ELEMENTS OF DEDUCTIVE LOGIC

AC PASEAU
Faculty of Philosophy
Oxford University

Lecture Notes
April 2022

Corrections to: paseau@maths.ox.ac.uk

These notes are an updated version of the previous set dating from Hilary 2019. They contain the great majority of the results from the Hilary 2022 lectures. I've omitted some of the accompanying 'chat': motivating remarks, informal discussion, and some examples. If you spot any infelicities—typos, thinkos, notational inconsistencies, etc.—*please* let me know by emailing me at: paseau@maths.ox.ac.uk. I'm grateful to past EDL students and colleagues for comments that improved an earlier draft.[1]

Lecture 1

In an ideal world, we would devote the first couple of lectures to a review of mathematical notation and basic proof techniques (e.g. argument by contraposition, argument by contradiction). There's no time for that, alas. That said, it's worth rehearsing notation we'll use fairly often: $A \subseteq B$ means that A is a (proper or improper) subset of B, in other words that every element of A is an element of B; $A \cup B = \{x : x \in A \text{ or } x \in B\}$ is the set of elements in A or B (or both); $A \cap B = \{x : x \in A \text{ and } x \in B\}$ is the set of elements in both A and B ; and $A \backslash B = \{x : x \in A \text{ and } x \notin B\}$ is the set of elements of A not in B. For simplicity, I'll be casual about use and mention, as I was in the last sentence by omitting quotation marks, and I'll certainly dispense with Quine quotes (if you don't know what they are, don't worry). We'll also assume knowlege of *The Logic Manual*, covered in the Introduction to Logic course.

The main mathematical tool we'll need is proof by induction. I expect most of you have seen this before, so I'll be relatively brief. One standard form of the PMI (Principle of Mathematical Induction), in premise–conclusion form, is:

$$\begin{array}{l} \Phi(0) \\ \forall n\big(\Phi(n) \rightarrow \Phi(n+1)\big) \\ \hline \forall n\Phi(n) \end{array}$$

The quantifiers here range over the natural numbers (0, 1, 2, 3, ...), and Φ is any numerical property. Observe that this statement of the PMI is informal, i.e. formulated in the language of informal mathematics. If/when you take a more advanced course in logic, you'll come across a formal statement of the

[1]By name: Emily Gray, Emma Baldassari, Joe Deakin, Raymond Douglas, Joel David Hamkins, Paolo Marimon and Beau Madison Mount. Apologies to anyone whose name I have inadvertently omitted.

PMI in an extension of $\mathcal{L}_=$ that contains function symbols. Here's another version of the PMI, sometimes called the strong form:

$$\frac{\begin{array}{l}\Phi(0)\\ \forall n\big(\forall k \leqslant n\Phi(k) \to \Phi(n+1)\big)\end{array}}{\forall n\Phi(n)}$$

It's an easy exercise to show that these two versions of the PMI are equivalent (and in particular that they are no different in strength). As you'll see, we'll tend to use the latter more than the former. In lectures, I also mentioned the Least Number Principle, another equivalent of the PMI, which states that if at least one natural number has property Φ then there is a least natural number with property Φ.

Let's now have two examples of the PMI's use, the first from arithmetic, the second from logic.

Proposition 1 *The sum of the natural numbers from* 0 *to* N *is* $\frac{N(N+1)}{2}$.

You may have come across non-inductive proofs of this proposition. Our proof will be by induction.

Proof. First, let's check that 0 has the property. Yes, it does, since the sum of the natural numbers from 0 to 0, namely 0, is indeed $\frac{0(0+1)}{2} = 0$. So we've checked the *base case.*

Next, suppose that the sum of the natural numbers from 0 to N is $\frac{N(N+1)}{2}$. This is the assumption in the *induction step.* On that assumption, we need to prove that the sum of the numbers from 0 to $N+1$ is $\frac{(N+1)((N+1)+1)}{2} = \frac{(N+1)((N+2)}{2}$. Now

$$\sum_{0 \leqslant i \leqslant N+1} i = \Big(\sum_{0 \leqslant i \leqslant N} i\Big) + N + 1 = \frac{N(N+1)}{2} + N + 1 = \frac{(N+1)(N+2)}{2}$$

The middle one of the three equations follows from the *induction hypothesis* and the last one summarises a short algebraic manipulation. Since we've proved the base case and the induction step, we can conclude by the PMI that all numbers have the stated property. ■

The next example of a proof by induction is much more typical of the kind of reasoning used in this course. Before that, though, we need some definitions.

Definition 1 *Let* ϕ *be an* $\mathcal{L}_1$*-formula.* $Conn(\phi)$ *is the set of connectives in* ϕ. $NConn(\phi)$ *is the number of occurrences of connectives in* ϕ. $Atom(\phi)$ *is the set of atoms in* ϕ *and* $SenLett(\phi)$ *is the set of sentence letters in* ϕ.

Suppose for example that ϕ is $\big(((P \wedge Q) \wedge R) \vee P_1\big)$. Then $Conn(\phi) = \{\wedge, \vee\}$, $NConn(\phi) = 3$, and $Atom(\phi) = SenLett(\phi) = \{P, Q, R, P_1\}$. In fact, $Atom(\phi) = SenLett(\phi)$ more generally for all $\mathcal{L}_1$-sentences ϕ. Note that $Conn(\phi)$ is a set of connective types, whereas $NConn(\phi)$ counts the number of connective tokens in ϕ.

Many of the proofs in this course show that all formulas have some property Ψ 'by induction on the complexity of ϕ'. This means that the induction is on $NConn(\phi)$: we prove that a formula with no connectives has the property, and that if a formula with n connectives has the property Ψ then a formula with $n + 1$ connectives also has the property Ψ; we then conclude that all formulas have the property. That gives us a handy way of proving facts about *all* formulas, by associating them with natural numbers in this fashion. In other words, what we're doing is proving that 0 is such that every formula with that many connectives has the property Ψ, and that every number n is such that every formula with that many connectives has the property Ψ, then $n + 1$ is also such that every formula with that many connectives has the property Ψ. So every number n is such that every formula with that many connectives has the property Ψ; so all formulas have the property.

Here's an example of this approach at work.

Proposition 2 *Suppose $Conn(\phi) \subseteq \{\leftrightarrow\}$ Then ϕ is not a contradiction, i.e. there's an $\mathcal{L}_1$-structure in which ϕ is true.*

Proof. We'll use induction on a stronger claim, for reasons that will become apparent. Let $\Phi(n)$ be the claim:

> Suppose $NConn(\phi) = n$ and $Conn(\phi) \subseteq \{\leftrightarrow\}$. Let $\mathcal{A}$ be an $\mathcal{L}_1$-structure such that $|\alpha|_{\mathcal{A}} = T$ for all $\alpha \in Atom(\phi)$. Then $|\phi|_{\mathcal{A}} = T$.

Clearly, the proposition we've set out to prove follows from $\forall n \Phi(n)$. So it will be sufficient to prove $\forall n \Phi(n)$, which we'll do using the second version of the PMI mentioned above.

The base case: suppose $NConn(\phi) = 0$. Then ϕ is an atom. So if $|\alpha|_{\mathcal{A}} = T$ for all $\alpha \in Atom(\phi)$, it follows that $|\phi|_{\mathcal{A}} = T$ since ϕ is an atom of ϕ (indeed the only atom of ϕ).

Now for the induction step. We assume that $\forall k \leqslant n \Phi(k)$. Suppose then that $NConn(\phi) = n + 1$ and $Conn(\phi) \subseteq \{\leftrightarrow\}$. Then $\phi = \phi_1 \leftrightarrow \phi_2$ for some ϕ_1, ϕ_2 such that $Conn(\phi_i) \subseteq \{\leftrightarrow\}$ and $NConn(\phi_i) \leqslant n$, for $i = 1, 2$. (In fact, we know that $NConn(\phi_1) + NConn(\phi_2) = n$.)

Suppose that $\mathcal{A}$ is an $\mathcal{L}_1$-structure such that $|\alpha|_{\mathcal{A}} = T$ for all $\alpha \in Atom(\phi)$. Since $Atom(\phi_1), Atom(\phi_2) \subseteq Atom(\phi)$, it follows that $|\alpha|_{\mathcal{A}} = T$ for all $\alpha \in$

$Atom(\phi_i)$, for $i = 1, 2$. Hence by the induction hypothesis we know that $|\phi_1|_{\mathcal{A}} = T$ and $|\phi_2|_{\mathcal{A}} = T$. The truth-table for $\leftrightarrow$ then yields, for this same structure $\mathcal{A}$:

$$|\phi|_{\mathcal{A}} = |\phi_1 \leftrightarrow \phi_2|_{\mathcal{A}} = T$$

By the Principle of Mathematical Induction, $\forall n \Phi(n)$, which proves our proposition. ■

Notice two facts about this proof. We had to use the second rather than the first version of the PMI because the inductive step involves two formulas ϕ_1 and ϕ_2 whose complexity is unknown. All we know about the formulas' complexity is that $NConn(\phi_1) + NConn(\phi_2) = n$, so our inductive hypothesis must perforce be that the required property holds for formulas of *any* complexity $\leqslant n$. The second point is that we proved the proposition from a stronger assumption: whereas the proposition states that a formula whose only connective if any is $\leftrightarrow$ is true in *some* $\mathcal{L}_1$-structure, our proof showed that such a formula is true in the structure that maps all the formula's atoms to T. If in the inductive step all we could draw on were the hypotheses that ϕ_1 is true in some structure $\mathcal{A}_1$ and that ϕ_2 is true in some structure $\mathcal{A}_2$ then we would have been stuck. For without knowing that $\mathcal{A}_1$ and $\mathcal{A}_2$ agree on the atoms of ϕ, we couldn't use them to define an $\mathcal{L}_1$-structure in which $\phi = \phi_1 \leftrightarrow \phi_2$ is true. Hence the moral: choose your inductive hypotheses wisely!

At the end of the first lecture, we covered some of the material to follow in the next section. But I'll end the section here since it's a natural break.

Lecture 2

Wherein we make a start on metalogic proper and prove the Interpolation Theorem for propositional logic. Henceforth we'll use 1 and 0 as truth-values rather than T and F, as we can algebraically manipulate the former more easily.

It's useful to consider a slight extension of $\mathcal{L}_1$ which I'll (unimaginatively) call $\mathcal{L}_1^+$. This logic expands the language of $\mathcal{L}_1$ with two atomic sentences $\top$ and $\bot$ such that $|\top|_{\mathcal{A}} = 1$ and $|\bot|_{\mathcal{A}} = 0$ for all structures $\mathcal{A}$. Note that $\top$ and $\bot$ are atoms but not sentence letters. So whereas in $\mathcal{L}_1$, $SenLett(\phi)$ is always identical to $Atom(\phi)$, these notions can come apart in $\mathcal{L}_1^+$; e.g. $SenLett(P \to \top) = \{P\}$ whereas $Atom(P \to \top) = \{P, \top\}$. We'll assume that $\mathcal{L}_1^+$-consequence respects $\mathcal{L}_1$-consequence, i.e. that a subset Γ of $\mathcal{L}_1$-sentences entails an $\mathcal{L}_1$-sentence ϕ in $\mathcal{L}_1^+$ iff Γ entails ϕ in $\mathcal{L}_1$.

There now follows a sequence of lemmas and definitions, of great use for the rest of the course. Intuitively, the next lemma tells us that the truth-value of a sentence in a structure is fixed by the truth-value of its sentence letters in that structure. Although that should seem obvious enough, it does require proof. The lemma, note, applies to both $\mathcal{L}_1$ and $\mathcal{L}_1^+$.

Lemma 3 (Relevance Lemma) *Suppose $|\alpha|_{\mathcal{A}} = |\alpha|_{\mathcal{B}}$ for all $\alpha \in SenLett(\phi)$. Then $|\phi|_{\mathcal{A}} = |\phi|_{\mathcal{B}}$.*

Proof We prove the result by induction on the complexity of ϕ, that is, by induction on $NConn(\phi)$. From now on, it's understood that that is what we mean by saying 'by induction on the complexity of ϕ'.

For the base case, we consider ϕ of complexity 0, i.e. sentences with no connectives. So ϕ is either a sentence letter or, in the case of $\mathcal{L}_1^+$, it could also be one of $\top$ or $\bot$. If $\mathcal{A}$ and $\mathcal{B}$ agree on the sentence letters of ϕ and ϕ *is* a sentence letter, then $\mathcal{A}$ and $\mathcal{B}$ must agree on ϕ. And if ϕ is $\top$ or $\bot$, then $\mathcal{A}$ and $\mathcal{B}$ agree on ϕ since all valuations agree on $\top$ or $\bot$ ($|\top|_{\mathcal{A}} = |\top|_{\mathcal{B}} = 1$ for all $\mathcal{A}$ and $\mathcal{B}$, and $|\bot|_{\mathcal{A}} = |\bot|_{\mathcal{B}} = 0$ for all $\mathcal{A}$ and $\mathcal{B}$).

For the inductive step, we use the 'strong' form of induction once more. So suppose the relevance property holds for all formulas of complexity $\leqslant n$. A formula ϕ of complexity $n + 1$ is of the form $\neg\psi$ or $\phi_1 \wedge \phi_2$ or $\phi_1 \vee \phi_2$ or $\phi_1 \to \phi_2$ or $\phi_1 \leftrightarrow \phi_2$, where ψ, ϕ_1 and ϕ_2 are all of complexity $\leqslant n$. If $\mathcal{A}$ and $\mathcal{B}$ agree on the atoms of ϕ then they must agree on the atoms of ψ, ϕ_1 and ϕ_2 in each of these five cases. By the inductive hypothesis, that means that the truth-value of ψ in $\mathcal{A}$ is the same as its truth-value in $\mathcal{B}$ (in the first case), and the truth-values of ϕ_1 and ϕ_2 in $\mathcal{A}$ are the same as their truth-values in $\mathcal{B}$ (in the other four cases). Since the truth-value of $\neg\psi$ is determined by that of ψ, and the truth-value of each of $\phi_1 \wedge \phi_2$ and $\phi_1 \vee \phi_2$ and $\phi_1 \rightarrow \phi_2$ and $\phi_1 \leftrightarrow \phi_2$ is likewise determined by those of ϕ_1 and ϕ_2, $\mathcal{A}$ and $\mathcal{B}$ agree on ϕ. In other words, $|\phi|_{\mathcal{A}} = |\phi|_{\mathcal{B}}$. ■

One way to understand the Relevance Lemma is that it underwrites the use of the usual, finite, truth-tables. For consider a truth-table such as:

P	Q	$P \wedge Q$
1	1	1
1	0	0
0	1	0
0	0	0

The usual way of doing things—including ours—has each of the rows representing not a structure or valuation but a class of stuctures, all of which agree on the truth-values of P and Q (but which may differ on the truth-values of other sentence letters not represented in the table). Yet how do we know that all structures that agree on the truth-values of P and Q agree on the truth-value of $P \wedge Q$? The answer: by the Relevance Lemma! So if you thought that the Relevance Lemma didn't require proof, it's perhaps because you hadn't fully appreciated that propositional structures are specified by their assignment of truth-values to *all* sentence letters.[2]

We now offer a sequence of definitions, in which Γ is a set of sentences and $\mathcal{A}$ is a structure. The definitions, which apply to both $\mathcal{L}_1$ and $\mathcal{L}_1^+$, set out some widely-used notational variants for concepts you've already come across.

Definition 2 *$\mathcal{A} \vDash \phi$ means that $|\phi|_{\mathcal{A}} = 1$; we say '$\mathcal{A}$ satisfies ϕ'. And we write $\mathcal{A} \vDash \Gamma$ to abbreviate: for all $\gamma \in \Gamma$, $\mathcal{A} \vDash \gamma$.*

[2](A non-examinable note for students who know about countable and uncountable sets.) Each row in fact represents uncountably many structures. To see that there are uncountably many $\mathcal{L}_1$-structures, observe that there is an onto map from the class of structures to the numbers in the closed interval $[0, 1]$, which is well known to be uncountable. Simply enumerate the sentence letters in a list $\alpha_1, \cdots, \alpha_n, \cdots$ and map structure $\mathcal{A}$ to the real number $\sum_{1 \leqslant i} |\alpha_i|_{\mathcal{A}} 2^{-i}$, which you can think of as the number's binary representation. As an aside, we note that this particular map is onto but not one-one since for example $1 \cdot 2^{-1} + 0 \cdot 2^{-2} + 0 \cdot 2^{-3} + \ldots = 0 \cdot 2^{-1} + 1 \cdot 2^{-2} + 1 \cdot 2^{-3} + \ldots$. (In other words, the real number corresponding to the binary decimal $0.1000000\ldots$ is the same as that corresponding to the binary decimal $0.0111111\ldots$, just as $0.5 = 0.4999999\ldots$.) There are in fact one-to-one *and* onto maps from the set of $\mathcal{L}_1$-structures to $[0, 1]$ but they are slightly trickier to describe than the fairly simple onto map just mentioned.

Γ is satisfiable iff it's semantically consistent, i.e. just when there's a structure $\mathcal{A}$ such that $\mathcal{A} \vDash \Gamma$. Γ is unsatisfiable otherwise, i.e. just when for all $\mathcal{A}$, $\mathcal{A} \nvDash \Gamma$.

ϕ is a tautology just when, for all $\mathcal{A}$, $\mathcal{A} \vDash \phi$. And ϕ is a contradiction just when, for all $\mathcal{A}$, $\mathcal{A} \nvDash \phi$.

$\Gamma \vDash \phi$ means: for all $\mathcal{A}$ if $\mathcal{A} \vDash \Gamma$ then $\mathcal{A} \vDash \phi$.

ϕ and ψ are logically equivalent just when: for all $\mathcal{A}$, $\mathcal{A} \vDash \phi$ iff $\mathcal{A} \vDash \psi$.

We note that the symbol '$\vDash$' is ambiguous in logic. '$\mathcal{A} \vDash \phi$' means that the structure $\mathcal{A}$ satisfies the sentence ϕ, whereas '$\Gamma \vDash \phi$' means that all structures that satisfy Γ also satisfy ϕ. In practice, it would be hard to confuse the two, as both notation and context should make it clear whether the symbol to the left of '$\vDash$' denotes a set of sentences or a structure. In lectures, I also used $\phi \dashv\vDash \psi$ to mean that ϕ and ψ are logically equivalent, but I won't do so here.

Let's now move on to substitution. We'd like to capture the idea that we can replace all occurrences of a sentence letter, or more generally sequence of letters, with some formula(s). This motivates the following definition.

Definition 3 *Let Sen($\mathcal{L}$) be the set of logic $\mathcal{L}$'s sentences. Let $\pi : SenLett(\mathcal{L}_1^+) \to Sen(\mathcal{L}_1^+)$, i.e. π is a function which maps the sentence letters of $\mathcal{L}_1^+$ to sentences of $\mathcal{L}_1^+$. Then we extend π to a map defined on* all *the sentences of $\mathcal{L}_1^+$ which by abuse of notation we also label π and which we represent by a superscript, as follows:*

$$
\begin{aligned}
&\alpha^\pi = \pi(\alpha) \text{ if } \alpha \in SenLett(\mathcal{L}_1^+)\\
&\top^\pi = \top\\
&\bot^\pi = \bot\\
&(\neg\phi)^\pi = \neg(\phi^\pi)\\
&(\phi \wedge \psi)^\pi = \phi^\pi \wedge \psi^\pi\\
&(\phi \vee \psi)^\pi = \phi^\pi \vee \psi^\pi\\
&(\phi \to \psi)^\pi = \phi^\pi \to \psi^\pi\\
&(\phi \leftrightarrow \psi)^\pi = \phi^\pi \leftrightarrow \psi^\pi
\end{aligned}
$$

We also extend π to sets of sentences in the obvious way: $\Gamma^\pi = \{\gamma^\pi : \gamma \in \Gamma\}$.

Given such a map $\pi : SenLett(\mathcal{L}_1^+) \to Sen(\mathcal{L}_1^+)$, we define a corresponding map on structures. Given any structure $\mathcal{A}$, let $\mathcal{A}^\pi$ be the structure defined by the following stipulation:

$$|\alpha|_{\mathcal{A}^\pi} = |\pi(\alpha)|_{\mathcal{A}} \text{ for all } \alpha \in SenLett(\mathcal{L}_1^+)$$

The value $|\phi|_{\mathcal{A}^\pi}$ for a complex sentence ϕ is then fixed by the values of $|\alpha|_{\mathcal{A}^\pi}$ for ϕ's sentence letters (by the Relevance Lemma). The following lemma shows that we can 'move the superscript down from the sentence to the structure'.

Lemma 4 (Substitution Lemma) $|\phi^\pi|_{\mathcal{A}} = |\phi|_{\mathcal{A}^\pi}$, *for all* ϕ *and* $\mathcal{A}$.

Proof By induction on the complexity of ϕ, as usual.

For the base case, suppose that ϕ is a sentence letter α. Then by the definition of α^π, $|\alpha^\pi|_{\mathcal{A}} = |\pi(\alpha)|_{\mathcal{A}}$. And by the definiton of $\mathcal{A}^\pi$, $|\alpha|_{\mathcal{A}^\pi} = |\pi(\alpha)|_{\mathcal{A}}$ also. The case of $\top$ is straightforward: $\top^\pi = \top$ and $\top$ is true in all structures; similarly, $\bot^\pi = \bot$ and $\bot$ is false in all structures.

For the inductive step, we'll need to consider five cases. I'll do two of them and leave the remaining three to you. The first case is when $\phi = \neg\psi$. Using the inductive hypothesis and the fact that $|\neg\chi|_{\mathcal{A}} = 1 - |\chi|_{\mathcal{A}}$ for any sentence χ and structure $\mathcal{A}$:

$$\begin{aligned}
|\phi^\pi|_{\mathcal{A}} &= |(\neg\psi)^\pi|_{\mathcal{A}} \\
&= |\neg\psi^\pi|_{\mathcal{A}} \\
&= 1 - |\psi^\pi|_{\mathcal{A}} \\
&= 1 - |\psi|_{\mathcal{A}^\pi} \\
&= |\neg\psi|_{\mathcal{A}^\pi} \\
&= |\phi|_{\mathcal{A}^\pi}
\end{aligned}$$

The second case, in which $\phi = \phi_1 \wedge \phi_2$, is very similar. We use the inductive hypothesis and the fact that $|\chi_1 \wedge \chi_2|_{\mathcal{A}} = |\chi_1|_{\mathcal{A}}|\chi_2|_{\mathcal{A}}$:

$$\begin{aligned}
|\phi^\pi|_{\mathcal{A}} &= |(\phi_1 \wedge \phi_2)^\pi|_{\mathcal{A}} \\
&= |\phi_1^\pi \wedge \phi_2^\pi|_{\mathcal{A}} \\
&= |\phi_1^\pi|_{\mathcal{A}}|\phi_2^\pi|_{\mathcal{A}} \\
&= |\phi_1|_{\mathcal{A}^\pi}|\phi_2|_{\mathcal{A}^\pi} \\
&= |\phi_1 \wedge \phi_2|_{\mathcal{A}^\pi} \\
&= |\phi|_{\mathcal{A}^\pi}
\end{aligned}$$

The other cases are entirely analogous. ■

We draw a couple of corollaries from the lemma. Their proof is left as an exercise.

Corollary 5 *If* $\Gamma \vDash \phi$ *then* $\Gamma^\pi \vDash \phi^\pi$.

Corollary 6 *Suppose* $\vDash \phi_i \leftrightarrow \psi_i$ *(i.e.* ϕ_i *and* ψ_i *are logical equivalents) for* $1 \leqslant i \leqslant N$. *Let* $\chi(\phi_1/\alpha_1, \cdots, \phi_N/\alpha_N)$ *be the formula obtained by replacing all occurrences (if any) of the sentence letter* α_i *in* χ *by* ϕ_i*, for* $1 \leqslant i \leqslant N$*; and let* $\chi(\psi_1/\alpha_1, \cdots, \psi_N/\alpha_N)$ *be the formula obtained by replacing all occurrences (if any) of the sentence letter* α_i *in* χ *by* ψ_i*, for* $1 \leqslant i \leqslant N$. *Note that this is simultaneous, not sequential, substitution in the sense of* π. *Then* $\vDash \chi(\phi_1/\alpha_1, \cdots, \phi_N/\alpha_N) \leftrightarrow \chi(\psi_1/\alpha_1, \cdots, \psi_N/\alpha_N)$.

We now come to the highlight of this lecture.

Theorem 7 (Interpolation Theorem for $\mathcal{L}_1^+$) *Suppose* $\phi \vDash \psi$. *Then there is a sentence* λ *such that: (i) Senlett*$(\lambda) \subseteq$ *Senlett*$(\phi) \cap$ *Senlett*(ψ)*; (ii)* $\phi \vDash \lambda$*; and (iii)* $\lambda \vDash \psi$. λ *is known as an* interpolant *for the sequent* $\phi \vDash \psi$.

Proof We may assume that *Senlett*$(\phi)\backslash$*Senlett*(ψ) is non-empty; for if it's empty then *Senlett*$(\phi) \subseteq$ *Senlett*(ψ), in which case we may take ϕ as our interpolant.

We define a set Π of substitutions as follows:

> $\Pi = \{\pi : \pi(\alpha) \in \{\top, \bot\}$ for all $\alpha \in$ *Senlett*$(\phi)\backslash$*Senlett*(ψ) and $\pi(\alpha) = \alpha$ if α is a sentence letter not in *Senlett*$(\phi)\backslash$*Senlett*$(\psi)\}$

Intuitively, a substitution π gets rid of the sentence letters in ϕ but not ψ by replacing each of them by either $\top$ or $\bot$. Π then consists of all the possible ways of doing so: it's the set of all substitutions of this kind.

Given Π, we define λ as follows:

$$\lambda = \bigvee_{\pi \in \Pi} \phi^\pi$$

Thus if *Senlett*$(\phi)\backslash$*Senlett*(ψ) is a set of size n, λ will be a disjunction of 2^n sentences. It remains to prove three things: (i) *Senlett*$(\lambda) \subseteq$ *Senlett*$(\phi) \cap$ *Senlett*(ψ); (ii) $\phi \vDash \lambda$; and (iii) $\lambda \vDash \psi$.

(i) That *Senlett*$(\lambda) \subseteq$ *Senlett*$(\phi) \cap$ *Senlett*(ψ) is immediate from the definition of λ, since *SenLett*$(\phi^\pi) \subseteq$ *Senlett*$(\phi) \cap$ *Senlett*(ψ) for each disjunct ϕ^π in λ.

(ii) Suppose $\mathcal{A} \models \phi$. Consider the substitution function π defined as follows on sentence letters α:

$$\pi(\alpha) = \begin{cases} \top & \text{if } \alpha \in \mathit{Senlett}(\phi) \backslash \mathit{Senlett}(\psi) \text{ and } |\alpha|_{\mathcal{A}} = 1 \\ \bot & \text{if } \alpha \in \mathit{Senlett}(\phi) \backslash \mathit{Senlett}(\psi) \text{ and } |\alpha|_{\mathcal{A}} = 0 \\ \alpha & \text{otherwise, i.e. if } \alpha \notin \mathit{Senlett}(\phi) \backslash \mathit{Senlett}(\psi) \end{cases}$$

Clearly, $\mathcal{A} = \mathcal{A}^{\pi}$ since they agree on all sentence letters (formally, an inductive argument would be required here). Now by the Substitution Lemma, $|\phi^{\pi}|_{\mathcal{A}} = |\phi|_{\mathcal{A}^{\pi}}$. So since $\mathcal{A} = \mathcal{A}^{\pi}$, $|\phi^{\pi}|_{\mathcal{A}} = |\phi|_{\mathcal{A}} = 1$. Since ϕ^{π} is a disjunct in λ, λ is true in $\mathcal{A}$. We conclude that $\phi \models \lambda$.

(iii) Suppose $\mathcal{A} \models \lambda$. Since $\lambda = \bigvee_{\pi \in \Pi} \phi^{\pi}$, $\mathcal{A} \models \phi^{\pi}$ for some $\pi \in \Pi$, i.e. $|\phi^{\pi}|_{\mathcal{A}} = 1$ for this π. By the Substitution Lemma, it follows that $|\phi|_{\mathcal{A}^{\pi}} = 1$ for this π. As $\phi \models \psi$, we deduce that $|\psi|_{\mathcal{A}^{\pi}} = 1$. Now since $\mathcal{A}^{\pi}$ and $\mathcal{A}$ agree on all the sentence letters in ψ, we deduce from the Relevance Lemma that $|\psi|_{\mathcal{A}} = 1$. Thus $\lambda \models \psi$. ■

It's worth thinking about the relative strength of interpolants for a given sequent. How does λ as defined in the theorem's proof compare to other interpolants for the sequent $\phi \models \psi$?

The interpolation theorem for $\mathcal{L}_1$ is a relatively straightforward corollary of the interpolation theorem for $\mathcal{L}_1^+$. To prove it, we require a couple of further lemmas.

Lemma 8 *Suppose $\phi \in Sen(\mathcal{L}_1^+)$ s.t. $Atom(\phi) \subseteq \{\top, \bot\}$. Then ϕ is logically equivalent (in $\mathcal{L}_1^+$) to $\top$ or to $\bot$.*

Proof. By induction on the complexity of ϕ; left as an exercise. Or prove this more directly from the Relevance Lemma, since any two structures must agree on the truth-value of ϕ. ■

Lemma 9 *Suppose $\phi \in Sen(\mathcal{L}_1^+)$ s.t. $SenLett(\phi)$ is non-empty. Then ϕ is logically equivalent (in $\mathcal{L}_1^+$) to a formula ψ such that $Atom(\psi) = SenLett(\phi)$.*

Proof. By induction on the complexity of ϕ; left as an exercise. ■

Theorem 10 (Interpolation Theorem for $\mathcal{L}_1$) *Suppose $\phi \models \psi$, with ϕ not a contradiction and ψ not a tautology. Then there is a λ such that: (i) $\phi \models \lambda \models \psi$; and (ii) $Senlett(\lambda) \subseteq Senlett(\phi) \cap Senlett(\psi)$.*

Proof. Suppose $\phi \vDash \psi$, with $\phi, \psi \in Sen(\mathcal{L}_1)$. Consider this as a claim about $\mathcal{L}_1^+$-sentences (since every $\mathcal{L}_1$-sentence is also an $\mathcal{L}_1^+$-sentence). By the Interpolation Theorem for $\mathcal{L}_1^+$, there is an $\mathcal{L}_1^+$-interpolant λ for the sequent $\phi \vDash \psi$ whose set of sentence letters $SenLett(\lambda)$ is a subset of the intersection of $SenLett(\phi)$ and $SenLett(\psi)$. There are two possibilies, depending on whether $SenLett(\lambda)$ is empty or not.

Suppose $SenLett(\lambda)$ is empty. By Lemma 8, λ is then logically equivalent to $\top$ or $\bot$. But if ϕ entails $\bot$ then ϕ is a contradiction; and if $\top$ entails ψ then ψ is a tautology; hence this case is excluded.

It follows then that $SenLett(\lambda)$ must be non-empty. So by Lemma 9, λ is equivalent to a sentence whose atom set is $SenLett(\lambda)$; the latter sentence is the required $\mathcal{L}_1$-interpolant. ■

The theorem is 'best possible' because we can't remove the restriction to ϕ not being a contradiction or ψ not being a tautology. To see this, consider for example the sequents $P \vDash Q \vee \neg Q$ and $P \wedge \neg P \vDash Q$.

Lecture 3

Last time we proved the Interpolation Theorem for $\mathcal{L}_1^+$, and, as a corollary, the Interpolation Theorem for $\mathcal{L}_1$. The topic of today's lecture is *duality.* It's a fairly self-contained topic, popular with EDL examiners, so well worth getting under your belt. To introduce it, it will be useful to know something about truth-functions.

Definition 4 *A truth-function is a function from an n-tuple of 1s and 0s to 1 or 0. A bit more formally, it is a function from $\{0,1\}^n$ to $\{0,1\}$; here $\{0,1\}^n$ is the n-fold Cartesian product of $\{0,1\}$ with itself: $\underbrace{\{0,1\} \times \cdots \times \{0,1\}}_{n}$.*

If c is an n-place propositional connective, the n-ary truth-function f_c associated with c is defined by:

> *for all structures $\mathcal{A}$ and formulas $\phi_1, \cdots \phi_n$, $f_c(|\phi_1|_\mathcal{A}, \cdots, |\phi_n|_\mathcal{A}) = |c(\phi_1, \cdots, \phi_n)|_\mathcal{A}$.*

1 and 0 can be considered as the constantly true and constantly false 0-ary truth-functions respectively.

Every formula ϕ is also associated with a unique truth-function f_ϕ defined by:

> *Let SenLett$(\phi) = \{\alpha_1, \cdots, \alpha_n\}$. Then for all structures $\mathcal{A}$, $f_\phi(|\alpha_1|_\mathcal{A}, \cdots, |\alpha_n|_\mathcal{A}) = |\phi|_\mathcal{A}$*

In the case of $\mathcal{L}_1^+$: if SenLett(ϕ) is empty, we associate it with the 0-ary truth-function 1 or the 0-ary truth-function 0, as the case may be. (Lemma 8 justifies this definition.)

For example, $f_\wedge(t_1, t_2) = t_1 t_2, f_\vee(t_1, t_2) = max\{t_1, t_2\} = t_1 + t_2 - t_1 t_2$ and $f_\neg(t) = 1 - t$. If $\phi = P \wedge Q$ then $f_\phi = f_\wedge$, and so on.

In lectures, we then motivated the definition of the dual of a connective by playing around with truth-tables. We'll dispense with the pictures here, but state the key idea. To find the dual of a connective c, take its truth-table and turn all the 1s into 0s and 0s into 1s (in all the 'input' columns as well as the 'output' column); this defines a new connective, c's dual, which we call c^*. Applying this procedure to $\wedge, \vee$ and $\neg$ we discover that $\wedge^* = \vee$, $\vee^* = \wedge$, and $\neg^* = \neg$; the last equation shows that $\neg$ is self-dual. We also observe that $P \rightarrow^* Q$ is logically equivalent to $\neg(Q \rightarrow P)$ and that $P \leftrightarrow^* Q$ is logically equivalent to $\neg(Q \leftrightarrow P)$. Flipping a single 1 turns it into a 0 and a single 0 into a 1, which also justifies: $\top^* = \bot$ and $\bot^* = \top$.

The most natural setting for duality is a propositional logic in which there is a unique way to represent each connective's dual. In the propositional logics $\mathcal{L}_1$ or $\mathcal{L}_1^+$, defining a connective's dual involves some arbitrary choices. For example, we can define $P \to^* Q$ as $\neg(Q \to P)$, or as $\neg P \wedge Q$, or in countless other ways. There is no such thing as the 'right' definition, only a choice to be made on pragmatic grounds. To avoid this, and to introduce you to a third propositional logic, we work in $\mathcal{L}_1^{++}$. We'll define this logic informally, as we did $\mathcal{L}_1^+$.

Definition 5 *$\mathcal{L}_1^{++}$ is a propositional logic with the same set of atoms as $\mathcal{L}_1^+$, i.e. the sentence letters of $\mathcal{L}_1$ as well as $\top$ and $\bot$. For each n-place truth-function f, $\mathcal{L}_1^{++}$ has a single n-place connective c s.t. $f = f_c$, i.e. f is the truth-function associated with c, for $n \geqslant 1$. The syntax and semantics of $\mathcal{L}_1^{++}$ is that of $\mathcal{L}_1$ with the obvious changes.*

As with $\mathcal{L}_1^+$, I encourage you to come up with $\mathcal{L}_1^{++}$'s full syntactic and semantic specification. You'll notice that in this case, the recursive syntactic and semantic clauses can't be individually specified, one per connective, since there are infinitely many such connectives. So the clauses must be schematic. No problem with this, of course, since even the specification of $\mathcal{L}_1$'s syntax and semantics is schematic when it comes to the base cases, seeing as there are infinitely many sentence letters.

In the definition of $\mathcal{L}_1^{++}$, we've taken $\top$ and $\bot$ as atoms; we could equally have taken them to be 0-place connectives, and occasionally we'll regard them as such. We'll also continue to use the symbols $\wedge, \vee, \neg, \to$ and $\leftrightarrow$ for the five connectives $\mathcal{L}_1^{++}$ has in common with $\mathcal{L}_1$ (and $\mathcal{L}_1^+$).

We now define three dual maps on $\mathcal{L}_1^{++}$. We begin with the dual of a connective; then that of a sentence; and we end with the dual of an $\mathcal{L}_1^{++}$-structure. The highlight of the lecture and its culmination will be the Duality Theorem, which links the last two notions.

Definition 6 (Dual of a connective) *Let c be an n-place connective in $\mathcal{L}_1^{++}$ with associated truth-function f_c. Then its dual c^* is the $\mathcal{L}_1^{++}$-connective whose associated truth-function f_{c^*} is defined by:*

$$f_{c^*}(t_1, \cdots, t_n) = 1 - f_c(1 - t_1, \cdots, 1 - t_n)$$

for all truth-values $t_1, \cdots t_n$.

We note that c^* exists and that it's unique. In languages such as $\mathcal{L}_1$ which lack a primitive symbol for some connectives' duals, we'd have to define c^*

in some arbitrary fashion, as mentioned, and of course the dual map would only be defined on the connectives the language happens to contain.

Let's have some examples, starting with negation. (In these examples, construe $\mathcal{A}$ as a variable ranging over $\mathcal{L}_1^{++}$-structures.) From our definitions: $|\neg^*\phi|_{\mathcal{A}} = 1 - (f_{\neg}(1 - |\phi|_{\mathcal{A}})) = 1 - (1 - (1 - |\phi|_{\mathcal{A}})) = 1 - |\phi|_{\mathcal{A}} = |\neg\phi|_{\mathcal{A}}$. This calculation shows that $\neg^* = \neg$, or in other words that $\neg$ is self-dual.

For our second example, let's check that $\wedge^* = \vee$. Here's the calculation: $f_{\wedge^*}(|\phi_1|_{\mathcal{A}}, |\phi_2|_{\mathcal{A}}) = 1 - f_{\wedge}(1 - |\phi_1|_{\mathcal{A}}, 1 - |\phi_2|_{\mathcal{A}}) = 1 - (1 - |\phi_1|_{\mathcal{A}})(1 - |\phi_2|_{\mathcal{A}}) = |\phi_1|_{\mathcal{A}} + |\phi_2|_{\mathcal{A}} - |\phi_1|_{\mathcal{A}}|\phi_2|_{\mathcal{A}} = f_{\vee}(|\phi_1|_{\mathcal{A}}, |\phi_2|_{\mathcal{A}})$. Thus $\wedge^* = \vee$. We may similarly check that $\vee^* = \wedge$. In fact, we can deduce that $\vee^* = \wedge$ from $\wedge^* = \vee$ and the lemma to follow.

Lemma 11 *For any connective c of $\mathcal{L}_1^{++}$, $c^{**} = c$.*

Proof

$$\begin{aligned} f_{c^{**}}(t_1, \cdots, t_n) &= 1 - f_{c^*}(1 - t_1, \cdots, 1 - t_n) \\ &= 1 - \big(1 - f_c(1 - (1 - t_1), \cdots, 1 - (1 - t_n))\big) \\ &= f_c(t_1, \cdots, t_n) \ \blacksquare \end{aligned}$$

We've defined the dual of a connective, so now let's have the dual of a sentence and of a structure.

Definition 7 *Given an $\mathcal{L}_1^{++}$-structure $\mathcal{A}$, we define its dual $\mathcal{A}^*$ by stipulating that $|\alpha|_{\mathcal{A}^*} = 1 - |\alpha|_{\mathcal{A}}$ for all sentence letters α.*

We also define the dual ϕ^* of an $\mathcal{L}_1^{++}$-sentence ϕ recursively. For any sentence letter α, $\alpha^* = \alpha$; also, $\top^* = \bot$, $\bot^* = \top$. If $\phi = c(\phi_1, \cdots, \phi_n)$ then $\big(c(\phi_1, \cdots, \phi_n)\big)^* = c^*(\phi_1^*, \cdots, \phi_n^*)$.

We note that the definition of a sentence's dual depends on that of a connective's dual. Note also that in the first definition we need only define the truth-value of sentence letters in $\mathcal{A}^*$, since the truth-values of all other sentences are determined by these. (We also know that $|\top|_{\mathcal{A}^*} = 1$ since this holds for all structures, and similarly $|\bot|_{\mathcal{A}^*} = 0$.)

We saw earlier that the double dual of any connective is itself. The double dual of any sentence is also itself.

Proposition 12 *For all $\mathcal{L}_1^{++}$-sentences ϕ, $\phi^{**} = \phi$.*

Proof. We prove this by induction on the complexity of ϕ. Base case: a sentence letter's dual is itself, so its double dual is also itself. $\top^{**} = \bot^* = \top$ and similarly $\bot^{**} = \top^* = \bot$.

As for the induction step:

$$\begin{aligned}\big(c(\phi_1, \cdots, \phi_n)\big)^{**} &= \big(c^*(\phi_1^*, \cdots, \phi_n^*)\big)^* \\ &= c^{**}(\phi_1^{**}, \cdots, \phi_n^{**}) \\ &= c^{**}(\phi_1, \cdots, \phi_n) \\ &= c(\phi_1, \cdots, \phi_n)\end{aligned}$$

The first two equations are definitional, the third uses the inductive hypothesis and the fourth the fact that $c^{**} = c$. ■

We note a trivial corollary: ϕ^{**} is logically equivalent to ϕ. In propositional languages other than $\mathcal{L}_1^{++}$, the double dual of a sentence may not be the sentence itself; but it should at least be equivalent to it. Time now for the most important result about duality.

Theorem 13 (Duality Theorem) *For all $\mathcal{L}_1^{++}$-structures $\mathcal{A}$ and $\mathcal{L}_1^{++}$-sentences ϕ,*

$$|\phi^*|_{\mathcal{A}} + |\phi|_{\mathcal{A}^*} = 1$$

Proof By induction on the complexity of ϕ. For the base case, we note from definitions that if α is a sentence letter then $|\alpha^*|_{\mathcal{A}} = |\alpha|_{\mathcal{A}} = 1 - |\alpha|_{\mathcal{A}^*}$. Also $|\top^*|_{\mathcal{A}} = |\bot|_{\mathcal{A}} = 0$ and $|\bot^*|_{\mathcal{A}} = |\top|_{\mathcal{A}} = 1$, so each of $\top$ and $\bot$ also has the noted property.

For the induction step, let ϕ be $c(\phi_1, \cdots, \phi_n)$. Then:

$$\begin{aligned}|\phi^*|_{\mathcal{A}} &= |\big(c(\phi_1, \cdots, \phi_n)\big)^*|_{\mathcal{A}} \\ &= |c^*(\phi_1^*, \cdots, \phi_n^*)|_{\mathcal{A}} \\ &= f_{c^*}(|\phi_1^*|_{\mathcal{A}}, \cdots, |\phi_n^*|_{\mathcal{A}}) \\ &= 1 - f_c(1 - |\phi_1^*|_{\mathcal{A}}, \cdots, 1 - |\phi_n^*|_{\mathcal{A}}) \\ &= 1 - f_c(|\phi_1|_{\mathcal{A}^*}, \cdots, |\phi_n|_{\mathcal{A}^*}) \\ &= 1 - |c(\phi_1, \cdots, \phi_n)|_{\mathcal{A}^*} \\ &= 1 - |\phi|_{\mathcal{A}^*}\end{aligned}$$

The main step, from the fourth to the fifth line, uses the induction hypothesis. This proves the result. ■

We note a corollary of the theorem.

Corollary 14 *If $\phi \vDash \psi$ then $\psi^* \vDash \phi^*$.*

Proof Assume $\phi \models \psi$. If $|\psi^*|_{\mathcal{A}} = 1$ then by the Duality Theorem, $|\psi|_{\mathcal{A}^*} = 0$. If $|\psi|_{\mathcal{A}^*} = 0$ then by the assumption, $|\phi|_{\mathcal{A}^*} = 0$. It follows from the Duality Theorem that $|\phi^*|_{\mathcal{A}} = 1$. Since $\mathcal{A}$ is any $\mathcal{L}_1^{++}$-structure, we have proved that $\psi^* \models \phi^*$. ■

Lecture 4

Last time was all about duality. Today, we'll look at Expressive Adequacy and the Compactness Theorem. Recall from the first lecture that we can't build a contradiction solely using the connective $\leftrightarrow$. Thus the set $\{\leftrightarrow\}$ is not expressively adequate. So which sets of connectives are? First, let's get clear on what the question is.

Definition 8 ($\mathcal{L}_1$, $\mathcal{L}_1^+$, $\mathcal{L}_1^{++}$) *ϕ expresses the truth-function f just when*

$$\text{for all structures } \mathcal{A}, \ |\phi(\alpha_1, \ldots, \alpha_n)|_{\mathcal{A}} = f(|\alpha_1|_{\mathcal{A}}, \ldots, |\alpha_n|_{\mathcal{A}})$$

where $\alpha_1, \ldots, \alpha_n$ are the sentence letters in ϕ.

The following definition is for $\mathcal{L}_1$, but is easily adapted to $\mathcal{L}_1^+$ or $\mathcal{L}_1^{++}$.

Definition 9 (Expressive Adequacy) *A set C of connectives is expressively adequate just when, for every truth-function f, there is a $\phi \in Sen(\mathcal{L}_1)$ with $Conn(\phi) \subseteq C$ that expresses f.*

We're now in position to prove a key theorem about expressive adequacy.

Theorem 15 ($\mathcal{L}_1$, $\mathcal{L}_1^+$, $\mathcal{L}_1^{++}$) *The set $\{\neg, \wedge, \vee\}$ is expressively adequate.*

Proof. The proof is for $\mathcal{L}_1$ but is easily adapted to $\mathcal{L}_1^+$ or $\mathcal{L}_1^{++}$. First, let's get 0-place truth-functions out of the way: they're respectively expressed by $\top$ and $\bot$ in $\mathcal{L}_1^+$ or $\mathcal{L}_1^{++}$, and by e.g. $P \vee \neg P$ and $P \wedge \neg P$ in $\mathcal{L}_1$.

Suppose then that f is an N-place truth-function, with $N \geqslant 1$. If $f(t_1, \cdots, t_N) = 0$ for all n-tuples of truth-values $\langle t_1, \cdots, t_N \rangle$, we can take ϕ as $P \wedge \neg P$.

In all other cases, we use f's truth table to define a formula ϕ that is a disjunction of formulas χ_j each of which is a conjunction of literals (sentence letters or their negations). To do so, first define α_{j_k} for each $j : \{1, \cdots, N\} \to \{0, 1\}$ (i.e. j is a function from the set $\{1, \cdots, N\}$ to $\{0, 1\}$) and $1 \leqslant k \leqslant N$ as follows:

$$\alpha_{j_k} = \begin{cases} P_k & \text{if } j(k) = 1 \\ \neg P_k & \text{if } j(k) = 0 \end{cases}$$

Intuitively, j is a row of f's truth-table, and α_{j_k} 'corresponds' to the k^{th} element of this row, sentence letters corresponding to 1s and negations of sentence letters to 0s. We define χ_j to be the conjunction of these α_{j_k}, i.e.

$$\chi_j = \bigwedge_{1 \leqslant k \leqslant N} \alpha_{j_k}$$

Next, let $T(f)$ be the set

$\{j : j \text{ is a function from } \{1, \cdots, N\} \text{ to } \{0,1\} \text{ s.t. } f(j(1), \cdots, j(N)) = 1\}$

By our previous assumption, $T(f)$ is non-empty. Putting everything together, we can define the formula expressing f:

$$\phi = \bigvee_{j \in T(f)} \chi_j$$

ϕ is thus a disjunction of conjunctions of sentence letters or their negations (or in the case in which $T(f)$ is of size 1, just a conjunction of literals). Notice that $SenLett(\phi) = \{P_1, \cdots, P_N\}$.

We claim that the $\mathcal{L}_1$-formula ϕ expresses the truth-function f. Suppose $\mathcal{A}$ is an $\mathcal{L}_1$-structure and that $j_{\mathcal{A}}$ is the particular function from $\{1, \cdots, N\}$ for which $j_{\mathcal{A}}(k) = |P_k|_{\mathcal{A}}$ for $1 \leqslant k \leqslant N$. $\mathcal{A}$ determines $j_{\mathcal{A}}$, so we need only focus on the row of the truth-table that corresponds to $j_{\mathcal{A}}$.

Then:

$$\begin{gathered} f(|P_1|_{\mathcal{A}}, \cdots, |P_n|_{\mathcal{A}}) = 1 \\ \Updownarrow \\ j_{\mathcal{A}} \in T(f) \\ \Updownarrow \\ \chi_{j_{\mathcal{A}}} \text{ is a disjunct in } \phi \text{ and } |\chi_{j_{\mathcal{A}}}|_{\mathcal{A}} = 1 \\ \Updownarrow \\ |\phi(P_1, \cdots, P_n)|_{\mathcal{A}} = 1 \end{gathered}$$

First biconditional: by the definition of $j_{\mathcal{A}}$ and $T(f)$.

Second biconditional: by the definition of ϕ in terms of the χ_j and the properties of conjunctions.

Third biconditional, top to bottom: if $\chi_{j_{\mathcal{A}}}$ is a disjunct in ϕ and is true in $\mathcal{A}$ then clearly so is ϕ.

Third biconditional, bottom to top: if $|\phi|_{\mathcal{A}} = 1$ then some disjunct χ_j in ϕ must be true in $\mathcal{A}$; by the properties of conjunctions $|\chi_j|_{\mathcal{A}} = 1$ only if $j = j_{\mathcal{A}}$. ■

The proofs of the next three corollaries and proposition are left as exercises.

Corollary 16 *The sets* $\{\wedge, \neg\}$*,* $\{\vee, \neg\}$ *and* $\{\rightarrow, \neg\}$ *are all expressively adequate.*

Definition 10 *A* literal *is a sentence letter or its negation. Any disjunction of conjunctions of literals is in* disjunctive normal form *and any conjunction of disjunctions of literals is in* conjunctive normal form*.*

Corollary 17 *By Theorem 15, every* $\mathcal{L}_1$*,* $\mathcal{L}_1^+$ *and* $\mathcal{L}_1^{++}$ *formula is equivalent to one in disjunctive normal form, and also equivalent to one in conjunctive normal form.*

Proposition 18 *In* $\mathcal{L}_1^{++}$ *there are exactly 2 expressively adequate binary connectives.*

We now move on to a key theorem, one of the most important theorems in logic. It's almost certainly the most important theorem mentioned in this course, even if it's not apparent why yet.

Definition 11 *Consider a logic* $\mathcal{L}$ *with entailment relation* $\vDash_{\mathcal{L}}$*.* $\mathcal{L}$ *is compact just when: for any set of formulas* Γ *and formula* ϕ*, if* $\Gamma \vDash_{\mathcal{L}} \phi$ *then* $\Gamma^{fin} \vDash_{\mathcal{L}} \phi$ *for some finite* $\Gamma^{fin} \subseteq \Gamma$*.*

When there is just one logic around, there will be no need to subscript the entailment relation; so we usually just write '$\vDash$' instead of '$\vDash_{\mathcal{L}}$'.

It is easy to show—you should check this—that, for any logic $\mathcal{L}$ containing negation, $\mathcal{L}$ is compact just when: if Γ is unsatisfiable then some finite subset Γ^{fin} of Γ is unsatisfiable. Another alternative characterisation of compactness: if every finite subset of Γ is satisfiable then so is Γ itself. This notion crops up sufficiently often that it deserves its own definition.

Definition 12 Γ *is* finitely satisfiable *just when all of its finite subsets are satisfiable.*

It's time now for the promised major theorem. (A more detailed and more advanced treatment of the theorem may be found in a 2022 encyclopaedia article I co-wrote with Rob Leek; see 'The Compactness Theorem'.) The theorem holds for $\mathcal{L}_1$, $\mathcal{L}_1^+$, $\mathcal{L}_1^{++}$, $\mathcal{L}_2$ and $\mathcal{L}_=$, but we'll prove it for $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_=$ only since we haven't given proof systems for $\mathcal{L}_1^+$ and $\mathcal{L}_1^{++}$. In fact, our argument proves a much more general result: any logic with a sound and complete procedure is compact.

Theorem 19 $\mathcal{L}_1$ $\mathcal{L}_2$ *and* $\mathcal{L}_=$ *are compact logics.*

Proof.

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

Here Γ^{fin} is some finite subset of Γ. The argument just given applies to any logic with a sound and complete proof procedure. In particular, it includes $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_=$, which we know from Introduction to Logic all satisfy the noted conditions on the right when $\vdash$ is interpreted as 'provable in ND'. ■

The only step in the proof requiring further comment is that from (2) to (3). Anything that deserves to be called a proof procedure usually satisfies a host of syntactic requirements. The relevant requirement here is that proofs draw only on finitely many premisses. That's why the step from (2) to (3) is valid.

This proof of the Compactness Theorem, magnificent though it is, is something of a cheat since we haven't yet proved the soundness and completeness of any of the logics we're interested in. So this argument is deceptively simple. We'll prove the soundness and completeness of $\mathcal{L}_1$ in the second half of this course, and thereby the compactness of $\mathcal{L}_1$. The proof of soundness and completeness of predicate logics such as $\mathcal{L}_2$ or $\mathcal{L}_=$ is more advanced (though not much more), and will have to await a later course.

As I mentioned in lectures, some logicians take exception to the proof just given, because they think that proofs of semantic facts such as the Compactness Theorem need not, and should not, invoke any syntactic notions. In lecture 8, we'll sketch an entirely semantic proof of the compactness of $\mathcal{L}_1$.

To give you a glimpse of the Compactness Theorem's importance, I'll mention a more advanced and non-examinable result. It shows that the logic $\mathcal{L}_=$ is not powerful enough to define the notion 'there are infinitely many things'. First, a definitional aside:

Definition 13 *An infinite $\mathcal{L}_=$-structure (or $\mathcal{L}_2$-structure) is one with an infinite domain; similarly for a finite $\mathcal{L}_=$-structure (or $\mathcal{L}_2$-structure).*

Proposition 20 (Expressive Limitation of $\mathcal{L}_=$) *There is no $\mathcal{L}_=$-sentence that is true in all and only infinite $\mathcal{L}_=$-structures.*

Sketch Proof. Suppose for reductio that ϕ were a sentence with this property. Then $\neg\phi$ would be true in all and only finite structures. For each

$n \geqslant 1$, let $\exists_{\geqslant n}$ be $\exists x_1 \cdots \exists x_n (\bigwedge_{1 \leqslant i < j \leqslant n} \neg x_i = x_j)$. Clearly, $\exists_{\geqslant n}$ is true in an $\mathcal{L}_{=}$-structure iff the structure has at least n elements in its domain.

Now consider $\Gamma = \{\neg\phi\} \cup \{\exists_{\geqslant n} : n \geqslant 1\}$.

Our first subclaim is that Γ is unsatisfiable, because $\neg\phi$ is satisfiable in all and only finite structures, and the set $\{\exists_{\geqslant n} : n \geqslant 1\}$ is satisfiable in all and only infinite structures.

The second subclaim is that any finite subset of Γ is satisfiable. You should convince yourself of this by thinking about what a finite subset of $\{\exists_{\geqslant n} : n \geqslant 1\}$ looks like.

Putting the two subclaims together shows that Γ contradicts the compactness of $\mathcal{L}_{=}$. Thus there is no such formula ϕ. ■

For much more on compactness, see the encyclopaedia article 'The Compactness Theorem'.

Lecture 5

The first half of the course consisted of semantic metatheory, principally that of $\mathcal{L}_1$. We were concerned with the semantic consequence relation $\vDash$, and with proving general results *about* the logic, of the form say 'If $\Gamma \vDash \phi$ then...', rather than with proving specific results such as say $\{P, P \to Q\} \vDash Q$. In the second half, we'll focus more on deductive metatheory. In this lecture, we'll think about deductive systems in a more abstract way than you've hitherto been used to. In later lectures the focus will be on the specific proof system you studied in Introduction to Logic, especially its propositional fragment.

A logic $\mathcal{L}$ may be characterised in general terms as consisting of a language and a consequence relation $\vDash_{\mathcal{L}}$. $Sen(\mathcal{L})$ is the set of sentences of the logic, and the consequence relation $\vDash_{\mathcal{L}}$ is almost always taken to relate subsets of $Sen(\mathcal{L})$ to elements of $Sen(\mathcal{L})$; that is, if $\Gamma \vDash_{\mathcal{L}} \phi$ then $\Gamma \subseteq Sen(\mathcal{L})$ and $\phi \in Sen(\mathcal{L})$. (Formally speaking, then, the relation $\vDash_{\mathcal{L}}$ is a subset of $\mathbb{P}(Sen(\mathcal{L})) \times Sen(\mathcal{L})$; don't worry if this way of putting things doesn't make sense.) The relation $\vDash_{\mathcal{L}}$ is not usually taken as primitive but rather defined using the notion of an $\mathcal{L}$-structure, as in the Introduction to Logic course,: $\Gamma \vDash_{\mathcal{L}} \phi$ just when, for all $\mathcal{L}$-structures $\mathcal{A}$, if $\mathcal{A}$ satisfies all the elements of Γ then $\mathcal{A}$ satisfies ϕ.

A deductive system D for a logic gives rise to a consequence relation $\vdash_D$, where $\Gamma \vdash_D \phi$ is usually taken to mean: there's a proof in D of ϕ from premises all of which are elements of Γ. It's tricky to say exactly what we require of a deductive system and what it means to be a proof. Part of the job description of a deductive system is that it be syntactic, i.e. concerned only with symbols and not their meanings; but spelling this out is not as straightforward as you might think. Proofs are also required to be decidable: there should be an algorithm (a computer programme if you like) that returns the answer YES when faced with a string of symbols that is a D-proof and NO when it isn't. A string of symbols may not be a D-proof either because it's not a sequence of $\mathcal{L}$-formulas, or because it is but the formulas are not combined together in the right way to make up a D-proof. We'll set to one side all these difficult questions about what a proof system is and just assume that it satisfies three conditions.

Definition 14 (Minimal assumptions on a proof system) *If* $\vdash_D$ *is a deductive consequence relation, we assume that*

i. For all $\Gamma \subseteq Sen(\mathcal{L})$ *and all* $\phi \in Sen(\mathcal{L})$*, if* $\Gamma \vdash_D \phi$ *then there is a finite* $\Gamma^{\text{fin}} \subseteq \Gamma$ *such that* $\Gamma^{\text{fin}} \vdash_D \phi$*.*

ii. For all $\Gamma \subseteq Sen(\mathcal{L})$ *and all* $\phi \in Sen(\mathcal{L})$*,* $\Gamma \vdash_D \phi$ *if* $\phi \in \Gamma$*.*

iii. For all $\Gamma, \Delta \subseteq Sen(\mathcal{L})$ *and all* $\phi \in Sen(\mathcal{L})$*, if* $\Gamma \vdash_D \phi$ *then* $\Gamma \cup \Delta \vdash_D \phi$*.*

We used the first condition at the end of lecture 4 to prove the compactness theorem. The condition states that proofs can only have finitely many premises. The second condition expresses the intuitive thought that one can prove anything one assumes. And the third condition expresses the idea that provability is monotonic: unlike inductive reasoning, in deductive reasoning the set of conclusions you can prove can't shrink by adding more premises.

These three requirements are fairly weak conditions to put on a proof system. I certainly wouldn't want to claim that satisfying them is sufficient for being a proof system. In fact, the conditions may not even be necessary: perhaps something deserves to be called a deductive system even whilst failing one or more of the conditions.

Moving on, let's consider how $\vDash_{\mathcal{L}}$ and $\vdash_D$ may be related.

Definition 15 *D is* strongly sound *with respect to* $\vDash_{\mathcal{L}}$ *just when for all* $\Gamma \subseteq Sen(\mathcal{L})$ *and all* $\phi \in Sen(\mathcal{L})$*, if* $\Gamma \vdash_D \phi$ *then* $\Gamma \vDash_{\mathcal{L}} \phi$*.*

D is strongly complete *with respect to* $\vDash_{\mathcal{L}}$ *just when for all* $\Gamma \subseteq Sen(\mathcal{L})$ *and all* $\phi \in Sen(\mathcal{L})$*, if* $\Gamma \vDash_{\mathcal{L}} \phi$ *then* $\Gamma \vdash_D \phi$*.*

D is weakly sound *with respect to* $\vDash_{\mathcal{L}}$ *just when for all* $\phi \in Sen(\mathcal{L})$*, if* $\vdash_D \phi$ *then* $\vDash_{\mathcal{L}} \phi$*.*

D is weakly complete *with respect to* $\vDash_{\mathcal{L}}$ *just when for all* $\phi \in Sen(\mathcal{L})$*, if* $\vDash_{\mathcal{L}} \phi$ *then* $\vdash_D \phi$*.*

$\vDash_{\mathcal{L}}$ *is* strongly completable *just when there is a deductive system D that is strongly (sound and) complete with respect to* $\vDash_{\mathcal{L}}$*.*

$\vDash_{\mathcal{L}}$ *is* weakly completable *just when there is a deductive system D that is weakly (sound and) complete with respect to* $\vDash_{\mathcal{L}}$*.*

I've put the soundness condition in brackets in the definitions of strong and weak completability because it's usually assumed that any proof system we might be interested in is sound. Note that $\vdash_D \phi$ means that ϕ is provable in D from the empty set, so could be alternatively written as $\varnothing \vdash_D \phi$; similarly $\vDash_{\mathcal{L}} \phi$ means that ϕ is semantically entailed from the empty set and could be written as $\varnothing \vDash_{\mathcal{L}} \phi$. It's immediate, then, that the strong versions of the theses imply the weak versions. Logicians often omit the adjectives 'weak' and 'strong', it being clear from context which they mean.

In Introduction to Logic, you saw that the ND system is sound and complete with respect to $\mathcal{L}_=$-consequence. It will be useful to have labels for the three systems you encountered there.

Definition 16 *Let ND_i be the system of rules in* The Logic Manual *for the $\mathcal{L}_i$-connectives, for* $i = 1, 2, =$*.*

Thus ND_1 consists of all and only the propositional rules, ND_2 extends ND_1 with the rules for $\forall$ and $\exists$, and $ND_=$ in turn extends ND_2 with the rules for $=$. As was mentioned in that course, ND_i is sound and complete with respect to $\mathcal{L}_i$-consequence, for $i = 1, 2$ and $=$. Caveat: it does not in general follow from the fact that logic $\mathcal{L}$ with deductive system D is a sublogic of $\mathcal{L}^*$ with deductive system D^* that if $\Gamma \vdash_{D*} \phi$ for $\Gamma \subseteq Sen(\mathcal{L})$ and $\phi \in Sen(\mathcal{L})$ then $\Gamma \vdash_D \phi$. When no more $\mathcal{L}$-sequents are proved by D^* than by D, we say that the former system is *conservative* over the latter. In more advanced proof theory, the notion of conservativeness will turn out to be of central importance.

Let's now do some elementary 'abstract proof theory'. In what follows, we assume that $Sen(\mathcal{L})$ is non-empty and that it is closed under negation, i.e. if $\phi \in Sen(\mathcal{L})$ then $\neg\phi \in Sen(\mathcal{L})$. We use $\Gamma \nvdash_D \phi$ to mean that it's not the case that $\Gamma \vdash_D \phi$.

Definition 17 Γ *is* consistent$_D$ *(or D-consistent) just when there is a* $\phi \in Sen(\mathcal{L})$ *such that* $\Gamma \nvdash_D \phi$.

Γ *is* negation-consistent$_D$ *just when there is no* $\phi \in Sen(\mathcal{L})$ *such that* $\Gamma \vdash_D \phi$ *and* $\Gamma \vdash_D \neg\phi$.

How are these two notions related? Consistency$_D$ is more general than negation-consistency$_D$, because it does not assume the existence of negation in the language. It's also weaker, even granted that assumption. To spell all this out, let's see first why negation-consistency$_D$ implies consistency$_D$. It's understood throughout that $\Gamma \subseteq Sen(\mathcal{L})$ and $\phi \in Sen(\mathcal{L})$.

Lemma 21 *If* Γ *is negation-consistent*$_D$ *then* Γ *is consistent*$_D$.

Proof Suppose Γ is negation-consistent$_D$, i.e. there is no ϕ such that $\Gamma \vdash_D \phi$ and $\Gamma \vdash_D \neg\phi$. Given any $\phi \in Sen(\mathcal{L})$ (a set we've assumed is non-empty), either $\Gamma \nvdash_D \phi$ or $\Gamma \nvdash_D \neg\phi$. So Γ is consistent$_D$. ■

Negation-consistency$_D$ is in fact equivalent to consistency$_D$ plus the following property.

Definition 18 *Deductive system D* underwrites EFQ *(Ex Falso Quodlibet)* from Γ *just when: if* $\Gamma \vdash_D \phi$ *and* $\Gamma \vdash_D \neg\phi$ *for some* ϕ *then* $\Gamma \vdash_D \psi$ *for all* ψ.

Deductive system D underwrites EFQ *(Ex Falso Quodlibet) just when it underwrites EFQ from all* Γ.

Proposition 22 Γ *is negation-consistent*$_D$ *iff* Γ *is consistent*$_D$ *and D underwrites EFQ from* Γ.

Proof Suppose Γ is negation-consistent$_D$. We saw in the previous lemma that Γ is consistent$_D$. Γ also vacuously satisfies the condition for underwriting EFQ since there is no ϕ such that $\Gamma \vdash_D \phi$ and $\Gamma \vdash_D \neg\phi$ (if the antecedent of a conditional begins 'there exists an F such that ...' then the conditional is true when there is no such F).

Suppose conversely that Γ is consistent$_D$ and that D underwrites EFQ from Γ. If Γ were negation-inconsistent$_D$, there would be a ϕ such that $\Gamma \vdash_D \phi$ and $\Gamma \vdash_D \neg\phi$. Because D underwrites EFQ from Γ, it would follow that $\Gamma \vdash \psi$ for all ψ. Hence Γ would be inconsistent$_D$, a contradiction. ■

I leave it as an exercise for you to prove that ND_1 underwrites EFQ from Γ for all $\Gamma \subseteq Sen(\mathcal{L}_1)$. Let's now have three more properties of a deductive system, which this time I'll phrase in terms of all premise sets Γ.

Definition 19 D underwrites Double Negation Introduction *(DNI) just when, for all* Γ *and* ϕ*, if* $\Gamma \vdash \phi$ *then* $\Gamma \vdash \neg\neg\phi$*.*
D underwrites Double Negation Elimination *(DNE) just when, for all* Γ *and* ϕ*, if* $\Gamma \vdash \neg\neg\phi$ *then* $\Gamma \vdash \phi$*.*
D underwrites Redundancy *(RED) just when, for all* Γ *and* ϕ*, if* $\Gamma \cup \{\phi\} \vdash \neg\phi$ *then* $\Gamma \vdash \neg\phi$*.*

Proposition 23 *Suppose D underwrites EFQ and RED. Then $\Gamma \cup \{\phi\}$ is consistent$_D$ iff $\Gamma \nvdash_D \neg\phi$.*

Proof. We'll prove the contrapositives. Suppose $\Gamma \vdash_D \neg\phi$. Then by the third of the conditions on D (Definition 14), $\Gamma \cup \{\phi\} \vdash_D \neg\phi$. But also, by the second of the conditions on D (Definition 14), $\Gamma \cup \{\phi\} \vdash_D \phi$. Thus $\Gamma \cup \{\phi\}$ is negation-inconsistent$_D$, so it's inconsistent$_D$, by Proposition 22.

For the other direction, suppose $\Gamma \cup \{\phi\}$ is inconsistent$_D$. As it proves everything, it proves $\neg\phi$ in particular: $\Gamma \cup \{\phi\} \vdash_D \neg\phi$. Since D underwrites RED, it follows that $\Gamma \vdash_D \neg\phi$. ■

Corollary 24 *Suppose D underwrites EFQ, RED, DNI and DNE. Then $\Gamma \cup \{\neg\phi\}$ is consistent$_D$ iff $\Gamma \nvdash_D \phi$.*

By the previous proposition, $\Gamma \cup \{\neg\phi\}$ is consistent$_D$ iff $\Gamma \nvdash_D \neg\neg\phi$. By the fact that D underwrites DNI and DNE, $\Gamma \nvdash_D \neg\neg\phi$ iff $\Gamma \nvdash_D \phi$. Hence $\Gamma \cup \{\neg\phi\}$ is consistent$_D$ iff $\Gamma \nvdash_D \phi$. ■

The global conditions DNI, DNE and RED are, as you may have guessed, overkill for these results.

You should check that ND_1 underwrites DNI, DNE and RED, as well as EFQ. Going forward, we'll assume all these facts about ND_1, e.g. we'll assume that $\Gamma \cup \{\phi\}$ is consistent$_{ND_1}$ iff $\Gamma \nvdash_{ND_1} \neg\phi$, for all $\Gamma \subseteq Sen(\mathcal{L}_1)$ and $\phi \in Sen(\mathcal{L}_1)$.

Lecture 6

Let's continue to think about deductive systems abstractly before focusing on a concrete proof sytem, that of ND_1 (the propositional fragment of the proof system in *The Logic Manual*). As usual, we assume that Γ is a set of sentences and ϕ a sentence, and in the following definitions we also assume that our logic $\mathcal{L}$ is closed under negation (if ϕ is in $Sen(\mathcal{L})$ then so is $\neg\phi$).

Definition 20 *Γ is* semantically complete *just when: for all ϕ, $\Gamma \vDash \phi$ or $\Gamma \vDash \neg\phi$ (or both).*

Γ *is* deductively complete *with respect to D (or D-complete) just when: for all ϕ, $\Gamma \vdash_D \phi$ or $\Gamma \vdash_D \neg\phi$ (or both).*

Γ *is* maximally consistent$_D$ *(or maximally D-consistent) just when Γ is consistent$_D$ (see Definition 17) and if $\Gamma \cup \{\phi\}$ is consistent$_D$ then $\phi \in \Gamma$.*

So when Γ is semantically complete it acts like an $\mathcal{L}$-structure by semantically deciding every claim, and when Γ is deductively complete, it does so deductively. Finally to say that Γ is *maximally consistent$_D$* is, informally, to say that it is full to the brim, almost bursting, as far as consistency$_D$ is concerned: add an extra sentence to it and it will no longer be consistent$_D$.

For the rest of this lecture, we take $D = ND_1$. We'll also assume that ND_1 underwrites the conditions EFQ, RED, DNI and DNE, something I asked you to prove earlier. It follows that all the results proved in lecture 5 apply to ND_1.

Lemma 25 *Suppose $\Gamma \subseteq Sen(\mathcal{L}_1)$ is maximally ND_1-consistent. Then Γ is ND_1-consistent and ND_1-complete.*

Proof Assume that Γ is maximally ND_1-consistent. Γ is ND_1-consistent by the definition of maximal ND_1-consistency.

Suppose Γ were ND_1-incomplete, i.e. $\Gamma \nvdash_{ND_1} \phi$ and $\Gamma \nvdash_{ND_1} \neg\phi$ for some ϕ. Using Corollary 24 applied to ND_1, we deduce from $\Gamma \nvdash_{ND_1} \phi$ that $\Gamma \cup \{\neg\phi\}$ is ND_1-consistent. By Γ's maximal ND_1-consistency, it follow that $\neg\phi \in \Gamma$.

Similarly, by Proposition 23 applied to ND_1, we deduce from $\Gamma \nvdash_{ND_1} \neg\phi$ that $\Gamma \cup \{\phi\}$ is ND_1-consistent. By Γ's maximal ND_1-consistency, it follows that $\phi \in \Gamma$.

Since both $\phi, \neg\phi$ are in Γ, Γ proves them both, and so is negation-inconsistent$_{ND_1}$ and hence inconsistent$_{ND_1}$. This contradicts our assumption, thereby proving that Γ is ND_1-complete. ■

We note that the lemma's converse fails. Consider for example $\Gamma = \{\alpha : \alpha$ is a sentence letter$\}$. You should check that Γ is ND_1-consistent and ND_1-complete. Yet Γ is patently not maximally ND_1-consistent; e.g. $\Gamma \vdash_{ND_1} P \wedge Q$, so $\Gamma \cup \{P \wedge Q\}$ is ND_1-consistent yet $P \wedge Q \notin \Gamma$.

The next lemma tell us that for maximally ND_1-consistent sets, membership coincides with derivability.

Lemma 26 *Suppose Γ is maximally ND_1-consistent. Then for any ϕ, $\Gamma \vdash_{ND_1} \phi$ iff $\phi \in \Gamma$.*

Proof If $\phi \in \Gamma$ then clearly $\Gamma \vdash_{ND_1} \phi$.

For the other direction, we invoke Proposition 23, applied to ND_1, which states that $\Gamma \cup \{\phi\}$ is ND_1-consistent iff $\Gamma \nvdash_{ND_1} \neg\phi$. From $\Gamma \vdash_{ND_1} \phi$ and Γ's ND_1-consistency (more precisely: its negation-consistency$_{ND_1}$), it follows that $\Gamma \nvdash_{ND_1} \neg\phi$. So by Proposition 23 applied to ND_1, $\Gamma \cup \{\phi\}$ is ND_1-consistent. So by Γ's maximal ND_1-consistency, $\phi \in \Gamma$. ■

Next, we show that consistency and completeness are sufficient for derivability to behave just like truth in a structure.

Lemma 27 (Consistency + Completeness Lemma) *Suppose Γ is both ND_1-consistent and ND_1-complete. Then for all ϕ and ψ,*

(i) $\Gamma \vdash_{ND_1} \neg\phi$ *iff* $\Gamma \nvdash_{ND_1} \phi$

(ii) $\Gamma \vdash_{ND_1} \phi \wedge \psi$ *iff* $\Gamma \vdash_{ND_1} \phi$ *and* $\Gamma \vdash_{ND_1} \psi$

(iii) $\Gamma \vdash_{ND_1} \phi \vee \psi$ *iff* $\Gamma \vdash_{ND_1} \phi$ *or* $\Gamma \vdash_{ND_1} \psi$ *(or both)*

(iv) $\Gamma \vdash_{ND_1} \phi \rightarrow \psi$ *iff* $\Gamma \nvdash_{ND_1} \phi$ *or* $\Gamma \vdash_{ND_1} \psi$ *(or both)*

(v) $\Gamma \vdash_{ND_1} \phi \leftrightarrow \psi$ *iff* $(\Gamma \vdash_{ND_1} \phi$ *and* $\Gamma \vdash_{ND_1} \psi)$ *or* $(\Gamma \nvdash_{ND_1} \phi$ *and* $\Gamma \nvdash_{ND_1} \psi)$

Proof I'll prove the first two parts and leave the rest to you. We assume thoughout that Γ is ND_1-consistent and ND_1-complete. We'll also assume the equivalence of consistency$_{ND_1}$ and negation-inconsistency$_{ND_1}$, previously proved.

For the left-to-right direction in (i), suppose $\Gamma \vdash_{ND_1} \neg\phi$. Since Γ is ND_1-consistent, it follows that $\Gamma \nvdash \phi$.

For the right-to-left-direction in (i), suppose $\Gamma \nvdash \phi$. Then by Γ's ND_1-completeness, $\Gamma \vdash \neg\phi$.

For the right-to-left direction in (ii), suppose $\Gamma \vdash_{ND_1} \phi$ and $\Gamma \vdash_{ND_1} \psi$. Let Π_1 be an ND_1-proof whose set of undischarged assumptions is a subset of Γ and whose conclusion is ϕ, and similarly let Π_2 be an ND_1-proof whose set of undischarged assumptions is a subset of Γ and whose conclusion is ψ. We let Π_3 be the proof obtained by putting Π_1 and Π_2 side by side and applying the $\wedge$-introduction rule to their respective conclusions ϕ and ψ to obtain $\phi \wedge \psi$. Π_3 is thus an ND_1-proof whose set of undischarged assumptions is a subset of Γ (since the sets of undischarged assumptions of Π_1 and Π_2 are subsets of Γ) and whose conclusion is $\phi \wedge \psi$. Pictorially:

$$\dfrac{\begin{array}{c}\Pi_1\\ \vdots\\ \phi\end{array} \qquad \begin{array}{c}\Pi_2\\ \vdots\\ \psi\end{array}}{\phi \wedge \psi} \wedge intro$$

For the left-to-right direction in (ii), suppose $\Gamma \vdash_{ND_1} \phi \wedge \psi$. A similar argument to the one just given shows that a proof of $\phi \wedge \psi$ from premises that are all elements of Γ can be extended by a single application of the first $\wedge$-elimination rule to a proof of ϕ from these same premises, and that this same proof can be extended by a single application of the other $\wedge$-elimination rule to a proof of ψ from these same premises. Hence $\Gamma \vdash_{ND_1} \phi$ and $\Gamma \vdash_{ND_1} \psi$.

Cases (iii), (iv) and (v) are entirely analogous. ■

Corollary 28 *Suppose Γ is maximally ND_1-consistent. Then for all ϕ and ψ,*

(i) $\neg\phi \in \Gamma$ *iff* $\phi \notin \Gamma$

(ii) $\phi \wedge \psi \in \Gamma$ *iff* $\phi \in \Gamma$ *and* $\psi \in \Gamma$

(iii) $\phi \vee \psi \in \Gamma$ *iff* $\phi \in \Gamma$ *or* $\psi \in \Gamma$ *(or both)*

(iv) $\phi \rightarrow \psi \in \Gamma$ *iff* $\phi \notin \Gamma$ *or* $\psi \in \Gamma$ *(or both)*

(v) $\phi \leftrightarrow \psi \in \Gamma$ *iff* $(\phi \in \Gamma$ *and* $\psi \in \Gamma)$ *or* $(\phi \notin \Gamma$ *and* $\psi \notin \Gamma)$ *(or both)*

Proof A consequence of the previous two lemmas. We'll do cases (i) and (ii). We assume that Γ is maximally ND_1-consistent throughout.

As for (i): by Lemma 26, $\neg\phi \in \Gamma$ iff $\Gamma \vdash_{ND_1} \neg\phi$. By Lemma 27, $\Gamma \vdash_{ND_1} \neg\phi$ iff $\Gamma \nvdash_{ND_1} \phi$. And by Lemma 26 again, $\Gamma \nvdash_{ND_1} \phi$ iff $\phi \notin \Gamma$. So $\neg\phi \in \Gamma$ iff $\phi \notin \Gamma$.

As for (ii): by Lemma 26, $\phi \wedge \psi \in \Gamma$ iff $\Gamma \vdash_{ND_1} \phi \wedge \psi$. By Lemma 27, $\Gamma \vdash_{ND_1} \phi \wedge \psi$ iff $\Gamma \vdash \psi$ and $\Gamma \vdash \phi$. And by Lemma 26 again, $\Gamma \vdash \phi$ and $\Gamma \vdash \psi$ iff $\phi \in \Gamma$ and $\psi \in \Gamma$. So $\phi \wedge \psi \in \Gamma$ iff $\phi \in \Gamma$ and $\psi \in \Gamma$.

Cases (iii), (iv) and (v) are entirely analogous. ■

Lecture 7

Last time, we familiarised ourselves with maximally consistent sets and their properties, as well as the properties of complete and consistent sets, of which maximally consistent sets are an example. Today, we'll prove the soundness of ND_1 with respect to $\mathcal{L}_1$-consequence, and make a start on proving its completeness. Without further ado, then, let's prove soundness.

Theorem 29 (Soundness of ND_1) *For all $\Gamma \subseteq Sen(\mathcal{L}_1)$, $\phi \in Sen(\mathcal{L}_1)$, if $\Gamma \vdash_{ND_1} \phi$ then $\Gamma \vDash \phi$.*

Proof We assign a natural number to each ND_1-proof given by the number of rule applications in the proof. Let's call this number the *size* of the proof. ('Length' is more usual when discussing proofs; but natural deduction proofs are trees rather than linear sequences.) For the avoidance of misunderstanding, rule applications are tokens not types, so that e.g. the proof

$$\dfrac{\dfrac{P}{P \vee Q}\ \vee \text{ intro } 1}{(P \vee Q) \vee R}\ \vee \text{ intro } 1$$

has size 2 since it applies the (first) $\vee$-introduction rule twice. We then prove the result by induction on the size of proofs, taking as our inductive hypothesis that if Π is an ND_1-proof of ϕ from Γ of size N then $\Gamma \vDash \phi$.

For the induction basis, suppose Π is a proof of size 0, i.e. containing no rule applications. Then Π must take the form: ϕ. Since this one-line proof is a proof from Γ of ϕ, ϕ must be an element of Γ. Clearly, then, $\Gamma \vDash \phi$ as $\phi \in \Gamma$.

For the inductive step, assume the inductive hypothesis for proofs of size $\leqslant N$. So let Π be a proof of size $N + 1$ of ϕ from a set of undischarged assumptions all of which are elements of Γ. Let's call the last rule used in this proof ρ (notice that every proof has a last rule). The argument now proceeds by considering all the possibilities for ρ. This involves a large number of cases, so I'll do one here and leave the rest to you.

Suppose ρ is the $\wedge$-introduction rule. So Π takes the following form:

$$\dfrac{\begin{matrix}\Pi_1 \\ \vdots \\ \phi_1\end{matrix} \qquad \begin{matrix}\Pi_2 \\ \vdots \\ \phi_2\end{matrix}}{\phi_1 \wedge \phi_2}\ \rho = \wedge intro$$

where $\phi = \phi_1 \wedge \phi_2$. Let's call $Prem(\Pi_i)$ the set of undischarged assumptions in subproof Π_i of ϕ_i, for $i = 1, 2$. Since $Prem(\Pi_1)$ and $Prem(\Pi_2)$ are both subsets of the set of undischarged assumptions in Π, which itself is a subset of Γ, it follows that $Prem(\Pi_1), Prem(\Pi_2) \subseteq \Gamma$. And since Π_1 and Π_2 are of size $\leqslant N$, it follows from the inductive hypothesis that

$$Prem(\Pi_1) \vDash \phi_1 \text{ and } Prem(\Pi_2) \vDash \phi_2$$

By the semantic rule for conjunction, we deduce that $Prem(\Pi_1) \cup Prem(\Pi_2) \vDash \phi_1 \wedge \phi_2$. And since $Prem(\Pi_1) \cup Prem(\Pi_2) \subseteq \Gamma$, we futher deduce

$$\Gamma \vDash \phi_1 \wedge \phi_2$$

This proves the inductive step for the case $\rho = \wedge$-introduction. I leave the cases in which ρ is some other rule as exercises for you, which you should endeavour to do. They are very similar to the above. (A detail that's unlikely to trip you up yet that's nevertheless worth mentioning: for rules that discharge assumptions you must be careful not to assume that Π's set of undischarged assumptions is the union of the set of undischarged assumptions of its subproofs.) ■

Having dealt with soundness, we now move on to completeness. To prove completeness, we'll need two important auxiliary lemmas. The first lemma states that any consistent set can be extended to a maximally consistent set; the second says that given any maximally consistent set we can define an $\mathcal{L}_1$-structure by equating truth in that structure with membership of the maximally consistent set.

Lemma 30 (First Auxiliary Lemma/Lindenbaum's Lemma) *If Γ is ND_1-consistent then there's a maximally ND_1-consistent set Γ^+ such that $\Gamma \subseteq \Gamma^+$.*

Proof Assume Γ is ND_1-consistent. Though I won't do it here, it's not hard to show that $Sen(\mathcal{L}_1)$ is a countably infinite set, which we may enumerate (without repetition) as $\phi_1, \cdots, \phi_n, \cdots$. Informally, the idea behind the proof is that we run though these sentences, adding a sentence to Γ if we can do so whilst preserving consistency; what we end up with must then be not just consistent, but maximally so.

More formally, we define Γ_n, for n a natural number, recursively. We set $\Gamma_0 = \Gamma$ and

$$\Gamma_{n+1} = \begin{cases} \Gamma_n \cup \{\phi_n\} & \text{if } \Gamma_n \cup \{\phi_n\} \text{ is } ND_1\text{-consistent} \\ \Gamma_n & \text{otherwise} \end{cases}$$

It's immediate from this definition that Γ_n is ND_1-consistent for all n. We can prove this by induction: $\Gamma_0 = \Gamma$ is ND_1-consistent by assumption, and Γ_{n+1} is clearly ND_1-consistent if Γ_n is.

Let's now define Γ^+ as the set of sentences that appear in any of the $\Gamma_n : \Gamma^+ = \bigcup_{0 \leqslant n} \Gamma_n$. Clearly, $\Gamma_n \subseteq \Gamma^+$ for all n, including the case $n = 0$, i.e. $\Gamma = \Gamma_0 \subseteq \Gamma^+$. We first prove that Γ^+ is ND_1-consistent before proving that it's maximally ND_1-consistent.

Suppose for reductio, then, that Γ^+ were ND_1-inconsistent, so that $\Gamma^+ \vdash_{ND_1} \phi$ and $\Gamma^+ \vdash_{ND_1} \neg\phi$. (As ever, we're taking negation-consistency$_{ND_1}$ and consistency$_{ND_1}$ as equivalent.) Thus, since proofs are finite:

$$\{\gamma_1, \cdots, \gamma_m\} \vdash_{ND_1} \phi \text{ for some } \{\gamma_1, \cdots, \gamma_m\} \subseteq \Gamma^+$$
$$\{\gamma_1^*, \cdots, \gamma_n^*\} \vdash_{ND_1} \neg\phi \text{ for some } \{\gamma_1^*, \cdots, \gamma_n^*\} \subseteq \Gamma^+$$

where m and n are natural numbers. The finitely many ($\leqslant m + n$) elements of $\{\gamma_1, \cdots, \gamma_m\} \cup \{\gamma_1^*, \cdots, \gamma_n^*\}$ have all appeared at some finite stage of the enumeration of $Sen(\mathcal{L}_1)$, so define

$$k = \text{the maximum } i \text{ such that } \phi_i \in \{\gamma_1, \cdots, \gamma_m\} \cup \{\gamma_1^*, \cdots, \gamma_n^*\}$$

It follows from this definition that $\{\gamma_1, \cdots, \gamma_m\} \cup \{\gamma_1^*, \cdots, \gamma_n^*\} \subseteq \Gamma_{k+1}$, so that $\Gamma_{k+1} \vdash_{ND_1} \phi$ and $\Gamma_{k+1} \vdash_{ND_1} \neg\phi$. This, however, contradicts Γ_{k+1}'s ND_1-consistency. Hence we may conclude that Γ^+ is ND_1-consistent.

It now remains to show that Γ^+ is maximally ND_1-consistent. So suppose $\Gamma^+ \cup \{\phi\}$ is ND_1-consistent, where $\phi = \phi_k$ in our enumeration of $Sen(\mathcal{L}_1)$. Since $\Gamma^+ \cup \{\phi_k\}$ is ND_1-consistent and $\Gamma_k \subseteq \Gamma^+$, it follows that $\Gamma_k \cup \{\phi_k\}$ is ND_1-consistent. (The third property in Definition 14 implies that a subset of a consistent set is also consistent.) Thus by the definition of Γ_{k+1}, $\phi_k \in \Gamma_{k+1}$ so that $\phi_k \in \Gamma^+$ since $\Gamma_{k+1} \subseteq \Gamma^+$. We conclude that Γ^+ is maximally ND_1-consistent.

To recap: we've shown how, given an ND_1-consistent set Γ, we may define a sequence of consistent extensions of Γ, $\Gamma = \Gamma_0 \subseteq \Gamma_1 \cdots \subseteq \Gamma_n \subseteq \cdots$. Letting Γ^+ be the union of these Γ_n, so that $\Gamma \subseteq \Gamma^+$, we then checked that Γ^+ is not just ND_1-consistent but maximally ND_1-consistent. This proves the lemma. ■.

The First Auxiliary Lemma (Lemma 30) is the first staging post on the way to proving ND_1-completeness. We now state and prove the second.

Lemma 31 (Second Auxiliary Lemma) *If Γ is maximally ND_1-consistent then there is an $\mathcal{L}_1$-structure $\mathcal{A}_\Gamma$ such that, for all $\phi \in Sen(\mathcal{L}_1)$, $\phi \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \phi$.*

Proof Assume Γ is maximally ND_1-consistent. We define $\mathcal{A}_\Gamma$ so that for *atomic* formulas α (i.e. sentence letters),

$$\mathcal{A}_\Gamma \vDash \alpha \text{ iff } \alpha \in \Gamma$$

We must now prove that this applies to *all* formulas ϕ, i.e. that $\phi \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \phi$ whatever ϕ may be. The proof is by induction on the complexity of ϕ.

Base case: ϕ is of complexity 0, i.e. ϕ is a sentence letter. There is nothing to prove, since $\phi \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \phi$ holds by the definition of $\mathcal{A}_\Gamma$.

For the induction step, suppose ϕ is of complexity $N + 1$. There are, as you would expect, five cases to consider.

The first case is $\phi = \neg\psi$, where ψ has complexity N. Since Γ is maximally ND_1-consistent, then by Corollary 28, $\neg\psi \in \Gamma$ iff $\psi \notin \Gamma$. By the induction hypothesis, $\psi \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \psi$, or equivalently, $\psi \notin \Gamma$ iff $\mathcal{A}_\Gamma \nvDash \psi$. And clearly, by the semantic rule for $\neg$, $\mathcal{A}_\Gamma \nvDash \psi$ iff $\mathcal{A}_\Gamma \vDash \neg\psi$. From these three biconditionals, we deduce that $\neg\psi \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \neg\psi$.

The second case is $\phi = \phi_1 \wedge \phi_2$, where ϕ_1 and ϕ_2 each has complexity $\leqslant N$. Since Γ is maximally ND_1-consistent, then by Corollary 28, $\phi_1 \wedge \phi_2 \in \Gamma$ iff $\phi_1 \in \Gamma$ and $\phi_2 \in \Gamma$. By the induction hypothesis, $\phi_i \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \phi_i$, for $i = 1, 2$. And clearly, by the semantic rule for $\wedge$, $\mathcal{A}_\Gamma \vDash \phi_1 \wedge \phi_2$ iff $\mathcal{A}_\Gamma \vDash \phi_1$ and $\mathcal{A}_\Gamma \vDash \phi_2$. From these three biconditionals, we deduce that $\phi_1 \wedge \phi_2 \in \Gamma$ iff $\mathcal{A}_\Gamma \vDash \phi_1 \wedge \phi_2$.

The third, fourth and fifth cases, dealing with $\vee$, $\rightarrow$ and $\leftrightarrow$ are entirely analogous and are left as exercises. ■

We now have all the ingredients for proving ND_1-completeness. But we've run out of time, so we'll prove it in the next and final lecture.

Lecture 8

Last time, we proved the soundness of ND_1 with respect to $\mathcal{L}_1$'s consequence relation. We also proved that any ND_1-consistent set can be extended to a maximally consistent set (First Auxiliary Lemma, i.e. Lemma 30) and that a maximally consistent set corresponds to a unique $\mathcal{L}_1$-structure (Second Auxiliary Lemma, i.e. Lemma 34). It now remains to put these pieces together to prove ND_1's completeness with respect to $\mathcal{L}_1$-consequence.

Theorem 32 (Completeness of ND_1) *For all $\Gamma \subseteq Sen(\mathcal{L}_1)$ and $\phi \in Sen(\mathcal{L}_1)$, if $\Gamma \vDash_{\mathcal{L}_1} \phi$ then $\Gamma \vdash_{ND_1} \phi$.*

Proof We prove the contrapositive: if $\Gamma \nvdash_{ND_1} \phi$ then $\Gamma \nvDash_{\mathcal{L}_1} \phi$. So suppose $\Gamma \nvdash_{ND_1} \phi$. By Corollary 24 applied to ND_1, *viz.* $\Gamma \cup \{\neg\phi\}$ is ND_1-consistent iff $\Gamma \nvdash_{ND_1} \phi$, it follows that $\Gamma \cup \{\neg\phi\}$ is ND_1-consistent. So by the First Auxiliary Lemma, there is a maximal ND_1-consistent set Γ^+ such that $\Gamma \cup \{\neg\phi\} \subseteq \Gamma^+$. And by the Second Auxiliary Lemma, there is an $\mathcal{L}_1$-structure $\mathcal{A}$ such that for all $\delta \in Sen(\mathcal{L}_1)$, $\mathcal{A} \vDash \delta$ iff $\delta \in \Gamma^+$. In particular, since $\Gamma \cup \{\neg\phi\} \subseteq \Gamma^+$, $\mathcal{A} \vDash \Gamma$ and $\mathcal{A} \vDash \neg\phi$. Hence $\Gamma \nvDash_{\mathcal{L}_1} \phi$, which was to be proved. ∎

The proof is deceptively simple, because most of the work has been packed into the First Auxiliary Lemma and the Second Auxiliary Lemma. Moreover, the Second Auxiliary Lemma itself depended on a host of lemmas about the properties of maximally ND_1-consistent sets.

The proof we gave at the end of Lecture 4 of the Compactness Theorem made use of Soundness and Completeness. Since we've now proved these for ND_1, we've discharged our obligations as far as $\mathcal{L}_1$ is concerned: we've proved that $\vDash_{\mathcal{L}_1}$ is compact. But as I mentioned in lectures, it would be better if possible to give a purely semantic proof of that semantic result than have to give a deductive argument. In the last part of today's lecture, we'll sketch a strictly semantic proof of $\mathcal{L}_1$'s compactness. The proof will effectively be a semantic version of the argument for ND_1's completeness. So it will also serve the purpose of highlighting the proof's more abstract features.

Recall Definition 12, which stated that a set of sentences is finitely satisfiable just when all its finite subsets are satisfiable. $\mathcal{L}_1$'s compactness is equivalent to the statement that if a set of sentences Γ is finitely satisfiable then it's satisfiable – this was left as an exercise back in lecture 4, which I urge you to do. As we're going to try to mimic the proof of ND_1's completeness as much as possible in proving the compactness of $\mathcal{L}_1$, we'll need the following definition.

Definition 21 Γ *is* maximally finitely satisfiable *just when* Γ *is finitely satisfiable and if* $\Gamma \cup \{\phi\}$ *is finitely satisfiable then* $\phi \in \Gamma$*.*

Mimicking ND_1's completeness proof, we lay down two auxiliary propositions.

Lemma 33 (First Auxiliary Lemma*) *If* Γ *is finitely satisfiable then there's a maximal finitely satisfiable set* Γ^+ *such that* $\Gamma \subseteq \Gamma^+$*.*

Lemma 34 (Second Auxiliary Lemma*) *If* Γ *is maximally finitely satisfiable then there is an* $\mathcal{L}_1$*-structure* $\mathcal{A}_\Gamma$ *such that, for all* $\phi \in Sen(\mathcal{L}_1)$*,* $\phi \in \Gamma$ *iff* $\mathcal{A}_\Gamma \models \phi$*.*

I'll give abbreviated proofs of each of these lemmas, since their proofs are so similar to that of their deductive counterparts.

Sketch proof of the First Auxiliary Lemma* We assume Γ is finitely satisfiable and enumerate the sentences of $Sen(\mathcal{L}_1)$ as $\phi_1, \cdots, \phi_n, \cdots$. Define $\Gamma_0 = \Gamma$ and

$$\Gamma_{n+1} = \begin{cases} \Gamma_n \cup \{\phi_n\} & \text{if } \Gamma_n \cup \{\phi_n\} \text{ is finitely satisfiable} \\ \Gamma_n & \text{otherwise} \end{cases}$$

It's immediate from the definition that Γ_n is finitely satisfiable for all n. Γ^+ is then defined as before by $\Gamma^+ = \bigcup\limits_{0 \leqslant n} \Gamma_n$. Clearly, $\Gamma_n \subseteq \Gamma^+$ for all n, including the case $n = 0$, i.e. $\Gamma = \Gamma_0 \subseteq \Gamma^+$.

If Γ^+ were not finitely satisfiable then one of its finite subsets would be unsatisfiable, and this finite subset must be entirely drawn from some Γ_n, which would contradict Γ_n's finite satisfiability. Hence Γ^+ is finitely satisfiable. And if ϕ_k in our enumeration is such that $\Gamma \cup \{\phi_k\}$ is finitely satisfiable then $\Gamma_{k+1} = \Gamma \cup \{\phi_k\}$ must be finitely satisfiable, so that $\phi_k \in \Gamma_{k+1} \subseteq \Gamma^+$. In other words, Γ^+ is maximally finitely satisfiable. ■

Sketch Proof of the Second Auxiliary Lemma* Assume Γ is maximally finitely satisfiable. Diverging a little from the proof of the Second Auxiliary Lemma, we first prove that Γ contains exactly one of $\phi, \neg\phi$, for every $\phi \in Sen(\mathcal{L}_1)$. Clearly, Γ cannot contain both ϕ and $\neg\phi$ since $\{\phi, \neg\phi\}$ is finite and unsatisfiable. And if it contains neither, then some finite subset F_1 of Γ is such that that $F_1 \cup \{\phi\}$ is unsatisfiable, and some finite subset F_2 of Γ is such that that $F_2 \cup \{\neg\phi\}$ is unsatisfiable. But then $F_1 \cup F_2$ is a finite and unsatisfiable subset of Γ (since $F_1 \models \neg\phi$ and $F_2 \models \phi$), contradicting Γ's finite satisfiability.

Given the fact that $\neg\phi \in \Gamma$ iff $\phi \notin \Gamma$, it's now easy to prove that (i) $\phi_1 \wedge \phi_2 \in \Gamma$ iff $\phi_1 \in \Gamma$ and $\phi_2 \in \Gamma$; (ii) $\phi_1 \vee \phi_2 \in \Gamma$ iff $\phi_1 \in \Gamma$ or $\phi_2 \in \Gamma$ (or both); (iii) $\phi_1 \to \phi_2 \in \Gamma$ iff $\phi_1 \notin \Gamma$ or $\phi_2 \in \Gamma$ (or both); and (iv) $\phi_1 \leftrightarrow \phi_2 \in \Gamma$ iff ($\phi_1 \in \Gamma$ and $\phi_2 \in \Gamma$) or ($\phi_1 \notin \Gamma$ and $\phi_2 \notin \Gamma$). For example, if $\phi_1 \wedge \phi_2 \in \Gamma$ and $\phi_1 \notin \Gamma$ then $\{\phi_1 \wedge \phi_2, \neg\phi_1\}$ would be a finite but unsatisfiable subset of Γ. We conclude that membership in Γ behaves just like truth in an $\mathcal{L}_1$-structure.

Using the above, we define $\mathcal{A}_\Gamma$ as in the proof of the Second Auxiliary Lemma, so that for *atomic* formulas α (i.e. sentence letters),

$$\mathcal{A}_\Gamma \models \alpha \text{ iff } \alpha \in \Gamma$$

It's now easy to prove that $\phi \in \Gamma$ iff $\mathcal{A}_\Gamma \models \phi$ for all *all* formulas ϕ (not just atomic ones). The proof is once more by induction on the complexity of ϕ, using the facts established in the previous paragraph. ■

Combining the First Auxiliary Lemma* and the Second Auxiliary Lemma* yields an alternative proof of the compactness of $\mathcal{L}_1$ (an instance of Theorem 19).

Alternative proof of Theorem 19 for $\mathcal{L}_1$. As noted, we prove an equivalent of the compactness theorem of $\mathcal{L}_1$: if Γ is finitely satisfiable then Γ is satisfiable. So suppose Γ is finitely satisfiable. By the First Auxiliary Lemma*, there's a maximal finitely satisfiable Γ^+ such that $\Gamma \subseteq \Gamma^+$. By the Second Auxiliary Lemma*, there's an $\mathcal{L}_1$-structure $\mathcal{A}$ such that $\mathcal{A} \models \phi$ iff $\phi \in \Gamma^+$ for all $\phi \in Sen(\mathcal{L}_1)$. In particular, $\mathcal{A} \models \Gamma$, since $\Gamma \subseteq \Gamma^+$. Hence Γ is satisfiable. ■

You may have been wondering about how our proofs of ND_1-completeness and $\mathcal{L}_1$'s compactness would go if the set of sentence letters were not countably infinite. Clearly, if there were only finitely many sentence letters the proofs would, if anything, be easier; an alternative proof of compactness could be given from the observation that there's a finite subset of sentences such that every sentence is logically equivalent to one of its members. So the question is what happens in the case in which the set of sentence letters is uncountable. In fact, one can give a more abstract version and (once one has got used to its initially dizzying level of abstraction) easier version of the arguments for the First Auxiliary Lemma and the First Auxiliary Lemma*. Here we'll give the argument for the latter lemma, easily amended to yield the argument for the former lemma. (The next paragraph is non-examinable, and uses some terminology about orders that I won't define but invite you to look up.)

Suppose Γ is finitely satisfiable. Order by inclusion the set F_Γ of finitely-satisfiable sets of sentences of the language containing Γ. F_Γ is non-empty, since it contains at least Γ. Any chain in F_Γ has an upper bound, obtained by taking the union of the elements in the chain: this union contains Γ as a subset since all the members of the chain do, and it is finitely satisfiable since any of its finite subsets must come from some element of the chain, which by hypothesis is finitely satisfiable. *Zorn's Lemma* states precisely that every partial order with the property that every chain has an upper bound has a maximal element. Since the conditions of Zorn's Lemma are satisfied, we deduce that F_Γ has a maximal element; that is, F_Γ is a maximal finitely satisfiable set extending Γ. Note that nowhere did we rely on the fact that the sentence letters are denumerably many, or on any assumption about the set of connectives. So this more general argument establishes the analogue of the First Auxiliary Lemma* for a propositional logic with atom set of any size.[3]

Finally, I leave you with a teaser. As we've seen, the three logics studied in Intro Logic are compact. A natural question is whether English itself is compact. To tackle this question, one must first clarify what compactness means for a non-formal language such as English. A reasonable definition might be that if an English argument $\Gamma \therefore \delta$ is valid, where Γ is a set of English sentences and δ is an English sentence, then $\Gamma^{fin} \therefore \delta$ is valid, where Γ^{fin} is a finite subset of Γ. What validity in English comes to is of course a vexed issue. But put sceptical doubts aside for a minute and assume that this notion is in good order. Consider now the following argument:

There is at least one planet.

There are at least two planets.

$\vdots$

There are at least n planets.

$\vdots$

There are infinitely many planets.

This argument *appears* to be valid; for if there are at least n planets for every finite n then there must be infinitely many planets. It also seems that

[3]The abstract argument just given invoked Zorn's Lemma, well known to be equivalent to the Axiom of Choice in standard set theory. In fact, a slightly weaker principle than Zorn's Lemma, the *Ultrafilter Lemma*, suffices.

no finite subset of the premiss set entails the conclusion; that there are at least $n_1, \cdots n_k$ planets for some finite $n_1, \cdots n_k$ does not entail that there are infinitely many planets. (This argument is examined in more detail in an article Owen Griffiths and I published in 2021 called 'Is English Consequence Compact?'; see also chapter 5 of our 2022 book *One True Logic*.)

If this is right—and I certainly haven't proved it—then English consequence is incompact, and none of $\mathcal{L}_1$, $\mathcal{L}_2$ and $\mathcal{L}_=$ is capable of capturing it, since these three logics are all compact. Indeed, as we saw in lecture 4, the strongest of the three, $\mathcal{L}_=$, isn't even capable of formulating the argument's conclusion. So which logic captures English validity?