi \wedge \psi}, \pi, w \Vdash \chi$
- \Leftrightarrow if $\mathfrak{M}_w^{\phi \wedge \psi}$ is defined and π is a timeline in $\mathfrak{M}_w^{\phi \wedge \psi}$, then $\mathfrak{M}_w^{\phi \wedge \psi}, \pi, w \Vdash \chi$
- \Leftrightarrow if $(\mathfrak{M}_w^\phi)^\psi$ is defined and π is a timeline in $(\mathfrak{M}_w^\phi)^\psi$, then $(\mathfrak{M}_w^\phi)^\psi, \pi, w \Vdash \chi$
- \Leftrightarrow if \mathfrak{M}_w^ϕ is defined, π is a timeline in \mathfrak{M}_w^ϕ , $(\mathfrak{M}_w^\phi)^\psi$ is defined and π is a timeline in $(\mathfrak{M}_w^\phi)^\psi$, then $(\mathfrak{M}_w^\phi)^\psi, \pi, w \Vdash \chi$
- \Leftrightarrow if $(\mathfrak{M}_w^\phi, \pi)$ is well-defined and $((\mathfrak{M}_w^\phi)^\psi, \pi)$ is well-defined, then $(\mathfrak{M}_w^\phi)^\psi, \pi, w \Vdash \chi$
- \Leftrightarrow if $(\mathfrak{M}_w^\phi, \pi)$ is well-defined, then if $((\mathfrak{M}_w^\phi)^\psi, \pi)$ is well-defined, then $(\mathfrak{M}_w^\phi)^\psi, \pi, w \Vdash \chi$
- \Leftrightarrow if $(\mathfrak{M}_w^\phi, \pi)$ is well-defined, then $\mathfrak{M}_w^\phi, \pi, w \Vdash [\psi]\chi$
- $\Leftrightarrow \mathfrak{M}, \pi, w \Vdash [\phi][\psi]\chi$

□

In the sequel, $[\phi \wedge \psi]\chi \leftrightarrow [\phi][\psi]\chi$ is called the axiom $\mathbf{Ax}_{[\wedge]}$.

Lemma 5.3 *Let $\phi \in \Phi_{X[\cdot]}$ and $\mathfrak{M}, \pi, i \Vdash X\phi$. Then the generated submodels of $\mathfrak{M}_{\pi(i)}^{X\phi}$ and $\mathfrak{M}_{\pi(i+1)}^\phi$ at $\pi(i+1)$ coincide.*

Proof. As $\mathfrak{M}, \pi, i \Vdash X\phi$, $\mathfrak{M}_{\pi(i)}^{X\phi}$ is defined and $\mathfrak{M}, \pi, i+1 \Vdash \phi$. Then $\mathfrak{M}_{\pi(i+1)}^\phi$ is defined. By Corollary 3.2, for every timeline τ passing through $\pi(i+1)$, we have the following equivalences: τ is a timeline in $\mathfrak{M}_{\pi(i)}^{X\phi} \Leftrightarrow \mathfrak{M}, \tau, i \Vdash X\phi \Leftrightarrow \mathfrak{M}, \tau, i+1 \Vdash \phi \Leftrightarrow \tau$ is a timeline in $\mathfrak{M}_{\pi(i+1)}^\phi$. From this the claim follows easily. □

Lemma 5.4 *The following equivalences are valid, where $\phi \in \Phi_{PC}$ and $X\psi \in \Phi_X$:*

$$\begin{aligned} \mathbf{Ax}_{[\cdot]p} &: [\phi \vee X\psi]p \leftrightarrow ((\phi \vee X\psi) \rightarrow p) \\ \mathbf{Ax}_{[\cdot]\top} &: [\phi \vee X\psi]\top \leftrightarrow \top \\ \mathbf{Ax}_{[\cdot]\neg} &: [\phi \vee X\psi]\neg\chi \leftrightarrow ((\phi \vee X\psi) \rightarrow \neg[\phi \vee X\psi]\chi) \\ \mathbf{Ax}_{[\cdot]\wedge} &: [\phi \vee X\psi](\chi \wedge \xi) \leftrightarrow ([\phi \vee X\psi]\chi \wedge [\phi \vee X\psi]\xi) \\ \mathbf{Ax}_{[\cdot]X} &: [\phi \vee X\psi]X\chi \leftrightarrow ((\phi \rightarrow X\chi) \wedge (\neg\phi \rightarrow X[\psi]\chi)) \\ \mathbf{Ax}_{[\cdot]A} &: [\phi \vee X\psi]A\chi \leftrightarrow ((\phi \vee X\psi) \rightarrow A[\phi \vee X\psi]\chi) \end{aligned}$$

Proof. The first 4 equivalences are straightforward.

$\mathbf{Ax}_{[\cdot]X}$. Assume $\mathfrak{M}, \pi, i \not\models [\phi \vee X\psi]X\chi$. Then $(\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi}, \pi)$ is well-defined and $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi}, \pi, i \not\models X\chi$. By Lemma 3.1, $\mathfrak{M}, \pi, i \models \phi \vee X\psi$. Assume $\mathfrak{M}, \pi, i \models \phi$. As ϕ is in Φ_{PC} , $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi} = \mathfrak{M}$. Then $\mathfrak{M}, \pi, i \not\models X\chi$. Then $\mathfrak{M}, \pi, i \not\models \phi \rightarrow X\chi$. Assume $\mathfrak{M}, \pi, i \models \neg\phi$. As ϕ is in Φ_{PC} , $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi} = \mathfrak{M}_{\pi(i)}^{X\psi}$. Then $\mathfrak{M}_{\pi(i)}^{X\psi}, \pi, i \not\models X\chi$. Then $\mathfrak{M}_{\pi(i)}^{X\psi}, \pi, i+1 \not\models \chi$. By Lemma 5.3, $\mathfrak{M}_{\pi(i+1)}^{\psi}, \pi, i+1 \not\models \chi$. Then $\mathfrak{M}, \pi, i+1 \not\models [\psi]\chi$. Hence $\mathfrak{M}, \pi, i \not\models X[\psi]\chi$. Therefore, $\mathfrak{M}, \pi, i \not\models \neg\phi \rightarrow X[\psi]\chi$.

Assume $\mathfrak{M}, \pi, i \not\models \phi \rightarrow X\chi$. Then $\mathfrak{M}, \pi, i \models \phi$ and $\mathfrak{M}, \pi, i \not\models X\chi$. As ϕ is in Φ_{PC} , $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi} = \mathfrak{M}$. Then $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi}, \pi, i \not\models X\chi$. Then $\mathfrak{M}, \pi, i \not\models [\phi \vee X\psi]X\chi$. Assume $\mathfrak{M}, \pi, i \not\models \neg\phi \rightarrow X[\psi]\chi$. Then $\mathfrak{M}, \pi, i \models \neg\phi$ and $\mathfrak{M}, \pi, i \not\models X[\psi]\chi$. As ϕ is in Φ_{PC} , $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi} = \mathfrak{M}_{\pi(i)}^{X\psi}$. Then $\mathfrak{M}_{\pi(i)}^{X\psi}, \pi, i+1 \not\models [\psi]\chi$. Then $\mathfrak{M}_{\pi(i+1)}^{\psi}, \pi, i+1 \not\models \chi$. By Lemma 5.3, $\mathfrak{M}_{\pi(i)}^{X\psi}, \pi, i+1 \not\models \chi$. Then $\mathfrak{M}_{\pi(i)}^{X\psi}, \pi, i \not\models X\chi$. Hence $\mathfrak{M}_{\pi(i)}^{\phi \vee X\psi}, \pi, i \not\models X\chi$. Therefore, $\mathfrak{M}, \pi, i \not\models [\phi \vee X\psi]X\chi$.

$\mathbf{Ax}_{[\cdot]A}$. Let $\alpha = \phi \vee X\psi$. Assume $\mathfrak{M}, \pi, i \not\models \alpha \rightarrow A[\alpha]\chi$. Then $\mathfrak{M}, \pi, i \models \alpha$ but $\mathfrak{M}, \pi, i \not\models A[\alpha]\chi$. By Lemma 3.1, π is in $\mathfrak{M}_{\pi(i)}^{\alpha}$. Then there is a timeline τ passing through $\pi(i)$ such that $\mathfrak{M}, \tau, i \not\models [\alpha]\chi$. Then $\mathfrak{M}_{\pi(i)}^{\alpha}, \tau, i \not\models \chi$. Then $\mathfrak{M}_{\pi(i)}^{\alpha}, \tau, i \not\models A\chi$. Then $\mathfrak{M}_{\pi(i)}^{\alpha}, \pi, i \not\models A\chi$. Then $\mathfrak{M}, \pi, i \not\models [\alpha]A\chi$.

Assume $\mathfrak{M}, \pi, i \not\models [\alpha]A\chi$. Then $(\mathfrak{M}_{\pi(i)}^{\alpha}, \pi)$ is well-defined and $\mathfrak{M}_{\pi(i)}^{\alpha}, \pi, i \not\models A\chi$. By Lemma 3.1, $\mathfrak{M}, \pi, i \models \alpha$. Then there is a timeline τ passing through $\pi(i)$ in $\mathfrak{M}_{\pi(i)}^{\alpha}$ such that $\mathfrak{M}_{\pi(i)}^{\alpha}, \tau, i \not\models \chi$. Then $\mathfrak{M}, \tau, i \not\models [\alpha]\chi$. Then $\mathfrak{M}, \pi, i \not\models A[\alpha]\chi$. Hence $\mathfrak{M}, \pi, i \not\models \alpha \rightarrow A[\alpha]\chi$. \square

Recall that Φ_{XA} denotes the language generated from Φ_0 under X and A .

Theorem 5.5 *There is an effective translation $\mathbf{t} : \Phi_{LTC} \rightarrow \Phi_{XA}$ preserving formulae up to equivalence, i.e. such that $\phi \equiv \mathbf{t}(\phi)$ for every $\phi \in \Phi_{LTC}$.*

Proof. Define a language $\Phi_{XA[\vee]}$ as follows, where $\eta \in \Phi_{PC}$ and $\chi \in \Phi_X$:

$$\phi ::= p \mid \top \mid \neg\phi \mid (\phi \wedge \phi) \mid X\phi \mid A\phi \mid [\eta \vee X\chi]\phi$$

There is a translation $f : \Phi_{LTC} \rightarrow \Phi_{XA[\vee]}$ preserving formulae up to equivalence. The reason is as follows. Take a formula $[\phi]\psi$ in Φ_{LTC} . By Lemma 3.3, ϕ is equivalent to some ϕ' in Φ_X . Then $[\phi]\psi$ is equivalent to $[\phi']\psi$. By Lemma 5.1, ϕ' is equivalent to a conjunction $\phi'_1 \wedge \dots \wedge \phi'_n$ for some $\phi'_i = \eta_i \vee X\chi_i$, where $\eta_i \in \Phi_{PC}$ and $\chi_i \in \Phi_X$, for $i = 1, \dots, n$. Then $[\phi']\psi$ is equivalent to $[\phi'_1 \wedge \dots \wedge \phi'_n]\psi$. By Lemma 5.2, $[\phi'_1 \wedge \dots \wedge \phi'_n]\psi$ is equivalent to $[\phi'_1] \dots [\phi'_n]\psi$.

Define a translation $g : \Phi_{\mathbf{XA}[v]} \rightarrow \Phi_{\mathbf{XA}}$ in the following way:

- (i) $p^g = p$
- (ii) $\top^g = \top$
- (iii) $(\neg\psi)^g = \neg\psi^g$
- (iv) $(\psi \wedge \chi)^g = \psi^g \wedge \chi^g$
- (v) $(X\psi)^g = X\psi^g$
- (vi) $(\mathbf{A}\psi)^g = \mathbf{A}\psi^g$
- (vii) (a) $([\eta \vee X\chi]p)^g = ((\eta \vee X\chi) \rightarrow p)^g$
- (b) $([\eta \vee X\chi]\top)^g = \top^g$
- (c) $([\eta \vee X\chi]\neg\psi)^g = ((\eta \vee X\chi) \rightarrow \neg[\eta \vee X\chi]\psi)^g$
- (d) $([\eta \vee X\chi](\psi \wedge \xi))^g = ([\eta \vee X\chi]\psi \wedge [\eta \vee X\chi]\xi)^g$
- (e) $([\eta \vee X\chi]X\psi)^g = ((\eta \rightarrow X\psi) \vee (\neg\eta \rightarrow X[\chi]\psi))^g$
- (f) $([\eta \vee X\chi]\mathbf{A}\psi)^g = ((\eta \vee X\chi) \rightarrow \mathbf{A}[\eta \vee X\chi]\psi)^g$
- (g) $([\eta \vee X\chi][\alpha \vee X\beta]\psi)^g = ([\eta \vee X\chi](\alpha \vee X\beta)\psi)^g$

Now we show by three layers of induction that for any $\phi \in \Phi_{\mathbf{XA}[v]}$, $\phi^g \in \Phi_{\mathbf{XA}}$ and $\phi^g \equiv \phi$. We consider only the former result; with Lemma 5.4, the argument for the latter result is similar. We put the first layer of induction on ϕ . The only non-trivial case is $\phi = ([\eta \vee X\chi]\gamma)$. Then we put the second layer of induction on γ . We can easily go through all the subcases except the one $\gamma = [\alpha \vee X\beta]\psi$. Assume $\gamma = [\alpha \vee X\beta]\psi$. By the inductive hypothesis for the second layer of induction, $([\alpha \vee X\beta]\psi)^g \in \Phi_{\mathbf{XA}}$. Then by the third layer of induction on $([\alpha \vee X\beta]\psi)^g$, we can get $([\eta \vee X\chi]\psi)^g \in \Phi_{\mathbf{XA}}$.

Now, the translation $\mathbf{t} : \Phi_{\text{LTC}} \rightarrow \Phi_{\mathbf{XA}}$ defined as $g \circ f$ satisfies the claim. \square

Example 5.6 This example demonstrates how the translation \mathbf{t} is computed in the non-trivial case of $\mathbf{t}([\phi]\psi)$. It also illustrates the exponential blow-up that the translation causes to the length of the input formula in this case.

$$\begin{aligned}
& [(p \vee Xq) \wedge (r \vee Xs)]Xt \\
& \equiv [p \vee Xq][r \vee Xs]Xt \\
& \equiv [p \vee Xq]((r \rightarrow Xt) \wedge (\neg r \rightarrow X[s]t)) \\
& \equiv [p \vee Xq]((r \rightarrow Xt) \wedge (\neg r \rightarrow X(s \rightarrow t))) \\
& \equiv [p \vee Xq](r \rightarrow Xt) \wedge [p \vee Xq](\neg r \rightarrow X(s \rightarrow t)) \\
& \equiv ([p \vee Xq]r \rightarrow [p \vee Xq]Xt) \wedge ([p \vee Xq]\neg r \rightarrow [p \vee Xq]X(s \rightarrow t)) \\
& \equiv (((p \vee Xq) \rightarrow r) \rightarrow [p \vee Xq]Xt) \wedge (((p \vee Xq) \rightarrow \neg r) \rightarrow [p \vee Xq]X(s \rightarrow t)) \\
& \equiv \left(((p \vee Xq) \rightarrow r) \rightarrow ((p \rightarrow Xt) \wedge (\neg p \rightarrow X[q]t)) \right) \wedge \\
& \quad \left(((p \vee Xq) \rightarrow \neg r) \rightarrow ((p \rightarrow X(s \rightarrow t)) \wedge (\neg p \rightarrow X[q](s \rightarrow t))) \right) \\
& \equiv \left(((p \vee Xq) \rightarrow r) \rightarrow ((p \rightarrow Xt) \wedge (\neg p \rightarrow X(q \rightarrow t))) \right) \wedge
\end{aligned}$$

$$\left(\left((p \vee Xq) \rightarrow \neg r \right) \rightarrow \left((p \rightarrow X(s \rightarrow t)) \wedge (\neg p \rightarrow X(q \rightarrow (s \rightarrow t))) \right) \right)$$

Lemma 5.7 *Let η be a state formula and χ a formula in Φ_{XA} . Then the following equivalences are valid:*

- (i) $\mathbf{A}(\eta \vee \chi) \leftrightarrow (\eta \vee \mathbf{A}\chi)$
- (ii) $\mathbf{A}X\chi \leftrightarrow \mathbf{A}X\mathbf{A}X\chi$

Theorem 5.8 *There is an effective translation \mathbf{u} from the class of state formulas of Φ_{XA} to Φ_{\square} preserving formulae up to equivalence.*

Proof. (Sketch.) Induction on the combined depth ϕ^{cd} of ϕ . When $\phi^{cd} = 0$ then ϕ is already in Φ_{\square} . Suppose $\phi^{cd} = k + 1$ and the claim holds for all state formulae θ such that $\theta^{cd} \leq k$.

Now, the formula ϕ is a boolean combination of propositional formulae and formulae of the type $\mathbf{A}\psi$. In the latter, ψ is equivalent to a CNF, with disjunctions of the type $\delta = \alpha \vee \mathbf{E}\theta \vee \mathbf{A}\theta_1 \vee \dots \vee \mathbf{A}\theta_k \vee X\chi$, where $\alpha \in \Phi_{PC}$. Since each of $\alpha, \mathbf{E}\theta, \mathbf{A}\theta_1, \dots, \mathbf{A}\theta_k$ is a state formula of combined depth $\leq k$, by using the inductive hypothesis we can assume that each of these is already replaced by an equivalent formula from Φ_{\square} . Then $\mathbf{A}\psi$ is equivalent to a conjunction of state formulae of the type $\mathbf{A}\delta = \mathbf{A}(\alpha \vee \mathbf{E}\theta \vee \mathbf{A}\theta_1 \vee \dots \vee \mathbf{A}\theta_k \vee X\chi)$, which, by Lemma 5.7, is equivalent to $\alpha \vee \mathbf{E}\theta \vee \mathbf{A}\theta_1 \vee \dots \vee \mathbf{A}\theta_k \vee \mathbf{A}X\chi$. Note that $\chi^{cd} \leq k - 1$. After transforming, likewise, χ to a CNF, distributing \mathbf{A} over the conjunctions in it and pulling out the state subformulae from these disjunctions outside \mathbf{A} , we reduce $\mathbf{A}X\chi$ above to an equivalent boolean combination of formulae from Φ_{\square} and formulae of the type $\mathbf{A}X\chi$. The latter is equivalent to $\mathbf{A}X\mathbf{A}X\chi$ by Lemma 5.7. Now, since $\chi^{cd} \leq k - 2$ then $(\mathbf{A}X\chi)^{cd} \leq k$, we can use the inductive hypothesis to replace $\mathbf{A}X\chi$ with an equivalent formula from Φ_{\square} . By doing that for all such subformulae $\mathbf{A}X\chi$ we obtain an equivalent from Φ_{\square} for $\mathbf{A}X\chi$, and therefore can define such an equivalent $\mathbf{u}(\mathbf{A}\delta)$ for $\mathbf{A}\delta$. Finally, we define $\mathbf{u}(\phi)$ to be the conjunction of all these $\mathbf{u}(\mathbf{A}\delta)$. \square

6 Complete axiomatization of LTC

Here we present a sound and complete axiomatic system \mathbf{AxSys}_{LTC} for the logic LTC, consisting of the following groups of axioms schemes and inference rules:

I. Any complete system of axioms for PC.

II. The usual axiom schemes for X:

$$\mathbf{Ax}_{K(X)}: X(\phi \rightarrow \psi) \rightarrow (X\phi \rightarrow X\psi)$$

$$\mathbf{Ax}_{D(X)}: \neg X\perp$$

$$\mathbf{Ax}_{Fun(X)}: \neg X\phi \rightarrow X\neg\phi$$

II. The S5 axioms for \mathbf{A} , plus the following:

$$\mathbf{Ax}_{St(\mathbf{A})}: \phi \rightarrow \mathbf{A}\phi, \text{ for any state formula } \phi^7.$$

⁷ It suffices to state that axiom for atomic propositions only, and the rest would be derivable,

$\mathbf{Ax}_{\mathbf{AXX}}: \mathbf{AXX}\phi \rightarrow \mathbf{AXAX}\phi$

IV. Axiom schemes for $[\cdot]$: the valid schemes $\mathbf{Ax}_{[\wedge]}$ from Lemma 5.2 and $\mathbf{Ax}_{[\cdot]p}$, $\mathbf{Ax}_{[\cdot]\top}$, $\mathbf{Ax}_{[\cdot]\neg}$, $\mathbf{Ax}_{[\cdot]\wedge}$, $\mathbf{Ax}_{[\cdot]X}$, $\mathbf{Ax}_{[\cdot]A}$ from Lemma 5.4.

V. Inference rules: Modus Ponens (MP) and the necessitation rules:

$$\mathbf{Nec}_A: \frac{\vdash \phi}{\vdash \mathbf{A}\phi}, \quad \mathbf{Nec}_X: \frac{\vdash \phi}{\vdash X\phi}, \quad \mathbf{Nec}_{[\cdot]}: \frac{\vdash \phi}{\vdash [\psi]\phi}, \quad \mathbf{Equiv}_{[\cdot]}: \frac{\vdash \phi \leftrightarrow \psi}{\vdash [\phi]\chi \leftrightarrow [\psi]\chi}$$

The subsystem $\mathbf{AxSys}_{\mathbf{AX}}$ consists of axioms I, II, III, and the rules in V.

Recall that $\Box\phi$ denotes $\mathbf{AX}\phi$ and Φ_{\Box} is the respective fragment of $\Phi_{\mathbf{XA}}$ built from Φ_0 only by using the propositional connectives and \Box . We define the subsystem \mathbf{AxSys}_{\Box} , consisting of the KD axioms for \Box :

$\mathbf{Ax}_{\mathbf{K}(\Box)}: \Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$,

$\mathbf{Ax}_{\mathbf{D}(\Box)}: \neg \Box \perp$,

plus the rules MP and $\mathbf{Nec}_{\Box}: \frac{\vdash \phi}{\vdash \Box\phi}$.

Note that:

- (i) The axioms and rules of \mathbf{AxSys}_{\Box} are derivable in $\mathbf{AxSys}_{\mathbf{AX}}$.
- (ii) \mathbf{AxSys}_{\Box} , being canonical, is sound and complete for the validities in Φ_{\Box} .

Now, we are going to establish a series of equivalent reductions of formulae of LTC, all derivable in $\mathbf{AxSys}_{\text{LTC}}$, which will enable us to eventually reduce the proof of completeness of $\mathbf{AxSys}_{\text{LTC}}$ to the completeness of the much simpler logic \mathbf{AxSys}_{\Box} . We begin with some useful derivations, used further in the proofs.

Lemma 6.1 *The following are derivable in $\mathbf{AxSys}_{\text{LTC}}$:*

- (i) $\vdash_{\text{LTC}} [\perp]\phi \leftrightarrow \top$
- (ii) $\vdash_{\text{LTC}} [\phi]\psi \leftrightarrow (\phi \rightarrow \psi)$, where ϕ is in Φ_{PC}

Proof. (i) By induction on ϕ . The only non-trivial case is $\phi = X\psi$. Since $\vdash_{\text{LTC}} \perp \leftrightarrow (\perp \vee X\perp)$, we obtain $\vdash_{\text{LTC}} [\perp]X\psi \leftrightarrow [\perp \vee X\perp]X\psi$ by $\mathbf{Equiv}_{[\cdot]}$. Then, using $\mathbf{Ax}_{[\cdot]X}$ we derive $\vdash_{\text{LTC}} [\perp \vee X\perp]X\psi \leftrightarrow ((\perp \rightarrow X\psi) \wedge (\neg\perp \rightarrow X[\perp]\psi))$. By the inductive hypothesis, $\vdash_{\text{LTC}} [\perp]\psi \leftrightarrow \top$. Then $\vdash_{\text{LTC}} [\perp \vee X\perp]X\psi \leftrightarrow ((\perp \rightarrow X\psi) \wedge (\neg\perp \rightarrow X\top))$. Therefore, $\vdash_{\text{LTC}} [\perp \vee X\perp]X\psi \leftrightarrow ((\perp \rightarrow X\psi) \wedge (\neg\perp \rightarrow \top))$. Hence, $\vdash_{\text{LTC}} [\perp \vee X\perp]X\psi \leftrightarrow \top$. Therefore, $\vdash_{\text{LTC}} [\perp]X\psi \leftrightarrow \top$.

(ii) By induction on ψ . The only non-trivial cases are $\psi = X\chi$ and $\psi = \mathbf{A}\chi$.

Case $\psi = X\chi$. As $\vdash_{\text{LTC}} \phi \leftrightarrow (\phi \vee X\perp)$, we derive $\vdash_{\text{LTC}} [\phi]X\chi \leftrightarrow [\phi \vee X\perp]X\chi$ by $\mathbf{Equiv}_{[\cdot]}$. Then, $\vdash_{\text{LTC}} [\phi \vee X\perp]X\chi \leftrightarrow ((\phi \rightarrow X\chi) \wedge (\neg\phi \rightarrow X[\perp]\chi))$, by $\mathbf{Ax}_{[\cdot]X}$. By claim (i), $\vdash_{\text{LTC}} [\perp]\phi \leftrightarrow \top$. Hence, $\vdash_{\text{LTC}} [\phi \vee X\perp]X\chi \leftrightarrow ((\phi \rightarrow X\chi) \wedge (\neg\phi \rightarrow X\top))$. Then $\vdash_{\text{LTC}} [\phi \vee X\perp]X\chi \leftrightarrow ((\phi \rightarrow X\chi) \wedge (\neg\phi \rightarrow \top))$, so $\vdash_{\text{LTC}} [\phi \vee X\perp]X\chi \leftrightarrow (\phi \rightarrow X\chi)$.

Case $\psi = \mathbf{A}\chi$. By $\mathbf{Ax}_{[\cdot]A}$ we have $\vdash_{\text{LTC}} [\phi]\mathbf{A}\chi \leftrightarrow (\phi \rightarrow \mathbf{A}[\phi]\chi)$. By the inductive hypothesis, $\vdash_{\text{LTC}} [\phi]\chi \leftrightarrow (\phi \rightarrow \chi)$. Then $\vdash_{\text{LTC}} [\phi]\mathbf{A}\chi \leftrightarrow (\phi \rightarrow$

but we prefer to streamline the deductive system.

$\mathbf{A}(\phi \rightarrow \chi)$). Note that $\phi \in \Phi_{\text{PC}}$. Then $\vdash_{\text{LTC}} \mathbf{A}(\phi \rightarrow \chi) \leftrightarrow (\phi \rightarrow \mathbf{A}\chi)$, hence $\vdash_{\text{LTC}} [\phi]\mathbf{A}\chi \leftrightarrow (\phi \rightarrow \mathbf{A}\chi)$. \square

Lemma 6.2 *For every LTC-formula ϕ , $\vdash_{\text{AxSys}_{\text{LTC}}} \mathbf{t}(\phi) \leftrightarrow \phi$.*

Proof. By induction on ϕ , by applying the axioms for $[\cdot]$, the derivations in Lemma 6.1, and the inference rules of $\text{AxSys}_{\text{LTC}}$, one can formalise the proof of Theorem 5.5 in $\text{AxSys}_{\text{LTC}}$. The details are routine and we leave them out. \square

Lemma 6.3 *For every LTC-formula ϕ , if ϕ is $\text{AxSys}_{\text{LTC}}$ -consistent then $\mathbf{t}(\phi)$ is AxSys_{AX} -consistent.*

Proof. Equivalently (since $\mathbf{t}(\neg\phi) = \neg\mathbf{t}(\phi)$), we have to prove that for every LTC-formula ϕ , if $\vdash_{\text{AxSys}_{\text{AX}}} \mathbf{t}(\phi)$ then $\vdash_{\text{AxSys}_{\text{LTC}}} \phi$. Indeed, take a derivation $\vdash_{\text{AxSys}_{\text{AX}}} \mathbf{t}(\phi)$. It is also a derivation in $\text{AxSys}_{\text{LTC}}$. Append to it the derivation $\vdash_{\text{AxSys}_{\text{LTC}}} \mathbf{t}(\phi) \leftrightarrow \phi$ from Lemma 6.2, to obtain $\vdash_{\text{AxSys}_{\text{LTC}}} \mathbf{t}(\phi) \rightarrow \phi$. Then, by MP we obtain $\vdash_{\text{AxSys}_{\text{LTC}}} \phi$. \square

Lemma 6.4 *Every state formula $\psi \in \Phi_{\text{XA}}$ is provably equivalent in AxSys_{AX} to $\mathbf{u}(\psi) \in \Phi_{\square}$.*

Proof. The equivalences used in defining the function \mathbf{u} in the proof for Theorem 5.8, including those in Lemma 5.7, are all derivable in AxSys_{AX} . \square

Lemma 6.5 *For every formula $\psi \in \Phi_{\text{XA}}$ there is an effectively computable formula $\mathbf{v}(\psi) \in \Phi_{\square}$ such that ψ and $\mathbf{v}(\psi)$ are equally satisfiable and equally consistent in AxSys_{AX} .*

Proof. (Sketch.) Here ‘consistency’ will mean ‘consistency in AxSys_{AX} ’. We will prove the claim for consistency; the proof for satisfiability is fully analogous.

The definition of $\mathbf{v}(\psi)$ extends that of $\mathbf{u}(\phi)$ for state formulae ϕ in Lemma 6.4 as follows. The formula ψ is provably in AxSys_{AX} equivalent to some, hereafter fixed, formula ψ' in DNF, where every disjunct is of the type $\delta = \theta \wedge X\chi$, for some state formula θ (note that χ may also be \top). Then, consistency of ψ is equivalent (provably in AxSys_{AX}) to consistency of some of these disjuncts. Note that $\theta \wedge X\chi$ is consistent iff $\theta \wedge \mathbf{E}X\chi$ is so. Indeed, $\vdash_{\text{AxSys}_{\text{AX}}} \theta \rightarrow \neg X\chi$ iff $\vdash_{\text{AxSys}_{\text{AX}}} \mathbf{A}(\theta \rightarrow \neg X\chi)$ iff $\vdash_{\text{AxSys}_{\text{AX}}} \theta \rightarrow \mathbf{A}\neg X\chi$. Lastly, $\theta \wedge \mathbf{E}X\chi$ is consistent iff $\mathbf{u}(\theta) \wedge \mathbf{u}(\mathbf{E}X\chi)$ is so. Now, if ψ is consistent we define $\mathbf{v}(\psi)$ to be $\mathbf{u}(\theta) \wedge \mathbf{u}(\mathbf{E}X\chi)$ for the first disjunct δ in ψ' that is consistent, else we define $\mathbf{v}(\psi)$ to be \perp . The claim of the lemma follows by construction. \square

Theorem 6.6 *$\text{AxSys}_{\text{LTC}}$ is sound and complete for the logic LTC.*

Proof. The soundness of the axioms for $[\cdot]$ follows from Lemmas 5.2, and 5.4. The soundness of Ax_{AX} is by Lemma 5.7. All others axioms are well-known and proving their soundness, as well as that of the inference rules, is routine.

For the completeness, we need to show that every $\text{AxSys}_{\text{LTC}}$ -consistent formula ϕ is satisfiable in a LTC-model. We prove that by a series of reductions, as follows. Let ϕ $\text{AxSys}_{\text{LTC}}$ -consistent. Then, by Lemma 6.3, $\mathbf{t}(\phi)$ is AxSys_{AX} -consistent. Next, by Lemma 6.5, $\mathbf{v}(\mathbf{t}(\phi))$ is AxSys_{AX} -consistent, and therefore AxSys_{\square} -consistent. Since AxSys_{\square} is complete for the logic KD, it follows that

$\mathbf{v}(\mathbf{t}(\phi))$ is satisfiable, hence $\mathbf{t}(\phi)$ is satisfiable by Lemma 6.5, and therefore ϕ is satisfiable, too, by Theorem 5.5. \square

Theorem 6.7 *The satisfiability problem for LTC is decidable, in EXPSpace.*

Proof. The translations and reductions defined here effectively reduce the satisfiability problem for LTC to that of KD, which is PSPACE-complete (see e.g. [11]). However, as illustrated by Example 5.6, the translation \mathbf{t} can cause an exponential blow-up to the length of the formula, hence the complexity upper bound that we provide based on the results obtained here is EXPSpace. However, the number of different subformulae need not grow exponentially, so it is conceivable that a more refined, on-the-fly application of the translation \mathbf{t} , or an alternative decision method may bring that complexity down. We leave the question of establishing the precise complexity of LTC open. \square

7 Concluding remarks

In this work we have proposed the logic LTC for formalising reasoning with and about temporal conditionals, for which we establish a series of technical results, eventually leading to a sound and complete axiomatization and a decision procedure. We show how formalization on LTC can provide a solution to the Sea Battle Puzzle, by showing that the underlying reasoning per cases argument is not sound in that logic.

The present work suggests a variety of interesting follow-up developments:

- On the conceptual side, we believe that the logic LTC or suitable variations of it can be applied to resolve other philosophical problems arising from reasoning about temporal conditionals. We also intend to adapt that logic for reasoning about temporal counterfactuals, as well as to related type of reasoning, that often occurs in the theory of extensive forms, about the players' behaviour on and off the equilibrium path, non-credible threats, etc., leading, inter alia, to the more refined notion of subgame-perfect equilibrium and myopic equilibrium.
- On the technical side, we note that the Nexttime fragment of the logic CTL*, which is precisely the fragment Φ_{XA} of LTC, seems not to have been studied so far. The present work provides, inter alia, a new complete axiomatization result for it. We intend to extend the present results to extensions of LTC, also involving long-term temporal operators, by relating them with more expressive, yet still manageable fragments of CTL*.
- In particular, an extension of the logic LTC with past and future temporal modalities can likewise resolve the problem raised by the Diodorean Master Argument, by rendering that argument unsound with respect to the resulting semantics.

Acknowledgment

Thanks go to Maria Aloni, Johan van Benthem, Alessandra Marra, Floris Roelofsen, Frank Veltman, and the audience of seminars or workshops at Delft University of Technology, Beijing Normal University and Tsinghua University, for their useful comments and suggestions. Valentin Goranko is grateful to the Department of Philosophy of Beijing Normal University for supporting his visit there, during which part of the present work was done. Thanks are also due to the three anonymous referees for helpful comments and some corrections.

References

- [1] Arlo-Costa, H. and P. Egré, *The logic of conditionals*, in: E. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, 2007.
- [2] Cantwell, J., *Changing the modal context*, *Theoria* **74** (2010), pp. 331–351.
- [3] Demri, S., V. Goranko and M. Lange, “Temporal Logics in Computer Science,” *Cambridge Tracts in Theoretical Computer Science* **58**, Cambridge University Press, 2016.
- [4] Edgington, D., *Indicative conditionals*, in: E. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, 2014.
- [5] Garrett, B., “What Is This Thing Called Metaphysics?” Taylor & Francis, 2017.
- [6] Kolodny, N. and J. MacFarlane, *Ifs and oughts*, *The Journal of Philosophy* **107** (2010), pp. 115–143.
- [7] Kratzer, A., *Conditionals*, *Chicago Linguistics Society* **22** (1986), pp. 1–15.
- [8] Lewis, D., *Adverbs of quantification*, in: E. Keenan, editor, *Formal Semantics of Natural Language*, Cambridge University Press, 1975 pp. 178–188.
- [9] MacFarlane, J., *Future contingents and relative truth*, *The Philosophical Quarterly* **53** (2003), pp. 321–336.
- [10] Marra, A., *When is reasoning by cases valid? Some remarks on persistence and information-sensitivity*, Manuscript.
- [11] Marx, M., *Complexity of modal logic*, in: P. Blackburn, J. van Benthem and F. Wolter, editors, *Handbook of Modal Logic*, Elsevier, 2007 pp. 139–179.
- [12] McCall, S., editor, “Polish Logic 1920-1939,” Oxford, 1967.
- [13] Øhrstrøm, P., *Problems regarding the future operator in an indeterministic tense logic*, *Danish Yearbook of Philosophy* **18** (1981), pp. 81–95.
- [14] Øhrstrøm, P. and F. Hasle, *Future contingents*, in: E. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Metaphysics Research Lab, Stanford University, 2015.
- [15] Prior, A., “Past, Present and Future,” Oxford, 1967.
- [16] Ramsey, F., *General propositions and causality*, in: D. Mellor, editor, *Philosophical Papers*, Cambridge University Press, 1990 pp. 145–163.
- [17] Thomason, R., *Indeterminist time and truth value gaps*, *Theoria* **36** (1970), pp. 264–281.
- [18] van Benthem, J., *Foundations of conditional logic*, *Journal of Philosophical Logic* **13** (1984), pp. 303–349.
- [19] Veltman, F., “Logics for Conditionals,” Ph.D. thesis, University of Amsterdam (1985).
- [20] Veltman, F., *Defaults in update semantics*, *Journal of Philosophical Logic* **25** (1996), pp. 221–261.