Logic and Philosophical Logic

Miscellaneous Results and Remarks[1]

A.C. Paseau
alexander.paseau@philosophy.ox.ac.uk
Monday 14 January 2019

In today's class we'll go over some formal results. Some of these will be required or helpful for later classes; others will not be directly related to the remaining classes but should be part of every philosophical logician's toolkit. I assume that students have taken an intermediate logic course (or perhaps even just a demanding introductory course, like the Oxford one), in which they have gained familiarity with propositional and first-order logic (also known as predicate logic). In any case, I've included a review of first-order logic and its semantics in section 2 below, for completeness.

1 Languages

Let's begin by distinguishing sentences of (a) natural language, (b) an interpreted formal language, and (c) an uninterpreted formal language. Examples of natural-language sentences:

- Everyone is self-identical.
- If Joe is tall then Joe is tall.
- No even numbers are odd.

Some examples of interpreted formal-language sentences:

- $\forall x(x = x)$ with the following interpretation: the domain is the set of people, the identity sign is interpreted as identity, and the universal quantifier symbol as the universal quantifier over all people.
- $p \to p$ with the interpretation: p means that Joe is tall and the arrow is interpreted as the truth-functional (material) conditional.
- $\forall x(Ex \to \neg Ox)$ with the interpretation: the domain is the set of all numbers, the first predicate symbol is interpreted as 'is even', the second as 'is odd', the arrow as the material conditional, the negation symbol as truth-functional (sentential) negation, and the universal quantifier symbol ranges over all numbers.

Some examples of uninterpreted formal-language sentences:

- $\forall x(x = x)$

[1] Some of the material in section 4 draws on my forthcoming encyclopaedia article on the Compactness Theorem co-written with Robert Leek and some of the material in section 5 on lecture notes co-written with Dan Isaacson.

- $p \to p$
- $\forall x(Ex \to \neg Ox)$

Natural-language sentences and interpreted formal sentences are capable of having a truth-value. In contrast, the third batch of sentences are not interpreted, so can't be true or false, still less logically true or false. Do not confuse a logical truth, which must be interpreted/meaningful, with a validity, which is a formal sentence satisfied in all interpretations.

2 First-order logic

We give the vocabulary, grammar and semantics of first-order logic with identity, culminating in the definition of model-theoretic consequence for this logic.

Vocabulary
The vocabulary of first-order logic consists of the following:

- Symbols for propositional connectives: $\wedge, \vee, \to, \leftrightarrow, \neg$ (or any other expressively adequate set).
- Object variables: $x_0, x_1, x_2, \cdots$
- Quantifiers: $\forall$ and $\exists$ (or one defined in terms of the other).
- The identity symbol: $=$.
- Left and right parentheses: (,)
- Non-logical constants: $c_0, c_1, c_2, \cdots$
- Non-logical function constants: $f_0, f_1, f_2, \cdots$ (infinitely many of each adicity)
- Non-logical predicate constants: $R_0, R_1, R_2, \cdots$ (infinitely many of each adicity)

One can be more precise than we have been here about how many object variables, non-logical constants, function constants and predicate constants each first-order logic contains.

Grammar
The grammar of first-order logic is given by specifying the set of terms and well-formed formulas (wffs). Terms:

- An object variable or a constant is a term. Terms of this form are known as *atomic*.
- If ζ is an n-place function constant and $\tau_1, \cdots, \tau_n$ are terms, then $\zeta(\tau_1, \cdots, \tau_n)$ is a term.

Well-formed formulas (wffs) are specified as follows:

- If τ_1 and τ_2 are terms, then $\tau_1 = \tau_2$ is a wff.
- If π is an n-place predicate constant and $\tau_1, \cdots, \tau_n$ are terms, then $\pi\tau_1 \cdots \tau_n$ is a wff.
- If ϕ is a wff so is $\neg\phi$; if ϕ and ψ are wffs, so are $(\phi \wedge \psi)$, $(\phi \vee \psi)$, $(\phi \to \psi)$ and $(\phi \leftrightarrow \psi)$.
- If ϕ is a wff and α is an object variable then $\forall\alpha\phi$ and $\exists\alpha\phi$ are wffs.

We often drop brackets in wffs.

Model and Term Assignment
To define the notion of satisfaction, we first need to define the notions of a model and of a term assignment. A model $\mathfrak{M}$ consists of a set M known as the model's domain, and an interpretation (here denoted by the superscript $\mathfrak{M}$) of the constants, function constants and predicate constants given by:

- If α is a constant, then $\alpha^{\mathfrak{M}}$ is an element of M.
- If ζ is an n-place function constant then $\zeta^{\mathfrak{M}}$ is an n-place function on M.
- If π is an n-place predicate constant then $\pi^{\mathfrak{M}}$ is an n-place relation on M (i.e. a set of n-tuples of M).

A variable assignment g is a function with domain the set of variables such that:

- If α is a variable, $g(\alpha)$ is an element of M.

We denote by $g^+_{\mathfrak{M}}$ the term assignment uniquely determined by the model $\mathcal{M}$ and variable assignment g. By definition, $g^+_{\mathfrak{M}}$ is a function with domain the set of terms that satisfies these conditions:

- If α is a constant, then $g^+_{\mathfrak{M}}(\alpha) = \alpha^{\mathfrak{M}}$.
- If α is a variable, then $g^+_{\mathfrak{M}}(\alpha) = g(\alpha)$.
- For any complex term $\tau = \zeta(\tau_1, \cdots, \tau_n)$, where ζ is a function constant,

$$g^+_{\mathfrak{M}}(\tau) = \zeta^{\mathfrak{M}}((g^+_{\mathfrak{M}}(\tau_1), \cdots, g^+_{\mathfrak{M}}(\tau_n))$$

Satisfaction
The notion of satisfaction is a ternary relation between models $\mathfrak{M}$, variable assignments g, and formulas ϕ. It is defined recursively on the complexity of formulas, as follows.[2] For any n-place predicate constant π and terms $\tau_1, \cdots, \tau_n$:

$$Sat(\mathfrak{M}, g, \pi\tau_1 \cdots \tau_n) \text{ iff } \langle g^+_{\mathfrak{M}}(\tau_1), \cdots, g^+_{\mathfrak{M}}(\tau_n)\rangle \in \pi^{\mathfrak{M}}$$

[2]For brevity, we omit the condition from the definitions' right-hand sides that $\mathfrak{M}$ is a model of some first-order language, g is a variable assignment, and ϕ is a wff in $\mathfrak{M}$'s language.

$Sat(\mathfrak{M}, g, \tau_1 = \tau_2)$ iff $g^+_{\mathfrak{M}}(\tau_1) = g^+_{\mathfrak{M}}(\tau_2)$

(As usual, we use '$=$' as a symbol in both the object language and the metalanguage.) For any wffs ϕ and ψ and variable α (dropping brackets in some wffs):

- $Sat(\mathfrak{M}, g, \neg\phi)$ iff it's not the case that $Sat(\mathfrak{M}, g, \phi)$.
- $Sat(\mathfrak{M}, g, \phi \wedge \psi)$ iff $Sat(\mathfrak{M}, g, \phi)$ and $Sat(\mathfrak{M}, g, \psi)$.
- $Sat(\mathfrak{M}, g, \phi \vee \psi)$ iff $Sat(\mathfrak{M}, g, \phi)$ or $Sat(\mathfrak{M}, g, \psi)$ (or both).
- $Sat(\mathfrak{M}, g, \phi \to \psi)$ iff it's not the case that $Sat(\mathfrak{M}, g, \phi)$ or $Sat(\mathfrak{M}, g, \psi)$ (or both).
- $Sat(\mathfrak{M}, g, \phi \leftrightarrow \psi)$ iff $Sat(\mathfrak{M}, g, \phi)$ and $Sat(\mathfrak{M}, g, \psi)$ or it's not the case that $Sat(\mathfrak{M}, g, \phi)$ and it's not the case that $Sat(\mathfrak{M}, g, \psi)$.
- $Sat(\mathfrak{M}, g, \forall\alpha\phi)$ iff $Sat(\mathfrak{M}, h, \phi)$ for any h that is a variable assignment agreeing with g on all variables, with the possible exception of variable α.

$Sat(\mathfrak{M}, g, \phi)$ is often written as $(\mathfrak{M}, g) \vDash \phi$, or $\mathfrak{M} \vDash_g \phi$. $Sat(\mathfrak{M}, g, \Gamma)$ means that $Sat(\mathfrak{M}, g, \gamma)$ for every γ in Γ.

Model-Theoretic Consequence
If Γ is a set of first-order formulas and δ is a first-order formula, we write

$$\Gamma \vDash \delta$$

to mean that for any model $\mathfrak{M}$ and variable assignment g, if $Sat(\mathfrak{M}, g, \Gamma)$ then $Sat(\mathfrak{M}, g, \delta)$. We are often interested in the case in which Γ is a set of first-order sentences (formulas with no free variables) and δ is a first-order sentence.

Remarks:

(i) It is unfortunate that the symbol $\vDash$ is used ambiguously in logic, for both the relation of satisfaction and that of model-theoretic consequence.

(ii) A further ambiguity is that $\vDash$ in either sense is a logic-relative notion: strictly speaking, we should write $\vDash_{\mathcal{L}}$. Usually, context makes it clear which logic is in question.

(iii) The specification just given is usually called 'a semantics' for first-order logic. Its relation to a semantics of natural language—an account of the meanings of natural-language sentences—is a delicate question. Shared terminology should not prejudice the answer.

(iv) We often say that

ϕ is true in an interpretation $\langle\mathfrak{M}, g\rangle$

instead of

$$Sat(\mathfrak{M}, g, \phi)$$

or

$$(\mathfrak{M}, g) \vDash \phi$$

When ϕ is a sentence, we say more simply that ϕ is true or false in $\mathfrak{M}$, since in the former case ϕ is true in $\langle\mathfrak{M}, g\rangle$ for all variable assignments g and in the latter case ϕ is false in $\langle\mathfrak{M}, g\rangle$ for all variable assignments g. Since satisfaction is a technical notion, so is the notion of 'truth in'. Its relation to truth outright is also a moot question, not to be settled by common terminology.

3 Second-order logic

3.1 Second-order logic

Second-order logic is the logic obtained by adding function and relation variables to first-order logic. I list *supplementary* clauses to the first-order ones previously specified.

Vocabulary
The extra vocabulary of second-order logic consists of:

- Function variables: $f_0, f_1, f_2, \cdots$ (infinitely many of each adicity)
- Predicate variables: $X_0, X_1, X_2, \cdots$ (infinitely many of each adicity)

Grammar
The supplementary clauses to the first-order ones are:

- If Z is an n-place function variable and $\tau_1, \cdots, \tau_n$ are terms, then $Z(\tau_1, \cdots, \tau_n)$ is a term.
- If Π is an n-place predicate variable and $\tau_1, \cdots, \tau_n$ are terms, then $\Pi\tau_1 \cdots \tau_n$ is a wff.
- If ϕ is a wff and Ξ is a predicate or function variable then $\forall\Xi\phi$ and $\exists\Xi\phi$ are wffs.

Model and Term Assignment
A second-order model $\mathcal{M}$ is identical to a first-order model. A second-order variable assignment g is a first-order variable assignment with the extra properties:

- If Z is an n-place function variable then $g(Z)$ is an n-place function on M.
- If Π is an n-place predicate variable then $g(\Pi)$ is an n-place relation on M (i.e. a set of n-tuples of M).

Let $g^+_{\mathfrak{M}}$ be the term assignment uniquely determined by the model $\mathfrak{M}$ and the second-order variable assignment g. By definition, $g^+_{\mathfrak{M}}$ agrees with g on variables of any kind (object, function or predicate), takes the value $\alpha^{\mathfrak{M}}$ for constant α, and for any complex term $\tau = \zeta(\tau_1, \cdots, \tau_n)$, where ζ is either a function constant or a function variable,

$$g^+_{\mathfrak{M}}(\tau) = \zeta^{\mathfrak{M}}(g^+_{\mathfrak{M}}(\tau_1)), \cdots, g^+_{\mathfrak{M}}(\tau_n))$$

Satisfaction
The supplementary clauses are as follows. For any n-place predicate variable Ψ and terms $\alpha_1, \cdots, \alpha_n$:

$$Sat(\mathfrak{M}, g, \Psi\tau_1 \cdots \tau_n) \text{ iff } \langle g^+_{\mathfrak{M}}(\tau_1), \cdots, g^+_{\mathfrak{M}}(\tau_n)\rangle \in \Psi^{\mathfrak{M}}$$

For any function variable Φ:

> $Sat(\mathfrak{M}, g, \forall\Phi\phi)$ iff $Sat(\mathfrak{M}, h, \phi)$ for any h that is a second-order variable assignment agreeing with g on all variables, with the possible exception of the function variable Φ.

For any predicate variable Ψ:

> $Sat(\mathfrak{M}, g, \forall\Psi\phi)$ iff $Sat(\mathcal{M}, h, \phi)$ for any h that is a second-order variable assignment agreeing with g on all variables, with the possible exception of the predicate variable Ψ.

Model-Theoretic Consequence
As in the first-order case.

3.2 Second-order logic and the Continuum Hypothesis

To appreciate some of the literature on logical consequence,[3] it will be useful to know that there is a sentence S_{CH} which is a second-order validity iff the Continuum Hypothesis is true. As it would take us too far afield to justify this claim in detail, we merely sketch some of the ideas behind its justification.

The Continuum Hypothesis is an interpreted first-order sentence in the language of set theory which we may take to be:

CH $2^{\aleph_0} = \aleph_1$; in words: the cardinality of the set of functions from $\aleph_0$ to $2 = \{0,1\}$ is the first uncountable cardinal.

An easy argument shows that $2^{\aleph_0}$ is the size of the real numbers, so that CH is equivalent to the claim that if a subset of the real numbers is strictly larger than the natural numbers then it is at least as large as the set of real numbers—and therefore exactly as large as this set.

[3] In particular, the 'overgeneration argument' advanced in Etchemendy (1990) and discussed in many articles since, including Griffiths & Paseau (2016).

We show that there is a second-formula $Nat(X)$ whose only free variable is the monadic second-order variable X as displayed with the following property: it is satisfied in an interpretation by a subset of the domain just when that subset is equinumerous with the natural numbers. We do this by listing a sequence of properties and their definitions.

1. $X^{\mathfrak{M}}$ is a subset of $Y^{\mathfrak{M}}$; $X^{\mathfrak{M}}$ is a proper subset of $Y^{\mathfrak{M}}$ (here X and Y are one-place predicate variables or constants).

   Definition: the former is $\forall x(Xx \to Yx)$, abbreviated as $X \subseteq Y$; the latter is $\forall x(Xx \to Yx) \wedge \exists y(\neg Xy \wedge Yy)$, abbreviated as $X \subsetneq Y$.

2. $f^{\mathfrak{M}}$ is a function with codomain $Y^{\mathfrak{M}}$ (here f is a one-place function variable or constant and Y is a one-place predicate variable or constant).

   Definition: $\forall x \forall y(fx = y \to Yy)$.

3. $f^{\mathfrak{M}}$ is injective on $X^{\mathfrak{M}}$ (f is a one-place function variable or constant and X is a one-place predicate variable or constant).

   Definition: $\forall x \forall y((Xx \wedge Xy) \to (fx = fy \to x = y))$.

4. $X^{\mathfrak{M}}$ is of cardinality smaller or equal to $Y^{\mathfrak{M}}$ (X and Y are one-place predicate variables or constants).

   Definition: $\exists f$(f is injective on X and has codomain Y). We abbreviate this as $X \leq Y$.

5. $X^{\mathfrak{M}}$ is infinite (X is a one-place predicate variable or constant).

   Definition: $\exists f \exists Y$($Y \subsetneq X \wedge f$ is injective on $X \wedge f$ has codomain Y).

6. $X^{\mathfrak{M}}$ is of the same size of the natural numbers (X is a one-place predicate variable or constant).

   Definition: X is infinite $\wedge \forall Y$(Y is infinite $\to X \leq Y$). We abbreviate this formula as $Nat(X)$.

The last clause shows that '$X^{\mathfrak{M}}$ is of the same size as the set of natural numbers' is definable by $Nat(X)$; so we have found the formula we were looking for.

It turns out that '$Y^{\mathfrak{M}}$ is of the same size as the set of real numbers' is also definable by a second-order formula $Real(Y)$. To see this in detail would require more mathematics than we have time to present. The idea is that the real numbers with the operations $+$ and $\times$ and the relation $<$ enjoy a uniqueness property: any other structure satisfying the conditions enjoyed by this structure must be isomorphic to it. The more precise statement is that the reals are the up-to-isomorphism unique ordered field with the least upper bound property. If this structure is defined by $\phi(Y, f_1, f_2, R)$ then $Real(Y)$ may be defined as $\exists f_1 \exists f_2 \exists R \phi(Y, f_1, f_2, R)$.

Using the two formulas $Nat(X)$ and $Real(Y)$, it is now easy to define the following second-order sentence S_{CH}:

$$S_{\mathsf{CH}} \quad \forall X \forall Y \forall Z((Nat(X) \wedge Real(Y) \wedge Z \subseteq Y \wedge X < Z) \to Y \leq Z)$$

Suppose S_{CH} is true in all second-order interpretations. It is true in particular in the model whose domain is the real numbers. It follows that any subset of the reals whose size is greater than that of the natural numbers is of the same size as the real numbers; i.e. it follows that CH is true. Conversely, if CH is true, then by a similar argument S_{CH} must be true in all interpretations. The upshot is that S_{CH} is a second-order validity iff CH is true.

4 Compactness

4.1 General properties

Suppose a logic consisting of a language, grammar, semantics and consequence relation $\vDash$ has been specified. As usual, and as just elaborated in the case of first- and second-order logic, if Γ is a set of sentences of the logic and δ a single sentence, $\Gamma \vDash \delta$ means that any model of Γ (i.e. of all the sentences in Γ) is a model of δ. The logic in question is said to be *compact* when one of the following three statements holds:[4]

- If $\Gamma \vDash \delta$ then $\Gamma^{fin} \vDash \delta$ for some finite subset Γ^{fin} of Γ.
- If Γ is an unsatisfiable set of sentences then so is Γ^{fin} for some finite subset Γ^{fin} of Γ.
- If every finite subset Γ^{fin} of Γ is satisfiable then so is Γ.

The equivalence of these three characterisations of compactness is immediate: the third statement is the contrapositive of the second, and in a logic containing negation the equivalence of the first and second statements follows from

$$\Gamma \vDash \delta \text{ if and only if } \Gamma \cup \{\neg\delta\} \text{ is unsatisfiable.}$$

The compactness theorem is said to hold for a logic precisely when the logic is compact.

Two important examples of compact logics are propositional logic and first-order logic. First-order logic's compactness is of tremendous importance, since to this day it remains the canonical logic within mathematics, the widespread interest in higher, supplementary and alternative logics notwithstanding. By 'first-order logic', we understand throughout first-order logic with identity; first-order logic without identity is of course also compact, as it is a sublogic of first-order logic with identity.

Second-order logic with standard or full semantics, in which second-order n-place predicate variables range over all the n-tuples from the domain of interpretation (and similarly for functional variables), is in contrast not compact. To see this, let $\exists_{\geq n}$ be a sentence of first-order logic satisfied in all and only models with domain of size $\geq n$; $\exists_{\geq 1}$ may be taken to be $\exists x(x = x)$, $\exists_{\geq 2}$ as $\exists x \exists y \neg(x = y)$, and so on. Since first-order logic is a sublogic of second-order logic, $\exists_{\geq n}$ is a sentence of second-order logic too. Consider next the sentence

[4] Some authors take the compactness of a logic to be its satisfaction of the statements' biconditional versions.

$$\exists R(R \text{ is functional } \wedge R \text{ is injective } \wedge \neg R \text{ is surjective})$$

with R a two-place predicate. The clause 'R is functional' abbreviates

$$\forall x \exists y Rxy \wedge \forall x \forall y_1 \forall y_2((Rxy_1 \wedge Rxy_2) \to y_1 = y_2),$$

'R is injective' abbreviates

$$\forall x_1 \forall x_2 \forall y((Rx_1y \wedge Rx_2y) \to x_1 = x_2)$$

and 'R is surjective' abbreviates

$$\forall y \exists x Rxy$$

Alternatively, use the sentence

$$\exists X \exists f \exists Y(Y \subsetneq X \wedge f \text{ is injective on } X \wedge f \text{ has codomain } Y),$$

which existentially quantifies the formula in the previous section. Any interpretation of either of these two sentences states that the domain is Dedekind infinite. The following second-order argument is then valid:

$$\begin{array}{c} \exists_{\geq 1} \\ \exists_{\geq 2} \\ \vdots \\ \exists_{\geq n} \\ \vdots \\ \hline \exists R(R \text{ is functional } \wedge R \text{ is injective } \wedge \neg R \text{ is surjective}) \end{array}$$

But no finite subset of the premisses entails the conclusion. For let the finite subset be $\{\exists_{\geq i_1}, \exists_{\geq i_2}, \ldots, \exists_{\geq i_k}\}$ and take $m \geq \max\{i_1, i_2, \ldots, i_k\}$. Then there is a model of size m in which the k premisses $\exists_{\geq i_1}, \exists_{\geq i_2}, \ldots, \exists_{\geq i_k}$ are true but the argument's conclusion is false. Hence second-order logic is not compact.

It follows from its incompactness that second-order logic is also incompletable. For a simple argument demonstrates that if a logic has a sound and complete proof procedure, then it must be compact:

(1)	$\Gamma \vDash \delta$	Assumption
(2)	$\Gamma \vdash \delta$	From (1) by Completeness
(3)	$\Gamma^{fin} \vdash \delta$	From (2) by the finiteness of proofs
(4)	$\Gamma^{fin} \vDash \delta$	From (3) by Soundness

Here Γ^{fin} is some finite subset of Γ. The validity of the inference from (2) to (3) follows from the requirement that proofs draw only on finitely many premisses. The argument just given therefore applies to any logic which has a sound and complete proof procedure in this liberal sense. It follows that second-order logic is incompletable (by a sound proof procedure) as well as incompact.

The compactness theorem also typically, but not invariably, fails for infinitary logics (see below). Any logic which allows infinite disjunctions, for example, is incompact, since the set of sentences $\{c \neq c_i : i \in \omega\} \cup \{\bigvee_{i \in \omega} c = c_i\}$ is finitely satisfiable but unsatisfiable.

Not all logics with higher-order (second-order or above) quantifiers are incompact. Second-order logic with Henkin semantics is compact.[5] A less familiar example is so-called pure second-order logic with identity.(Beware: 'pure second-order logic' is also used to mean second-order logic without any non-logical vocabulary, with no non-logical constants, function constants or predicate constants.) Pure second-order logic with identity lacks functional and first-order variables, but has predicate variables and quantifiers as well as both second-order and first-order identity. It may be thought of as the complement of first-order logic relative to second-order logic in the following sense: first-order logic has object but not predicate quantifiers; pure second-order logic has predicate but not object quantifiers; and second-order logic combines the two.[6] In other words second-order logic merges first-order and pure second-order logic. Pure second-order logic with identity is also known to be compact.[7] The moral is that the incompactness of second-order logic is not owed solely to the presence of second-order quantifiers but to the combination of both first- and second-order quantifiers.

What about natural language? Is it compact? First, let's clarify what the question means. Assume there is such a thing as the relation of logical consequence in natural example. For example, consider these two natural-language arguments:

Hypatia is a woman.	Hypatia is mortal.
All women are mortal.	All women are mortal.
Hypatia is mortal.	Hypatia is a woman.

The argument on the left is logically valid, whereas the argument on the right is invalid. Let's say that a natural language N is compact just when, for any logically valid N-argument, there is a logically valid argument whose premiss set is a finite subset of the orginal argument's premiss set and whose conclusion is the same as the original argument's conclusion. This definition is the analogue of the first definition of compactness above for a formal language (*viz.* if $\Gamma \vDash \delta$ then $\Gamma^{fin} \vDash \delta$ for some finite subset Γ^{fin} of Γ).

Consider now the following English analogue of the second-order-logic argument presented earlier:

[5]For an account of Henkin semantics, see e.g. chapter 4 of Enderton (2001).

[6]Of course, standard second-order logic also has functional variables and quantifiers, but in the presence of predicate variables and quantifiers these are dispensable.

[7]See my (2010). Denyer (1992) gives an argument that applies to pure second-order logic without identity.

$$
\begin{array}{c}
\text{There is at least one thing.} \\
\text{There are at least two things.} \\
\vdots \\
\text{There are at least } n \text{ things.} \\
\vdots \\
\hline
\text{There are infinitely many things.}
\end{array}
$$

If this argument is valid, as some philosophers believe,[8] then the consequence relation in English is incompact, since no finite subset of the premiss set entails the conclusion. We shall return to the compactness or otherwise of natural language in next week's class.

4.2 Semantic and deductive completeness

It is worth briefly pausing to explain what a complete *theory* (as opposed to a complete proof system) is, in case anyone hasn't encountered this notion before. A *theory* T in a formal language $\mathcal{L}$ is understood as a set of $\mathcal{L}$-sentences. Then:

T is semantically complete $=_{df}$ for any $\mathcal{L}$-sentence ϕ, $T \models \phi$ or $T \models \neg\phi$.

T is deductively complete $=_{df}$ for any $\mathcal{L}$-sentence ϕ, $T \vdash \phi$ or $T \vdash \neg\phi$.

Note that these definitions apply to sentences and not open formulas. (A notion of completeness that took in open formulas as well as sentences would not be very useful, since any theory T satisfiable by a model of domain size greater than 1 is such that $T \not\models x = y$ and $T \not\models \sim x = y$.) It is immediate from the definitions just given that in any logic $\mathcal{L}$ with a sound and complete proof system (e.g. first-order logic), T is semantically complete iff it is deductively complete.

4.3 First-order logic's expressive limitations

This subsection briefly draws some implications of the compactness of first-order logic. The sample below is a tiny selection from a list that could fill volumes. We assume knowledge of elementary model theory.[9]

Any compact logic extending first-order logic cannot express the notions of finitude or infinitude (of a model). Suppose towards a contradiction that ϕ_F is satisfied by all and only finite models. Then the set $\{\phi_F\} \cup \{\exists_{\geq n} : n \in \omega\}$ is unsatisfiable, and hence by compactness must have an unsatisfiable finite subset, which must be a subset of $\{\phi_F\} \cup \{\exists_{\geq i_1}, \ldots, \exists_{\geq i_k}\}$ for some $i_1 < \ldots < i_k$. But any finite model with domain of size $\geq i_k$ satisfies (any subset of) $\{\phi_F\} \cup \{\exists_{\geq i_1}, \ldots, \exists_{\geq i_k}\}$, thereby

[8] See for example Oliver & Smiley (2013, p. 238), or Yi (2006, p. 262) for a similar claim.

[9] Standard references are Chang & Keisler (1990), Hodges (1993) and Hodges (1997).

contradicting our hypothesis. And if there were a sentence ϕ_I satisfied by all and only infinite models then $\neg\phi_I$ would be satisfied by all and only finite models, a hypothesis we have just refuted. Thus no such sentence ϕ_I exists either.

This application of the compactness theorem is entirely typical. Schematically, one shows by contradiction that the class of models with some ω-property expressible by a set of first-order sentences is not definable by a single sentence ϕ. In that case, $\neg\phi$ and the union of the set whose sentences are 'the model has the n-property' for each finite n (plus any background assumptions) is unsatisfiable, but any of its finite subsets is satisfiable, contradicting compactness. Informally speaking, in these applications the ω-property is the conjunction of all the n-properties; in our example, the n-property is having size $\geq n$ and the ω-property is having infinite size.

The compactness theorem may be used to show that any first-order theory of arithmetic T_{AR} satisfied by the standard model has a non-standard model.[10] Assuming that each numeral $\overline{n}$ is definable in T_{AR}, consider

$$T_{AR}^{+} = T_{AR} \cup \{c \neq \overline{n} : n \in \omega\}$$

where c is any constant not in T_{AR}'s language. Any finite subset of T_{AR}^{+} is satisfied by the standard model, because we may interpret c as a number larger than the largest n such that the sentence $c \neq \overline{n}$ is in the given finite subset. Hence by compactness, T_{ar}^{+} has a model $\mathfrak{M}$. The reduct of $\mathfrak{M}$ to the language of T_{AR} is non-standard since it contains an element not identical to any natural number, viz. $c^{\mathfrak{M}}$, the denotation of c in $\mathfrak{M}$.

Extending the argument just given with an appeal to the downward Löwenheim-Skolem Theorem shows that any such first-order theory of arithmetic T_{AR} is not even $\aleph_0$-categorical, since it contains a countably infinite non-standard model. Beware: this does *not* imply that arithmetic is incomplete. Gödel (1931) proved that for any consistent first-order theory of arithmetic T with a recursively enumerable set of theorems, there is a sentence ϕ such that neither ϕ nor $\neg\phi$ is a theorem of T. That a theory fails to be $\aleph_0$-categorical is, however, compatible with its being (deductively or semantically) complete. To see this, run the previous argument supplemented by an application of the downward Löwenheim-Skolem Theorem for a *complete* theory of arithmetic T_{AR}, for example the first-order theory of all true arithmetical sentences in the first-order language which contains the constant $\overline{0}$ (whose intended interpretation is 0), the one-place symbol S (whose intended interpretation is the successor function), the two-place symbol $+$ (whose intended interpretation is addition), and the two-place symbol $\times$ (whose intended interpretation is multiplication). By the above argument, the theory T_{AR} has a non-standard model, complete though it may be.

The same general idea can be used to demonstrate the existence of non-standard models of real analysis. Let T_{AN} be a first-order theory of analysis satisfied by the standard model (the ordered field of real numbers). As above, consider $T_{AN}^{+} =$

[10] By the standard model of arithmetic (for the theory T_{AR} in question), we mean the structure of natural numbers with the standard interpretation of the non-logical symbols in the language of T_{AR}: the constant $\overline{0}$ denotes 0, the two-place symbol $+$ denotes addition, $\times$ denotes multiplication, etc.

$T_{AN} \cup \{0 < c < \frac{1}{n} : n = 1, 2, \ldots\}$ where c is any constant not in T_{AN}'s language. Any finite subset of T^+_{AN} is satisfied by the standard model, because we may interpret c as a positive real number smaller than $\frac{1}{n}$, where n is the largest number for which the sentence $0 < c < \frac{1}{n}$ is in the given finite subset. Hence by compactness, T^+_{AN} has a model $\mathfrak{M}$. The reduct of $\mathfrak{M}$ to the language of T_{AN} is non-standard since it contains an element not identical to any real number, viz. $c^{\mathfrak{M}}$, the denotation of c in $\mathfrak{M}$. Indeed, this element must be a positive infinitesimal, meaning that it is a number greater than 0 but smaller than every fraction $\frac{1}{n}$. As well as infinitesimals, our non-standard model also contains infinite elements, since the model satisfies $\forall x \neq 0 \exists y (x \cdot y = 1)$ and thus any non-zero element has an inverse. From these foundations, a consistent version of the calculus that revives to a fashion the use of infinitesimals in early modern mathematics may be constructed.[11]

Since second-order logic with standard semantics is incompact, the arguments just given fail for second-order theories of arithmetic and analysis. Indeed, there are categorical second-order axiomatisations of arithmetic and real analysis, as mentioned above. The standard second-order axiomatisation of arithmetic is given by the following three axioms:

- $\neg \exists x (Sx = 0)$
- $\forall x \forall y (Sx = Sy \rightarrow x = y)$
- $\forall X[(X0 \wedge \forall x(Xx \rightarrow XSx)) \rightarrow \forall x Xx]$

This axiomatisation, usually known as second-order Peano Arithmetic, is categorical.

Assuming the downward Löwenheim-Skolem theorem, another corollary to the compactness of first-order logic is the upward Löwenheim-Skolem theorem. This upward version of the theorem states that if a first-order language $\mathcal{L}$ has cardinality $\leq \lambda$ and $\mathfrak{M}$ is an infinite model with domain of cardinality $\leq \lambda$ then $\mathfrak{M}$ has an elementary extension of cardinality λ. For the proof, we consider the set of sentences consisting of the elementary diagram of $\mathfrak{M}$ and each sentence in $\{c_\alpha \neq c_\beta : \alpha, \beta \in \lambda \text{ s.t. } \alpha \neq \beta\}$, where the c_α are new constants. This set is finitely satisfiable (because the infinite model $\mathfrak{M}$ satisfies any finite subset), and hence by compactness it is satisfiable, satisfied by a model $\mathfrak{N}$ say, which must be of size $\geq \lambda$ as it satisfies $\{c_\alpha \neq c_\beta : \alpha, \beta \in \lambda \text{ s.t. } \alpha \neq \beta\}$. Since $\mathfrak{N}$ also satisfies the elementary diagram of $\mathfrak{M}$, an elementary embedding of $\mathfrak{M}$ into $\mathfrak{N}$ exists, and thus there is an elementary extension $\mathfrak{O}$ of $\mathfrak{M}$ with domain of size $\geq \lambda$ ($\mathfrak{O}$ is an isomorph of $\mathfrak{N}$ whose domain includes that of $\mathfrak{M}$). To find an elementary extension of $\mathfrak{M}$ of size exactly λ, now apply the downward Löwenheim-Skolem theorem to $\mathfrak{O}$.

The upward Löwenheim-Skolem theorem may be applied to show not only that theories of arithmetic and analysis satisfied by their respective standard models have non-standard models, but also that they have non-standard models of every infinite cardinality. More generally, any first-order theory in a countable language satisfied by an infinite model has models of every infinite cardinality.

The compactness theorem for first-order logic has a great many other applications to model theory—as Keisler has put it, 'The most useful theorem in model theory

[11] See Goldblatt (1998) or the original Robinson (1966).

is probably the compactness theorem' (Keisler 1965, p. 113)—as well as to many other parts of mathematics; see chapter 6 of Hodges (1993) for more.

5 The ω-rule

The ω-rule is an infinitary rule that may be added to systems of arithmetic. If added to sufficiently strong first-order systems, it allows them to overcome incompleteness. In this section, we introduce the ω-rule and sketch why its addition to an arithmetical theory that enjoys a weak completeness property turns the latter into a complete theory.

Suppose our theory of arithmetic T_{Ar} has numerals of the form $\overline{0}, \overline{1}, \cdots, \overline{n}, \cdots$. Typically, such a system will generate these numerals by taking $\overline{0}$ as a constant and by applying the successor function symbol S to it, so that $\overline{1} = S(\overline{0})$, $\overline{2} = SS(\overline{0})$, etc. The first-order systems of Peano Arithmetic or Robinson Arithmetic are of this form. Suppose that $\phi(x)$ is a formula with one free variable, x. The ω-rule is the following inference:

$$\frac{\phi(\overline{0}), \phi(\overline{1}), \ldots, \phi(\overline{n}), \ldots}{\forall x \phi(x)}$$

This rule is obviously sound with respect to truth in the standard model of arithmetic: if the premisses are all true in that model so is the conclusion. We define T^{ω}_{Ar} as the system T_{Ar} augmented with the ω-rule. A derivation using the ω-rule has infinitely many premisses so is an infinite object, unlike a typical formal proof or our usual informal idea of proof. Indeed, derivability from the axioms of T^{ω}_{Ar} for a typical first-order system T_{Ar} is tantamount to truth in the standard model.

To see this in a bit more (though far from complete) detail, we must appreciate that in a typical arithmetical system, every formula is equivalent to one of a special form. For x any variable, and $\overline{n}$ any numeral, quantification in either of the forms $\forall x(x \leq \overline{n} \rightarrow \phi)$ or $\exists x(x \leq \overline{n} \wedge \phi)$ is called *bounded quantification*, abbreviated as $(\forall x \leq \overline{n})\phi$ and $(\exists x \leq \overline{n})\phi$ respectively. Also, for x and y distinct variables and ϕ any formula, quantification in either of the forms $\forall x(x \leq y \rightarrow \phi)$ and $\exists x(x \leq y \wedge \phi)$ is also called bounded quantification, abbreviated as $(\forall x \leq y)\phi$ and $(\exists x \leq y)\phi$, respectively.[12]

The definition of a Σ_0-formula is then: every atomic formula of the language of T_{Ar} is Σ_0; and Σ_0-formulas are closed under Boolean operations and bounded existential and universal quantification. Σ_0-formulas are also known as Π_0-formulas (as well as Δ_0-formulas, though this won't be relevant here). In typical systems of arithmetic, we can effectively decide (compute) the truth or falsity in the standard model of each Σ_0-sentence, i.e. closed Σ_0-formula.

A Σ_1-formula is any formula of the form $\exists x \phi$ where ϕ is a Σ_0-formula, and a Π_1-formula is any formula of the form $\forall x \phi$ where ϕ is a Σ_0-formula. More generally,

[12] The restriction that the variables x and y be distinct when the bound on the quantification is a variable is essential, since $\forall x(x \leq x \rightarrow \phi)$ is logically equivalent to $\forall x \phi$, which is unbounded quantification.

a Σ_{n+1}-formula prefaces a Π_n-formula with an existential quantifier, and a Π_{n+1}-formula prefaces a Π_n-formula with an existential quantifier. In a typical theory T_{AR}, one can use prenex normal form theorems plus some arithmetical coding to show that every formula is T_{AR}-equivalent to a Σ_n-formula or a Π_n-formula.

Given these facts, it is easy to show that if T_{Ar} is Σ_0-complete, then T^{ω}_{Ar} is complete. To say that T_{Ar} is Σ_0-complete (equivalent to its being Π_0-complete) is to say that if a Σ_0 sentence ϕ is true in the standard model of arithmetic then $T_{Ar} \vdash \phi$. The result is proved by induction, the base case being the argument's assumption.

Assume for the Induction Hypothesis that the result holds for Σ_n- and Π_n-sentences. First, let X be a true-in-the-standard-model Σ_{n+1}-sentence $\exists x\phi(x)$, where $\phi(x)$ is a Π_n-formula. Then for some natural number m, $\phi(\overline{m})$ is a true Π_n-sentence. By the Induction Hypothesis, $T^{\omega}_{Ar} \vdash \phi(\overline{m})$. Since $T_{Ar} \vdash \phi(\overline{m}) \rightarrow \exists x\phi(x)$, $T^{\omega}_{Ar} \vdash \exists x\phi(x)$. Second, let X be a true-in-the-standard-model Π_{n+1}-sentence $\forall x\phi(x)$. Then for each number n, $\phi(\overline{n})$ is a true-in-the-standard-model Σ_n-sentence. Then by the Induction Hypothesis, for each n, $T^{\omega}_{Ar} \vdash \phi(\overline{n})$. So by one application of the ω-rule, $T^{\omega}_{Ar} \vdash \forall x\phi(x)$.

6 The undecidability of first-order logic

6.1 Decision Problems and Decidability

A problem is effectively decidable (or decidable or solvable) iff there is an algorithm (or effective procedure) for resolving whether any given x of kind K has or lacks proprty P: it delivers YES or NO within a finite number of steps and finite amount of time. N.B. The algorithm need not be known.

What is an algorithm? A mechanical procedure for deciding a problem. It can be given as a finite set of instructions which are executed in a stepwise manner, without appeal to random processes or ingenuity. We ignore 'accidental' limits on the amount of time, speed of computation and matter in the universe, requiring all of them to be finite but setting no finite upper bound. The notion of algorithm, note, is informal.

Some example problems, all of which are effectively decidable:

- Is a given sentence of the propositional calculus a tautology?
- Is a given sentence of the English language a palindrome?
- Is a positive integer expressed as an Arabic numeral prime?
- Is a given finite sequence of symbols of propositional calculus a sentence (wff)?
- Is a given finite sequence of finite sequences of symbols of propositional calculus a proof of its last member?

The first problem, for instance, is effectively decidable because we may use the truth-table test to determine whether the sentence is a tautology. If the sentence has n distinct sentence letters, we will have to check 2^n rows. If n is very large, we may in practice never complete this test; indeed, we may not even get started on it because it would take too long even to process the sentence—no computer can even store it. For the third problem, we can use any of the well-known tests for primality, e.g. the sieve of Eratosthenes which runs through each of the numbers up to $\sqrt{N}$ and checks whether it divides N. What this third problem illustrates is that the problem's presentation affects its decidability. The following problem, for example, is *not* effectively decidable (recall the definition of the Continuum Hypothesis–CH–from earlier):

Determine whether N is prime, where N is defined by

$$N = \begin{cases} 3 & \text{if } \mathsf{CH} \text{ is true} \\ 4 & \text{if } \mathsf{CH} \text{ is false} \end{cases}$$

The property P mentioned within a problem is effectively decidable (or decidable or effective) iff the problem is; the effective procedure deciding P is a decision procedure for P. Similarly, a set is decidable iff there is a decision procedure for determining whether a given object is or is not a member of the set. A relation is decidable iff there is a decision procedure for determining whether given objects (in an order) stand in the relation or do not so stand.

6.2 Church's Theorem

What does it mean to say that a logic is or is not decidable? Associated with any logic (including a semantics) is a salient set: the set of its validities. A logic is decidable iff the set of its validities is decidable. As just mentioned, we know that the propositional calculus is decidable. In 1936, Alonzo Church proved that first-order logic with or without identity is not decidable. Do not confuse Church's Theorem with Church's Thesis (equivalently: Turing's Thesis), although the latter is used in the proof of the former.

The behaviour of logics which lie between propositional calculus and first-order logic with identity is curious. Let monadic and dyadic first-order logic be first-order logic with only monadic (one-place) and dyadic (two-place) predicates respectively. First-order logic is undecidable. Monadic first-order logic, with or without identity, is *decidable*. Dyadic first-order logic is *undecidable*. It follows that any logic which includes first-order logic is also undecidable, assuming that the property of being among the first-order sentences of this particular logic is decidable, as it usually is.

A standard way of proving first-order logic's undecidabilty is to derive it from the undecidability of the Halting Problem. We sketch a proof of the latter, and mention how to use it to prove the former. Suppose we equate computability with computability by a Turing Machine (Turing's Thesis). It is not hard to see that the set of Turing Machines (TMs) can be algorithmically generated in a list. Suppose TM_i is the i^{th} Turing Machine and that the Halting Problem is decidable, i.e. there is a TM that decides whether TM_i with input j halts or not. We may then define another TM—TM_k for some k—which *halts* if TM_i with input i does *not* halt and does *not* halt if TM_i with input i *does* halt. TM_k then halts on input k iff it does not; a contradiction. The conclusion is that the Halting Problem is undecidable.

First-order logic's undecidability may be demonstrated by expressing (a version of) the Halting Problem in first-order terms. For example, for each TM M we can come up with a first-order formula ϕ_M that is valid iff M halts with input 0 (say). The formula ϕ_M can be constructed by describing what M does in first-order terms. This is not too hard to see since M's states are finite, its list of symbols it writes is finite, and its rules (instructions for what to do in a given state when scanning a particular symbol) are also finite. So we can come up with a first-order sentence expressing the following claim: 'If the input to machine M with rules ... and states ... is 0 then M halts (i.e. is in the halting state)'. First-order logic's undecidablity then follows from the Halting Problem's undecidability. Of course, there's a lot of detail to fill in, but that's the basic strategy. More detailed discussion of these matters may be found in, for example, chapter 21 of Boolos, Burgess and Jeffrey (2007).

6.3 Decidability of English validity

At least as far back as Leibniz, logicians have aspired to find a mechanical means of testing an argument's validity. In the usual jargon, we say that argumentative validity is *decidable* iff such a mechanical means exists. Suppose you take validity in English to be correctly modelled by semantic consequence in logic $\mathcal{L}$; for example W.V. Quine equated first-order validity with validity *simpliciter*. Let's call $Form_{\mathcal{L}}(s)$ the formalisation of an English sentence s into $\mathcal{L}$, and $Form_{\mathcal{L}}(S)$ the formalisation of a set S of English sentences into $\mathcal{L}$, which is the same as the set of formalisations of the sentences in S. The standpoint of the $\mathcal{L}$-logician (who takes $\mathcal{L}$ to correctly model English logical consequence) is captured by the following biconditional:

S logically entails s iff $Form_{\mathcal{L}}(S) \vDash_{\mathcal{L}} Form_{\mathcal{L}}(s)$.

From this standpoint, two conditions sufficient for the decidability of argumentative validity are:

- that the formalisation of natural-language sentences into $\mathcal{L}$-sentences be a mechanical procedure;
- that $Form_{\mathcal{L}}(S) \vDash_{\mathcal{L}} Form_{\mathcal{L}}(s)$ be decidable for any set of natural-language sentences S and any natural-language sentence s.

Whether the first condition obtains depends on $\mathcal{L}$; for the usual choices of $\mathcal{L}$, it's at best moot. However that may be, Church's result shows that the second condition fails for any $\mathcal{L}$ that incorporates first-order logic,[13] since being a first-order validity is undecidable.

[13] Assuming once more that the sentences of first-order logic are 'detachable' from those of $\mathcal{L}$, as above.

7 Notation, infinitary logics and Lindenbaum algebras

7.1 Notation

We may distinguish between propositional, first-order and second-order logics whose sentence sets have different cardinalities. For κ a cardinal, PL_κ is any propositional logic with κ-many sentence letters; for example, PL_1 has one sentence letter. There are of course lots of propositional logics with κ sentence letters, the two dimensions of variation being the logic's set of sentence letters (e.g. $p_1, \cdots p_n, \cdots$ or $q_1, \cdots q_n, \cdots$) and its truth-functional connectives. The differences between these are usually immaterial if the propositional logic has a truth-functionally complete set of connectives, as is almost required. Accordingly, we happily speak of *the* logic PL_κ for a given κ, assuming unless otherwise stated that it comes with a truth-functionally complete set of connectives.

PL_κ has a countable infinity of sentences when $1 \leq \kappa \leq \omega$, and κ-many sentences when κ is infinite (these two cases overlap when $\kappa = \omega$).[14] Notice that when $\kappa = n$ is finite, the compactness of the propositional logic PL_n is a trivial consequence of the fact that any sentence of PL_n is logically equivalent to a sentence drawn from a fixed set of size no greater than 2^{2^n}.[15] For example, PL_1 is the propositional logic with a truth-functionally complete set of connectives and a single sentence letter p. Each of PL_1's infinitely many sentences is logically equivalent to one of the following four: $p, \neg p, p \vee \neg p, p \wedge \neg p$.

Similarly, FOL_κ is first-order logic with κ-many variables, constants, predicate and function symbols of all arities; a countable and expressively adequate set of truth-functional connectives; standard formation rules; and its standard consequence relation. If it were of relevance, which it won't be for us, one could draw further cardinality distinctions between first-order logics that have κ_1 variables, κ_2 constants, κ_3 truth-functional connectives, λ_i predicate symbols of arity i, and μ_j function symbols of arity j.

Turn now to second-order logic, very briefly. As well as a countable and expressively adequate set of connectives, the language of SOL_κ has κ many: first-order variables, second-order predicate and function variables of all arities, non-logical predicates of all arities (including arity 0, i.e. constants), and function symbols of all arities.

7.2 Infinitary logics

Infinitary logics play an important in contemporary discussions of logical constants. One simple way to extend finitary logics is to allow infinitary conjunctions and disjunctions and well-formed formulas. Take for example the much-studied logic $\mathcal{L}_{\omega_1\omega} = \mathsf{FOL}_{\omega_1\omega}$.[16] Let FOL^* be first-order logic with a countable infinity of relation, function and constant symbols, ω_1-many variables, as well as parentheses and

[14] Assuming PL_κ has at least one n-ary connective for some $n > 0$.

[15] The set is of size exactly 2^{2^n} just when the set of connectives is truth-functionally complete.

[16] For much more on $\mathcal{L}_{\omega_1\omega}$, see Keisler (1971).

the usual logical vocabulary. The latter includes a countable and truth-functionally complete set of connectives containing inter alia the negation symbol $\neg$, the conjunction symbol $\wedge$ and the disjunction symbol $\vee$. The last two symbols are usually written as $\bigwedge$ and $\bigvee$ when they conjoin statements in a set (see below). The syntax of $\mathcal{L}_{\omega_1\omega}$ is given by:

Vocabulary The vocabulary of $\mathcal{L}_{\omega_1\omega}$ is that of FOL^*.

Grammar The class of all well-formed formulas of $\mathcal{L}_{\omega_1\omega}$ is the least class $\mathcal{C}$ such that:

(i) Each atomic FOL^*-formula is in $\mathcal{C}$;

(ii) If ϕ is in $\mathcal{C}$ and v is a variable, then $\neg\phi$ is in $\mathcal{C}$, $\forall v\phi$ is in $\mathcal{C}$, and $\exists v\phi$ is in $\mathcal{C}$;

(iii) If Φ is a countable (finite or countably infinite) non-empty subset of $\mathcal{C}$, then $\bigwedge\Phi$ is in $\mathcal{C}$, and $\bigvee\Phi$ is in $\mathcal{C}$.

As for its semantics, we may define satisfaction for the pure language of $\mathcal{L}_{\omega_1\omega}$ as follows (here σ is a variable assignment):[17]

$(\mathcal{M}, \sigma) \vDash x_i = x_j$ iff $\sigma(x_i) = \sigma(x_j)$;

$(\mathcal{M}, \sigma) \vDash \bigvee\{\phi_i : i < \kappa\}$ iff $(\mathcal{M}, \sigma) \vDash \phi_i$ for some ϕ_i in the sequence $\{\phi_i : i < \kappa\}$;

$(\mathcal{M}, \sigma) \vDash \neg\phi$ iff it is not the case that $(\mathcal{M}, \sigma) \vDash \phi$;

$(\mathcal{M}, \sigma) \vDash \exists x\phi$ iff there is a variable assignment ρ that differs from σ at most over the variable x such that $(\mathcal{M}, \rho) \vDash \phi$.

A full $\mathcal{L}_{\omega_1\omega}$-interpretation consists not just of a domain, but also of an interpretation of the non-logical constants over that domain. The satisfaction relation for full $\mathcal{L}_{\omega_1\omega}$ is these clauses' obvious extension to sentences containing non-logical constants.

The logic $\mathcal{L}_{\omega_1\omega}$ is only the tip of the infinitary iceberg. The logic $\mathcal{L}_{\kappa\lambda}$, where κ and λ are infinite cardinals, extends first-order logic by allowing $< \kappa$ conjunctions and disjunctions of well-formed formulas for any κ, and by allowing quantification over $< \lambda$ variables. In this terminology, first-order logic is then $\mathcal{L}_{\omega\omega}$, as it allows only finitary conjunction/disjunction and quantification. The logic $\mathcal{L}_{\infty\infty}$ extends FOL by allowing κ-ary conjunctions and disjunctions of well-formed formulas for any κ, and by allowing κ-ary quantification over variables for any infinite κ. Its semantics also extends that of first-order logic and $\mathcal{L}_{\omega_1\omega}$ in the obvious way.

7.3 Lindenbaum algebras

The *Lindenbaum algebra* of a logic that contains a truth-functionally complete set of connectives is a particular kind of Boolean algebra. Its elements are the equivalence classes of the logic's sentences under the relation of logical equivalence. We give a brief and informal characterisation of Lindenbaum algebras by developing a particular example, referring you to Halmos & Durant (2009) for a complete introduction

[17] The formulas in question are assumed to be well-formed.

to Boolean algebras, and Hinman (2005, pp. 74-9) for a bit more on Lindenbaum algebras in particular.

The Lindenbaum algebra of FOL_ω is the Boolean algebra of FOL_ω quotiented by $\models_{\mathsf{FOL}_\omega}$-equivalence. Two elements $[\gamma_1]$ and $[\gamma_2]$ of this Boolean algebra are equal if and only if $\models_{\mathsf{FOL}_\omega} \gamma_1 \leftrightarrow \gamma_2$. A Boolean algebra's three operations are *join*, *meet* and *complement*. Intuitively, join is disjunction, meet is conjunction and complementation is negation. More precisely, the three operations are respectively defined by:

$$[\gamma_1] \vee [\gamma_2] = [\gamma_1 \vee \gamma_2]$$
$$[\gamma_1] \wedge [\gamma_2] = [\gamma_1 \wedge \gamma_2]$$
$$\overline{[\gamma]} = [\neg\gamma],$$

In these equations, we are using the symbols $\wedge, \vee$ ambiguously, as is customary: on the left, $\wedge$ and $\vee$ are used as Boolean algebra operators, and on the right (within the square brackets) as connectives of FOL_ω. To check that the left-hand side symbols are well-defined we must check that they don't depend on the choice of representatives; for example, if γ_1 and δ_1 are equivalent, and γ_2 and δ_2 are also equivalent, then $\gamma_1 \wedge \delta_1$ is equivalent to $\gamma_2 \wedge \delta_2$.

The Boolean algebra's top element $[\top]$ is the equivalence class of any validity $\top$, and the bottom element $[\bot]$ the equivalence class of any contradiction $\bot$. For elements α and β of a Boolean algebra, we write $\beta \leq \alpha$ if $\alpha \wedge \beta = \beta$, and $\beta < \alpha$ if $\beta \leq \alpha$ and $\alpha \neq \beta$. An *atom* of a Boolean algebra is an element α distinct from the bottom element $\bot$ such that there is no β with the property $\bot < \beta < \alpha$. A Boolean algebra is *atomless* if it has no atoms.

It is easily seen that FOL_ω's Lindenbaum algebra contains countably many elements, since for example $[\exists x F_1 x], \cdots [\exists x F_n x], \cdots$ are all distinct elements. It is also atomless because if α is any element of FOL_ω's Lindenbaum algebra distinct from the bottom element, there is another element β such that $\bot < \beta < \alpha$. To put it in terms of FOL_ω-sentences: given any non-contradictory sentence γ, there is a non-contradictory δ that is stronger than γ (i.e. δ implies γ but not the other way around). Given such a non-contradictory γ, we may for example take δ to be $\gamma \wedge \exists x F x$ where F is a monadic predicate constant not appearing in γ. That

$$\bot \models \gamma \wedge \exists x F x \models \gamma$$

and

$$\gamma \not\models \gamma \wedge \exists x F x \not\models \bot$$

is easily checked.

A useful exercise is to check that PL_n's Lindenbaum algebra consists of 2^{2^n} elements and that if the sentence letters of PL_n are $p_1, \cdots, p_n$, the algebra's atoms are of the form $[\pm p_1 \wedge \cdots \wedge \pm p_n]$, where $\pm p_i$ is either p_i or $\neg p_i$.

8 MPL, S5 in particular

Modal propositional logic MPL is (classical) propositional logic plus the propositional operator $\Box$.

Vocabulary
The vocabulary of MPL consists of:

- Sentence letters: $p_1, p_2, \cdots, p_n, \cdots$
- Boolean operator symbols: $\wedge, \vee, \rightarrow, \leftrightarrow, \neg$
- The modal operator: $\Box$

We often abbreviate $\neg\Box\neg$ as $\Diamond$.

Grammar
MPL's grammar is given by specifying the set of terms and well-formed formulas (wffs).

- Any sentence letter is a wff.
- If ϕ is a wff so is $\neg\phi$; if ϕ_1 and ϕ_2 are wffs, so are $(\phi_1 \wedge \phi_2)$, $(\phi_1 \vee \phi_2)$, $(\phi_1 \rightarrow \phi_2)$ and $(\phi_1 \leftrightarrow \phi_2)$.
- If ϕ is a wff so is $\Box\phi$.

As above, we often drop brackets in wffs.

8.1 S5

There are various systems of modal propositional logic. We present the best-known system: S5.

Theorems of S5

- Any substitution instance (in the language of MPL) of a propositional tautology is an axiom.
- Any instance of the K-schema is an axiom: $\Box(\phi \rightarrow \psi) \rightarrow (\Box\phi \rightarrow \Box\psi)$
- Any instance of the T-schema is an axiom: $\Box\phi \rightarrow \phi$
- Any instance of the S4-schema is an axiom: $\Box\phi \rightarrow \Box\Box\phi$
- Any instance of the S5-schema is an axiom: $\neg\Box\phi \rightarrow \Box\neg\Box\phi$
- If ϕ and $\phi \rightarrow \psi$ are theorems so is ψ (Modus Ponens). If ϕ is a theorem so is $\Box\phi$ (Necessitation).

It is explicit in the first clause and implicit in all the others that an acceptable substitution instance must be a wff in the language of MPL.

8.2 Possible worlds semantics for S5

There is a possible world semantics for S5 for which the proof system just given is sound and complete. We give the semantics in terms of *universal* models. Another sound and complete semantics may be given in terms of models based on frames in which there is an accessibility relation and this relation is an equivalence relation on the set of worlds. This latter semantics generalises better to other modal logics, but our very limited purposes do not require it.

S5 models
An S5 model is an ordered pair $\langle \mathcal{W}, \mathcal{I} \rangle$:

- $\mathcal{W}$ is a non-empty set of objects usually known as possible worlds;
- $\mathcal{I}$ is a two-place function that assigns 0 or 1 to each sentence letter relative to a world, i.e. $\mathcal{I}(\alpha, w) = 0$ or 1, for any sentence letter α and world w in $\mathcal{W}$.

S5 valuation
Where $\langle \mathcal{W}, \mathcal{I} \rangle$ is an S5 model, the valuation $\mathcal{V}_{\mathcal{I}}$ based on it is defined as follows, where α is any sentence letter, ϕ and ψ are any MPL-wffs, and w any member of $\mathcal{W}$:[18]

- $\mathcal{V}_{\mathcal{I}}(\alpha, w) = \mathcal{I}(\alpha, w)$;
- $\mathcal{V}_{\mathcal{I}}(\neg\phi, w) = 1 - \mathcal{V}_{\mathcal{I}}(\alpha, w)$;
- $\mathcal{V}_{\mathcal{I}}(\phi \wedge \psi, w) = \min\{\mathcal{V}_{\mathcal{I}}(\phi, w), \mathcal{V}_{\mathcal{I}}(\psi, w)\}$;
- $\mathcal{V}_{\mathcal{I}}(\phi \vee \psi, w) = \max\{\mathcal{V}_{\mathcal{I}}(\phi, w), \mathcal{V}_{\mathcal{I}}(\psi, w)\}$;
- $\mathcal{V}_{\mathcal{I}}(\Box\phi, w) = 1$ iff for each $w \in \mathcal{W}$, $\mathcal{V}_{\mathcal{I}}(\phi, w) = 1$.

We write $\Gamma \vDash \delta$ iff for every S5-model $\langle \mathcal{W}, \mathcal{I} \rangle$, for every w in $\mathcal{W}$, if $\mathcal{V}_{\mathcal{I}}(\Gamma, w) = 1$ then $\mathcal{V}_{\mathcal{I}}(\delta, w) = 1$. (As usual, we say that $\mathcal{V}_{\mathcal{I}}(\Gamma, w) = 1$ iff $\mathcal{V}_{\mathcal{I}}(\gamma, w) = 1$ for every $\gamma \in \Gamma$.) In particular, we say that δ is an S5 validity iff $\mathcal{V}_{\mathcal{I}}(\delta, w) = 1$ for each w in any W that makes up an S5-model $\langle \mathcal{W}, \mathcal{I} \rangle$. One can then show that this semantics is sound and complete with respect to S5-theoremhood, i.e. that ϕ is an S5-theorem iff ϕ is an S5-validity. (This is sometimes known as *weak* soundness and completeness.) Moreover, S5 has the finite model property: the soundness and completeness claim just given holds for the class of finite models, i.e. models with finite set of worlds $\mathcal{W}$.[19]

[18] We give the clauses for $\neg, \wedge, \vee$, which determine those for the other propositional connectives.
[19] For more on the finite model property, see ch. 8 of Hughes & Cresswell (1996).

References

[1] G.S. Boolos, J.P. Burgess & R.C. Jeffrey (2007), *Computability and Logic* (5th ed.), Cambridge University Press.

[2] J. P. Burgess (1999), 'Which Modal Logic is the Right One?', *Notre Dame Journal of Formal Logic* 40, pp. 81-93.

[3] C. C. Chang & H. J. Keisler (1990), *Model theory* (3rd ed.), North-Holland.

[4] N. Denyer (1992), 'Pure second-order logic', *Notre Dame Journal of Formal Logic*, 33(2), pp. 220-4.

[5] H.B. Enderton (2001), *A mathematical introduction to logic* (2nd ed.). Harcourt/Academic Press.

[6] J. Etchemendy (1990), *The Concept of Logical Consequence*, Harvard University Press.

[7] S. Givant & P. Halmos (2009), *Introduction to Boolean Algebras*, Springer.

[8] K. Gödel, 'Über formal unentscheidbare Sätze der *Principia Mathematica* und verwandter Systeme I', *Monatshefte für Mathematik und Physik*, 38, pp. 173-98, transl. by J. van Heijenoort as 'On formally undecidable propositions of *Principia Mathematica* and related systems I' and repr. in Gödel's *Collected Works Vol. 1: Publications 1929-1936*, S. Feferman *et al.* eds, 1986, Oxford University Press, pp. 144-195.

[9] R. Goldblatt (1998), *Lectures on the hyperreals*, Springer-Verlag.

[10] O. Griffiths & A. Paseau (2016), 'Isomorphism Invariance and Overgeneration', *Bulletin of Symbolic Logic* 22, pp. 482-503.

[11] P.G. Hinman (2005), *Fundamentals of Mathematical Logic*, A.K. Peters.

[12] W. Hodges (1993), *Model theory*, Cambridge University Press.

[13] W.Hodges (1997), *A shorter model theory*. Cambridge University Press.

[14] G.E. Hughes & M.J. Cresswell (1996), *A New Introduction to Modal Logic*, Routledge.

[15] H. J. Keisler (1965), 'A survey of ultraproducts', in *Logic, Methodology and Philosophy of Science*, North-Holland, pp. 112-26.

[16] A.D. Oliver & T. Smiley (2013), *Plural Logic*, Oxford University Press.

[17] A.C. Paseau (2010), 'Pure second-order logic with second-order identity', *Notre Dame Journal of Formal Logic*, 51(3), pp. 351-60.

[18] A. Robinson (1966), *Non-standard analysis*, North-Holland.

[19] S. Shapiro (1991), *Foundations without Foundationalism*, The Clarendon Press.

[20] B-U. Yi (2006), 'The Logic and Meaning of Plurals. Part II', *Journal of Philosophical Logic* 35, pp. 239-88.