Handout

Propositional Logic

Isaac Wilhelm

Summary

The material in this handout is taken from an introductory logic textbook—that I am currently writing—on the relationship between (i) propositional and first-order logic, and (ii) social justice. Since the textbook is written for freshman undergraduates enrolled in their very first logic course, it should be relatively familiar to you. So if this material feels pretty basic and elementary, as you start reading, feel free to skim it. Just be sure you know the following:

- the symbols which the language of propositional logic contains,
- what counts as a sentence in propositional logic, and
- how to construct truth tables for those sentences (and more generally, the account of truth in propositional logic).

This material is concentrated in p. 44 (Section 3.1.8), pp. 47-48 (the start of Section 3.4), and pp. 54-73 (Sections 4.1 and 4.2). So make sure that you are familiar with the content of those pages.

Chapter 3

The Language of Propositional Logic

In this chapter, I introduce the formal language of propositional logic. This language can be used to represent natural language sentences and natural language arguments.

3.1   The Formal Language

We will now study the language of propositional logic; call it '$\mathcal{P}$'. This language has just a few different kinds of symbols. But those symbols can be used to formulate an extremely large range of different kinds of sentences.

To give you a feel for $\mathcal{P}$, and to explain why it is useful, I will connect it to a language with which you are more familiar. In particular, I will show how various symbols in $\mathcal{P}$ can be used to represent English sentences. So to start, I will introduce $\mathcal{P}$ slowly, one symbol—or type of symbol—at a time. After

that, I will give a quick and complete characterization of all the symbols of $\mathcal{P}$. Finally, I will say what it takes for a sequence of symbols in $\mathcal{P}$ to be a grammatical, well-formulated sentence.

An important aside: as will become clear, the very same expressions in $\mathcal{P}$—the very same strings of symbols—can be used to represent many, many different sentences of English. This is a pretty common feature of languages in general. For instance, many different English sentences can be used to represent—or translate—a German sentence like "Es regnet": this sentence can be represented by "It is raining" or "It rains", for instance. So when working with $\mathcal{P}$, keep in mind that one and the same expression can represent, or be represented by, multiple English sentences.

$\mathcal{P}$ will probably seem quite different from other languages which you might know. It is, in fact, pretty formal and technical. But do not be fooled: $\mathcal{P}$ is, in many respects, a language like any other. It is like English, or Navajo, or Spanish. It is like the language of 0s and 1s that runs your computer. And like any other language, there are two main ways to learn $\mathcal{P}$: (i) translate it into a language that you already know, like English, or (ii) just work with it a bunch. Do (i) and (ii) enough, and you will become fluent in $\mathcal{P}$.

3.1.1 Sentence Letters

The basic symbols of $\mathcal{P}$ are letters: 'p', 'q', 'r', 's', and many more. Call these symbols 'sentence letters', since they are used to represent full English sentences. Here are some examples.

Example

The symbol 'p' can be used to represent the English sentence "The Civil

Rights Act outlaws racial discrimination". In other words, we can use $\mathcal{P}$ to translate "The Civil Rights Act outlaws racial discrimination" as 'p'.

Example

The symbol 'p' can be used to represent the English sentence "The Civil Rights Act passed in 1964". In other words, we can use $\mathcal{P}$ to translate "The Civil Rights Act passed in 1964" as 'p'.

Example

The symbol 'p' can be used to represent the English sentence "The Civil Rights Act passed in 1931", and the symbol 'q' can be used to represent the English sentence "The Civil Rights Act passed in 1964". In other words, we can use $\mathcal{P}$ to translate "The Civil Rights Act passed in 1931" as 'p' and "The Civil Rights Act passed in 1964" as 'q'.

An important point about the last example: because the two sentences of English—namely, "The Civil Rights Act passed in 1931" and "The Civil Rights Act passed in 1964"—mean different things, the sentence letters used to represent them should be different as well. It would be bad to translate "The Civil Rights Act passed in 1931" as 'p', and then to translate "The Civil Rights Act passed in 1964" as 'p' too. For then the very same symbol—namely, 'p'—would have multiple meanings. All else equal, it is better for expressions of a language to mean just one thing: that cuts down on ambiguity and confusion. So here is a rule to remember: if two English sentences have different meanings, then use different sentence letters to represent those sentences.

$\mathcal{P}$ contains infinitely many sentence letters: 'p', 'q', 'r', 's', 't', and many, many more. In all of the examples to come, I will never use more than five

sentence letters. But it is worth pointing out that $\mathcal{P}$ contains many more sentence letters than just five.[1]

3.1.2 Parentheses

$\mathcal{P}$ includes the parentheses '(' and ')'. Parentheses are not used to translate English words. Instead, parentheses are used to organize the other symbols of $\mathcal{P}$, in order to be completely precise about exactly what the sentence in question says. In the later subsections, it will become clear how parentheses are used, and why their organizational role is so important.

3.1.3 Negation

$\mathcal{P}$ includes a negation symbol: '$\neg$'. This symbol represents the English word 'not'. Think of negation as the formal translation of this English-language construction: "It is not the case that ...".

Here is an example of how '$\neg$' can be used to represent English sentences.

Example

The expression '$\neg p$' can be used to represent the English sentence "The Civil Rights Act was not passed in 1931". In other words, we can use $\mathcal{P}$ to translate "The Civil Rights Act was not passed in 1931" as '$\neg p$'.

Note that in this example, the translation of "The Civil Rights Act was not passed in 1931" can be divided into two steps. First, we used 'p' to represent

---

[1] If you ever need to use more than five sentence letters, when translating something, just appeal to subscripts. For instance, instead of just 'p', use 'p_1', 'p_2', 'p_3', and so on.

"The Civil Rights Act passed in 1931". Second, we negated 'p', in order to represent the negation of that English sentence. As a result, we used '$\neg p$' to represent "The Civil Rights Act was not passed in 1931".

Here is another example.

Example

The expression '$\neg\neg p$' can be used to represent the English sentence "It is not the case that it is not the case that the Civil Rights Act passed in 1931". In other words, we can use $\mathcal{P}$ to translate "It is not the case that it is not the case that the Civil Rights Act passed in 1931" as '$\neg\neg p$'.

This translation can be divided into three steps. First, we used 'p' to represent "The Civil Rights Act passed in 1931". Second, we negated 'p' once, in order to represent the negation of that English sentence. As a result, we implicitly took '$\neg p$' to represent "It is not the case that the Civil Rights Act passed in 1931". Third, we negated '$\neg p$', in order to represent the negation of the English sentence "It is not the case that the the Civil Rights Act passed in 1931". As a result, we used '$\neg\neg p$' to represent "It is not the case that it is not the case that the the Civil Rights Act passed in 1931".

Think of the English sentence, in this example, as a more complicated way of saying "The Civil Rights Act was not *not* passed in 1931". Given that, we can think of "The Civil Rights Act was not *not* passed in 1931" as represented by '$\neg\neg p$'.

3.1.4 Conjunction

$\mathcal{P}$ includes a conjunction symbol: '$\wedge$'. This symbol represents the English word 'and'. Think of conjunction as the formal translation of this linguistic construction: "...and ...".

Here is an example of how '$\wedge$' can be used to represent English sentences.

Example

The expression '$p \wedge q$' can be used to represent the English sentence "The Civil Rights Act passed in 1931 and Kimberlé Crenshaw writes about intersectionality".

In this translation, the 'p' stands for "The Civil Rights Act passed in 1931" and the 'q' stands for "Kimberlé Crenshaw writes about intersectionality". The '$\wedge$' is used to translate the 'and'.

This example illustrates why different English sentences should not be represented with the same sentence letter. If we did that here, the result would be, say, '$p \wedge p$', which clearly is not what the English sentence means. The sentence '$p \wedge p$' means something more like "The Civil Rights Act passed in 1931 and the Civil Rights Act passed in 1931".

Here are two more examples.

Example

The expression '$p \wedge q$' can be used to represent the English sentence "The Civil Rights Act passed in 1931 and the Civil Rights Act helps people".

Example

The expression '$p \land \neg q$' can be used to represent the English sentence "The Civil Rights Act passed in 1931 and not 1964".

In the second of these examples, 'p' represents "The Civil Rights Act passed in 1931" and 'q' represents "The Civil Rights Act passed in 1964". And so '$\neg q$' represents "The Civil Rights Act was not passed in 1964", which means the same thing as the "not 1964" clause in the example.

An alternative—but worse—translation would be this: '$p \land q$', where this time, 'q' represents "The Civil Rights Act was not passed in 1964". This is a suboptimal translation of the original English sentence, because it misses out on some of that sentence's *logical structure*. In particular, it misses out on the structure captured by the '$\neg$'. Whenever it is possible to translate an English sentence in a way that uses '$\neg$', '$\land$', or some other logical symbol of $\mathcal{P}$, do it. For the resulting translation gives you a more precise representation of the logical features of the original English sentence.

Here is a final example.

Example
The expression '$\neg(p \land q)$' can be used to represent the English sentence "It is not the case that the Civil Rights Act passed in 1931 and in 1964".

Note that this example includes some parentheses. Those parentheses are extremely important: they play a crucial role in correctly capturing the meaning of the English sentence "It is not the case that the Civil Rights Act passed in 1931 and in 1964". To see why, just observe what happens when we remove them. The resulting expression is '$\neg p \land q$'. But '$\neg p \land q$' does not mean the same thing as '$\neg(p \land q)$'. In particular, '$\neg p \land q$' does not represent the English

sentence "It is not the case that the Civil Rights Act passed in 1931 and in 1964". Rather, '$\neg p \wedge q$' represents the English sentence "It is not the case that the Civil Rights Act passed in 1931, and the Civil Rights Act passed in 1964".

It is worth driving this subtle point home. So just to be clear, here are two sentences in $\mathcal{P}$ which (i) are distinguished only by how they use parentheses, and (ii) represent different sorts of English sentences.

1. '$\neg(p \wedge q)$': this represents sentences like the following.

   - "It is not the case that the Civil Rights Act passed in 1931 and in 1964".

   - "It is not the case that the Civil Rights Act passed both in 1931 and in 1964".

   - "The Civil Rights Act was not passed in both the year 1931 and also the year 1964".

2. '$\neg p \wedge q$': this represents sentences like the following.

   - "It is not the case that the Civil Rights Act passed in 1931, and the Civil Rights Act passed in 1964".

   - "It is not the case that the Civil Rights Act passed in 1931, but the Civil Rights Act passed in 1964".

   - "It is not the case that the Civil Rights Act passed in 1931, but also, by the way, the Civil Rights Act passed in 1964".

In short, '$\neg(p \wedge q)$' says that the Civil Rights Act was not passed in *both* years, while '$\neg p \wedge q$' says that the Civil Rights Act was *not* passed in the first year (1931) but *was* passed in the second year (1964).

3.1.5 Disjunction

$\mathcal{P}$ includes a disjunction symbol: '$\vee$'. This symbol represents the English word 'or'. Think of disjunction as the formal translation of this linguistic construction: "...or ...".

Here are some examples of how '$\vee$' can be used to represent English sentences.

> Example
>
> The expression '$p \vee q$' can be used to represent the English sentence "Frances Harper lived in Philadelphia or Cool Papa Bell was fast".

> Example
>
> The expression '$p \vee \neg q$' can be used to represent the English sentence "Frances Harper lived in Philadelphia or Cool Papa Bell was not fast".

> Example
>
> The expression '$p \vee (q \wedge r)$' can be used to represent the English sentence "Either the Civil Rights Act passed in 1964, or the Civil Rights Act passed in 1931 and Cool Papa Bell was fast".

In the last example above, 'p' represents "The Civil Rights Act passed in 1964", 'q' represents "The Civil Rights Act passed in 1931", and 'r' represents "Cool Papa Bell was fast".

Here is a final example.

Example

The expression '$p \lor (\neg q \land r)$' can be used to represent "The Civil Rights Act passed in 1964, or Kimberlé Crenshaw does not write about intersectionality and Jan Morris wrote *Conundrum*".

In this example, 'p' represents "The Civil Rights Act passed in 1964". The sentence letter 'q' represents "Kimberlé Crenshaw writes about intersectionality"; so '$\neg q$' represents "Kimberlé Crenshaw does not write about intersectionality". Finally, 'r' represents "Jan Morris wrote *Conundrum*"; so '$\neg q \land r$' represents "Kimberlé Crenshaw does not write about intersectionality and Jan Morris wrote *Conundrum*".

3.1.6   Conditional

$\mathcal{P}$ includes a conditional symbol: '$\rightarrow$'. This symbol represents the following English-language construction: "If ... then ...".

Here is an example of how '$\rightarrow$' can be used to represent English sentences.

Example

The expression '$p \rightarrow q$' can be used to represent "If Toni Morrison wrote *Beloved*, then Toni Morrison wrote novels".

Example

The expression '$\neg q \rightarrow \neg p$' can be used to represent "If Toni Morrison did not write novels, then Toni Morrison did not write *Beloved*".

In both of the examples above, 'p' represents "Toni Morrison wrote *Beloved*", and 'q' represents "Toni Morrison wrote novels".

Here is a final example.

Example

The expression '$(p \land \neg q) \to r$' can be used to represent "If Ta-Nehisi Coates speaks at my graduation and I do not attend, then I will be sad".

In this example, 'p' represents "Ta-Nehisi Coates speaks at my graduation". The sentence letter 'q' represents "I attend"; so '$\neg q$' represents "I do not attend". And so '$p \land \neg q$' represents "Ta-Nehisi Coates speaks at my graduation and I do not attend". Finally, 'r' represents "I will be sad".

3.1.7 Biconditional

$\mathcal{P}$ includes a biconditional symbol: '$\leftrightarrow$'. This symbol represents the following English-language construction: "...if and only if ...".

Here is an example of how '$\leftrightarrow$' can be used to represent English sentences.

Example

The expression '$p \leftrightarrow q$' can be used to represent "Kimberlé Crenshaw will speak at the convention if and only if bell hooks will".

In this example, 'p' represents "Kimberlé Crenshaw will speak at the convention" and 'q' represents "bell hooks will speak at the convention".

Here is a final example.

Example

The expression '$p \leftrightarrow \neg q$' can be used to represent "Kimberlé Crenshaw will speak at the convention if and only if Barack Obama will not".

In this example, 'p' represents "Kimberlé Crenshaw will speak at the convention". The sentence letter 'q' represents "Barack Obama will speak at the convention". As a result, '$\neg q$' represents "Barack Obama will not speak at the convention".

3.1.8 Summary

Here is a summary of all the symbols which $\mathcal{P}$ contains.

1. Sentence letters: 'p', 'q', 'r', 's', 't', and so on.[2]

2. Parentheses: '(' and ')'.

3. Negation: '$\neg$'.

4. Conjunction: '$\wedge$'.

5. Disjunction: '$\vee$'.

6. Conditional: '$\rightarrow$'.

7. Biconditional: '$\leftrightarrow$'.

The symbols '$\neg$', '$\wedge$', '$\vee$', '$\rightarrow$', and '$\leftrightarrow$' are called 'logical connectives'. This is because they *connect* to sentences of the *logical* language $\mathcal{P}$ to form other sentences of $\mathcal{P}$. Sometimes they are also called 'connectives', 'logical symbols', and 'logical constants'.

Here is the most important feature of these symbols: they always mean the same thing. Whereas 'p' can mean different things, depending on what it is used to translate, the meanings of '$\neg$', '$\wedge$', '$\vee$', '$\rightarrow$', and '$\leftrightarrow$', are constant.

---

[2]There are infinitely many of these.

3.2 Translation Tips

Throughout Section 3, I explained the meaning of each symbol '¬', '∧', '∨', '→', and '↔', using just one English word or phrase. For instance, I wrote that '¬' represents the English word 'not', and I wrote that '∧' represents the English word 'and'.

But actually, each of these symbols can be used to represent many more English words. Here is an incomplete list of the sorts of English words, phrases, expressions, and so on, which symbols in $\mathcal{P}$ can be used to represent.

1. '¬': 'not', 'it is not the case that...', 'neither...nor...', and any other word or expression which suggests that something is being denied.

2. '∧': 'and', 'both...and...', 'but', 'though..., ...', and any other word or expression which suggests that some things have been conjoined.

3. '∨': 'or', 'either...or...', 'neither...nor...', and any other word or expression which suggests that some things have been disjoined.

4. '→': 'if ... then ...', 'when ..., ...', 'since ..., ...', and any other word or expression which suggests that one thing follows from another thing.

5. '↔': '...if and only if ...', '...just in case ...', '...exactly when ...', and any other word or expression which suggests that (i) one thing follows from another thing, and (ii) that other thing follows from the first thing.

There is no fully precise, fully rigorous way to say exactly which English expressions can be represented by symbols of $\mathcal{P}$. That is a common feature of all languages: there is no fully precise, fully rigorous way to say exactly which Hindi expressions can be represented using the English word 'not', for instance. But as you work with $\mathcal{P}$, you will get a better and better sense for how its symbols relate to various English language constructions. You will develop an

intuition for how to translate English into $\mathcal{P}$.

3.3   Translating English Arguments

Just as $\mathcal{P}$ can be used to translate English sentences, $\mathcal{P}$ can also be used to translate English arguments. The method is pretty straightforward: in order to translate an English argument into $\mathcal{P}$, just translate each individual sentence of that argument into a sentence of $\mathcal{P}$. In other words, use $\mathcal{P}$ to translate each English sentence individually. The resulting sequence of sentences in $\mathcal{P}$ is an argument: it is the translation of the English argument in question.

For example, recall the Civil argument.

1. The Civil Rights Act outlaws racial discrimination.
2. If the Civil Rights Act outlaws racial discrimination, then the Civil Rights Act helps people.
3. The Civil Rights Act helps people.

Let 'p' represent "The Civil Rights Act outlaws racial discrimination" and let 'q' represent "The Civil Rights Act helps people". Then the translation of the Civil argument, into $\mathcal{P}$, is this.

1. p
2. $p \rightarrow q$
3. q

This translation proceeds sentence-by-sentence. The first line of the translated argument represents the first line of the original Civil argument. The second line of the translated argument represents the second line of the original Civil

argument. And the third line of the translated argument represents the third line of the original Civil argument.

3.4 Sentences of $\mathcal{P}$

In this section, I give a fully rigorous definition of what counts as a sentence of $\mathcal{P}$. Then I give a series of examples of that definition at work.

The definition of a sentence of $\mathcal{P}$ has roughly three parts. In the first and simplest part, the definition says what the basic, simplest sentences of $\mathcal{P}$ are. In the second and most complicated part, the definition says how to use simpler sentences—along with symbols like '$\neg$', '$\wedge$', and so on—to construct more complicated sentences. In the third part, the definition says that there are no other ways to get sentences in $\mathcal{P}$: the only sentences are the ones that can be built using the first part and the second part.

Here is the definition; call it the 'Recursive Definition' of a sentence of $\mathcal{P}$.[3]

Definition 3.1: Sentence of $\mathcal{P}$

The 'sentences' of $\mathcal{P}$ are defined as follows.

1. First part

   - Every sentence letter is a sentence.

2. Second part

   - If 'ϕ' is a sentence, then '$\neg\phi$' is a sentence.
   - If 'ϕ' and 'ψ' are sentences, then '$(\phi \wedge \psi)$' is a sentence.
   - If 'ϕ' and 'ψ' are sentences, then '$(\phi \vee \psi)$' is a sentence.

---

[3]This is a 'recursive' definition because it uses sentences to define sentences: in particular, it defines more complicated sentences in terms of simpler sentences which have already been defined.

- If 'ϕ' and 'ψ' are sentences, then '$(\phi \rightarrow \psi)$' is a sentence.

- If 'ϕ' and 'ψ' are sentences, then '$(\phi \leftrightarrow \psi)$' is a sentence.

3. <u>Third part</u>

- Nothing else is a sentence.

In general, the sentences of a language—whether that language is English, $\mathcal{P}$, or something else—are the grammatical, well-formulated, coherent expressions of that language. In any language, many expressions are not grammatical: for instance, "Red that coffee is and" is an English expression, since it contains only English words; yet it is not grammatical, since those words do not conform to the English rules for sentence formation. Just like English, $\mathcal{P}$ has rules for forming sentences. And just like English, the sentences of $\mathcal{P}$ are the expressions of $\mathcal{P}$ which conform to those rules.[4]

For a particularly simple example, note that 'p' is a sentence of $\mathcal{P}$. This follows from (i) the first part of the Recursive Definition, and (ii) the fact that 'p' is a sentence letter.

For a more complicated example, consider '$(p \wedge q)$'. This too is a sentence of $\mathcal{P}$. To see why, note that since both 'p' and 'q' are sentence letters, the first part of the Recursive Definition implies that both 'p' and 'q' are sentences. Now consider the *second* bullet point in the *second* part of the Recursive Definition: it says that if 'ϕ' and 'ψ' are sentences, then '$(\phi \wedge \psi)$' is a sentence. Substituting 'p' for 'ϕ' and 'q' for 'ψ', we get the following: if 'p' and 'q' are

---

[4]There is a notational subtlety in the Recursive Definition: the language $\mathcal{P}$ does not contain the symbols 'ϕ' and 'ψ'. Those Greek letters are neither symbols nor sentences of $\mathcal{P}$. Instead, in the Recursive Definition, those Greek letters are more like placeholders. They stand for any sentence of $\mathcal{P}$ at all; they are not, themselves, part of $\mathcal{P}$. Throughout this book, I use those Greek letters to talk about $\mathcal{P}$; I use them to *mention* symbols of $\mathcal{P}$, as it were. So just to be perfectly clear: the English letters 'p', 'q', and so on, are symbols of $\mathcal{P}$; the Greek letters 'ϕ', 'ψ', and so on, are not. For more on these sorts of issues, concerning how best to talk about the terms in languages, see (Sider, 2010).

sentences, then '$(p \wedge q)$' is a sentence. As I just showed, 'p' and 'q' are indeed sentences of $\mathcal{P}$. So it follows that '$(p \wedge q)$' is a sentence.

For an even more complicated example, consider '$((p \wedge q) \to r)$'. This is a sentence of $\mathcal{P}$ too. As I just showed, '$(p \wedge q)$' is a sentence. Since 'r' is a sentence letter, the first part of the Recursive Definition implies that 'r' is a sentence of $\mathcal{P}$. Now consider the *fourth* bullet point in the *second* part of the Recursive Definition: it says that if 'ϕ' and 'ψ' are sentences, then '$(\phi \to \psi)$' is a sentence. Substituting '$(p \wedge q)$' for 'ϕ' and 'r' for 'ψ', we get the following: if '$(p \wedge q)$' and 'r' are sentences, then '$((p \wedge q) \to r)$' is a sentence. As I just showed, '$(p \wedge q)$' and 'r' are indeed sentences. Therefore, '$((p \wedge q) \to r)$' is a sentence.

Plenty of expressions of $\mathcal{P}$—plenty of strings of symbols in $\mathcal{P}$, that is—are *not* sentences. Here is one: 'pq'. This is not a sentence because (i) it is not a sentence letter, and (ii) it cannot be built from sentences using the five bullet points in the second part of the Recursive Definition. Here is another: '$p \to$'. Again, this is not a sentence because (i) it is not a sentence letter, and (ii) it cannot be built from sentences using the five bullet points in the second part of the Recursive Definition. Here is yet another: '$p \to (\neg \vee r)$'. This is not a sentence for the same reason: it is not a sentence letter, and it cannot be built from sentences using those five bullet points.

3.5   Notation and Terminology

Before closing, it is worth introducing some standard notation and terminology. Here is one notational shorthand: following standard practice, I will often omit certain parentheses—typically, the outermost parentheses—of sentences

49

of $\mathcal{P}$. For instance, instead of writing '$(p \wedge q)$', I will write '$p \wedge q$'. And instead of writing '$((p \wedge q) \rightarrow r)$', I will write '$(p \wedge q) \rightarrow r$'. This is standard practice because it increases readability: after you work with $\mathcal{P}$ for a while, it becomes easier to parse expressions like '$\Big(\big((p \wedge q) \rightarrow r\big) \vee \big(s \leftrightarrow \neg(q \rightarrow p)\big)\Big)$' if the outermost parentheses are dropped, and thus, if that expression is rewritten as '$\big((p \wedge q) \rightarrow r\big) \vee \big(s \leftrightarrow \neg(q \rightarrow p)\big)$'. So in this book, I will count expressions like '$p \wedge q$' as sentences, even though—contrary to the second bullet point in the Recursive Definition—they do not contain outermost parentheses.

Now for some terminology. In a sentence of the form '$\phi \wedge \psi$', 'ϕ' and 'ψ' are called the 'conjuncts' of '$\phi \wedge \psi$'. For they are the basic parts of that conjunction. Likewise, in an English sentence such as "The Civil Rights Act outlaws racial discrimination and the Civil Rights Act helps people", the sentences on either side of 'and' are called 'conjuncts' of that English sentence too. "The Civil Rights Act outlaws racial discrimination" is a conjunct of that sentence, for instance, as is "the Civil Rights Act helps people".

Similarly, in a sentence of the form '$\phi \vee \psi$', 'ϕ' and 'ψ' are called 'disjuncts' of '$\phi \vee \psi$'. For they are the basic parts of that disjunction. Likewise, in an English sentence such as "The Civil Rights Act outlaws racial discrimination or the Civil Rights Act helps people", the sentences on either side of 'or' are called 'disjuncts' of that English sentence too. "The Civil Rights Act outlaws racial discrimination" is a disjunct of that sentence, for instance, as is "The Civil Rights Act helps people".

Finally, in a sentence of the form '$\phi \rightarrow \psi$', 'ϕ' is the 'antecedent' of '$\phi \rightarrow \psi$', and 'ψ' is the 'consequent' of '$\phi \rightarrow \psi$'. Likewise, in an English sentence such as "If the Civil Rights Act outlaws racial discrimination, then the Civil Rights Act helps people", "The Civil Rights Act outlaws racial discrimination" is the antecedent of that sentence and "The Civil Rights Act helps people" is the

consequent of that sentence.

3.6   Problems

Here are some problems based on the ideas presented in this chapter.

Problem 3.1. *Which of the following are sentences of* $\mathcal{P}$*?*

  1. p

  2. $(p \rightarrow q)$

  3. $(p \vee q)$

  4. $(p \rightarrow (q \vee (\neg r \leftrightarrow s\neg)))$

  5. $\neg\neg p \wedge q$

Problem 3.2. *Which of the following are sentences of* $\mathcal{P}$*?*

  1. $q \vee p)$

  2. q

  3. $p \rightarrow (\neg(q \leftrightarrow p)) \vee r$

  4. $\neg p$

  5. $\vee p$

Problem 3.3. *Translate the sentence below into* $\mathcal{P}$*.*

  "Frederick Douglass edited The North Star and Susan Stryker attended
  Berkeley."

Problem 3.4. *Translate the sentence below into* $\mathcal{P}$*.*

"Ibn al-Shatir reformed the Ptolemaic model of the Sun."

Problem 3.5. *Translate the sentence below into* $\mathcal{P}$*.*

"Mahatma Gandhi and Abdul Ghaffar Khan will not both attend."

Problem 3.6. *Translate the sentence below into* $\mathcal{P}$*.*

"Neither Martin Luther King Jr. nor Malcom X will be forgotten."

Problem 3.7. *Translate the sentence below into* $\mathcal{P}$*.*

"If you apologize, and if you never visit New Orleans again, and if you either start donating to BLM or help my friend paint their sign, then I might take you to the movies and buy you some popcorn."

Problem 3.8. *Translate the argument below into* $\mathcal{P}$*.*

1. *Jan Morris was not born in the U.S.*
2. *If Jan Morris was not born in the U.S., then Jan Morris was not born in New York.*
3. *Jan Morris was not born in New York.*

Problem 3.9. *Translate the argument below into* $\mathcal{P}$*.*

1. *The Civil Rights Act does not outlaw racial discrimination.*
2. *If the Civil Rights Act outlaws racial discrimination, then the Civil Rights Act helps people.*
3. *The Civil Rights Act does not help people.*

Problem 3.10. *Translate the argument below into $\mathcal{P}$.*

1. *If Huck leaves, then Aunt Polly cannot 'civilize' Huck.*

2. *If Aunt Polly cannot 'civilize' Huck, then Huck will probably help Jim.*

3. *If Huck leaves, then Huck will probably help Jim.*

Chapter 4

Truth and Validity in $\mathcal{P}$

In this chapter, I present the accounts of truth and validity in $\mathcal{P}$. The ultimate goal, recall, is to use $\mathcal{P}$ to provide a rigorous theory of what it takes for natural language arguments to be good. The account of truth in $\mathcal{P}$, and the account of validity in $\mathcal{P}$, are the ingredients needed to do exactly that.

So do not be discouraged by the fact that in this chapter, things get pretty formal. That formalism represents one of the main theories in all of logic: the theory of validity for arguments in $\mathcal{P}$. And learning it pays off: it will help us understand what makes certain English arguments good.

4.1 Truth

In this section, I give a complete account of truth in $\mathcal{P}$. That is, I say exactly what it takes for any given sentence in $\mathcal{P}$ to be true. The basic idea is as follows. Sentence letters in $\mathcal{P}$ do not have fixed truth values: they can be true, or false, depending on what they represent. But once the truth values of

the sentence letters are fixed, the truth values of all other sentences get fixed as well. In other words, the truth values of more complicated sentences—sentences, that is, which have been built out of sentence letters and logical connectives—are completely determined by the truth values of the simplest sentences: namely, sentence letters.

For convenience, we will use what are called 'truth tables' to study truth in $\mathcal{P}$. For any given truth table, and for each way of assigning truth values to certain sentences, the table specifies what the truth value of a certain complicated sentence must be. So truth tables provide a clear, concise specification of all the ways in which the truth values of more complicated sentences depend on the truth values of simpler sentences.

To start, consider sentence letters. Each sentence letter can be either true or false. The truth value of a sentence letter is not really determined by anything. In particular, the truth value of a sentence letter is not determined by the truth values of any simpler sentences which that sentence letter contains. The reason is simple: sentence letters do not contain simpler sentences. They are the simplest sentences of $\mathcal{P}$. Because of this, there are no truth tables for sentence letters.

For example, take the sentence letter 'p'. This sentence letter can be either true or false, depending on what we take it to represent. For instance, if we use 'p' to represent the English sentence "The Civil Rights Act passed in 1964", then 'p' is true: for that is, indeed, when the Civil Rights Act was passed. If we use 'p' to represent the English sentence "The Civil Rights Act passed in 1931", then 'p' is false: for the Civil Rights Act was not passed until much later.

Now for negation. For any sentence 'ϕ', '$\neg\phi$' is true if and only if 'ϕ' is false, and '$\neg\phi$' is false if and only if 'ϕ' is true. Here is the truth table.

ϕ	$\neg\phi$
T	F
F	T

Read this truth table as follows. The leftmost column lists all possible truth values of 'ϕ': it can be true or false. 'T' represents 'true', of course, and 'F' represents 'false'; 'T' and 'F' are what I have been calling 'truth values'. For each truth value of 'ϕ', the rightmost column says what the truth value of '$\neg\phi$' is. For example, the first row[1] says the following: if 'ϕ' is true, then '$\neg\phi$' is false. And the second row says the following: if 'ϕ' is false, then '$\neg\phi$' is true.

So for example, suppose 'ϕ' is the sentence 'p'. And suppose that 'p' is true. Then by the first row in the truth table, the sentence '$\neg p$' is false. Or suppose 'ϕ' is the sentence '$p \lor q$'. And suppose that this sentence is false; I will explain how disjunctions can be false below. Then by the second row in the truth table, '$\neg(p \lor q)$' is true.

The truth table for '$\neg$' is pretty intuitive. Roughly, it says that a sentence of the form "It is not the case that X" is true just in case 'X' is false. And that seems pretty plausible. For example, consider the English sentence "It is not the case that the Civil Rights Act passed in 1931". That sentence seems true. Why? Because the sentence "The Civil Rights Act passed in 1931" is false. The above truth table basically says the same thing, just for sentence of $\mathcal{P}$ rather than sentences of English.

Now for conjunction. For any sentences 'ϕ' and 'ψ', '$\phi \land \psi$' is true if and only if both 'ϕ' is true and 'ψ' is true. Otherwise, '$\phi \land \psi$' is false. Here is the

---

[1]For brevity, throughout this book, I take the first row of a truth table to be the first row which contains truth values. So the row which contains sentences—like the row which contains 'ϕ' and '$\neg\phi$' in the truth table for negation—will not count as the first row. In addition, other rows are numbered off in the natural way, starting from the first. For instance, in the truth table for negation, the row containing 'F' and then 'T' is the second row.

truth table.

ϕ	ψ	$\phi \wedge \psi$
T	T	T
T	F	F
F	T	F
F	F	F

Read this truth table as follows. Taken together, the two columns on the left list all possible *combinations* of assignments of truth values—call each of these a 'truth value assignment'—to 'ϕ' and 'ψ'. There are four:

1. 'ϕ' could be true and 'ψ' could be true,

2. 'ϕ' could be true and 'ψ' could be false,

3. 'ϕ' could be false and 'ψ' could be true, and

4. 'ϕ' could be false and 'ψ' could be false.

For each of those combinations, the rightmost column of the truth table says what the truth value of '$\phi \wedge \psi$' is. For example, the first row says the following: if 'ϕ' is true and 'ψ' is true, then '$\phi \wedge \psi$' is true. The second row says the following: if 'ϕ' is true and 'ψ' is false, then '$\phi \wedge \psi$' is false. The third row says the following: if 'ϕ' is false and 'ψ' is true, then '$\phi \wedge \psi$' is false. And the fourth row says the following: if 'ϕ' is false and 'ψ' is false, then '$\phi \wedge \psi$' is false.

So for example, suppose 'ϕ' is the sentence 'p' and 'ψ' is the sentence 'q'. Furthermore, suppose that 'p' is true and 'q' is false. Then by the second row, the sentence '$p \wedge q$' is false. Or suppose 'ϕ' is the sentence '$p \vee q$', and 'ψ' is the sentence 'r'. Suppose that '$p \vee q$' and 'r' are both true. Then by the first row, '$(p \vee q) \wedge r$' is true.

Truth tables are helpful and convenient. They compress all that information—all the information in the paragraph above—into a very small space. That is why books and classes on logic often cover truth tables.

The truth table for '∧' is pretty intuitive. Roughly, it says that a sentence of the form "X and Y" is true just in case 'X' is true and 'Y' is true. If either 'X' or 'Y' is false, then 'X and Y' is false. And that all seems pretty plausible. For example, consider the English sentence "The Civil Rights Act outlaws racial discrimination and the Civil Rights Act helps people". That sentence seems true. Why? Because "The Civil Rights Act outlaws racial discrimination" is true, and "The Civil Rights Act helps people" is true. As another example, consider the English sentence "The Civil Rights Act outlaws racial discrimination and the Civil Rights Act passed in 1931". That sentence seems false. Why? Because "The Civil Rights Act passed in 1931" is false. The above truth table basically says the same thing, just for sentence of $\mathcal{P}$ rather than sentences of English.

Note what these truth tables are doing. They say how the truth values of more complicated sentences are built out of the truth values of less complicated sentences. The truth table for '¬', for instance, illustrates how the truth value of the more complicated sentence '$\neg\phi$' is built out of the truth value of the less complicated sentence 'ϕ'. Similarly, the truth table for '∧' illustrates how the truth value of the more complicated sentence '$\phi \wedge \psi$' is built out of the truth values of less complicated sentences: the less complicated sentence 'ϕ', and the less complicated sentence 'ψ'. In other words, truth tables describe how the truth values of simpler sentences determine, or ground, or account for, or explain, the truth values of more complicated sentences. Truth flows from the simplest sentences to the ever more complex ones, like a river splitting into ever more distributaries.

Now for disjunction. For any sentences 'ϕ' and 'ψ', '$\phi \vee \psi$' is true if and only if either 'ϕ' is true or 'ψ' is true. Otherwise, '$\phi \vee \psi$' is false. Here is the truth table.

ϕ	ψ	$\phi \vee \psi$
T	T	T
T	F	T
F	T	T
F	F	F

As before, read this truth table like so: the two columns on the left list all possible combinations of truth value assignments to 'ϕ' and 'ψ', while the column on the right lists the corresponding truth values of the more complicated sentence '$\phi \vee \psi$'.

So for example, suppose 'ϕ' is the sentence 'p' and 'ψ' is the sentence 'q'. Furthermore, suppose that 'p' is false and 'q' is true. Then by the third row, the sentence '$p \vee q$' is true. Or suppose 'ϕ' is the sentence '$p \rightarrow q$', and 'ψ' is the sentence 'r'. Suppose that '$p \rightarrow q$' and 'r' are both false. Then by the fourth row, '$(p \rightarrow q) \vee r$' is false.

The truth table for '$\vee$' is pretty intuitive. Roughly, it says that sentences of the form "X or Y" are true just in case either 'X' is true or 'Y' is true. If both 'X' and 'Y' are false, then 'X and Y' is false. And that all seems pretty plausible. For example, consider the English sentence "The Civil Rights Act outlaws racial discrimination or the Civil Rights Act passed in 1931". That sentence seems true. Why? Because "The Civil Rights Act outlaws racial discrimination" is true. Of course, "The Civil Rights Act passed in 1931" is false. But the sentence "The Civil Rights Act outlaws racial discrimination or

the Civil Rights Act passed in 1931" is still true, because one of its disjuncts—namely, the disjunct "The Civil Rights Act outlaws racial discrimination"—is true.

Now for the conditional. For any sentences 'ϕ' and 'ψ', '$\phi \rightarrow \psi$' is true if and only if either 'ϕ' is false or 'ψ' is true. Otherwise, '$\phi \rightarrow \psi$' is false. Here is the truth table.

ϕ	ψ	$\phi \rightarrow \psi$
T	T	T
T	F	F
F	T	T
F	F	T

As before, read this truth table like so: the two columns on the left list all possible combinations of truth value assignments to 'ϕ' and 'ψ', while the column on the right lists the corresponding truth values of the more complicated sentence '$\phi \rightarrow \psi$'.

So for example, suppose 'ϕ' is the sentence 'p' and 'ψ' is the sentence 'q'. Furthermore, suppose that 'p' is true and 'q' is true. Then by the first row, the sentence '$p \rightarrow q$' is true. Or suppose 'ϕ' is the sentence '$\neg q$', and 'ψ' is the sentence '$r \wedge p$'. Suppose that '$\neg q$' is true while '$r \wedge p$' is false. Then by the second row, '$\neg q \rightarrow (r \wedge p)$' is false.

Of all the truth tables for all the logical connectives of $\mathcal{P}$, this one is probably the weirdest. There are several reasons why: here, I will discuss just one. The third row says that in order for a conditional to be true, it suffices for the antecedent of that conditional to be false. But that is not really how the English construction corresponding to the symbol '$\rightarrow$'—the construction "If ..., then ..."—seems to work. To see why, consider the sentence "If the

Civil Rights Act passed in 1931, then the Civil Rights Act outlaws racial discrimination". It is not clear whether, intuitively, that sentence is true or false: some people think it is false, some think it is true, and some think that it simply lacks a truth value. But according to the truth table above, the corresponding sentence of $\mathcal{P}$—the sentence '$p \rightarrow q$', where 'p' represents "The Civil Rights Act passed in 1931" and 'q' represents "The Civil Rights Act outlaws racial discrimination"—is true, since 'p' is false and 'q' is true. So the symbol '$\rightarrow$', in $\mathcal{P}$, is quite different from the corresponding English expression "If ..., then ...".

Many philosophers, linguists, and psychologists study this. In particular, many philosophers, linguistics, and psychologists explore different ways of setting up formal languages, so that (i) the resulting formal language has many of the nice properties that $\mathcal{P}$ has, but (ii) the resulting formal language has a symbol which does a better job than '$\rightarrow$' of representing the English construction "If ..., then ...". For lack of space, I will not investigate that here. But it is a fascinating subject to study.[2]

Now for the biconditional. For any sentences 'ϕ' and 'ψ', '$\phi \leftrightarrow \psi$' is true if and only if 'ϕ' and 'ψ' have the same truth value. In particular, here is the truth table for '$\leftrightarrow$'.

ϕ	ψ	$\phi \leftrightarrow \psi$
T	T	T
T	F	F
F	T	F
F	F	T

---

[2]For discussion, see (Bennett, 2003; Evans & Over, 2004).

As before, read this truth table like so: the two columns on the left list all possible combinations of truth value assignments to 'ϕ' and 'ψ', while the column on the right lists the corresponding truth values of the more complicated sentence '$\phi \leftrightarrow \psi$'.

So for example, suppose 'ϕ' is the sentence 'p' and 'ψ' is the sentence 'q'. Furthermore, suppose that 'p' is false and 'q' is false. Then by the fourth row, the sentence '$p \leftrightarrow q$' is true. Or suppose 'ϕ' is the sentence '$\neg q$', and 'ψ' is the sentence '$r \wedge p$'. Suppose that '$\neg q$' is true while '$r \wedge p$' is false. Then by the second row, '$\neg q \leftrightarrow (r \wedge p)$' is false.

The truth table for '$\leftrightarrow$' is perhaps more intuitive than the truth table for '$\rightarrow$', but it is still less intuitive than the other truth tables. Roughly, it says that a sentence of the form "X if and only if Y" is true just in case either (i) 'X' is true and 'Y' is true, or (ii) 'X' is false and 'Y' is false. So "X if and only if Y" is false whenever the truth values of 'X and Y' differ. And that all seems at least somewhat plausible. For example, consider the English sentence "The Civil Rights Act outlaws racial discrimination if and only if the Civil Rights Act passed in 1931". That sentence seems false. Why? Because "The Civil Rights Act outlaws racial discrimination" is true, but "the Civil Rights Act passed in 1931" is false. Or consider the sentence "The Civil Rights Act outlaws racial discrimination if and only if the Civil Rights Act helps people". Since both sides of the '...if and only if ...' construction are true, that sentence seems—at least to some extent—true as well.[3]

---

[3]Some people might disagree. To my ear at least, it is not totally intuitive that this sentence is true. Again, this is the sort of thing that many philosophers, linguistics, and psychologists study.
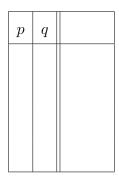
4.2 Truth for Sentences of $\mathcal{P}$: A Detailed Example

Now for the first big upshot of the above theory of truth: using the truth tables for the five logical connectives, one can construct truth tables for *any sentence of $\mathcal{P}$ whatsoever*. In this section, I explain how. I focus on a single example: '$p \vee \neg q$'. In later subsections, I work through a few more examples.

Consider the sentence '$p \vee \neg q$'. When is it true, and when is it false? Or a little more precisely: under what assignments of truth values to the sentence letters in '$p \vee \neg q$' does that sentence come out true, and under what assignments of truth values to the sentence letters in '$p \vee \neg q$' does that sentence come out false? The truth table for '$p \vee \neg q$' provides a complete answer. Let us see how to construct it.

The construction proceeds in roughly two parts. The first part is pretty simple. The second is quite complex.

The first part is as follows. To start, determine which sentence letters are in '$p \vee \neg q$'. There are two: the sentence letter 'p', and the sentence letter 'q'.

Now put those sentence letters into a truth table, like so.

p	q	

Then list out all the combinations of truth value assignments to those sentence letters. Note that for each sentence letter, there are two possible assignments of truth values to that sentence letter: it could be assigned 'T', and it could

be assigned 'F'. So for two sentence letters, there are four different truth value combinations: they could both be true, the first could be true while the second is false, the second could be false while the first is true, and they could both be false. So there are four combinations of assignments—or for short, there are four 'assignments'—of truth values to 'p' and 'q' taken together. List them in the truth table, like so.

p	q	
T	T	
T	F	
F	T	
F	F	

Now take the sentence in question—that is, '$p \vee \neg q$'—and put it in the upper right corner, like so.

p	q	p $\vee$ $\neg$ q
T	T	
T	F	
F	T	
F	F	

This completes the first part of the construction. It will soon become clear why I created so much space in the rightmost column.

The second part is as follows. Look at *each* assignment of truth values to sentence letters. That is, look at *each* row in the above truth table. For each such assignment, use it to determine the truth value of the rightmost sentence.

Let us begin with the first assignment. According to that assignment, 'p' is true and 'q' is true. What is the truth value of '$p \lor \neg q$', given that assignment? To answer this question, let us break it up into pieces.

For starters, what is the truth value of '$\neg q$', given that assignment? The truth table for negation provides the answer: since 'q' is true on this assignment, the first row of the truth table for negation implies that '$\neg q$' is false. And so what is the truth value of '$p \lor \neg q$'? The second row of the truth table for disjunction provides the answer: since 'p' is true on this assignment, and '$\neg q$' is false on this assignment, the second row of the truth table for disjunction implies that '$p \lor \neg q$' is true.

All that reasoning, in the paragraph above, can be carried out using the truth table. Let us see how. While doing so, keep in mind that right now, we are only looking at the first assignment of truth values; that is, we are only looking at the first row in the truth table.

To start, under each sentence letter in the rightmost sentence, list that sentence letter's truth value *for that first assignment.*

p	q	p	$\lor$	$\neg$	q
T	T	T			T
T	F				
F	T				
F	F				

This amounts, basically, to just taking the truth values of the sentence letters on the left—for the first assignment—and copying them under the sentence letters on the right.

Now ask a somewhat vague but very important question: given that 'p' is true and 'q' is true—on this first assignment—what truth values of what

sentences can be immediately determined? Note that to determine the truth value of '$p \lor \neg q$', we need to know the truth value of 'p' and *also* the truth value of '$\neg q$'. On this first assignment, the truth value of 'p' is clear: it is 'T'; I just wrote that in the rightmost column of the table above. We do not yet know the truth value of '$\neg q$', however. So before determining the truth value of '$p \lor \neg q$', we first need to determine the truth value of '$\neg q$'.

And what is that truth value? Recall that on this first assignment, 'p' is true and 'q' is true. We can use this, it turns out, to determine the truth value of '$\neg q$'. For since 'q' is true, the first row of the truth table for negation implies that '$\neg q$' is false. So put an 'F' under the '$\neg$' in the first row of the above truth table, like so.

p	q	p	$\lor$	$\neg$	q
T	T	T		F	T
T	F				
F	T				
F	F				

The 'F' under the '$\neg$' represents the fact that on the first assignment, '$\neg q$' comes out false.

Now ask: given that 'p' is true and '$\neg q$' is false—on this first assignment— what is the truth value of '$p \lor \neg q$'? By the second row of the truth table for disjunction, '$p \lor \neg q$' true. So put a 'T' under the '$\lor$' in the second row of the above truth table, like so.

p	q	p	$\vee$	$\neg$	q
T	T	T	**T**	F	T
T	F				
F	T				
F	F				

Note that the 'T' is bolded. This makes clear to everyone involved—namely, you and I—exactly which truth value in the rightmost column is **the** truth value for the sentence at the top of that column.

We are still not done with the second part of the construction. But before continuing, here is a summary of what we have accomplished. We asked the following question: what is the truth value of '$p \vee \neg q$' on the first assignment of truth values to the sentence letters in that sentence? We answered this question in step-by-step fashion. We used that first assignment to determine the truth value of '$\neg q$' (for that assignment), and then we used *that* to determine the truth value of '$p \vee \neg q$' (for that assignment).

To complete the truth table—and in so doing, the second part of the construction—we repeat the above reasoning for *all* the truth value assignments. So we do it for the second assignment – which takes 'p' to be true and 'q' to be false – the third assignment – which takes 'p' to be false and 'q' to be true – and the fourth assignment – which takes 'p' to be false and 'q' to be false. For each assignment, we first determine the truth value of '$\neg q$' (for that assignment), and then we use that to determine the truth value of '$p \vee \neg q$' (for that assignment). As a result, we get the following truth table.

p	q	p	$\vee$	$\neg$	q
T	T	T	**T**	F	T
T	F	T	**T**	T	F
F	T	F	**F**	F	T
F	F	F	**T**	T	F

This truth table provides a *complete* answer to the question with which we began. Recall the question: when is '$p \vee \neg q$' true, and when is '$p \vee \neg q$' false? Or a little more precisely: under what assignments of truth values to the sentence letters in '$p \vee \neg q$' does that sentence come out true, and under what assignments of truth values to the sentence letters in '$p \vee \neg q$' does that sentence come out false? The above truth table provides the answer. The sentence '$p \vee \neg q$' is true on the first, second, and fourth assignments. On all other assignments—namely, the third one—'$p \vee \neg q$' is false.

Note that to determine the truth value of '$p \vee \neg q$', we first had to determine the truth values of the simpler sentences—namely, 'p', 'q', and '$\neg q$'—that it contains. This might have been confusing. People sometimes struggle to figure out which simpler sentences' truth values need to be determined first, before determining the truth value of a given larger sentence X. There is always, in fact, a particular order in which you must proceed: first you must determine the truth values of the simplest sentences, then you must determine the truth values of the 'next simplest' sentences, and so on, until finally you can determine X's truth value. But people sometimes struggle to see what the simplest sentences are, what sentences count as 'next simplest', and so on.

If you ever get confused about that, ask the following question: how would I build up X, using simpler sentences? More precisely, how would I build up X, using the Recursive Definition of a sentence of $\mathcal{P}$ from Chapter 3.4? The

answer to that question tells you more than just how X is built from simpler sentences. The answer *also* tells you the order in which you must proceed, to determine the truth value of X. For the order in which you must proceed, to determine the truth value of a sentence, is *exactly the same* as the order in which you must proceed, to build that sentence using the Recursive Definition.

For example, consider the sentence '$p \vee \neg q$'. Here is how you would build that sentence, using the Recursive Definition from Chapter 3.4:

1. 'p' is a sentence, since it is a sentence letter;
2. 'q' is a sentence, since it is a sentence letter;
3. '$\neg q$' is a sentence, since 'q' is a sentence;
4. '$p \vee \neg q$' is a sentence, since both 'p' and '$\neg q$' are sentences.

And here is the order in which you must proceed, to determine the truth value of '$p \vee \neg q$' for the first truth value assignment:

1. 'p' is true, since it is true on that first assignment;
2. 'q' is true, since it is true on that first assignment;
3. '$\neg q$' is false—by the first row of the truth table for negation—since 'q' is true;
4. '$p \vee \neg q$' is true—by the second row of the truth table for disjunction—since 'p' is true and '$\neg q$' is false.

As you can see, the order in which we build the sentence '$p \vee \neg q$', using the Recursive Definition, is exactly the same as the order in which we determine the truth values of the simplest sentences, the 'next simplest' sentences, and so on, which '$p \vee \neg q$' contains. So if you ever get confused about the order in which you must proceed, to determine the truth value of some complicated sentence X, just figure out the order in which you must proceed, to build X from simpler sentences. The latter is a guide to the former.

All this, in fact, is why we went through such pains to rigorously define what the sentences of $\mathcal{P}$ are. That rigorous definition acts as a guide for determining the truth value of any sentence of $\mathcal{P}$ whatsoever. In other words, the Recursive Definition contains the materials needed to give a complete account of truth in $\mathcal{P}$.[4]

Let us now consider two more examples of truth tables for sentences of $\mathcal{P}$. The first is a truth table for the sentence '$p \rightarrow (p \vee \neg p)$'. The second is a truth table for the sentence '$p \leftrightarrow (r \wedge (\neg p \vee q))$'.

To start, consider '$p \rightarrow (p \vee \neg p)$'. Under what assignments of truth values to its sentence letters does this sentence come out true, and under what assignments of truth values to its sentence letters does this sentence come out false? We will answer this question by constructing a truth table. Note that in what follows, I skip several steps; if I wrote out every step, then this chapter would be ridiculously long.

Recall that the construction proceeds in roughly two parts. In the first, we (i) identify all the sentence letters in '$p \rightarrow (p \vee \neg p)$', (ii) list those letters in columns of a table, (iii) list all possible assignments of truth values to those letters in those columns, and then (iv) put the sentence '$p \rightarrow (p \vee \neg p)$' in a column of its own. The result is below.

p	p	$\rightarrow$	$(p$	$\vee$	$\neg$	$p)$
T						
F						

---

[4]And in this respect, $\mathcal{P}$ is a much simpler language than, say, English. No one knows an analogously rigorous definition for what the English sentences are, because compared to $\mathcal{P}$, English is extremely complicated. And because of that, no one has proposed a successful, rigorous account of what makes any given English sentence true. That is one of the foundational open projects in linguistics, computer science, psychology, and philosophy.

Note that unlike the example from before, there are just two truth value assignments to the sentence letters in '$p \rightarrow (p \vee \neg p)$'. The reason is straightforward: there is only one sentence letter in '$p \rightarrow (p \vee \neg p)$'—namely, the sentence letter 'p'—so there are only two ways of assigning truth values to the sentence letters which '$p \rightarrow (p \vee \neg p)$' contains.

Now for the second part: we look at *each* assignment, and determine the truth value of '$p \rightarrow (p \vee \neg p)$' under it. As before, we start with the first assignment. And as before, we work in step-by-step fashion. In particular, we (i) write down the truth value of each sentence letter in '$p \rightarrow (p \vee \neg p)$' under the instances of that letter in the sentence at the top of the rightmost column, (ii) determine the truth values of the sentences which '$p \rightarrow (p \vee \neg p)$' contains—in particular, '$\neg p$', and then '$p \vee \neg p$'—and then (iii) determine the truth value '$p \rightarrow (p \vee \neg p)$' itself. The result is below.

p	p	$\rightarrow$	$(p$	$\vee$	$\neg$	$p)$
T	T	**T**	T	T	F	T
F						

Here is how I completed the first row in the above truth table. For starters, I determined the truth value of '$\neg p$'. To do so, I observed that since 'p' is true, the first row of the truth table for negation implies that '$\neg p$' is false. So I wrote an 'F' under the '$\neg$' in the table above. Then I observed that since 'p' is true and '$\neg p$' is false, the second row of the truth table for disjunction implies that '$p \vee \neg p$' is true. So I wrote a 'T' under the '$\vee$' in the table above. Finally, I observed that since 'p' is true and '$(p \vee \neg p)$' is true, the first row of the truth table for the conditional implies that '$p \rightarrow (p \vee \neg p)$' is true. So I wrote a bolded 'T' under the '$\rightarrow$' in the table above. And then I was done:

for that bolded 'T' represents the fact that on the first assignment of truth values, '$p \rightarrow (p \vee \neg p)$' comes out true.

Finally, repeat for each other assignment. In this case, there is just one.

p	p	$\rightarrow$	$(p$	$\vee$	$\neg$	$p)$
T	T	**T**	T	T	F	T
F	F	**T**	F	T	T	F

This truth table provides a complete answer to the question with which we began. Recall the question: under what assignments of truth values to the sentence letters in '$p \rightarrow (p \vee \neg p)$' does that sentence come out true, and under what assignments of truth values to the sentence letters in '$p \rightarrow (p \vee \neg p)$' does that sentence come out false? The above truth table provides the answer. The sentence '$p \rightarrow (p \vee \neg p)$' is true on *each* assignment. It is, in other words, always true.

It is worth doing one more example. This one involves a particularly complicated sentence: '$p \leftrightarrow \left(r \wedge (\neg p \vee q) \right)$'. Here is the truth table before the truth value of '$p \leftrightarrow \left(r \wedge (\neg p \vee q) \right)$' has been determined for each assignment.

p	q	r	p	$\leftrightarrow$	$(r$	$\wedge$	$(\neg$	p	$\vee$	$q))$
T	T	T								
T	T	F								
T	F	T								
T	F	F								
F	T	T								
F	T	F								
F	F	T								
F	F	F								

Note that there are eight truth value assignments, because there are eight different combinations of truth value assignments to the three sentence letters 'p', 'q', and 'r' which show up in '$p \leftrightarrow (r \wedge (\neg p \vee q))$'.[5]

And here is the completed truth table for '$p \leftrightarrow (r \wedge (\neg p \vee q))$'.

p	q	r	p	↔	(r	∧	(¬	p	∨	q))
T	T	T	T	**T**	T	T	F	T	T	T
T	T	F	T	**F**	F	F	F	T	T	T
T	F	T	T	**F**	T	F	F	T	F	F
T	F	F	T	**F**	F	F	F	T	F	F
F	T	T	F	**F**	T	T	T	F	T	T
F	T	F	F	**T**	F	F	T	F	T	T
F	F	T	F	**F**	T	T	T	F	T	F
F	F	F	F	**T**	F	F	T	F	T	F

So the sentence '$p \leftrightarrow (r \wedge (\neg p \vee q))$' is true on the first, sixth, and eighth assignment; for in the corresponding rows, the bolded letter is 'T'. On all other assignments, '$p \leftrightarrow (r \wedge (\neg p \vee q))$' is false.

Truth tables are among the most important things which you will learn in this book. For they play a crucial role in the definition of validity for arguments in $\mathcal{P}$; I discuss this in the next section. And because they can be used to define validity in $\mathcal{P}$, truth tables also play a crucial role in the definition of validity for arguments in English; I discuss this in the next chapter. So make sure that you are very comfortable with truth tables.

---

[5]There is a formula which, given the number of distinct sentence letters in a sentence, specifies the total number of different combinations of truth value assignments to the sentence letters in that sentence. Here it is: if there are n sentence letters in a given sentence, then there are 2^n truth value assignments for that sentence.