

A Guide to TFL Proof Rules (for Worksheets 5, onward)

In this lesson sheet, I will be doing things slightly differently. The objective will be to walk through all of our basic TFL deduction rules and to make sense of *why* they work in the way that they do. The key to learning natural deduction is first and foremost to learn the rules--it will be vital that you remember them for the exam! For now, it's okay if you need to refer back to the text, but you will find the proofs will be easier to complete the more familiar you are with all of the relevant rules. To this end, in what follows, I will explain the rules in a way that connects them with the truth-tables from early last term. Hopefully, this will prove helpful, both for committing the rules to memory, but also for working with the rules in any given proof.

CONJUNCTION:

I'll begin with conjunction. Let's start by recalling the truth-table for conjunction.

A	B	$A \wedge B$
T	T	T
T	F	F
F	T	F
F	F	F

In general, our primary concern will, of course, be with the line(s) of the truth-table on which the proposition (in which the relevant connective is the main operator) is true. Our proofs have to **truth-preserving**, and thus we do not want rules that allow us to 'prove' propositions that are false. For instance, we do not want a conjunction rule that would allow us to infer $(A \wedge B)$ when only A is true.

With this in mind, let's consider the first line of the truth-table above. It shows us that conjunctions are true when and only when both conjuncts are true. From this, we can make sense of our **Conjunction Elimination** rule (abbreviated as ' $\wedge E$ ' in our proofs).

Since conjunctions are true if and only if both conjuncts are true, *whenever* we have some conjunction $A \wedge B$, we know that A is true and also that B is true. Hence:

m	$A \wedge B$	
	A	$\wedge E m$

and equally:

m	$A \wedge B$	
	B	$\wedge E m$

Similarly, since we know that whenever A is true and also B is true, their conjunction will be true (i.e. $A \wedge B$), we have the following rule for **Conjunction Introduction** (abbreviated as ' $\wedge I$ ' in our proofs):

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

CONDITIONAL:

Let's start with the relevant truth-table again.

A	B	$A \rightarrow B$
T	T	T
T	F	F
F	T	T
F	F	T

Here, the only possibility we want to rule out is that in which the antecedent is true and the consequent false. But, the truth of the conditional does not *demand* the truth of the antecedent, or the truth of the consequent. The antecedent could be true or false and the conditional still true; similarly, the consequent could be true or false, and the conditional still true.

Enter the rule for **Conditional Introduction** (abbreviated as ' $\rightarrow I$ ' in our proofs).

i	<table style="border-collapse: collapse;"> <tr> <td style="border-right: 1px solid black; padding: 5px;">\mathcal{A}</td> <td></td> </tr> <tr> <td style="border-right: 1px solid black; padding: 5px;">\mathcal{B}</td> <td style="border-top: 1px solid black;"></td> </tr> </table>	\mathcal{A}		\mathcal{B}		
\mathcal{A}						
\mathcal{B}						
	$\mathcal{A} \rightarrow \mathcal{B}$	$\rightarrow I\ i-j$				

This rule tells us that we can prove a conditional by showing that, if we *assumed* A we could prove B. In other words, we could show that if A is true, then B is also true. BUT, importantly, we do so without saying whether A or B is *in fact* true. This is in keeping with our observations about the truth-table from above. We don't know for sure if A and/or B are true; but the sub-proof shows us that IF A is true, THEN B is also true. Hence we derive the conditional $A \rightarrow B$

At this point, it's worth taking a brief moment to talk about **assumptions** in proofs. In a proof, we use vertical lines to indicate *scope* of a given argument. The left-most line is the one best thought of as the "reality" line. Propositions on the left-most line -- the reality line -- are the things we know to be true (in the context of that proof). When we make an assumption, we are *supposing* "for the sake of argument" that something is the case (e.g. A, in the example from above), to see what follows (e.g. B in the example above). We can always move things from the reality line to an assumption; however (this is important!) there are only very specific circumstances in which we can take things from the assumption to the reality line. Consider the rule for Conditional Introduction again. I could not, for instance, write the proposition A on the reality line -- A only holds *within* the scope of the assumption. I don't know whether, *in reality*, A holds (unless, of course, I can prove it from the initial premises). What's important is that, just because a proposition holds in an assumption, it does not necessarily mean that it holds *simpliciter*, i.e. it does not necessarily mean that that proposition can show up on the reality line. The proof rules will provide an exhaustive list of circumstances in which you are allowed to move propositions from an assumption to the reality line.

Let's turn now to the **Conditional Elimination** rule (abbreviated ' $\rightarrow E$ ' in our proofs). This is basically the rule for Modus Ponens.

m	$\mathcal{A} \rightarrow \mathcal{B}$	
n	\mathcal{A}	
	\mathcal{B}	$\rightarrow E\ m, n$

We can make sense of this rule by thinking about the truth-table again. If a conditional is true, and if its antecedent is also true, then it simply *must* be the case that the consequent is true. After all, Line 1 of the truth table is the only line on which the antecedent is true and the entire conditional is also true.

That exhausts the basic rules involving conditionals, however there's a further derived rule **Modus Tollens** (abbreviated 'MT').

m	$\mathcal{A} \rightarrow \mathcal{B}$	
n	$\neg \mathcal{B}$	
	$\neg \mathcal{A}$	$MT\ m, n$

This too can be understood in terms of our truth-table. Consider the lines on which consequent of the conditional is false, but the whole conditional is still true. There is, in fact, only one line on which this is

the case: the last one, i.e. the one on which the antecedent is also false. So, from a conditional $A \rightarrow B$ and the proposition $\neg B$, we can prove $\neg A$.

BICONDITIONAL

Unsurprisingly, the rules for biconditional are very similar to the rules for conditional. Let's start with the truth-table:

A	B	$A \leftrightarrow B$
T	T	T
T	F	F
F	T	F
F	F	T

The rule for **Biconditional Introduction** (abbreviated ' $\leftrightarrow I$ ') is basically the same as the rule for $\rightarrow I$, except that you need to do it twice! Once in each direction. Recall that a biconditional $A \leftrightarrow B$ is equivalent to $(A \rightarrow B) \wedge (B \rightarrow A)$; essentially, when we use a biconditional, we are saying that the conditional goes in both directions. So, that means that we need to show that (a) whenever A is true, B is also true; but also (b) whenever B is true, A is also true. Thus:

i	\mathcal{A}	
j	\mathcal{B}	
k	\mathcal{B}	
l	\mathcal{A}	
	$\mathcal{A} \leftrightarrow \mathcal{B}$	$\leftrightarrow I\ i-j, k-l$

Notice once again that we do this by using assumptions and sub-proofs. We don't know for sure whether A and B are true or false, but the rule relies on the fact that, in order for the biconditional to be true, it must be true that (a) whenever A is true, B is also true; and (b) whenever B is true, A is also true. This reflects the fact that the biconditional is true on both lines 1 and 4 of the truth-table, so it might well be that both A and B are false. It also reflects that *whenever* one of the two is true, so to is the other. Just as the truth-table shows.

Similarly, **Biconditional Elimination** (abbreviated ' $\leftrightarrow E$ ') is very like $\rightarrow E$. Except, in this case, we know that whenever the antecedent is true, so too is the consequent AND whenever the consequent is true, so too is the antecedent. In other words, it's as though we can perform modus ponens in both directions. Hence:

m	$\mathcal{A} \leftrightarrow \mathcal{B}$	
n	\mathcal{A}	
	\mathcal{B}	$\leftrightarrow E\ m, n$

and equally:

m	$\mathcal{A} \leftrightarrow \mathcal{B}$	
n	\mathcal{B}	
	\mathcal{A}	$\leftrightarrow E\ m, n$

Now, you might be thinking something like the following: “If all this is true about the biconditional, then I should also be able to perform *modus tollens* in both directions!” And you wouldn’t be misguided in thinking that. *However* (and this is very important) as a matter of fact, Tim has not given us such a rule in our system. So, we won’t be allowed to infer from $A \leftrightarrow B$ and $\neg B$ that $\neg A$. Or, at least, you can’t do this in a single step, because we don’t have a rule for doing so. That said, you could still prove this using your basic rules. [If you like, try to do this as an exercise. Give a proof showing that: $A \leftrightarrow B, \neg B \therefore \neg A$]

DISJUNCTION:

As above, we’ll start with the relevant truth-table:

A	B	A \vee B
T	T	T
T	F	T
F	T	T
F	F	F

Notice that it is both necessary and sufficient for the truth of a disjunction that *at least one* of the disjuncts is true. I.e. A disjunction is true if and only if at least one (but maybe both!) disjuncts are true; this will encompass line 1 of the above truth table as well as lines 2 and 3, since, when BOTH disjuncts are true it is true that at least one of them is true. Based on this, we can know that, if one disjunct is true *regardless of the truth-value of the other* (whether the other is true or false), the entire disjunction is true.

Here, then, is our rule for **Disjunction Introduction** (abbreviated ‘ $\vee I$ ’):

m	\mathcal{A}	$\mathcal{A} \vee \mathcal{B}$	$\forall I\ m$
-----	---------------	--------------------------------	----------------

and

m	\mathcal{A}	$\mathcal{B} \vee \mathcal{A}$	$\forall I\ m$
-----	---------------	--------------------------------	----------------

According to this rule, whenever some proposition \mathcal{A} is true, we can disjoin it with ANY other proposition we like. And, because of the truth-table for disjunction, we know that *regardless* of what we choose to disjoin \mathcal{A} with, that disjunction $\mathcal{A} \vee __$ will be true. We know this because the truth of just one of the disjuncts was enough (i.e. sufficient) for the truth of the disjunction. The only line on which the disjunction is false is that on which BOTH disjuncts are false (as in line 4 of the truth-table).

Let's turn now to **Disjunction Elimination** (abbreviated 'vE'). We won't look at the formal rule just yet, though. First, consider some disjunction that I tell you is true. To make things concrete, let's use the disjunction 'Either it is raining or it is snowing' -- $\text{Rain} \vee \text{Snow}$. What do we know about this disjunction in this context? Well, since we're taking this as our premise, we know that it's true, so we know, based on the truth-table, that either Rain is true or Snow is true. BUT, we don't know which.

Now let's imagine that later in our conversation, I say that "If it's raining, then I'll wear my boots". "Well that's nice," you might think to yourself, but there's nothing you can do with this, since you don't know whether it's raining! It's just as you're thinking this that I tell you, "And, if it's *snowing*, then I'll wear my boots." Is there anything we can do with this, now?

Well, let's take stock of what we know. We know that the disjunction ($\text{Rain} \vee \text{Snow}$) is true. Based on the truth-table, that means we know ONE of these (i.e. Rain, or Snow) has to be true. And from the rest of the conversation, we know both that (a) if Rain, then I'll wear my boots, and (b) if Snow, then I'll wear my boots. But, since it *HAS* to be the case that either Rain is true or Snow is true, and either way I'll wear my boots, we know for sure that I'll wear my boots! Notice that, we can know I'll wear my boots, even though we still don't know which of Rain or Snow is true. Just that the disjunction ($\text{Rain} \vee \text{Snow}$) is true.

Let's look at the formal rule now:

m	$\mathcal{A} \vee \mathcal{B}$	
i	\mathcal{A}	
j	\mathcal{C}	
k	\mathcal{B}	
l	\mathcal{C}	
	\mathcal{C}	$\vee E\ m, i-j, k-l$

This is effectively a formalisation of the natural-language reasoning we just went through about me and my boots. It can be a bit confusing to think of this as disjunction elimination, since we don't end up proving one of the disjuncts. In this case, what we are 'eliminating', so-to-speak, is the entire disjunction. (I say "so-to-speak" since the disjunction still exists in our proof, and could be used again later, if we needed it.) The rule turns on the fact that (as we learn from the truth-table) at least one of the disjuncts in a disjunction must be true, whenever a disjunction is true. After that, it obeys the following reasoning: if something follows from both of the disjuncts of a disjunction, then it also follows from the disjunction.

Tertium Non Datur (TND)

Before moving on to our last connective, I want to mention TND (a.k.a Tertium Non Datur, a.k.a No Third Way), since it is effectively a version of Disjunction Elimination. In the system in which we are working, ALL propositions are either true or false. In other words, it is *always* the case that, for any A, $A \vee \neg A$. That disjunction is a *tautology*--it is always true. Now, given that this is a disjunction, it follows the same rules, so from the disjunction $A \vee \neg A$, if we can derive some proposition B from A and also derive it from $\neg A$, then we know B is true as well. Hence, TND:

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

Notice, though, that there is no line on which we have $A \vee \neg A$ in this rule. The thought is, since $A \vee \neg A$ is a tautology, it can be proven from no premises at all; and, as this is the case, we forgo having to prove it in order to use the TND rule. Whenever you can prove something from a proposition and from its negation, then you are allowed to infer that that something is true.

One interesting thing to note is that this rule would not work in a system that does not have a Law of Excluded Middle (e.g. intuitionistic logic). Law of Excluded Middle (LEM) demands that there only be two truth-values -- there is *no third way*.

Disjunctive Syllogism

Staying with the disjunction, the next rule we'll look at is **Disjunctive Syllogism** (abbreviated 'DS'). We are still concerned with the truth-table for disjunction, so I'll repeat it here just to refresh our memories.

A	B	$A \vee B$
T	T	T
T	F	T
F	T	T
F	F	F

The line's we're interested in here are lines 2 and 3 of the table. There is only one way a disjunction can be true if one of the disjuncts is false: the other disjunct *must* be true. Remember *at least one* disjunct must be true in order for the disjunction to be true. And based on this reasoning, we have the rule for DS:

m	$\mathcal{A} \vee \mathcal{B}$	
n	$\neg \mathcal{A}$	
	\mathcal{B}	DS m, n

and

m	$\mathcal{A} \vee \mathcal{B}$	
n	$\neg \mathcal{B}$	
	\mathcal{A}	DS m, n

NEGATION

Before I turn to the proof rules for negation, I need to explain the rule for **Contradiction Introduction** (abbreviated ' $\perp I$ '). This one is quite straightforward. It is simply this: whenever you have some proposition A and its negation $\neg A$, you can derive \perp (i.e. a contradiction). Hence:

m	A	
n	$\neg A$	
	\perp	$\perp I\ m, n$

With this rule at our disposal, we can look at the rule for **Negation Introduction** (abbreviated ' $\neg I$ '). This is the rule we use for proving the negation of some proposition A ; i.e. for proving $\neg A$. The kind of reasoning involved in this rule will be familiar to you as *reductio ad absurdum*. To prove a negation $\neg A$, we assume the opposite A , and show (within the scope of that assumption) that this leads to contradiction. Absurdity. But, if we used perfectly acceptable reasoning to get from the assumption A to a contradiction, then there must have been something wrong with the assumption! So, we discharge the assumption, and derive $\neg A$.

i	<table style="border-collapse: collapse;"> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">A</td> </tr> <tr> <td style="border-left: 1px solid black; padding-left: 10px;">\perp</td> </tr> </table>	A	\perp	
A				
\perp				
j	$\neg A$	$\neg I\ i-j$		

Remember: nothing from inside the sub-proof can come out to the left-most line--the "reality" line. Once we get to the contradiction, only then can we discharge the assumption, and write the *negation of the assumption* on the reality line.

Now, this rule will come in handy when you're proving atomic propositions and their negation. Any time you're trying to prove a proposition without some other main operator, and you can't see anything else to do, give Negation Introduction a shot; assume the opposite of the conclusion and try to prove a contradiction.

But, say I wanted to prove A . I start by assuming $\neg A$, and derive a contradiction.

i	$\neg A$	
j	\perp	
k	??	$\neg I\ i-j$

What should I conclude on line k? Answer: $\neg\neg A$. You *cannot* conclude that A. While it's true that A and $\neg\neg A$ are tautologically equivalent, the rule we used was for negation *introduction*, so when we finish executing the rule, we need to introduce a negation.

i	$\neg A$	
j	\perp	
k	$\neg\neg A$	$\neg I$ i-j

However, in order to derive A from this, we can use the derived rule **Double-Negation Elimination** (abbreviated 'DNE'). This rule is as simple as this: whenever you have $\neg\neg A$, you can derive A.

m	$\neg\neg A$	
	A	DNE m

There's just a few more rules to cover. I'll start with **Contradiction Elimination** (abbreviated ' $\perp E$ '). This one is also a straightforward one, but an easy one to forget when doing a proof. The idea behind this rule is that, from a contradiction *anything follows* -- in Latin, *ex falso quodlibet*. So, whenever we have a contradiction, we can derive *any proposition we like*.

m	\perp	
	A	$\perp E$ m

(Interestingly, this too is a controversial rule. It works in our system because TFL adheres to the *Law of Non-Contradiction*; but not all systems accept this law. E.g. paraconsistent logics)

Next, **Reiteration** (abbreviated, 'R'). This rule allows us to repeat propositions in our proofs, when necessary. Often we have to do this so that a proposition from the main proof can be used in some sub-proof. HOWEVER, remember the rules for moving propositions over scope lines! You *cannot* use R to move a proposition from *inside* a sub-proof out to the reality line. In general, remember that you can move a proposition RIGHTWARD over scope-lines whenever you like using R, not leftward. (To move propositions leftward over scope-lines, you must use one of the rules previously given.) With all that in mind, here's the reiteration rule:

m	A	
	A	R m

Finally, **De Morgan Rules** (abbreviated, 'DM'). The key to the De Morgan Rules lies in truth-tables for conjunction and disjunction. But this time, we're interested in the lines on which the con/disjunction is FALSE. Here's a truth-table that will be helpful for understanding DM rules.

1	2	3	4	5	6
A	B	$\neg A \vee \neg B$	$\neg(A \wedge B)$	$\neg A \wedge \neg B$	$\neg(A \vee B)$
T	T	F	F	F	F
T	F	T	T	F	F
F	T	T	T	F	F
F	F	T	T	T	T

DM rules give us a way of moving back and forth between Columns 3 and 4 (green), and also for moving back and forth between Columns 5 and 6 (purple). Predictably, there are four DM rules:

The first De Morgan rule is:

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

The second De Morgan is the reverse of the first:

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

The third De Morgan rule is the *dual* of the first:

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

And the fourth is the reverse of the third:

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

And that exhausts all of our basic and derived rules for TFL.

A rough guide to approaching proofs:

Before closing off, I want to leave you with a rough-and-ready guide to approaching proofs. Like I've said before, there isn't a set of one-size-fits-all instructions I can give you for tackling proofs, since (a) there's almost always more than one way to complete any given proof, and (b) the approach differs in every case. However, all of that being said, there are some rules of thumb that you might find helpful. In general, I'd recommend the following rough procedure:

0. Prepare for Dead-ends! (← plural!)

This part is perhaps the most important. You are *going* to hit dead-ends when trying to do proofs. This is okay! We all go down blind alleys, including those of us who've been doing these for a while. Don't be disheartened--it doesn't mean that you're no good at this, or that you can't do it. If you're really stuck on one proof, walk away from it for a while and come back to it, but keep hacking away.

1. Check the Conclusion -- Where are you going? What's the main operator?

One indicator for how you might start a proof lies in the main operator of the conclusion you're after. Find the main operator, then try starting by using the introduction rule for that operator.

- Is the main operator a negation?
 - Assume the opposite, derive a contradiction
- Is the main operator a disjunction?
 - try to prove one of the disjuncts
- Is the main operator a conjunction?
 - try to prove each of the conjuncts
- Is the main operator a conditional?
 - Assume the antecedent, derive the consequent
- Is the main operator a biconditional?
 - Assume the antecedent/consequent, derive the consequent/antecedent, and vice versa
- Is the proposition an atomic proposition?
 - Assume the opposite, derive a contradiction, DNE

These aren't going to be fool-proof ways to start. You might hit one of those dead-ends trying these, but they give you something to try if you can't see anything else.

2. Check Your Premises -- What's the main operator? What can you get?

Sometimes your premises will require "unpacking". There may be things that will later be helpful in getting to your conclusion. Don't worry if you can't tell yet whether you need a given proposition; just unpack everything you can. What do I mean by this? Well, go through all of your elimination rules.

- Are there any conjunctions?
 - Derive the conjuncts with $\wedge E$
- Are there any (bi)conditionals? And their antecedent (/consequent)?
 - Derive the consequent (/antecedent) with $\rightarrow E$ ($\leftarrow \rightarrow E$)
- Are there any disjunctions? And the negation of one of the disjuncts?
 - Derive the other disjunct with DS
- Are there any disjunctions? But nothing else to do with them?
 - Assume a disjunct, prove the same thing from each disjunct for $\vee E$

3. Break the Proof into Steps -- Don't try to do/see it all at once!

There are precious few of us who can just look at a set of premises and see how to get to the conclusion. Think about smaller steps. Just think about the next few propositions. If you're trying to prove a conjunction ($A \wedge B$), you know you need to prove each of the conjuncts A and then B. Great! Once you've worked that out, make a note to yourself (either physically on the page, or a mental note) that your intermediate conclusion is A, and then try to prove that. At that point, you don't even need to worry about the conjunction--just about proving A. Maybe, to prove A, you need to prove some other ($C \vee D$). Well, then just repeat Steps 1 and 2, treating ($C \vee D$) as your conclusion. I always make a note on my page of my main conclusion and my intermediate conclusions so I always know where I'm going. If you're visual thinker, it can be really helpful just to be able to see what you're shooting for.

4. When all else fails, assume the opposite of the conclusion for $\neg I$.

If you can't see any way to start based on the main operator of the conclusion, and the premises, try assuming the negation of the conclusion, and deriving a contradiction.

5. When *that* fails, TND.

TND is the rule that everyone forgets to try. It's easy to forget since it's not based on introducing or eliminating one of our connectives. If you can't see any other way in, and if assuming the negation of the conclusion doesn't work, then remember to give TND a go. What should you assume for your disjunction? I.e. what should you use for your disjuncts? Have a look at your premises; see if there's a way of generating contradictions, of eliminating conditionals, of using disjunctive syllogism. In general, look for something that is going to interact with your other propositions in some way.

6. Rinse and Repeat!

If at first you don't succeed, try, try, try..... etc.

When you're doing a proof, don't worry how long your proof is getting -- if it's getting long, it doesn't necessarily mean it's wrong. In fact, it doesn't matter how long your proof is; as long as each step is permissible, and you arrive at the desired conclusion, the proof is correct! That remains true, even if there was a shorter way to do it. (Note: This means, even if your answer doesn't match the answer key, it may not necessarily be wrong. Ask us if you're ever unsure!)

Finally, I can't emphasise enough the benefit of practice here. The more you work at proofs, the better you'll get at it. It's not enough to study the proof rules; you have to *do* them. Knowledge of doing proofs is as much knowledge-how as it is knowledge-that!