

An Intuitionistic Logic Primer*

Robert Trueman

1 Natural Deduction for Intuitionistic Logic

So far, we have looked at two new logics: Modal Logic and Second-Order Logic. They were both extensions of Classical Logic; they took Classical Logic, and added some extra resources to it. Now we are going to look at *Intuitionistic Logic* (IL), which we get by restricting Classical Logic, not extending it.

The language of IL is exactly the same as the language of FOL.¹ The difference between IL and FOL shows up in their natural deduction systems. The system for IL incorporates all of the basic rules for FOL, *apart from* Tertium Non Datur (TND):

i			\mathcal{A}	
			\mathcal{B}	
j			\mathcal{B}	
			$\neg\mathcal{A}$	
k			$\neg\mathcal{A}$	
			\mathcal{B}	
l			\mathcal{B}	
			\mathcal{B}	
				TND, $i-j, k-l$

And that's it, that's the basic difference between IL and FOL! Of course, this basic difference has all sorts of interesting ramifications. First off, if we reject the basic rule of TND, then we also have to reject all sorts of *derived* rules for FOL. Most obviously, we have to reject Double Negation Elimination (DNE):

m		$\neg\neg\mathcal{A}$	
		\mathcal{A}	DNE, m

*Thanks to Tim Button for kindly letting me steal from some of his old lecture notes while preparing this primer.

¹We could have started with SOL, rather than FOL, if we liked. However, I want to keep things simple, and so I will keep them first-order.

And it doesn't end there. Although we can keep three of the De Morgan Rules, we have to reject this one:

$$\boxed{
 \begin{array}{c|l}
 m & \neg(\mathcal{A} \wedge \mathcal{B}) \\
 & \neg\mathcal{A} \vee \neg\mathcal{B} \quad \text{DeM, } m
 \end{array}
 }$$

And while we're at it, we also have to reject the following rule for Converting Quantifiers:

$$\boxed{
 \begin{array}{c|l}
 m & \neg\forall x\mathcal{A} \\
 & \exists x\neg\mathcal{A} \quad \text{CQ, } m
 \end{array}
 }$$

But aside from that, all of the other rules for FOL listed in *forall* χ , basic and derived, carry over to IL. As before, we will use the symbol \vdash to express *provability*, but we will also add subscripts to indicate whether we are dealing with IL or classical FOL:

$$\boxed{
 \begin{array}{l}
 \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash_{\text{IL}} \mathcal{C} \text{ iff } \mathcal{C} \text{ can be proved from } \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n, \text{ using only} \\
 \text{the rules of IL.} \\
 \\
 \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n \vdash_{\text{C}} \mathcal{C} \text{ iff } \mathcal{C} \text{ can be proved from } \mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_n, \text{ using any} \\
 \text{of the rules of classical FOL.}
 \end{array}
 }$$

2 Rejecting the Law of the Excluded Middle

It is often said that intuitionists reject the Law of the Excluded Middle (LEM): $\mathcal{A} \vee \neg\mathcal{A}$. This is quite right, but we have to be very clear about what this means.

The first thing to note is that LEM is a *schematic* law of Classical Logic. That means that every *instance* of LEM is a theorem of Classical Logic. To build an instance of LEM, we just have to substitute the same sentence for both of the \mathcal{A} s in the law. Here are some examples:

$$P \vee \neg P, (P \vee Q) \vee \neg(P \vee Q), \exists y\forall x(Fy \leftrightarrow x = y) \vee \neg\exists y\forall x(Fy \leftrightarrow x = y)$$

At this point it should be clear that you can reject LEM *without* accepting the negation of LEM as a new law (\neg LEM): $\neg(\mathcal{A} \vee \neg\mathcal{A})$. Plainly, you could deny that every instance of LEM was a logical theorem without accepting that every instance

of \neg LEM is a logical theorem! But what might come as more of a surprise is that intuitionists do not accept that *any* instance of \neg LEM is a logical theorem. In fact, you can prove that \neg LEM is a full blown *contradiction* in IL.

So, the intuitionists reject LEM, even though they know it would be a contradiction to assert any instance of \neg LEM. What is going on? Well, when an intuitionist rejects LEM, all they are doing is denying that all of its instances are *logical theorems*. In other words, intuitionists deny that it is always possible to prove an instance of LEM without the help of any premises. And that is quite right, in Intuitionistic Logic: $\not\vdash_1 \mathcal{A} \vee \neg\mathcal{A}$.

The crucial point, then, is that every instance of LEM is a *theorem* of classical FOL, but not every instance is a *theorem* of IL. And in fact, these are not the only theorems of classical FOL which cannot be proven in IL. One particularly noteworthy example is Peirce's Law: $((\mathcal{A} \rightarrow \mathcal{B}) \rightarrow \mathcal{A}) \rightarrow \mathcal{A}$. You can prove this in classical FOL (see §4.2, exercise 1), but not in IL.

Right, that's enough by way of introduction to IL. If you would like to get the hang of what it's like working without TND, here are some exercises for you to try (solutions in §7):

Provide proofs for the following:

1. $\neg(A \vee \neg A) \vdash_1 \perp$
2. $A \vee B \vdash_1 \neg\neg A \vee \neg\neg B$
3. $\neg\neg\neg A \vdash_1 \neg A$
4. $\neg\forall x\neg\neg Px \vdash_1 \neg\forall x\neg\neg\neg Px$
5. $\vdash_1 ((\neg\neg A \rightarrow \neg\neg B) \rightarrow \neg\neg A) \rightarrow \neg\neg A$

3 Translating Classical Theorems

I just emphasised that there are theorems of classical FOL which are not theorems of IL. But having said all that, I would also like to take a moment to sketch a neat result, proved independently by Gödel, Gentzen and Kolmogorov. The result is that there is a way of *translating* classical theorems into intuitionistic theorems (and back again). Here is Kolmogorov's translation manual:

$$\begin{aligned}
\mathcal{A}^K &= \neg\neg\mathcal{A}, \text{ if } \mathcal{A} \text{ is atomic} \\
(\neg\mathcal{A})^K &= \neg(\mathcal{A}^K) \\
(\mathcal{A} \wedge \mathcal{B})^K &= \neg\neg(\mathcal{A}^K \wedge \mathcal{B}^K) \\
(\mathcal{A} \vee \mathcal{B})^K &= \neg\neg(\mathcal{A}^K \vee \mathcal{B}^K) \\
(\mathcal{A} \rightarrow \mathcal{B})^K &= \neg\neg(\mathcal{A}^K \rightarrow \mathcal{B}^K) \\
(\mathcal{A} \leftrightarrow \mathcal{B})^K &= \neg\neg((\mathcal{A} \rightarrow \mathcal{B})^K \wedge (\mathcal{B} \rightarrow \mathcal{A})^K) \\
(\forall\chi\mathcal{A})^K &= \neg\neg\forall\chi(\mathcal{A}^K) \\
(\exists\chi\mathcal{A})^K &= \neg\neg\exists\chi(\mathcal{A}^K)
\end{aligned}$$

Now, it is important not to get carried away by this talk of ‘translation’. Ordinarily, we expect translation to preserve meaning: if I ask you to translate an English sentence into French, I expect you to write out a French sentence which *means the same thing* as the English sentence. But Kolmogorov’s translation manual doesn’t preserve meaning in this way: in general, \mathcal{A}^K will mean something quite different from \mathcal{A} , at least from an intuitionistic point of view.

So, if \mathcal{A}^K and \mathcal{A} aren’t guaranteed to mean the same thing, then what is the relation between them? Well, Kolmogorov proved the following:

$\vdash_C \mathcal{A} \text{ iff } \vdash_I \mathcal{A}^K.$

To prove this result, you need to prove both directions of the biconditional. It is fairly easy to prove the right-to-left reading: if $\vdash_I \mathcal{A}^K$ then $\vdash_C \mathcal{A}$. The trickier direction is the left-to-right reading: if $\vdash_C \mathcal{A}$ then $\vdash_I \mathcal{A}^K$. To prove this direction, you have to go through all of the basic inference rules of classical FOL, and show that if you translate all of the sentences in accordance with Kolmogorov’s manual, then they become derived inference rules of IL. I won’t go through all of the rules here, but I will pause briefly on TND.

It turns out that we can make our lives considerably easier by swapping TND for DNE. (This isn’t a cheat: in *forall* χ we took TND as a basic classical rule and then derived DNE, but we could have taken DNE as the basic rule and then derived TND.) The crucial observation is then just this. You can tell by looking at the translation manual that for any sentence \mathcal{A} , there is some sentence \mathcal{B} such that $\mathcal{A}^K = \neg\mathcal{B}$. So when we translate the premise and conclusion of a DNE inference, we end up with something of this form: $\neg\neg\neg\mathcal{B} \therefore \neg\mathcal{B}$. And you have already shown that this inference is intuitionistically valid in §2 exercise 3.

Hold on, I thought that intuitionists reject DNE! That’s right, they do, when it is taken as a general rule. But intuitionists also allow you to apply DNE whenever doing so will leave at least one negation at the front of the sentence.

4 Proof-Theoretic Arguments for IL

You now have a fairly good sense of how IL works. But you may have absolutely no idea why anyone would want to use IL, rather than full classical FOL. In this section, I will sketch out one famous argument for IL. But be warned: the path has some twists and turns in it.

4.1 Inferentialism and ‘tonk’

We start with a short paper by Prior, called ‘The Runabout Inference-Ticket’ (1960, *Analysis* 21, pp. 38–9). Prior’s paper wasn’t really about intuitionism at all. Prior was interested in an approach to logic known as *inferentialism*. According to inferentialism, the logical connectives are somehow defined by the inferential rules which govern them. Consider \wedge . It is governed by one introduction rule, and two elimination rules:

m		\mathcal{A}	
n		\mathcal{B}	
		$\mathcal{A} \wedge \mathcal{B}$	$\wedge\text{I}, m, n$

m		$\mathcal{A} \wedge \mathcal{B}$	
		\mathcal{A}	$\wedge\text{E}, m$

m		$\mathcal{A} \wedge \mathcal{B}$	
		\mathcal{B}	$\wedge\text{E}, m$

How are these rules related to the meaning of ‘ \wedge ’? There are two different answers we could give. The first answer goes like this: ‘ \wedge ’ has a meaning which is entirely independent of these rules. (Perhaps this meaning is given by a truth-table.) These rules are justified because they conform to this independent meaning in the appropriate way. The second answer goes like this: ‘ \wedge ’ does not have any independent meaning. The rules *define* the meaning of ‘ \wedge ’. We do not need to justify these rules by showing that they conform to the antecedent meaning of ‘ \wedge ’. ‘ \wedge ’ gets its meaning from these rules. This second answer is the inferentialist answer.

Prior did not like this inferentialist answer. He thought that it threatened to trivialise our whole system of logic. His argument was based on the assumption that *if* inferentialism is true, *then* we can define a new logical connective with any combination of inferential rules. And that seems like a plausible assumption: if the inferential rules really do *define* the logical connective, then who is to stop us from defining a connective with any rules we like? Prior then imagined defining a new connective, ‘tonk’, with the following rules:

m	\mathcal{A}	
	\mathcal{A} tonk \mathcal{B}	tonk-I, m

m	\mathcal{A} tonk \mathcal{B}	
	\mathcal{B}	tonk-E, m

In effect, ‘tonk’ has one of the introduction rules for ‘ \vee ’, and one of the elimination rules for ‘ \wedge ’. Now, on the face of it, you would expect an inferentialist to be happy with these rules for ‘tonk’: they simply *define* a new meaning for that word. Unfortunately, however, adding ‘tonk’ to our natural deduction system will immediately trivialise it: it will let us prove any sentence we like from any other sentence we like:

1	\mathcal{A}	
2	\mathcal{A} tonk \mathcal{B}	tonk-I, 1
3	\mathcal{B}	tonk-E, 2

Prior took it that this amounted to a refutation of inferentialism: if inferentialism were true, then ‘tonk’ would be a perfectly good logical connective; ‘tonk’ is not a perfectly good logical connective; so inferentialism is false.

4.2 Conservativeness

Belnap replied to Prior in his brilliantly titled paper, ‘Tonk, plonk and plink’ (1962, *Analysis* 22, pp. 130–4). Belnap denied that inferentialism implies that ‘tonk’ should be a perfectly good logical connective. According to an inferentialist, the meaning

of a logical connective is defined by the inferential rules that govern it. But that doesn't automatically mean that any arbitrary collection of rules will do serve to define a coherent meaning. There are restrictions on which rules can be coherently combined.

Suppose we wanted to add a new connective, $\$$, to a well behaved logical system Δ . (We will refer to the result of adding this connective to this system as: $\Delta + \$$.) According to Belnap, we must select rules for $\$$ in such a way as to guarantee that $\Delta + \$$ is a *conservative extension* of Δ :

Let Γ be any set of sentences in the language of $\Delta + \$$, and let \mathcal{A} be any sentence in that language which does *not* contain any occurrences of the new connective $\$$.

$\Delta + \$$ is a CONSERVATIVE EXTENSION of Δ iff (if $\Gamma \vdash_{\Delta+\$} \mathcal{A}$, then $\Gamma \vdash_{\Delta} \mathcal{A}$).

In other words, $\Delta + \$$ is a conservative extension of Δ iff the only new things we can prove with $\Delta + \$$ are sentences containing the new connective $\$$; if $\Delta + \$$ allows us prove a sentence not containing $\$$, then Δ must already have allowed us to prove it.

Clearly, adding 'tonk' to a well behaved logical system will not result in a conservative extension of that system: adding 'tonk' will allow us to prove any sentence from any assumption, including sentences which do not feature the connective 'tonk'. So, by Belnap's conservativeness requirement, the rules governing 'tonk' are not admissible, and do not define a meaning for 'tonk'.

On the face of it, this seems like a pretty good solution to Prior's 'tonk' problem. But it has an interesting upshot. Consider a new system, called *Negation-Free Logic* (NFL). This system is exactly like FOL, except it is missing the negation sign and the absurdity sign. I assume that NFL is a well behaved logical system. It may be a little expressively limited, but the system behaves perfectly well within those limits. If we wanted, we could expand NFL by adding a ' \neg ' and ' \perp '. Let's suppose we do that, and that we use the full set of classical rules: \neg I, \perp I, \perp E and TND. What we end up with is classical FOL. But classical FOL is not a conservative extension of NFL. There are negation-free sentences which we can prove in classical FOL, but not in NFL.

Provide proofs for the following (solutions in §7):

1. $\vdash_C ((A \rightarrow B) \rightarrow A) \rightarrow A$
2. $\vdash_C (S \rightarrow T) \vee (T \rightarrow S)$

$$3. \vdash_C (P \leftrightarrow Q) \vee ((P \leftrightarrow R) \vee (Q \leftrightarrow R))$$

Thus it seems that if you accept Belnap’s conservativeness requirement, you should declare the classical rules for negation to be as inadmissible as the rules for ‘tonk’. But interestingly, **IL** *is* a conservative extension of **NFL**. So as far as Belnap’s conservativeness requirement is concerned, there is nothing wrong with intuitionistic negation.

Does this give us any reason to prefer **IL** to classical **FOL**? Well, that partly depends on whether you think that Belnap’s conservativeness requirement is the best way of ruling out connectives like ‘tonk’. In fact, you might take the fact that Belnap’s requirement also rules out classical negation as a proof that Belnap’s requirement was too strong. But if we do say that, how will we rule out ‘tonk’?

4.3 Harmony

Intuitively, the trouble with ‘tonk’ is that its elimination rule is not properly balanced by its introduction rule. A connective’s introduction and elimination rules should be in *harmony* with each other: you shouldn’t be able to get any more out of eliminating a connective than you have to put in to introduce it. (You also shouldn’t get any *less* out by eliminating a connective than you have to put in to introduce it.) Clearly, the rules for ‘tonk’ are not harmonious: you can get a lot more out than you put in.

The idea that good rules are harmonious rules has become very popular amongst inferentialists. Unfortunately, it is not entirely clear what this talk of ‘harmony’ actually means. It is easy enough to get a handle on the intuitive idea, but it is very hard to offer any kind of rigorous definition. What we would really like to do is come up some necessary and sufficient conditions for harmony, i.e. conditions that are met by all and only harmonious sets of rules. Unfortunately, no one has been able to do quite that yet. Happily, however, philosophers and logicians do seem to have settled on a *necessary* condition for harmony.

Here’s an initial thought: if the introduction and elimination rules for connective \$ are in harmony with each other, then you shouldn’t be able to prove anything new just by introducing \$ and then eliminating it. We can make this initial thought rigorous. The first thing we need to do is introduce the concept of a *local peak*:

A LOCAL PEAK for \$ is a use of \$-I followed immediately by a use of \$-E (where this use of \$-E is eliminating the occurrence of \$ introduced in the immediately preceding line).

Here is an example of a local peak for ‘ \rightarrow ’:

1	P	
2		
		P
		P \vee Q
3		\vee I, 2
4	P \rightarrow (P \vee Q)	\rightarrow I, 2-3
5	P \vee Q	\rightarrow E, 4, 1

The local peak appeared at lines 4 and 5: at line 4 we introduced an occurrence of ‘ \rightarrow ’, and then we immediately eliminated it at line 5.

If the introduction and elimination rules for \$ are harmonious, then local peaks for \$ should be entirely redundant. In other words, it should always be possible to *level local peaks* for \$: if there is a proof that contains a local peak for \$, then it should be possible to re-write that proof without that local peak. Here is the official statement of this necessary condition for harmony:

If \$ is governed by harmonious introduction and elimination rules, then there is a procedure for levelling any local peak for \$.

It is fairly easy to show that all of the connectives of NFL meet this necessary condition. Let’s stick with ‘ \rightarrow ’ as our example. A local peak for ‘ \rightarrow ’ always has the same form. At some line in a proof, we prove \mathcal{A} . Then at some later line, we assume \mathcal{A} , and derive \mathcal{B} from that assumption. After that, we introduce $\mathcal{A} \rightarrow \mathcal{B}$ via \rightarrow I, and then immediately apply \rightarrow E, leaving us with \mathcal{B} :

i	\mathcal{A}	
j		\mathcal{A}
		...
		\mathcal{B}
l	$\mathcal{A} \rightarrow \mathcal{B}$	\rightarrow I, $j-k$
$l+1$	\mathcal{B}	\rightarrow E, l, i

The procedure for levelling local peaks for ‘ \rightarrow ’ is pretty straightforward. Simply

take whatever appears within the subproof between lines j and k , and copy it out in the main proof, immediately below \mathcal{A} at line i :

$$\begin{array}{c|c} i & \mathcal{A} \\ \dots & \dots \\ j & \mathcal{B} \end{array}$$

If our original proof of \mathcal{B} from \mathcal{A} , which includes a local peak for ‘ \rightarrow ’, works, then this proof, without that local peak, is guaranteed to work too. So local peaks for ‘ \rightarrow ’ can be levelled. Does this show that the rules for ‘ \rightarrow ’ are harmonious? No, since the requirement that it be possible to level local peaks is only a *necessary* condition for harmony, not a sufficient one. But it does show that the rules for ‘ \rightarrow ’ pass an important test for harmony.

So what about ‘tonk’? Does it pass this test for harmony? No, it doesn’t. The local peaks for ‘tonk’ all look like this:

$$\begin{array}{c|c} m & \mathcal{A} \\ n & \mathcal{A} \text{ tonk } \mathcal{B} \quad \text{tonk-I, } m \\ n + 1 & \mathcal{B} \quad \text{tonk-E, } n \end{array}$$

There is no general procedure for levelling these local peaks for ‘tonk’. After all, \mathcal{A} and \mathcal{B} can be *any* two sentences we like! So ‘tonk’ does not meet this necessary condition for harmony. Thus, if admissible rules are harmonious rules, then the rules for ‘tonk’ come out as inadmissible, just as they should.

But now, given what happened in §4.2, we have to ask: Does classical negation meet this necessary condition for harmony? And interestingly, the answer is: No. But before explaining why, it will be useful to change the classical rules for negation a little bit. In *forall* χ , the Introduction and Elimination rules for negation were all bound up with the Contradiction symbol, \perp . That way of presenting the rules was very convenient, but it is not very helpful in this context: when asking whether the rules for a connective are harmonious, it is much better if those rules only involve that connective. So, let’s swap the old *forall* χ rules with these new variations:

m	\mathcal{A}	
n	\mathcal{B}	
	$\neg\mathcal{A}$	$\neg\text{I}, m-n$

Where \mathcal{B} is an arbitrary atomic sentence, i.e. an atom that does not appear in any undischarged assumptions.

m	\mathcal{A}	
n	$\neg\mathcal{A}$	
	\mathcal{C}	$\neg\text{E}, m, n$

Where \mathcal{C} is an *any* sentence (atomic or complex).

I'll leave it as an exercise for the reader to show that adding these new rules to NFL is equivalent to adding $\perp\text{I}$, $\perp\text{E}$ and the version of $\neg\text{I}$ given in *forall* χ . (But I would also recommend that you read to the end of this section first.)

As well as changing the introduction and elimination rules for negation, it will also be helpful to swap TND for DNE. That's because DNE is a Negation Elimination rule, and harmony is all about balancing introduction and elimination rules. And as I mentioned earlier, there is nothing wrong with this swap: in *forall* χ we used TND as a basic classical rule and then derived DNE, but we could have used DNE as the basic classical rule and then derived TND.

So we now have one Negation Introduction rule, and two Negation Elimination rules. That means that there are two kinds of local peak for ' \neg '. The kind which cause trouble use DNE, and they all look like this:

i	$\neg\mathcal{A}$	
\dots	\dots	
j	\mathcal{B}	
k	$\neg\neg\mathcal{A}$	$\neg\text{I}, i-j$
$k+1$	\mathcal{A}	DNE, k

(where \mathcal{B} is an arbitrary atom). There is no general procedure for levelling these kinds of local peaks. So the full classical rules for negation are not harmonious. So

if good rules are harmonious rules, the classical rules for negation must be dismissed along with the rules for ‘tonk’.

What about the intuitionistic rules for negation? Well, as I already said, we don’t have any *sufficient* condition for being harmonious, but the intuitionistic rules do satisfy our *necessary* condition. For present purposes, we can take the intuitionistic rules for negation to be the versions of \neg I and \neg E given above. Since there is only one introduction rule and one elimination rule, all the local peaks look the same:

$$\begin{array}{r|l|l}
 i & \mathcal{A} & \\
 j & \begin{array}{|l} \mathcal{A} \\ \hline \dots \\ \mathcal{B} \end{array} & \\
 \dots & & \\
 k & & \\
 l & \neg\mathcal{A} & \neg\text{I, } j-k \\
 l+1 & \mathcal{C} & \neg\text{E, } i, l
 \end{array}$$

(where \mathcal{B} is an arbitrary atom, and \mathcal{C} is any sentence). We can level these local peaks in much the same way that we levelled the local peaks for ‘ \rightarrow ’. All we need to do is copy whatever appears within the subproof between lines j and k , and paste it into the main proof immediately below \mathcal{A} at line i , but with every occurrence of \mathcal{B} swapped for \mathcal{C} :

$$\begin{array}{r|l}
 i & \mathcal{A} \\
 \dots & \dots \\
 j & \mathcal{C}
 \end{array}$$

You might be wondering why this is guaranteed to work. After all, \mathcal{B} was an arbitrary *atom*, and \mathcal{C} can be *any* sentence at all, atomic or complex. The important thing to note, though, is that atoms are never easier to prove than complex sentences: any reasoning that works on an atom will work on a complex sentence too; you can just ignore the inner complexity of the complex sentence!

What we have here, then, is another interesting argument in favour of IL. Does it succeed? Philosophers and logicians disagree. But if you are interested, the philosopher who developed this argument furthest was Michael Dummett, in his brilliant but *difficult* book *The Logical Basis of Metaphysics* (1991, Duckworth).

The argument is spread right through the book, but it comes to a head in Chapters 8–13. I really should warn, though, that this book is not easy to read, and no one should feel bad if they struggle to follow it!

5 Semantics for Intuitionistic Logic

So far we have focussed on the natural deduction system for \mathbf{IL} . Now it is time to turn to the semantics.

5.1 Adding Extra Truth-Values

Here’s an initially promising idea. We have already seen that intuitionists reject LEM. And there seems to be a pretty intimate relationship between LEM and the *Principle of Bivalence*, according to which every sentence is true or false (but not both). So it would be natural to try to construct a semantics for \mathbf{IL} which abandons Bivalence. In particular, it is tempting to introduce a new truth-value: a sentence can be true (T), false (F), or neither (N). You might then suggest that the reason we cannot always prove $\mathcal{A} \vee \neg\mathcal{A}$ is that \mathcal{A} might be neither true nor false.

This is all very tempting, but Gödel proved that it won’t work. Or at least, it won’t work if we make these two assumptions: (i) a disjunction is true if at least one of its disjuncts is true; (ii) $\mathcal{A} \leftrightarrow \mathcal{B}$ is true if \mathcal{A} and \mathcal{B} have the same truth-value. To see why, consider the following long disjunction:

$$(A \leftrightarrow B) \vee (A \leftrightarrow C) \vee (A \leftrightarrow D) \vee (B \leftrightarrow C) \vee (B \leftrightarrow D) \vee (C \leftrightarrow D)$$

This long sentence is built out of four atomic sentences: A , B , C and D . Now, since there are only three truth-values to go around, at least two of these sentences *must* have the same truth-value as each other. So by principle (ii), at least one of the biconditionals connecting these atomic sentences must be true. And so by (i), the long disjunction above must also be true. In other words, this disjunction must be a tautology: it is true no matter how we distribute the three truth-values between the four atoms. However, this disjunction is not a theorem of intuitionistic logic.

This argument can be generalised: if we hold onto principles (i) and (ii), then we cannot construct a semantics for \mathbf{IL} which has finitely many truth-values. We could try dodging this result by rejecting (i) or (ii), but that doesn’t look like a very good move: without (i) and (ii), it would be hard to know how to validate \mathbf{IL} ’s rules for disjunction and the biconditional. We could just try living with the result, and adding adding *infinitely many* truth-values, but that doesn’t look very promising either: no two atoms would be allowed to take the same truth-value, and that would

be very strange! Our only other option is to fundamentally change the whole way we think about semantics. That is the option we will try in this primer.²

5.2 The BHK Semantics

In the course of this module, we have looked at a number of different semantic theories for a number of different logics. But while these semantic theories have been different from each other, they have also all had something in common. They always treat TRUTH as the fundamental semantic concept. The point of the semantic theories was to give us a way of determining the truth-values of the sentences in the relevant language.

However, many intuitionists reject the assumption that TRUTH is the fundamental semantic concept. They take WARRANTED ASSERTIBILITY to be the fundamental semantic concept, instead. So rather than trying to explain the meaning of a sentence in terms of its truth-conditions, we should explain its meaning in terms of its assertibility-conditions: once you have told me the circumstances under which I would be warranted to assert a sentence, you have told me everything there is to know about what that sentence means.

What exactly does it mean to say that we are warranted to assert a given sentence? This is a very difficult question. In all likelihood, there is no single answer. Different areas of discourse seem to be governed by different rules for assertion. A science like physics, for example, has quite high standards on assertion: you cannot assert a sentence unless you have good (experimental or theoretical) evidence in its favour. In everyday contexts, by contrast, the standards are much lower: if we are just talking about how our days went, you are warranted to assert any sentence if it fits with your memory of how things were.

Clearly, then, all this talk of warranted assertibility is a little bit slippery. Happily, however, there is at least one context where the rules for assertion are pretty clear: in mathematics, you are warranted to assert a sentence iff you have a proof for that sentence. (At this point it is worth mentioning that intuitionism started off life as a philosophy of mathematics.) So according to intuitionists, if we are presenting a semantics for a mathematical language, we should take PROOF to be our fundamental semantic notion: once you have told me what it would take to *prove* a mathematical sentence, you have told me everything there is to know about what that sentence means.

This kind of semantics was rigorously developed by Heyting, and independently

²A road not taken: we could give a possible-world-style semantics for IL. See Priest's *An Introduction to Non-Classical Logic* (2008, Cambridge University Press, 2nd ed), ch. 6.

by Kolmogorov. It is now known as the BHK semantics: the ‘H’ is for Heyting, the ‘K’ is for Kolmogorov, and the ‘B’ is for Brouwer, the inventor of intuitionism. The semantics gives us a method of describing proofs of more complex sentences in terms of proofs of simpler sentences. (It is taken for granted that we already know what it takes to prove the simplest sentences of the language, i.e. the atoms. In the mathematical case, we might take the atoms to be simple equalities and inequalities, like $2 + 3 = 5$ and $3^2 > 2^3$; you already know how to prove those by elementary computations.) Here is the BHK semantics:

- (1) A proof of $\mathcal{A} \wedge \mathcal{B}$ consists of a proof of \mathcal{A} and a proof of \mathcal{B} .
- (2) A proof of $\mathcal{A} \vee \mathcal{B}$ consists of a proof of \mathcal{A} or a proof of \mathcal{B} .
- (3) A proof of $\mathcal{A} \rightarrow \mathcal{B}$ consists of a method for converting any proof of \mathcal{A} into a proof of \mathcal{B} .
- (4) A proof of $\mathcal{A} \leftrightarrow \mathcal{B}$ consists of a proof of $\mathcal{A} \rightarrow \mathcal{B}$ and a proof of $\mathcal{B} \rightarrow \mathcal{A}$.
- (5) A proof of $\neg \mathcal{A}$ consists of a proof of $\mathcal{A} \rightarrow \perp$.
- (6) A proof of $\exists x \mathcal{A}(x)$ consists of a proof of $\mathcal{A}(c)$, for some element of the domain, c .
- (7) A proof of $\forall x \mathcal{A}(x)$ consists of a method which acts on any element in the domain, c , and delivers a proof that $\mathcal{A}(c)$.

5.3 Intuitionism, Infinity and the BHK Semantics

The BHK semantics fits very well with IL. We have already seen that LEM is not a theorem of intuitionistic logic. Now let’s look at what the semantics has to say about it. Let’s start by focussing on a particular instance of LEM: $G \vee \neg G$. We will take G to be a statement of *Goldbach’s Conjecture*, according to which every even number greater than 2 is the sum of two primes.³ According to the BHK semantics, a proof of $G \vee \neg G$ would consist in a proof of G or a proof of $\neg G$. But as things stand, no one has a proof of either of these things. So, assuming that a mathematical sentence is assertible iff it is provable, no one is currently in a position to assert $G \vee \neg G$.

Now, you might not be too impressed just yet. Sure, no one alive today has a proof of G or a proof of $\neg G$, but maybe that is just because no one has been smart or lucky enough to hit on one just yet! Maybe G or $\neg G$ is provable in the sense

³In formal symbols: $\forall x((x > 2 \wedge \exists y(x = 2y)) \rightarrow \exists y \exists z (Py \wedge Pz \wedge x = y + z))$.

that there is an ideal proof out there waiting to be discovered, and so $G \vee \neg G$ still counts as assertible in this sense.

This may be right. But the question is: What guarantees that there must either be a proof out there for G , or else there must be a proof out there for $\neg G$? If we cannot supply such a guarantee, then we cannot assert $G \vee \neg G$. Of course, this is not to say that we would then have to assert $\neg(G \vee \neg G)$. As you proved in Exercise 1 at the end of §2, $\neg(G \vee \neg G)$ is an intuitionistic contradiction. Rather, the point is that, absent a guarantee that G is provable or that $\neg G$ is, we have to remain absolutely silent on G .

It is important to spot the role that *infinity* has to play in this failure of LEM. Take any *instance* of Goldbach's conjecture, for example: if 2^{100} is an even number greater than 2, then 2^{100} is the sum of two primes. Now, 2^{100} is an absolutely huge even number, and I have no idea if anyone has yet checked whether it is the sum of two primes. But we could check if we liked: just go through all of the primes smaller than 2^{100} , and see if any pair of them add up to 2^{100} .⁴ In technical terms, we say that it is *decidable* whether 2^{100} is the sum of two primes; this means that we have a finite procedure for proving or refuting the claim that 2^{100} is the sum of two primes. As a result, the BHK semantics accepts that the following disjunction is assertible: 2^{100} is the sum of two primes $\vee \neg(2^{100}$ is the sum of two primes). We might not have a proof of either disjunct right now, but we do know that one of these disjunctions is *provable*, and so the disjunction is *assertible*.

The same goes for any other instance of Goldbach's Conjecture, G : every instance is decidable, and so LEM holds for every instance. We only lose our guarantee for LEM when we stop considering instances, and look at the fully general version of Goldbach's Conjecture: *every* even number greater than 2 is the sum of two primes. The trouble is that there are *infinitely* many even numbers. As a result, we cannot go through each even number and check whether or not it is the sum of two primes. Of course, if we start going through the even numbers, we might find one which isn't the sum of two primes. That would be a great result, since that would amount to a proof of $\neg G$. But it might also be that no matter how far we work through the even numbers, we never find a counterexample to Goldbach's Conjecture. In that case, we would keep proving instance after instance of Goldbach's Conjecture, but that would not add up to a proof of the infinite generalisation itself: no matter how far through the even numbers we go, there might be a counterexample waiting around the corner.

⁴There is also a procedure for checking whether a number, n , is prime: just go through all of the numbers below n , and ask whether multiplying any combination of them results in n .

Things would be completely different if we were only dealing with a finite domain. Suppose that \mathcal{F} is a decidable predicate, meaning that we have a finite procedure that we can apply to each object in the domain, and each time it either proves that the object satisfies \mathcal{F} , or it proves that the object does not satisfy \mathcal{F} . If the domain were known to be finite, an intuitionist would be happy to accept this instance of LEM: $\forall x \mathcal{F} x \vee \neg \forall x \mathcal{F} x$. That's because we would have a finite procedure for proving this disjunction: go through each object in the finite domain; if each one satisfies \mathcal{F} , then that is proof of $\forall x \mathcal{F} x$; if at least one does not satisfy \mathcal{F} , then that is proof of $\neg \forall x \mathcal{F} x$.

Things only get interesting when we start dealing with infinite domains. If the domain is infinite, then we might not have any guarantee that $\forall x \mathcal{F} x$ is decidable, even if we already know that \mathcal{F} itself is decidable. Goldbach's Conjecture is an example of this: 'If x is an even number, then x is the sum of two primes' is decidable, but what guarantee do we have that 'Every even number greater than 2 is the sum of two primes' must be decidable? It is in these kinds of cases that the intuitionists reject LEM.

6 Semantic Arguments for IL

Let's take it as given that the BHK semantics fits perfectly with IL: if you accept the BHK semantics (for mathematical languages), then you should accept IL (for mathematical languages).

This gives us a new way of arguing for IL. If someone could convince you that the fundamental semantic notion for mathematical languages is PROOF, not TRUTH, then you should be an intuitionist about mathematics. This is exactly the kind of argumentative strategy that Michael Dummett pursues in a number of different places. As a starting point, I would recommend 'The philosophical basis of intuitionistic logic', reprinted in Putnam and Benacerraf (eds) *Philosophy of Mathematics: Selected Readings*.

6.1 The Manifestation Argument

Dummett presents two versions of the argument. The first has become known as the *Manifestation Argument*. This argument has the form of a *reductio ad absurdum*. Dummett starts off by assuming that the fundamental semantic notion is TRUTH, and derives an absurd result from that assumption. He thus concludes that TRUTH is not the fundamental semantic notion, and suggests that we replace it with PROOF (at least in mathematical discourse).

So, let's start the *reductio* off, and assume that TRUTH is the fundamental semantic notion. Presumably, this means that to understand a sentence is to know its truth-conditions. So when we say that you understand Goldbach's Conjecture, G , we are saying that you know the truth-conditions of G .

Now for Dummett's big idea. Whatever exactly your understanding of G consists in, it is essential that this understanding be fully *manifestable*. That is, you must be able to demonstrate that you have the knowledge which would underlie an understanding of G . That's because meaning is fundamentally *public*: the meaning of our sentences is what we communicate to each other when we use those sentences, and so it must be possible to make that meaning publicly available. So if we assume that to understand G is to know its truth-conditions, we must say this: somehow, the way that you use G manifests your knowledge of the truth-conditions of G .

Now suppose that G is undecidable: no finite procedure will ever prove G or $\neg G$. (It may turn out that G is decidable, but never mind: we can just pick another example!) In that case, the truth-conditions of G are *verification transcendent*, meaning it is beyond our means to verify whether G is true. The heart of Dummett's *reductio* is this question: How would you manifest knowledge of G 's verification-transcendent truth-conditions?

You might have thought that you could do it by *explicitly* stating what those truth-conditions are, like this: G is true iff every even number is the sum of two primes. The trouble with this strategy is that you end up *using* a sentence which has verification-transcendent truth-conditions to state G 's verification-transcendent truth-conditions. This will not be much help to you if you were worried about how you could manifest an understanding of sentences with verification-transcendent truth-conditions!

Ultimately, then, your understanding of a sentence can only be manifested through the way that you *use* it. So if to understand G is to know its truth-conditions, then you must be able to manifest this knowledge through the way that you use G . But what use could manifest knowledge of *verification-transcendent* truth-conditions?

Dummett does not think that there is any good answer to this question, and so concludes that TRUTH cannot be the fundamental semantic concept. (Or at least, it cannot be in the fundamental semantic concept for mathematical discourse.) Dummett recommends that we take PROOF as the fundamental semantic concept instead. While the *truth-conditions* for G may be verification-transcendent, the *proof-conditions* are not. We know what it would take to prove G , and we can manifest that knowledge in various ways. (For example, we can look over putative

proofs of G , and judge whether they are successful.)

6.2 The Acquisition Argument

The second version of Dummett's semantic argument for **IL** has become known as the *Acquisition Argument*. It has the same structure as his Manifestation Argument. He starts by assuming for *reductio* that TRUTH is the fundamental semantic concept, and then tries to derive an absurd result. But this time, rather than asking how you would *manifest* your knowledge of verification-transcendent truth-conditions, he asks how you would *acquire* that knowledge to begin with.

So, let's kick things off again by supposing that to understand G is to know its truth-conditions. And let's again assume that the truth-conditions of G are verification-transcendent: we do not have a finite procedure for proving G or $\neg G$.

Now, you weren't born understanding G . You somehow acquired that understanding over the course of your life. How? Well, in all likelihood you learnt what G meant by observing other people. You might have had a maths teacher who spoke about Goldbach's Conjecture at school. But even if you didn't, you definitely had teachers who used sentences that were related in important ways to Goldbach's Conjecture: you had teachers who explained to you what sums of numbers are, what even numbers are, and what prime numbers are.

So far this all seems well and good, but now we run into trouble with the idea that your understanding of G consists in knowledge of its verification-transcendent truth-conditions. How could watching people use G and related sentences ever lead to you acquiring knowledge of its *verification-transcendent* truth-conditions? There does not seem to be a good answer to this question. It would seem that you would have to observe how people use their mathematical sentences, and then make an educated guess about what the unmanifested, verification-transcendent truth-conditions of G are. But that is an absurd picture of how you acquire linguistic understanding. We can see that it is absurd by noting that on this picture, it would be a kind of semantic good luck if we all guessed at the same truth-conditions for G .

How else might you have come to acquire your understanding of G ? Well, rather than learning it from other people, you might have come up with some or all of the concepts yourself. (You might, for example, be the first person to come up with the concept of a *prime number*.) But how would you go about cooking up your understanding of G all by yourself? Presumably, you would develop these concepts by attempting to *use* them in various contexts. (If you were coming up with the concept of a *prime number*, you would presumably start labelling various numbers as 'prime' when you found that they could only be divided by themselves and by 1.)

Again, this is all well and good, but it is hard to see how this sort of process would ever lead you to knowledge of verification-transcendent truth-conditions.

In short, then, it is hard to see how you would ever come to acquire knowledge of verification-transcendent truth-conditions. Dummett takes this to be another demonstration that TRUTH cannot be the fundamental semantic concept. And again, he suggests that we replace TRUTH with PROOF. After all, acquiring knowledge of what it would take to prove G seems a lot less mysterious than acquiring knowledge of verification-transcendent truth-conditions.

These, then, are Dummett's semantic arguments for IL, squeezed into nutshells. As you can imagine, they are *very* controversial. But we can discuss that in the seminars!

Further Reading

If you are looking for a textbook on intuitionistic logic, then you might want to consider:

- Priest, G. (2008) *An Introduction to Non-Classical Logic*, 2nd ed., Cambridge: Cambridge University Press.
- Dummett, M. (1977) *Elements of Intuitionism*, Oxford: Clarendon Press.

Here are the two papers which kicked off the whole 'tonk' debate:

- Prior, A.N. (1960) 'The runabout inference-ticket', *Analysis* 21, 38–9.
- Belnap, N.D. (1962) 'Tonk, plonk and plink', *Analysis* 22, 130–4.

For a detailed study of different conceptions of harmony, I strongly recommend:

- Steinberger, F. (2011) 'What Harmony Could and Could Not Be', *Australasian Journal of Philosophy* 89, 617–39.

And here are some of Dummett's classic papers on intuitionism:

- Dummett, M. (1958–9) 'Truth', *Proceedings of the Aristotelian Society* 59, 141–62.
- ——— 'The philosophical basis of intuitionistic logic', reprinted in Putnam & Benacerraf eds (1983) *Philosophy of Mathematics: Selected Readings*, 2nd Edition, Cambridge: Cambridge University Press.
- ——— (1982) 'Realism', *Synthese* 52, 55–112.

7 Solutions

Solutions for §2

Provide proofs for the following:

- $\neg(A \vee \neg A) \vdash \perp$

1		$\neg(A \vee \neg A)$	
2			A
3			$A \vee \neg A$ $\vee I, 2$
4			\perp $\perp I, 3, 1$
5		$\neg A$	$\neg I, 2-4$
6		$A \vee \neg A$	$\vee I, 5$
7		\perp	$\perp I, 6, 1$

- $A \vee B \vdash \neg\neg A \vee \neg\neg B$

1		$A \vee B$	
2			A
3			$\neg A$
4			\perp $\perp I, 2, 3$
5		$\neg\neg A$	$\neg I, 3-4$
6		$\neg\neg A \vee \neg\neg B$	$\vee I, 5$
7			B
8			$\neg B$
9			\perp $\perp I, 7, 8$
10		$\neg\neg B$	$\neg I, 8-9$
11		$\neg\neg A \vee \neg\neg B$	$\vee I, 10$
12		$\neg\neg A \vee \neg\neg B$	$\vee E, 1, 2-6, 7-11$

3. $\neg\neg A \vdash_1 \neg A$

1	$\neg\neg A$	
2	A	
3	$\neg A$	
4	\perp	\perp I, 2, 3
5	$\neg\neg A$	\neg I, 3-4
6	\perp	\perp I, 5, 1
7	$\neg A$	\neg I, 2-6

4. $\neg\neg\neg\neg Px \vdash_1 \neg\neg\neg\neg Px$

1	$\neg\neg\neg\neg Px$	
2	$\forall x\neg\neg\neg Px$	
3	$\neg\neg\neg Pa$	\forall E, 2
4	$\neg Pa$	
5	$\neg\neg Pa$	
6	\perp	\perp I, 4, 5
7	$\neg\neg\neg Pa$	\neg I, 5-6
8	$\neg Pa \rightarrow \neg\neg\neg Pa$	\rightarrow I, 4-7
9	$\neg\neg Pa$	MT, 8, 3
10	$\forall x\neg\neg Px$	\forall I, 9
11	\perp	\perp I, 10, 1
12	$\neg\neg\neg\neg Pa$	\neg I, 2-11

5. $\vdash_1 ((\neg A \rightarrow \neg B) \rightarrow \neg A) \rightarrow \neg A$

1		$(\neg A \rightarrow \neg B) \rightarrow \neg A$	
2		$\neg A$	
3		$\neg A$	
4		\perp	$\perp I, 3, 2$
5		$\neg \neg A$	$\neg I, 3-4$
6		$\neg(\neg A \rightarrow \neg B)$	$MT, 1, 5$
7		$\neg A$	
8		\perp	$\perp I, 2, 7$
9		$\neg B$	$\perp E, 8$
10		$\neg A \rightarrow \neg B$	$\rightarrow I, 7-9$
11		\perp	$\perp I, 10, 6$
12		$\neg \neg A$	$\neg I, 2-11$
13		$((\neg A \rightarrow \neg B) \rightarrow \neg A) \rightarrow \neg A$	$\rightarrow I, 1-12$

Solutions for §4

Provide proofs for the following:

1. $\vdash_C ((A \rightarrow B) \rightarrow A) \rightarrow A$

1			$(A \rightarrow B) \rightarrow A$	
2			$\neg A$	
3		$\neg(A \rightarrow B)$		MT, 1, 2
4		A		
5		\perp		\perp I, 4, 2
6		B		\perp E, 5
7		$A \rightarrow B$		\rightarrow I, 4–6
8		\perp		\perp I, 7, 3
9		$\neg\neg A$		\neg I, 2–8
10		A		DNE, 9
11		$((A \rightarrow B) \rightarrow A) \rightarrow A$		\rightarrow I, 1–10

2. $\vdash_C (S \rightarrow T) \vee (T \rightarrow S)$

1		S	
2		T	
3		S	R, 1
4		$T \rightarrow S$	\rightarrow I, 2–3
5		$(S \rightarrow T) \vee (T \rightarrow S)$	\vee I, 4
6		$\neg S$	
7		S	
8		\perp	\perp I, 7, 6
9		T	\perp E, 8
10		$S \rightarrow T$	
11		$(S \rightarrow T) \vee (T \rightarrow S)$	\vee I, 10
12		$(S \rightarrow T) \vee (T \rightarrow S)$	TND, 1–5, 6–11

3. $\vdash_C (P \leftrightarrow Q) \vee ((P \leftrightarrow R) \vee (Q \leftrightarrow R))$

1	$\neg((P \leftrightarrow Q) \vee ((P \leftrightarrow R) \vee (Q \leftrightarrow R)))$	
2	$\neg((P \leftrightarrow Q) \wedge \neg((P \leftrightarrow R) \vee (Q \leftrightarrow R)))$	DeM, 1
3	$\neg((P \leftrightarrow R) \vee (Q \leftrightarrow R))$	\wedge E, 2
4	$\neg(P \leftrightarrow R) \wedge \neg(Q \leftrightarrow R)$	DeM, 3
5	$\neg(P \leftrightarrow Q)$	\wedge E, 2
6	$\neg(P \leftrightarrow R)$	\wedge E, 4
7	$\neg(Q \leftrightarrow R)$	\wedge E, 4
8	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> P </div>	
9	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> Q </div>	
10	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> P </div> </div>	
11	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> Q </div> </div>	R, 9
12	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> Q </div> </div>	
13	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> P </div> </div>	R, 8
14	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> $P \leftrightarrow Q$ </div>	\leftrightarrow I, 10–11, 12–13
15	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> \perp </div>	\perp I, 14, 5
16	$\neg Q$	\neg I, 9–15
17	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> R </div>	
18	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> P </div> </div>	
19	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> R </div> </div>	R, 17
20	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> R </div> </div>	
21	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> <div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> P </div> </div>	R, 8
22	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> $P \leftrightarrow R$ </div>	\leftrightarrow I, 18–19, 20–21
23	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> \perp </div>	\perp I, 22, 6
24	$\neg R$	\neg I, 17–23
25	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> Q </div>	
26	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> \perp </div>	\perp I, 25, 16
27	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> R </div>	\perp E, 26
28	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> R </div>	
29	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> \perp </div>	\perp I, 28, 24
30	<div style="border-left: 1px solid black; padding-left: 5px; margin-left: 5px;"> Q </div>	\perp E, 29
31	$Q \leftrightarrow R$	\leftrightarrow I, 25–27, 28–30

32	\perp	\perp I, 31, 7
33	$\neg P$	\neg I, 8–32
34	$\neg Q$	
35	P	
36	\perp	\perp I, 35, 33
37	Q	\perp E, 36
38	Q	
39	\perp	\perp I, 38, 34
40	P	\perp E, 39
41	$P \leftrightarrow Q$	\leftrightarrow I, 35–37, 38–40
42	\perp	\perp I, 5
43	$\neg\neg Q$	\neg I, 34–42
44	R	
45	Q	
46	R	R, 44
47	R	
48	Q	DNE, 43
49	$Q \leftrightarrow R$	\leftrightarrow I, 45–46, 47–48
50	\perp	\perp I, 49, 7
51	$\neg R$	\neg I, 44–50
52	P	
53	\perp	\perp I, 52, 33
54	R	\perp E, 53
55	R	
56	\perp	\perp I, 55, 51
57	P	\perp E, 56
58	$P \leftrightarrow R$	\leftrightarrow I, 52–54, 55–57
59	\perp	\perp I, 58, 6
60	$\neg\neg((P \leftrightarrow Q) \vee ((P \leftrightarrow R) \vee (Q \leftrightarrow R)))$	\neg I, 1–59
61	$(P \leftrightarrow Q) \vee ((P \leftrightarrow R) \vee (Q \leftrightarrow R))$	DNE, 60