On the Complexity of Conditional Logics*

Nir Friedman
Department of Computer Science
Stanford University
Stanford, CA 94305-2140
nir@cs.stanford.edu

Joseph Y. Halpern
IBM Almaden Research Center
650 Harry Road
San Jose, CA 95120–6099
halpern@almaden.ibm.com

Abstract

Conditional logics, introduced by Lewis and Stalnaker, have been utilized in artificial intelligence to capture a broad range of phenomena. In this paper we examine the complexity of several variants discussed in the literature. We show that, in general, deciding satisfiability is PSPACE-complete for formulas with arbitrary conditional nesting and NP-complete for formulas with bounded nesting of conditionals. However, we provide several exceptions to this rule. Of particular note are results showing that (a) when assuming *uniformity* (i.e., that all worlds agree on what worlds are possible), the decision problem becomes EXPTIME-complete even for formulas with bounded nesting, and (b) when assuming *absoluteness* (i.e., that all worlds agree on all conditional statements), the decision problem is NP-complete for formulas with arbitrary nesting.

INTRODUCTION

The study of *conditional* statements of the form "If ... then ..." has a long history in philosophy (Sta68; Lew73; Che80; Vel85). In recent years these logics have been applied in artificial intelligence to capture nonmonotonic inference (Del88; Bel89; KLM90; Bou92), belief change (Gra91; Bou92), counterfactual reasoning (Gin86), qualitative probabilities (Pea89; GP92), and intentions and desires (Pea93; Bou94). In general, conditional logics provide a logical language to reason about structures that contain some sort of ordering. In this paper we present complexity results for a family of conditional logics introduced by Lewis (Lew73; Lew74). We also provide an overview of a completeness proof which substantially simplifies previous proofs in the literature (Bur81).

Lewis's construction starts with a set W of *possible worlds*, each one describing a possible way the world might be. We associate with each possible world $w \in W$ a preorder $\preceq_w$ over a subset W_w of W. Intuitively, W_w is the set of worlds considered possible at w. There are a number of differing intuitions for what is being represented by the $\preceq_w$ relation. For example, in *counterfactual* reasoning, $\preceq_w$ is viewed as capturing a measure of distance from w, so that $w' \preceq_w w''$ if w' is more similar or closer to w than w'' is. In this variant it is usually assumed (Lew73) that the real world is closest to itself. In *nonmonotonic* reasoning the $\preceq_w$ relation captures an agent's plausibility ordering on the worlds, so that $w' \preceq w''$ if w' is more plausible than w'' according to the agent's beliefs in w. Typically (although not, for example, in (FH94a)) it is assumed that the agent's beliefs are the same in all the worlds in W, so that $\preceq_w$ is independent of w. The $\preceq_w$ relation is used to give semantics to conditional formulas of the form $\varphi \rightarrow \psi$; such a formula is taken to be true at a world w if all the $\preceq_w$-minimal worlds satisfying φ also satisfy ψ.

As these examples suggest, we can construct a number of different logics, depending on the assumptions we make about $\preceq_w$. In this paper, we focus on the following assumptions (all of which have been considered before (Lew73; Bur81; Gra91; KS91)), which apply to all $w \in W$:[1]

N *Normality*: $W_w \neq \emptyset$.

R *Reflexivity*: $w \in W_w$.

T *Centering*: w is a minimal element in W_w, i.e., for all $w' \in W_w$, we have $w \preceq_w w'$.[2]

U *Uniformity*: W_w is independent of w, i.e., for all $w' \in W_w$, $W_{w'} = W_w$.

A *Absoluteness*: $\preceq_w$ is independent of w, i.e., for all $w' \in W_w$, $W_{w'} = W_w$ and for all $w_1, w_2 \in W_w$, we have $w_1 \preceq_{w'} w_2$ if and only if $w_1 \preceq_w w_2$.[3]

*This paper appears in *Proceedings of the 4th International Conference on Principles of Knowledge Representation and Reasoning (KR '94)*, 1994, pp. 202–213

[1] Whenever possible we adopt the naming scheme used by Lewis (Lew73, pp. 120).

[2] Our notion of centering is that used by Lewis (Lew73). Other authors (KS91; Gra91) assume the stronger condition of *strict centering*, that is w is the only minimal world in $\preceq_w$. Our results for centering apply with minor technical modifications to strict centering.

[3] Lewis (Lew73) distinguishes between a *local* definition of uniformity and absoluteness and a *global* one. We adopt the local one (i.e., "for all $w' \in W_w$ ...", rather than "for

C *Connectedness*: all worlds in W_w are comparable according to $\preceq_w$; i.e., for all $w_1, w_2 \in W_w$, either $w_1 \preceq_w w_2$ or $w_2 \preceq_w w_1$.

Notice that centering implies reflexivity, which in turn implies normality. Normality is a minimal assumption, typically made in almost all applications of conditional logics. As we mentioned earlier, centering is typically assumed in counterfactual reasoning, while absoluteness is typically assumed in nonmonotonic reasoning. Uniformity is assumed when, for example, the set of possible worlds is taken to be the set of all logically possible worlds (i.e., the set of all truth assignments). Combinations of these conditions are used in the various applications of conditional logics. For example, Boutilier's (Bou92) work in nonmonotonic reasoning assumes absoluteness and considers variants satisfying connectedness; similar assumptions are made in (KLM90; GP92; Bel89). Works on counterfactuals (such as Grahne's (Gra91)) typically assume centering and uniformity. Katsuno and Satoh (KS91) consider variants satisfying absoluteness, centering and connectedness.

Completeness results have been obtained for the logics corresponding to various combinations of these constraints (Lew73; Lew74; Bur81). While we do present completeness proofs here, using a proof that is substantially simpler than that of (Bur81), our focus is on complexity-theoretic issues.

Burgess (Bur81) shows that any satisfiable conditional formula is satisfiable in a finite structure. The structures he obtains are of nonelementary size.[4] To obtain our complexity results we prove that if a formula is satisfiable at all, it can be satisfied in a much smaller structure. We start by showing that a formula without nested conditionals is satisfiable if and only if it is satisfiable in a polynomial-sized structure. Applying the construction for formulas without nested conditionals recursively, we show that, in general, a satisfiable formula with bounded nesting depth is satisfiable in a polynomial-sized structure, and an arbitrary satisfiable formula is satisfiable in an exponential-sized structure. In most variants, this structure takes the form of a tree, where each level of the tree corresponds to one level of nesting. We show that checking whether such a tree-like structure exists can be done in polynomial space, without explicitly storing the whole tree in memory. This gives a PSPACE upper bound for the satisfiability problem for most variants of the logic.[5]

Can we do better? In general, no. We show that an appropriate modal logic (either K,D or T depending on the variant in question) can be embedded in most variants of the logic.[6] The result then follows from results of Ladner (Lad77; HM92) on the complexity of satisfiability for these logics. There are exceptions to the PSPACE results. For one thing, it already follows from our "small model" results that for bounded-depth formulas (in particular, depth-one formulas) satisfiability is NP-complete. Moreover, in the presence of absoluteness, every formula is equivalent to one without nesting, so we can again get NP-completeness. Interestingly, the appropriate modal logic in the presence of absoluteness in the lower bound construction mentioned above is S5, whose satisfiability problem is also NP-complete (Lad77). On the other hand, while the assumption of uniformity seems rather innocuous, and much in the spirit of absoluteness, assuming uniformity without absoluteness leads to an EXPTIME-complete satisfiability problem, even for formulas with bounded nesting.

Our results form an interesting contrast to those of Eiter and Gottlob (EG92; EG93) and Nebel (Neb91) for a framework for *counterfactual* queries defined by Ginsberg (Gin86), using an approach that goes back to Fagin, Ullman, and Vardi (FUV83). In this framework a conditional query $p>q$ is evaluated by modifying the knowledge base to include p and then checking whether q is entailed. As shown by Nebel (Neb91) and Eiter and Gottlob (EG92), for formulas without nested conditionals, evaluating such a query is Π_2^p-complete.[7] Roughly speaking, the reason for the higher complexity is that once we prove an analogous small model theorem for this more syntactic approach, checking that a formula is entailed by a theory is co-NP hard, while in our case, checking that a formula is satisfied in a small structure can be done in polynomial time. Eiter and Gottlob (EG93) show that if we restrict to right-nested formulas, without negations of nested conditionals, then queries are still Π_2^p complete. Finally, Eiter and Gottlob show that once we move beyond simple right-nesting, the problem becomes PSPACE-hard; the complexity of queries for the full language is not known. In contrast to these results, we show that the language of simple right-nested conditionals is NP-complete, and when negations are allowed, it becomes PSPACE-complete.

The rest of the paper is organized as follows: In Section 2 we formally define the logical language and its semantics. In Section 3 we prove small model theorems for the different variants. In Section 4 we prove

all $w' \in W \ldots$"), but it is easy to see that all our results, including the axiomatization, also apply to the global definition with essentially no change.

[4] Roughly speaking, a nonelementary function of n is of the form $2^{2^{2^{\cdots}}}$, where the height of the stack of 2's is on the order of n.

[5] We assume some familiarity with complexity theory, especially with the complexity classes NP, PSPACE, and EXPTIME. See Section for a review of these complexity-theoretic notions.

[6] We assume familiarity with modal logic, especially the logics K, D and T. See (HM92) for an overview of these logics and their axiomatizations.

[7] Π_2^p is the complexity class that is characterized by decision problems that can be determined in polynomial time given an NP oracle. This class is believed to be harder than NP, but simpler than PSPACE.

the complexity results. In Section 5 we provide an axiomatization for each of the logics we consider and sketch a completeness proof.

CONDITIONAL LOGIC

The syntax of the logic is simple: we start with a set Φ of primitive propositions, and close off under $\wedge$, $\neg$, and $\rightarrow$ (where $\rightarrow$ is the conditional operator). We call the resulting language $\mathcal{L}^C$. We denote by $\mathcal{L}^C_k$ the sublanguage of $\mathcal{L}^C$ with *bounded nesting*, i.e., formulas in $\mathcal{L}^C$ with no more than k level of nested conditionals. For example, $\mathcal{L}^C_0$ contains propositional formulas without any conditional sentences, and $\mathcal{L}^C_1$ contains $p \rightarrow q$ but not $p \rightarrow (q \rightarrow r)$. Of course, we define the propositional connectives $\vee$, $\Rightarrow$ (material implication), and $\Leftrightarrow$ (logical equivalence) in terms of $\wedge$ and $\neg$ in the standard way.

We use the semantic representation suggested by Lewis to capture conditionals (Lew73; Bur81): A *structure* M is a tuple (W, π, R), such that W is a set of *possible worlds*, π maps each possible world to a truth assignment over Φ, and R is a ternary relation over W. We think of the possible worlds as different ways the world could be, or the different situations we might be in. The relation R is a preorder on worlds: $(w, u, v) \in R$ if u is as close/preferred/plausible as v when the real world is w. We use the notation $u \preceq_w v$ to denote that $(w, u, v) \in R$. We define $W_w = \{u | u \preceq_w v \in R \text{ for some } v \in W\}$; thus, the worlds in W_w are those that are at least as plausible as some world in W according to $\preceq_w$. We require that $\preceq_w$ be a preorder, i.e., a reflexive and transitive relation, on W_w. As usual, we define $u \prec_w v$ if $u \preceq_w v$ and not $v \preceq_w u$.

We now provide semantics for formulas in $\mathcal{L}^C$. The truth of a propositional formula in a world w is determined by the truth assignment $\pi(w)$. The truth of a conditional formula is determined by the ordering $\preceq_w$. The intuition is that $\varphi \rightarrow \psi$ holds at w if all the minimal (e.g., closest, most plausible) φ-worlds satisfy ψ (where a φ-world, of course, is a world where φ is true). Unfortunately, if W is infinite, it may not have minimal φ-worlds. Thus, the actual definition we use, which is standard in the literature (see (Lew73; Bur81; Bou92)), is more complicated. Roughly speaking, $\varphi \rightarrow \psi$ is true if, from a certain point on, whenever φ is true, so is ψ. More precisely, $\varphi \rightarrow \psi$ is true at w if for every φ-world u in W_w, there is another world v such that (a) v is at least as plausible as u, (b) v satisfies $\varphi \wedge \psi$, and (c) each φ-world that is at least as plausible as v is also a ψ-world. It is easy to see that if W_w is finite, then this is equivalent to saying that the minimal φ-worlds in W_w satisfy ψ.

Formally, we define the truth of $\varphi \in \mathcal{L}^C$ at a world w in a structure $M = (W, \pi, R)$ recursively:

- $(M, w) \models p$, when $p \in \Phi$, if $\pi(w)(p) = \mathbf{true}$.
- $(M, w) \models \varphi \wedge \psi$ if $(M, w) \models \varphi$ and $(M, w) \models \psi$.
- $(M, w) \models \neg\varphi$ if it is not the case that $(M, w) \models \varphi$.
- $(M, w) \models \varphi \rightarrow \psi$ if for any world $u \in W_w$, if $(M, u) \models \varphi$ then there is a world v, such that $v \preceq_w u$ and $(M, v) \models \varphi \wedge \psi$ and there is no $v' \preceq_w v$ such that $(M, v) \models \varphi \wedge \neg\psi$.

We say that φ is *valid in M* (resp., *satisfiable in M*) if $(M, w) \models \varphi$ for all worlds w (resp., some world w) in M.

We define the set of all possible structures as $\mathcal{M}$. For each combination of the constraints defined in the introduction, we define the corresponding class of structures satisfying them. For example, $\mathcal{M}^{N,T,U}$ is the class of all structures satisfying normality, centering and uniformity. For $\mathcal{A} \subseteq \{N, R, T, A, U, C\}$, we say that a formula φ is *valid* with respect to $\mathcal{M}^{\mathcal{A}}$, written $\mathcal{M}^{\mathcal{A}} \models \varphi$, if φ is valid in every structure $M \in \mathcal{M}^{A}$. Similarly, we say that φ is satisfiable in $\mathcal{M}^{\mathcal{A}}$ if it is satisfiable in some structure $M \in \mathcal{M}^{\mathcal{A}}$.

SMALL MODEL THEOREMS

In this section we provide small model theorems for the logics we examine, showing that if a formula φ is satisfiable, than it is satisfiable in a structure of bounded size. These results play a crucial role in our complexity considerations.

We start with some definitions. Given a formula $\varphi \in \mathcal{L}^C$, we define $Sub(\varphi)$ to be the set of all subformulas of φ and $Sub^+(\varphi) = Sub(\varphi) \cup \{\neg\psi \mid \psi \in Sub(\varphi)\}$. Finally, let $Sub_C(\varphi)$ consist of all formulas in $Sub(\varphi)$ of the form $\varphi \rightarrow \psi$. It is easy to verify that $|Sub(\varphi)|$ (the number of formulas in $Sub(\varphi)$) is at most $|\varphi|$ (the length of φ, viewed as a string of symbols).

We begin by examining formulas without nested conditionals. The first case is when φ is a conjunction of a number of (non-negated) conditional statements and one negated conditional statement. This case will serve as a basis for the general case.

Lemma 0.1: *Let $\varphi = \neg(\psi_0 \rightarrow \psi'_0) \wedge \bigwedge_{i=1}^{k} (\psi_i \rightarrow \psi'_i)$ where $\psi_i, \psi'_i \in \mathcal{L}^C_0$. If φ is satisfiable in $\mathcal{M}$, then φ is satisfiable in a structure in $\mathcal{M}$ with at most $k+1$ worlds which are totally ordered by $\preceq$.*

Proof: Assume we are given $M \in \mathcal{M}$ and w such that $(M, w) \models \varphi$. From the results of (Bur81), it follows that, without loss of generality, we can assume that M is finite (i.e., that M has only finitely many worlds). Since $(M, w) \models \neg(\psi_0 \rightarrow \psi'_0)$, there is a a world w_0 such that w_0 is a minimal ψ_0-world in $\preceq_w$ and satisfies $\neg\psi'_0$. Let $\leq$ be a total order over W_w that is compatible with $\preceq_w$, in that if $w_1 \prec_w w_2$ then $w_1 < w_2$, such that $w_0 < w'$ for any $w' \neq w_0$ satisfying ψ_0. (Since $\leq$ is a total order, if $w_1 \neq w_2$, then either $w_1 < w_2$ or $w_2 < w_1$.) Let w_i be the minimal ψ_i-world in W_w according to $\leq$, if there is a ψ_i-world in W_w, and w_0 otherwise.

We now construct a new structure $M' = (W', \pi', R')$. Let $W' = \{w_0, \ldots, w_k\}$, let π' be the restriction of π to W', and let R' be such that for all $w' \in W'$, we have $W'_{w'} = W'$ and $w_i \preceq'_{w'} w_j$ if and only if $w_i \leq w_j$. It is easy to verify that $(M', w') \models \varphi$ for all $w' \in W'$, since if w_i is the minimal ψ_i-world according to $\leq$, then by the construction of $\leq$, it must be a minimal ψ_i-world in M according to $\preceq_{w'}$ and thus must also satisfy ψ'_i, while w_0 is the minimal ψ_0-world and thus $(M', w') \models \neg(\psi_0 \rightarrow \psi'_0)$. ∎

We now use this construction to prove that any formula without nesting is satisfiable if and only if it is satisfiable in a polynomial structure.

Proposition 0.2: *Let $\varphi \in \mathcal{L}_1^C$. If φ is satisfiable in $\mathcal{M}$, then φ is satisfiable in a structure in $\mathcal{M}$ with at most $O(|\mathrm{Sub}_C(\varphi)|^2)$ worlds.*

Proof: Suppose that $M \in \mathcal{M}$ and that $(M, w) \models \varphi$. Again, by Burgess's result, we can assume without loss of generality that M is finite. Our goal is to construct a small structure M' such that for each formula $\psi \in Sub(\varphi)$, we have $(M, w) \models \psi$ if and only if $(M', w) \models \psi$. It clearly suffices to do this for the primitive propositions and the formulas in $Sub_C(\varphi)$. We cannot use the construction of the previous lemma directly, because we may now have to deal with more than one negated conditional. For example, if $(M, w) \models \neg(p \rightarrow q) \wedge \neg(p \rightarrow \neg q)$, the structure M' we construct must have a minimal p-world satisfying q and a minimal p-world satisfying $\neg q$. This cannot be done by using one total order, as was done in the previous lemma.

We solve the problem by considering the union of several total orders, one for each negated conditional. Let $Neg = \{\psi \rightarrow \xi \in Sub_C(\varphi) : (M, w) \models \neg(\psi \rightarrow \xi)\}$ and let $Pos = \{\psi \rightarrow \xi \in Sub_C(\varphi) : (M, w) \models \psi \rightarrow \xi\}$. Suppose $Neg = \{\psi_1^n \rightarrow \xi_1^n, \ldots, \psi_k^n \rightarrow \xi_k^n\}$. From Lemma 0.1, it follows that for each formula $\psi_i^n \rightarrow \xi_i^n \in Neg$, we can construct a structure M_i whose set of worlds W_i has size at most $|Sub_C(\varphi)|$, such that M_i satisfies $\neg(\psi_i^n \rightarrow \xi_i^n)$ and all the formulas in Pos. This gives us $|Neg|$ structures, one for each formula in Neg. Without loss of generality, we can assume that the sets W_i are disjoint and do not contain w. The structure we are interested in is essentially the disjoint union of the structures M_i. More precisely, we take $M' = (W', \pi', R')$, where W' is the union of the sets W_i for $1 \leq i \leq |Neg|$, together with w. We define π' to be such that for each world in $w' \in W'$, the truth assignment $\pi'(w')$ is the same as the truth assignment in the structure that w' was drawn from. Finally, we define R' so that for all $w' \in W'$, we have $W'_{w'} = W' - \{w\}$, and $\preceq_{w'}$ is the union of the orderings in the structure $M_{\psi \in \psi'}$. (We have defined $\preceq_{w'} = \preceq_w$ for all $w' \in W'$, but this was not necessary. Since we are dealing with depth-one nesting here, all that matters in the proof is the definition of $\preceq_w$. We can redefine $\preceq_{w'}$ for $w' \neq w$ arbitrarily, without changing the truth value of any formula in $\mathcal{L}_1^C$ at w.) A straightforward induction on the structure of formulas shows that for each formula $\psi \in Sub(\varphi)$, we have $(M, w) \models \psi$ if and only if $(M', w) \models \psi$. In particular, because negated conditionals have an existential nature (i.e., $\neg(p \rightarrow q)$ holds if there is a minimal p world satisfying $\neg q$), each negated conditional in Neg is satisfied at (M', w) because it is satisfied in one of the total orders. On the other hand, the conditionals in Pos hold at (M', w) since they hold in each of the total orders. ∎

With minor changes the same construction applies to all the variants we consider.

Corollary 0.3: *Let $\varphi \in \mathcal{L}_1^C$ and let $\mathcal{A}$ be a subset of $\{N, T, U, A, R, C\}$. If φ is satisfiable in $\mathcal{M}^{\mathcal{A}}$, then φ is satisfiable in a structure in $\mathcal{M}^{\mathcal{A}}$ with at most $O(|\mathrm{Sub}_C(\varphi)|^2)$ worlds.*

Proof: Suppose $M \in \mathcal{M}^{\mathcal{A}}$ and $(M, w) \models \varphi$. We now build a structure of the appropriate size satisfying φ. If $\mathcal{A} \subseteq \{U, A\}$, then we can just use the construction of Proposition 0.2, since the structure M' constructed in that proof already satisfies absoluteness (and thus uniformity). If $C \notin \mathcal{A}$, then we can easily modify M' so that it also satisfies whichever of N, T, or R is in $\mathcal{A}$. For example, if $N \in \mathcal{A}$, then M satisfies normality, so W_w is nonempty and we can choose a minimal world in W_w and add it to W'_w as one of the minimal worlds. If $T \in \mathcal{A}$, then we can always choose w as the world to add. If $R \in \mathcal{A}$ but $T \notin \mathcal{A}$, we add w as a maximal world in W'_w.

If $C \in \mathcal{A}$, then we use a different construction. For each formula $\psi \rightarrow \psi' \in Sub_C(\varphi)$,

- if $(M, w) \models \psi \rightarrow \psi'$, then let $w_{\psi \rightarrow \psi'}$ be a minimal ψ-world in W_w if there are ψ-worlds in W_w; otherwise take $w_{\psi \rightarrow \psi'}$ to be w.
- if $(M, w) \models \neg(\psi \rightarrow \psi')$, then let $w_{\psi \rightarrow \psi'}$ be a minimal ψ-world in W_w that satisfies $\neg\psi'$. (There must be such a world since $(M, w) \models \neg(\psi \rightarrow \psi')$.)

Let $W' = \{w\} \cup \{w_{\psi \rightarrow \psi'} : \psi \rightarrow \psi' \in Subc(\varphi)\}$, and let $M' = (W', \pi', R')$, where π' is the restriction of π to W' and R' is the restriction of R to W'. By construction, M' has at most $|Sub_C(\varphi)| + 1$ worlds. We leave it to the reader to check that $(M', w) \models \varphi$. This simple construction depends on the properties of connected preorders. In particular, we need the property that any minimal φ-world is strictly more plausible than *all* the non-minimal φ-worlds. This is not true in the general case. ∎

What happens with formulas that have nested conditionals? It turns out that the answer depends on whether we assume absoluteness and/or uniformity. We first consider the situation where we assume absoluteness. The key observation here is that if we assume absoluteness, since the ordering is the same at all worlds, we can get rid of nested conditionals. For example, in structures satisfying absoluteness, the formula $r \rightarrow (q \rightarrow p))$ is equivalent to $((q \rightarrow p) \wedge (r \rightarrow true)) \vee (\neg(q \rightarrow p) \wedge (p \rightarrow false))$. In general, the denested for-

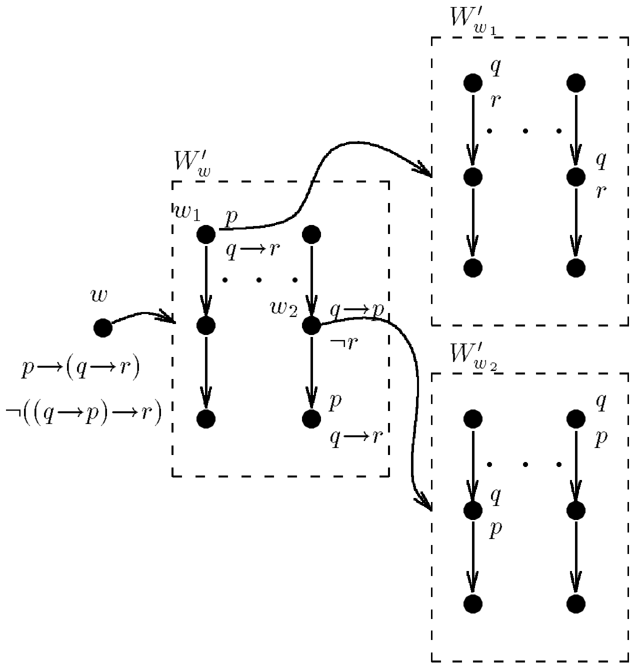

Figure 1: The structure for Proposition 0.6.

mula may be of length exponential in the original formula, but it can be rewritten as a disjunction of formulas each of which does not have too many new conditional subformulas. And since we need to construct a structure satisfying only one of these disjuncts, we conclude that a small structure suffices. More precisely, we have:

Proposition 0.4: *Let $\mathcal{A}$ be a subset of $\{N, T, U, R, C\}$. Given a formula $\varphi \in \mathcal{L}^C$, there are formulas $\varphi_1, \ldots, \varphi_k \in \mathcal{L}^C_1$ such that $\mathcal{M}^{\{A\} \cup \mathcal{A}} \models \varphi \Leftrightarrow \vee_{i=1}^k \varphi_i$. Moreover, for $i = 1, \ldots, k$, we have $|\mathrm{Sub}_C(\varphi_i)| \leq 5|\mathrm{Sub}_C(\varphi)|$.*

From Proposition 0.4 and Corollary 0.3, we immediately get:

Corollary 0.5: *Let $\mathcal{A}$ be a subset of $\{N, T, U, R, C\}$. If $\varphi \in \mathcal{L}^C$ is satisfiable in $\mathcal{M}^{\{A\} \cup \mathcal{A}}$, then φ is satisfiable in a structure in $\mathcal{M}^{\{A\} \cup \mathcal{A}}$ with at most $O(|Subc(\varphi)|^2)$ worlds.*

In structures that do not satisfy absoluteness, we can still extend the ideas of Proposition 0.2 recursively to get polynomial-sized structures for formulas of bounded-depth nesting, where the polynomial depends on the depth of nesting *provided we do not also assume uniformity.*

Proposition 0.6: *Let $\varphi \in \mathcal{L}^C_k$ and let $\mathcal{A}$ be a subset of $\{N, T, R, C\}$. If φ is satisfiable in $\mathcal{M}^{\mathcal{A}}$, then φ is satisfiable in a structure in $\mathcal{M}^{\mathcal{A}}$ with at most $O(|\mathrm{Sub}(\varphi)|^{2k})$ worlds.*

Proof: We apply the construction of Proposition 0.2 recursively. Roughly speaking, at the top level of the recursion, we treat all nested conditionals as new primitive propositions. Applying the construction of Proposition 0.2 we get the set W'_w. For each $w' \in W'_w$, let $\varphi_{w'}$ be the conjunction of all the propositions (including the nested conditionals) that hold at w'. We note that $Sub_C(\varphi_{w'}) \subseteq Sub_C(\varphi)$. We now apply the procedure recursively to w' and $\varphi_{w'}$ to construct $W'_{w'}$. We proceed in this manner, constructing a tree-like structure (as shown in Figure 1), dealing with conditionals nested i deep at the ith level of the recursion. Thus, we can stop at the kth level. Note that for w', w'' in the structure, $W'_{w'}$ is disjoint from $W'_{w''}$ if $w' \neq w''$. Thus, this structure does not satisfy uniformity.

We now give a formal description of the construction. We define $Basic_i(\varphi) \subseteq \Phi \cup Sub_C(\varphi)$ as the set of primitive propositions and conditional statements that are subformulas of φ and appear inside exactly i levels of conditional nesting. For example, if φ is $(p \rightarrow (q \rightarrow r)) \wedge \neg((r \rightarrow q) \rightarrow r)$, then $Basic_0(\varphi) = \{p \rightarrow (q \rightarrow r), (r \rightarrow q) \rightarrow r\}$, $Basic_1(\varphi) = \{p, q \rightarrow r, r \rightarrow q, r\}$ and $Basic_2(\varphi) = \{q, r\}$. We treat formulas in $Basic_{i+1}(\varphi)$ as primitive propositions during the construction of the orderings at level i.

We construct the tree-like structure in the following fashion. The procedure gets as input a structure M, a world w, and a formula φ such that $(M, w) \models \varphi$, and returns a structure M' such that $(M', w) \models \varphi$. Moreover, M' contains at most $O(|Sub_C(\varphi)|^{2k})$ worlds, where k is the depth of nesting in φ. If φ is propositional then the structure M' consists of the single worlds w. If φ contains conditional formulas, then we construct a tree with w as the root. The truth value of any primitive proposition $p \in Basic_0(\varphi)$ is determined at w by π. Thus, we only need to satisfy conditional formulas in $Basic_0(\varphi)$. We apply the procedure described in the proof of Proposition 0.2, treating every formula in $Basic_1(\varphi)$ as a primitive proposition. We get a structure M^w of size $O(|Sub_C(\varphi)|^2)$, such that π^w maps each world $w' \in W^w_w$ to a truth assignment over $Basic_1(\varphi)$. Recall that the construction of M^w is such that each $w' \in W^w_w$ corresponds to a world $f(w')$ in W_w. Moreover, for all $\psi \in Basic_1(\varphi)$, $pi^w(w')(\psi) =$ **true** if and only $(M, f(w')) \models \psi$.

For each $w' \in W'_w$ we define $\varphi_{w'}$ so that it describes the truth value of all formulas in $Basic_1(\varphi)$ at w'. However, since we want to capture conditionals holding in w', we have to be careful; we use the corresponding world $f(w)$ in M to evaluate these conditionals. Formally, $\varphi_{w'}$ is defined as $\bigwedge_{\psi \in Basic_1(\varphi),\ (M, f(w')) \models \psi} \psi \wedge \bigwedge_{\psi \in Basic_1(\varphi),\ (M, f(w')) \models \neg\psi} \neg\psi$. We note that $Sub_C(\varphi_{w'}) \subseteq Sub_C(\varphi)$, $Basic_i(\varphi_{w'}) \subseteq Basic_{i+1}(\varphi)$, and that $\varphi_{w'}$ contains at most $k-1$ levels of conditional nesting. We now recursively apply the tree construction procedure on $(M, f(w'))$ and $\varphi_{w'}$ and get a structure $M^{w'}$ such that $(M^{w'}, w') \models \varphi_{w'}$.

We now construct M'. Let W' contain w and all the worlds in $M^{w'}$ for all $w' \in W_w^w$. Without loss of generality we can assume that the sets of worlds in $M^{w'}$ are disjoint and do not contain w. We define π' to be such that for each world $w' \in W'$, the truth assignment $\pi'(w')$ is the same as the truth assignment in the structure w' is taken from. Finally, we define $\preceq_w$ according to the construction of M^w, and $\preceq_{w'}$ for all $w' \in W' - \{w\}$ to be the same as the ordering $\preceq_{w'}$ in the structure w' is taken from.

It is easy to see that this recursive procedure is well-defined. At i level of recursion the depth of the formula is at most $k - i$, and thus the procedure must terminate. It is also easy to verify, by induction, that $(M', w') \models \varphi$. Finally, we show that the structure M' is not too large. The size of M' is $O(|Sub_C(\varphi)|^2 \cdot |M^{w'}|)$. According to the recursive construction, $|M^{w'}| = O(|Sub_C(\varphi)|^{2(k-1)})$. Thus, the size of M' is $O(|Sub_C(\varphi)|^{2k})$.

The procedure we described constructs structures in $\mathcal{M}$. If $\mathcal{A}$ is not empty we have to modify M' to satisfy the constraints in $\mathcal{A}$. This is done locally at each world in the manner described in the proof of Corollary 0.3. ∎

What happens if we have no bound on the nesting depth? In this case we can get an exponential-sized structure. The result without uniformity follows immediately from Proposition 0.6, since the depth of nesting in a formula φ is clearly bounded by $|\varphi|$. With uniformity, we have to work a little harder; we leave details to the full paper.

Proposition 0.7: *Suppose $\varphi \in \mathcal{L}_k^C$ and $\mathcal{A}$ is a subset of $\{N, T, R, C, U\}$. If φ is satisfiable in $\mathcal{M}^{\mathcal{A}}$, then φ is satisfiable in a structure in $\mathcal{M}^{\mathcal{A}}$ with at most $O(2^{2|Sub_C(\varphi)|})$ worlds.*

The natural question to ask is whether this the best we can guarantee. The answer is yes. Since the technique for proving this, which depends on the observation that we can embed various modal logics into conditional logic, is also useful for proving lower bounds on complexity, we go into a little detail here.

Let $\mathcal{L}^K$ be the language with a single modal operator K (which intuitively stands for knowledge). As usual, we capture the semantics of knowledge in terms of an accessibility relation, on which we can place various restrictions. Thus, an *epistemic structure* N has the form $(W, \pi, \mathcal{K})$, where W is a set of worlds, π maps each possible world to a truth assignment, and $\mathcal{K}$ is a binary relation. We define $\models$ in the standard way; in particular,

- $(N, w) \models K\varphi$ if $(N, w') \models \varphi$ for all w' such that $(w, w') \in \mathcal{K}$.

Let $\mathcal{N}$ be the class of all epistemic structures. We add superscripts r, s, t, and e, respectively, to denote restrictions on the $\mathcal{K}$ relation to reflexive, serial, transitive, and Euclidean relations, respectively.[8] For each subset $\mathcal{B}$ of $\{r, s, t, e\}$, we let $\mathcal{N}^{\mathcal{B}}$ denote the class of epistemic structures where the $\mathcal{K}$ relation satisfies the appropriate restrictions.

We can also define modal operators in the context of conditional logic. Let $\boxdot\varphi$ be an abbreviation for $true \rightarrow \varphi$, and let $\Box\varphi$ be an abbreviation for $(\neg\varphi) \rightarrow false$. It is easy to verify that $\boxdot\varphi$ holds at w exactly if all the minimal worlds according to $\preceq_w$ satisfy φ and that $\Box\varphi$ holds at w if all worlds in W_w satisfy φ. Traditionally (Lew73), $\boxdot$ has been called the *inner* modality and $\Box$ has been called the *outer* modality.

As we now show, the inner modality $\boxdot$ corresponds in a precise sense to K. Under this correspondence, conditions on $\preceq_w$ correspond to conditions on the binary relation $\mathcal{K}$. In particular, conditions N and R both correspond to $\mathcal{K}$ being serial, T corresponds to $\mathcal{K}$ being reflexive, and A corresponds to $\mathcal{K}$ being both transitive and Euclidean. This intuition is made precise by the theorem below.

Proposition 0.8: *Given a formula $\varphi \in \mathcal{L}^K$, let φ^* be the result of replacing each K operator by $\boxdot$. Let $\mathcal{A}$ be a (possibly empty) subset of $\{N, R, T, A\}$, and let $\mathcal{B}$ be the corresponding subset of $\{e, r, s, t\}$, where s corresponds to N and R, r corresponds to T, and both e and t correspond to A. Finally, let $\mathcal{A}'$ be a subset of $\{C\}$. Then φ is satisfiable in $\mathcal{N}^{\mathcal{B}}$ in a structure of size k if and only if φ^* is satisfiable in $\mathcal{M}^{\mathcal{A} \cup \mathcal{A}'}$ in a structure of size k.*

Proof: We show how to map epistemic structures satisfying φ to structures satisfying φ^* and vice versa.

Recall that $\boxdot\psi$ holds at w exactly when ψ is true in all the minimal worlds in W_w. Assume $(M, w) \models \varphi^*$ for $M = (W, \pi, R)$. Let $N = (W, \pi, \mathcal{K})$ such that $(w_1, w_2) \in \mathcal{K}$ if w_2 is minimal in W_{w_1}. It is easy to check that $(N, w) \models \varphi$. Moreover, if M satisfies normality or reflexivity, then for every w_1 there is at least one minimal w_2, and thus $\mathcal{K}$ is serial. If M satisfies centering, then w_1 is minimal in W_{w_1}, and thus $\mathcal{K}$ is reflexive. If M satisfies absoluteness, then if $w_2 \in W_{w_1}$, then $\preceq_{w_2}$ is the same as $\preceq_{w_1}$. This implies that if $(w_1, w_2) \in \mathcal{K}$ and $(w_2, w_3) \in \mathcal{K}$ then $(w_1, w_3) \in \mathcal{K}$ since w_3 must be minimal in W_{w_1}. Similarly, if $(w_1, w_2), (w_1, w_3) \in \mathcal{K}$ then (w_2, w_3) is also in $\mathcal{K}$. Thus, $\mathcal{K}$ is transitive and Euclidean.

Now assume $(N, w) \models \varphi$ for $N = (W, \pi, \mathcal{K})$. We construct a structure $M = (W, \pi, R)$, where R is such that for each world w, the set W_w consists of all worlds accessible from w according to $\mathcal{K}$, and each of these worlds is equally plausible. This ensures that the minimal worlds according to $\preceq_w$ are precisely the worlds accessible from w, and guarantees that $(N, w) \models K\varphi$ if and only if $(M, w) \models \boxdot\varphi$. Moreover, if $\mathcal{K}$ is serial, then $W_w \neq \emptyset$ for all w, and thus M satisfies reflexivity (and

[8] $\mathcal{K}$ is *serial* if for all w, there exists some w' such that $(w, w') \in \mathcal{K}$; $\mathcal{K}$ is Euclidean if for all u, v, w, if $(u, v) \in \mathcal{K}$ and $(u, w) \in \mathcal{K}$, then $(v, w) \in \mathcal{K}$.

normality). If $\mathcal{K}$ is reflexive, then w is accessible from w. Thus, $w \in W_w$ in M, and hence minimal (since all worlds in W_w are minimal). Finally, if $\mathcal{K}$ is both transitive and Euclidean, then then it is well known (see (HM92)) that we can assume without loss of generality that the same set of worlds is accessible from each $w \in W$. This implies that the ordering at each world is the same. Thus, M satisfies absoluteness. ∎

In the presence of uniformity we can get similar results. However the reduction is less natural. Since such a reduction does not play a role in our treatment of structures satisfying uniformity, we omit the details here.

Halpern and Moses (HM92) describe formulas in $\mathcal{L}^K$ that can be satisfied only in exponential-sized structures in $\mathcal{N}$, $\mathcal{N}^r$, and $\mathcal{N}^e$. (However, they can be satisfied in polynomial-sized structures in $\mathcal{N}^{et}$.) They also show that once *common knowledge* is added to the language, then there are formulas that have depth of nesting two and can be satisfied only in exponential-sized structures in $\mathcal{N}$, $\mathcal{N}^r$, and $\mathcal{N}^e$. It turns out that the outer modality behaves very much like common knowledge in the presence of uniformity. More precisely, the statement $C\varphi$ ("it is common knowledge that φ") holds exactly when every world that is accessible through repeated applications of K satisfies φ. Similarly, $\Box\varphi$ holds at w when all worlds in W_w satisfy φ. If we assume uniformity, then $\Box\varphi$ implies that all worlds that are accessible by arbitrary level of conditional nesting must satisfy ψ. This is close enough to common knowledge to get the behavior needed for constructing a proof similar to their construction for common knowledge.

When we do not require uniformity, we immediately get the following from the results of Halpern and Moses and Proposition 0.8:

Corollary 0.9: *Let $\mathcal{A}$ be a subset of $\{N, R, T, C\}$. Then for each n, there is a formula $\varphi_n^{\mathcal{A}}$ of size $O(n^2)$ such that $\varphi_n^{\mathcal{A}}$ that is satisfiable in $\mathcal{M}^{\mathcal{A}}$, but only in structures of size at least 2^n.*

When we require uniformity we have to work a bit harder. We can modify the construction Halpern and Moses use for common knowledge to get the following result; we leave the details for the full paper.

Proposition 0.10: *Let $\mathcal{A}$ be a subset of $\{N, R, T, C\}$. Then for each n, there is a formula $\varphi_n^{\mathcal{A}}$ of size $O(n^2)$ and using only depth-two nesting of conditionals such that $\varphi_n^{\mathcal{A}}$ that is satisfiable in $\mathcal{M}^{\mathcal{A}\cup\{U\}}$, but only in structures of size at least 2^n.*

COMPLEXITY RESULTS

In this section we examine the inherent difficulty of deciding whether a formula is satisfiable. Checking validity is closely related since φ is valid if and only if $\neg\varphi$ is not satisfiable. We start with an overview of the complexity-theoretic notions we need. For a more detailed treatment of the topic, see (GJ79; HU79).

Table 1: The complexity of the satisfiability problem for $\mathcal{M}^{\mathcal{A}}$.

	$A \in \mathcal{A}$	$A, U \notin \mathcal{A}$	$U \in \mathcal{A}$, $A \notin \mathcal{A}$
$\mathcal{L}_1^C$	NP	NP	NP
$\mathcal{L}_k^C$	NP	NP	EXPTIME
$\mathcal{L}^C$	NP	PSPACE	EXPTIME

Complexity theory examines the difficulty of determining membership in a set as a function of the input size. In our case we check if a formula φ is in the set of satisfiable formulas. Difficulty is measured in terms of the time or space required to decide if a formula φ is satisfiable as a function of $|\varphi|$, the length of the formula. The complexity classes we are interested in are NP, PSPACE, and EXPTIME. These classes contain sets such that deciding membership can be done in nondeterministic polynomial time, polynomial space, and exponential time, respectively.

To show that a set is in a complexity class we usually describe a procedure that determine membership in the set and conforms to the time or space restriction of the class. Usually, we also want to show that a set is not in an easier class. To do we show that the set is *hard* in the class. A set A is hard in a class $\mathcal{C}$ if for every set $B \in \mathcal{C}$, an algorithm deciding membership in B can be easily obtained from an algorithm deciding membership in A. A set is *complete* with respect to a complexity class $\mathcal{C}$ if it is both in $\mathcal{C}$ and $\mathcal{C}$-hard.

We now turn to the complexity results. These results are summarized in Table 1 (where each problem is complete for the complexity class listed). For most classes of structures of interest to us, deciding satisfiability is NP-complete for $\mathcal{L}_k^C$ and PSPACE-complete $\mathcal{L}^C$. However, there are several exceptions to this rule: absoluteness makes the problem easier and uniformity makes it harder. Notice that all the other semantic variants do not affect the complexity.

All the logical variants we examine contain the propositional calculus and thus checking satisfiability is NP-hard. For the variants with polynomial-sized structures we see that deciding satisfiability is in NP: We simply nondeterministically choose a structure and then verify that it satisfies the formula. The verification step is easily shown to be in polynomial time, provided the structure is polynomial-sized. Using Proposition 0.2, Corollaries 0.3 and 0.5, and Proposition 0.6 we get the following theorem:

Theorem 0.11: *Let $\mathcal{A}$ be a subset of $\{N, R, T, U, A, C\}$. Then the following problems are NP-complete:*

(a) the problem of deciding whether a formula in $\mathcal{L}_1^C$ is satisfiable in $\mathcal{M}^{\mathcal{A}}$,

(b) *the problem of deciding whether a formula in $\mathcal{L}^C$ is satisfiable in $\mathcal{M}^{\mathcal{A}}$, if $\mathcal{A}$ contains A,*

(c) *for a fixed $k > 0$, the problem of deciding whether a formula in $\mathcal{L}_k^C$ is satisfiable in $\mathcal{M}^{\mathcal{A}}$, if $\mathcal{A}$ does not contain U.*

We now turn to the harder cases. As we showed in Corollary 0.9 and Proposition 0.10, in all the remaining variants there are formulas that are satisfiable only in exponential-sized structures. We show that most of these variants, except the ones satisfying uniformity, are PSPACE-complete.

Theorem 0.12: *If $\mathcal{A}$ is a subset of $\{N, T, R, C\}$, the problem of deciding if a formula in $\mathcal{L}^C$ is satisfiable in $\mathcal{M}^{\mathcal{A}}$ is PSPACE-complete.*

Proof: The lower bound is an immediate corollary of Proposition 0.8 and the fact (proved by Ladner (Lad77; HM92)) that checking whether a formula in $\mathcal{L}^K$ is satisfiable in $\mathcal{N}$ (resp., $\mathcal{N}^r$, $\mathcal{N}^e$) is PSPACE-hard.

For the upper bound we use the construction in Proposition 0.7. We describe a polynomial space algorithm that essentially searches through all the tree-like structures of the form described in the proof of Proposition 0.7. In order to simplify the description of this algorithm we rely on the fact that NPSPACE (nondeterministic polynomial space) is equivalent to PSPACE (HU79). Thus, we describe an algorithm that uses nondeterministic choices and polynomial space.

The algorithm **check-tree** is given a world w and formula φ and returns *true* if there is a tree-like structure containing w such that $w \models \varphi$.

check-tree(w, φ)
  Guess a truth assignment at w to propositions in Φ
  If $\varphi \in \mathcal{L}_0^C$, then
    Let $W_w = \emptyset$
  Else,
    Let $n = |Sub_C(\varphi)|$
    Let $W_w = \{w_{1,1}, \ldots, w_{1,n}, \ldots, w_{n,n}\}$
    Let $w_{i,j} \preceq_w w_{i,k}$ exactly if $j \leq k$
    For each $w_{i,j}$,
      Guess $T_{i,j} \subseteq Basic_1(\varphi)$
      Let $\varphi_{w_{i,j}} = \bigwedge_{\psi \in T_{i,j}} \psi \wedge \bigwedge_{\psi \in Basic_1(\varphi) - T_{i,j}} \neg\psi$
      If **check-tree**$(w_{i,j}, \varphi_{w_{i,j}}) = $ *false*, then
        return *false*.
  Return the evaluation of φ at w (using the ordering $\preceq_w$
    and assuming $\varphi_{w_{i,j}}$ is true at $w_{i,j}$).
end.

This algorithm emulates the construction that we used in the proof of Proposition 0.6. It guesses a structure and then checks that φ evaluates to *true* in this structure. It starts by guessing a truth assignment at w. If φ contains conditionals, then the algorithm guesses a structure that contains $|Sub_C(\varphi)|^2$ worlds and defines an ordering $\preceq_w$ over these worlds which is a disjoint union of $|Sub_C(\varphi)|$ total orders. It then guesses a truth assignment in each of these $|Sub_C(\varphi)|^2$ worlds to the formulas in $Basic_1(\varphi)$. According to the proof of Proposition 0.2, if φ is satisfiable (when we consider formulas in $Basic_1(\varphi)$ as propositions), it must be satisfiable in such as structure. The algorithm then verifies that the formulas assigned to each $w_{i,j}$ can be satisfied using a recursive call. Finally, the algorithm verifies that φ evaluates to true at w according to the truth assignment at w and W_w (using $T_{i,j}$ to evaluate formulas at each $w_{i,j} \in W_w$).

We note that the space requirements of the algorithm are the space requirements of all the instances that are active at once. The maximal number of active instances is exactly the recursion depth, i.e., the conditional nesting depth in φ. The space requirements in each instance are $O(|Sub(\varphi)|^3)$ for storing the sets $T_{i,j}$. Thus, the space requirements for **check-tree**(φ, w) are $O(|\varphi|^4)$. ∎

The remaining cases are those satisfying uniformity but not absoluteness. Somewhat surprisingly, these variants are harder than all the others. Roughly speaking, this is because in the presence of uniformity, the outer modality essentially allows us to express common knowledge.

Theorem 0.13: *If $\mathcal{A}$ is a subset of $\{N, T, R, C\}$, the problem of deciding if a formula in $\mathcal{L}^C$ is satisfiable in $\mathcal{M}^{\mathcal{A} \cup \{U\}}$ is EXPTIME-complete.*

Proof: The lower bound is constructed in a similar manner to the lower bound for logics of knowledge and common knowledge of Halpern and Moses (HM92). The basic idea is that we can simulate the execution of an alternating polynomial-space Turing machine by a sentence φ in $\mathcal{L}^C$, such that φ is satisfiable if only if the machine accepts the input, and φ is of polynomial size.[9] We leave the details of this construction to the full paper.

We prove the upper bound by modifying the algorithm **check-tree** we described in the proof of Theorem 0.12. The basic idea is straightforward: We try to modify the tree-like structure M constructed by **check-tree** to a structure M' over the same set of worlds that satisfies uniformity. The idea is to modify the preference relation so that at each world w, the set W'_w of worlds considered possible consists of all worlds in the tree except the root, and defining $\preceq'_w$ so that the minimal worlds in W'_w are exactly those in W_w. This modification guarantees that if $(M, w) \models \neg(\psi \rightarrow \psi')$ then $(M', w) \models \neg(\psi \rightarrow \psi')$. Since there is a minimal ψ-world in W_w that satisfies $\neg\psi'$, the same world is also a minimal ψ-world in W'_w. Moreover, if $(M, w) \models \psi \rightarrow \psi'$ and there are some ψ-worlds in W_w, then $(M', w) \models \varphi \rightarrow \psi'$ for the same reasons. Unfortunately, this approach runs into problems if there are no ψ-worlds in W_w, so that $\psi \rightarrow \psi'$ holds vacuously at world w in structure M. In that

[9] The class of sets recognizable by alternating polynomial-space Turing machines is equal to EXPTIME (CKS81).

case, if there are some ψ-worlds in W'_w (which is possible), then the conditional $\psi \rightarrow \psi'$ may not be true at (M', w).

To avoid this problem we can decide in advance which conditionals in $Sub_C(\varphi)$ will be satisfied vacuously in M. We initially nondeterministically choose a subset V of $Sub_C(\varphi)$. We then modify **check-tree** so that it searches structures where the only conditionals that hold vacuously are those in V. The modified **check-tree** ensures that no world satisfies ψ for each formula $\psi \rightarrow \psi' \in V$. One side-effect of this change is that we may get new conditionals at each level of recursion, so the algorithm may not terminate. We avoid this by using the fact that there are only an exponential number of formulas of the form $\varphi_{w_{i,j}}$ that can be given as an argument to **check-tree**. We leave details to the full paper. Note that the modified **check-tree** is no longer guaranteed to be in PSPACE. In the full paper we show that it is guaranteed to be in EXPTIME. ∎

RIGHT-NESTED FORMULAS

As mentioned in the introduction, a similar approach to conditional logic is the framework of counterfactual queries of (FUV83; Gin86). Eiter and Gottlob (EG93) show that the complexity of evaluating a query of the form $p_1 > (p_2 > \ldots (p_n > q) \ldots)$ is Π_2^p-complete, and the complexity of queries that allow negation on the right-hand side is PSPACE-complete. Since right-nested conditionals also appear in the conditional logic literature (Bou93; FH94b), it seems worth understanding if right-nesting simplifies things here too.

We now define the language $\mathcal{L}_s^C$ of simple right-nested conditionals and the language $\mathcal{L}_r^C$ of (possibly negated) right-nested conditionals. Let $\mathcal{L}_s^C$ be the least language such that if $\varphi, \varphi' \in \mathcal{L}_s^C$ and $\psi, \psi_1, \ldots, \psi_n \in \mathcal{L}_0^C$ $(n \geq 0)$, then $\varphi \wedge \varphi'$, $\neg\varphi$, and $\psi_1 \rightarrow \cdots \rightarrow \psi_n \rightarrow \psi$ are in $\mathcal{L}_s^C$. Let $\mathcal{L}_r^C$ be the minimal language such that if $\varphi, \varphi' \in \mathcal{L}_r^C$ and $\psi \in \mathcal{L}_0^C$ then $\psi, \varphi \wedge \varphi', \neg\varphi, \psi \rightarrow \varphi$ are in $\mathcal{L}_r^C$. Thus $p \rightarrow q \rightarrow r$ is in both languages, and $p \rightarrow \neg(q \rightarrow r)$ is in $\mathcal{L}_r^C$ but not in $\mathcal{L}_s^C$.

Things are considerably simpler for $\mathcal{L}_s^C$. It is easy to show that the satisfiability problem for $\mathcal{L}_s^C$ is NP-complete for all variants of the logic:

Theorem 0.14: *Let $\mathcal{A}$ be a subset of $\{N, R, T, U, A, C\}$. Then the problem of deciding whether a formula in $\mathcal{L}_s^C$ is satisfiable is NP-complete.*

Proof: Using techniques similar to these of Proposition 0.2, it is easy to show that a formula in $\mathcal{L}_s^C$ is satisfiable if and only if it is satisfiable in a linear-size structure. Thus we get the NP upper bound. The NP-hardness is a result of the fact that $\mathcal{L}_s^C$ contains the propositional calculus. ∎

Things get more complicated when we consider the language $\mathcal{L}_r^C$. In many cases this fragment is already as complex as the full language. Recall that the PSPACE lower bound in Theorem 0.12 is proved by a reduction from modal logic. This reduction substitutes the $\boxdot$ modality for the modal operator K. However $\boxdot\varphi$ is defined as $true \rightarrow \varphi$. Thus, the reduction maps a modal formula into a formula in $\mathcal{L}_r^C$. (Because modal formulas may be negated, the resulting formula may not be in $\mathcal{L}_s^C$.) Thus, we get the following corollary:

Corollary 0.15: *If $\mathcal{A}$ is a subset of $\{N, T, R, C\}$, the problem of deciding if a formula in $\mathcal{L}_r^C$ is satisfiable in $\mathcal{M}^{\mathcal{A}}$ is PSPACE-complete.*

However, when we consider structures that satisfy uniformity, the satisfiability problem for formulas in $\mathcal{L}_r^C$ is easier than satisfiability of formulas in the full language.

Theorem 0.16: *If $\mathcal{A}$ is a subset of $\{N, T, R, C\}$, the problem of deciding if a formula in $\mathcal{L}_r^C$ is satisfiable in $\mathcal{M}^{\mathcal{A} \cup \{U\}}$ is PSPACE-complete.*

AXIOMATIZATION

Several axiom systems for variants of conditional logics appear in the literature (Lew73; Lew74; Che80; Bur81; Bel89; Gra91; KS91). We present an axiom system for all the variants we introduced based on Burgess's (Bur81) axiomatization. In the full paper, we provide a full completeness proof based on Burgess's techniques, but substantially simpler.

The basic axiom system, AX, contains the following axiom schemata:

A0 All the propositional tautologies

A1 $\varphi \rightarrow \varphi$

A2 $((\varphi \rightarrow \psi_1) \wedge (\varphi \rightarrow \psi_2)) \Rightarrow (\varphi \rightarrow (\psi_1 \wedge \psi_2))$

A3 $(\varphi \rightarrow (\psi_1 \wedge \psi_2)) \Rightarrow (\varphi \rightarrow \psi_1)$

A4 $((\varphi_1 \rightarrow \varphi_2) \wedge (\varphi_1 \rightarrow \psi)) \Rightarrow ((\varphi_1 \wedge \varphi_2) \rightarrow \psi)$

A5 $((\varphi_1 \rightarrow \psi) \wedge (\varphi_2 \rightarrow \psi)) \Rightarrow ((\varphi_1 \vee \varphi_2) \rightarrow \psi)$

and the following inference rules:

MP From φ and $\varphi \Rightarrow \psi$ infer ψ.

RPE From $\varphi_1 \Leftrightarrow \varphi_2$ and ψ infer ψ', where ψ' differs from ψ only by replacing some subformulas of φ of the form φ_1 by φ_2.

The completeness proof works as follows. Given φ and an extension AX$'$ of AX, we consider all the *maximal consistent* subsets, according to AX$'$, of $Sub^+(\varphi)$ (where a maximal consistent set is an AX$'$-consistent set which is not a strict subset of any other AX$'$-consistent subset of $Sub^+(\varphi)$). We call a such a maximal consistent set an *AX$'$-atom.* (We henceforth omit AX$'$ unless it is relevant to the discussion.) It is easy to verify that each atom is *complete* in the sense that for $\psi \in Sub(\varphi)$, either ψ or $\neg\psi$ must be in the atom. For example, if φ is $p \wedge (q \rightarrow r)$ then $Sub^+(\varphi) = \{p, \neg p, q, \neg q, r, \neg r, q \rightarrow r, \neg(q \rightarrow r), p \wedge (q \rightarrow r), \neg(p \wedge (q \rightarrow r))\}$. The set $\{p, q, \neg r, \neg(q \rightarrow r), \neg(p \wedge (q \rightarrow r))\}$ might be an atom (depending on AX$'$), but $\{\neg p, q, \neg r, q \rightarrow r, p \wedge (q \rightarrow r)\}$ cannot be an atom since $\neg p$ and $p \wedge (q \rightarrow r)$ are inconsistent. Similarly, $\{p, q\}$ is not atom since it is not maximal. In the following discussion let α, β, and γ stand for atoms and A stand for a set of atoms. We slightly abuse notation and use α both as a set (e.g., $\psi \in \alpha$) and as a formula (e.g., $\alpha \Rightarrow \psi$) which is the conjunction of all members of α.

Given α, β, and A we define $\mathit{Prefer}_{\mathrm{AX}'}(\beta, \alpha, A)$ if $\beta \wedge \neg(\alpha \vee \bigvee A \rightarrow \bigvee A)$ is consistent according to AX$'$. The intuitive account is that a world where β holds is consistent with an ordering that makes worlds where α holds strictly preferred to worlds satisfying one of the atoms in A. We will use this definition to construct all the preorders that are consistent with each possible world.

Given AX$'$ and φ, we construct a structure $M = (W, \pi, R)$ as follows:

- We set W to be the set of tuples (γ, A) where $\gamma \notin A$. Given $w = (\gamma, A)$ we define $\gamma(w) = \gamma$.
- We set $\pi(w)(p) = \mathbf{true}$ if and only if $p \in \gamma(w)$.
- For any world w, we construct $\preceq_w$ by setting

$$W_w = \{(\gamma, A) \in W \mid \mathit{Prefer}_{\mathrm{AX}'}(\gamma(w), A, \gamma)\}$$

  and setting $w' \prec_w w''$ if $w' = (\gamma', A')$, $w'' = (\gamma'', A'')$, and $A'' \cup \{\gamma''\} \subseteq A'$.

The intuition is simple: A world $w = (\gamma, A)$ represents a world satisfying γ that is intended to be strictly preferred to all worlds that satisfy one of the atoms in A. We define π so that it assigns truth values to primitive propositions according their values in γ. The set W_w contains all the worlds (γ', A') such that $\mathit{Prefer}_{\mathrm{AX}'}(\gamma, \gamma', A')$, i.e., it is consistent with γ that (γ', A') is strictly preferred to worlds satisfying one of the atoms in A. The definition of $\preceq$ implements this intuition: if $(\gamma', A') \prec_w (\gamma'', A'')$ then $\gamma'' \in A'$. This matches our intuition since (γ', A') is intended to be preferred to worlds satisfying atoms in A'. We also demand that $A'' \subset A'$, this ensures that $\preceq_w$ will be transitive. It implies that if $(\gamma'', A'') \prec_w (\gamma''', A''')$, then $\gamma''' \in A'' \subset A'$ and also $A''' \subset A'' \subset A'$. Thus, $(\gamma', A') \prec_w (\gamma''', A''')$.

We now show that each world w in M satisfies $\gamma(w)$. Since the details of this proof are essentially the same as Burgess's proof (Bur81, p. 82), we leave the details to the full paper.

Lemma 0.17: *Let AX$'$ be an extension of AX and let $\varphi \in \mathcal{L}^C$. Let M be the structure constructed above. For any $w \in W$ and $\psi \in \mathrm{Sub}^+(\varphi)$, $\psi \in \gamma(w)$ if and only if $(M, w) \models \psi$.*

Using this lemma it is easy to prove the following theorem:

Theorem 0.18: *If $\varphi \in \mathcal{L}^C$, then φ is valid in $\mathcal{M}$ if and only if $\vdash_{AX} \varphi$.*

Proof: It is easy to check the soundness of AX in $\mathcal{M}$. Thus, if $\vdash_{\mathrm{AX}} \varphi$, then φ must be valid in $\mathcal{M}$. For the other direction, assume that φ is consistent with AX. Then there is an atom α such that $\varphi \in \alpha$, and from Lemma 0.17 we get that $(M, (\alpha, \emptyset)) \models \varphi$. ∎

We note that this construction is much simpler than Burgess's even though the proof of Lemma 0.17 is almost identical to Burgess's proof. The main difference is that Burgess constructs a tree-like structure of finite but nonelementary size. Our construction, on the other hand, uses the same stock of worlds to construct the ordering for each world. The resulting structure is of doubly-exponential size. (We note that our results from Section 2 show that this can be improved, since only an exponential-sized structure is needed for satisfiability.) The fact that the structure is not tree-like allows us to give completeness proofs for properties such as uniformity and reflexivity that cannot be satisfied in tree-like structures.

The following axioms characterize the various semantic conditions we have considered. These axioms appeared originally in (Lew73) and (Bur81).

AN (Normality) $\neg(\mathit{true} \rightarrow \mathit{false})$

AR (Reflexivity) $\Box\varphi \Rightarrow \varphi$

AT (Centering) $\boxdot\varphi \Rightarrow \varphi$[10]

AU (Uniformity) $(\Box\varphi \Rightarrow \Box\Box\varphi) \wedge (\neg\Box\varphi \Rightarrow \Box\neg\Box\varphi)$

AA (Absoluteness) $(\varphi \rightarrow \psi \Rightarrow \Box(\varphi \rightarrow \psi)) \wedge (\neg(\varphi \rightarrow \psi) \Rightarrow \Box\neg(\varphi \rightarrow \psi))$

AC (Connectedness)$(\varphi_1 \vee \varphi_2) \rightarrow \neg\varphi_2 \Rightarrow ((\varphi_1 \vee \psi) \rightarrow \neg\psi) \vee ((\psi \vee \varphi_2) \rightarrow \neg\varphi_2)$

The next results shows that each axiom captures exactly the corresponding condition:

Theorem 0.19: *Let $\varphi \in \mathcal{L}$ and let $\mathcal{A}$ be a subset of $\{N, R, T, U, A, C\}$ and A the corresponding subset of $\{AN, AR, AT, AU, AA, AC\}$. Then φ is valid in $\mathcal{M}^{\mathcal{A}}$ if and only if $\vdash_{\mathrm{AX} \cup A} \varphi$.*

Proof: In the full paper we provide the details of this proof. The essence of the proof is showing that each axiom forces the constructed structure to satisfy the semantic condition. This is straightforward in the case of absoluteness, uniformity and normality. The other cases require a little more care; we leave details to the full paper. ∎

CONCLUSIONS

In this paper we analyzed the complexity problem for conditional logics. As we observed in the introduction, such logics are now being used in many areas of artificial intelligence. The techniques we have introduced in this paper (especially the results in Section) can applied to frameworks that combine conditional logics with other modalities. For example, in (FH94a) we use these results to derive complexity results for a logic that contains both conditionals and epistemic modalities.

[10] The axiom for strict centering is $\boxdot\varphi \Leftrightarrow \varphi$.

We did not attempt, in this work, to isolate tractable fragments of the logic. This is certainly an important aspect of any analysis of formal method in artificial intelligence (Lev86; Lev88). We note that all the logics we examined are intractable because they contain the propositional calculus. It is certainly feasible that there are nontrivial fragments that do not contain the propositional calculus that are tractable (e.g., results in the style of Kautz and Selman's analysis of default logic (KS89)). We plan to pursue this issue in the future. We note that the methods used in this paper are certainly relevant to such an investigation.

Acknowledgements The authors are grateful to Phil Hubbard, Daphne Koller, Moshe Vardi and the anonymous referees for comments on drafts of this paper.

References

J. A. Allen, R. Fikes, and E. Sandewall, editors. *Principles of Knowledge Representation and Reasoning: Proc. Second International Conference (KR '91).* Morgan Kaufmann, San Francisco, Calif., 1991.

J. Bell. The logic of nonmonotonicity. *Artificial Intelligence*, 41:365–374, 1989.

C. Boutilier. Normative, subjective and autoepistemic defaults: adopting the Ramsey test. In *Principles of Knowledge Representation and Reasoning: Proc. Third International Conference (KR '92)*, pages 685–696. Morgan Kaufmann, San Francisco, Calif., 1992.

C. Boutilier. Revision sequences and nested conditionals. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)*, pages 519–525, 1993.

C. Boutilier. Toward a logic for qualitative decision theory. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pages 75–86. Morgan Kaufmann, San Francisco, Calif., 1994.

J. Burgess. Quick completeness proofs for some logics of conditionals. *Notre Dame Journal of Formal Logic*, 22:76–84, 1981.

B. F. Chellas. Basic conditional logic. *Journal of Philosophical Logic*, 4:133–153, 1980.

A. K. Chandra, D. Kozen, and L. J. Stockmeyer. Alternation. *Journal of the ACM*, 28:114–133, 1981.

J. P. Delgrande. An approach to default reasoning based on a first-order conditional logic: revised report. *Artificial Intelligence*, 36:63–90, 1988.

T. Eiter and Gottlob G. The complexity of propositional knowledge base revision. *Artificial Intelligence*, 57:227–270, 1992.

T. Eiter and Gottlob G. The complexity of nested counterfactuals and iterated knowledge base revisions. In *Proc. Thirteenth International Joint Conference on Artificial Intelligence (IJCAI '93)*, pages 526–531, 1993.

N. Friedman and J. Y. Halpern. A knowledge-based framework for belief change. Part I: foundations. In R. Fagin, editor, *Theoretical Aspects of Reasoning about Knowledge: Proc. Fifth Conference*, pages 44–64. Morgan Kaufmann, San Francisco, Calif., 1994.

N. Friedman and J. Y. Halpern. A knowledge-based framework for belief change. Part II: revision and update. In J. Doyle, E. Sandewall, and P. Torasso, editors, *Principles of Knowledge Representation and Reasoning: Proc. Fourth International Conference (KR '94)*, pages 190–201. Morgan Kaufmann, San Francisco, Calif., 1994.

R. Fagin, J. D. Ullman, and M. Y. Vardi. On the semantics of updates in databases. In *Proc. 2nd ACM Symp. on Principles of Database Systems*, pages 352–365, 1983.

M. L. Ginsberg. Counterfactuals. *Artificial Intelligence*, 30:35–79, 1986.

M. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-completeness.* W. Freeman and Co., San Francisco, Calif., 1979.

M. Goldszmidt and J. Pearl. Rank-based systems: A simple approach to belief revision, belief update and reasoning about evidence and actions. In *Principles of Knowledge Representation and Reasoning: Proc. Third International Conference (KR '92)*, pages 661–672. Morgan Kaufmann, San Francisco, Calif., 1992.

G. Grahne. Updates and counterfactuals. In *kr91*, pages 269–276. Morgan Kaufmann, San Francisco, Calif., 1991.

J. Y. Halpern and Y. Moses. A guide to completeness and complexity for modal logics of knowledge and belief. *Artificial Intelligence*, 54:319–379, 1992.

J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation.* Addison-Wesley, New York, 1979.

S. Kraus, D. Lehmann, and M. Magidor. Nonmonotonic reasoning, preferential models and cumulative logics. *Artificial Intelligence*, 44:167–207, 1990.

H. A. Kautz and B. Selman. Hard problems for simple default logics. In R. J. Brachman, H. J. Levesque, and R. Reiter, editors, *Proc. First International Conference on Principles of Knowledge Representation and Reasoning (KR '89)*, pages 189–197. Morgan Kaufmann, San Francisco, Calif., 1989.

H. Katsuno and K. Satoh. A unified view of consequence relation, belief revision and conditional logic. In *Proc. Twelfth International Joint Conference on Artificial Intelligence (IJCAI '91)*, pages 406–412, 1991.

R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

H. J. Levesque. Knowledge representation and reasoning. In J. F. Traub, B. J. Grosz, B. W. Lampson, and N. J. Nilsson, editors, *Annual Review of Computer Science, Vol. 1*, pages 255–287. Annual Reviews Inc., Palo Alto, CA, 1986.

H. J. Levesque. Logic and the complexity of reasoning. *Journal of Philosophical Logic*, 17(4):355–389, 1988.

D. K. Lewis. *Counterfactuals*. Harvard University Press, Cambridge, Mass., 1973.

D. K. Lewis. Intensional logics without iterative axioms. *Journal of Philosophical Logic*, 3:457–466, 1974.

B. Nebel. Belief revision and default reasoning: syntax-based approaches. In Allen et al. (AFS91), pages 417–428.

J. Pearl. Probabilistic semantics for nonmonotonic reasoning: a survey. In R. J. Brachman, H. J. Levesque, and R. Reiter, editors, *Proc. First International Conference on Principles of Knowledge Representation and Reasoning (KR '89)*, pages 505–516, 1989. Reprinted in *Readings in Uncertain Reasoning*, G. Shafer and J. Pearl (eds.), Morgan Kaufmann, San Francisco, Calif., 1990, pp. 699–710.

J. Pearl. From conditional oughts to qualitative decision theory. In *Proc. Ninth Conference on Uncertainty in Artificial Intelligence (UAI '93)*, pages 12–20. Morgan Kaufmann, San Francisco, Calif., 1993.

R. C. Stalnaker. A theory of conditionals. In N. Rescher, editor, *Studies in logical theory*, number 2 in American Philosophical Quarterly monograph series. Blackwell, Oxford, 1968. Also appears in *Ifs*, (ed., by W. Harper, R. C. Stalnaker and G. Pearce), Reidel, Dordrecht, Netherlands, 1981.

F. Veltman. *Logics for conditionals*. PhD thesis, Universiteit van Amsterdam, 1985.