Arguments and Proofs

For the next section of this course, we will study PROOFS. A proof can be thought of as the formal representation of a process of reasoning. Proofs are comparable to arguments, since they can be seen as having premises and conclusions. That is only part of the story, however.

Let's begin by returning to the definition of an argument. An argument is two things:

1. A set of sentences (the **premises**)

2. A sentence (the **conclusion**)

Arguments Made of WFFs

Our language for sentential logic also contains sentences, so this definition can be applied to it. A sentence in the formal language is a **well-formed formula** (**wff**). So, we can also say that an argument **in sentential logic** is:

1. A set of wffs. (the **premises**)

2. A wff. (the **conclusion**)

An example would be:

Premises: $(P \vee \sim S)$

$\sim(Q \leftrightarrow (R \& S))$

$(P \rightarrow (Q \& R))$

Conclusion: $(Q \vee (P \leftrightarrow \sim R))$

Sequents and Arguments

For convenience, and just because it's what we want to do, we will write arguments in another way, and when we do that we will call them by another name: **SEQUENTS**. A sequent is just an argument written on one line, with the premises first (separated by commas) and the conclusion last (separated from the premises by the symbol '$\vdash$'). So, the argument on the previous page, written as a sequent, looks like this:

$$(P \vee \sim S), \sim(Q \leftrightarrow (R\&S)), (P \rightarrow (Q\&R)) \vdash (Q \vee (P \leftrightarrow \sim R))$$

The symbol '$\vdash$' is called the **TURNSTILE**.

A Small Confession

In general, I have not required that you learn the details of the parenthesis-dropping convention. However, in these notes the **outside pair of parentheses** for wffs will often be dropped. That is, you will find such forms as $P \rightarrow \sim(Q \& R)$ instead of the strictly correct $(P \rightarrow \sim(Q \& R))$.

Valid Arguments and Reasoning

An argument is a set of premises and a conclusion. If the premises and conclusion are related in such a way that it's impossible for the conclusion to be false if the premises are all true, then we call it **valid**. How do we tell if an argument is valid?

One way, and one that we often use, is to see if we can **reason deductively** from the premises to the conclusion. It's best to explain this with an example. Consider the following argument:

$P1$	My kitchen is ten feet long and twelve feet wide.
$P2$	George's kitchen is eight feet long and fourteen feet wide.
C	Therefore, my kitchen is bigger than George's kitchen.

Why is it valid? You might defend this by reasoning as follows:

$P1$	My kitchen is ten feet long and twelve feet wide.
$P2$	George's kitchen is eight feet long and fourteen feet wide.
$R3$	*If my kitchen is ten feet long and twelve feet wide, then its area is one hundred twenty square feet*
$R4$	*So, the area of my kitchen is one hundred twenty square feet*
$R5$	*If George's kitchen is eight feet long and fourteen feet wide, then its area is one hundred twelve square feet*
$R6$	*So, the area of George's kitchen is one hundred twelve square feet*
$R7$	*One hundred twenty square feet is a larger area than one hundred twelve square feet*
C	Therefore, my kitchen is bigger than George's kitchen.

The additional sentences in red are a series of steps that start with the premises of the argument and eventually reach the conclusion. Each step follows necessarily from previous steps, and the last step is the conclusion. Therefore, these additional sentences **show that this argument is valid**. Let's call these intermediate steps **REASONING**.

Reasoning and Proof

A **proof** is like an argument with reasoning added to show that it is valid. However, proofs are in the formal language of sentental logic, not English. We use proofs to model the process of reasoning.

In the example of reasoning above, I didn't explain how we know that the reasoning actually does show that the argument is valid. Instead, I just relied on "intuition": it seems obvious that this reasoning works. But what makes it seem obvious?

One of the purposes of logic is to give an answer to that question. We will give a precise definition of a **proof** in the language of sentential logic that corresponds to the intuitive idea of a valid argument with reasoning added to show that it is valid.

Just as we defined the sentences of sentential logic exactly with a list of basic vocabulary and a set of rules, we'll also define a proof using a precise set of rules. These rules will be easier to understand if we start by looking at a proof.

A Proof

Here is a **proof of the sequent** $P \rightarrow Q, Q \rightarrow R \vdash P \rightarrow R$:

1	(1)	$P \rightarrow Q$	A
2	(2)	$Q \rightarrow R$	A
3	(3)	P	A
1,3	(4)	Q	1,3 $\rightarrow$E
1,2,3	(5)	R	2,4 $\rightarrow$E
1,2	(6)	$P \rightarrow R$	5 $\rightarrow$I (3)

Notice that this proof is made up of **lines**. Each line has a certain structure. It begins with one or more numbers, separated by commas. After that, there is a number in parentheses. Next comes a wff. Finally, at the right side of the line, there is a mysterious and very short string. Each of these parts has a name:

$$\overbrace{1,2,3}^{Assumption\ set} \quad \overbrace{(5)}^{Line\ no.} \quad \overbrace{R}^{Sentence} \quad \overbrace{2,4 \rightarrow \mathrm{E}}^{Annotation}$$

Each of these different parts has a function:

Assumption set: The line numbers of all the lines on which this line depends.

Line number: Which line this is (lines are numbered in order, starting with one).

Sentence: The contents of this line.

Annotation: The rule that allows us to add the sentence, and the number(s) of the line(s) we applied that rule to.

Whenever we add a line to a proof, we add a **sentence** by applying a **rule** to **previous lines**. The **assumption set** we put at the front of the line is determined by the assumption set(s) of the line(s) we used and the particular rule we used.

Primitive Rules of Proof

There are ten basic rules of proof which we call **primitive rules** because they are the fundamental rules in terms of which proofs are defined. (Later on, we'll see how we can add more rules to the system.) These include:

1. The **assumption** rule (the simplest of them all)

2. For each of the connectives $\&$, $\vee$, $\rightarrow$, $\leftrightarrow$, two rules: an **intro** rule and an **elim** rule.

3. The rule called **reductio ad absurdum**

Assumption

assumption	Assume any sentence

This is the simplest rule of all, and it is also the rule with which a proof usually starts. You may assume any sentence you like, no restrictions. This can correspond to taking something as a premise in an argument, but there are also cases where you will want to make assumptions for other purposes.

When you add a line using **assumption**, you give it **an assumption set containing just its own line number**. An assumption doesn't depend on anything else, but other things may depend on it.

When you add a line using **assumption**, you add an **annotation** consisting simply of the letter **A**.

Ampersand-Intro (&I)

ampersand-intro	Given two sentences (at lines *m* and *n*, conclude a conjunction of them.

What this rule permits you to do is **add a conjunction**, that is, add a sentence with & as its main connective. The two lines can be any previous lines of the proof.

.	.	.	.
1,2,3	(*m*)	ϕ	
.	.	.	
2,4	(*n*)	ψ	
1,2,3,4	(*k*)	$(\phi \& \psi)$	m, n&I

The assumption set for the line we add includes all the numbers that are in either of the lines we used. Notice that we include each number only once. (By the way, the numbers in this example are just made up.)

A Closer Look at Annotations

The annotation for &I has certain form:

.	.	.	.	
1,2,3	(*m*)	ϕ		
.	.	.		
2,4	(*n*)	ψ		
1,2,3,4	(*k*)	$(\phi\&\psi)$	m,n	&I
			line numbers	**Rule name**

The annotation begins with **the numbers (or number) of the lines (or line) to which the rule was applied**, separated by a comma (if there is more than one number). Following that is **the name of the rule**. We will see this same structure for the annotations in several rules: &E, $\vee$I, $\vee$E, $\rightarrow$E, $\leftrightarrow$I, and $\leftrightarrow$E. Some of these rules apply to two lines at a time: &I, $\vee$E, $\rightarrow$E, $\leftrightarrow$I, $\leftrightarrow$E. The rest apply to only one line: &E, $\vee$I.

A closer Look at Assumption Sets

Every rule specifies how to determine the **assumption set** for the line it adds, based on the assumptions set(s) of the line(s) it is applied to. **Getting this part right is critical**. In the case of &E, the assumption set is *the union of the assumption sets* of the lines to which the rule is applied. ‘Union’ is a term from set theory: the union of two sets is the set of all the things that are in either of them. In other words, the assumption set here contains **all the numbers that are in either of the assumption sets** of the lines that the rule applies to:

.	.	.	.
1,2,3	(*m*)	ϕ	
.	.	.	
2,4	(*n*)	ψ	
1,2,3,4	(*k*)	$(\phi \& \psi)$	m,n &I

Ampersand-Elim (&E)

ampersand-elim	Given a sentence that is a conjunction (at line *m*), conclude either conjunct.

What this rule permits you to do is **get either part of a conjunction**, that is, add either the left or the right conjunct of a conjuction as a line.

.	.	.	.
2,4	(*m*)	$(\phi \& \psi)$	
.	.	.	
2,4	(*k*)	ϕ	m&E

We could equally well have added ψ.

The assumption set here is the same as the assumption set in front of the line containing the conjunction we used.

Wedge-Intro ($\vee$I)

wedge-intro	Given a sentence (at line *m*), conclude any sentence having it as a disjunct.

What this rule permits you to do is **take any line you like and add a line that contains the sentence you get by combining that line with any other sentence whatsoever, using the wedge**. The sentence you combine it with can be anything: it doesn't have to occur already in the proof as a line.

.	.	.	.
1,2,3	(*m*)	ϕ	
.	.	.	
1,2,3	(*k*)	$(\phi \vee \psi)$	$m\vee$I

We could equally well have added $(\psi \vee \psi)$. The sentence ψ can be anything

The assumption set for the line we add is the same as the assumption set for the line we used, that is, *m*.

Wedge-Elim ($\vee$E)

wedge-elim Given a sentence (at line *m*) that is a disjunction and another sentence (at line *n*) that is a denial of one of its disjuncts, conclude the other disjunct.

What this rule permits you to do is **get one of the disjuncts of a disjunction**. For the definition of a **denial**, see p. 7 of the text. Briefly, if there are two sentences, one of which is the other sentence with a $\sim$ in front of it, then the two sentences are denials of each other.

.	.	.	.
1,2,3	(*m*)	$(\phi \vee \psi)$	
.	.	.	
2,4	(*n*)	$\sim \phi$	
.	.	.	
1,2,3,4	(*k*)	ψ	$m, n \vee$E

The assumption set here includes all the numbers on either of the lines *m* or *n*.

Arrow-Elim ($\rightarrow$E)

arrow-elim	Given a conditional sentence (at line *m*) and another sentence that is its antecedent (at line *n*), conclude the consequent of the conditional.

What this rule permits you to do is **get the consequent (right side) of a conditional**. To apply it, you need the conditional on one line and its **antecedent** (left side) on the other. **NOTE: there is no rule that lets you get the antecedent from a conditional and its consequent**.

.	.	.	.
1,2,3	(*m*)	$(\phi \rightarrow \psi)$	
.	.	.	
2,4	(*n*)	ϕ	
.	.	.	
1,2,3,4	(*k*)	ψ	$m, n \rightarrow$E

The assumption set here includes all the numbers on either of the lines *m* or *n*.

Proofs

Before we go on to $\rightarrow$I and RAA, which introduce a new feature of rules, let's pause to consider a proof that uses many of the rules we've introduced. We can start by defining **PROOF** precisely (see p. 17 of the text):

> A **PROOF** is a sequence of lines containing sentences. Each sentence is either an assumption or the result of applying a rule of proof to earlier sentences in the sequence.

The meaning of this should now be clear. Notice that it adds one important requirement: **every proof starts with one or more assumptions**. (An assumption, in a proof, is a line that was introduced by the rule **assumption**).

Proofs for Arguments

> A **PROOF FOR A GIVEN ARGUMENT** is a proof whose last sentence is the argument's conclusion depending on nothing other than the argument's premises.

To rephrase this, here is how we construct a proof for a given argument:

1. Begin by **adding all of the premises of the argument**, each on one line, using the rule **assumption**.

2. Add more lines to the proof, using the rules, until you reach **a line containing the conclusion** of the argument.

3. If the **assumption set** of the line containing the conclusion includes **only the numbers of lines containing premises of the argument**, then then entire proof is a proof of the given argument.

Working through an example will help. Here is an argument (sequent):

$$P\& \sim Q, (P \vee S) \rightarrow (Q \vee R) \vdash P\&R$$

We start by assuming each of the premises:

1	(1)	$P\& \sim Q$	A
2	(2)	$(P \vee S) \rightarrow (Q \vee R)$	A

Our goal in constructing a proof is then to add lines in accordance with the rules until we arrive at a line that meets two criteria:

1. Its sentence is the conclusion, $P\&R$

2. Its assumption set includes only the assumption numbers of the premises (1 and 2).

What do we do next? We will disucss strategies for completing proofs a little later, but for the present let's concentrate on what a proof is. Here are two steps to get our proof started:

1	(1)	$P\& \sim Q$	A
2	(2)	$(P \vee S) \rightarrow (Q \vee R)$	A
1	(3)	P	1 &E
1	(4)	$P \vee S$	3 $\vee$I

The first step (line 3) was to add P by applying &E to line 1. Next, we used $\vee$I to add $P \vee S$ as line 4. In each case, the assumption set includes only 1 (why?).

1	(1)	$P \& \sim Q$	A
2	(2)	$(P \vee S) \rightarrow (Q \vee R)$	A
1	(3)	P	1 &E
1	(4)	$P \vee S$	3 $\vee$I
1,2	(5)	$Q \vee R$	2,4 $\rightarrow$E
1	(6)	$\sim Q$	1 &E

Next, we added line 5 using $\rightarrow$E. Notice that the assumption set includes the assumption sets from lines 2 and 4. We **don't** put down 4 as part of the assumption set (line 4 is not an assumption); instead, we include the **assumption set of** line 4.

Line 6 is exactly like line 3 except that here we added the right conjunct of line 1, not the left one.

1	(1)	$P\& \sim Q$	A
2	(2)	$(P \vee S) \rightarrow (Q \vee R)$	A
1	(3)	P	1 &E
1	(4)	$P \vee S$	3 $\vee$I
1,2	(5)	$Q \vee R$	2,4 $\rightarrow$E
1	(6)	$\sim Q$	1 &E
1,2	(7)	R	5,6 $\vee$E
1,2	(8)	$P\&R$	3,7 &I

LIne 7 applies $\vee$E to lines 5 and 6. Notice that the sentence on line 6, $\sim Q$, is the denial of one of the disjuncts of line 5.

Line 8 applies &I to lines 3 and 7. Just to make a point clear, notice that this is the second time we applied a rule to line 3. You can apply rules to a line or lines as often as you need to, in a proof: they don't get used up once you've used them.

Line 8 is the conclusion of the argument. Its assumption set only includes the numbers of the premises of the argument (1 and 2). Therefore, this is a proof of the sequent in question.

Rules for Biconditionals: $\leftrightarrow$I and $\leftrightarrow$E

Double-Arrow-Elim ($\leftrightarrow$E)

double-arrow-elim	Given a biconditional sentence $(\phi \leftrightarrow \psi)$ (at line *m*), conclude either $(\phi \rightarrow \psi)$ or $(\psi \rightarrow \phi)$.

What this rule permits you to do is **get a conditional** having one of the constitutents of a biconditional as antecedent and the other as consequent.

.	.	.	.
7,8	(*m*)	$(\phi \leftrightarrow \psi)$	
.	.	.	
7,8	(*k*)	$(\phi \rightarrow \psi)$	m $\leftrightarrow$E

We could equally well have added $\psi \rightarrow \phi$.

The assumption set here is the same as the assumption set in front of the line containing the biconditional we used.

Double-Arrow-Intro (↔I)

double-arrow-intro	Given two conditionals having the forms $(\phi \rightarrow \psi)$ and $(\psi \rightarrow \phi)$ (at lines *m* and *n*), conclude a biconditional with ϕ on one side and ψ on the other.

What this rule permits you to do is **get a conditional** having one of the constitutents of a biconditional as antecedent and the other as consequent.

5	(*m*)	$(\phi \rightarrow \psi)$	
.	.	.	
7	(*n*)	$(\psi \rightarrow \phi)$	
.	.	.	
5,7	(*k*)	$(\phi \leftrightarrow \psi)$	m, n ↔I

We could equally well have added $(\psi \leftrightarrow \phi)$.

The assumption set includes all the numbers in the assumption sets for the two conditionals used.

Assumption Rules: $\rightarrow$I and RAA

There are two remaining rules, $\rightarrow$I and RAA. These two rules introduce a new feature: each of them **discharges an assumption**. The critical point about these rules is that **an assumption number is dropped from the assumption set** on the line added by either of these rules.

The rule **assumption** . Although we have used it up to now for assuming the premises of an argument, the rule itself is broader in scope: you can assume anything you like at any time. Of course, if you use that assumption to get other lines, then its assumption number will be carried through to those lines. Consequently, you cannot get a proof of a sequent just by adding some extra premises to get the conclusion, since you'll have numbers in the assumption set that don't correspond to premises—.

—except under certain circumstances. One way to argue for a conditional conclusion is to **assume its antecedent** and then see if you can deduce its consequent. If you can, then what that shows is this: **if** you assume the antecedent, **then** you get the consequent. In other words, you have proved the conditional *If (the antecedent) then (the consequent)*.

Another strategy in argument is to assume **the negation of what you want to prove** and then try to get an inconsistent pair of sentences using that assumption. If you can, then what you've shown is this: *If (the assumption) were true, then something impossible would follow*. Since impossible things can't happen, the assumption must be false, and we can conclude its negation.

The rules $\rightarrow$I and RAA represent these lines of reasoning formally.

Arrow-Intro (→I)

arrow-intro Given an assumption (at line *m*) and a sentence (at line *n*), conclude the conditional having the assumption as its antecedent and the sentence as its consequent.

What this rule permits you to do is **get a conditional if you have already got its consequent and have assumed its antecedent**. The antecedent **must be an assumption**: it cannot be a line obtained using any other rule.

The assumption set here includes all the numbers on *n* **except for the line number of the assumption**. This last part of the rule is critical:

m	(*m*)	ϕ	
.	.	.	
1,3,*m*	(*n*)	ψ	
$\uparrow\uparrow\uparrow$	.	.	
1,3	(*k*)	$\phi \rightarrow \psi$	n →I (*m*)
$\uparrow\uparrow$			

Getting the annotation and the assumption set right here is crucial. Notice that we only cite one line, the line *n* on which the **consequent** of the conditional is found. However, we also cite the line number of the **antecedent** after the rule, and in parentheses.

An example will show how this works.

Sequent: $S \rightarrow P, P \rightarrow Q \vdash S \rightarrow Q$

Proof:

1	(1)	$S \rightarrow P$	A
2	(2)	$P \rightarrow Q$	A
3	(3)	S	A
1,3	(4)	P	1,3 $\rightarrow$E
1,2,3	(5)	Q	2,4 $\rightarrow$E
1,2	(6)	$S \rightarrow Q$	5 $\rightarrow$I (3)

Notice three things: (1) line 3 is an assumption; (2) the line number of line 3 does **not** appear in the assumption set of line 6; (3) in the annotation for line 6, the line number of line 3 is placed **after the name of the rule, in parentheses**.

Discharging assumptions

When we use $\rightarrow$I in this way to get a line that does not contain the number of an assumption in its assumption set, we say that we have **discharged** that assumption. This terminology is just a little misleading, since it implies that we can only do this one time with a given assumption. In fact, there's no such limitation. Here is an example:

Sequent:

$P \rightarrow (Q\&S), P \rightarrow R, R \rightarrow T, S\vee \sim Q, \vdash (P \rightarrow S)\&(P \rightarrow T)$

1	(1)	$P \rightarrow Q$	A
2	(2)	$P \rightarrow R$	A
3	(3)	$R \rightarrow T$	A
4	(4)	$S\vee \sim Q$	A
5	(5)	P	A
1,5	(6)	Q	1,5 $\rightarrow$E
2,5	(7)	R	2,5 $\rightarrow$E
1,4,5	(8)	S	4,6 $\vee$E
1,4	(9)	$P \rightarrow S$	8 $\rightarrow$I (5)
2,3,5	(10)	T	3,7 $\rightarrow$E
2,3	(11)	$P \rightarrow T$	10 $\rightarrow$I (5)
1,2,3,4	(12)	$(P \rightarrow S)\&(P \rightarrow T)$	9,11 &I

Reductio ad Absurdum (RAA)

Reductio ad absurdum	Given an assumption (at line *k*) and two sentences (at lines *m* and *n*) one of which is a negation of the other, conclude the negation of the assumption.

What this rule lets you do is conclude **the negation of an assumption** if you are able to get a **pair** of lines, one of which is the negation of the other.

Like $\rightarrow$I, this rule discharges an assumption: the assumption set for the line added includes all the numbers in the assumption sets of the pair of lines, **except for** the line number of the assumption.

.	.	.	.
k	(*k*)	ϕ	
.	.	.	
1,2,*k*	(*m*)	ψ	
.	.	.	
2,3	(*n*)	$\sim\psi$	
.	.	.	
1,2,3	(*x*)	$\sim\phi$	*m*,*n* RAA (*k*)

In this example, we've shown line *m* as depending on assumption *k* but line *n* not depending on it. There is actually no requirement that either of the lines containing the contradictory pair contain the assumption number of the assumption in its assumption set: both may, or one may, or neither may.

RAA corresponds to a familiar kind of argument: we prove that something is **not** so by assuming that it is and deducing an impossibility from that. The particular form of impossibility is a pair consisting of a sentence and its negation. Such a pair is often called a **contradiction**, and the pair of sentences are said to be **inconsistent**.

As with $\rightarrow$I, it is critical to get the assumption set right when using **RAA**.

Here's an example.

Sequent:

$P \rightarrow Q, Q \rightarrow \sim R, R \vdash \sim P$

1	(1)	$P \rightarrow Q$	A
2	(2)	$Q \rightarrow \sim R$	A
3	(3)	R	A
4	(4)	P	A
1,4	(5)	Q	1,4 $\rightarrow$E
1,2,4	(6)	$\sim R$	2,5 $\rightarrow$E
1,2,3	(8)	$\sim P$	3,6 RAA (4)