



Counterfactuals Modulo Temporal Logics*

Bernd Finkbeiner and Julian Siber

CISPA Helmholtz Center for Information Security, Saarbrücken, Germany
{finkbeiner,julian.siber}@cispa.de

Abstract

Lewis’ theory of counterfactuals is the foundation of many contemporary notions of causality. In this paper, we extend this theory in the temporal direction to enable symbolic counterfactual reasoning on infinite sequences, such as counterexamples found by a model checker and trajectories produced by a reinforcement learning agent. In particular, our extension considers a more relaxed notion of similarity between worlds and proposes two additional counterfactual operators that close a semantic gap between the previous two in this more general setting. Further, we consider versions of counterfactuals that minimize the distance to the witnessing counterfactual worlds, a common requirement in causal analysis. To automate counterfactual reasoning in the temporal domain, we introduce a logic that combines temporal and counterfactual operators, and outline decision procedures for the satisfiability and trace-checking problems of this logic.

1 Introduction

Evaluating counterfactual statements is a fundamental problem for many approaches to causal reasoning [40]. Such reasoning can for instance be used to explain erroneous system behavior with a counterfactual statement such as ‘If the input i at the first position of the observed computation π had not been enabled then the system would not have reached an error e .’ which can be formalized using the counterfactual operator $\square \rightarrow$ and the temporal operator F :

$$\pi \models (\neg i) \square \rightarrow (\neg F e) .$$

Since the foundational work by Lewis [38] on the formal semantics of counterfactual conditionals, many applications for counterfactuals [28, 5, 34, 46, 3, 15] and some theoretical results on the decidability of the original theory [37] and related notions [20, 2] have been discovered. Still, certain domains have proven elusive for a long time, for instance, theories involving higher-order reasoning and an infinite number of variables. In this paper, we consider a domain that combines both of these aspects: temporal reasoning over infinite sequences. In particular, we consider counterfactual conditionals that relate two properties expressed in temporal logics, such as the temporal property $\neg F e$ from the introductory example. Temporal logics are used ubiquitously

*This work was partially supported by the DFG in project 389792660 (Center for Perspicuous Systems, TRR 248), and by the ERC Grant HYPER (No. 101055412).

as high-level specifications for verification [21, 4] and synthesis [22, 41], and recently have also found use in specifying reinforcement learning tasks [32, 39]. Our work lifts the language of counterfactual reasoning to similar high-level expressions. We consider Quantified Propositional Temporal Logic (QPTL) because it can characterize the full class of ω -regular properties and in this way subsumes popular specification languages. This results in our logic QPTL_{cf} , which mixes QPTL with counterfactual conditionals and can be used to check counterfactual dependencies between ω -regular properties. We abstract away from any concrete causal models in this paper but refer to recent works by Halpern and Peters [29], and Coenen et al. [16] on extending these models to infinitely many variables.

Focusing on the core counterfactual reasoning inherent to causality allows us to study several key problems arising in the temporal domain. We believe the main reason higher-order reasoning and an infinite number of variables induce difficulties for counterfactual reasoning is ultimately tied to Lewis’ rejection of the *Limit Assumption*, which stipulates that for any world W and property φ , there is a unique set of worlds minimally close to W that satisfy φ . If the assumption holds in some domain, counterfactual reasoning is simple: A statement such as ‘If φ had held then ψ would have held, too.’, formally expressed by the formula $\varphi \Box \rightarrow \psi$, then would only need to compute the set of worlds minimally close to W that satisfy φ and check whether all of them satisfy ψ , too. However, as Lewis points out, the assumption is generally not true in any continuous domain, and as we will see later, it is also not generally true in the temporal domain. Previous works on defining notions of causality in settings with infinitely many variables sidestep the issue by restricting to similarity relations [27, 16] or logics for cause and effect [35] that satisfy the Limit Assumption, but this imposes a significant toll on the precision of the inferred causes, since this requires coarsely overapproximating the set of closest traces.

One of the key insights of this work is that it is possible to reason about counterfactuals in the temporal domain even when rejecting the Limit Assumption. Without the assumption, evaluating counterfactual conditionals requires complex quantification over the possibly infinite chains of worlds ever closer to W . To solve these quantified statements in our domain, we use recent advances in the study of hyperproperties [14] and their corresponding temporal logics [13, 43], which originate in the verification of information-flow policies and allow to relate multiple traces of a system to another. This pushes their decidability to the edge, for instance, the satisfiability problem of **HyperQPTL**, which extends QPTL with quantification over traces, is undecidable. While our counterfactual-temporal logic QPTL_{cf} has inherently relational semantics, its models are still traces, in contrast to the sets of traces modeled by **HyperQPTL**. Further, in QPTL_{cf} the trace-quantification is guarded by the counterfactual conditionals. We show that together this yields decidability for the satisfiability, model-checking and trace-checking problems of QPTL_{cf} by encoding them into the decidable model-checking problem of **HyperQPTL**.

We address several limitations of Lewis’ original theory of counterfactuals, with the goal that our logic QPTL_{cf} can be practically used to specify notions of temporal causality. For instance, the original theory requires a similarity relation between worlds that is total, i.e., for any two worlds W_1 and W_2 it needs to be possible to assess which one is closer to the reference world W . This turns out to be far too restrictive for reasonable similarity relations between infinite traces, and hence we extend Lewis’ theory to non-total similarity relations. However, since in such relations, there may be several incomparable sets of worlds that may count as the worlds minimally close to W , this opens a semantic gap between Lewis’ two proposed counterfactual operators. The crux is that, for instance, a naive extension of the ‘Would’ counterfactual $\Box \rightarrow$ quantifies existentially over these sets of closest worlds, meaning its enough if there is one path from W to worlds satisfying φ where the closest worlds then also satisfy ψ . But there may be other paths to φ where the closest worlds do not satisfy ψ . A similar problem exists with Lewis’

semantics of the ‘*Might*’ counterfactual. We argue that for both operators, the naive extensions to non-total similarity relations do not match the intended semantics of the counterfactual statements they are supposed to formalize, and propose fixed semantics for non-total similarity relations. Another common requirement for causes is a notion of minimality [26, 25, 30, 16, 15], such that negating the cause describes the minimal changes necessary to avoid the effect. This notion is not covered by Lewis’ original theory. To enable our logic to express the minimality condition, we introduce *minimal* counterfactual operators. The intuition behind a minimal counterfactual such as $\varphi \Box \rightarrow_{\min} \psi$ is that it is meant to minimize the path from the reference world to the counterfactual worlds satisfying ψ . This boils down to a second-order requirement that quantifies over properties φ' to see whether some of them characterize a superset of φ and still qualify in the counterfactual. Since the second-order quantification is guarded by the minimal counterfactual operators, we can eliminate it by giving equisatisfiable first-order formulas that only quantify over traces.

We show that with an appropriate choice of underlying universe, QPTL_{cf} can express several notions of causality proposed in previous literature, and use it to forge an interesting link between event-based actual causation and property-based counterfactual causation.

Contributions. In summary, our contributions are as follows:

- We extend Lewis’ theory of counterfactual conditionals to non-total similarity relations (Section 3.1) by proposing two additional counterfactual conditionals that capture the intended semantics of ‘*Would*’ and ‘*Might*’ on these relations (Section 3.3).
- We study a minimality criterion for counterfactuals that captures necessary reasoning for causal analysis and introduce minimal counterfactual operators (Section 3.4).
- We build a logic that mixes the classic counterfactuals due to Lewis and our newly proposed counterfactual conditionals with temporal properties expressed in QPTL, and show that the corresponding satisfiability and trace-checking problems are decidable (Section 4).

Necessary preliminaries on temporal logics are introduced in the following section (Section 2), and Lewis’ original counterfactual conditionals are discussed in Section 3.2.

Related Work. Our theory is an extension of Lewis’ theory of counterfactuals [38, 37] both in terms of the language of the antecedents and consequents, as well as for more general similarity relations and reasoning about minimality. Previous works in the context of axiomatizing causal modeling have extended the language of consequents to arbitrary Boolean formulas [24] as well as to counterfactual consequents [10]. Our work lifts the language of cause and effect to restricted first- and second-order reasoning in an infinite domain and is in this way related to recent efforts by Halpern and Peters [29] on causal reasoning with infinitely many variables.

While to our knowledge there has been no previous work extending Lewis’ logic of counterfactuals to temporal reasoning, there are works that have proposed some notion of temporal causality in, e.g., Markov Decision Processes [47] and reactive systems [16]. Several previous works have made a connection between causality and hyperproperties [1, 18, 16, 15].

2 Temporal Logics

We consider temporal logics whose models are infinite *traces* $t = t[0]t[1]\dots \in (2^{AP})^\omega$ over some finite set of atomic propositions AP . As a basic temporal logic we consider Linear-time

Temporal Logic (LTL) [42]. LTL formulas are built with the following grammar, where $a \in AP$:

$$\varphi ::= a \mid \neg\varphi \mid \varphi \wedge \varphi \mid \mathbf{X}\varphi \mid \varphi \mathbf{U}\varphi .$$

The semantics of LTL are given by the following satisfaction relation, which recurses over the positions i of the trace.

$$\begin{aligned} t, i \models a & \quad \text{iff} \quad a \in t[i] \\ t, i \models \neg\varphi & \quad \text{iff} \quad t, i \not\models \varphi \\ t, i \models \varphi \wedge \psi & \quad \text{iff} \quad t, i \models \varphi \text{ and } \pi_i \models \psi \\ t, i \models \mathbf{X}\varphi & \quad \text{iff} \quad t, i+1 \models \varphi \\ t, i \models \varphi \mathbf{U}\psi & \quad \text{iff} \quad \exists j \geq i. t, j \models \psi \wedge \forall i \leq k < j. t, k \models \varphi \end{aligned}$$

We say a trace t satisfies a formula φ iff the formula holds at the first position: $t, 0 \models \varphi$, we also write $t \models \varphi$ to denote this. We also consider the usual derived Boolean (\vee , \rightarrow , \leftrightarrow) and temporal operators ($\varphi \mathbf{R}\psi \equiv \neg(\neg\varphi \mathbf{U}\neg\psi)$, $\mathbf{F}\varphi \equiv \text{true} \mathbf{U}\varphi$, $\mathbf{G}\varphi \equiv \text{false} \mathbf{R}\varphi$).

Example 1. *To illustrate how LTL can specify the dynamics of a system, consider an elevator that moves up (u) and down (d) between three floors bottom (b), middle (m), and top (t). We have a set of atomic propositions $AP = A \cup S$ composed of two subsets $A = \{u, d\}$ for actions and $S = \{b, m, t\}$ for states. The dynamics of the system starting at the lowest floor can be specified in an LTL formula:*

$$\begin{aligned} \varphi_{\text{elevator}} \equiv & b \wedge \mathbf{G} \left((b \wedge u \rightarrow \mathbf{X}m) \wedge (b \wedge d \rightarrow \mathbf{X}g) \wedge (m \wedge u \rightarrow \mathbf{X}t) \wedge (m \wedge d \rightarrow \mathbf{X}g) \wedge \right. \\ & \left. (t \wedge u \rightarrow \mathbf{X}t) \wedge (t \wedge d \rightarrow \mathbf{X}m) \wedge (t \not\leftrightarrow b) \wedge (b \not\leftrightarrow m) \wedge (m \not\leftrightarrow t) \right) . \end{aligned}$$

The ‘Globally’ operator \mathbf{G} universally quantifies over all time points in one sequence, requiring that all of the conjuncts in its body hold. The conjuncts themselves encode the dynamics, e.g., $b \wedge u \rightarrow \mathbf{X}m$ ensures that when the elevator moves up from the bottom floor, it reaches the middle floor in the next state (which the ‘Next’ operator \mathbf{X} accesses). The formulas in the last line encode that the elevator can only be on one floor at the same time, together with the others they also ensure that only one action can be done at any time point. The traces that satisfy the formula then describe all the valid dynamics of the system, e.g.:

$$t = \{b, u\}\{m, d\}(\{b, u\}\{m, d\})^\omega$$

is a trace where the elevator cycles between the bottom and the middle floor. The ω -superscript symbolizes that this part of the trace is repeated infinitely often.

LTL is of practical significance because its corresponding decision procedures are of comparatively low complexity. However, this comes at a cost in expressivity, such that it cannot specify that, e.g., an atomic proposition eventually holds at an odd position. To make our results applicable to as many properties as possible, we therefore consider Quantified Propositional Temporal Logic (QPTL), introduced by Sistla [45], throughout the technical sections of this paper. QPTL extends LTL by quantification over atomic propositions. Its syntax is built atop of LTL as follows, where $q \notin AP$ is a fresh atomic proposition, and φ is an LTL formula:

$$\psi ::= \exists q. \psi \mid \forall q. \psi \mid \varphi .$$

As presented by Finkbeiner et al. [23], the semantics of the formulas quantifying over propositions can be stated using a replacement function $t[q \mapsto t_q]$ that given a trace $t \in (2^{AP})^\omega$ and

a trace $t_q \in (2^{\{q\}})^\omega$ sets the occurrences of q in t to the ones in t_q , i.e., $t[q \mapsto t_q] =_{\{q\}} t_q$ and $t[q \mapsto t_q] =_{AP \setminus \{q\}} t$, where $t =_E t'$ means the truth value of the subset $E \subseteq AP$ agrees on all positions of the two traces.

$$\begin{aligned} t_i \models \exists q. \psi & \text{ iff } \exists t_q \in (2^{\{q\}})^\omega. t[q \mapsto t_q] \models \psi \\ t_i \models \forall q. \psi & \text{ iff } \forall t_q \in (2^{\{q\}})^\omega. t[q \mapsto t_q] \models \psi \end{aligned}$$

For a formula φ in QPTL (and hence also in the fragment LTL), we denote by $\mathcal{L}(\varphi)$ the set of traces that satisfy it.

Example 2. *Quantification over propositions allows limited forms of counting in QPTL, such as in the following formula φ_{odd} that tracks the parity of positions with q and hence can express that the atomic proposition b eventually holds at an odd position:*

$$\varphi_{\text{odd}} \equiv \exists q. \neg q \wedge \mathbf{G}(\neg q \rightarrow \mathbf{X} q) \wedge \mathbf{F}(q \wedge b) .$$

Since on trace π from Example 1 b only holds at even positions we have that $\pi \not\models \varphi_{\text{odd}}$.

The semantics of the counterfactual conditionals we consider in this work require quantification over the worlds described by their antecedents and consequents. In our case, the worlds are traces. To express quantification over traces, we make use of hyperlogics [13, 17], temporal logics that originated in information-flow security and relate multiple traces to one another. The hyper-counterpart to QPTL is HyperQPTL [43] and extends the syntax with trace quantifiers over a set of trace variables \mathcal{V} , where $\pi \in \mathcal{V}$. Since HyperQPTL is only used for decidability proofs, and all proofs are found in the appendix due to space reasons, we also state the syntax and semantics of HyperQPTL only in Appendix A.1.

3 Counterfactuals

We now outline our extended version of Lewis' theory of counterfactual conditionals. First, we extend Lewis' notion of a similarity relation to non-total orders (Section 3.1). Then, we introduce the classic counterfactuals (Section 3.2) and two new operators (Section 3.3), and lastly, we consider a minimality criterion for counterfactual antecedents (Section 3.4). Since the concepts discussed in this section are not only applicable to the traces and temporal logics discussed in the previous section, we will adopt Lewis' modal nomenclature and speak of worlds and properties. We assume a set of worlds \mathcal{U} called *universe*. If not stated explicitly otherwise, all quantifiers in this section will quantify over \mathcal{U} . Further, we assume some logic \mathcal{L} and a satisfaction relation that tells us for any world $W \in \mathcal{U}$, whether it satisfies some property $\varphi \in \mathcal{L}$, which we denote with $W \models \varphi$. We will make the connection to the previous section clear by using traces as worlds and LTL formulas as properties in our concrete examples.

3.1 The Distance Between Worlds

The semantics of counterfactuals rely on reasoning about the relative similarity of worlds with respect to the reference world in which the counterfactual is evaluated. Directly mapping Lewis' counterfactuals to our setting would necessitate a total preorder over the set of traces to tell us which of any two given traces is closer to our original trace. However, in practice, such a total order is unrealistic not just in our trace-based context, as changes between two worlds may simply be incomparable. For instance, consider that changing atomic proposition u at position 0 has the same quantitative distance as changing it at position 1, but since the identity

of these changes differs, the direction in the space of worlds is different. Many instances of counterfactual reasoning, therefore, base their notion of distance on subset relationships between changes, i.e., some world is further away than another if the changes necessary to obtain the former are a superset of the changes necessary to obtain the latter [26, 25, 16]. Then, if no subset relationship in either directions holds between the changes manifesting in two worlds, their distance is incomparable. The underlying spatial structure of such a similarity relation is a lattice over the equivalence classes of some preorder. We generalize Lewis' counterfactual reasoning to preorders to account for these more general similarity relations. Formally, we require a *comparative similarity relation* \leq_W , which is a preorder on \mathcal{U} such that W is a minimum: $\forall W'. W' \not\leq_W W$.¹

Example 3. *In the context of trace logics as outlined in the previous section, worlds correspond to infinite traces such as t in Example 1. Our universe may be given by the language of some QPTL formula such as the one describing the dynamics of the elevator system in Example 1, so we have $\mathcal{U} = \mathcal{L}(\varphi_{\text{elevator}})$. A similarity relation may track the changes with respect to the reference trace t over a subset X of atomic propositions and can be formalized as follows:*

$$\leq_t(X) = \{(t_1, t_2) \mid \forall n \in \mathbb{N}. \forall x \in X. t[n] \neq_{\{x\}} t_1[n] \Rightarrow t[n] \neq_{\{x\}} t_2[n]\} .$$

To illustrate with the elevator system from Example 1, the similarity relation $\leq_t(A)$ orders a trace $t'' = \{b, d\}\{b, d\}(\{b, u\}\{m, d\})^\omega$, that changes the first two actions in trace t but keeps the other actions the same, as closer to t than trace $t' = (\{b, d\})^\omega$, which changes them on the whole sequence. We have $t'' \leq_t(A) t'$. Note that the above similarity relation is not total and would hence not be covered by Lewis' original theory.

3.2 Classic Counterfactuals

We start this section by recalling the semantics of Lewis' counterfactual operators $\Box \rightarrow$ and $\Diamond \rightarrow$, based on the reformulation for similarity relations [38]. In Lewis' theory, the operator $\Box \rightarrow$ is a formalization of 'Would' counterfactual statements such as 'If the elevator had eventually moved up two times in a row, then it would have reached the top floor.' The operator $\Diamond \rightarrow$ is a formalization of 'Might' counterfactual statements such as 'If the elevator had eventually moved up two times in a row, then it might have reached the top floor two steps after the start.' Intuitively, both of the statements should be true on trace t from Example 1. No matter where we change the trace such that the elevator moves upwards twice, it will in all cases end up at the top floor. And there is one instance, i.e., when changing the first two actions appropriately, that it will be at the top floor at the third position. Formally speaking, the distinction between the two statements stems from different quantification over the closest worlds. 'Would' counterfactuals are statements that quantify universally over all closest worlds that satisfy the antecedent, while 'Might' counterfactuals quantify existentially. There are further subtle differences in the vacuous case which we will discuss after giving the formal semantics.

Definition 4 (Semantics of 'Would'). *A world W satisfies $\varphi \Box \rightarrow \psi$ iff:*

$$\forall W_1. W_1 \not\models \varphi (1), \text{ or } \exists W_1. W_1 \models \varphi \wedge \forall W_2. W_2 \leq_W W_1 \Rightarrow W_2 \models \varphi \rightarrow \psi (2) .$$

¹Lewis refines his similarity relation based on a notion of accessibility. As our similarity relation is not necessarily total, accessibility can be easily encoded by not relating in \leq_W the models inaccessible from W . We could then express accessibility of a world W' from W by requiring $W \leq_W W'$, but we will abstract away from this concept for simplicity and assume that W is the unique minimum: $\forall W'. W \leq_W W'$.

It is worth pointing out that the complex expression in Condition 2 of Definition 4 mainly stems from Lewis' rejection of the Limit Assumption, which poses that for any antecedent φ and world W there exists a unique set of (equally) closest worlds satisfying φ . Then it would be easy to simply quantify over this set universally and require the consequent to hold in all of the worlds. However, in many scenarios, this assumption, unfortunately, does not hold. Instead, there may in fact be an infinite chain of ever closer worlds that satisfy φ . In these instances, what we are rather interested in is finding a 'threshold world' after which all closer worlds satisfying the antecedent φ also satisfy the consequent ψ .

Example 5. $\Box \rightarrow$ over QPTL formulas allows us to express the mix of counterfactual and temporal statements that is 'If the elevator had eventually moved up two times in a row, then it would have reached the top floor.' It yields the following formula:

$$\varphi_{top} \equiv \mathbf{F}(u \wedge \mathbf{X} u) \Box \rightarrow \mathbf{F} t .$$

Interpreted in the universe given by $\mathcal{L}(\varphi_{elevator})$ and with the similarity relation $\leq_t(A)$, we have that trace $t = \{b, u\}\{m, d\}(\{b, u\}\{m, d\})^\omega$ from Example 1 satisfies φ_{top} , because no matter where we change actions to ensure two moves upward in a row, we always end up in the top floor (this in fact holds for any trace in the universe).

For a more complex formula illustrating that the Limit Assumption does not hold in this setting, consider the statement 'If the elevator would eventually only move downwards, then it would eventually stay on the bottom floor.' This corresponds to the formula:

$$\varphi_{bottom} \equiv \mathbf{F}(\mathbf{G} d) \Box \rightarrow \mathbf{F}(\mathbf{G} b) .$$

Trace $t' = (\{b, d\})^\omega$ satisfies $\mathbf{F}(\mathbf{G} d)$, but there is an infinite chain of traces closer to t that also satisfy $\mathbf{F}(\mathbf{G} d)$:

$$t'' = \{b, u\}\{m, d\}(\{b, d\})^\omega, \quad t''' = \{b, u\}\{m, d\}\{b, u\}\{m, d\}(\{b, d\})^\omega, \dots$$

and so on. Hence, we cannot avoid Lewis' complex quantification over traces to evaluate temporal counterfactuals in the general case.

As one can see from Condition 1 in Definition 4, a 'Would' counterfactual is vacuously satisfied by a world if there are no related worlds that satisfy the antecedent. In contrast, a 'Might' counterfactual strictly requires a world that satisfies the antecedent, mainly because Lewis bases the semantics of the two operators on the following duality law: $\varphi \Box \rightarrow \psi \equiv \neg(\varphi \Diamond \rightarrow \neg\psi)$ [38]. This yields the following semantics for the 'Might' counterfactual.

Definition 6 (Semantics of 'Might'). A world W satisfies $\varphi \Diamond \rightarrow \psi$ iff all of the following holds:

$$\exists W_1. W_1 \models \varphi \quad (1), \quad \text{and} \quad \forall W_1. W_1 \models \varphi \Rightarrow \exists W_2. W_2 \leq_W W_1 \wedge W_1 \models \varphi \wedge \psi \quad (2) .$$

Again, significant complexity is introduced into the definition based on the rejection of the Limit Assumption. Here, however, the idea is not to find a 'threshold world', but instead to find for any world in the chain of ever closer worlds another one that is closer (or equally close) such that both antecedent and consequent are true.

3.3 Counterfactuals Over Non-total Similarity Relations

The semantics proposed by Lewis' work well if the similarity relation is a total order. However, as we can see in the following example, the semantics do not match the intuitive meaning of the operators when the similarity relation is not total, as in our setting.

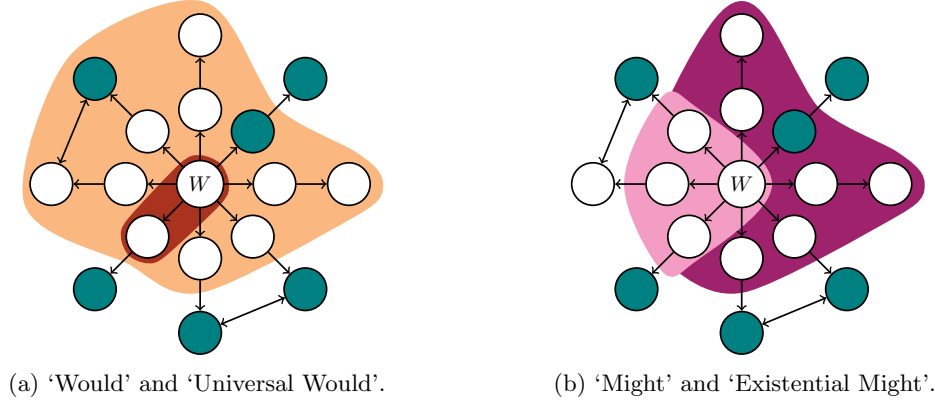


Figure 1: The set of worlds characterized by the counterfactuals on some similarity relation \leq_W . Nodes represent worlds (e.g., $W_{1,2,\dots}$) where the center node is the reference world W . An edge from W_1 to W_2 means $W_1 \leq_W W_2$. For simplicity, we omit edges that can be inferred from reflexivity or transitivity of \leq_W . Nodes of worlds that satisfy ψ are colored in teal.

Example 7. Consider the statement ‘If the elevator had eventually moved up two times in a row, then it would have reached the top floor two steps after the start.’, which corresponds to the following formula:

$$\varphi''_{top} \equiv \mathbf{F}(u \wedge \mathbf{X}u) \square \rightarrow \mathbf{X}(\mathbf{X}t) .$$

Intuitively, this statement should not be satisfied by trace t from the previous examples. After all, there is trace

$$t'_1 = \{b, u\}\{m, d\}\{b, u\}\{m, u\}\{t, d\}\{m, d\}t[6]t[7] \dots$$

that is a closest trace satisfying $\mathbf{F}(u \wedge \mathbf{X}u)$ and it does not satisfy $\mathbf{X}(\mathbf{X}t)$. Yet, we can simply instantiate the existential quantifier in Condition 2 of Definition 4 with

$$t'_2 = \{b, u\}\{m, u\}\{t, d\}\{m, d\}t[4]t[5] \dots .$$

Since there are no traces between t'_2 and t , the condition is satisfied, and we obtain that $t \models \varphi''_{top}$. Now, consider

$$\varphi'''_{top} \equiv \mathbf{F}(u \wedge \mathbf{X}u) \diamond \rightarrow \mathbf{X}(\mathbf{X}t) .$$

Intuitively, this statement should be satisfied, as there are indeed closest worlds that satisfy the consequent, namely t'_2 . However, since the semantics of the operator quantifies over all worlds satisfying the antecedent and requires a smaller one that satisfies the consequent for each, the existence of t'_1 means that $t \not\models \varphi'''_{top}$.

Example 7 shows that on preorders the semantics of $\square \rightarrow$ is too weak and that of $\diamond \rightarrow$ too strong to account for the intended meaning of the operators. Intuitively, ‘Would’ should express that in the set of closest possible worlds, all worlds satisfy the consequent. However, in preorders there may exist multiple incomparable classes of worlds on different chains along the similarity relation. The focal question is how to quantify over these classes, existentially such that one class is enough (like in the naive extension) or universally such that in all classes all models have to satisfy the consequent.

Figure 1 abstractly illustrates the problem with the ‘Would’ counterfactual in Subfigure 1a and the problem with the ‘Might’ counterfactual in Subfigure 1b. The areas colored in the dark

colors represent the complement of the set characterized by an antecedent that satisfies the corresponding counterfactual in this universe and similarity relation, where the worlds colored in teal satisfy ψ . For example $W \models \varphi_{\text{Would}} \square \rightarrow \psi$ if the worlds satisfying $\neg \varphi_{\text{Would}}$ are the ones included in the dark brown area in Subfigure 1a. Here we can see that there are paths out of this area that lead to closest φ_{Would} -worlds that do not satisfy ψ , which clashes with the intended semantics of ‘If φ_{Would} were true, then ψ would be true as well.’ What we instead want to capture is something like the bright brown area: No matter where we leave this area along the similarity relation, we always end up in a ψ -world. We capture this intention with the ‘Universal Would’ counterfactual defined in the following, which requires all closest φ -worlds on any chain in the similarity relation to satisfy ψ .

Definition 8 (Semantics of ‘Universal Would’). *A world W satisfies $\varphi \square \rightarrow \psi$ iff either of the following holds:*

$$\begin{aligned} & \forall W_1. W_1 \not\models \varphi, \text{ or} \\ & \forall W_1. W_1 \models \varphi \Rightarrow \exists W_2. W_2 \leq_W W_1 \wedge W_2 \models \varphi \wedge \forall W_3. W_3 \leq_W W_2 \Rightarrow W_3 \models \psi \end{aligned}$$

To ground our definition, we can show that it is equivalent to Lewis’ classic ‘Would’ counterfactual on total orders.

Proposition 9. *If \leq_W is a total order, then ‘Universal Would’ and ‘Would’ are equivalent, i.e., for all pairs of properties φ and ψ we have that $\varphi \square \rightarrow \psi \equiv \varphi \square \rightarrow \psi$.*

We now introduce a similar counterpart to the ‘Might’ counterfactual. We derive the semantics from a duality law similar to the original one: $\varphi \square \rightarrow \psi \equiv \neg(\varphi \diamond \rightarrow \neg\psi)$.

Definition 10 (Semantics of ‘Existential Might’). *A world W satisfies $\varphi \diamond \rightarrow \psi$ iff:*

$$\exists W_1. W_1 \models \varphi \wedge \forall W_2. W_2 \leq_W W_1 \wedge W_2 \models \varphi \Rightarrow \exists W_3. W_3 \leq_W W_2 \wedge W_3 \models \psi$$

Since we derived the ‘Existential Might’ counterfactual from a similar duality law as Lewis, we can deduce from Proposition 9 that our ‘Existential Might’ and the original ‘Might’ are also equivalent for total similarity relations.

Corollary 11. *If \leq_W is a total order, then ‘Existential Might’ and ‘Might’ are equivalent, i.e., for all pairs of properties φ and ψ we have that $\varphi \diamond \rightarrow \psi \equiv \varphi \diamond \rightarrow \psi$.*

Subfigure 1b illustrates how this definition captures the intended meaning of ‘Might’ counterfactuals on non-total similarity relations. To illustrate the problem with Lewis’ original semantics for ‘Might’: The complement of some satisfying antecedent φ_{Might} for a classic ‘Might’ counterfactual is colored in dark violet. For satisfaction it is required that no matter where we leave this area, for any world there exists an equally close world satisfying ψ . This means for example, that we can include the worlds furthest to the west and north-west in φ_{Might} . This is quite strict, our ‘Existential Might’ counterfactual illustrated in bright violet instead allows the closest worlds on a chain to not satisfy ψ , as long as on some chain there exists a closest world satisfying ψ .

3.4 Minimal Counterfactuals

When reasoning about causation, we are most often interested in some notion of minimality to characterize the minimal changes necessary to avoid a given effect [26, 25, 30, 15, 16]. From a

counterfactual point of view, minimality formulates an additional condition on the antecedent φ such that the property defines the largest set possible. The question of whether to term this notion minimality or maximality is a matter of perspective. On the one hand, the criterion maximizes the language of φ , but on the other hand, this in fact minimizes the amount of changes necessary to ensure φ , since more worlds in φ mean more opportunities to move into φ earlier when moving along the similarity relation. Since for causation this abstract criterion is usually called minimality, we adopt the same name here.

Definition 12 (Minimal Counterfactuals). *Given a world W and a counterfactual conditional $\rightsquigarrow \in \{\Box \rightarrow, \Diamond \rightarrow, \Box \rightarrow, \Diamond \rightarrow\}$, the minimal counterfactual conditional $\varphi \rightsquigarrow_{\min} \psi$ is true iff all of the following holds:*

$$W \models \varphi \rightsquigarrow \psi \text{ (1), and } \nexists \varphi'. (\varphi \rightarrow \varphi') \wedge (\varphi' \not\rightarrow \varphi) \wedge W \models \varphi' \rightsquigarrow \psi \text{ (2) .}$$

Example 13. *Minimal counterfactuals ensure that the antecedent does not overspecify the changes necessary to get to the consequent. For instance, consider that both of the following statements hold:*

$$t \models u \wedge X u \Box \rightarrow X(Xt), \text{ and } t \models X u \Box \rightarrow X(Xt) .$$

While the antecedent in the lower statement is more concise because t already has u at the first position, neither is a minimal antecedent, because minimality additionally ensures that all possible antecedents are included, we have:

$$t \models (X u) \vee X(Xt) \Box \rightarrow_{\min} X(Xt) .$$


For a short argument, consider why there cannot be a trace in the minimal antecedent that does not satisfy $(X u) \vee X(Xt)$. Such a trace would have no upwards movement at the second position, and all traces closer to t do neither. Since all traces start at the bottom floor, we know that none of the traces in between satisfy $X(Xt)$.

While it may seem odd that the consequent can be a necessary part of the antecedent like in the above example, we note that it is common for the effect to be in the set of its counterfactual causes [26, 25, 9]. We believe there is a quite direct connection between minimal counterfactuals, and conjunctive and disjunctive causes from the causal modeling literature. We expand on this later in Section 4.2.

As one can see in Definition 12, minimality is a second-order statement that quantifies over properties. Hence, the minimal counterfactual conditionals $\Box \rightarrow_{\min}$, $\Diamond \rightarrow_{\min}$, $\Box \rightarrow_{\min}$, and $\Diamond \rightarrow_{\min}$ are essentially guarded second-order quantifiers. However, one of our main insights is that this second-order quantification inherent to these operators can be eliminated such that for any minimal counterfactual conditional there exists an equisatisfiable formula without quantification over properties. We prove this in the following lemma and will use this in the following section for concrete decidability results.

Lemma 14. *For a minimal counterfactual statement $\varphi \rightsquigarrow_{\min} \psi$, with $\rightsquigarrow \in \{\Box \rightarrow, \Diamond \rightarrow, \Box \rightarrow, \Diamond \rightarrow\}$, there exists a parameterized formula $\varphi_{FO}(W)$, with $W \in \mathcal{U}$, that quantifies only over worlds and not properties such that for all worlds $W' \in \mathcal{U}$: $W' \models \varphi \rightsquigarrow_{\min} \psi$ iff $\varphi_{FO}(W')$ is valid.*

To illustrate minimal counterfactuals consider the abstract scenario in Figure 2. In Subfigure 2a colored areas correspond to the complements of a satisfying antecedent of a ‘Universal Would’ counterfactual with reference world W . However, this antecedent is not minimal, because the complement of the dashed area is equally part of a ‘Universal Would’ counterfactual



(a) ‘Universal Would’ and its minimal version. (b) ‘Existential Might’ and its minimal version.

Figure 2: The set of worlds characterized by non-minimal counterfactuals (colored areas) compared to the set of worlds characterized by minimal counterfactuals (dashed line) in the same setting as in Figure 1.

satisfied by W . The illustration sums up the following intuition: The worlds encountered by leaving the dashed area that satisfy ψ are now overall closer to W than before. The same idea holds for the more minimal antecedent of the ‘Existential Might’ counterfactual illustrated with the dashed area in Subfigure 2b. However, since for the ‘Existential Might’ counterfactual it is enough to find just one closest antecedent-world satisfying ψ , we can actually remove all worlds except W , because there is a ψ -world directly adjacent to it.

4 Temporal Counterfactuals

We now leverage the semantic insights garnered on counterfactuals in the previous section to design a logic for expressing notions of counterfactual causality in the temporal domain. For this, it is not enough to solely evaluate the truth values of counterfactual conditionals or even their minimal versions, since causality commonly places additional conditions for a causal relationship to hold. For instance, cause and effect may have to be satisfied in the reference world [26, 25, 16]. The syntax and semantics of our logic are introduced in the following section, where we also establish decidability of satisfiability and trace checking. In Section 4.2, we close by illustrating how our logic can express definitions of causality proposed in previous literature.

4.1 QPTL With Counterfactuals

We consider a logic that builds Boolean combinations of QPTL formulas and the proposed counterfactual conditionals relating QPTL statements. Nesting of counterfactuals is not allowed, although may be interesting to explore in future work to study temporally structured interventions. We call our logic QPTL_{cf} , for *QPTL with counterfactuals*, and its formulas are built according to the following grammar:

$$\begin{aligned} \xi ::= & \psi \square \rightarrow \psi \mid \psi \diamond \rightarrow \psi \mid \psi \blacksquare \rightarrow \psi \mid \psi \blacklozenge \rightarrow \psi \mid \psi \square \rightarrow_{\min} \psi \mid \\ & \psi \diamond \rightarrow_{\min} \psi \mid \psi \blacksquare \rightarrow_{\min} \psi \mid \psi \blacklozenge \rightarrow_{\min} \psi \mid \xi \wedge \xi \mid \neg \xi \mid \psi , \end{aligned}$$

where ψ is a QPTL formula. The semantics of the counterfactual conditionals are as discussed in the previous section where worlds are now infinite traces and the satisfaction relation is as for QPTL, and the semantics of the QPTL-formulas are as described in Section 2.

Now that we have fixed the syntax and semantics of our logic QPTL_{cf} , we show that important decision problems can be automatically decided. Note that previous results on Lewis' original two counterfactuals and related works on causal analysis either consider theories with a finite set of models [37], endorse the problematic limit assumption [16], or both [25, 2]. In this sense, our results extend previous work in several orthogonal directions.

We start by considering the satisfiability problem of QPTL_{cf} . Satisfiability denotes the problem of deciding whether there exists a trace that satisfies a given formula. Deciding the satisfiability of counterfactuals of course depends on whether the similarity relation can be expressed in a decidable logic. We, therefore, show satisfiability of minimal counterfactuals with respect to the following family of QPTL-expressible similarity relations:

$$\begin{aligned} \leq^{\text{QPTL}} = \{ R \subseteq ((2^{AP})^\omega)^3 \mid \exists \text{ QPTL-formula } \varphi_R \text{ over } AP_{\{\pi_1, \pi_2, \pi_3\}} \cdot \forall t_1, t_2, t_3 \in (2^{AP})^\omega. \\ R(t_1, t_2, t_3) \text{ iff } t_1 \cup t_2 \cup t_3 \models \varphi_R \} , \end{aligned}$$

where $AP_P = \{a_\pi \mid a \in AP \wedge \pi \in P\}$ indexes the atomic propositions with trace variables and $t_1 \cup t_2 \cup t_3$ denotes the disjoint fusion of the three traces: We have for all positions $n \in \mathbb{N}$ that the following holds: $(t_1 \cup t_2 \cup t_3)[n] = \{a_{\pi_1} \mid a \in t_1[n]\} \cup \{a_{\pi_2} \mid a \in t_2[n]\} \cup \{a_{\pi_3} \mid a \in t_3[n]\}$. Note that the similarity relations $\leq_t^{\text{QPTL}}(X)$ introduced in Example 3 are subsumed by this family, and also the distance metric considered by Coenen et al. [16]. We consider three-place relations here because the first place models the W -subscript of the relation \leq_W .

Satisfiability of formulas in QPTL_{cf} can then be decided with an idea roughly similar to the reduction of LTL satisfiability to LTL model checking proposed by Rozier and Vardi [44]. However, since there is no corresponding trace-based logic that can express the semantics of counterfactual conditionals, we instead encode the problem as a **HyperQPTL** model-checking problem over a model containing all possible traces over our set of atomic propositions AP . In the end, since **HyperQPTL** model checking is decidable [43], so is the satisfiability of QPTL_{cf} -formulas. This approach is interesting from a theoretical point of view, because **HyperQPTL** satisfiability checking itself is undecidable. We leverage the fact that models of QPTL_{cf} formulas are traces and not sets of traces. While the logic does have relational semantics, these are effectively guarded through the counterfactual conditionals, allowing us to encode into model checking and avoid the undecidable **HyperQPTL** satisfiability problem.

Theorem 15. *For any QPTL_{cf} formula φ , it is decidable to check whether φ is satisfiable when the similarity relation is from the family \leq^{QPTL} and the universe is defined by a QPTL formula.*

Since our logic contains negation, this also covers the problem of checking validity of a formula in the usual way, as a formula is valid if and only if its negation is not satisfiable. Similarly, because our logic subsumes QPTL, we can also model-check QPTL_{cf} -properties on any system that can be expressed in QPTL, which covers many practically relevant finite-state systems. This can be done by checking satisfiability of the conjunction of the system formula and the negation of the property formula (and possibly choosing the system as the universe).

Lastly, we consider the problem of checking whether some QPTL_{cf} -formula holds on a given trace. Since we need a finite representation of the infinite trace to feasibly compute this truth value, we consider lasso-shaped traces, i.e., traces of the form $t = t_0 \dots t_i (t_{i+1} \dots t_j)^\omega$ that ultimately repeat some loop in the infinite ω -part. The proof of the theorem then follows a similar idea as the satisfiability proof, except that we do not search for any satisfying trace and instead fix the corresponding trace variable to the lasso-shaped trace.

Theorem 16. *For a lasso-shaped trace t and a QPTL_{cf} formula φ , it is decidable to check whether $t \models \varphi$ when the similarity relation is from the family \leq^{QPTL} and the universe is defined by a QPTL formula.*

Theorem 15 and Theorem 16 show that it is possible to build tools that automatically check whether causal relationships hold in a given finite-state system, or whether causal relationships expressed in our logic are present on a given trace. While an exact complexity analysis is out of scope of this paper, QPTL has non-elementary complexity that scales in the number of propositional quantifier alternations. The complexity of our decision procedure for QPTL_{cf} as of now additionally scales in the number of counterfactual conditionals, since these introduce trace quantifier alternations and the complexity of HyperQPTL model checking scales with this number. We believe there may be a more efficient encoding, as there is no dependence between the trace variables resulting from different counterfactual conditionals and, hence, quantifiers can be rearranged to avoid alternations. However, the purpose of this work is to explore how expressive our temporal logic with counterfactuals can be while retaining decidability. In practice, it may in fact be more feasible to consider fragments of QPTL that have a more practical complexity, such as LTL . There exist efficient model checkers [6] for its counterpart HyperLTL [13] and, since recently, also for HyperQPTL [7], such that implementing a model checker for QPTL_{cf} or a fragment modulo LTL using our outlined encoding is feasible.

4.2 Examples

We close by illustrating how QPTL_{cf} can be used to express several notions of counterfactual causality proposed in previous literature. We start with Lewis’ account of counterfactual causation to illustrate how a basic, but limited, notion of causation can be expressed solely with counterfactuals. We then outline at the example of Halpern’s (non-temporal) actual causality the important role of proper modeling in causal analysis, and lastly we show how previous work of Coenen et al. on extending Halpern’s actual causality to temporal properties is subsumed by our work on temporal-counterfactual reasoning with QPTL_{cf} .

Lewis’ Counterfactual Causation [36]. Based on counterfactual conditionals, Lewis introduced a definition of causality in order to capture Hume’s informal account that “we may define a cause to be an object followed by another, where, if the first object had not been, the second never had existed” [31]. Lewis only draws causal relationships between two events, which are rather informally defined objects, but which in previous literature on traces are commonly interpreted to mean the value of an atomic proposition at a certain position [5, 35, 15]. Lewis definition stipulates that an event c is the cause for an event e if the following condition holds: $(c \Box \rightarrow e) \wedge (\neg c \Box \rightarrow \neg e)$. The intuition behind the formula is that either the cause c and e appear together (and hence the left conjunct is also satisfied), and then moving to any closest $\neg c$ -world is sufficient to avoid the effect e (corresponding to the left conjunct), or, neither cause and effect appear, and then moving to a c -world is sufficient to bring about the effect (the mirrored case). Consequently, this causal relation holds even when cause and effect are not present in a world, and hence may be model checked on a system to infer whether it holds globally.

With QPTL_{cf} , we can improve on Lewis’ original formulation in several ways. First, we can use the ‘Universal Would’ counterfactual and hence need not assume a total similarity relation. Further, Lewis’ original logic lacks facilities for temporal reasoning and, hence, cannot express that a cause is “an object followed by another”. In QPTL_{cf} , it is possible to express that the cause has to precede the effect, and the remaining temporal requirements, in the following way:

$$((\mathbf{F}c \wedge \mathbf{F}e) \rightarrow (\neg e \mathbf{U}c)) \wedge (\mathbf{F}c \square \rightarrow \mathbf{F}e) \wedge (\neg \mathbf{F}c \square \rightarrow \neg \mathbf{F}e) .$$

We can deduce causal relationships by checking this formula for validity in a universe of interest. It ensures that if the cause and effect appear on some trace, then, particularly, the cause happens before the effect. While this solves the lacking temporal expressivity in Lewis' logic, it still shares the idiosyncratic features of this definition of causality which have been raised in the literature since. For instance, if a cause has two effects, the earlier effect is also considered a cause for the later effect.

These peculiarities cannot be solved through counterfactual reasoning alone and over the years several solutions in the causal modeling literature have emerged, for instance *interventions* and *contingencies*. In the following paragraph, we show that our framework can emulate these concepts by modeling their mechanics in the universe and the similarity relation.

Halpern's Actual Causality [25]. We outline a direct correspondence of our approach to Halpern's modified version [25] of actual causation [26], for binary variables. The exact definition of this notion of causality is given in Appendix A.2 due to space reasons. The general idea is that the dynamics of the universe are defined by a structural equation f_x for every variable x , which may depend on other variables and a set of external variables which are defined by a context. We assume these to be Boolean functions in the case of binary variables. Halpern restricts the analysis to acyclic dependencies. Based on some model $\mathcal{M} = (\mathcal{S}, \mathcal{F})$ and context c we may assume an evaluation function $U_c^{\mathcal{M}}(x)$ which tells us for any variable $x \in V$ whether it evaluates to 1 or 0. To establish a correspondence to QPTL_{cf} , we construct a universe for a specific context under analysis based on basic equivalences for every variable: $x \leftrightarrow ((f_x \wedge \neg i_x) \vee c_x)$ if $U_c^{\mathcal{M}}(x) = 1$, else $x \leftrightarrow ((f_x \vee i_x) \wedge \neg c_x)$. If some f_x depends on external variables, we substitute them by the appropriate Boolean constant depending on the context. Then, $\mathcal{U}(\mathcal{M}, c)$ denotes the conjunction of all of these equivalences together with $\mathbf{XG}(\bigwedge_{v \in V} \neg v \wedge \neg i_v \wedge \neg c_v)$, i.e., all traces have only empty sets after the first position. This suffices because actual causation has no particular temporal aspect, so we may effectively model all outcomes of the structural equations on a finite trace prefix of length one. Let $t_c^{\mathcal{M}} \in \mathcal{L}(\mathcal{U}(\mathcal{M}, c))$ be the unique trace where for all $n > 0 : t_c^{\mathcal{M}}[n] = \emptyset$ and for all $x \in V : \{i_x, c_x\} \cap t_c^{\mathcal{M}}[0] = \emptyset$, i.e., no interventions or contingencies are present. It may seem problematic that we have to construct the universe based on given observations, but note that even in structural equations, certain equations have to be manipulated for modeling interventions and contingencies.

Theorem 17. $X_0 = x_0^0 \wedge \dots \wedge x_0^k, \dots, X_n = x_n^0 \wedge \dots \wedge x_n^j$ are actual causes of φ in (\mathcal{M}, \vec{u}) , iff $t_c^{\mathcal{M}} \models \varphi_X \wedge \varphi \wedge \neg \varphi_X \diamond_{\rightarrow_{\min}} \neg \varphi$ in the universe $\mathcal{U}(\mathcal{M}, c)$ with respect to the similarity relation $\leq_t(\{i_x \mid x \in V\})$ that tracks only the active interventions between traces, where

$$\neg \varphi_X = (\neg x_0^0 \wedge \dots \wedge \neg x_0^k) \vee \dots \vee (\neg x_n^0 \wedge \dots \wedge \neg x_n^j) \text{ is in Blake canonical form [8].}$$

The intuition behind the proof is that the closest traces to $t_c^{\mathcal{M}}$ that satisfy $\neg \varphi_X$ require negating one of the disjuncts, and hence require flipping the values of all of the variables appearing there. This simulates interventions on these variables. Existential search for contingencies is taken care of by the 'Might' counterfactual modality, since all possible contingencies for some intervention form an equivalence class under $\leq_t(\{i_x \mid x \in V\})$. The minimality of the counterfactual ensures that $\neg \varphi_X$ cannot be enlarged. This means neither can a conjunct be added to φ_X , and hence indeed all causes are described, nor can a disjunct be removed, and hence all causes are minimal. We give an illustrative example of this encoding in the following.

Example 18. Consider the classic example of determining whether lightning (l) or a dropped match (m) caused a forest fire (f). Assume that only the match was dropped and that this suffices to cause the fire because the equation for the fire is $f := l \vee m$ and the other two variables are determined by outside factors. We construct our universe based on the concrete observations $f = 1$, $l = 0$, and $m = 1$ and model interventions and contingencies accordingly:

$$\begin{aligned} \mathcal{U}_{\text{fire}}(\mathcal{M}, c) = & (l \leftrightarrow (i_l \wedge \neg c_l)) \wedge (m \leftrightarrow (\neg i_m \vee c_m)) \\ & \wedge (f \leftrightarrow (((l \vee m) \wedge \neg i_f) \vee c_f)) . \end{aligned}$$

Note that the actual causes according to Halpern’s definition in this scenario are $f = 1$, i.e., the effect itself, and $m = 1$. We have that the formula $f \wedge m$ qualifies as a cause for the effect f on the trace $t_{\text{fire}} = \{f, m\}\emptyset^\omega$ describing the above scenario. This is because we have $t_{\text{fire}} \models f \wedge m \wedge (\neg f \vee \neg m) \diamondrightarrow_{\min} \neg f$ in $\mathcal{U}_{\text{fire}}(\mathcal{M}, c)$.

All conjunctions appearing in the Blake canonical form in Theorem 17 are prime implicants. Dubsiaff et al. [19] have used these before to compute feature causes, which are counterfactual in nature, in configurable software systems. We use prime implicants here to establish an intriguing formal connection between property-based causes and event-based causes in structural equation models. This result is valuable because event-based causes have a restrictive and explicit syntax. The formula of Theorem 17 allows characterizing the same counterfactual reasoning with more expressive languages for causes and provides a formal basis for property-based extensions of actual causality. We will discuss such an extension by Coenen et al. in the following paragraph.

Coenen et al.’s Temporal Causality [16]. Coenen et al. extend interventions and contingencies to reactive systems described by Moore automata, and in this way lift Halpern’s actual causality to temporal properties on traces. The key idea is that, if the reference trace t is given in a lasso-shape, then the behavior of contingencies can be modeled in a finite-state machine called counterfactual automaton $\mathcal{C}_t^\mathcal{T}$, where \mathcal{T} is the original Moore automaton. This corresponds to the construction for finite settings in the above paragraph. Interventions need not be modeled explicitly because Coenen et al.’s causality only characterizes causes on the input sequence of the traces and as there are no dependencies between inputs, just changing them outright suffices to model interventions. Therefore, the concrete distance metric used in that work is a modification of $\leq_t(\mathcal{I})$, where \mathcal{I} is the set of inputs. The modification is done in order to satisfy the Limit Assumption, but results in a coarse over-approximation of the set of closest traces. For instance, for $t = \{a\}^\omega$, the closest traces satisfying $\mathbf{FG} \neg a$ with Coenen et al.’s similarity relation are the whole set $\mathcal{L}(\mathbf{FG} \neg a)$, while our work does not require the Limit Assumption to hold and can use the unmodified $\leq_t(\mathcal{I})$, which has an infinite chain of ever closer traces: $\{\}^\omega$, $\{a\}\{\}^\omega$, $\{a\}\{a\}\{\}^\omega$, and so on.

It turns out that when we characterize Coenen et al.’s definition of causality in \mathbf{QPTL}_{cf} , we see a minor divergence from Halpern’s actual causality. We can encode Coenen et al.’s causality in this way: φ is a temporal cause for ψ on t , iff

$$t \models \varphi \wedge \psi \wedge ((\neg \varphi \boxrightarrow_{\min} \neg \psi) \vee (\neg \varphi \diamondrightarrow_{\min} \neg \psi)) .$$

in the universe defined by the traces of $\mathcal{C}_t^\mathcal{T}$. We include the ‘Universal Would’ counterfactual because we need to make use of its vacuous satisfaction mechanics, i.e., \top may qualify as a temporal cause according to Coenen et al.’s definition by virtue of quantification over an empty set, but \perp will never work as an antecedent in $\diamondrightarrow_{\min}$, since there needs to be a trace satisfying it. We can use the compositional nature of \mathbf{QPTL}_{cf} to emulate this by using \boxrightarrow_{\min} , because it is,

except for the vacuous case, a stronger condition, i.e., whenever $\neg\varphi \Box \rightarrow_{\min} \neg\psi$ holds on a trace t in a universe and a $\neg\varphi$ -trace exists, also $\neg\varphi \Diamond \rightarrow_{\min} \neg\psi$ holds on trace t . We believe that the compositionality of a logic like $\text{QP}TL_{cf}$ can be a useful tool for comparing different definitions of causality, as demonstrated here between Halpern’s actual causality and Coenen et al.’s temporal causality. This process can even be automated with the outlined $\text{QP}TL_{cf}$ decision procedures.

5 Conclusion

In this paper, we study a fusion of two prominent flavors of modal logic: counterfactual and temporal reasoning. Our theoretical results are a step towards the automatic evaluation of temporal counterfactual conditionals on infinite sequences, such as counterexample traces returned by a model checker or trajectories of a reinforcement learning agent. Further, our extension of Lewis’ theory of counterfactual conditionals to non-total similarity relations and our minimal counterfactual operators are relevant to the theory of counterfactuals beyond the temporal reasoning considered in this work. In the future, we plan on using our logic to define system-level and trace-level causation in reactive systems, which correspond to the notions of global and actual causation. An interesting question here is, whether system-level semantics of counterfactuals should be a universal application of the trace semantics, or should counterfactually relate different system models. The latter approach may utilize previous work on system mutations [33] studied in the area of coverage [12], which has a tight relationship to counterfactual causality [11]. We are also interested in automating the discovery of causal relationships between temporal properties on infinite sequences. This problem can be framed as synthesizing satisfying antecedents of counterfactuals.

References

- [1] E. Ábrahám and B. Bonakdarpour. HyperPCTL: A temporal logic for probabilistic hyperproperties. In A. McIver and A. Horváth, editors, *Quantitative Evaluation of Systems - 15th International Conference, QEST 2018, Beijing, China, September 4-7, 2018, Proceedings*, volume 11024 of *LNCS*, pages 20–35. Springer, 2018.
- [2] G. Aleksandrowicz, H. Chockler, J. Y. Halpern, and A. Ivrii. The computational complexity of structure-based causality. *J. Artif. Intell. Res.*, 58:431–451, 2017.
- [3] C. Baier, N. Coenen, B. Finkbeiner, F. Funke, S. Jantsch, and J. Siber. Causality-based game solving. In A. Silva and K. R. M. Leino, editors, *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I*, volume 12759 of *LNCS*, pages 894–917. Springer, 2021.
- [4] C. Baier and J. Katoen. *Principles of model checking*. MIT Press, 2008.
- [5] I. Beer, S. Ben-David, H. Chockler, A. Orni, and R. J. Treffler. Explaining counterexamples using causality. *Formal Methods Syst. Des.*, 40(1):20–40, 2012.
- [6] R. Beutner and B. Finkbeiner. Autohyper: Explicit-state model checking for HyperLTL. In S. Sankaranarayanan and N. Sharygina, editors, *Tools and Algorithms for the Construction and Analysis of Systems - 29th International Conference, TACAS 2023, Paris, France, April 22-27, 2023, Proceedings, Part I*, volume 13993 of *LNCS*, pages 145–163. Springer, 2023.
- [7] R. Beutner and B. Finkbeiner. Model checking omega-regular hyperproperties with AutoHyperQ. In A. Voronkov and R. Piskac, editors, *24th International Conference on Logic for Programming, Artificial Intelligence and Reasoning, LPAR 2023, Manizales, Colombia, June 4-9, 2023*, EPiC Series in Computing. EasyChair, 2023.
- [8] A. Blake. *Canonical expressions in Boolean algebra*. PhD thesis, The University of Chicago, 1937.

- [9] A. Bochman. Actual causality in a logical setting. In J. Lang, editor, *Proceedings of the Twenty-Seventh International Joint Conference on Artificial Intelligence, IJCAI 2018, July 13-19, 2018, Stockholm, Sweden*, pages 1730–1736. ijcai.org, 2018.
- [10] R. Briggs. Interventionist counterfactuals. *Philosophical Studies*, 160(1):139–166, 2012.
- [11] H. Chockler, J. Y. Halpern, and O. Kupferman. What causes a system to satisfy a specification? *ACM Trans. Comput. Log.*, 9(3):20:1–20:26, 2008.
- [12] H. Chockler, O. Kupferman, and M. Y. Vardi. Coverage metrics for temporal logic model checking. In T. Margaria and W. Yi, editors, *Tools and Algorithms for the Construction and Analysis of Systems, 7th International Conference, TACAS 2001, Genova, Italy, April 2-6, 2001, Proceedings*, volume 2031 of *LNCS*, pages 528–542. Springer, 2001.
- [13] M. R. Clarkson, B. Finkbeiner, M. Koleini, K. K. Micinski, M. N. Rabe, and C. Sánchez. Temporal logics for hyperproperties. In M. Abadi and S. Kremer, editors, *Principles of Security and Trust - Third International Conference, POST 2014, Grenoble, France, April 5-13, 2014, Proceedings*, volume 8414 of *Lecture Notes in Computer Science*, pages 265–284. Springer, 2014.
- [14] M. R. Clarkson and F. B. Schneider. Hyperproperties. *J. Comput. Secur.*, 18(6):1157–1210, 2010.
- [15] N. Coenen, R. Dachsel, B. Finkbeiner, H. Frenkel, C. Hahn, T. Horak, N. Metzger, and J. Siber. Explaining hyperproperty violations. In S. Shoham and Y. Vizel, editors, *Computer Aided Verification - 34th International Conference, CAV 2022, Haifa, Israel, August 7-10, 2022, Proceedings, Part I*, volume 13371 of *LNCS*, pages 407–429. Springer, 2022.
- [16] N. Coenen, B. Finkbeiner, H. Frenkel, C. Hahn, N. Metzger, and J. Siber. Temporal causality in reactive systems. In A. Bouajjani, L. Holik, and Z. Wu, editors, *Automated Technology for Verification and Analysis - 20th International Symposium, ATVA 2022, Virtual Event, October 25-28, 2022, Proceedings*, volume 13505 of *LNCS*, pages 208–224. Springer, 2022.
- [17] N. Coenen, B. Finkbeiner, C. Hahn, and J. Hofmann. The hierarchy of hyperlogics. In *34th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2019, Vancouver, BC, Canada, June 24-27, 2019*, pages 1–13. IEEE, 2019.
- [18] R. Dimitrova, B. Finkbeiner, and H. Torfah. Probabilistic hyperproperties of markov decision processes. In D. V. Hung and O. Sokolsky, editors, *Automated Technology for Verification and Analysis - 18th International Symposium, ATVA 2020, Hanoi, Vietnam, October 19-23, 2020, Proceedings*, volume 12302 of *LNCS*, pages 484–500. Springer, 2020.
- [19] C. Dubslaff, K. Weis, C. Baier, and S. Apel. Causality in configurable software systems. In *44th IEEE/ACM 44th International Conference on Software Engineering, ICSE 2022, Pittsburgh, PA, USA, May 25-27, 2022*, pages 325–337. ACM, 2022.
- [20] T. Eiter and T. Lukasiewicz. Complexity results for structure-based causality. *Artif. Intell.*, 142(1):53–89, 2002.
- [21] I.-C. et al. IEEE standard for property specification language (PSL). *IEEE Std 1850-2010 (Revision of IEEE Std 1850-2005)*, pages 1–182, 2010.
- [22] B. Finkbeiner. Synthesis of reactive systems. *Dependable Software Systems Engineering*, 45:72–98, 2016.
- [23] B. Finkbeiner, C. Hahn, J. Hofmann, and L. Tentrup. Realizing omega-regular hyperproperties. In S. K. Lahiri and C. Wang, editors, *Computer Aided Verification - 32nd International Conference, CAV 2020, Los Angeles, CA, USA, July 21-24, 2020, Proceedings, Part II*, volume 12225 of *LNCS*, pages 40–63. Springer, 2020.
- [24] J. Y. Halpern. Axiomatizing causal reasoning. *J. Artif. Intell. Res.*, 12:317–337, 2000.
- [25] J. Y. Halpern. A modification of the halpern-pearl definition of causality. In Q. Yang and M. J. Wooldridge, editors, *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence, IJCAI 2015, Buenos Aires, Argentina, July 25-31, 2015*, pages 3022–3033. AAAI Press, 2015.
- [26] J. Y. Halpern and J. Pearl. Causes and explanations: A structural-model approach: Part 1: Causes. In J. S. Breese and D. Koller, editors, *UAI '01: Proceedings of the 17th Conference*

- in Uncertainty in Artificial Intelligence, University of Washington, Seattle, Washington, USA, August 2-5, 2001*, pages 194–202. Morgan Kaufmann, 2001.
- [27] J. Y. Halpern and J. Pearl. Causes and explanations: A structural-model approach. part i: Causes. *The British Journal for the Philosophy of Science*, 56(4):843–887, 2005.
- [28] J. Y. Halpern and J. Pearl. Causes and explanations: A structural-model approach. part ii: Explanations. *The British Journal for the Philosophy of Science*, 56(4):889–911, 2005.
- [29] J. Y. Halpern and S. Peters. Reasoning about causal models with infinitely many variables. In *Thirty-Sixth AAAI Conference on Artificial Intelligence, AAAI 2022 Virtual Event, February 22 - March 1, 2022*, pages 5668–5675. AAAI Press, 2022.
- [30] J. Harbecke. Counterfactual theories of causation and the problem of large causes. *Philosophical Studies*, 178(5):1647–1668, 2021.
- [31] D. Hume. *An Enquiry Concerning Human Understanding*. London, 1748.
- [32] K. Jothimurugan, R. Alur, and O. Bastani. A composable specification language for reinforcement learning tasks. In *NeurIPS 2019, December 8-14, 2019, Vancouver, BC, Canada*, pages 13021–13030, 2019.
- [33] O. Kupferman, W. Li, and S. A. Seshia. A theory of mutations with applications to vacuity, coverage, and fault tolerance. In A. Cimatti and R. B. Jones, editors, *Formal Methods in Computer-Aided Design, FMCAD 2008, Portland, Oregon, USA, 17-20 November 2008*, pages 1–9. IEEE, 2008.
- [34] M. J. Kusner, J. R. Loftus, C. Russell, and R. Silva. Counterfactual fairness. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017, December 4-9, 2017, Long Beach, CA, USA*, pages 4066–4076, 2017.
- [35] F. Leitner-Fischer and S. Leue. Causality checking for complex system models. In R. Giacobazzi, J. Berdine, and I. Mastroeni, editors, *Verification, Model Checking, and Abstract Interpretation, 14th International Conference, VMCAI 2013, Rome, Italy, January 20-22, 2013. Proceedings*, volume 7737 of *LNCS*, pages 248–267. Springer, 2013.
- [36] D. Lewis. Causation. *Journal of Philosophy*, 70(17):556–567, 1973.
- [37] D. K. Lewis. Completeness and decidability of three logics of counterfactual conditionals. *Theoria*, 37(1):74–85, 1971.
- [38] D. K. Lewis. *Counterfactuals*. Cambridge, MA, USA: Blackwell, 1973.
- [39] X. Li, C. I. Vasile, and C. Belta. Reinforcement learning with temporal logic rewards. In *2017 IEEE/RSJ International Conference on Intelligent Robots and Systems, IROS 2017, Vancouver, BC, Canada, September 24-28, 2017*, pages 3834–3839. IEEE, 2017.
- [40] P. Menzies and H. Beebe. Counterfactual Theories of Causation. In E. N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*. Metaphysics Research Lab, Stanford University, Winter 2020 edition, 2020.
- [41] P. J. Meyer, S. Sickert, and M. Luttenberger. Strix: Explicit reactive synthesis strikes back! In H. Chockler and G. Weissenbacher, editors, *Computer Aided Verification - 30th International Conference, CAV 2018, Oxford, UK, July 14-17, 2018, Proceedings, Part I*, volume 10981 of *LNCS*, pages 578–586. Springer, 2018.
- [42] A. Pnueli. The temporal logic of programs. In *18th Annual Symposium on Foundations of Computer Science, Providence, Rhode Island, USA, 31 October - 1 November 1977*, pages 46–57. IEEE Computer Society, 1977.
- [43] M. N. Rabe. *A temporal logic approach to information-flow control*. PhD thesis, Saarland University, 2016.
- [44] K. Y. Rozier and M. Y. Vardi. LTL satisfiability checking. *Int. J. Softw. Tools Technol. Transf.*, 12(2):123–137, 2010.
- [45] A. P. Sistla. *Theoretical Issues in the Design and Verification of Distributed Systems*. PhD thesis, Harvard University, 1983.

- [46] S. Wachter, B. D. Mittelstadt, and C. Russell. Counterfactual explanations without opening the black box: Automated decisions and the GDPR. *CoRR*, abs/1711.00399, 2017.
- [47] R. Ziemek, J. Piribauer, F. Funke, S. Jantsch, and C. Baier. Probabilistic causes in markov chains. *Innov. Syst. Softw. Eng.*, 18(3):347–367, 2022.

A Additional Preliminaries

A.1 HyperQPTL

The syntax of **HyperQPTL** [43, 17] is given as follows, where ψ is a QPTL formula. For sake of simplicity, we only present a fragment of **HyperQPTL** without alternation of propositional and trace quantification. This fragment suffices for the results presented in this work.

$$\chi ::= \exists\pi. \chi \mid \forall\pi. \chi \mid \psi .$$

The semantics of **HyperQPTL** are not defined over single traces, but sets of traces. The alphabet of atomic propositions $AP_{\text{Hyper}} = \{a_\pi \mid a \in AP \wedge \pi \in \mathcal{V}\}$, therefore, is indexed by trace variables. The satisfaction relation is defined with respect to a time point i , a set of traces T and a trace assignment $\Pi : \mathcal{V} \rightarrow T$ that maps trace variables to traces. To update the trace assignment so that it maps trace variable π to trace t , we write $\Pi[\pi \mapsto t]$. We further lift our replacement function $t[q \mapsto t_q]$ to sets of traces such that $T[q \mapsto t_q] = \{t[q \mapsto t_q] \mid t \in T\}$.

$$\begin{array}{lll}
\Pi, i \models_T a_\pi & \text{iff} & a \in \Pi(\pi)[i] \\
\Pi, i \models_T q & \text{iff} & \forall t \in T. q \in t[i] \\
\Pi, i \models_T \neg\varphi & \text{iff} & \Pi, i \not\models_T \varphi \\
\Pi, i \models_T \varphi \wedge \psi & \text{iff} & \Pi, i \models_T \varphi \text{ and } \Pi, i \models_T \psi \\
\Pi, i \models_T \mathbf{X}\varphi & \text{iff} & \Pi, i+1 \models_T \varphi \\
\Pi, i \models_T \varphi \mathbf{U}\psi & \text{iff} & \exists j \geq i. \Pi, j \models_T \psi \wedge \forall i \leq k < j. \Pi, k \models_T \varphi \\
\Pi, i \models_T \forall\pi. \varphi & \text{iff} & \forall t \in T. \Pi[\pi \mapsto t], i \models_T \varphi \\
\Pi, i \models_T \exists\pi. \varphi & \text{iff} & \exists t \in T. \Pi[\pi \mapsto t], i \models_T \varphi \\
\Pi, i \models_T \forall q. \varphi & \text{iff} & \forall t_q \in (2^{\{q\}})^\omega. \Pi, i \models_{T[q \mapsto t_q]} \varphi \\
\Pi, i \models_T \exists q. \varphi & \text{iff} & \exists t_q \in (2^{\{q\}})^\omega. \Pi, i \models_{T[q \mapsto t_q]} \varphi
\end{array}$$

A set of traces T satisfies a **HyperQPTL** formula φ iff $\emptyset, 0 \models_T \varphi$, which we also denote by $T \models \varphi$.

Example 19. *With **HyperQPTL**, we can relate traces to one another. For instance, the following formula characterizes sets where all pairs of traces (and hence all traces) have the same action sequence:*

$$\forall\pi. \forall\pi'. \mathbf{G}(u_\pi \leftrightarrow u_{\pi'}) \wedge \mathbf{G}(d_\pi \leftrightarrow d_{\pi'}) .$$

It is satisfied by the singleton set $\{t\}$ with trace t from Example 1, but not by $\{t, t'\}$ where $t' = (\{b, d\})^\omega$, since, e.g., the actions at the first position of the traces differ.

A.2 Actual Causality

Actual causality was originally proposed by Halpern and Pearl [26]. Several improvements have been appeared since, we consider the latest of these proposed by Halpern [25]. A *causal model* $\mathcal{M} = (\mathcal{S}, \mathcal{F})$ is defined by a *signature* \mathcal{S} and set of *structural equations* \mathcal{F} . A signature \mathcal{S} is a tuple (E, V, R) , where E is a set of *exogenous* variables, V is a set of *endogenous* variables,

and R defines the *range* of possible values $R(Y)$ for all variables $Y \in E \cup V$. For some context \vec{u} , the value of an exogenous variable is determined by factors outside of the model, while the value of some endogenous variable X is defined by the associated structural equation $f_X \in \mathcal{F}$.

Definition 20. $\vec{X} = \vec{x}$ is an actual cause of φ in (\mathcal{M}, \vec{u}) , if the following holds.

AC1: $(\mathcal{M}, \vec{u}) \models \vec{X} = \vec{x}$ and $(\mathcal{M}, \vec{u}) \models \varphi$, i.e., cause and effect are true in the actual world, and

AC2: There is a set \vec{W} of variables in V and a setting \vec{x}' of the variables in \vec{X} such that if $(\mathcal{M}, \vec{u}) \models \vec{W} = \vec{w}$, then $(\mathcal{M}, \vec{u}) \models [\vec{X} \leftarrow \vec{x}', \vec{W} \leftarrow \vec{w}] \neg \varphi$, and

AC3: \vec{X} is minimal, i.e. no subset of \vec{X} satisfies AC1 and AC2.

In the case of binary variables, we may denote a cause $\vec{X} = \vec{x}$ with $X = x^0 \wedge \dots \wedge x^n$, where for $0 \leq i \leq n$ a literal x^i is positive if it evaluates to 1 in \vec{x} , and negative if not. We denote by $V(x^i)$ the variable $v \in V$ associated with literal x^i .

B Proofs

Proposition 9. If \leq_W is a total order, then ‘Universal Would’ and ‘Would’ are equivalent, i.e., for all pairs of properties φ and ψ we have that $\varphi \boxrightarrow \psi \equiv \varphi \square \rightarrow \psi$.

Proof. We prove the equivalence by proving the entailment in each direction separately. So first assume that there exists a world W that satisfies $\varphi \boxrightarrow \psi$. We show that this world also satisfies $\varphi \square \rightarrow \psi$. Consider two cases: First let us assume $\varphi \boxrightarrow \psi$ is vacuously satisfied by W such that there exists no world W_1 with $W_1 \models \varphi$. Then W also vacuously satisfies $\varphi \square \rightarrow \psi$. As the second case, consider the non-vacuous case such that there exists at least one world W' that satisfies φ . Then, $\varphi \boxrightarrow \psi$ is satisfied because, for all worlds W_1 that satisfy φ , there exists an at least equally close world W_2 that also satisfies φ such that all closer worlds W_3 satisfy $\varphi \rightarrow \psi$. Hence, such a W_2 world exists in particular for W' , and can serve as a witness for the existential quantifier in the semantics of $\square \rightarrow$ (Condition 2 of Definition 5). Hence, $W \models \varphi \square \rightarrow \psi$.

For the entailment in the other direction, assume that there is a world W satisfying $\varphi \square \rightarrow \psi$. We show that this world also satisfies $\varphi \boxrightarrow \psi$. Again, if W satisfies $\varphi \square \rightarrow \psi$ vacuously it also satisfies $\varphi \boxrightarrow \psi$ vacuously. We, therefore, assume that there is a world W_1 such that $W_1 \models \varphi$ and all at least equally close worlds W_2 satisfy $\varphi \rightarrow \psi$. We now show that this is enough to satisfy Condition 2 of Definition 9. Pick any world W' as an instantiation of the outermost universal quantifier in that condition. If $W' \not\models \varphi$ it trivially satisfies the implication in the quantifiers body, so assume $W' \models \varphi$. Since \leq_W is a total order, we know that either $W' \leq_W W_1$ or $W_1 \leq_W W'$. If $W' \leq_W W_1$, we know that all smaller worlds $W'' \leq_W W'$ satisfy $\varphi \rightarrow \psi$ due to transitivity of \leq_W . Hence, W' is a witness for the existential quantifier in Condition 2 of Definition 9. Lastly, assume $W_1 \leq_W W'$. Then W_1 is the witness for the existential quantifier. In any case, we have $W \models \varphi \boxrightarrow \psi$ which closes this direction. \square

Corollary 11. If \leq_W is a total order, then ‘Existential Might’ and ‘Might’ are equivalent, i.e., for all pairs of properties φ and ψ we have that $\varphi \blacklozenge \rightarrow \psi \equiv \varphi \diamond \rightarrow \psi$.

Proof. The result follows directly from the two duality laws $\varphi \square \rightarrow \psi \equiv \neg(\varphi \diamond \rightarrow \neg\psi)$ and $\varphi \boxrightarrow \psi \equiv \neg(\varphi \blacklozenge \rightarrow \neg\psi)$, and Proposition 10, through substitution. \square

Lemma 14. *For a minimal counterfactual statement $\varphi \rightsquigarrow_{\min} \psi$, with $\rightsquigarrow \in \{\Box \rightarrow, \Diamond \rightarrow, \Box \rightarrow, \Diamond \rightarrow\}$, there exists a parameterized formula $\varphi_{FO}(W)$, with $W \in \mathcal{U}$, that quantifies only over worlds and not properties such that for all worlds $W' \in \mathcal{U}$: $W' \models \varphi \rightsquigarrow_{\min} \psi$ iff $\varphi_{FO}(W')$ is valid.*

Proof. The high-level idea of the proof is that, iff φ is not the minimal antecedent and there, therefore, exists a formula φ' characterizing a superset of the worlds characterized by φ , then there exists a set of worlds S that can be added to φ without changing the truth value of the counterfactual conditional. Such worlds fall into two categories: they are further away from the reference world W than the closest φ -worlds and can be added without considering whether they satisfy ψ , or they are at least equally close to W than the closest φ -world but also satisfy ψ so they could be added without harm and would constitute new closest φ -worlds. The exact relationships these worlds have to satisfy with respect to the closest φ -worlds differ between the operators, further, a difference is whether the reasoning has to be extended to incomparable worlds on other chains in the similarity relation. Therefore, we split the argumentation for the four minimal operators from now. We mark the world-variables that correspond to the closest φ -worlds of interest with W^c in each, and else use variables W' for W_1 to establish a close correspondence to definitions of the semantics. All other world variables are specific to the minimality reasoning. If not stated otherwise, quantifiers quantify over worlds from the universe \mathcal{U} . We have that W satisfies $\varphi \Box \rightarrow_{\min} \psi$ iff

$$((\forall W'. W' \not\models \psi) \wedge (\forall W^h \in \bar{\mathcal{U}}. \forall W' \in \mathcal{U}. W^h \models \varphi \wedge W' \not\models \varphi)) \vee (\exists W^c. W^c \models \varphi \wedge \forall W''). \quad (1)$$

$$(W'' \leq_W W^c \Rightarrow W'' \models \varphi \leftrightarrow \psi \wedge (W'' \models \psi \Rightarrow \forall W^i. W^i \not\leq_W W'' \Rightarrow W^i \models \varphi)) \wedge \quad (2)$$

$$(W'' \not\leq_W W^c \wedge W'' \models \psi \wedge (\exists W^n. W^n \leq_W W^c \wedge W^n \not\leq_W W'' \wedge W^n \models \neg \varphi) \Rightarrow \quad (3)$$

$$\exists W^p. W^p \leq_W W'' \wedge W^p \models \varphi) \quad (4)$$

is valid. The first disjunct in Line 1 encodes the fact that iff ψ is unsatisfiable with respect to the universe, then φ should characterize exactly the complement of the universe. The right disjunct encodes that we can enlarge the set characterized by φ based on the following conditions in relation to the counterfactual world W^c (corresponds to W_1 in Condition 2 of Definition 5):

- There is an at least equally close or closer world W'' such that ψ holds but this world is not included in φ yet (strengthens the semantic by $W'' \models \psi \rightarrow \varphi$ as in Line 2), or in W'' ψ holds but further away or incomparable worlds are not included in φ but could be (Rest of Line 2).
- There is a world at least as close as W^c which could take its place in a more minimal antecedent, i.e., W'' satisfies ψ , all its at-least-as-close worlds are a proper subset of the ones of W^c (Line 3) and are not included in φ (Line 4).

Next, we have that W satisfies $\varphi \Diamond \rightarrow_{\min} \psi$ iff

$$(\exists W'. W' \models \varphi) \wedge (\forall W'. W' \models \varphi \Rightarrow \exists W^c. W^c \leq_W W' \wedge (W^c \models \varphi \wedge \psi) \wedge \quad (5)$$

$$\forall W'''. (W^c \not\leq_W W''' \Rightarrow W''' \models \neg \varphi \rightarrow \neg \psi) \wedge (W^c \leq_W W''' \Rightarrow W''' \models \varphi)) \quad (6)$$

is valid. Here, we have a fairly direct strengthening of the semantics of regular ‘Might’ (Line 5): We now additionally require for φ -world W' on an infinite chain not only an at least equally closest world W^c (this time corresponds to W_2 in Condition 2 of Definition 7) that satisfies $\varphi \wedge \psi$, we also require that φ includes all smaller ψ -worlds and hence place $\neg \varphi \rightarrow \psi$ as a requirement on them. If this did not hold, we could include them in the property. Additionally, we could

include any world on a different chain (and which is hence not comparable to W^c) if they are not yet included in φ but satisfy ψ . Hence, such worlds also have to satisfy $\neg\varphi \rightarrow \psi$. Both is covered by the left conjunct in Line 6. Additionally, we again require that worlds further away than W^c are included in the property (right disjunct in Line 6). If they were not we could include them while still retaining W^c as a witness for all of them.

Next, we have that W satisfies $\varphi \square \rightarrow_{\min} \psi$ iff

$$((\forall W'. W' \not\models \psi) \wedge (\forall W^h \in \bar{\mathcal{U}}. \forall W' \in \mathcal{U}. W^h \models \varphi \wedge W' \not\models \varphi)) \vee \quad (7)$$

$$(\exists W^h. W^h \models \varphi \wedge \forall W'. W' \models \varphi \Rightarrow \exists W^c. W^c \leq_W W' \wedge W^c \models \varphi \wedge \forall W'''. \quad (8)$$

$$(W''' \leq_W W^c \Rightarrow W''' \models \varphi \leftrightarrow \psi) \wedge (W^c \leq_W W''' \Rightarrow W''' \models \varphi) \quad (9)$$

$$\wedge (W''' \not\leq_W W^c \wedge W^c \not\leq_W W''' \Rightarrow W''' \models \neg\varphi \rightarrow \neg\psi)) \quad (10)$$

is valid. This is now because of a combination from the reasoning of the previous two operators. Like for $\square \rightarrow_{\min}$, we again have that if ψ is unsatisfiable φ has to characterize the complement of the universe (Line 7). If this is not the case, then there has to be at least one W^h world satisfying φ (Line 8). We have to ensure this because φ may only characterize an effectively empty set if there exists no ψ -world. Further, for any φ -world there exists a φ -world W^c closest to W such that on any closer worlds $\varphi \rightarrow \psi$. Again, like for $\square \rightarrow_{\min}$, we add the inverse direction $\psi \rightarrow \varphi$ to ensure that φ cannot be enlarged in this direction (left in Line 9). Further, any world further away than W^c has to be included in φ as they can all use W^c as a closest φ -world (right in Line 9). Lastly, like for $\diamond \rightarrow_{\min}$, if some world incomparable to W^c which hence is on a different chain satisfies ψ , it also has to be included in φ which we ensure by requiring $\neg\varphi \rightarrow \neg\psi$ (Line 10).

Lastly, we have that W satisfies $\varphi \diamond \rightarrow_{\min} \psi$ iff

$$\exists W'. W' \models \varphi \wedge \forall W''. W'' \leq_W W' \wedge W'' \models \varphi \Rightarrow \exists W^c. W^c \leq_W W'' \wedge (W^c \models \varphi \wedge \psi) \wedge \quad (11)$$

$$\forall W^h. (W^h \leq_W W^c \Rightarrow W^h \models \neg\varphi \rightarrow \neg\psi) \wedge (W^h \not\leq_W W^c \Rightarrow W^h \models \varphi \wedge W^h \models \psi) \wedge \quad (12)$$

$$(\exists W^n. W^n \leq_W W^c \wedge W^n \not\leq_W W^h \wedge W^n \models \neg\varphi) \Rightarrow \exists W^p. W^p \leq_W W^h \wedge W^p \models \varphi) \quad (13)$$

is valid. In Line 11, we have the usual semantics of ‘Existential Might’. We now additionally require that worlds at least as close as the closest φ -worlds W^c that satisfy ψ , which correspond to W_3 in Definition 10, are included in φ if they satisfy ψ (left conjunct in Line 12). Further, we require all worlds farther away or incomparable to W^c to be included in φ (right in conjunct in Line 12). Similar to the reasoning for $\square \rightarrow_{\min}$, we have to ensure that these worlds W^h do not induce a more minimal antecedent by qualifying as a W^c by satisfying ψ (end of Line 12) and having a proper subset of at-least-equally-close worlds (Line 13). \square

Theorem 15. *For any QPTL_{cf} formula φ , it is decidable to check whether φ is satisfiable when the similarity relation is from the family \leq^{QPTL} and the universe is defined by a QPTL formula.*

Proof. We sketch how to encode the satisfiability of our QPTL_{cf}-formulas in a HyperQPTL model-checking problem over the most general model \mathcal{M} that contains all traces over our alphabet 2^{AP} . Let $\varphi^{\mathcal{U}}$ be the formula encoding the universe. The semantics of the non-minimal operators contain quantification over all possible traces over our alphabet 2^{AP} and hence over the traces from the model \mathcal{M} . In Lemma 14 we showed that the semantics of the minimal operators can similarly be expressed by quantification over traces from \mathcal{M} and in $\varphi^{\mathcal{U}}$ as well. We now illustrate at the example of $\diamond \rightarrow_{\min}$ how such a formula can be encoded into a HyperQPTL-formula

$\varphi_{FO}^{\text{QPTL}}(\pi)$ in prenex normal form, as even if $AP = \emptyset$, there exists at least one trace. Let φ_R be the QPTL-formula for the similarity metric $R \in \leq^{\text{QPTL}}$. For some formula φ and trace variable $\pi \in \mathcal{V}$, φ_π denotes the same formula where all atomic propositions are indexed with trace variable π . Then $\varphi_{FO}^{\text{QPTL}}(\pi)$ for some $\varphi \diamond \rightarrow_{\min} \psi$ is of the following form:

$$\varphi_{FO}^{\text{QPTL}}(\pi) = \exists \pi_1. \forall \pi_2. \exists \pi_3. \forall \pi_4. \varphi_{\pi_1}^{\mathcal{U}} \wedge \varphi_{\pi_1} \wedge \left(\varphi_{\pi_2}^{\mathcal{U}} \rightarrow \left(\varphi_{\pi_2} \rightarrow \varphi_{\pi_3}^{\mathcal{U}} \wedge \varphi_R(\pi, \pi_3, \pi_2) \wedge \varphi_{\pi_3} \wedge \psi_{\pi_3} \wedge \left(\varphi_{\pi_4}^{\mathcal{U}} \rightarrow (\neg \varphi_R(\pi, \pi_3, \pi_4) \rightarrow (\neg \varphi_{\pi_4} \rightarrow \neg \psi_{\pi_4})) \wedge (\varphi_R(\pi, \pi_3, \pi_4) \rightarrow \varphi_{\pi_4}) \right) \right) \right).$$

We restrict quantification to the universe with conjunctions and implications containing $\varphi^{\mathcal{U}}$. Quantification over the complement of the universe uses $\neg \varphi^{\mathcal{U}}$ instead.

For the original QPTL_{cf}-formula φ , we apply this transformation to any counterfactual conditional appearing in it and again transform the result into prenex normal form. This is possible while retaining equisatisfiability because temporal operators appear only on the lowest level, i.e., no quantifier appears in the scope of a temporal operator. We denote the resulting formula with $\varphi_{FO}^{\text{full}}(\pi)$ and put a π -subscript on any top-level QPTL formula (i.e., that was not in the body of a counterfactual operator).

We can then model check the formula HyperQPTL-formula $\exists \pi. \varphi_{FO}^{\text{full}}(\pi)$ on the model \mathcal{M} to check whether there exists a trace satisfying our original QPTL_{cf}-formula. Since HyperQPTL model checking is decidable, so is the satisfiability checking of QPTL_{cf}. This approach has similarities to the idea of reducing the satisfiability problem of linear-time temporal logic (LTL) to model checking of the same logic [44]. Note that, crucially, our approach is possible because the models of our logic QPTL_{cf} are traces and not sets of traces as in HyperQPTL, which allows an encoding into the decidable model checking problem while avoiding the undecidable problem of HyperQPTL satisfiability checking. \square

Theorem 16. *For a lasso-shaped trace t and a QPTL_{cf} formula φ , it is decidable to check whether $t \models \varphi$ when the similarity relation is from the family \leq^{QPTL} and the universe is defined by a QPTL formula.*

Proof. (Sketch) Since the trace t is lasso-shaped, we encode it in a formula φ^t that characterizes the set $\{t\}$. We can then use the ideas of the satisfiability proof above to solve the trace checking problem. However, we now model check the formula $\forall \pi. \varphi_\pi^t \rightarrow \varphi_{FO}^{\text{full}}(\pi)$ on the model \mathcal{M} , not searching for any trace satisfying the formula but instead fixing the reference trace to t . \square

Theorem 17. *$X_0 = x_0^0 \wedge \dots \wedge x_0^k, \dots, X_n = x_n^0 \wedge \dots \wedge x_n^j$ are actual causes of φ in (\mathcal{M}, \vec{u}) , iff $t_c^{\mathcal{M}} \models \varphi_X \wedge \varphi \wedge \neg \varphi_X \diamond \rightarrow_{\min} \neg \varphi$ in the universe $\mathcal{U}(\mathcal{M}, c)$ with respect to the similarity relation $\leq_t(\{i_x \mid x \in V\})$ that tracks only the active interventions between traces, where*

$$\neg \varphi_X = (\neg x_0^0 \wedge \dots \wedge \neg x_0^k) \vee \dots \vee (\neg x_n^0 \wedge \dots \wedge \neg x_n^j) \text{ is in Blake canonical form [8].}$$

Proof. This is effectively a finite setting modeled on the first position of a trace, such that the Limit Assumption holds and for any formula $\neg \varphi_X$ without temporal operators, there exists a unique (finite) set of closest traces from $t_c^{\mathcal{M}}$ that satisfy it. Let $I = \{i_x \mid x \in V\}$ denote the set of intervention variables. We will use the following auxiliary results throughout the proof:

1. For any actual cause $X_0 = x_0^0 \wedge \dots \wedge x_0^k$ for effect φ , there is a contingency W such that $(\mathcal{M}, \vec{u}) \models [X \leftarrow \neg x_i^0 \wedge \dots \wedge \neg x_i^k, W \leftarrow w_p^0 \wedge \dots \wedge w_p^q] \neg \varphi$, i.e., the intervention that flips all the variables is a witnessing intervention for AC2. With Boolean variables, this follows directly from minimality, as any variable in an intervention that is not negated could be moved to the contingency W instead, yielding a more minimal cause.

2. There is a direct mapping from interventions that flip all values to intervention variables on traces: By induction over the number of structural equations we can show that for all traces $t \in \mathcal{U}(\mathcal{M}, c)$ and variables $v \in V$, we have $(\mathcal{M}, \vec{u}) \models [X \leftarrow x_i^0 \wedge \dots \wedge x_i^k, W \leftarrow w_p^0 \wedge \dots \wedge w_p^q]v$ where for all $x_i^r \in X : (\mathcal{M}, \vec{u}) \models \neg x_i^r$, iff $v \in t[0]$, where t is the unique trace that has all the contingency variables of W enabled at the first position, i.e., for all $0 \leq r \leq q : V(w_p^r) \in t[0]$, and all intervention variables corresponding to literals in X are also enabled: for all $0 \leq r \leq k : (x_p^k \leftrightarrow V(x_p^k) \in t_c^{\mathcal{M}}) \rightarrow i_{V(x_p^k)} \in t[0]$.
3. Let t_{cf} be a closest $\neg\varphi_X$ -trace to $t_c^{\mathcal{M}}$ such that $t_{cf} \models (\neg x_z^0 \wedge \dots \wedge \neg x_z^j)$ for some conjunction and $t_{cf} \models \neg\varphi$. We can use induction over the length of the conjunction to show that for all $i \in I : i \in t_{cf}[0]$ iff there is some $0 \leq r \leq z : i_{V(x_z^r)} = i \wedge c_{V(x_z^r)} \notin t_{cf}[0]$, i.e., exactly the intervention variables corresponding to literals in the conjunction are enabled on t_{cf} , while all corresponding contingency variables are disabled.

We now proceed to proof the two directions of the equivalence separately.

“ \Rightarrow ”: We know that no intervention or contingency variables are true on $t_c^{\mathcal{M}}$ by construction. With Result 2 and the fact that X_0, \dots, X_n are actual causes of φ in (\mathcal{M}, \vec{u}) , it follows from AC1 that $t_c^{\mathcal{M}} \models \varphi_X \wedge \varphi$. To show that $\neg\varphi_X \diamond_{\rightarrow_{\min}} \neg\varphi$ is satisfied, consider that from Result 2 and Result 3, it follows the closest traces from satisfying $\neg\varphi_X$ have to differ in all of the values in at least one of the conjuncts of φ_X , w.l.o.g. assume this to be $(x_0^0 \vee \dots \vee x_0^n)$. With Result 1 we have $(\mathcal{M}, \vec{u}) \models [X_0 \leftarrow \neg x_i^0 \wedge \dots \wedge \neg x_i^k, W_0 \leftarrow w_p^0 \wedge \dots \wedge w_p^q] \neg\varphi$, and with Result 2 it then follows that there is a trace t such that $t \models \neg\varphi$, and from Result 3 it follows that t is a closest $\neg\varphi_X$ -trace. Hence, We have $t_c^{\mathcal{M}} \models \neg\varphi_X \diamond_{\rightarrow} \varphi$. To show that $\neg\varphi_X$ is also a minimal antecedent, consider what would happen if a trace was added to $\mathcal{L}(\neg\varphi_X)$. Such a trace t'' has to satisfy φ_X , and hence satisfies one of the literals in each of that formula’s conjuncts. That means that the changes between $t_c^{\mathcal{M}}$ and t'' are either incomparable to any other closest $\neg\varphi_X$ -trace, or a proper subset. In the former case, this induces another actual cause, in the latter case, this would mean one of the causes is not minimal. In both cases, we have a contradiction. Note that $\neg\varphi_X$ is in Blake canonical form because all of the causes are minimal, and hence prime implicants, i.e., no X_v implies another X_w , and this also holds for their negations in $\neg\varphi_X$. This closes this direction

“ \Leftarrow ”: We first show AC1 holds for all of the causes. We know that $(\mathcal{M}, \vec{u}) \models \varphi$ from $t_c^{\mathcal{M}} \models \varphi$ and Result 2. To show that $t_c^{\mathcal{M}} \models (x_i^0 \wedge \dots \wedge x_i^u)$ for all X_i , consider t_{cf} as defined in Result 3. Since we know from the proof of Result 3 that on t_{cf} exactly the intervention variables corresponding to the literals in the conjunction are set, and intervention variables flip the value of variables with respect to $t_c^{\mathcal{M}}$, we can deduce that all the literals occurring in the conjunction have the inverted value in $t_c^{\mathcal{M}}$, which proves AC1 for all of the causes using Result 2. AC2 follows directly from $t_c^{\mathcal{M}} \models \neg\varphi_X \diamond_{\rightarrow_{\min}} \varphi$, Result 3 which states on some closest trace $t_{cf} \models \neg\varphi$ all literals $x_z^0 \wedge \dots \wedge x_z^j$ are intervened upon, and Result 2 which relates the intervention and contingency variables $w_z^0 \wedge \dots \wedge w_z^q$ of t_{cf} to the causal model (\mathcal{M}, \vec{u}) . Then, we can deduce that $(\mathcal{M}, \vec{u}) \models [X \leftarrow \neg x_z^0 \wedge \dots \wedge \neg x_z^j, W \leftarrow w_z^0 \wedge \dots \wedge w_z^q] \neg\varphi$. Lastly, we can show AC3 by contradiction: If one of the causes was not minimal, then Result 2 and Result 3 imply there is closer $\neg\varphi$ -trace, and therefore $\neg\varphi_X$ is not a minimal antecedent, a contradiction. \square