What is an inference rule?*

Ronald Fagin
Joseph Y. Halpern
Moshe Y. Vardi

IBM Research
Almaden Research Center
650 Harry Road
San Jose, California 95120-6099

email: {fagin, halpern, vardi}@almaden.ibm.com

Abstract

What is an inference rule? This question does not have a unique answer. One usually finds two distinct standard answers in the literature: validity inference ($\sigma \vdash_v \varphi$ if for every substitution τ, the validity of $\tau[\sigma]$ entails the validity of $\tau[\varphi]$), and truth inference ($\sigma \vdash_t \varphi$ if for every substitution τ, the truth of $\tau[\sigma]$ entails the truth of $\tau[\varphi]$). In this paper we introduce a general semantic framework that allows us to investigate the notion of inference more carefully. Validity inference and truth inference are in some sense the extremal points in our framework. We investigate the relationship between various types of inference in our general framework, and consider the complexity of deciding if an inference rule is sound, in the context of a number of logics of interest: classical propositional logic, a nonstandard propositional logic, various propositional modal logics, and first-order logic.

*This version of the paper is essentially identical to one that appears in *Journal of Symbolic Logic* **57**:3, 1992, pp. 1018–1045. A preliminary version appears in *Proc. Jerusalem Conference on Information Technology*, Oct. 1990, pp. 391-401.

1 Introduction

What is logic? Quine implies that logic is a science of truths [Qui50]. Indeed, an approach sometimes taken to defining a logic is to specify its set of theorems, or valid formulas. This was, for example, the original approach taken to defining various modal logics [LL59]. However, logic originally started as the study of sound arguments. In fact, Hacking defines logic as a science of deduction [Hac79]. We would argue that in order to study the reasoning patterns appropriate to a logic, it is not sufficient to specify just the valid formulas. Indeed, there are well-known *relevance logics* [AB75, Dun86], which were specifically designed to capture legitimate patterns of reasoning about implications, that have no valid formulas at all.

In our view, it is the rules of inference of a logic that capture the patterns of reasoning that are appropriate for that logic. But this leads us to the next question: What is an inference rule?

As Avron noted [Avr91], this question does not have a unique answer. One usually finds two distinct standard answers in the literature. Assume that we have a set $\mathcal{F}$ of formulas, a class $\mathcal{S}$ of structures, and a notion of what it means for a formula in $\mathcal{F}$ to be true in a structure in $\mathcal{S}$. We write $S \models \varphi$ if φ is true in structure S. As usual, we say that a formula φ is *valid* in $\mathcal{S}$, written $\mathcal{S} \models \varphi$, if φ is true in every structure in $\mathcal{S}$.[1] We write $\sigma \vdash_v \varphi$ if for all substitution instances $\tau[\sigma]$ of σ, if $\tau[\sigma]$ is valid in $\mathcal{S}$, then the corresponding substitution instance $\tau[\varphi]$ of φ is valid in $\mathcal{S}$.[2] [3] Note that this definition is language-sensitive, in the sense that it matters what the possible substitution instances are. This is one type of inference considered in the literature. In this paper, we call this *validity inference* (with respect to $\mathcal{S}$). Some standard rules such as *universal generalization* of first-order logic ($\varphi \vdash_v \forall x\varphi$) and *necessitation* in modal logic ($\varphi \vdash_v \Box\varphi$) are validity inferences.

Notice that in defining validity inference, we actually consider *schematic rules*, i.e., inference rule *schemes*, since we quantify over all substitution instances. The distinction between schematic and nonschematic inference rules is typically unimportant when considering axioms. In all cases of which we are aware, a formula is valid iff all substitution instances of the formula are valid. However, when it comes to inference, this distinction is crucial. For example, if we did not consider substitution instances, then $p \vdash_v \mathit{false}$ would hold, where p is a primitive proposition, since p is not valid.[4] However, when viewed as

[1] If $\mathcal{S}$ is understood or not important, we may write simply $\models \varphi$ for $\mathcal{S} \models \varphi$.

[2] We provide a formal definition of *substitution instance* later in the text. For a propositional logic, it simply means replacing each primitive proposition by an arbitrary formula in the language.

[3] Throughout this paper we restrict attention to the case where there is only one formula σ on the left side of $\vdash_v$. Assuming that there is some notion of conjunction in $\mathcal{F}$ (as is the case in all the languages we consider), then a finite set can be replaced by the conjunction of its elements. We do not consider the case of an infinite set of formulas here, since the issues that will interest us here already arise with finite sets.

[4] Throughout, it is convenient to take true to be an abbreviation for some tautology (such as $p \vee \neg p$ if we are working with propositional or modal logic). We define false to be $\neg \mathit{true}$.

an inference rule scheme, this is clearly not sound: we can substitute *true* for p to get an instance of the rule where the antecedent is valid, although the conclusion is not. Note that all the familiar inference rules, such as modus ponens, universal generalization in first-order logic, and necessitation in modal logic, are actually inference rule schemes. It is interesting to note that historically the notion of axiom schemes is quite modern; Church [Chu56, page 158] says that it was von Neumann [Neu27] who is responsible for the idea of using axiom schemes rather than axioms. In contrast, inference rules were always taken to be schematic.

Another type of inference that has been considered is what we shall call *truth inference* (with respect to $\mathcal{S}$), and has usually been called *logical implication* in the literature. We write $\sigma \vdash_t \varphi$ if, for all structures $S \in \mathcal{S}$ and all substitutions τ, if $S \models \tau[\sigma]$ then $S \models \tau[\varphi]$. An axiom can be viewed as a special case of both a truth inference and a validity inference, namely, one with no antecedent (or equivalently, an antecedent that is vacuously true). If we restrict attention to axioms, then there is no difference between $\vdash_v$ and $\vdash_t$; it is easy to see that $\vdash_v \varphi$ holds iff $\vdash_t \varphi$ holds.[5] In classical propositional logic, validity inference and truth inference also coincide (see Proposition 3.1 below). However, in general, the two are different. While $\sigma \vdash_t \varphi$ implies $\sigma \vdash_v \varphi$, the converse does not always hold. For example, in first-order logic, universal generalization is a sound validity inference rule, but not a sound truth inference rule.

In this paper we introduce a general semantic framework that allows us to investigate more carefully the notion of an "appropriate pattern of reasoning", and allows us to study the inference problem in more detail. As we will see, truth inference and validity inference are in some sense the extremal points in our framework.

The basic idea underlying our framework is that in many logics of interest, the object with respect to which truth is defined is typically broken up into several components. For example, in first-order logic, if we consider formulas with free variables, then we need both a relational structure M and a valuation w to determine the truth of a formula. (A *valuation* is an assignment that associates a member of the universe of M with each variable.) In modal logic, truth is typically defined with respect to a pair (M, w) consisting of a Kripke structure $M = (W, R, \pi)$ (where W is a set of possible worlds, R is an accessibility relation between these worlds, and π associates a truth assignment with each world) and a particular world w in M. As can be seen from these examples, w depends on M. Thus, we assume that M ranges over some set $\mathcal{M}$, and that for each $M \in \mathcal{M}$, we have a set W_M such that $w \in W_M$. We then define $\mathcal{S} = \{(M, w) \mid w \in W_M\}$. Assume that truth is defined relative to a pair $(M, w) \in \mathcal{S}$. We now write $\mathcal{S} \models \varphi$ if $(M, w) \models \varphi$ for all $(M, w) \in \mathcal{S}$. The definitions of validity inference and truth inference remain unchanged. We still define $\sigma \vdash_v \varphi$ iff $\mathcal{S} \models \tau[\sigma]$ implies $\mathcal{S} \models \tau[\varphi]$ for all substitutions τ, while $\sigma \vdash_t \varphi$ iff $(M, w) \models \tau[\sigma]$ implies $(M, w) \models \tau[\varphi]$ for all $(M, w) \in \mathcal{S}$ and all substitutions τ. Considering pairs (M, w) allows us to focus on what we call $\mathcal{M}$ *inference* (with respect to $\mathcal{S}$). We say that φ is *valid in* M, and write $M \models \varphi$, if $(M, w) \models \varphi$ for all

[5] We use $\vdash \varphi$ to denote axiomhood, i.e., inference from the empty set of premises.

$w \in W_M$. We then define $\sigma \vdash_{\mathcal{M}} \varphi$ to hold iff $M \models \tau[\sigma]$ implies $M \models \tau[\varphi]$ for all $M \in \mathcal{M}$ and all substitutions τ.

More generally, we can think of a Kripke structure $M = (W, R, \pi)$ along with a world w of M to be a 4-tuple (W, R, π, w). Rather than "splitting" the tuple into the two pieces (W, R, π) and w, we could split it in other ways, which gives us other useful notions of $\mathcal{M}$ inference. We discuss this issue in depth in Section 5.

By appropriately partitioning the structures in $\mathcal{S}$ into pairs (M, w), we can view $\mathcal{M}$ inference as a generalization of both validity inference and truth inference. Essentially, if we take W_M to be degenerate, we recapture truth inference, while if we take $\mathcal{M}$ to be degenerate, we recapture validity inference. More formally, given a set $\mathcal{S}$ of structures, let $\mathcal{S}^{\mathbf{0}}$ consist of the pairs $(S, \mathbf{0})$ such that $S \in \mathcal{S}$ and $\mathbf{0}$ is some distinguished symbol, and define $(S, \mathbf{0}) \models \varphi$ iff $S \models \varphi$. It is easy to see that truth inference with respect to $\mathcal{S}$ is identical to $\mathcal{S}$ inference with respect to $\mathcal{S}^{\mathbf{0}}$. Similarly, let $\mathcal{S}^{\mathbf{1}}$ consist of the pairs $(\mathbf{1}, S)$ such that $S \in \mathcal{S}$, where $\mathbf{1}$ is some distinguished symbol, and take $(\mathbf{1}, S) \models \varphi$ iff $S \models \varphi$. Then validity inference with respect to $\mathcal{S}$ is identical to $\{\mathbf{1}\}$ inference with respect to $\mathcal{S}^{\mathbf{1}}$. No matter how we choose the $\mathcal{M}$ according to which we partition the structures in $\mathcal{S}$ into pairs, it is easy to check that truth inference is at least as strong as $\mathcal{M}$ inference, which in turn is at least as strong as validity inference. That is, $\sigma \vdash_t \varphi$ implies $\sigma \vdash_{\mathcal{M}} \varphi$, and $\sigma \vdash_{\mathcal{M}} \varphi$ implies $\sigma \vdash_v \varphi$. Thus, if we view $\vdash_t$, $\vdash_{\mathcal{M}}$, and $\vdash_v$ as binary relations on formulas, we always have $\vdash_t \subseteq \vdash_{\mathcal{M}} \subseteq \vdash_v$.

Which type of inference should be used depends on the application. Validity inference is right if the goal is to prove the validity of a given formula. It is the type of inference that allows the most inference rules. However, in order to deduce what is true in a particular situation given that other facts are true in that situation, truth inference is the appropriate type of inference. In other situations $\mathcal{M}$ inference might be more appropriate, for a particular class $\mathcal{M}$. For example, if we use tools of reasoning about knowledge to analyze protocols in distributed computing systems [Hal87], then we identify a protocol with a certain type of Kripke structure. If we are interested in studying a class of protocols which are characterized by some properties (intuitively, these properties hold in each Kripke structure corresponding to a protocol in the class), then $\mathcal{M}$ inference is the appropriate type of inference, where $\mathcal{M}$ consists of all Kripke structures corresponding to protocols.

Many of the inference rules that arise in practice are actually $\mathcal{M}$ rules for some natural choice of $\mathcal{M}$. In fact, the way one typically checks that a validity inference rule is sound is to check that it is sound when viewed as an $\mathcal{M}$ inference rule. For example, the argument for showing that universal generalization is sound runs as follows: Fix a relational structure M. If $M \models \varphi$, that is, $(M, w) \models \varphi$ for all valuations w, then it must be the case that $M \models \forall x \varphi$. The proof of soundness of necessitation in modal logic is analogous.

There is a sense in which both validity inference and truth inference give us no more information about a logic than that which is already contained in the set of valid formulas of a logic. Thus, a validity inference $\sigma \vdash_v \varphi$ is sound iff, for every substitution τ, either $\tau[\sigma]$ is not valid or $\tau[\varphi]$ is valid. In principle, this information can be obtained from looking at a list of all valid formulas. Similarly, in a logic with a notion of material implication, the truth inference $\sigma \vdash_t \varphi$ is sound iff the formula $\sigma \Rightarrow \varphi$ is valid. Again, we have reduced truth inference to a question about validity. This last reduction depends on there being a notion of material implication in the logic. The use of material implication in this reduction is actually essential. This follows from the fact, noted earlier, that in the case of certain relevance logics [AB75, Dun86] without the notion of material implication, there are no valid formulas, but there are nontrivial truth inference rules.

Sometimes $\mathcal{M}$ inference also gives us no more information than that which is already contained in the set of valid formulas. If we consider first-order logic, where $\mathcal{M}$ consists of all relational structures, then $\sigma \vdash_{\mathcal{M}} \varphi$ iff $\sigma^{\forall} \Rightarrow \varphi^{\forall}$ is valid, where we define $\psi^{\forall}$ to be the universal closure of the first-order formula ψ. This phenomenon does not seem, however, to be the case in general for $\mathcal{M}$ inference. As we shall see, there are choices of $\mathcal{M}$ for which $\mathcal{M}$ inference tells us more about the patterns of reasoning in a logic than we can obtain by simply looking at the valid formulas of the logic.

In this paper, we investigate the relationship between various types of inference in our general framework, and consider their complexity, in the context of a number of logics of interest: classical propositional logic, the nonstandard propositional logic NPL introduced in [FHV90], various propositional modal logics, and first-order logic. In some case, we have only partial results; some of the questions appear to be quite difficult. Nevertheless, we present some new techniques for answering these questions, and prove some surprising results.

Some previous work has been done in considering some of these questions for particular logics. The difference between validity inference and truth inference has been well studied, particularly in the context of intuitionistic propositional logic (IPL) (see, for example, [Min76, Tsi77, Ryb89]). In particular, Rybakov [Ryb89] showed that soundness of validity inference rules for IPL is decidable, answering a question posed by Harvey Friedman [Fri75]. He did this by showing that soundness of validity inference rules is decidable for the modal logic S4, and exploiting Gödel's well-known translation from IPL to S4.

Although we consider our notion of inference to be the most interesting one, there are two other important notions of inference that have been studied in the literature. The first notion, which we call *axiomatic inference* and denote by $\vdash^{ax}$, considers what can be inferred from a given set of axioms [Ben79]. It differs from our notion of inference in the order of quantification. For example, we define *axiomatic validity inference*, denoted $\vdash^{ax}_v$, by taking $\sigma \vdash^{ax}_v \varphi$ to mean

"$\models \tau[\sigma]$ for all substitutions τ" implies "$\models \tau[\varphi]$ for all substitutions τ".

(This is somewhat related to Avron's "extension method" [Avr91].) We can similarly define notions of *axiomatic truth inference* and *axiomatic $\mathcal{M}$ inference*. The choice of name comes from the observation that if $\sigma \vdash^{ax}_{\mathcal{M}} \varphi$, then if we take σ as an axiom, that is, if we restrict our attention to structures in $\mathcal{M}$ where every substitution instance of σ holds, then φ is a theorem, that is, every substitution instance of φ holds. It is easy to see that $\vdash_v \subseteq \vdash^{ax}_v$, and similarly for truth inference and $\mathcal{M}$ inference. In general, the inclusion is proper: for example, in propositional logic it is easy to see that if p is a primitive proposition, then $p \vdash^{ax}_v \mathit{false}$ is sound, whereas $p \vdash_v \mathit{false}$ is not sound.

The second notion is that of *nonschematic* inference rules, i.e., rules where we do not consider substitution instances [Tho75a, Tho75b, Kap87].[6] Thus, taking $\vdash^{ns}_v$ to denote *n*on*s*chematic validity inference, we have $\sigma \vdash^{ns}_v \varphi$ if $\models \sigma$ implies $\models \varphi$. As we mentioned above, nonschematic validity inference is quite different from validity inference. For example, in propositional logic, $p \vdash^{ns}_v \mathit{false}$ is sound, although $p \vdash_v \mathit{false}$ is not. Another example (due to Frege [Fre79]) of a nonschematic validity inference rule is (for an arbitrary substitution τ) the rule $\varphi \vdash^{ns}_v \tau[\varphi]$. A sound nonschematic validity inference rule can be viewed as corresponding to a *specific instance* of sound reasoning, rather than a general pattern. We can safely use a sound nonschematic validity inference rule in our reasoning, but we cannot use substitution instances of it. A powerful collection of nonschematic rules may allow us to shorten proofs. For example, it is possible that by adding Frege's inference rule to a standard axiomatization of propositional logic, we may be able to shorten proofs in propositional logic. (See [CR79] for a discussion of these issues.) Interestingly, a notion that combines axiomatic and nonschematic inference appears in Meyer, Streett, and Mirkowska [MSM81]. We say more about their work in Section 7.

The concept of inference rules, as well as related concepts such as that of *consequence relations*, has been fairly extensively studied by logicians (cf. [Gab76, Gab81]). The focus of these studies is, however, usually on the algebraic properties of these inference rules (or consequence relations) rather than on their relationship to a given semantics, which is the focus of this work. A more recent work on consequence relations is that of Avron [Avr91].[7] Avron's main focus is also quite different from ours; his focus is on various syntactic methods for classifying and representing consequence relations, while our focus is semantical in nature.

The rest of the paper is organized as follows. In Section 2, we give a brief introduction to complexity theory. In Sections 3, 4, 5, and 6 we consider validity inference, truth inference, and $\mathcal{M}$ inference in the context of classical propositional logic, the nonstandard propositional logic NPL, modal logic, and first-order logic. In Section 7, we briefly consider the two other notions of inference (axiomatic and nonschematic) discussed above.

[6] More generally, of course, we can consider a situation where we have some variables for which we can substitute, and others for which we cannot. For simplicity, we do not consider this extension here.

[7] We remark that Avron is somewhat inconsistent as to whether he considers schematic or nonschematic rules. He considers schematic rules in the case of propositional logic, and nonschematic rules for first-order logic and modal logic.

2 Definitions in complexity theory

This section is intended as a brief introduction to complexity theory, where we provide some of the basic definitions. Those readers who are familiar with complexity theory can skip it. For a more comprehensive introduction to complexity theory, we refer the reader to books by Hopcroft and Ullman [HU79] and by Garey and Johnson [GJ79], and a paper by Stockmeyer [Sto87].

Formally, we view everything in terms of the difficulty of determining membership in a set. For example, the problem of deciding whether $\sigma \vdash_v \varphi$ in propositional logic is viewed as the problem of determining whether a given possible validity inference $\sigma \vdash_v \varphi$ is a member of the set of all sound validity inferences in propositional logic. The difficulty of determining set membership is usually measured by the amount of time and/or space (memory) required to do this, as a function of the input size. For example, in the case of validity inference, we define the *size* of a formula ψ, denoted $|\psi|$, to be its length over an alphabet made up of the primitive propositions and the logical connectives, and the size of $\sigma \vdash_v \varphi$ to be $1 + |\sigma| + |\varphi|$. We will typically be interested in the difficulty of determining whether the possible validity inference $\sigma \vdash_v \varphi$ is sound, as a function of the size of $\sigma \vdash_v \varphi$. We are sometimes interested in *deterministic* computations, where at any point in a computation, the next step of the computation is uniquely determined. However, thinking in terms *nondeterministic* computations—ones where the program may "guess" which of a finite number of steps to take—has been very helpful in classifying the intrinsic difficulty of a number of problems. The complexity classes we will be most concerned with here are P, PSPACE, EXPTIME, and NP: those sets such that determining whether a given element x is a member of the set can be done in deterministic polynomial time, deterministic polynomial space, deterministic exponential time, and nondeterministic polynomial time, respectively (as a function of the size of x). It is not hard to show that P $\subseteq$ NP $\subseteq$ PSPACE $\subseteq$ EXPTIME; it is also known that P $\neq$ EXPTIME. While it is conjectured that all the other inclusions are strict, proving this remains elusive. The P = NP problem is currently considered the most important open problem in the field of computational complexity. It is perhaps worth mentioning that it is known [Sav70] that PSPACE = NPSPACE; that is, set membership can be determined in deterministic polynomial space if and only if it can be determined in nondeterministic polynomial space. Nondeterminism does not add any power at the level of polynomial space.

Roughly speaking, a set A is said to be *hard* with respect to a complexity class $\mathcal{C}$ (e.g., NP-hard, PSPACE-hard, etc.) if every set in $\mathcal{C}$ can be "efficiently" *reduced* to A (where "efficiently" is defined with respect to the class $\mathcal{C}$); i.e., for any set B in $\mathcal{C}$, an algorithm deciding membership in B can be easily obtained from an algorithm for deciding membership in A. A set is *complete* with respect to a complexity class $\mathcal{C}$, or *$\mathcal{C}$-complete*, if it is both in $\mathcal{C}$ and $\mathcal{C}$-hard.

A well-known result due to Cook [Coo71] shows that the problem of determining whether a formula of propositional logic is satisfiable (i.e., the problem of determining whether a given propositional formula is in the set of satisfiable propositional formulas) is

NP-complete. In particular, this means that if we could find a polynomial-time algorithm for deciding satisfiability for propositional logic, we would also have polynomial-time algorithms for all other NP problems. This is considered highly unlikely.

Given a complexity class $\mathcal{C}$, the class co-$\mathcal{C}$ consists of all of the sets whose *complement* is a member of $\mathcal{C}$. Notice that if we have a deterministic algorithm M for deciding membership in a set A, then it is easy to convert it to an algorithm M' for deciding membership in the complement of A that runs in the same space and/or time bounds: M' accepts an input x iff M rejects. It follows that $\mathcal{C} = \text{co}-\mathcal{C}$ must hold for every deterministic complexity class $\mathcal{C}$. This is not necessarily the case for a nondeterministic algorithm, since in this case we say that the algorithm accepts an input if it accepts for some appropriate sequence of guesses. There is no obvious way to construct an algorithm M' that will accept an element of the complement of A by an appropriate sequence of guesses. Thus, in particular, it is not known whether NP = co-NP. Clearly, if P = NP, then it would immediately follow that NP = co-NP, but it is conjectured that in fact NP $\neq$ co-NP. By way of contrast, since PSPACE = NPSPACE, it follows that NPSPACE = co-NPSPACE.

3 Inference in propositional logic

Before we consider the notion of inference in propositional logic, we define the notion of *substitution* for arbitrary propositional languages. Assume that we have a propositional language $\mathcal{F}$ with a countable set $\{p_1, p_2, \ldots\}$ of primitive propositions. A *substitution* τ is a function that associates with each primitive proposition a formula in $\mathcal{F}$. Given a formula φ, we write $\tau[\varphi]$ to indicate the result of replacing each primitive proposition p_i that appears in φ by $\tau(p_i)$.

It seems to be well known that in the context of propositional logic, truth inference and validity inference are identical. We give the proof here, partly because we could not find it in the literature, and partly because it illustrates some important ideas.

Proposition 3.1: *In propositional logic,* $\vdash_t = \vdash_v$.

Proof: As we have noted, $\sigma \vdash_t \varphi$ implies $\sigma \vdash_v \varphi$. For the converse, suppose $\sigma \vdash_v \varphi$. Let w be a truth assignment such that $w(\sigma) = \mathbf{true}$. We want to show that $w(\varphi) = \mathbf{true}$. Let τ be the substitution such that $\tau(p_i) = true$ if $w(p_i) = \mathbf{true}$ and $\tau(p_i) = false$ if $w(p_i) = \mathbf{false}$. It is easy to see that for any formula ψ, the formula $\tau[\psi]$ is valid iff $w(\psi) = \mathbf{true}$. Since $w(\sigma) = \mathbf{true}$, it follows that $\tau[\sigma]$ is valid. Since $\sigma \vdash_v \varphi$, we must have that $\tau[\varphi]$ is valid. Hence $w(\varphi) = \mathbf{true}$, as desired. ∎

Corollary 3.2: *In propositional logic, the problem of deciding whether* $\sigma \vdash_v \varphi$ *(or equivalently* $\sigma \vdash_t \varphi$*) holds is co-NP-complete.*

Proof: Since $\sigma \vdash_v \varphi$ holds iff $\sigma \vdash_t \varphi$ holds, and $\sigma \vdash_t \varphi$ holds iff $\sigma \Rightarrow \varphi$ is valid, the result follows from the co-NP-completeness of the validity problem for propositional logic. ∎

Note that we do not need to consider any notion of $\mathcal{M}$ inference in propositional logic, since by our general inclusions $\vdash_t \subseteq \vdash_{\mathcal{M}} \subseteq \vdash_v$, and by Proposition 3.1, any such notion $\vdash_{\mathcal{M}}$ would be equivalent to $\vdash_t$ and $\vdash_v$.

The proof of Proposition 3.1 shows that we can simulate a truth assignment in propositional logic by an appropriate substitution. That is precisely why we get the equivalence between truth inference and validity inference in this case. In general, even for some quite simple propositional logics, we cannot find substitutions that completely simulate truth assignments in this way. The logic NPL considered in the next section provides perhaps the simplest example of this phenomenon. In particular, $\vdash_t$ and $\vdash_v$ do not coincide in NPL.

4 Inference in NPL

NPL, a nonstandard propositional logic introduced in [FHV90], is somewhat akin to several well-known relevance logics. The major difference between NPL and classical propositional logic is in the treatment of negation. In NPL, it is consistent that both p and $\neg p$ are true, and it is consistent that neither one is true. We consider NPL here, since it provides a simple framework in which we can distinguish various types of inference.

The syntax of NPL is just like that of propositional logic, except that we introduce one new binary connective $\hookrightarrow$, which is meant to represent material implication. It turns out that $\varphi \Rightarrow \psi$, which we take to be an abbreviation for $\neg\varphi \vee \psi$, does not represent material implication because of the nonstandard semantics of negation. An *NPL structure* consists of a pair (s, t) of classical truth assignments to the primitive propositions. It is sometimes convenient to refer to s and t as *states* of the NPL structure (s, t). We take * to be a function that maps a truth assignment in a structure to the other truth assignment in that structure. Thus, if $S = (s, t)$, then $s^* = t$ and $t^* = s$. Truth in NPL is defined relative to a pair (S, u) consisting of a structure $S = (s, t)$ and a truth assignment $u \in \{s, t\}$. We define $\neg\varphi$ to be true with respect to a pair (S, u) if φ is not true at u^*; thus, we use the second truth assignment in order to define negation. (This technique for defining the semantics of negation was introduced in [RR72].) Intuitively, one could think of u as the knowledge base of all facts that are guaranteed to be true, while u^* is the knowledge base of all formulas whose negations are not guaranteed to be true. Very roughly, u could be thought of as the lower bound on the true facts, while u^* could be thought of the the upper bound on the true facts. More formally, given a structure $S = (s, t)$, and $u \in \{s, t\}$, we define:

- $(S, u) \models p$ iff $u(p) = \textbf{true}$ for a primitive proposition p
- $(S, u) \models \varphi \wedge \psi$ iff $(S, u) \models \varphi$ and $(S, u) \models \psi$

- $(S,u) \models \neg\varphi$ iff $(S,u^*) \not\models \varphi$
- $(S,u) \models \varphi \hookrightarrow \psi$ iff $(S,u) \models \psi$ whenever $(S,u) \models \varphi$.

(That is, $(S,u) \models \varphi \hookrightarrow \psi$ iff either $(S,u) \not\models \varphi$ or $(S,u) \models \psi$.)

It is not hard to show [FHV90] that if we take NPL$^-$ to be the language without $\hookrightarrow$, then there are no valid formulas in NPL$^-$. There are still some truth inferences that hold; for example, we have $\varphi \vdash_t \varphi$ and $\varphi \wedge \psi \vdash_t \varphi$. With $\hookrightarrow$ in the language, we do have valid formulas: an example is $\varphi \hookrightarrow \varphi$. It is not hard to see that $(S,u) \models \varphi \hookrightarrow \mathit{false}$ iff $(S,u) \not\models \varphi$. Thus, $\varphi \hookrightarrow \mathit{false}$ simulates classical negation. (All these results are discussed in [FHV90]; the interested reader is referred there for more intuition and detail.)

If $S = (s,t)$, then we write $S \models \varphi$, and say that φ is *valid in* S, if $(S,s) \models \varphi$ and $(S,t) \models \varphi$. Let $\mathcal{M}$ be the set of all NPL structures. This gives rise to a natural notion of $\mathcal{M}$ inference for NPL. The $\mathcal{M}$ inference $\sigma \vdash_{\mathcal{M}} \varphi$ holds in NPL if, for all structures S and all substitutions τ, we have $S \models \tau[\sigma]$ implies $S \models \tau[\varphi]$.

Truth inference, validity inference, and $\mathcal{M}$ inference are all distinct in the context of NPL. As before, viewing $\vdash_t$, $\vdash_v$, and $\vdash_{\mathcal{M}}$ as binary relations on formulas, and taking $\subset$ to denote proper containment, we have

Proposition 4.1: *In NPL,* $\vdash_t \subset \vdash_{\mathcal{M}} \subset \vdash_v$.

Proof: We already know that $\vdash_t \subseteq \vdash_{\mathcal{M}} \subseteq \vdash_v$. To see that $\vdash_t \neq \vdash_{\mathcal{M}}$, notice that $(p \hookrightarrow \mathit{false}) \vdash_t \neg p$ does not hold. For example, if we take s and t to be truth assignments such that $s(p) = \mathbf{true}$ and $t(p) = \mathbf{false}$, and take $S_0 = (s,t)$, then $(S_0,t) \models (p \hookrightarrow \mathit{false})$ and $(S_0,t) \not\models \neg p$. On the other hand, it is easy to check that we do have $(p \hookrightarrow \mathit{false}) \vdash_{\mathcal{M}} \neg p$: if $p \hookrightarrow \mathit{false}$ is true at both truth assignments in a structure, then p cannot be true at either one, so $\neg p$ is true at both. (The inference rule "from $\varphi \hookrightarrow \mathit{false}$ infer $\neg\varphi$" is the *negation replacement rule* for NPL in [FHV90].)

In order to show that $\vdash_{\mathcal{M}} \neq \vdash_v$, let φ_0 be the formula $(p \hookrightarrow \neg p) \wedge (\neg p \hookrightarrow p)$. For the structure S_0 constructed above, it is easy to check that $S_0 \models \varphi_0$. Since $S_0 \not\models \mathit{false}$, the $\mathcal{M}$ inference $\varphi_0 \vdash_{\mathcal{M}} \mathit{false}$ does not hold. However, we now show that $\varphi_0 \vdash_v \mathit{false}$ does hold, so $\vdash_{\mathcal{M}} \neq \vdash_v$. We first show that no substitution instance of the formula φ_0 can be valid. Define a structure S to be *standard* if $S = (s,s)$ for some truth assignment s. Given an NPL formula φ, let φ^t be the classical formula that results by replacing all occurrences of $\hookrightarrow$ by $\Rightarrow$. A straightforward argument by induction on the structure of ψ shows that if $S = (s,s)$, then for each formula ψ, we have $(S,s) \models \psi$ iff $s(\psi^t) = \mathbf{true}$ (using the standard semantics of propositional logic). Since $(\tau[\varphi_0])^t$ is equivalent to false for any substitution τ, it follows that $\tau[\varphi_0]$ is not valid in any standard structure S. Thus, no substitution instance of φ_0 can be valid. As a consequence, the validity inference $\varphi_0 \vdash_v \mathit{false}$ holds. ∎

Despite the fact that the three types of inference are all distinct, as we now show, they have the same complexity. We describe the proofs of these results in some detail here, since the techniques used extend to other, more complicated, logics. We make use of the following result on the validity problem for NPL.

Proposition 4.2: [FHV90] *The validity problem for NPL is co-NP-complete.*

We now show that in NPL, each of the three types of inference have the same complexity as the validity problem for NPL, namely, co-NP-complete.

Theorem 4.3: *The problem of deciding whether $\sigma \vdash_t \varphi$ (resp. $\sigma \vdash_{\mathcal{M}} \varphi$, $\sigma \vdash_v \varphi$) holds for NPL is co-NP-complete.*

Proof: The lower bound follows in each case from Proposition 4.2, since we can encode the validity problem for NPL by taking σ to be $true$.

The upper bound for $\vdash_t$ follows since $\sigma \vdash_t \varphi$ holds iff $\sigma \hookrightarrow \varphi$ is valid. Thus, the truth inference problem for NPL is co-NP-complete.

In order to deal with $\vdash_{\mathcal{M}}$, given a formula φ, let φ^* be the formula $\neg(\varphi \hookrightarrow false)$. It is easy to check that $(S, s) \models \varphi^*$ iff $(S, s^*) \models \varphi$. By using φ^*, we have a way of saying that φ is valid in S, since $S \models \varphi$ iff $(S, s) \models \varphi \wedge \varphi^*$. It follows that $\sigma \vdash_{\mathcal{M}} \varphi$ holds iff $(\sigma \wedge \sigma^*) \hookrightarrow (\varphi \wedge \varphi^*)$ is valid. Thus, the $\mathcal{M}$ inference problem for NPL is co-NP-complete.

It might seem that dealing with validity inference should be straightforward. To show that $\sigma \vdash_v \varphi$ it seems to be enough to show that either φ is valid or that σ is not valid. This is not quite true, because we must consider substitution instances. For example, although p is not valid, it is not the case that the inference $p \vdash_v q$ holds. While having to deal with substitution instances may seem to complicate things (in that there are infinitely many substitution instances to consider), it turns out that being able to substitute gives us some control over the problem, and in fact, as we shall see, might make it easier.

In dealing with validity inference, it is useful to consider the dual problem, which bears the same relationship to validity inference as satisfiability does to validity. We say that φ is *compatible with σ being valid* if there is a substitution τ such that $\tau[\sigma]$ is valid and $\tau[\varphi]$ is satisfiable. Note that φ is compatible with σ being valid iff $\sigma \not\vdash_v (\varphi \hookrightarrow false)$. Similarly, $\sigma \vdash_v \varphi$ iff it is not the case that $(\varphi \hookrightarrow false)$ is compatible with σ being valid.

We shall show that the compatibility problem (the problem of deciding whether φ is compatible with σ being valid) for NPL is in NP. It then follows that the problem of deciding validity inference is in co-NP, since as we noted, $\sigma \vdash_v \varphi$ iff it is not the case that $(\varphi \hookrightarrow false)$ is compatible with σ being valid. The proof technique that we shall use to prove that the compatibility problem for NPL is in NP will be used again later for other logics. In each case, we shall show that φ is compatible with σ being valid precisely if (a) there is some structure where σ is valid and where φ is satisfiable, and (b) σ is valid in some "special" structure or small set of special structures. In the case of NPL, "special" means standard. It is immediate that (a) and (b) are in NP, so the compatibility problem is in NP. The reason that (a) and (b) imply that φ is compatible with σ being valid is that we define a substitution τ which, intuitively, makes each structure behave as if it were one of the structures described in (a) and (b), so that in particular, $\tau[\varphi]$ is satisfiable and $\tau[\sigma]$ is valid.

In order to make this intuition work, we need to be able to understand the effect of a substitution on the truth value of φ. Clearly, the effect of the substitution depends only on the truth values (at both states) of the formulas being substituted for the primitive propositions. This is made precise in Lemma 4.5 below. We define the *truth status* of a formula φ in state u of structure S to be T if $(S, u) \models \varphi$, and F otherwise. More generally, we define the truth status of a vector $\langle \varphi_1, \ldots, \varphi_n \rangle$ of formulas in state u of structure S to be the vector $\langle L_1, \ldots, L_n \rangle$, where L_i is the truth status of φ_i in state u of S. The truth status of $\langle \varphi_1, \ldots, \varphi_n \rangle$ in structure $S = (s, t)$, written $ts(\langle \varphi_1, \ldots, \varphi_n \rangle, S)$, is the pair (v_1, v_2), where v_1 is the truth status in state s and v_2 is the truth status in state t. We say that a truth status (v_1, v_2) is *standard* if $v_1 = v_2$. Define $TS(\langle \varphi_1, \ldots, \varphi_n \rangle)$, the *truth set* of $\langle \varphi_1, \ldots, \varphi_n \rangle$, to be the set consisting of $\{ts(\langle \varphi_1, \ldots, \varphi_n \rangle, S)\}$, as S ranges over all structures. Let TS_n be the set of all conceivable truth sets for vectors of length n, that is, all sets consisting of pairs (v_1, v_2) where v_1 and v_2 are (not necessarily distinct) vectors $\langle L_1, \ldots, L_n \rangle$ of truth values T or F. We now investigage which members of TS_n are actually attained as the truth set of some vector of formulas.

To understand the issue, consider the case of one formula. For ease in notation, here we write, say, TF for $(\langle T \rangle, \langle F \rangle)$. In this case, it is easy to check that we have $TS(true) = \{TT\}$, $TS(false) = \{FF\}$, $TS((p \hookrightarrow \neg p) \wedge (\neg p \hookrightarrow p)) = \{TT, FF\}$, $TS(p \vee \neg p) = \{TT, TF, FT\}$, $TS(p \wedge \neg p) = \{TF, FT, FF\}$, and $TS(p) = \{TT, TF, FT, FF\}$. Note that whenever TF appears in a truth set, so does FT. This is not an accident: this is because if $ts(\langle \varphi \rangle, (s, t))$ is TF, then $ts(\langle \varphi \rangle, (t, s))$ is FT. Note also that we do not have an example of a formula φ such that $TS(\varphi) = \{TF, FT\}$. This is also not an accident. It is clear that the truth set of every formula must include at least one standard truth status, since if S is a standard structure, then $ts(\varphi, S)$ must be standard. The next lemma says that these are the only restrictions we have on truth sets.

Lemma 4.4: *Let A be a member of TS_n such that*

1. *whenever $(v_1, v_2) \in A$, then $(v_2, v_1) \in A$, and*
2. *A contains at least one standard truth status.*

Then there exist formulas $\varphi_1, \ldots, \varphi_n$ such that $TS(\langle \varphi_1, \ldots, \varphi_n \rangle) = A$.

Proof: Let $p_1, \ldots, p_k$ be primitive propositions. It is easy to see that if the number k of primitive propositions is picked to be sufficiently large, then there is a function f that associates with every structure (s, t), where s and t are each truth assignments over $p_1, \ldots, p_k$, a member $f(s, t) \in A$ such that

1. if $f(s, t) = (v_1, v_2)$, then $f(t, s) = (v_2, v_1)$, and

2. f is onto A, that is, for each $(v_1, v_2) \in A$, there is (s, t) such that $f(s, t) = (v_1, v_2)$.

(The number k of primitive propositions must be sufficiently large to make condition (2) possible.) Note that condition (1) implies that for each standard structure (s, s) we must have $f(s, s)$ be a standard truth status (v_1, v_1). This is possible, since by assumption A contains at least one standard truth status.

Define an *atomic description* to be a formula $\psi_1 \wedge \ldots \wedge \psi_k \wedge \gamma_1 \wedge \ldots \wedge \gamma_k$, where ψ_i is either p_i or $(p_i \hookrightarrow false)$, for $1 \leq i \leq k$, and where γ_i is either $\neg(p_i \hookrightarrow false)$ or $\neg p_i$, for $1 \leq i \leq k$. If $S = (s, t)$ is a structure, then it is easy to see that $(S, s) \models \alpha$ for exactly one atomic description $\alpha(s, t)$, namely the atomic description $\psi_1 \wedge \ldots \wedge \psi_k \wedge \gamma_1 \wedge \ldots \wedge \gamma_k$ where ψ_i is p_i precisely if p_i is true under the truth assignment s, and where γ_i is $\neg(p_i \hookrightarrow false)$ precisely if p_i is true under the truth assignment t, for $1 \leq i \leq k$.

We are now ready to define the formulas $\varphi_1, \ldots, \varphi_n$. We let φ_i (for $1 \leq i \leq n$) be the disjunction of all atomic descriptions $\alpha(s, t)$ such that if $f(s, t) = (v_1, v_2)$, then the ith component of the vector v_1 is T. We now show that if $S = (s, t)$, and if $f(s, t) = (v_1, v_2)$, then $(S, s) \models \varphi_i$ iff the ith component of the vector v_1 is T. If $(S, s) \models \varphi_i$, then $\alpha(s, t)$ is a disjunct of φ_i, so the ith component of the vector v_1 is T. Conversely, if the ith component of the vector v_1 is T, then $\alpha(s, t)$ is a disjunct of φ_i, so $(S, s) \models \varphi_i$. From what we just showed, it follows that the truth status of $\langle \varphi_1, \ldots, \varphi_n \rangle$ in state s of structure S is v_1. Since by construction $f(t, s) = (v_2, v_1)$, it follows identically that the truth status of $\langle \varphi_1, \ldots, \varphi_n \rangle$ in state t of structure (t, s) is v_2. But this latter truth status is the same as the truth status in state t of structure $S = (s, t)$. So $ts(\langle \varphi_1, \ldots, \varphi_n \rangle, S) = (v_1, v_2)$. Hence, $TS(\langle \varphi_1, \ldots, \varphi_n \rangle)$ is the range of f, that is, A. ∎

The next lemma is straightforward, and the proof is left to the reader.

Lemma 4.5: *Let ψ be an NPL formula and $S = (s, t)$ an NPL structure. Assume that the substitution τ replaces the primitive propositions $p_1, \ldots, p_n$ by $\varphi_1, \ldots, \varphi_n$, and that the truth status of $\langle \varphi_1, \ldots, \varphi_n \rangle$ in S is (v_1, v_2). Let $S' = (s', t')$ be the structure where the truth status of $\langle p_1, \ldots, p_n \rangle$ in S' is (v_1, v_2). If $(S, s) \models \tau[\psi]$, then $(S', s') \models \psi$. Furthermore, if S is standard, then so is S'.*

Corollary 4.6: *Let ψ be an NPL formula, and let τ be a substitution. If $\tau[\psi]$ is satisfiable, then so is ψ. Furthermore, if $\tau[\psi]$ is satisfiable in a standard structure, then so is ψ.*

The next lemma is the key step in showing that compatibility problem is in NP.

Lemma 4.7: *φ is compatible with σ being valid iff (a) $\sigma \wedge \sigma^* \wedge \varphi$ is satisfiable and (b) σ is satisfiable in some standard structure.*

Proof: Let $p_1, \ldots, p_n$ be the primitive propositions that appear in φ or σ. Assume first that φ is compatible with σ being valid. Then there is a substitution τ such that $\tau[\sigma]$ is valid and $\tau[\varphi]$ is satisfiable. Since $\tau[\varphi]$ is satisfiable, there is a structure $S = (s,t)$ such that $(S,s) \models \tau[\varphi]$. Since $\tau[\sigma]$ is valid, we know that $(S,s) \models \tau[\sigma]$ and $(S,t) \models \tau[\sigma]$, that is $(S,s) \models (\tau[\sigma])^*$. It is easy to see that that $(\tau[\sigma])^*$ equals $\tau[\sigma^*]$, so $(S,s) \models \tau[\sigma \wedge \sigma^* \wedge \varphi]$. Hence, $\tau[\sigma \wedge \sigma^* \wedge \varphi]$ is satisfiable. By Corollary 4.6, $\sigma \wedge \sigma^* \wedge \varphi$ is satisfiable. Therefore, (a) holds. Since $\tau[\sigma]$ is valid, it is certainly satisfiable in some standard structure. By Corollary 4.6, σ is satisfiable in some standard structure. This proves (b).

For the converse, assume that $S = (s,t)$ is a structure and $S' = (s', s')$ is a standard structure such that $(S,s) \models \sigma \wedge \sigma^* \wedge \varphi$, and $(S', s') \models \sigma$. Suppose that the truth status of $\langle p_1, \ldots, p_n \rangle$ in S is (v_1, v_2), and the truth status of $\langle p_1, \ldots, p_n \rangle$ in S' is (v_1', v_1'). Let $A = \{(v_1, v_2), (v_2, v_1), (v_1', v_1')\}$. By Lemma 4.4, there exist formulas $\varphi_1, \ldots, \varphi_n$ such that $TS(\langle \varphi_1, \ldots, \varphi_n \rangle) = A$. Let τ be the substitution that substitutes $\varphi_1, \ldots, \varphi_n$ for $p_1, \ldots, p_n$ respectively. Since σ is true at (S,s), (S,t), and (S', s'), it follows easily that $\tau[\sigma]$ is valid. Since $(S,s) \models \varphi$, it follows easily that $\tau[\varphi]$ is satisfiable. So φ is compatible with σ being valid, as desired. ∎

Using Lemma 4.7, it is easy to see that the compatibility problem for NPL is in NP: we simply have to guess the appropriate structures. Thus, the validity inference problem for NPL is co-NP-complete. ∎

5 Inference in propositional modal logic

We assume that the reader is familiar with the basic semantics of propositional modal logic. We provide a brief review here, referring the reader to one of the standard modal logic texts (e.g., [Che80, HC68]) for more details.

A *frame* F is a pair (W, R), where W is a set of *worlds* and R is a binary relation on W. A *(Kripke) structure* M is a triple (W, R, π) where (W, R) is a frame and π is a mapping associating with each world $w \in W$ a truth assignment to the primitive propositions. We say that a structure $M = (W, R, \pi)$ is *based on* the frame F if $F = (W, R)$.

The language for the modal logic consists of propositional logic augmented by one modal operator $\Box$. Truth is defined relative to a structure M and world w of M We take $\Box\varphi$ to be true at a world w if φ is true at every world accessible from w. We define $\Diamond\varphi$ to be an abbreviation for $\neg\Box\neg\varphi$; thus, $\Diamond\varphi$ is true at w exactly if φ is true at some world accessible from w.

- $(M, w) \models p$ iff $\pi(w)(p) = \textbf{true}$ for a primitive proposition p
- $(M, w) \models \varphi \wedge \psi$ iff $(M, w) \models \varphi$ and $(M, w) \models \psi$
- $(M, w) \models \neg\varphi$ iff $(M, w) \not\models \varphi$
- $(M, w) \models \Box\varphi$ iff $(M, w') \models \varphi$ for all w' such that $(w, w') \in R$.

We write $M \models \varphi$ if $(M, w) \models \varphi$ for all worlds w in M, and write $F \models \varphi$ if $M \models \varphi$ for all structures M based on F. In the latter case, we say that φ is *valid in the frame* F.

The modal logic just described is typically referred to as the modal logic K. We can obtain various other modal logics by restricting to certain types of frames. We focus here on K and S5. To obtain the modal logic S5, we allow only those frames (W, R) where R is an equivalence relation.

The complexity of the validity problem for modal logic has been well studied. We will need the following result:

Theorem 5.1: [Lad77] *The validity problem for K is PSPACE-complete, while the validity problem for S5 is co-NP-complete.*

Once we move to modal logics, there are several natural notions of $\mathcal{M}$ inference. An individual world w of a specific Kripke structure (W, R, π) can be thought of as a tuple (W, R, π, w). As we now describe, we can obtain different types of inference rules by holding fixed various parts of the tuple (W, R, π, w). Thus, in truth inference we fix the whole tuple (W, R, π, w), so that we are considering only one world of one structure. Intuitively, we care about truth inference when we are trying to decide whether a formula is true at a *specific* world of a *specific* Kripke structure. In validity inference, nothing is fixed, so that we are considering all worlds of all structures. Intuitively, we care about validity inference when we are trying to decide whether a formula is true at *every* world of *every* Kripke structure. As we now describe, there are other natural choices.

One choice, which we call *structure inference* and denote by $\vdash_s$, corresponds to fixing (W, R, π) and varying w, so that we are considering all the worlds in one structure. Formally, $\sigma \vdash_s \varphi$ is a sound structure inference rule precisely if $M \models \tau[\sigma]$ implies $M \models \tau[\varphi]$ for all structures M and substitutions τ. Intuitively, we care about structure inference when we are trying to decide whether a formula is true at *every* world of a *specific* Kripke structure.

Another choice, which we call *frame inference* and denote by $\vdash_f$, corresponds to fixing (W, R) and varying π and w. Formally, $\sigma \vdash_f \varphi$ is a sound frame inference rule precisely if $F \models \tau[\sigma]$ implies $F \models \tau[\varphi]$ for all frames F and substitutions τ. Intuitively, we care about structure inference when we are trying to decide whether a formula is true at *every* world of a *specific set* of Kripke structures, namely those Kripke structures with a given frame. We remark that the notion of a sequent being *valid* in modal logic, as defined in [Avr91], corresponds to frame inference; there is no notion in [Avr91] corresponding to structure inference.

It is easy to check that, when viewed as relations on formulas, we have $\vdash_t \subseteq \vdash_s \subseteq \vdash_f \subseteq \vdash_v$. As we shall see, in general the containments are proper. We remark that Humberstone [Hum86] has considered $\vdash_t$, $\vdash_s$, and $\vdash_f$ in the context of modal logic (calling them *inferential consequence*, *model consequence*, and *frame consequence*, respectively), as well as a fourth type of $\mathcal{M}$ inference, which he called *point consequence*, where $\mathcal{M}$ consists of pairs (F, w), where F is a frame and w is a world. He shows that inference rules provide a more general way of distinguishing between modal logics than axiom systems. In particular, he shows that there are classes of frames that agree on the axioms that they satisfy, but not on the inference rules that they satisfy.

5.1 Inference in K

We start our investigation with the modal logic K. In K, the various types of inference are in fact distinct.

Theorem 5.2: *In the modal logic K, we have* $\vdash_t \subset \vdash_s \subset \vdash_f \subset \vdash_v$.

Proof: To see that $\vdash_t \neq \vdash_s$, notice that $p \vdash_s \Box p$ holds, while $p \vdash_t \Box p$ does not.

To see that $\vdash_s \neq \vdash_f$, we first show that no substitution instance of the formula $\Diamond p \wedge \Diamond \neg p$ can be valid in any frame. For suppose $F \models \Diamond\varphi \wedge \Diamond\neg\varphi$ for some frame $F = (W, R)$. Clearly there cannot be a world w in F such that no worlds w' are accessible from w, for then $\Diamond\varphi \wedge \Diamond\neg\varphi$ must surely be false at w, no matter what truth assignment we choose. Thus, we can suppose that every world in F has some world accessible from it. Choose a mapping π that associates the same truth assignment with every world w in F. Now an easy induction on the structure of formulas shows that every formula ψ has the same truth value at every world w in F under π. Thus, one of $\Box\varphi$ or $\Box\neg\varphi$ must be valid in the structure $M = (W, R, \pi)$. It follows that $\Diamond\varphi \wedge \Diamond\neg\varphi$ is not valid in F. Therefore, $\Diamond p \wedge \Diamond\neg p \vdash_f \mathit{false}$ holds; however, it is easy to check that $\Diamond p \wedge \Diamond\neg p \vdash_s \mathit{false}$ does not hold.

Finally, to show that $\vdash_f \neq \vdash_v$, note that $\Box p \vdash_v p$ holds, but $\Box p \vdash_f p$ does not,[8] as we can see by considering a frame (W, R) where $R = \emptyset$. ∎

Turning to complexity, since we can again reduce truth inference to validity, we can show:

Theorem 5.3: *The truth inference problem for K is PSPACE-complete.*

As we did with $\mathcal{M}$ inference for NPL, we would like to deal with structure inference for K by reducing the question of whether the structure inference $\sigma \vdash_s \varphi$ holds to a question of the validity of a related formula. We cannot quite do this, since the language of K is too weak. Roughly speaking, we would like a formula φ^* such that φ^* is true at a world s in structure M exactly if φ is true at every world of M. We cannot quite achieve this; we can come close by slightly extending K. Let R^+ be the transitive closure of the binary relation R. Suppose we extend K by adding a new modal operator $\Box^+$, and define

- $(M, w) \models \Box^+\varphi$ iff $(M, t) \models \varphi$ for all t such that $(s, t) \in R^+$.

[8] We are indebted to Arnon Avron for pointing out this last example to us.

The next lemma shows that structure inference can be reduced to validity in the enriched language. The observation that such a reduction is possible was made first by Goranko and Passy [GP89].

Proposition 5.4: *$\sigma \vdash_s \varphi$ holds for K iff $\Box^+\sigma \Rightarrow \Box^+\varphi$ is valid for K.*

Proof: Assume first that $\sigma \vdash_s \varphi$ holds for K, that $M = (W, R, \pi)$ is a structure, and that $(M, w) \models \Box^+\sigma$; we must show that $(M, w) \models \Box^+\varphi$. Let $W' = \{w' \mid (w, w') \in R^+\}$, let R' and π' be the restrictions of R and π, respectively, to W', and let $M' = (W', R', \pi')$. By a straightforward induction on the structure of formulas, it follows that for every formula ψ and for every world $s \in W'$, we have

$$(M, s) \models \psi \text{ iff } (M', s) \models \psi. \tag{1}$$

Since $(M, w) \models \Box^+\sigma$, we know that $(M, s) \models \sigma$ for every $s \in W'$, so by (1), it follows that $(M', s) \models \sigma$ for every $s \in W'$. Since $\sigma \vdash_s \varphi$, it follows that $(M', s) \models \varphi$ for every $s \in W'$. By (1) again, $(M, s) \models \varphi$ for every $s \in W'$. Therefore, $(M, w) \models \Box^+\varphi$, as desired.

Conversely, assume that $\Box^+\sigma \Rightarrow \Box^+\varphi$ is valid for K; we must show that $\sigma \vdash_s \varphi$ holds for K. Assume that $M \models \tau[\sigma]$ for some structure $M = (W, R, \pi)$ and substitution τ; we must show that $M \models \tau[\varphi]$. Define π' on W by letting $\pi'(s)$ be the truth assignment that makes the primitive proposition p true iff $(M, s) \models \tau[p]$. Let $M' = (W, R, \pi')$. By a straightforward induction on the structure of formulas, it follows that for every formula ψ and for every world $s \in W$, we have

$$(M, s) \models \tau[\psi] \text{ iff } (M', s) \models \psi. \tag{2}$$

Since $M \models \tau[\sigma]$, it follows from (2) that $M' \models \sigma$. Let a be a new world not in W, let $W'' = W \cup \{a\}$, and let $R'' = R \cup \{(a, s) \mid s \in W\}$. Intuitively, we are adding a new world that has an edge to every old world. Define π'' by letting $\pi''(s) = \pi'(s)$ if $s \in W$, and letting $\pi''(a)$ be an arbitrary truth assignment. Let $M'' = (W'', R'', \pi'')$. By a straightforward induction on the structure of formulas, it follows that for every formula ψ and for every world $s \in W$, we have

$$(M', s) \models \psi \text{ iff } (M'', s) \models \psi. \tag{3}$$

Since $M' \models \sigma$, it follows from (3) that $(M'', s) \models \sigma$ for every $s \in W$, so $(M'', a) \models \Box^+\sigma$. Since $\Box^+\sigma \Rightarrow \Box^+\varphi$ is valid for K, it follows that $(M'', a) \models \Box^+\varphi$. So $(M'', s) \models \varphi$ for every $s \in W$. By (3) again, $M' \models \varphi$. By (2), $M \models \tau[\varphi]$, as desired. ∎

Formulas of the form $\Box^+\sigma \Rightarrow \Box^+\varphi$ can be viewed as a fragment of propositional dynamic logic (PDL). The fact that there is an exponential-time decision procedure for validity in PDL [Pra79] gives us an exponential-time decision procedure for validity of these formulas. Furthermore, the proof of [FL79] that the validity problem for PDL is exponential-time hard applies also to this fragment. Putting these results together, we see that the validity problem for formulas of the form $\Box^+\sigma \Rightarrow \Box^+\varphi$ is EXPTIME-complete. Therefore, from Proposition 5.4, we get:

Theorem 5.5: *The structure inference problem for K is EXPTIME-complete.*

Note that in this case, the structure inference problem is harder than the truth inference problem (provided PSPACE $\neq$ EXPTIME). Moreover, although we have reduced structure inference for K to a question of validity, it is validity in a more expressive logic than K.

We do not know whether either the validity inference problem or the frame inference problem is decidable for K. We note that Rybakov [Ryb89], who showed that the validity inference problem for S4 is decidable, left the validity inference problem for K as an open problem; it seems to be quite difficult.

5.2 Inference in S5

The situation for S5 is quite different from that for K. For one thing, $\vdash_f$ and $\vdash_v$ turn out to be identical in S5.

Theorem 5.6: *In the modal logic S5, we have* $\vdash_t \subset \vdash_s \subset \vdash_f = \vdash_v$.

Proof: Again, to see that $\vdash_t \neq \vdash_s$, note that $p \vdash_s \Box p$ holds, while $p \vdash_t \Box p$ does not. The proof that $\vdash_s \neq \vdash_f$ is essentially the same as that for K, again using the formula $\Diamond p \wedge \Diamond \neg p$. We leave details to the reader. The fact that $\vdash_f = \vdash_v$ follows from the proof of Theorem 5.7 below. ∎

Theorem 5.7: *The problem of deciding whether $\sigma \vdash_t \varphi$ (resp. $\sigma \vdash_s \varphi$, $\sigma \vdash_f \varphi$, $\sigma \vdash_v \varphi$) holds for S5 is co-NP-complete.*

Proof: The lower bound follows from Theorem 5.1, since we can encode the validity problem for S5 in all cases by taking σ to be *true*.

For the upper bound in the case of truth inference, note that $\sigma \vdash_t \varphi$ holds iff $\sigma \Rightarrow \varphi$ is valid in S5; thus the result follows from Theorem 5.1. In the case of structure inference, it is easy to show that $\sigma \vdash_s \varphi$ iff $\Box\sigma \Rightarrow \Box\varphi$ is valid in S5. Again, the desired co-NP-completeness now follows from Theorem 5.1.

In the case of frame inference and validity inference, we need to use techniques of the same flavor as those used in the validity inference case in the proof of Theorem 4.3. As in the NPL case earlier, we say that φ is *compatible with σ being valid* if there is a substitution τ such that $\tau[\sigma]$ is valid and $\tau[\varphi]$ is satisfiable. Similarly to before, φ is compatible with σ being valid if $\sigma \not\vdash_v \neg\varphi$, and $\sigma \vdash_v \varphi$ iff it is not the case that $\neg\varphi$ is compatible with σ being valid. Similarly, we say that φ is *frame compatible with σ being valid* if there is a substitution τ and a frame F such that $\tau[\sigma]$ is valid in the frame F and $\tau[\varphi]$ is satisfiable in the frame F (by the latter, we mean that there is a structure M based on F and a world s of M such that $(M, s) \models \tau[\varphi]$). As before, φ is frame compatible with σ being valid if $\sigma \not\vdash_f \neg\varphi$, and $\sigma \vdash_f \varphi$ iff it is not the case that $\neg\varphi$ is frame compatible with σ being valid.

Similarly to before, we define the *truth status* of a formula φ in world s of Kripke structure M to be T if $(M, s) \models \varphi$, and F otherwise. More generally, we define the truth status of a vector $\langle \varphi_1, \ldots, \varphi_n \rangle$ of formulas in world s of structure M to be the vector $\langle L_1, \ldots, L_n \rangle$, where L_i is the truth status of φ_i in world s of M. We need analogs of Lemma 4.5 and Corollary 4.6. Again, the proofs are straightforward, and are left to the reader.

Lemma 5.8: *Let ψ be a formula and $M = (W, R, \pi)$ be a Kripke structure. Assume that the substitution τ replaces the primitive propositions $p_1, \ldots, p_n$ by $\varphi_1, \ldots, \varphi_n$, and that for each world t, the truth status of $\langle \varphi_1, \ldots, \varphi_n \rangle$ in t is v_t. Let $M' = (W, R, \pi')$ be a Kripke structure with the same set W of worlds and the same accessibility relation R as M, but where for each world t, the truth status of $\langle p_1, \ldots, p_n \rangle$ in world t is v_t. If $(M, s) \models \tau[\psi]$, then $(M', s) \models \psi$.*

Corollary 5.9: *Let ψ be a formula, and let τ be a substitution. If $\tau[\psi]$ is satisfiable, then so is ψ. Furthermore, ψ is satisfiable in a structure with the same frame as $\tau[\psi]$.*

We shall show that the (frame) compatibility problem, that is, the problem of deciding whether φ is (frame) compatible with σ being valid, is in NP. It then follows that the problem of deciding validity inference and frame inference is in co-NP, as before. We use a similar proof technique to that we used to prove that the compatibility problem for NPL is in NP: we show that φ is (frame) compatible with σ being valid precisely if (a) there is some structure where σ is valid and where φ is satisfiable, and (b) σ is valid in some "special" structure. Here, "special" means a structure with only one world. Thus, we show the following lemma, somewhat analogous to Lemma 4.7.

Lemma 5.10: *The following are equivalent:*

(a) *φ is compatible with σ being valid.*

(b) *φ is frame compatible with σ being valid.*

(c) *(i) $\varphi \wedge \Box\sigma$ is satisfiable and (ii) σ is satisfiable in some structure with only one world.*

Proof: (a) $\Rightarrow$ (b): Let τ be a substitution such that $\tau[\sigma]$ is valid and $\tau[\varphi]$ is satisfiable. Let F be a frame in which $\tau[\varphi]$ is satisfiable. Since $\tau[\sigma]$ is valid, it is certainly valid in the frame F. It then follows immediately from the definition that φ is frame compatible with σ being valid.

(b) $\Rightarrow$ (c): Let τ be a substitution and F a frame such that $\tau[\sigma]$ is valid in the frame F and $\tau[\varphi]$ is satisfiable in the frame F. Let M be a structure based on F and s a world of M such that $(M, s) \models \tau[\varphi]$. So $(M, s) \models \tau[\varphi] \wedge \Box\tau[\sigma]$, that is, $(M, s) \models \tau[\varphi \wedge \Box\sigma]$. So $\tau[\varphi \wedge \Box\sigma]$ is satisfiable. By Corollary 5.9, $\varphi \wedge \Box\sigma$ is satisfiable.

Let $M = (W, R, \pi)$ be a structure with the frame F as given above, such that for every $s \in W$, the truth assignment $\pi(s)$ makes every primitive proposition true. Then $M \models \tau[\sigma]$, since $\tau[\sigma]$ is valid in the frame F. Now let $M' = (W', R', \pi')$ be a structure where W' is a singleton set $\{s'\}$, where $R' = \{(s', s')\}$, and where $\pi'(s')$ is a truth assignment that makes every primitive proposition true. By a simple induction on formulas, it follows that for each formula ψ, we have $M \models \psi$ iff $M' \models \psi$. Therefore, $M' \models \tau[\sigma]$. Hence, $\tau[\sigma]$ is satisfiable in a structure with only one world. By Corollary 5.9, σ is satisfiable in a structure with only one world.

(c) $\Rightarrow$ (a): Assume that the primitive propositions are $p_1, \ldots, p_k$. Similarly to before, let us define an *atomic description* to be a formula $\psi_1 \wedge \ldots \wedge \psi_k$, where ψ_i is either p_i or $\neg p_i$, for $1 \leq i \leq k$. Let $\mathcal{A}$ be the set of all of the 2^k atomic descriptions. If $\mathcal{S}$ is a set of atomic descriptions, then define $\Box!\mathcal{S}$ to be the formula

$$\bigwedge_{\alpha \in \mathcal{S}} \Diamond\alpha \quad \wedge \quad \bigwedge_{\alpha \in \mathcal{A} - \mathcal{S}} \Box\neg\alpha.$$

We now explain why we are interested in formulas of the form $\Box!\mathcal{S}$. Let $M = (W, R, \pi)$ be such that R is *universal* (that is, $R = W \times W$). If $s \in W$, then $(M, s) \models \Box!\mathcal{S}$ precisely if the set $\{\pi(s') \mid s' \in W\}$ of truth assignments are those described by members of $\mathcal{S}$. If the truth assignment $\pi(s)$ is described by the atomic description α, and if the set $\{\pi(s') \mid s' \in W\}$ of truth assignments are those described by members of $\mathcal{S}$, then it is well-known that the formula $\alpha \wedge \Box!\mathcal{S}$ completely characterizes (M, s), in the sense that if $(M', s') \models \alpha \wedge \Box!\mathcal{S}$ for some (M', s'), then for every formula δ, we have $(M, s) \models \delta$ iff $(M', s') \models \delta$. Intuitively, α tells what the truth assignment is in the current world, and $\Box!\mathcal{S}$ tells what all the truth assignments are in the worlds of the structure. It follows easily that in S5, every formula ψ is equivalent to a disjunction of formulas of the form $\alpha \wedge \Box!\mathcal{S}$, where $\alpha \in \mathcal{S}$. Intuitively, this disjunction describes the set of structures where ψ is satisfied. Without loss of generality, we can assume that φ and σ are both of this form.

Since $\varphi \wedge \Box\sigma$ is satisfiable, it is straightforward to show that there is a nonempty set $\mathcal{S}$ of atomic descriptions and some $\alpha_0 \in \mathcal{S}$ such that φ has as a disjunct $\alpha_0 \wedge \Box!\mathcal{S}$, and σ has as disjuncts each of the formulas $\alpha \wedge \Box!\mathcal{S}$ where $\alpha \in \mathcal{S}$. Since σ is satisfiable in some structure with only one world, it follows easily that there is a singleton set $\mathcal{B}$ consisting of a single atomic description β such that σ has as a disjunct $\beta \wedge \Box!\mathcal{B}$. Assume that β is $\psi_1 \wedge \ldots \wedge \psi_k$, where ψ_i is either p_i or $\neg p_i$, for $1 \leq i \leq k$. Let τ be the substitution which substitutes $p_i \vee \neg\Box!\mathcal{S}$ for p_i if ψ_i is p_i, and which substitutes $p_i \wedge \Box!\mathcal{S}$ for p_i if ψ_i is $\neg p_i$.

We now show that $\tau[\sigma]$ is valid and $\tau[\varphi]$ is satisfiable. This proves that φ is compatible with σ being valid.

Let $M = (W, R, \pi)$ be an arbitrary structure. To show that $\tau[\sigma]$ is valid, we must show that $M \models \tau[\sigma]$. It is well-known that without loss of generality, we can assume that R is universal. There are two cases, depending on whether $M \models \neg\Box!\mathcal{S}$ or $M \models \Box!\mathcal{S}$.

Case 1: $M \models \neg\Box!\mathcal{S}$. Now $\tau[\beta]$ is $\tau[\psi_1] \wedge \ldots \wedge \tau[\psi_k]$. If ψ_i is p_i, then $\tau[\psi_i]$ is $p_i \vee \neg\Box!\mathcal{S}$, and if ψ_i is $\neg p_i$, then $\tau[\psi_i]$ is equivalent to $\neg p_i \vee \neg\Box!\mathcal{S}$, for $1 \leq i \leq k$. Since $M \models \neg\Box!\mathcal{S}$, it follows easily that $M \models \tau[\beta]$. It is also the case that $M \models \tau[\Box!\mathcal{B}]$. Possibly the easiest way to see this is that $\Box!\mathcal{B}$ is equivalent to $\Box\beta$, and since $M \models \tau[\beta]$, also $M \models \tau[\Box\beta]$. We have shown that $M \models \tau[\beta \wedge \Box!\mathcal{B}]$. Since $\tau[\beta \wedge \Box!\mathcal{B}]$ is a disjunct of $\tau[\sigma]$, it follows that $M \models \tau[\sigma]$.

Case 2: $M \models \Box!\mathcal{S}$. Let α be an arbitrary atomic description $\gamma_1 \wedge \ldots \wedge \gamma_k$, where γ_i is either p_i or $\neg p_i$, for $1 \leq i \leq k$. It is easy to see that $\tau[\alpha]$ is $\gamma_1' \wedge \ldots \wedge \gamma_k'$, where γ_i' is $\gamma_i \vee \neg\Box!\mathcal{S}$. Since $M \models \Box!\mathcal{S}$, it follows that $(M, s) \models \tau[\alpha]$ iff $(M, s) \models \alpha$. From this and the fact that $M \models \Box!\mathcal{S}$, it is fairly straightforward to show that $M \models \tau[\Box!\mathcal{S}]$. Let s be a world of M, and let $\alpha \in \mathcal{S}$ be the atomic description such that $(M, s) \models \alpha$. From what we showed, it follows that $(M, s) \models \tau[\alpha] \wedge \tau[\Box!\mathcal{S}]$, and so $(M, s) \models \tau[\alpha \wedge \Box!\mathcal{S}]$. Since σ has $\alpha \wedge \Box!\mathcal{S}$ as a disjunct, it follows that $\tau[\sigma]$ has as a disjunct $\tau[\alpha \wedge \Box!\mathcal{S}]$. So $(M, s) \models \tau[\sigma]$. Since s was arbitrary, $M \models \tau[\sigma]$, as desired.

We have shown that $\tau[\sigma]$ is valid. We must also show that $\tau[\varphi]$ is satisfiable. Let M be a structure that satisfies $\Box!\mathcal{S}$, as in Case 2 above. Since $\alpha_0 \wedge \Box!\mathcal{S}$ is a disjunct of φ that is satisfiable in M, it follows from what we showed in the discussion of Case 2 that $\tau[\alpha_0 \wedge \Box!\mathcal{S}]$, and hence $\tau[\varphi]$, is satisfiable in M. ∎

Co-NP-completeness (as well as the missing part of the proof of Theorem 5.6) follow from Lemma 5.10. This is because the problem of deciding if $\varphi \wedge \Box\sigma$ is in NP, by Theorem 5.1, and the problem of deciding if σ is satisfiable in some structure with only one world is in NP, since we need only guess the structure and verify that it satisfies σ. ∎

6 Inference in first-order logic

We now discuss the situation for inference in first-order logic. We begin with some definitions. Our definitions are somewhat informal; see, for example, Enderton [End72] for a careful development.

We assume for convenience that there are no function or constant symbols in the language, but only predicate symbols. At the end of this section, we consider the situation when function and constant symbols are allowed. Recall that if φ is a first-order formula, with free variables $x_1, \ldots, x_r$, then the *universal closure* of φ, written $\varphi^\forall$, is the formula $\forall x_1 \cdots \forall x_r \varphi$.[9] When we say that the first-order formulas φ and ψ are *equivalent*, we mean that the formula $(\varphi \Leftrightarrow \psi)^\forall$ is valid.

We have to slightly modify the notion of substitution in the first-order case; a substitution τ is now a mapping that maps predicate symbol P of arity n to a sequence $\langle \psi_P, x_1, \ldots, x_n \rangle$, where ψ_P is a formula whose free variables are among $x_1, \ldots, x_n$. We take $\tau[\varphi]$ to be the result of replacing each occurrence of an atomic formula $P(t_1, \ldots, t_n)$

[9] To make the universal closure unique, we might assume that $x_1, \ldots, x_r$ are written in some fixed order, such as lexicographical order.

in φ by $\psi_P[x_1/t_1, \ldots, x_n/t_n]$, where the latter is the result of substituting t_i for the free variable x_i, for $1 \leq i \leq n$. We might loosely say that τ replaces $Px_1 \cdots x_n$ by $\psi_P(x_1, \ldots, x_n)$.

As we discussed earlier, in the case of first-order logic, for $\mathcal{M}$ inference we let $\mathcal{M}$ consist of all relational structures, and we consider pairs (M, w), where M is a relational structure and w is a valuation. Then $M \models \varphi$ iff $(M, w) \models \varphi$ for every valuation w (and thus for every assignment to the free variables of φ) iff $M \models \varphi^\forall$, where as before, $\varphi^\forall$ is the universal closure of φ.

The next proposition is analogous to Proposition 5.4. It reduces $\mathcal{M}$ inference to validity of a certain formula. Unlike the case of Proposition 5.4, this formula is within our language.

Proposition 6.1: *$\sigma \vdash_\mathcal{M} \varphi$ holds in first-order logic iff $\sigma^\forall \Rightarrow \varphi^\forall$ is valid.*

Proof: By definition, $\sigma \vdash_\mathcal{M} \varphi$ holds iff $M \models \tau[\sigma]$ implies $M \models \tau[\varphi]$ for every structure M and every substitution τ. Now if ψ is σ or φ, we have $M \models \tau[\psi]$ iff (by our comments just before this proposition) $M \models (\tau[\psi])^\forall$ iff $M \models \tau[\psi^\forall]$ (since $(\tau[\psi])^\forall$ equals $\tau[\psi^\forall]$). So $\sigma \vdash_\mathcal{M} \varphi$ holds iff $M \models \tau[\sigma^\forall]$ implies $M \models \tau[\varphi^\forall]$ for every structure M and every substitution τ, that is, $M \models \tau[\sigma^\forall] \Rightarrow \tau[\varphi^\forall]$ for every structure M and every substitution τ. This holds iff $\tau[\sigma^\forall] \Rightarrow \tau[\varphi^\forall]$ is valid for every substitution τ, which holds iff $\tau[\sigma^\forall \Rightarrow \varphi^\forall]$ is valid for every substitution τ, which holds iff $\sigma^\forall \Rightarrow \varphi^\forall$ is valid (since it is not hard to see that a formula ψ is valid in first-order logic iff $\tau[\psi]$ is valid for every substitution τ). ∎

Proposition 6.2: *For first-order logic, $\vdash_t \subset \vdash_\mathcal{M} \subset \vdash_v$.*

Proof: We already know that $\vdash_t \subseteq \vdash_\mathcal{M} \subseteq \vdash_v$. To see that $\vdash_t \neq \vdash_\mathcal{M}$, note that $\varphi \vdash_\mathcal{M} \varphi^\forall$ (universal generalization) is a sound $\mathcal{M}$ inference rule, but $\varphi \vdash_t \varphi^\forall$ is not a sound truth inference rule (for example, it is easy to see that a counterexample occurs when φ is Px, where P is a unary relation symbol).

To see that $\vdash_\mathcal{M} \neq \vdash_v$, note that $(\forall x \forall y(x = y)) \vdash_v \mathit{false}$ is a sound validity inference rule, since $\tau[\forall x \forall y(x = y)]$ is simply $\forall x \forall y(x = y)$, which is not valid. However, it is false that $(\forall x \forall y(x = y)) \vdash_\mathcal{M} \mathit{false}$ is a sound $\mathcal{M}$ inference rule, since $\forall x \forall y(x = y)$ can certainly hold for some structure. ∎

Since the valid formulas are r.e. (recursively enumerable) in first-order logic (this is Church's Theorem [End72]), we know that truth inference is r.e. By Proposition 6.1, it follows that $\mathcal{M}$ inference is also r.e. for first-order logic. In fact, truth inference and $\mathcal{M}$ inference are both r.e.-hard, since the special case of deciding if $\mathit{true} \vdash_t \varphi$ (resp. $\mathit{true} \vdash_\mathcal{M} \varphi$), that is, if φ is valid, is r.e.-hard. We do not know the precise complexity of validity inference. By the same argument as for truth inference and $\mathcal{M}$ inference, we see that validity inference is r.e.-hard. Furthermore, it is straightforward to show[10] that validity inference is in Π^0_2. (This follows from the fact that $\sigma \vdash_v \varphi$ iff $\forall\tau((\exists \text{ proof of } \tau[\sigma]) \Rightarrow (\exists \text{ proof of } \tau[\varphi]))$. The claim then follows by arithmetization.) It would be interesting to close this gap.

Let us now consider first-order logic interpreted over *finite structures* (that is, structures where the universe is finite; see [Fag93] for a survey of finite-model theory). We first see that we have the following analog to Proposition 6.2.

Proposition 6.3: *For first-order logic over finite structures, $\vdash_t \subset \vdash_{\mathcal{M}} \subset \vdash_v$.*

Proof: This has exactly the same proof as Proposition 6.2. ∎

What about the complexity of deciding whether an inference rule is sound?

Theorem 6.4: *The problem of deciding whether $\sigma \vdash_t \varphi$ (resp. $\sigma \vdash_{\mathcal{M}} \varphi$, $\sigma \vdash_v \varphi$) holds for first-order logic over finite structures is co-r.e.-complete.*

Proof: In showing this, the only real difficulty lies in proving that validity inference is co-r.e. To show that validity inference over finite structures is co-r.e., we shall prove a lemma analogous to Lemmas 4.7 and 5.10. As before, we say that φ is *compatible with σ being valid over finite structures* if there is a substitution τ such that $\tau[\sigma]$ is valid over finite structures and $\tau[\varphi]$ is satisfiable in some finite structure.

We need a few more definitions. Let $(a_1, \ldots, a_k)$ and $(b_1, \ldots, b_k)$ be tuples of the same arity k. We say that these tuples *have the same equality pattern* if for each i and j with $1 \leq i \leq k$ and $1 \leq j \leq k$, we have $a_i = a_j$ iff $b_i = b_j$. Let $\mathcal{C} = \langle U; R_1, \ldots, R_m \rangle$ be a finite structure, where U is the finite universe and where the R_i's are relations over the universe. We say that $\mathcal{C}$ is *uniform* if for each of the relations R_i, whenever $(a_1, \ldots, a_k)$ and $(b_1, \ldots, b_k)$ are tuples of members of U that have the same equality pattern, then $(a_1, \ldots, a_k) \in R_i$ iff $(b_1, \ldots, b_k) \in R_i$. It is straightforward to verify that a structure is uniform precisely if every permutation of the universe is an automorphism of the structure.

We shall make use of the following simple lemma, which is analogous to Lemmas 4.5 and 5.8, and whose proof is left to the reader.

Lemma 6.5: *Let ψ be a first-order sentence and $\mathcal{C}$ a structure. Assume that the substitution τ replaces the atomic formula $Px_1 \cdots x_k$ by $\psi_P(x_1, \ldots, x_k)$, for each predicate symbol P. Let $\mathcal{C}'$ be a structure with the same universe as $\mathcal{C}$, but where the interpretation $P^{\mathcal{C}'}$ of P in $\mathcal{C}'$ is the interpretation of ψ_P in $\mathcal{C}$. If $\mathcal{C}$ satisfies $\tau[\psi]$, then $\mathcal{C}'$ satisfies ψ. Furthermore, if $\mathcal{C}$ is uniform, then so is $\mathcal{C}'$.*

Corollary 6.6: *Let ψ be a first-order sentence, and let τ be a substitution. If $\tau[\psi]$ is satisfiable in a finite structure, then so is ψ. Furthermore, if $\tau[\psi]$ is satisfiable in a uniform finite structure, then so is ψ.*

[10] We thank Martín Abadi for pointing this out to us.

To show that validity inference over finite structures is co-r.e., we use a similar approach to that we used for NPL and S5, but where "r.e." is used in place of "NP". We shall show that the compatibility problem, that is, the problem of deciding whether φ is compatible with σ being valid, is r.e. It then follows that the problem of deciding validity inference is co-r.e. Analogously to before, we show that φ is compatible with σ being valid precisely if (a) there is some structure where σ is valid and where φ is satisfiable, and (b) σ is valid in some set of "special" structures. Here, "special" means uniform and with the size of the universe being comparable to the size of σ. Thus, we show the following lemma, analogous to Lemmas 4.7 and 5.10.

Lemma 6.7: *Let φ and σ be first-order formulas. Then φ is compatible with σ being valid over finite structures iff (a) $\sigma^\forall \wedge \varphi$ is satisfiable in some finite structure, and (b) if r is the number of quantifiers that appear in $\sigma^\forall$, then $\sigma^\forall$ is satisfiable in some uniform structure with universe of cardinality m, for each m with $1 \leq m \leq r$.*

Proof: Assume first that φ is compatible with σ being valid over finite structures. Then there is a substitution τ such that $\tau[\sigma]$ is valid over finite structures and $\tau[\varphi]$ is satisfiable in some finite structure. Let $\mathcal{C}$ be a finite structure that satisfies $\tau[\varphi]$. Then $\mathcal{C}$ satisfies $\tau[\sigma^\forall \wedge \varphi]$, since $\tau[\sigma^\forall]$ is valid over finite structures. Since $\tau[\sigma^\forall \wedge \varphi]$ is satisfiable in some finite structure, it follows by Corollary 6.6 that $\sigma^\forall \wedge \varphi$ is satisfiable in some finite structure.

For each m, let $\mathcal{C}_m$ be an arbitrary uniform structure with universe of cardinality m. Since $\tau[\sigma^\forall]$ is valid over finite structures, it follows that $\mathcal{C}_m$ satisfies $\tau[\sigma^\forall]$. Let $\mathcal{C}'_m$ be obtained from $\mathcal{C}_m$ as $\mathcal{C}'$ is obtained from $\mathcal{C}$ in Lemma 6.5. Since $\mathcal{C}_m$ is uniform, it follows from Lemma 6.5 that $\mathcal{C}'_m$ is uniform. By Lemma 6.5, $\mathcal{C}'_m$ satisfies $\sigma^\forall$. Since the cardinality of the universe of $\mathcal{C}'_m$ is the same as that of $\mathcal{C}_m$, namely, m, it follows that $\sigma^\forall$ is satisfiable in a uniform structure with universe of cardinality m. But m was arbitrary. This proves the "only if" direction.

We now prove the "if" direction. Assume that (a) and (b) in the statement of the lemma hold. Let $\mathcal{A}$ be a finite structure that satisfies $\sigma^\forall \wedge \varphi$. It is well-known (see [Fag93] for a discussion) that there is a first-order sentence $\zeta_\mathcal{A}$ that characterizes $\mathcal{A}$ up to isomorphism (that is, an arbitrary structure $\mathcal{B}$ over the same language is isomorphic to $\mathcal{A}$ iff $\mathcal{B} \models \zeta_\mathcal{A}$). It is also easy to see that for each positive integer i, there is a first-order sentence (which we denote by $\chi_{\geq i}$) that says "There are at least i points." For example, we can take $\chi_{\geq 3}$ to be

$$\exists x_1 \exists x_2 \exists x_3 ((x_1 \neq x_2) \wedge (x_1 \neq x_3) \wedge (x_2 \neq x_3)).$$

Let us denote the sentence $\chi_{\geq i} \wedge \neg \chi_{\geq i+1}$ by $\chi_{=i}$. Thus, $\chi_{=i}$ says "There are exactly i points."

Let $\mathcal{C}_i$ be a uniform structure with universe of cardinality i that satisfies $\sigma^\forall$, for $1 \leq i \leq r$. For each predicate symbol P and each i with $1 \leq i \leq r$, if P is k-ary, then let $\lambda_{P,i}$ be a formula that describes the equality pattern among tuples $(x_1, \ldots, x_k)$ such that $(x_1, \ldots, x_k)$ is in $P^{\mathcal{C}_i}$ (the interpretation of P in $\mathcal{C}_i$). For example, if P is 3-ary, and if a tuple (a_1, a_2, a_3) of members of the universe of $\mathcal{C}_i$ is in $P^{\mathcal{C}_i}$ precisely if either a_1, a_2, a_3 are all distinct or else $a_1 = a_2$ and $a_1 \neq a_3$, then we would take $\lambda_{P,i}$ to be the formula

$$((x_1 \neq x_2) \wedge (x_1 \neq x_3) \wedge (x_2 \neq x_3)) \vee ((x_1 = x_2) \wedge (x_1 \neq x_3)).$$

For each i with $i > r$, define $\mathcal{C}_i$ to be a uniform structure with universe of cardinality i, where for each predicate P, we define $P^{\mathcal{C}_i}$ by letting the equality pattern of the tuples in the universe of $\mathcal{C}_i$ that satisfy $P^{\mathcal{C}_i}$ be precisely the same as the equality pattern of the tuples in the universe of $\mathcal{C}_r$ that satisfy $P^{\mathcal{C}_r}$. It follows easily by an Ehrenfeucht-Fraïssé game argument [Ehr61, Fra54] that each $\mathcal{C}_i$ with $i > r$ agrees with $\mathcal{C}_r$ on each sentence with at most r quantifiers. That is, each such sentence is true in $\mathcal{C}_i$ with $i > r$ precisely if it is true in $\mathcal{C}_r$. In particular, since $\mathcal{C}_r$ satisfies $\sigma^\forall$, so does each $\mathcal{C}_i$ with $i > r$.

We are now ready to define a substitution τ which shows that φ is compatible with σ being valid over finite structures (that is, we will show that $\tau[\sigma]$ is valid over finite structures and $\tau[\varphi]$ is satisfiable in some finite structure). The substitution τ replaces the atomic formula $Px_1 \ldots x_k$ by the conjunction of the following formulas:

$\zeta_{\mathcal{A}} \Rightarrow Px_1 \ldots x_k$

$(\neg\zeta_{\mathcal{A}} \wedge \chi_{=1}) \Rightarrow \lambda_{P,1}(x_1, \ldots, x_k)$

$\ldots$

$(\neg\zeta_{\mathcal{A}} \wedge \chi_{=(r-1)}) \Rightarrow \lambda_{P,(r-1)}(x_1, \ldots, x_k)$

$(\neg\zeta_{\mathcal{A}} \wedge \chi_{\geq r}) \Rightarrow \lambda_{P,r}(x_1, \ldots, x_k)$

Since $\mathcal{A}$ satisfies φ, it is easy to see that $\mathcal{A}$ satisfies $\tau[\varphi]$. So $\tau[\varphi]$ is satisfiable in some finite structure. We now show that $\tau[\sigma]$ is valid over finite structures. Since $\mathcal{A}$ satisfies $\sigma^\forall$, it is easy to see that $\mathcal{A}$ satisfies $\tau[\sigma^\forall]$. If $\mathcal{B}$ is a finite structure not isomorphic to $\mathcal{A}$ with universe of cardinality i, where i is arbitrary, then from the fact that $\mathcal{C}_i$ satisfies $\sigma^\forall$, it is straightforward to verify that $\mathcal{B}$ satisfies $\tau[\sigma^\forall]$. Therefore, every finite structure satisfies $\tau[\sigma^\forall]$, so $\tau[\sigma]$ is valid over finite structures, as desired. ∎

To show that validity inference over finite structures is co-r.e., we need only show that (a) and (b) of Lemma 6.7 are r.e. As for (a), it is well-known that deciding if a formula is satisfiable in some finite structure is r.e. This is because to find out whether ψ is satisfiable, it is possible to consider systematically every finite structure $\mathcal{A}$ over the language of ψ to see whether ψ is satisfiable in $\mathcal{A}$. This makes it possible to list all the formulas that are satisfiable. Finally, it is easy to see that the condition (b) is not only r.e., but even recursive. ∎

Throughout this section, we have assumed that there are no function or constant symbols in the language, but only predicate symbols. While this is typically an innocuous assumption (since a function of arity k can be encoded by a predicate of arity $k+1$), here the assumption seems to make a difference, because of the effects of substitutions. If we have function symbols in the language, then we can define a substitution τ to be a mapping that not only maps each atomic predicate P of arity n to a sequence $\langle \psi_P, x_1, \ldots, x_n \rangle$, where ψ_P is a formula whose free variables are among $x_1, \ldots, x_n$ (as before), but also maps each function symbol f of arity n to a sequence $\langle s_f, x_1, \ldots, x_n \rangle$, where s_f is a term whose free variables are among $x_1, \ldots, x_n$ (we treat a constant symbol as a function symbol of arity 0). Given a substitution τ, we first define $\tau[t]$ for an arbitrary term t inductively: $\tau[x] = x$ for a variable x and if $\tau(f) = \langle s_f, x_1, \ldots, x_n \rangle$, then $\tau[f(t_1, \ldots, t_n)] = s_f(x_1/\tau[t_1], \ldots, x_n/\tau[t_n])$. We now take $\tau[\varphi]$ to be the result of replacing each occurrence of an atomic formula $P(t_1, \ldots, t_n)$ by $\psi_P[x_1/\tau[t_1], \ldots, x_n/\tau[t_n]]$.

We remark that in the presence of function and constant symbols, we do not know whether validity inference over finite structures is still co-r.e. The proof of Lemma 6.7 breaks down, because in the presence of function symbols, we can no longer define "conditional" substitutions ("if condition 1 holds, then do this substitution, but if condition 2 holds, then do that substitution") as we did in defining the substitution τ in the proof of Lemma 6.7.

The main reason why we are interested in considering the effect of allowing function and constant symbols is that they arise in Skolemization. We close this section by considering how Skolemization fits into our framework. See Enderton [End72] for a more detailed discussion of Skolemization. A first-order formula is in *prenex normal form* if it is of the form $Q_1 x_1 \ldots Q_r x_r \gamma$, where each Q_i is a quantifier (either $\forall$ or $\exists$), each x_i is a variable, and γ is quantifier-free. Thus, a formula is in prenex normal form if "all quantifiers are in front". There is a well-known mechanical procedure for converting a formula into an equivalent formula in prenex normal form. A formula is *existential* if it is in prenex normal form, and all of the quantifiers are existential ($\exists$). *Skolemization* is the process of converting a prenex normal form formula φ into an existential formula (in an expanded language, that contains additional constant and function symbols) that is valid iff φ is valid. For example, let φ be the formula $\exists x \exists y \forall z \gamma(z)$, where $\gamma(z)$ is a quantifier-free formula whose free variables include z. The Skolemization is the formula $\exists x \exists y \gamma(f(x, y))$, where f is a new binary function symbol. As another example, let φ be the formula $\forall y_1 \exists x_1 \exists x_2 \forall y_2 \exists x_3 \forall y_3 \gamma(y_1, y_2, y_3)$, where $\gamma(y_1, y_2, y_3)$ is a quantifier-free formula whose free variables include y_1, y_2, y_3. The Skolemization is the formula $\exists x_1 \exists x_2 \exists x_3 \gamma(c, f(x_1, x_2), g(x_1, x_2, x_3))$, where c is a new constant symbol, f is a new binary function symbol, and g is a new ternary function symbol. Intuitively, the Skolemization of a prenex formula is obtained by replacing each universally quantified variable by a function of the previous existentially quantified variables.

It is straightforward to verify that the Skolemization is an existential formula that is valid iff the original formula is valid (this is perhaps easiest to see by taking negations, and showing that the negation of the Skolemization is satisfiable iff the negation of the original formula is satisfiable). Let us denote the Skolemization of a formula φ by φ^S. Thus, φ is valid iff φ^S is valid.

It is easy to see that $\varphi \vdash_t \varphi^S$ is a sound truth inference rule, so $\varphi \vdash_v \varphi^S$ is a sound validity inference rule (since $\vdash_t \subseteq \vdash_v$). However, $\varphi^S \vdash_v \varphi$ is not a sound validity inference rule (and all the more so, $\varphi^S \vdash_t \varphi$ is not a sound truth inference rule). The problem is the effect of substitutions. For example, let φ be $\forall x Px$, so that φ^S is Pc for a constant symbol c. Let τ be the substitution that replaces Px by $x = c$. Then $\tau[\varphi^S]$ is simply $c = c$, which is valid, but $\tau[\varphi]$ is $\forall x(x = c)$, which is not valid. This shows that $\varphi^S \vdash_v \varphi$ is not a sound validity inference rule. However, if we were to restrict the substitutions τ by not allowing them to make use of the constant and function symbols that arise by Skolemization, then $\varphi^S \vdash_v \varphi$ would become sound.

What does our Skolemization example tell us? Throughout this paper, we have chosen for simplicity to allow arbitrary substitutions. However, as we just saw, if we want the Skolemization rule $\varphi^S \vdash_v \varphi$ to be sound, then we must restrict our substitutions by side conditions. This suggests that there are some subtleties that must be dealt with when considering substitutions in first-order logic. Since the issues involved are somewhat orthogonal to the main thrust of this paper, we have chosen to ignore them here.

7 Other Notions of Inference

As we mentioned in the introduction, there has been work in the literature on two other notions of inference: axiomatic inference and nonschematic inference.

Consider first nonschematic $\mathcal{M}$ inference. We say that $\sigma \vdash^{ns}_{\mathcal{M}} \varphi$ holds iff $M \models \sigma$ implies $M \models \varphi$ for all $M \in \mathcal{M}$. Thus, we do not consider substitutions. A sound nonschematic validity inference rule can be viewed as corresponding to a specific instance of sound reasoning, rather than a general pattern. As we discussed earlier, a powerful collection of nonschematic rules may allow us to shorten proofs. It is easy to see that $\vdash_{\mathcal{M}} \subseteq \vdash^{ns}_{\mathcal{M}}$. However, $\vdash_{\mathcal{M}}$ and $\vdash^{ns}_{\mathcal{M}}$ may be distinct: as an example, we already saw that in propositional logic, $p \vdash^{ns}_v \mathit{false}$ holds while $p \vdash_v \mathit{false}$ does not. There are some cases where $\vdash_{\mathcal{M}}$ and $\vdash^{ns}_{\mathcal{M}}$ coincide. For example, in propositional logic, as well as all the other logics we consider, $\vdash_t$ and $\vdash^{ns}_t$ are identical; we have $\sigma \vdash_t \varphi$ iff $\sigma \vdash^{ns}_t \varphi$. As another example, it is easy to see that $\vdash_s$ and $\vdash^{ns}_s$ are the same. Just as with schematic inference rules, we have $\vdash^{ns}_t \subseteq \vdash^{ns}_{\mathcal{M}} \subseteq \vdash^{ns}_v$; in modal logic we also have $\vdash^{ns}_s \subseteq \vdash^{ns}_f$. In propositional logic, we saw (Proposition 3.1) that $\vdash_v = \vdash_t$. However, $\vdash^{ns}_v \neq \vdash^{ns}_t$, since, we just saw, $\vdash_v \neq \vdash^{ns}_v$, whereas $\vdash_t = \vdash^{ns}_t$.

We now turn our attention to axiomatic $\mathcal{M}$ inference. We say that $\sigma \vdash^{ax}_{\mathcal{M}} \varphi$ holds iff for all $M \in \mathcal{M}$, whenever we have that $M \models \tau[\sigma]$ for all substitutions τ, we also have that $M \models \tau[\varphi]$ for all substitutions τ. Thus, in axiomatic inference, we change the order of quantification. Intuitively, we are restricting our attention to structures where (every substitution instance of) σ holds, and asking whether (every substitution instance of) φ holds. This corresponds to asking whether taking σ as an axiom implies that φ is a theorem. For example, in the modal logic K, we have

$$(\neg\Box p \Rightarrow \Box\neg\Box p) \wedge (\Box p \Rightarrow p) \vdash^{ax}_f (\Box p \Rightarrow \Box\Box p).$$

Intuitively, taking $\neg\Box p \Rightarrow \Box\neg\Box p$ and $\Box p \Rightarrow p$ as axioms implies that $\Box p \Rightarrow \Box\Box p$ is also an axiom.

It is easy to see that $\vdash_{\mathcal{M}} \subseteq \vdash_{\mathcal{M}}^{ax}$. However, $\vdash_{\mathcal{M}}$ and $\vdash_{\mathcal{M}}^{ax}$ may be distinct: as an example that we noted earlier, in propositional logic, if p is a primitive proposition, then $p \vdash_v^{ax} false$ holds, whereas $p \vdash_v false$ does not. There are some cases where $\vdash_{\mathcal{M}}$ and $\vdash_{\mathcal{M}}^{ax}$ coincide. For example, in propositional logic, as well as all the other logics we consider, $\vdash_t$ and $\vdash_t^{ax}$ are identical. Again we have $\vdash_t^{ax} \subseteq \vdash_{\mathcal{M}}^{ax} \subseteq \vdash_v^{ax}$. It is easy to show that in the case of propositional logic we have $\vdash_t^{ax} = \vdash_v^{ax}$, just as we had $\vdash_t = \vdash_v$. Similarly to before, in modal logic we have $\vdash_s^{ax} \subseteq \vdash_f^{ax}$. A result of van Benthem [Ben79] (see also [HC84, page 57]) implies that $\vdash_s^{ax}$ and $\vdash_f^{ax}$ are distinct in the modal logic K.

There is an interesting connection between axiomatic and nonschematic inference. Recall that Frege's inference rule is (for an arbitrary substitution τ) the rule $\varphi \vdash_{\mathcal{M}}^{ns} \tau[\varphi]$. Axiomatic $\mathcal{M}$ inference is equivalent to nonschematic $\mathcal{M}$ inference as long as Frege's inference rule is sound. That is:

Proposition 7.1: *Assume $\varphi \vdash_{\mathcal{M}}^{ns} \tau[\varphi]$ for every formula φ and every substitution τ. Then axiomatic $\mathcal{M}$ inference is equivalent to nonschematic $\mathcal{M}$ inference.*

Proof: We first show that $\sigma \vdash_{\mathcal{M}}^{ax} \varphi$ implies $\sigma \vdash_{\mathcal{M}}^{ns} \varphi$. Assume that $\sigma \vdash_{\mathcal{M}}^{ax} \varphi$, and $M \models \sigma$ for some $M \in \mathcal{M}$. By Frege's inference rule, $M \models \tau[\sigma]$ for every substitution τ. So, since $\sigma \vdash_{\mathcal{M}}^{ax} \varphi$, it follows that $M \models \tau[\varphi]$ for every substitution τ. In particular, $M \models \varphi$. So $\sigma \vdash_{\mathcal{M}}^{ns} \varphi$.

We now show that $\sigma \vdash_{\mathcal{M}}^{ns} \varphi$ implies $\sigma \vdash_{\mathcal{M}}^{ax} \varphi$. Assume that $\sigma \vdash_{\mathcal{M}}^{ns} \varphi$, and $M \models \tau[\sigma]$ for every substitution τ. In particular, $M \models \sigma$. So, since $\sigma \vdash_{\mathcal{M}}^{ns} \varphi$, it follows that $M \models \varphi$. By Frege's inference rule, $M \models \tau[\varphi]$ for every substitution τ. So $\sigma \vdash_{\mathcal{M}}^{ax} \varphi$. ∎

There are several interesting cases where Frege's inference rule holds for validity (that is, $\varphi \vdash_v^{ns} \tau[\varphi]$ is sound, for every substitution τ), and hence, Proposition 7.1 is applicable. For each of the logics considered in this paper (propositional logic, NPL, modal logic, and first-order logic), we have $\varphi \vdash_v^{ns} \tau[\varphi]$, for every substitution τ. For these logics, we therefore obtain from Proposition 7.1 the following corollary:

Corollary 7.2: *Axiomatic validity inference is equivalent to nonschematic validity inference.*

Furthermore, for modal logic, Frege's inference rule is sound when we consider frame inference. Hence:

Corollary 7.3: *Axiomatic frame inference is equivalent to nonschematic frame inference.*

For modal logic, Frege's inference rule is *not* sound when we consider structure inference. As an example, let p be a primitive proposition, and let τ be a substitution that replaces p by $false$; it is easy to see that we do not have $p \vdash_s^{ns} \tau[p]$. Furthermore, the same counterexample shows that the conclusion of Proposition 7.1 fails in the case of structure inference, since $p \vdash_s^{ax} false$ holds, but $p \vdash_s^{ns} false$ does not.

We briefly comment on complexity issues. In Proposition 3.2, we saw that in propositional logic, the problem of deciding whether $\sigma \vdash_v \varphi$ (or equivalently $\sigma \vdash_t \varphi$) holds is co-NP-complete. Since $\vdash_t^{ns} = \vdash_t$, it of course follows that the problem of deciding whether $\sigma \vdash_t^{ns} \varphi$ holds is co-NP-complete. However, the situation is different for validity inference. We can encode both satisfiability and validity using $\vdash_v^{ns}$: φ is valid iff $true \vdash_v^{ns} \varphi$, and φ is satisfiable iff $\neg\varphi \vdash_v^{ns} false$. Thus, the problem of deciding if $\sigma \vdash_v^{ns} \varphi$ is both NP-hard and co-NP-hard. In fact, it is easily seen to be complete for co-D^P, the set of all problems that can be expressed as the union of an NP and co-NP problem [PY82] (the problem of deciding if $\sigma \vdash_v^{ns} \varphi$ is in co-D^P, since $\sigma \vdash_v^{ns} \varphi$ holds iff either σ is not valid (an NP problem) or φ is valid (a co-NP problem)). If NP $\neq$ co-NP (which is widely believed to be the case), then co-D^P is a complexity class that is "higher" than NP and co-NP, and in this sense, deciding if $\sigma \vdash_v^{ns} \varphi$ is "harder" than deciding if $\sigma \vdash_v \varphi$. This is what we meant in Section 4 when we said that "being able to substitute gives us some control over the problem, and in fact might make it easier".

The situation for NPL and S5 is somewhat similar to that for propositional logic, but for the logic K, the situation is drastically different. It is not hard to show that the nonschematic versions of all four types of inference we considered for K are distinct. As expected, the PSPACE-completeness result for truth inference carries over to the nonschematic case, as does the EXPTIME-completeness result for structure inference, since we can reduce structure inference to validity. In the case of validity inference, it is not hard to show using Theorem 5.1 that the problem is PSPACE-complete. However, it follows from results of Thomason [Tho75a, Tho75b] that the second-order theory of a binary relation can be reduced to nonschematic frame inference. This tells us that the complexity of nonschematic frame inference is at least as high as that of full type theory. This means that the complexity is higher than any level of the arithmetic or analytical hierarchy [Rog67]. Further, because of the the equivalence between axiomatic frame inference and nonschematic frame inference (Corollary 7.3), axiomatic frame inference has this same enormously high complexity.

We also note that Meyer, Streett, and Mirkowska [MSM81] considered a notion of inference that combines axiomatic and nonschematic inference in the context of PDL. They showed that their notion is undecidable (in fact, Π_1^1-complete, where Π_1^1 is the first level of the analytical hierarchy).

8 Conclusions

The notion of inference is one that seems to have been taken largely for granted by logicians. Issues such as the precise definition and complexity have not received the attention that it seems to us they deserve. We have tried to raise a number of issues regarding inference that we believe to be important. We hope that this paper provides the impetus for further study of these issues.

Acknowledgments: We are grateful to Martín Abadi, Arnon Avron, Jon Barwise, Johan van Benthem, Steve Cook, Solomon Feferman, Kit Fine, Haim Gaifman, Paris Kanellakis, Daniel Leivant, John Mitchell, Gordon Plotkin, Steve Thomason, and Alasdair Urquhart, who helped us in various ways, including useful discussions on the subject of inference, helpful comments, and assistance in tracking down references.

References

[AB75] A. Anderson and N. D. Belnap. *Entailment: The Logic of Relevance and Necessity.* Princeton University Press, Princeton, N.J., 1975.

[Avr91] A. Avron. Simple consequence relations. *Information and Computation*, 92(1):105–139, 1991.

[Ben79] J. F. A. K. van Benthem. Syntactic aspects of modal incompleteness theorems. *Theoria*, 45:63–77, 1979.

[Che80] B. F. Chellas. *Modal Logic.* Cambridge University Press, Cambridge, U.K., 1980.

[Chu56] A. Church. *Introduction to Mathematical Logic*, volume I. Princeton University Press, Princeton, N.J., 1956.

[Coo71] S. A. Cook. The complexity of theorem proving procedures. In *Proc. 3rd ACM Symp. on Theory of Computing*, pages 151–158, 1971.

[CR79] S. A. Cook and R. A. Reckhow. The relative efficiency of propositional proof systems. *Journal of Symbolic Logic*, 44(1):36–50, March 1979.

[Dun86] J. M. Dunn. Relevance logic and entailment. In D. Gabbay and F. Guenthner, editors, *Handbook of Philosophical Logic, Vol. III*, pages 117–224. Reidel, Dordrecht, Netherlands, 1986.

[Ehr61] A. Ehrenfeucht. An application of games to the completeness problem for formalized theories. *Fund. Math.*, 49:129–141, 1961.

[End72] H. B. Enderton. *A Mathematical Introduction to Logic.* Academic Press, New York, 1972.

[Fag93] R. Fagin. Finite-model theory—a personal perspective. *Theoretical Computer Science*, 116:3–31, 1993.

[FHV90] R. Fagin, J. Y. Halpern, and M. Y. Vardi. A nonstandard approach to the logical omniscience problem. In R. Parikh, editor, *Theoretical Aspects of Reasoning about Knowledge: Proc. Third Conference*, pages 41–55. Morgan Kaufmann, San Francisco, Calif., 1990. To appear in *Artificial Intelligence*.

[FL79] M. J. Fischer and R. E. Ladner. Propositional dynamic logic of regular programs. *Journal of Computer and System Sciences*, 18(2):194–211, 1979.

[Fra54] R. Fraïssé. Sur quelques classifications des systèmes de relations. *Publ. Sci. Univ. Alger. Sér. A*, 1:35–182, 1954.

[Fre79] G. Frege. *Begriffsschrift, eine der Arithmetischen Nachgebildete Formelsprache des Reinen Denkens*. Halle, 1879. English translation in *From Frege to Gödel, a Source Book in Mathematical Logic* (J. van Heijenoort, Editor), Harvard University Press, Cambridge, 1967, pp. 1–82.

[Fri75] H. Friedman. One hundred and two problems in mathematical logic. *Journal of Symbolic Logic*, 40(3):27–32, 1975.

[Gab76] D. Gabbay. *Investigations in Modal and Tense Logics*. Reidel, Dordrecht, Netherlands, 1976.

[Gab81] D. Gabbay. *Semantical Investigations in Heyting's Intuitionistic Logic*. Reidel, Dordrecht, Netherlands, 1981.

[GJ79] 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.

[GP89] V. Goranko and S. Passy. Using the universal modality: profits and questions. Manuscript, 1989.

[Hac79] I. Hacking. What is logic? *Journal of Philosophy*, 76(6):285–319, 1979.

[Hal87] J. Y. Halpern. Using reasoning about knowledge to analyze distributed systems. In J. F. Traub, B. J. Grosz, B. W. Lampson, and N. J. Nilsson, editors, *Annual Review of Computer Science, Vol. 2*, pages 37–68. Annual Reviews Inc., Palo Alto, Calif., 1987.

[HC68] G. E. Hughes and M. J. Cresswell. *An Introduction to Modal Logic*. Methuen, London, 1968.

[HC84] G. E. Hughes and M. J. Cresswell. *A Companion to Modal Logic*. Methuen, London, 1984.

[HU79] J. E. Hopcroft and J. D. Ullman. *Introduction to Automata Theory, Languages and Computation*. Addison-Wesley, New York, 1979.

[Hum86] I. L. Humberstone. A more discriminating approach to modal logic. *Journal of Symbolic Logic*, 51(2):503–504, 1986. (Abstract only.) There is also an expanded, but unpublished, manuscript.

[Kap87] B. Kapron. Modal sequents and definability. *Journal of Symbolic Logic*, 52(3):756–762, 1987.

[Lad77] R. E. Ladner. The computational complexity of provability in systems of modal propositional logic. *SIAM Journal on Computing*, 6(3):467–480, 1977.

[LL59] C. I. Lewis and C. H. Langford. *Symbolic Logic*. Dover, New York, 2nd edition, 1959.

[Min76] G. Mints. The derivability of admissible rules. *Journal of Soviet Mathematics*, 6:417–421, 1976.

[MSM81] A. R. Meyer, R. S. Streett, and G. Mirkowska. The deducibility problem in propositional dynamic logic. In E. Engeler, editor, *Logic of Programs, Proceedings, 1979*, Lecture Notes in Computer Science, Vol. 125. Springer-Verlag, Berlin/New York, 1981.

[Neu27] J. von Neumann. Zur Hilbertschen beweistheorie. *Mathematische Zeitschrift*, 26:1–46, 1927.

[Pra79] V. R. Pratt. Models of program logics. In *Proc. 20th IEEE Symp. on Foundations of Computer Science*, pages 115–122, 1979.

[PY82] C. H. Papadimitriou and M. Yannakakis. The complexity of facets (and some facets of complexity). *Journal of Computer and System Sciences*, 28(2):244–259, 1982.

[Qui50] W. V. O. Quine. *Methods of Logic*. Holt, New York, 1950.

[Rog67] H. Rogers, Jr. *Theory of Recursive Functions and Effective Computability*. McGraw-Hill, New York, 1967.

[RR72] R. Routley and V. Routley. Semantics of first degree entailment. *Noûs*, 6:335–359, 1972.

[Ryb89] V. V. Rybakov. Problems of admissibility and substitution, logical equations and restricted theories of free algebras. In J. E. Fenstad, I. T. Timofeevich, and R. Hilpinen, editors, *Logic, Methodology and Philosophy of Science VIII*, pages 121–139. Elsevier Science Publishers, Amsterdam, 1989.

[Sav70] W. J. Savitch. Relationships between nondeterministic and deterministic tape complexities. *Journal of Computer and System Sciences*, 4:177–192, 1970.

[Sto87] L. J. Stockmeyer. Classifying the computational complexity of problems. *Journal of Symbolic Logic*, 52:1–43, 1987.

[Tho75a] S. Thomason. The logic consequence relation of propositional tense logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:29–40, 1975.

[Tho75b] S. Thomason. Reduction of second-order logic to modal logic. *Zeitschrift für Mathematische Logik und Grundlagen der Mathematik*, 21:107–114, 1975.

[Tsi77] A. I. Tsitkin. On admissible rules of intuitionistic propositional calculus. *Matematicheskii Sbornik*, 102(2):314–323, 1977. An English translation appears in *Soviet Math. Sbornic*.