

Classical Logic

<http://plato.stanford.edu/entries/logic-classical/>

Typically, a logic consists of a formal or informal language together with a deductive system and/or a model-theoretic semantics. The language is, or corresponds to, a part of a natural language like English or Greek. The deductive system is to capture, codify, or simply record which inferences are correct for the given language, and the semantics is to capture, codify, or record the meanings, or truth-conditions, or possible truth conditions, for at least part of the language.

The following sections provide the basics of a typical logic, sometimes called "classical elementary logic" or "classical first-order logic". Section 2 develops a formal language, with a rigorous syntax and grammar. The formal language is a recursively defined collection of strings on a fixed alphabet. As such, it has no meaning, or perhaps better, the meaning of the formulas is given by the deductive system and the semantics. Some of the symbols have counterparts in ordinary language. We define an argument to be a non-empty collection of formulas in the formal language, one of which is designated to be the conclusion. The other formulas (if any) in an argument are its premises. Section 3 sets up a deductive system for the language, in the spirit of natural deduction. An argument is derivable if there is a deduction from some of its premises to its conclusion. Section 4 provides a model-theoretic semantics. An argument is valid if there is no interpretation (in the semantics) in which its premises are all true and its conclusion false. This reflects the longstanding view that a valid argument is truth-preserving.

In Section 5, we turn to relationships between the deductive system and the semantics, and in particular, the relationship between derivability and validity. We show that an argument is derivable only if it is valid. This pleasant feature, called soundness, entails that no deduction takes one from true premises to a false conclusion. Thus, deductions preserve truth, and there aren't too many deductions. Then we establish a converse, called completeness, that an argument is valid only if it is derivable. This establishes that the deductive system is rich enough to provide a deduction for every valid argument. There are enough deductions. All and only valid arguments are derivable. We briefly indicate other features of the logic, some of which are corollaries to soundness and completeness.

1. Introduction

2. Language

3. Deduction

4. Semantics

5. Meta-theory

Bibliography

Other Internet Resources

Related Entries

1. Introduction

Today, logic is both a branch of mathematics and a branch of philosophy. In most large universities, both departments offer sequences of courses in logic, and there is usually a lot of overlap between them. Formal languages, deductive systems, and model-theoretic semantics are mathematical objects and, as such, the logician is interested in their mathematical properties and relations. Soundness, completeness, and most of the other results reported below are typical examples. Philosophically, logic is the study of correct reasoning. Reasoning is an epistemic, mental activity. This raises questions concerning the philosophical relevance of the mathematical aspects of logic. How do deducibility and validity, as properties of formal languages--sets of strings on a fixed alphabet--relate to correct reasoning? What do the mathematical results reported below have to do with the original philosophical issue? This is an instance of the philosophical problem of explaining how mathematics applies to non-mathematical reality.

Typically, ordinary reasoning takes place in a natural language, or perhaps a natural language augmented with some mathematical symbols. So our question begins with the relationship between a natural language and a formal language. Without attempting to be comprehensive, it may help to sketch several options on this matter.

One view is that the formal languages accurately exhibit actual features of certain fragments of a natural language. Some philosophers claim that declarative sentences of natural language have underlying logical forms and that these forms are displayed by formulas of a formal language. Other writers hold that (successful) declarative sentences express propositions; and formulas of formal languages somehow display the forms of these propositions. On views like this, the components of a logic provide the underlying deep structure of correct reasoning. A chunk of reasoning in natural language is correct if

the forms underlying the sentences constitute a valid or deducible argument. See for example, Montague [1974], Davidson [1984], Lycan [1984].

Another view, held at least in part by Gottlob Frege and Wilhelm Leibniz, is that because natural languages are vague and ambiguous, they should be replaced by formal languages. A similar view, held by W. V. O. Quine (e.g., [1960], [1986]), is that a natural language should be regimented, cleaned up for serious scientific and metaphysical work. One desideratum of the enterprise is that the logical structures in the regimented language should be transparent. It should be easy to "read off" the logical properties of each sentence. A regimented language is similar to a formal language regarding, for example, the explicitly presented rigor of its syntax and its truth conditions.

On a view like this, deducibility and validity represent idealizations of correct reasoning in natural language. A chunk of reasoning is correct to the extent that it corresponds to, or can be regimented by, a valid or deducible argument in a formal language.

When mathematicians and many philosophers reason, they occasionally invoke formulas in a formal language to help disambiguate, or otherwise clarify what they mean. In other words, sometimes formulas in a formal language are used in ordinary reasoning. This suggests that one might think of a formal language as an addendum to a natural language. Then our present question concerns the relationship between this addendum and the original language. What do deducibility and validity, as sharply defined on the addendum, tell us about correct reasoning in general?

Another view is that a formal language is a mathematical model of a natural language in roughly the same sense as, say, a collection of point masses is a model of a system of physical objects, and the Bohr construction is a model of an atom. In other words, a formal language displays certain features of natural languages, or idealizations thereof, while ignoring or simplifying other features. The purpose of mathematical models is to shed light on what they are models of, without claiming that the model is accurate in all respects or that the model should replace what it is a model of. On a view like this, deducibility and validity represent mathematical models of (perhaps different aspects of) correct reasoning in natural languages. Correct chunks of reasoning correspond, more or less, to valid or deducible arguments; incorrect chunks of reasoning roughly correspond to invalid or non-deducible arguments. See, for example, Corcoran [1973] or Shapiro [1998].

There is no need to adjudicate this matter here. Perhaps the truth lies in a combination of the above options, or maybe some other option is the correct, or most illuminating one. I raise the matter only to lend some philosophical perspective to the formal treatment that follows.

2. Language

Here we develop the basics of a formal language, or to be precise, a class of formal languages. Again, a formal language is a recursively defined set of strings on a fixed alphabet. Some aspects of the formal languages correspond to, or have counterparts in, natural languages like English. Technically, this "counterpart relation" is not part of the formal development, but I will mention it from time to time, to motivate some of the features and results.

Building blocks

We begin with analogues of singular terms, linguistic items whose function is to denote a person or object. We call these terms. We assume a stock of individual constants. These are lower-case letters, near the beginning of the Roman alphabet, with or without numerical subscripts: a , a_1 , b_{23} , c , d_{22} , etc.

We envisage a potential infinity of individual constants. In the present system each constant is a single character, and so individual constants do not have an internal syntax. Thus we have an infinite alphabet. This last could be avoided by taking a constant like d_{22} , for example, to consist of three characters, a lowercase "d" followed by a pair of subscript "2"s.

We also assume a stock of individual variables. These are lower-case letters, near the end of the alphabet, with out without numerical subscripts: w , x , y_{12} , z , z_4 , etc.

Variables serve a dual function. Sometimes a variable is used as a singular term to denote a specific, but unspecified (or arbitrary) object. For example, a mathematician might start a derivation: "Let x be a natural number". Variables are also used to express generality, as in the mathematical assertion that for any natural number x , there is a natural number y , such that $y > x$ and y is prime. Some logicians employ different symbols for unspecified objects (sometimes called "individual parameters") and variables used to express generality.

Constants and variables are the only terms in our formal language, so all of our terms are simple, corresponding to proper names and pronouns. Some authors also introduce function letters, which allow complex terms corresponding to: "7+4" and "the wife of Bill Clinton", or complex terms containing variables, like "the father of x" and "x/y". Logic books aimed at mathematicians are likely to contain function letters, probably due to the centrality of functions to mathematical discourse. Books aimed at a more general audience (or at philosophy students), may leave out function letters, since it simplifies the syntax and theory. We follow the latter route here. This is an instance of a general tradeoff between presenting a system with greater expressive resources, at the cost of making its formal treatment more complex.

For each natural number n , we introduce a stock of n -place predicate letters. These are upper-case letters at the beginning or middle of the alphabet. A superscript indicates the number of places, and there may or may not be a subscript. For example, A_3 , B_{32} , P_3 , etc.

are three-place predicate letters. We often omit the superscript, when no confusion will result. We also add a special two-place symbol "=" for identity.

Zero-place predicate letters are sometimes called "sentence letters". They correspond to free-standing sentences whose internal structure does not matter. One-place predicate letters, called "monadic predicate letters", correspond to linguistic items denoting properties, like "being a man", "being red", or "being a prime number". Two-place predicate letters, "binary predicate letters", correspond to linguistic items denoting binary relations, like "is a parent of" or "is greater than". Three-place predicate letters correspond to three-place relations, like "lies on a straight line between". And so on.

The non-logical terminology of the language consists of its individual constants and predicate letters. The symbol "=", for identity, is not a non-logical symbol. In taking identity to be logical, we provide explicit treatment for it in the deductive system and the model-theoretic semantics. Most authors do the same, but there is some controversy over the issue (Quine [1986, Chapter 5]). If K is a set of constants and predicate letters, then we give the fundamentals of a language $1K=$ built on this set of non-logical terminology. It may be called the first-order language with identity on K . A similar language that lacks the symbol for identity (or which takes identity to be non-logical) may be called $1K$, first-order logic without identity.

Atomic formulas

If V is an n -place predicate letter in K , and t_1, \dots, t_n are terms of K (i.e., constants in K or variables), then $\forall t_1 \dots t_n$ is an atomic formula of $1K=$. Notice that the terms t_1, \dots, t_n need not be distinct. Examples of atomic formulas include:

$P4xaab, C1x, C1a, D0, A3abc.$

The last one is an analogue of a statement that a certain relation (A) holds between three objects (a, b, c). If t_1 and t_2 are terms, then $t_1=t_2$ is an atomic formula of $1K=$. It corresponds to an assertion that t_1 is identical to t_2 .

If an atomic formula has no variables, then it is called an atomic sentence. If it does have variables, it is called an open formula. In the above list of examples, the first and second are open; the rest are sentences.

Compound formulas

We now introduce the final items of the lexicon:

$\neg, \&, \vee, \rightarrow, (,)$

We give a recursive definition of a formula of $1K=$:

1. All atomic formulas of $1K=$ are formulas of $1K=$.
2. If ϕ is a formula of $1K=$, then so is $\neg\phi$.

Asserting a sentence corresponding to $\neg\phi$ is tantamount to denying the sentence corresponding to ϕ . The symbol " \neg " is called "negation", and is a unary connective.

3. If ϕ and ψ are formulas of $1K=$, then so is $(\phi \& \psi)$.

The ampersand "&" corresponds to the English "and" (when "and" is used to connect sentences). So $(\phi \& \psi)$ can be read " ϕ and ψ ". The formula $(\phi \& \psi)$ is called the "conjunction" of ϕ and ψ .

4. If ϕ and ψ are formulas of $1K=$, then so is $(\phi \vee \psi)$.

The wedge " \vee " corresponds to "either . . . or . . . or both", so $(\phi \vee \psi)$ can be read " ϕ or ψ ". The formula $(\phi \vee \psi)$ is called the "disjunction" of ϕ and ψ .

5. If ϕ and ψ are formulas of $1K=$, then so is $(\phi \rightarrow \psi)$.

The arrow " \rightarrow " corresponds to "if . . . then . . .", so $(\phi \rightarrow \psi)$ can be read "if ϕ then ψ " or " ψ only if ϕ ".

The symbols "&", " \vee ", and " \rightarrow " are called "binary connectives", since they serve to "connect"

two sentences into one. Some authors introduce \leftrightarrow as an abbreviation of $(\rightarrow) \& (\leftarrow)$. The symbol \leftrightarrow is an analogue of the locution "if and only if".

6. If ϕ is a formula of $L_{K=}$ and v is a variable, then $\forall v \phi$ is a formula of $L_{K=}$.

The symbol \forall is called a universal quantifier, and is an analogue of "for all"; so $\forall v \phi$ can be read "for all v , ϕ ".

7. If ϕ is a formula of $L_{K=}$ and v is a variable, then $\exists v \phi$ is a formula of $L_{K=}$.

The symbol \exists is called an existential quantifier, and is an analogue of "there exists" or "there is"; so $\exists v \phi$ can be read "there is a v such that ϕ ".

8. That's all folks. That is, all formulas are constructed in accordance with rules (1)-(7).

Clause (8) allows us to do inductions on the complexity of formulas. If a certain property holds of the atomic formulas and is closed under the operations presented in clauses (2)-(7), then the property holds of all formulas.

We next define the notion of an occurrence of a variable being free or bound in a formula. All variables that occur in an atomic formula are free. If a variable occurs free (or bound) in ϕ or in ψ , then that same occurrence is free (or bound) in $\neg \phi$, $(\phi \& \psi)$, $(\phi \vee \psi)$, and $(\phi \rightarrow \psi)$. That is, the (unary and binary) connectives do not change the status of variables that occur in them. All occurrences of the variable v in $\forall v \phi$ are bound in $\forall v \phi$ and ϕ . Any free occurrences of v in ϕ are bound by the initial quantifier. All other variables that occur in $\forall v \phi$ are free or bound in ϕ and v , as they are in ϕ . A variable that immediately follows a quantifier (as in "x" and "y") is neither free nor bound. We do not think of those as occurrences of the variable.

For example, in the formula $(\forall x(Axy \rightarrow Bx) \& Bx)$, the occurrences of "x" in Axy and in the first Bx are bound by the quantifier. The occurrence of "y" and last occurrence of "x" are free. In $\forall x(Ax \rightarrow \exists x Bx)$, the "x" in Ax is bound by the initial universal quantifier, while the other occurrence of x is bound by the existential quantifier. The above syntax allows this "overlap" of bound variables, and it does not create an ambiguity, but we will avoid such formulas, as a matter of taste and clarity.

Free variables correspond to place-holders, while bound variables are used to express generality. If a formula has no free variables, then it is called a sentence. If a formula has free variables, it is called open.

Features of the syntax

Before turning to the deductive system and semantics, I mention a few features of the language, as developed so far. This helps draw the contrast between formal languages and natural languages like English.

We assume at the outset that all of the categories are disjoint. For example, no connective is also a quantifier or a variable, and the non-logical terms are not also parentheses or connectives. Also, the items within each category are distinct. For example, the sign for disjunction does not do double-duty as the negation symbol, and perhaps more significantly, no two-place predicate is also a one-place predicate.

Theorem 1. Every formula of $1K=$ has the same number of left and right parentheses. Moreover, each left parenthesis corresponds to a unique right parenthesis, which occurs to the right of the left parenthesis. Similarly, each right parenthesis corresponds to a unique left parenthesis, which occurs to the left of the given right parenthesis. If a parenthesis occurs between a matched pair of parentheses, then its mate also occurs within that matched pair. In other words, parentheses that occur within a matched pair are themselves matched.

Proof: By clause (8), every formula is built up from the atomic formulas using clauses (2)-(7). The atomic formulas have no parentheses (by the policy that the categories are disjoint). Parentheses are introduced only in clauses (3)-(5), and each time they are introduced as a matched set. So at any stage in the construction of a formula, the parentheses are paired off.

One difference between natural languages like English and formal languages like $1K=$ is that the latter are not supposed to have any ambiguities. Our policy that the different categories of symbols do not overlap, and that no symbol does double-duty helps avoid the kind of ambiguity, sometimes called "equivocation", that occurs when a single word has two meanings: "I'll meet you at the bank." But there are other kinds of ambiguity. Consider the English sentence:

John is married, and Mary is single, or Joe is crazy.

It can mean that John is married and either Mary is single or Joe is crazy, or else it can mean that either both John is married and Mary is single, or else Joe is crazy. An ambiguity like this, due to different ways to parse the same sentence, is sometimes called an "amphiboly". If our formal language did not have the parentheses in it, it would have amphibolies. For example, there would be a "formula" $A \& B \ C$. Is this supposed to be $((A \& B) \ C)$, or is it $(A \& (B \ C))$? The parentheses resolve what would be an amphiboly. Can we be sure that there are no other amphibolies in our language? That is, can we be

sure that each formula of $1K=$ can be put together in only one way? Showing this is our next task.

Let us temporarily use the term "unary marker" for the negation symbol (\neg) or a quantifier followed by a variable (e.g., x, z).

Lemma 2. Each formula consists of a string of zero or more unary markers followed by either an atomic formula or a formula produced using a binary connective, via one of clauses (3)-(5).

Proof: We proceed by induction on the complexity of the formula or, in other words, on the number of formation rules that are applied. The Lemma clearly holds for atomic formulas. Let n be a natural number, and suppose that the Lemma holds for any formula constructed from n or fewer instances of clauses (2)-(7). Let ϕ be a formula constructed from $n+1$ instances. The Lemma holds if the last clause used to construct ϕ was either (3), (4), or (5). If the last clause used to construct ϕ was (2), then ϕ is $\neg \psi$. Since ψ was constructed with n instances of the rule, the Lemma holds for ψ (by the induction hypothesis), and so it holds for ϕ . Similar reasoning shows the Lemma to hold for ϕ if the last clause was (6) or (7). By clause (8), this exhausts the cases, and so the Lemma holds for ϕ , by induction.

Lemma 3. If a formula ϕ contains a left parenthesis, then it ends with a right parenthesis, which matches the leftmost left parenthesis in ϕ .

Proof: Here we also proceed by induction on the number of instances of (2)-(7) used to construct the formula. Clearly, the Lemma holds for atomic formulas, since they have no parentheses. Suppose, then, that the Lemma holds for formulas constructed with n or fewer instances of (2)-(7), and let ϕ be constructed with $n+1$ instances. If the last clause applied was (3)-(5), then the Lemma holds since ϕ itself begins with a left parenthesis and ends with the matching right parenthesis. If the last clause applied was (2), then ϕ is $\neg \psi$, and the induction hypothesis applies to ψ . Similarly, if the last clause applied was (6) or (7), then ϕ consists of a quantifier, a variable, and a formula to which we can apply the induction hypothesis. It follows that the Lemma holds for ϕ .

Lemma 4. Each formula contains at least one atomic formula.

The proof proceeds by induction on the number of instances of (2)-(7) used to construct the formula, and we leave it as an exercise.

Theorem 5. Let α, β be nonempty sequences of characters on our alphabet, such that $\alpha(\beta)$ (i.e. α followed by β) is a formula. Then α is not a formula.

Proof: By Theorem 1 and Lemma 3, if α contains a left parenthesis, then the right parenthesis that matches the leftmost left parenthesis in α comes at the end of β , and so the matching right parenthesis is in β . So, α has more left parentheses than right parentheses. By Theorem 1, α is not a formula. So now suppose that α does not contain any left parentheses. By Lemma 2, α consists of a string of zero or more unary markers followed by either an atomic formula or a formula produced using a binary connective, via one of clauses (3)-(5). If the latter formula was produced via one of clauses (3)-(5), then it begins with a left parenthesis. Since α does not contain any parentheses, it must be a string of unary markers. But then α does not contain any atomic formulas, and so by Lemma 4, α is not a formula. The only case left is where α consists of a string of unary markers followed by an atomic formula, either in the form $t_1=t_2$ or $Pt_1 \dots t_n$. Again, if α just consisted of unary markers, it would not be a formula, and so α must consist of the unary markers that start α , followed by either t_1 by itself, $t_1=$ by itself, or the predicate letter P , and perhaps some (but not all) of the terms t_1, \dots, t_n . In the first two cases, α does not contain an atomic formula, by the policy that the categories do not overlap. Since P is an n -place predicate letter, by the policy that the predicate letters are distinct, P is not an m -place predicate letter for any $m \neq n$. So the part of α that consists of P followed by the terms is not an atomic formula. In all of these cases, then, α does not contain an atomic formula. By Lemma 4, α is not a formula.

We are finally in position to show that there is no amphiboly in our language.

Theorem 6. Let α be any formula of $1K=$. If α is not atomic, then there is one and only one among (2)-(7) that was the last clause applied to construct α . That is, α could not be produced by two different clauses. Moreover, no formula produced by clauses (3)-(7) is atomic.

Proof: By Clause (8), either α is atomic or it was produced by one of clauses (2)-(7). Thus, the first symbol in α must be either a predicate letter, a term, a unary marker, or a left parenthesis. If the first symbol in α is a predicate letter or term, then α is atomic. In this case, α was not produced by any of (2)-(7), since all such formulas begin with something other than a predicate letter or term. If the first symbol in α is a negation sign " \neg ", then α was produced by clause (2), and not by any other clause (since the other clauses produce formulas that begin with either a quantifier or a left parenthesis). Similarly, if α begins with a universal quantifier, then it was produced by clause (6), and not by any other clause, and if α begins with an existential quantifier, then it was produced by clause (7), and not by any other clause. The only case left is where α begins with a left parenthesis. In this case, it

must have been produced by one of (3)-(5), and not by any other clause. We only need to rule out the possibility that ϕ was produced by more than one of (3)-(5). To take an example, suppose that ϕ was produced by (3) and (4). Then ϕ is $(1 \ \& \ 2)$ and ϕ is also $(3 \ 4)$, where 1, 2, 3, and 4 are themselves formulas. That is, $(1 \ \& \ 2)$ is the very same formula as $(3 \ 4)$. By Theorem 5, 1 cannot be a proper part of 3, nor can 3 be a proper part of 1. So 1 must be the same formula as 3. But then "&" must be the same symbol as "", and this contradicts the policy that all of the symbols are different. So ϕ was not produced by both Clause (3) and Clause (4). Similar reasoning takes care of the other combinations.

This result is sometimes called "unique readability". It shows that each formula is produced from the atomic formulas via the various clauses in exactly one way. If ϕ was produced by clause (2), then its main connective is the initial "-". If ϕ was produced by clauses (3), (4), or (5), then its main connective is the introduced "&", "", or "", respectively. If ϕ was produced by clauses (6) or (7), then its main connective is the initial quantifier. I apologize for the tedious details. I included them to indicate the level of precision and rigor for the syntax.

3. Deduction

We now introduce a deductive system, D , for our languages. As above, we define an argument to be a non-empty collection of formulas in the formal language, one of which is designated to be the conclusion. If there are any other formulas in the argument, they are its premises. By convention, we use $\{ \}$, $\{ \}$, $\{ \}$, etc, to range over sets of formulas, and we use the letters ϕ , ψ , χ , uppercase or lowercase, with or without subscripts, to range over single formulas. We write \cup , \cup for the union of $\{ \}$ and $\{ \}$, and \cup , \cup for the union of $\{ \}$ with $\{ \}$.

We write an argument in the form $\langle \{ \}, \phi \rangle$, where $\{ \}$ is the set of premises and ϕ is the conclusion. Remember that $\{ \}$ may be empty. We write \vdash to indicate that ϕ is deducible from $\{ \}$, or, in other words, that the argument $\langle \{ \}, \phi \rangle$ is deducible in D . We may write \vdash_D to emphasize the deductive system D . We write \vdash or \vdash_D to indicate that ϕ can be deduced (in D) from the empty set of premises.

The rules in D are chosen to match logical relations concerning the English analogues of the logical terminology in the language. Again, we define the deducibility relation by recursion. We start with a rule of assumptions:

(As) If ϕ is a member of $\{ \}$, then $\vdash \phi$.

We thus have that $\{ \} \vdash \phi$; each premise follows from itself. We next present two clauses for

each connective and quantifier. The clauses indicate how to "introduce" and "eliminate" formulas in which each symbol is the main connective.

First, recall that "&" is an analogue of the English connective "and". Intuitively, one can deduce a formula in the form (&) if one has deduced and one has deduced . Conversely, one can deduce from (&) and one can deduce from (&):

(&I) If 1 and 2 , then 1, 2 (&).

(&E) If (&) then ; and if (&) then .

The name "&I" stands for "&-introduction"; "&E" stands for "&-elimination".

Since, the symbol "" corresponds to the English "or", () should be deducible from , and () should also be deducible from :

(I) If then (); if then ().

The elimination rule is a bit more complicated. Suppose that " or " is true. Suppose also that follows from and that follows from . One can reason that if is true, then is true. If instead is true, we still have that is true. So either way, must be true.

(E) If 1 (), 2, and 3, , then 1, 2, 3 .

For the next clauses, recall that the symbol "" is an analogue of the English "if . . . then . . ." construction. If one knows, or assumes () and also knows, or assumes , then one can conclude . Conversely, if one deduces from an assumption , then one can conclude that ().

(I) If , , then ().

(E) If 1 () and 2 , then 1, 2 .

Our next clauses are for the negation sign, "¬". The underlying idea is that a formula is inconsistent with its negation ¬. They cannot both be true. We call a pair of formulas , ¬ contradictory opposites. If one can deduce such a pair from an assumption , then one can conclude that is false, or, in other words, one can conclude ¬.

(¬I) If 1, and 2, ¬, then 1, 2 ¬.

There is some controversy over the other rule for the negation sign.

By (As), we have that {A, ¬A} A and {A, ¬A} ¬A. So by ¬I we have that {A} ¬¬A. However, we do not have the converse yet. Intuitively, ¬¬ corresponds to "it is not the case that it is not the case that" . One might think that this last is equivalent to , and we have a rule to that effect:

(DNE) If $\neg\neg A$, then A .

The name DNE stands for "double-negation elimination". This inference is rejected by philosophers and mathematicians who do not hold that each meaningful sentence is either true or not true. Intuitionistic logic does not sanction the inference in question (see, for example Dummett [1977]), but, again, classical logic does.

To illustrate the parts of the deductive system D presented thus far, I show that $(A \rightarrow \neg A)$:

(i) $\{\neg(A \rightarrow \neg A), A\} \neg(A \rightarrow \neg A)$, by (As)

(ii) $\{\neg(A \rightarrow \neg A), A\} A$, by clause (As).

(iii) $\{\neg(A \rightarrow \neg A), A\} (A \rightarrow \neg A)$, by (I), from (ii).

(iv) $\{\neg(A \rightarrow \neg A)\} \neg A$, by (\neg I), from (i) and (iii).

(v) $\{\neg(A \rightarrow \neg A), \neg A\} \neg(A \rightarrow \neg A)$, by (As)

(vi) $\{\neg(A \rightarrow \neg A), \neg A\} \neg A$, by (As)

(vii) $\{\neg(A \rightarrow \neg A), \neg A\} (A \rightarrow \neg A)$, by (I), from (vi).

(viii) $\{\neg(A \rightarrow \neg A)\} \neg\neg A$, by (\neg I), from (v) and (vii).

(ix) $\neg\neg(A \rightarrow \neg A)$, by (\neg I), from (iv) and (viii).

(x) $(A \rightarrow \neg A)$, by (DNE), from (ix).

The principle $(\neg\neg A \rightarrow A)$ is sometimes called the law of excluded middle. It is not valid in intuitionistic logic.

Let A, B be a pair of contradictory opposites, and let C be any formula at all. By (As) we have $\{A, \neg A\} C$ and $\{B, \neg B\} C$. So by (\neg I), $\{A, \neg A\} \neg C$. So, by (DNE) we have $\{A, \neg A\} C$. That is, anything at all follows from a pair of contradictory opposites. Some logicians introduce a rule to codify a similar inference:

If 1 and 2 \neg , then for any formula C , 1, 2

The inference is sometimes called *ex falso quodlibet*. Some call it " \neg -elimination", but perhaps this stretches the notion of "elimination" a bit. We do not officially include *ex*

falso quodlibet as a separate rule in D, but as will be shown below (Theorem 10), each instance of it is derivable.

Some logicians object to ex falso quodlibet, on the ground that the formula ϕ may be irrelevant to any of the premises in Γ . Suppose, for example, that one starts with some premises Γ about human nature and facts about certain people, and then deduces both the sentence "Clinton had extra-marital sexual relations" and "Clinton did not have extra-marital sexual relations". One can surely conclude that there is something wrong with premises Γ . But should we be allowed to then deduce anything at all from Γ ? Should we be allowed to deduce "The economy is sound"?

Deductive systems that demur from ex falso quodlibet are part of relevance logic. See Anderson and Belnap [1975], Anderson, Belnap, and Dunn [1992], and Tennant [1987]. Deep philosophical issues concerning the nature of logical consequence are involved. Far be it for an article in a philosophy encyclopedia to avoid philosophical issues, but space considerations preclude a fuller treatment of this issue here. Suffice it to note that the inference is sanctioned in systems of classical logic, the subject of this article. It is essential to establishing the balance between the deductive system and the semantics (see §5 below).

The next pieces of D are the clauses for the quantifiers. Let ϕ be a formula, v a variable, and t a term (i.e., a variable or a constant). We define $\phi(v|t)$ to be the result of substituting t for each free occurrence of v in ϕ . So, if ϕ is $(\forall x \phi(x))$, then $\phi(c)$ is $(\forall x \phi(x))$. The last occurrence of x is not free (but recall that we avoid using formulas like this).

We have one other nicety to attend to. Suppose that v_1 and v_2 are variables. It may happen that some of the substituted instances of v_2 are bound in $\phi(v_1|v_2)$. When this happens, we say that there is a clash of the variables. Suppose, for example, that ϕ is $y(\neg x = y)$, and so $\phi(y)$ is $y(\neg y = y)$. We say that a term t is free for a variable v in ϕ if either t is a constant or there is no clash of variables in $\phi(v|t)$. The idea is that no substituted instance of t should become a bound variable in $\phi(v|t)$.

A formula in the form $\forall v \phi$ is an analogue of the English "for every v , ϕ holds". So one should be able to infer $\phi(v|t)$ from $\forall v \phi$:

(E) If $\forall v \phi$, then $\phi(v|t)$, provided that t is free for v in ϕ .

The idea here is that if $\forall v \phi$ is true, then ϕ should hold of t , no matter what t is. We can

illustrate the restriction on (E) as follows: The sentence $\forall x \exists y (\neg x=y)$ corresponds to an assertion that for every object x , there is an object different from x . This is a coherent, plausible assertion. It is true if and only if the universe has at least two objects. It should follow that no matter what object t may be, something is different from t . However, if we were allowed to substitute the variable y for x , we would conclude $\exists y (\neg y=y)$, which says that there is something which is different from itself, a blatant falsehood.

The introduction clause for the universal quantifier is a bit more complicated. Suppose that a formula ϕ has a variable v free, and that ϕ has been deduced from a set of premises Γ . If the variable v does not occur free in any member of Γ , then $\forall v \phi$ will hold no matter which object v may denote. That is, $\forall v \phi$ follows.

(I) If Γ and the variable v does not occur free in any member of Γ , then $\forall v \phi$.

This introduction rule corresponds to a common inference in mathematics. Suppose that a mathematician says "let n be a natural number" and goes on to show that n has a certain property P , without assuming anything about n (except that it is a natural number). She then reminds the reader that n is "arbitrary", and concludes that P holds for all natural numbers. The condition that the variable v not occur in any premise is what guarantees that it is indeed "arbitrary". It could be any object, and so anything we conclude about it holds for all objects.

The existential quantifier is an analogue of the English expression "there exists", or perhaps just "there is". If we have established (or assumed) that a given object t has a given property, then it follows that there is something that has that property. Again, we have to be careful with the syntax, and avoid clashes of variables.

(I) If $\Gamma, \phi(t)$, then $\exists v \phi(v)$, where $\phi(v)$ is obtained from $\phi(t)$ by substituting the variable v for zero or more occurrences of a term t , provided that (1) if t is a variable, then all of the replaced occurrences of t are free in $\phi(t)$, and (2) all of the substituted occurrences of v are free in $\phi(v)$.

The provision (1) keeps us from replacing bound variables. The provision (2) comes up only if v is bound by another quantifier in $\phi(t)$. As noted above, we avoid such formulas (since they appear to bind the same occurrence twice).

The elimination rule for \exists is not quite as simple:

(E) If $\Gamma, \exists v \phi$ and Δ, ψ , then Γ, Δ , provided that v does not occur free in Γ , nor in any member of Δ .

This elimination rule also corresponds to a common inference. Suppose that a mathematician asserts that there is a natural number with a given property P . She then

says "let n be such a natural number, so that Pn ", and goes on to establish a sentence ϕ , which does not mention the number n . If the derivation of ϕ does not invoke anything about n (other than the assumption that it has the given property P), then n could have been any number that has the property P . That is, n is an arbitrary number with property P . It does not matter which number n is. Since ϕ does not mention n , it follows from the assertion that something has property P . The provisions added to (E) are to guarantee that x is "arbitrary".

As noted in the previous section, some authors introduce different letters for bound variables and (what amounts to) free variables. This makes the syntax slightly more complex, but simplifies the provisions on some of the rules of inference. Writers of logic books often face tradeoffs like this.

The final items are the rules for the identity sign " $=$ ". The introduction rule is about as simple as can be:

(=I) $t=t$, where t is any term.

This "inference" corresponds to the truism that everything is identical to itself. The elimination rule corresponds to a principle that if a is identical to b , then anything true of a is also true of b , again paying attention to clashes of variables.

(=E) If ϕ $t_1=t_2$ and ψ , then ϕ , ψ , where ψ is obtained from ϕ by replacing zero or more occurrences of t_1 with t_2 , provided that no bound variables are replaced, and if t_2 is a variable, then all of its substituted occurrences are free.

The rule (=E) indicates a certain restriction in the expressive resources of our language. Suppose, for example, that Harry is identical to Donald (since his mischievous parents gave him two names). It would not follow from this and "Dick knows that Harry is wicked" that "Dick knows that Donald is wicked", for the reason that Dick might not know that Harry is identical to Donald. Contexts like this, in which identicals cannot safely be substituted for each other, are called "opaque". We assume that our language $L_{K=}$ has no opaque contexts.

One final clause completes the description of the deductive system D :

(*) That's all folks. ϕ only if ϕ follows from members of Γ by the above rules.

Again, this clause allows proofs by induction on the rules used to establish an inference. If a property of arguments holds of all instances of (As) and (=I), and if the other rules preserve the property, then every argument that is deducible in D enjoys the property in question.

Before moving on to the model theory for $1K=$, we pause to note a few features of the deductive system.

Lemma 7. Suppose that $D \vdash \phi$, and let v be a variable that does not occur free in ϕ or in any member of Γ . Assume that v is free for v in ϕ and in every member of Γ . Let ϕ' be $\{\phi(v|v) \mid \phi \in \Gamma\}$. That is, ϕ' is the result of replacing every free occurrence of a variable v with v in every member of Γ . Then $D \vdash \phi(v|v)$.

Proof: The proof of this lemma is tedious, but we give its essentials. We proceed by induction on the number of rules that were used to arrive at ϕ . Suppose that $n > 0$ is a natural number, and that the lemma holds for any argument that was derived using fewer than n rules, and suppose that $\Gamma \vdash \phi$ using exactly n rules. If $n=1$, then the rule applied is either (As) or (=I). In this case, $\phi(v|v)$ by the same rule. If the last rule applied is (&I), then ϕ has the form $(\& \phi_1 \phi_2)$, and we have $\Gamma \vdash \phi_1$ and $\Gamma \vdash \phi_2$, with $n = 1, 2$. We apply the induction hypothesis to the deductions of ϕ_1 and ϕ_2 , and then apply (&I) to the result. If the last rule applied was (&E), we have two sub-cases, but they are symmetric. We have $\Gamma \vdash (\& \phi_1 \phi_2)$. There are two slight complications here: the new variable v may occur free in ϕ_1 and it may not be free for v in ϕ_2 . In either case, first pick a new variable u that does not occur (free or bound) in $(\& \phi_1 \phi_2)$ or in any member of Γ . Now apply the induction hypothesis, substituting u for v in the deduction $(\& \phi_1 \phi_2)$. Since v does not occur free in ϕ_1 or in any member of Γ , those formulas are left unchanged. The maneuver removes any free occurrences of v from the subformula ϕ_2 . Now apply the induction hypothesis to the result, substituting v for u , and then apply (&E). The remaining cases are similar.

Theorem 8. The rule of Weakening. If $\Gamma \vdash \phi$ and $\Gamma \vdash \psi$, then $\Gamma \vdash \phi \wedge \psi$.

Proof: Again, we proceed by induction on the number of rules that were used to arrive at $\phi \wedge \psi$. Suppose that $n > 0$ is a natural number, and that the theorem holds for any argument that was derived using fewer than n rules. Suppose that $\Gamma \vdash \phi \wedge \psi$ using exactly n rules. If $n=1$, then the rule is either (As) or (=I). In these cases, $\Gamma \vdash \phi$ and $\Gamma \vdash \psi$ by the same rule. If the last rule applied was (&I), then $\phi \wedge \psi$ has the form $(\& \phi_1 \phi_2)$, and we have $\Gamma \vdash \phi_1$ and $\Gamma \vdash \phi_2$, with $n = 3, 4$. We apply the induction hypothesis to the deductions of ϕ_1 and ϕ_2 , to get $\Gamma \vdash \phi_1$ and $\Gamma \vdash \phi_2$, and then apply (&I) to the result to get $\Gamma \vdash \phi \wedge \psi$. Most of the other cases are exactly like this. Slight complications arise only in the rules (I) and (E), because there we have to pay attention to the conditions for the rules. Starting with (E), we have $\Gamma \vdash \phi \wedge \psi$ and $\Gamma \vdash \chi$, with n being $3, 4$, and v not free in χ , nor in any member of Γ . We apply the induction hypothesis to get $\Gamma \vdash \phi \wedge \psi$, and then (E) to end up with $\Gamma \vdash \chi$. Suppose that the last rule applied to get $\phi \wedge \psi$ is (I). So $\phi \wedge \psi$ is a formula in the form $\phi \wedge \psi$, and

we have 1 and the variable v does not occur free in any member of 1 . The problem is that v may occur free in a member of 2 , and so we cannot just invoke the induction hypothesis and apply (I) to the result. Let v be a variable that does not occur (free or bound) in 1 or in any member of 2 , and let $2'$ be the result of substituting v for every free occurrence of v in 2 . Since v does not occur free in any member of 1 , we still have 1 . The induction hypothesis gives us $2'$, and now we apply (I) to get 2 . We now apply Lemma 7, substituting v for the new variable v . The result is 2 .

Theorem 8 allows us to add on premises at will. It follows that $\Gamma \vdash \phi$ if and only if there is a subset Δ such that $\Delta \vdash \phi$. By clause (*), all derivations are established in a finite number of steps. So we have

Theorem 9. $\Gamma \vdash \phi$ if and only if there is a finite Δ such that $\Delta \vdash \phi$.

Theorem 10. The rule of ex falso quodlibet is a "derived rule" of D. That is, if 1 and $2 \neg$, then $1, 2 \vdash \phi$, for any formula ϕ .

Proof: Suppose that 1 and $2 \neg$. Then by Theorem 8, $1, \neg$, and $2, \neg$. So by (\neg I), $1, 2 \neg \neg$. By (DNE), $1, 2 \vdash \phi$.

Theorem 11. The rule of Cut. If $1 \vdash \phi$ and $2, \phi \vdash \psi$, then $1, 2 \vdash \psi$.

Proof: Suppose $1 \vdash \phi$ and $2, \phi \vdash \psi$. We proceed by induction on the number of rules used to establish $2, \phi \vdash \psi$. Suppose that n is a natural number, and that the theorem holds for any argument that was derived using fewer than n rules. Suppose that $2, \phi \vdash \psi$ was derived using exactly n rules. If the last rule used was ($=$ I), then $1, 2 \vdash \psi$ is also an instance of ($=$ I). If $2, \phi \vdash \psi$ is an instance of (As), then either ϕ is ψ , or ϕ is a member of 2 . In the former case, we have 1 by supposition, and get $1, 2 \vdash \psi$ by Weakening (Theorem 8). In the latter case, $1, 2 \vdash \psi$ is itself an instance of (As). Suppose that $2, \phi \vdash \psi$ was obtained using ($\&$ E). Then we have $2, \phi \vdash \psi$ ($\&$). The induction hypothesis gives us $1, 2 \vdash \psi$ ($\&$), and ($\&$ E) produces $1, 2 \vdash \psi$. The remaining cases are similar.

Theorem 11 allows us to chain together inferences. This fits the practice of establishing theorems and lemmas and then using those theorems and lemmas later, at will. The cut principle is, I think, essential to reasoning. In some logical systems, the cut principle is a deep theorem. The system here was designed to make the proof of Theorem 11 straightforward.

If D , then we say that the formula ϕ is a deductive consequence of the set of formulas Γ , and

that the argument \langle, \rangle is deductively valid. A formula ϕ is a logical theorem, or a deductive logical truth, if $D \vdash \phi$. That is, ϕ is a logical theorem if it is a deductive consequence of the empty set. A set Γ of formulas is consistent if there is no formula ϕ such that $D \vdash \phi$ and $D \vdash \neg \phi$. That is, a set is consistent if it does not entail a pair of contradictory opposite formulas.

Theorem 12. A set Γ is consistent if and only if there is a formula ϕ such that it is not the case that $D \vdash \phi$ and $D \vdash \neg \phi$.

Proof: Suppose that Γ is consistent and let ϕ be any formula. Then either it is not the case that $D \vdash \phi$ or it is not the case that $D \vdash \neg \phi$. For the converse, suppose that Γ is inconsistent and let ϕ be any formula. We have that there is a formula ψ such that both $D \vdash \psi$ and $D \vdash \neg \psi$. By ex falso quodlibet (Theorem 10), $D \vdash \phi$ and $D \vdash \neg \phi$.

Define a set Γ of formulas of the language L_K to be maximally consistent if Γ is consistent and for every formula ϕ of L_K , if ϕ is not in Γ , then Γ, ϕ is inconsistent. In other words, Γ is maximally consistent if Γ is consistent, and adding any formula in the language not already in Γ renders it inconsistent. Notice that if Γ is maximally consistent then $\Gamma \vdash \phi$ if and only if ϕ is in Γ .

Theorem 13. The Lindenbaum Lemma. Let Γ be any consistent set of formulas of L_K . Then there is a set Σ of formulas of L_K such that $\Gamma \subseteq \Sigma$ and Σ is maximally consistent.

Proof: Although this theorem holds in general, we assume here that the set K of non-logical terminology is either finite or denumerably infinite (i.e., the size of the natural numbers, usually called \aleph_0). It follows that there is an enumeration $\phi_0, \phi_1, \phi_2, \dots$ of the formulas of L_K , such that every formula of L_K eventually occurs in the list. Define a sequence of sets of formulas, by recursion, as follows: Σ_0 is Γ ; for each natural number n , if Σ_n, ϕ_n is consistent, then let $\Sigma_{n+1} = \Sigma_n, \phi_n$. Otherwise, let $\Sigma_{n+1} = \Sigma_n$. Let Σ be the union of all of the sets Σ_n . Intuitively, the idea is to go through the formulas of L_K , throwing each one into Σ if doing so produces a consistent set. Notice that each Σ_n is consistent. Suppose that Σ is inconsistent. Then there is a formula ϕ such that $\Sigma \vdash \phi$ and $\Sigma \vdash \neg \phi$. By Theorem 9 and Weakening (Theorem 8), there is finite subset Γ' of Σ such that $\Gamma' \vdash \phi$ and $\Gamma' \vdash \neg \phi$. Because Γ' is finite, there is a natural number n such that every member of Γ' is in Σ_n . So, by Weakening again, $\Sigma_n \vdash \phi$ and $\Sigma_n \vdash \neg \phi$. So Σ_n is inconsistent, which contradicts the construction. So Σ is consistent. Now suppose that a formula ψ is not in Σ . We have to show that Σ, ψ is inconsistent. The formula ψ must occur in the aforementioned list of formulas; say that it is ϕ_m . Since ϕ_m is not in Σ , then it is not in Σ_{m+1} . This happens only if Σ_m, ϕ_m is inconsistent. So a pair of contradictory opposites can be deduced from Σ_m, ϕ_m . By Weakening, a pair of contradictory opposites can be deduced from Σ, ψ . So Σ, ψ is inconsistent. Thus, Σ is maximally consistent.

Notice that this proof uses a principle corresponding to the law of excluded middle. In the construction of ω , we assumed that, at each stage, either n is consistent or it is not. Intuitionists, who demur from excluded middle, do not accept the Lindenbaum lemma (see Shapiro [1988]).

4. Semantics

Let K be a set of non-logical terminology. An interpretation for the language $1K=$ is a structure $M = \langle d, I \rangle$, where d is a non-empty set, called the domain-of-discourse, or simply the domain, of the interpretation, and I is an interpretation function. Informally, the domain is what we interpret the language $1K=$ to be about. It is what the variables range over. The interpretation function assigns appropriate extensions to the non-logical terms. In particular,

If c is a constant in K , then $I(c)$ is a member of the domain d .

If P_0 is a zero-place predicate letter in K , then $I(P_0)$ is a truth value, either truth or falsehood.

If Q_1 is a one-place predicate letter in K , then $I(Q)$ is a subset of d . Intuitively, $I(Q)$ is the set of members of the domain that the predicate Q holds of. If Q represents "red", then $I(Q)$ might be the red members of the domain.

If R_2 is a two-place predicate letter in K , then $I(R)$ is a set of ordered pairs of members of d . Intuitively, $I(R)$ is the set of pairs of members of the domain that the relation R holds between. If R represents "love", then $I(R)$ might consist of the pairs $\langle a, b \rangle$, such that a loves b .

In general, if S_n is an n -place predicate letter in K , then $I(S)$ is a set of ordered n -tuples of members of d .

Define s to be a variable-assignment, or simply an assignment, on an interpretation M , if s is a function from the variables to the domain d of M . The role of variable-assignments is to assign denotations to the free variables of open formulas. (In a sense, the quantifiers determine the "meaning" of the bound variables.) Logical systems that dispense with free variables do not need variable-assignments, but some other device is employed.

We now define a relation of satisfaction between interpretations, variable-assignments, and formulas of $1K=$. If ϕ is a formula of $1K=$, M is an interpretation for $1K=$, and s is a variable-assignment on M , then we write $M, s \models \phi$ for M satisfies ϕ under the assignment s . The

idea is that M, s is an analogue of " comes out true when interpreted as in M via s ".

Let t be a term of $1K=$. We define the denotation of t in M under s , in terms of the interpretation function and variable-assignment:

If c is a constant, then $DM, s(c)$ is $I(c)$, and if v is a variable, then $DM, s(v)$ is $s(v)$.

That is, the interpretation M assigns denotations to the constants, while the variable-assignment assigns denotations to the (free) variables. If the language contained function symbols, the denotation function would be defined by recursion.

We now proceed by recursion on the complexity of the formulas of $1K=$.

If t_1 and t_2 are terms, then $M, s \models t_1=t_2$ if and only if $DM, s(t_1)$ is the same as $DM, s(t_2)$.

This is about as straightforward as it gets. An identity $t_1=t_2$ comes out true if and only if the terms t_1 and t_2 denote the same thing.

If P_0 is a zero-place predicate letter in K , then $M, s \models P$ if and only if $I(P)$ is truth.

If S_n is an n -place predicate letter in K and t_1, \dots, t_n are terms, then $M, s \models S t_1 \dots t_n$ if and only if the n -tuple $\langle DM, s(t_1), \dots, DM, s(t_n) \rangle$ is in $I(S)$.

This takes care of the atomic formulas. We now proceed to the compound formulas of the language, following the meanings of the English counterparts of the logical terminology.

$M, s \models \neg \phi$ if and only if it is not the case that $M, s \models \phi$.

$M, s \models (\phi \ \& \ \psi)$ if and only if both $M, s \models \phi$ and $M, s \models \psi$.

$M, s \models (\phi \ \vee \ \psi)$ if and only if either $M, s \models \phi$ or $M, s \models \psi$.

$M, s \models \neg \phi$ if and only if either it is not the case that $M, s \models \phi$, or $M, s \models \phi$.

$M, s \models \forall v \phi$ if and only if $M, s \models \phi$, for every assignment s' that agrees with s except possibly at the variable v .

The idea here is that $\forall v \phi$ comes out true if and only if ϕ comes out true no matter what is assigned to the variable v . The final clause is similar.

$M, s \models \exists v \phi$ if and only if $M, s \models \phi$, for some assignment s' that agrees with s except possibly at the variable v .

So v comes out true if there is an assignment to v that makes ϕ true.

Theorem 6, unique readability, assures us that this definition is coherent. At each stage in breaking down a formula, there is exactly one clause to be applied, and so we never get contradictory verdicts concerning satisfaction.

As indicated, the role of variable-assignments is to give denotations to the free variables. We now show that variable-assignments play no other role.

Theorem 14. For any formula ϕ , if s_1 and s_2 agree on the free variables in ϕ , then $M, s_1 \models \phi$ if and only if $M, s_2 \models \phi$.

Proof: We proceed by induction on the complexity of the formula ϕ . The theorem clearly holds if ϕ is atomic, since in those cases only the values of the variable-assignments at the variables in ϕ figure in the definition. Assume, then, that the theorem holds for all formulas less complex than ϕ . And suppose that s_1 and s_2 agree on the free variables of ϕ . Assume, first, that ϕ is a negation, $\neg \psi$. Then, by the induction hypothesis, $M, s_1 \models \psi$ if and only if $M, s_2 \models \psi$. So, by the clause for negation, $M, s_1 \models \neg \psi$ if and only if $M, s_2 \not\models \neg \psi$. The cases where the main connective in ϕ is a binary connectives are also straightforward. Suppose that ϕ is $\forall v \psi$, and that $M, s_1 \models \forall v \psi$. Then there is an assignment s_1 that agrees with s_1 except possibly at v such that $M, s_1 \models \psi$. Let s_2 be the assignment that agrees with s_2 on the free variables not in ψ and agrees with s_1 on the others. Then, by the induction hypothesis, $M, s_2 \models \psi$. Notice that s_2 agrees with s_2 on every variable except possibly v . So $M, s_2 \models \forall v \psi$. The converse is the same, and the case where ϕ begins with a universal quantifier is similar.

Recall that a sentence is a formula with no free variables. So by Theorem 14, if ϕ is a sentence, and s_1, s_2 , are any two variable-assignments, then $M, s_1 \models \phi$ if and only if $M, s_2 \models \phi$. So we can just write $M \models \phi$ if $M, s \models \phi$ for some, or all, variable-assignments s .

Suppose that K_1, K_2 are two sets of non-logical terms. If $M = \langle d, I \rangle$ is an interpretation of L_{K_1} , then we define the restriction of M to L_{K_2} be the interpretation $M|_{K_2} = \langle d, I|_{K_2} \rangle$ such that $I|_{K_2}$ is the restriction of I to K_2 . That is, M and $M|_{K_2}$ have the same domain and agree on the non-logical terminology in K_2 . A straightforward induction establishes the following:

Theorem 15. If $M|_{K_2}$ is the restriction of M to L_{K_2} , then for every formula ϕ of L_{K_2} , if s is any variable-assignment, $M|_{K_2}, s \models \phi$ if and only if $M, s \models \phi$.

Theorem 16. If two interpretations M_1, M_2 have the same domain and agree on the non-logical terminology of a formula ϕ , then if s is any variable-assignment, $M_1, s \models \phi$ if and only if $M_2, s \models \phi$.

In short, the satisfaction of a formula ϕ only depends on the domain of discourse, the interpretation of the non-logical terminology in \mathcal{L} , and the assignments to the free variables in ϕ .

We say that an argument $\langle \Gamma, \phi \rangle$ is semantically valid, or just valid, written $\Gamma \models \phi$, if for every interpretation M of the language and any variable-assignment s on M , if $M, s \models \Gamma$, for every member ψ of Γ , then $M, s \models \psi$. If $\Gamma \models \phi$, we also say that ϕ is a logical consequence, or semantic consequence, or model-theoretic consequence of Γ . The definition corresponds to the informal idea that an argument is valid if it is not possible for its premises to all be true and its conclusion false. Our definition of logical consequence also sanctions the common thesis that a valid argument is truth-preserving--to the extent that satisfaction represents truth. Officially, an argument in $1K=$ is valid if its conclusion comes out true under every interpretation of the language in which the premises are true. Validity is the model-theoretic counterpart to deducibility.

A formula ϕ is logically true, or valid, if $M, s \models \phi$, for every interpretation M and assignment s . A formula is logically true if and only if it is a consequence of the empty set. If ϕ is logically true, then for any set Γ of formulas, $\Gamma \models \phi$. Logical truth is the model-theoretic counterpart of theoremhood.

A formula ϕ is satisfiable if there is an interpretation M and a variable-assignment s on M such that $M, s \models \phi$. That is, ϕ is satisfiable if there is an interpretation and assignment that satisfies it. A set Γ of formulas is satisfiable if there is an interpretation M and a variable-assignment s on M such that $M, s \models \psi$, for every formula ψ in Γ . If Γ is a set of sentences and if $M \models \phi$ for each sentence ϕ in Γ , then we say that M is a model of Γ . So a set of sentences is satisfiable if it has a model. Satisfiability is the model-theoretic counterpart to consistency.

Notice that $\Gamma \models \phi$ if and only if the set $\Gamma, \neg \phi$ is not satisfiable. It follows that if a set Γ is not satisfiable, then if ϕ is any formula, $\Gamma \models \phi$. This is a model-theoretic counterpart to *ex falso quodlibet* (see Theorem 10). We have the following, as an analogue to Theorem 12:

Theorem 17. Let Γ be a set of formulas. The following are equivalent: (a) Γ is satisfiable; (b) there is no formula ϕ such that both $\Gamma \models \phi$ and $\neg \phi$; (c) there is some formula ϕ such that it is not the case that $\Gamma \models \phi$.

Proof: (a)(b): Suppose that Γ is satisfiable and let ϕ be any formula. There is an

interpretation M and assignment s such that $M, s \models \phi$ for every member ϕ of Γ . By the clause for negations, we cannot have both $M, s \models \phi$ and $M, s \models \neg\phi$. So either $\langle \phi, \neg\phi \rangle$ is not valid or else $\langle \phi, \neg\phi \rangle$ is not valid. (b)(c): This is immediate. (c)(a): Suppose that it is not the case that Γ is satisfiable. Then there is an interpretation M and an assignment s such that $M, s \models \phi$ for every formula ϕ in Γ and it is not the case that $M, s \models \psi$. A fortiori, M, s satisfies every member of Γ , and so Γ is satisfiable.

5. Meta-theory

We now present some results that relate the deductive notions to their model-theoretic counterparts. The first one is probably the most straightforward. We motivated both the various rules of the deductive system D and the various clauses in the definition of satisfaction in terms of the meaning of the English counterparts to the logical terminology. So one would expect that an argument is deducible, or deductively valid, only if it is semantically valid.

Theorem 18. Soundness. For any formula ϕ and set Γ of formulas, if $\Gamma \vdash \phi$, then $\Gamma \models \phi$.

Proof: We proceed by induction on the number of clauses used to establish $\Gamma \vdash \phi$. So let n be a natural number, and assume that the theorem holds for any argument established as deductively valid with fewer than n steps. And suppose that $\Gamma \vdash \phi$ was established using exactly n steps. If the last rule applied was $(=I)$ then ϕ is a formula in the form $t=t$, and so ϕ is logically true. A fortiori, $\Gamma \models \phi$. If the last rule applied was (As) , then ϕ is a member of Γ , and so of course any interpretation and assignment that satisfies every member of Γ also satisfies ϕ . Suppose the last rule applied is $(\&I)$. So ϕ has the form $(\& \phi_1 \phi_2)$, and we have $\Gamma \vdash \phi_1$ and $\Gamma \vdash \phi_2$, with $\phi_1 = \phi_1$, $\phi_2 = \phi_2$. The induction hypothesis gives us $\Gamma \models \phi_1$ and $\Gamma \models \phi_2$. Suppose that M, s satisfies every member of Γ . Then M, s satisfies every member of ϕ_1 , and so M, s satisfies ϕ_1 . Similarly, M, s satisfies every member of ϕ_2 , and so M, s satisfies ϕ_2 . Thus, by the clause for "&" in the definition of satisfaction, M, s satisfies ϕ . So $\Gamma \models \phi$. Suppose the last clause applied was (E) . So we have $\Gamma \vdash \phi_1 \vee \phi_2$ and $\Gamma \vdash \phi$, where $\phi = \phi_1 \vee \phi_2$, and ϕ does not occur free in ϕ_1 , nor in any member of Γ . By the induction hypothesis, we have $\Gamma \models \phi_1 \vee \phi_2$ and $\Gamma \models \phi$. Let M be an interpretation and s an assignment such that M, s satisfies every member of Γ . Then M, s satisfies every member of ϕ_1 , and so $M, s \models \phi_1$. So there is an assignment s' that agrees with s on every variable except possibly ϕ such that $M, s' \models \phi$. We have that M, s satisfies every member of Γ . Since ϕ does not occur free in any member of Γ , and s' agrees with s on everything else, we have that M, s' satisfies every member of Γ , by Theorem 14. So $M, s' \models \phi$. Since ϕ does not occur free in Γ , and s' agrees with s on everything else, we have that $M, s \models \phi$, also by Theorem 14. So, in this case, $\Gamma \models \phi$. Notice the role of the restrictions on (E) here. The other cases are about as straightforward.

Corollary 19. Let Σ be a set of formulas. If Σ is satisfiable, then Σ is consistent.

Proof: Suppose that Σ is satisfiable. So let M be an interpretation and s an assignment such that M, s satisfies every member of Σ . Assume that Σ is inconsistent. Then there is a formula ϕ such that $\phi \in \Sigma$ and $\neg \phi \in \Sigma$. By soundness (Theorem 18), $M, s \models \phi$ and $M, s \models \neg \phi$. So we have that $M, s \models \phi$ and $M, s \models \neg \phi$. But this is impossible, given the clause for negation in the definition of satisfaction.

Even though the deductive system D and the model-theoretic semantics were developed with the meanings of the logical terminology in mind, one should not automatically expect the converse to soundness (or Corollary 19) to hold. For all we know so far, we may not have included enough rules of inference to deduce every valid argument. The converses to soundness and Corollary 19 are among the most interesting results in contemporary mathematical logic. We begin with the latter.

Theorem 20. Completeness. Gödel [1930]. Let Σ be a set of formulas. If Σ is consistent, then Σ is satisfiable.

Proof: The proof of completeness is rather complex. We only sketch it here. Let Σ be a consistent set of formulas of $1K=$. Again, we assume for simplicity that the set K of non-logical terminology is either finite or countably infinite (although the theorem holds even if K is uncountable). The task at hand is to find an interpretation M and a variable-assignment s on M , such that M, s satisfies every member of Σ . Consider the language obtained from $1K=$ by adding a denumerably infinite stock of new individual constants c_0, c_1, \dots . We stipulate that the constants, c_0, c_1, \dots , are all different from each other and none of them occur in K . We build an interpretation of the language from the language itself, using some of the constants as members of the domain of discourse. Let ϕ_0, ϕ_1, \dots be an enumeration of the formulas of the expanded language, so that each formula occurs in the list eventually. Let x be any variable, and define a sequence ϕ_0, ϕ_1, \dots of sets of formulas (of the expanded language) by recursion as follows: $\phi_0 = \Sigma$; and $\phi_{n+1} = \phi_n, (x_n \neg \phi_n(x|c_i))$, where c_i is the first constant in the above list that does not occur in ϕ_n or in any member of ϕ_n . The underlying idea here is that if x_n is true, then c_i is to be one such x . Let Σ' be the union of the sets ϕ_n .

I sketch a proof that Σ' is consistent. Suppose that Σ' is inconsistent. By Theorem 9, there is a finite subset of Σ' that is inconsistent, and so one of the sets ϕ_n is inconsistent. By hypothesis, $\phi_0 = \Sigma$ is consistent. Let n be the smallest number such that ϕ_n is consistent, but $\phi_{n+1} = \phi_n, (x_n \neg \phi_n(x|c_i))$ is inconsistent. By (\neg I), we have that

(1) $n \neg(xn \ n(x|ci))$.

By ex falso quodlibet (Theorem 10), $n, \neg xn, xn \ n(x|ci)$. So by (I), $n, \neg xn \ (xn \ n(x|ci))$. From this and (1), we have $n \neg\neg xn$, by (\neg I), and by (DNE) we have

(2) $n \ xn$.

By (As), $n, n(x|ci), xn \ n(x|ci)$. So by (I), $n, n(x|ci) \ (xn \ n(x|ci))$. From this and (1), we have $n \neg n(x|ci)$, by (\neg I). Let v be a variable that does not occur (free or bound) in n or in any member of n . By uniform substitution of v for ci , we can turn the derivation of $n\neg n(x|ci)$ into $n\neg n(x|v)$. By (I), we have

(3) $n \ v\neg n(x|v)$.

By (As) we have $\{v\neg n(x|v), n\} \ n$ and by (E) we have $\{v\neg n(x|v), n\} \ \neg n$. So $\{v\neg n(x|v), n\}$ is inconsistent. Let β be any sentence of the language (so that β has no free variables). By ex falso quodlibet (Theorem 10), we have that $\{v\neg n(x|v), n\} \ \beta$ and $\{v\neg n(x|v), n\} \ \neg\beta$. So with (2), we have that $n, v\neg n(x|v)$ and $n, v\neg n(x|v) \ \neg\beta$, by (E). By Cut (Theorem 11), n and $n \ \neg\beta$. So n is inconsistent, contradicting the assumption. So β is consistent.

Applying the Lindenbaum Lemma (Theorem 13), let β be a maximally consistent set of sentences (of the expanded language) that contains β . So, of course, β contains β . We define an interpretation M , and a variable-assignment s on M , such that M, s satisfies every member of β .

If we did not have a sign for identity in the language, we would let the domain of M be the collection of new constants $\{c_0, c_1, \dots\}$. But as it is, there may be a sentence in the form $ci=cj$, with $i < j$, in β . If so, we cannot have both ci and cj in the domain of the interpretation. So we define the domain d of M to be the set $\{ci \mid \text{there is no } j < i \text{ such that } ci=cj \text{ is in } \beta\}$. In other words, a constant ci is in the domain of M if β does not declare it to be identical to an earlier constant in the list. Notice that for each new constant ci , there is exactly one ji such that cj is in d and the sentence $ci=cj$ is in β .

We now define the interpretation function I . Let a be any constant in the expanded language. By ($=$ I) and (I), $x \ x=a$, and so $x \ x=a \ \beta$. By the construction of d , there is a sentence in the form $(x \ x=a \ ci=a)$ in β . We have that $ci=a$ is in β . As above, there is exactly one cj in d such that $ci=cj$ is in β . Let $I(a)=cj$. Notice that if ci is a constant in the domain d , then $I(ci)=ci$. That is each ci in d denotes itself.

Let P be a one-place predicate letter in K . Let $I(P)$ be the set of constants $\{ci \mid ci \text{ is in } d \text{ and the formula } Pc \text{ is in } \beta\}$. Let R be a binary predicate letter in K . Let $I(R)$ be the set of

pairs of constants $\{ \langle c_i, c_j \rangle \mid c_i \text{ is in } d, c_j \text{ is in } d, \text{ and the formula } R c_i c_j \text{ is in } \Sigma \}$. Three-place predicates, etc. are interpreted similarly. In effect, I interprets the non-logical terminology as they are in \mathcal{M} .

The variable-assignment is similar. If v is a variable, then $s(v) = c_i$, where c_i is the first constant in d such that $c_i = v$ is in Σ .

The final item in this proof is a tedious lemma that for every formula ϕ in the expanded language, $\mathcal{M}, s \models \phi$ if and only if ϕ is in Σ . This proceeds by induction on the complexity of ϕ . The case where ϕ is atomic follows from the definitions of \mathcal{M} (i.e., the domain d and the interpretation function I) and the variable-assignment s . The other cases follow from the various clauses in the definition of satisfaction.

Since Σ , we have that \mathcal{M}, s satisfies every member of Σ . By Theorem 15, the restriction of \mathcal{M} to the original language \mathcal{L} and s also satisfies every member of Σ . Thus Σ is satisfiable.

A converse to Soundness (Theorem 18) is a straightforward corollary:

Theorem 21. For any formula ϕ and set Σ of formulas, if Σ, ϕ , then $D \models \phi$.

Proof: Suppose that Σ, ϕ . Then there is no interpretation \mathcal{M} and assignment s such that \mathcal{M}, s satisfies every member of Σ but does not satisfy ϕ . So the set $\Sigma, \neg \phi$ is not satisfiable. By Completeness (Theorem 20), $\Sigma, \neg \phi$ is inconsistent. So there is a formula ψ such that Σ, ψ and $\Sigma, \neg \psi$. By (\neg I), $\neg \neg$, and by (DNE).

Our next item is a corollary of Theorem 9, Soundness (Theorem 18), and Completeness:

Corollary 22. Compactness. A set Σ of formulas is satisfiable if and only if every finite subset of Σ is satisfiable.

Proof: If \mathcal{M}, s satisfies every member of Σ , then \mathcal{M}, s satisfies every member of each finite subset of Σ . For the converse, suppose that Σ is not satisfiable. Then we show that some finite subset of Σ is not satisfiable. By Completeness (Theorem 20), Σ is inconsistent. By Theorem 9 (and Weakening), there is a finite subset Σ_0 such that Σ_0 is inconsistent. By Corollary 19, Σ_0 is not satisfiable.

Soundness and completeness together entail that an argument is deducible if and only if it is valid, and a set of formulas is consistent if and only if it is satisfiable. So we can go back and forth between model-theoretic and proof-theoretic notions, transferring

properties of one to the other. Compactness holds in the model theory because all derivations use only a finite number of premises.

Recall that in the proof of Completeness (Theorem 20), we made the simplifying assumption that the set K of non-logical constants is either finite or denumerably infinite. The interpretation we produced was itself either finite or denumerably infinite. Thus, we have the following:

Corollary 23. Löwenheim-Skolem Theorem. Let Σ be a satisfiable set of sentences of the language \mathcal{L}_K . If K is either finite or denumerably infinite, then Σ has a model whose domain is either finite or denumerably infinite.

In general, let Σ be a satisfiable set of sentences of \mathcal{L}_K , and let κ be the larger of the size of K and denumerably infinite. Then Σ has a model whose domain is at most size κ .

There is a stronger version of Corollary 23. Let $M_1 = \langle d_1, I_1 \rangle$ and $M_2 = \langle d_2, I_2 \rangle$ be interpretations of the language \mathcal{L}_K . Define M_1 to be a submodel of M_2 if $d_1 \subseteq d_2$, $I_1(c) = I_2(c)$ for each constant c , and I_1 is the restriction of I_2 to d_1 . For example, if R is a binary relation letter in K , then for all a, b in d_1 , the pair $\langle a, b \rangle$ is in $I_1(R)$ if and only if $\langle a, b \rangle$ is in $I_2(R)$. If we had included function letters among the non-logical terminology, we would also require that d_1 be closed under their interpretations in M_2 . Notice that if M_1 is a submodel of M_2 , then any variable-assignment on M_1 is also a variable-assignment on M_2 .

Say that two interpretations $M_1 = \langle d_1, I_1 \rangle$, $M_2 = \langle d_2, I_2 \rangle$ are elementarily equivalent if one of them is a submodel of the other, and for any formula of the language and any variable-assignment s on the submodel, M_1, s if and only if M_2, s . Notice that if two interpretations are elementarily equivalent, then they satisfy the same sentences.

Theorem 25. Downward Löwenheim-Skolem Theorem. Let $M = \langle d, I \rangle$ be an interpretation of the language \mathcal{L}_K . Let d_1 be any subset of d , and let κ be the maximum of the size of K , the size of d_1 , and denumerably infinite. Then there is a submodel $M' = \langle d_1, I' \rangle$ of M such that (1) d_1 is not larger than κ , and (2) M and M' are elementarily equivalent. In particular, if the set K of non-logical terminology is either finite or denumerably infinite, then any interpretation has an elementarily equivalent submodel whose domain is either finite or denumerably infinite.

Proof: Like completeness, this proof is complex, and we rest content with a sketch. The downward Löwenheim-Skolem theorem invokes the axiom of choice, and indeed, is equivalent to the axiom of choice. So let C be a choice function on the powerset of d , so

that for each non-empty subset e of d , $C(e)$ is a member of e . We stipulate that if e is the empty set, then $C(e)$ is $C(d)$.

Let s be a variable-assignment on M , let ϕ be a formula of $1K=$, and let v be a variable. Define the v -witness of ϕ over s , written $wv(\phi, s)$, as follows: Let q be the set of all elements c of d such that there is a variable-assignment s' on M that agrees with s on every variable except possibly v , such that M, s' and $s'(v)=c$. Then $wv(\phi, s) = C(q)$. Notice that if $M, s \models \phi$, then q is the set of elements of the domain that can go for v in ϕ . Indeed, $M, s \models \phi$ if and only if q is non-empty. So if $M, s \models \phi$, then $wv(\phi, s)$ (i.e., $C(q)$) is a chosen element of the domain that can go for v in ϕ . In a sense, it is a "witness" that verifies $M, s \models \phi$.

If e is a non-empty subset of the domain d , then define a variable-assignment s to be an e -assignment if for all variables u , $s(u)$ is in e . That is, s is an e -assignment if s assigns an element of e to each variable. Define $sk(e)$, the Skolem-hull of e , to be the set:

$$e \cup \{wv(\phi, s) \mid \phi \text{ is a formula in } 1K=, v \text{ is a variable, and } s \text{ is an } e\text{-assignment}\}.$$

That is, the Skolem-Hull of e is the set e together with every v -witness of every formula over every e -assignment. Roughly, the idea is to start with e and then throw in enough elements to make each existentially quantified formula true. But we cannot rest content with the Skolem-hull, however. Once we throw the "witnesses" into the domain, we need to deal with $sk(e)$ assignments. In effect, we need a set which is its own Skolem-hull, and also contains the given subset d_1 .

We define a sequence of non-empty sets e_0, e_1, \dots as follows: if the given subset d_1 of d is empty and there are no constants in K , then let e_0 be $C(d)$, the choice function applied to the entire domain; otherwise let e_0 be the union of d_1 and the denotations under I of the constants in K . For each natural number n , e_{n+1} is $sk(e_n)$. Finally, let d be the union of the sets e_n , and let I be the restriction of I to d . Our interpretation is $M = \langle d, I \rangle$.

Clearly, d_1 is a subset of d , and so M is a submodel of M . Let \aleph_n be the maximum of the size of K , the size of d_1 , and denumerably infinite. A calculation reveals that the size of d is at most \aleph_n , based on the fact that there are at most \aleph_n -many formulas, and thus, at most \aleph_n -many witnesses at each stage. Notice, incidentally, that this calculation relies on the fact that a denumerable union of sets of size at most \aleph_n is itself at most \aleph_n . This also relies on the axiom of choice.

The final item is to show that M is elementarily equivalent to M : For every formula ϕ and

every variable-assignment s on M ,

$M, s \models \phi$ if and only if $M, s \models \psi$.

We proceed by induction on the complexity of ϕ . If ϕ is atomic, then the definition of satisfaction entails the equivalence. So let ϕ be non-atomic, and assume that $M, s \models \psi$ if and only if $M, s \models \theta$, for all assignments s on M and all formulas ψ less complex than ϕ . Let s be any such assignment. If the main connective of ϕ is the negation sign or a binary connective, then the induction hypothesis entails that $M, s \models \phi$ if and only if $M, s \models \theta$. The remaining cases are those in which ϕ begins with a quantifier, i.e., ϕ is either $\forall v \psi$ or $\exists v \psi$. Suppose that $M, s \models \forall v \psi$. Then there is a variable-assignment s_1 that agrees with s except possibly at v such that $M, s_1 \models \psi$. By the induction hypothesis, $M, s \models \psi$ and so $M, s_1 \models \psi$. The converse is a bit tricky, and amounts to showing that the Skolem-hull of d is d . Assume that $M, s \models \exists v \psi$. We are given that s is a variable-assignment on d . Since there are only finitely many free-variables in ψ , let n be any natural number such that for all variables u that occur free in ψ , $s(u)$ is in e_n . Let s_1 be an e_n -assignment that agrees with s on all of the free variables in ψ . Then, by Theorem 14, $M, s_1 \models \psi$. Let c be $wv(\psi, s_1)$, the v -witness of ψ over s_1 . Notice that c is in e_{n+1} and so c is in d . Let s_1 agree with s , except possibly at v , and let $s_1(v) = c$. So s_1 is a variable-assignment on M . By the definition of the witness function, $M, s_1 \models \psi$. By the induction hypothesis, $M, s_1 \models \psi$, and so $M, s \models \exists v \psi$. By Theorem 14, $M, s \models \psi$. The final case, where ϕ has the form $\forall v \psi$, is similar.

Another corollary to Compactness (Corollary 22) is the opposite of the Löwenheim-Skolem theorem:

Theorem 26. Upward Löwenheim-Skolem Theorem. Let Σ be any set of formulas of $1K=$, such that for each natural number n , there is an interpretation $M_n = \langle d_n, I_n \rangle$, and an assignment s_n on M_n , such that d_n has at least n elements, and M_n, s_n satisfies every member of Σ . In other words, Σ is satisfiable and there is no finite upper bound to the size of the interpretations that satisfy every member of Σ . Then for any infinite cardinal \aleph_α , there is an interpretation $M = \langle d, I \rangle$ and assignment s on M , such that the size of d is at least \aleph_α and M, s satisfies every member of Σ . In particular, if Σ is a set of sentences, then it has arbitrarily large models.

Proof: Add a collection of new constants $\{c_i \mid i < \aleph_\alpha\}$, of size \aleph_α , to the language, so that if c is a constant in K , then c is different from c_i , and if $i < j$, then c_i is a different constant than c_j . Consider the set of formulas Σ' consisting of Σ together with the set $\{c_i = c_j \mid i < j < \aleph_\alpha\}$. That is,

consists of Σ together with statements to the effect that any two different new constants denote different objects. Let Σ_0 be any finite subset of Σ , and let m be the number of new constants that occur in Σ_0 . Then expand the interpretation M_{Σ_0} to an interpretation M_m of the new language, by interpreting each of the new constants in Σ_0 as a different member of the domain d_m . By hypothesis, there are enough members of d_m to do this. One can interpret the other new constants at will. So M_m is a restriction of M_{Σ_0} . By hypothesis (and Theorem 15), M_m, s_m satisfies every member of Σ_0 . Also M_m, s_m satisfies the members of $\{\neg c=c \mid c \text{ is a new constant}\}$ that are in Σ_0 . So M_m, s_m satisfies every member of Σ_0 . By compactness, there is an interpretation $M = \langle d, I \rangle$ and an assignment s on M such that M, s satisfies every member of Σ . Since Σ contains every member of $\{\neg c=c \mid c \text{ is a new constant}\}$, the domain d of M must be of size at least \aleph_m , since each of the new constants must have a different denotation. By Theorem 15, the restriction of M to the original language $1K=$ satisfies every member of Σ_0 , with the variable-assignment s .

The proofs of the downward and upward Löwenheim-Skolem theorems can be combined to show that for any satisfiable set Σ of sentences, if there is no finite bound on the models of Σ , then for any infinite cardinal \aleph_m , there is a model of Σ whose domain has size exactly \aleph_m . Moreover, if M is any interpretation whose domain is infinite, then for any infinite cardinal \aleph_m , there is an interpretation M' whose domain has size exactly \aleph_m such that M and M' are elementarily equivalent.

These results indicate a weakness in the expressive resources of first-order languages like $1K=$. No satisfiable set of sentences can guarantee that its models are all denumerably infinite, nor can any satisfiable set of sentences guarantee that its models are uncountable. So in a sense, first-order languages cannot express the notion of "denumerably infinite", at least not in the model theory.

Let A be any set of sentences in a first-order language $1K=$, where K includes terminology for arithmetic, and assume that every member of A is true of the natural numbers. We can even let A be the set of all sentences in $1K=$ that are true of the natural numbers. Then A has uncountable models, indeed models of any infinite cardinality. Such interpretations are sometimes called unintended, or non-standard models of arithmetic. Let B be any set of first-order sentences that are true of the real numbers, and let C be any first-order axiomatization of set theory. Then if B and C are satisfiable (in infinite interpretations), then each of them has denumerably infinite models. That is, any first-order, satisfiable set theory or theory of the real numbers, has (unintended) models the size of the natural

numbers. This is despite the fact that a sentence (seemingly) stating that the universe is uncountable is provable in most set-theories. This situation, known as the Skolem paradox, has generated much discussion, but we must refer the reader elsewhere for a sample of it.

Bibliography

Cited Works

Anderson, A. and N. Belnap [1975], *Entailment: The logic of relevance and necessity I*, Princeton: Princeton University Press.

Anderson, A. and N. Belnap, and M. Dunn [1992], *Entailment: The logic of relevance and necessity II*, Princeton: Princeton University Press.

Corcoran, J. [1973], "Gaps between logical theory and mathematical practice", *The methodological unity of science*, ed. by M. Bunge, Dordrecht: Holland, D. Reidel, 23-50.

Church, A. [1956], *Introduction to mathematical logic*, Princeton: Princeton University Press. Classic textbook.

Davidson, D. [1984], *Inquiries into truth and interpretation*, Oxford: Clarendon Press.

Dummett, M. [1977], *Elements of intuitionism*, Oxford: Oxford University Press.

Gödel, K. [1930], "Die Vollständigkeit der Axiome des logischen Funktionenkalküls", *Montatshefte für Mathematik und Physik* 37, 349-360; translated as "The completeness of the axioms of the functional calculus of logic", in van Heijenoort [1967], 582-591.

Lycan, W. [1984], *Logical form in natural language*, Cambridge, Massachusetts: The MIT Press.

Montague, R. [1974], *Formal philosophy*, ed. by R. Thomason, New Haven: Yale University Press.

Quine, W. V. O. [1960], *Word and object*, Cambridge, Massachusetts: The MIT Press.

Quine, W. V. O. [1986], *Philosophy of logic*, second edition, Englewood Cl