Elementary logic for lawyers

Jaap Hage *

This text provides an explanation of some basic concepts of logic and an introduction to propositional logic and class logic.

* The original version of this text was written by the author in Dutch. Sascha Hardt made the translation into English which was in turn edited by the author. The author thanks Sascha for the marvellous job he has done.

I. WHAT IS LOGIC AND WHAT IS ITS USE?

1. BASIC CONCEPTS

What is logic? This is not an easy question to answer, but the following is a good start: *Logic is the study of good reasoning.*[1] This definition is still ambiguous, because the term reasoning itself is ambiguous. Reasoning can mean a process of thought, an argument as it is uttered or written, but also the content of such a process of thought or argument. In the following, we are not primarily concerned with the form in which this content is cast, so not with thoughts and words. Instead, we are concerned with the *content* of arguments, with what is being argued and with the substance of the reasons which are put forward in support of a conclusion.

1.1 What is an argument?

An argument consists of one or more *premises* and exactly one *conclusion*. The premises are the starting points of the argument; the conclusion is what is assumed to follow from these premises. Examples of arguments are:

1. Thieves are punished. Jean is a thief. Therefore, Jean is punished.
2. When it rains, you need an umbrella. It is raining. Therefore, you need an umbrella.
3. When it rains, you need an umbrella. You don't need an umbrella, since it isn't raining.
4. Herding dogs are dogs and dogs are animals. Therefore, herding dogs are animals.
5. Paris is the capital of France. Therefore, Paris is not located in France.
6. André is taller than Timmy and Timmy is taller than his father. And Timmy's father is taller than Timmy's mother. Hence, Timmy's mother is taller than André.
7. If Berlin is located in Flanders, Brussels is not the capital of England. Brussels is the capital of England. Therefore, Berlin is not located in Flanders.

Arguments 1 to 4 in the list above all have two premises and one conclusion. However, they do not appear in the same order, as the conclusion of argument 3 is placed between the two premises. Argument 4 shows that two premises may be found in one sentence. Arguments 5 and 6 show that there can be arguments with less or more than two premises. Moreover, arguments 3 and 6 show that there can also be invalid arguments, in which the conclusion does not follow from the premises. Finally, argument 7 illustrates that arguments with nonsensical premises may still be valid: the conclusion follows from the premises, even though this is perhaps not immediately clear.

The arguments above all consist of only one argumentative step. An argument may also consist of two or more steps. Where this is the case, the conclusion of the first step constitutes a premise for the second step, or the conclusion of the second step constitutes a premise for the third step, and so on. The following are examples of arguments with two and three argumentative steps, respectively:

[1] This characterisation deviates from the more common one which holds that logic is the study of *valid* reasoning. The reason for this deviant definition is that the boundary between validity and soundness (see sections 1.3 and 1.4) can only be clearly defined where reasoning is formalised (see section 3.1).

1. Danny is six years old, and therefore younger than 18 years. A person younger than 18 years is a minor for the purposes of criminal law. Therefore, Danny is a minor for the purposes of criminal law. Minors are excluded from the application of regular criminal law. Hence, Danny is excluded from the application of regular criminal law.

2. Elly has bought the book from Petra. Therefore, Elly is the owner of the book and may forbid Anouk to read it. Now that Elly has indeed forbidden Anouk to read the book, Anouk may not read Elly's book.

The first of these two arguments is hopefully self-explaining. The second is somewhat more complicated, since a number of premises is not explicitly mentioned.

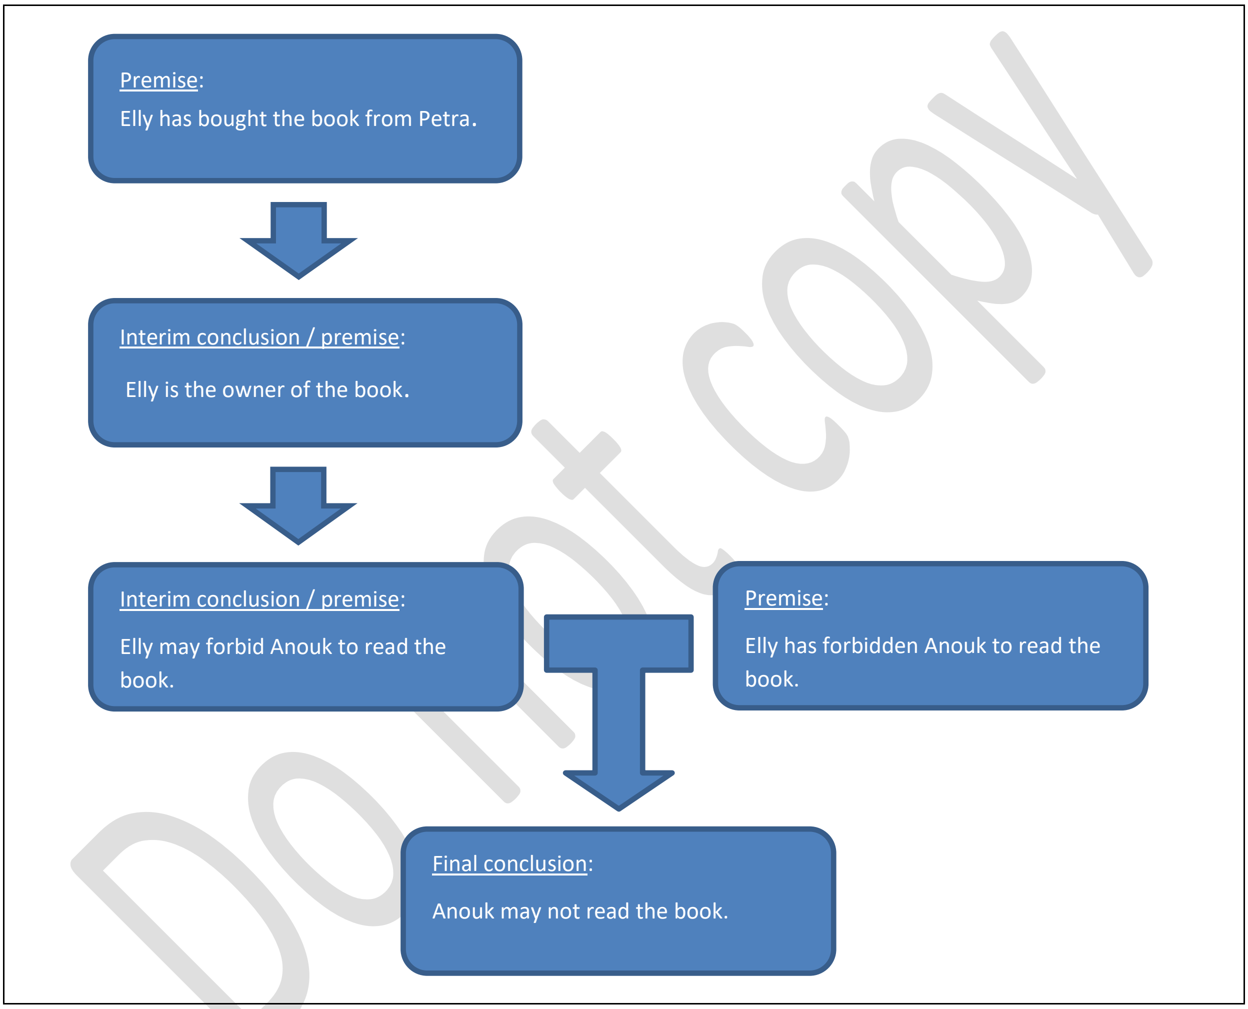

Although it is possible that one argument has two or more conclusions, we will assume – for the sake of convenience – that one argument can have only one conclusion. [2] If more than one conclusion is present, we assume that there are several arguments, one for each conclusion. For example:

> Mary has smeared syrup into Anne's hair out of jealousy. Therefore, Mary must bear the costs of Anne's visit to the hairdresser and is punishable for assault.

We analyse this as two separate arguments:

[2] This is more convenient, as it would otherwise be possible that an argument has one conclusion which does follow from the premises and one which does not. Is such an argument valid or not?

1. Mary has smeared syrup into Anne’s hair out of jealousy. Therefore, Mary must bear the costs of Anne’s visit to the hairdresser.
2. Mary has smeared syrup into Anne’s hair out of jealousy. Therefore, Mary is punishable for assault.

1.2 Propositions

Premises and conclusions are *propositions*. A proposition is anything that is expressed through a declarative sentence. The sentences ‘It is raining’, ‘Het regent’ and ‘Il pleut’ express the same proposition. The same is true for the declarative sentences ‘Pierre drives the car’ and ‘The car is driven by Pierre’.

The sentence ‘I am hungry’, uttered by Sascha, expresses another proposition than the same sentence uttered by Jaap. The same sentence uttered at 8 o’clock in the morning expresses another proposition than when it is uttered at half past four in the afternoon.

Therefore, it is possible that different declarative sentences all express the same proposition, while the same declarative sentence, uttered by someone else or at a different time or location, expresses each time a different proposition.

By the way, there are many sentences which are not declarative:

Run to the moon!
I’m Jaap. (in introducing oneself)
A cup of coffee, please. (at the café)
Asshole!
Could you please open the window?
I christen this ship the *President van Rompuy.*

Propositions have a ‘*truth value*’. This truth value is always either ‘*true*’ or ‘*false*.[3] The truth value of a declarative sentence is the truth value of the proposition expressed by this sentence. Since a declarative sentence can express different propositions depending on its context, it is therefore possible that the declarative sentence ‘I am hungry’, uttered by Johnny, has a different truth value from that of the same sentence uttered by Maurice. The first sentence means that Johnny is hungry and is true if he actually is. The second sentence means that Maurice is hungry and is true if Maurice is hungry.

1.3 Validity

An argument is the inference of a conclusion from one or more premises. We will primarily focus on deductive arguments, or – better – arguments which are valid according to standards of deduction (deductive validity).[4]

[3] ‘*Unknown*’ is not a truth value, but relates our knowledge about what is true. Nevertheless, ‘unknown’ could play a role in logic. We would then have three categories: knowing that it is true, unknown, and knowing that it is not true. We will not go further into this here, even though this form of logic is relevant to lawyers. Think, for instance, of proof in criminal law.

[4] Other forms of validity, where the premises of an argument make its conclusion acceptable but do not warrant its truth, will be left out of consideration. Nevertheless, these forms of validity are sometimes of great importance for lawyers. Think, for instance, of reasoning by analogy, *a fortiori* and *e contrario*.

An argument is deductively valid if – and only if – it is logically impossible that all premises of the argument are true while the conclusion is not.

The validity of an argument relates only to the question whether the conclusion follows from the premises. Validity does not say anything about the *truth* of the premises or of the conclusion as such. Hence, a valid argument may very well have false premises or an false conclusion. It may also be the case that a valid argument has one or more false premises while its conclusion is still true. What is impossible *in combination* is the following:

- that an argument is deductively valid;
- that *all* premises of the argument are true; and
- that the conclusion of the argument is not true.

All other combinations of truth and falsity are possible in a valid argument. The following arguments are all valid[5]:

All horses can fly.
Birds cannot fly.

Birds are not horses.

Jan is a thief.
Thieves deserve a reward.

Jan deserves a reward.

Either André still has my book or Juliette has burned it.
Juliette has not burned my book.

André still has my book.

1.4 Soundness

Where an argument is valid, that does not guarantee that its conclusion is true. After all, it may be the case that at least one of the premises of a valid argument is false. In such a case, the conclusion *may* also be false. In order to be certain that the conclusion is true, we need a valid argument with true premises. Such an argument is called *sound*.

An argument is sound if it is valid and has only true premises.

If an argument is valid and has true premises, its conclusion must logically also be true. Therefore every sound argument has a true conclusion (check this).

[5] The horizontal line separates the premises from the conclusion of an argument. Above the line are the premises (one or more), while the conclusion is below the line. The line makes no claim as to whether the conclusion actually follows from the premises.

The following arguments are sound:

Not a single prime number can be divided by three.
103 is a prime number.
----
103 cannot be divided by three.

The state of New York is situated in the United States.
No state is situated both in the United States and in Europe.

The state of New York is not situated in Europe.

Exercises

1. Give definitions of
   - an argument;
   - a proposition;
   - (deductive) validity;
   - soundness.

2. Explain why it is possible that one sentence, depending on context, can express different propositions.
3. Explain why it is possible that different sentences express the same proposition.

4. For each of the following arguments, determine if it is sound:

   a. Cows are animals.
   Animals are born.

   Cows are born.

   b. Lawyers are good at arguing.
   Donkeys are not lawyers.

   Donkeys are not good at arguing.

   c. Even numbers are divisible by four or become divisible by four after two has been added to them.
   The number 14 is not divisible by four and also does not become so after two has been added.
   ----
   14 is not an even number.

2. LOGICAL FORM

2.1 An example

Consider the following argument:

If the social democrats win the elections, the VAT will be reduced in combination with the introduction of a tax on an increase in the value of shares. The social democrats are not going to win the elections. --- It is not the case that the VAT will be reduced in combination with the introduction of a tax on an increase in the value of shares.

Many will still find it difficult to determine if this argument is valid. But now look at the following argument:

If Belle is a dog, she is an animal. Belle is not a dog. -- Belle is not an animal.

The latter argument is clearly invalid. After all, Belle may be another animal than a dog. In that case, both premises would be true, but the conclusion would be false. But if the latter argument is invalid, then the former argument about taxes must be equally invalid, since it is the same *kind* of argument. Only the sentences of which it is composed are a bit longer and the argument concerns a subject about which most people know less than about dogs and other animals. Therefore, many people recognise the invalidity of the former argument less easily.

2.2 Why logical form is important

The above thought – that, if one argument is of 'the same kind' as another, they must either both be valid or both be invalid – illustrates the importance of what logicians call the 'form' of an argument. What an argument is about, its content, is not relevant for its validity; only its form is. The invalidity of the first of the two arguments above has nothing to do with the social democrats or with taxes, just as the invalidity of the second argument has nothing to do with dogs and other animals. All of that is content. What is important, however, is that both arguments have the same form:

If A then B.

Not A.

Therefore not B.

The general *form* of the argument is invalid, and *that* is why all particular arguments which share this form are also invalid.[6]

[6] This is not a completely exact way of putting it, since an argument does not necessarily have only one logical form but can potentially have two or more (see chapter III.1). An argument is still logically valid if at least one of

The validity of an argument is determined by its logical form. If a logical form is valid, all arguments of this form are valid. If a form invalid, so are all arguments of this form.[7]

2.3 Some more examples

We already discussed an example of a logical form which is invalid and two arguments which have this form and are therefore invalid as well. In the following sections, we will examine valid and invalid argumentative forms in a systematic manner. Some examples may already be useful at this point.

One logical form is:

If A then B.

A.

Therefore B.

This is a valid logical form, which is known as 'Modus (Ponendo) Ponens'.[8] The following two arguments take this form and are therefore valid:

If the minister is a criminal, he must resign.
The minister is a criminal.

The minister must resign.

The weather lady forecasts rain and there was no error in the weather computer.
If the weather lady forecasts rain and there was no error in the weather computer, it is going to rain.

It is going to rain.

The latter argument illustrates that what comes after 'If…' may be rather complicated and that the order in which the premises appear does not determine the logical form.

The second example concerns the following logical form:

No A is B.

X is an A.

Therefore, X is no B.

its logical forms is valid. For now, however, we may simply assume that an argument has precisely one logical form and that this form determines the validity of the argument.

[7] This close entanglement of validity and logical form is what makes it possible to study logic as an independent academic discipline. If we had to study the validity of each individual argument, the complexity of the subject would render logic as a discipline impossible. But since validity is determined by the form and not by the content of an argument, logic 'only' has to deal with logical forms, which makes it a bit easier.

[8] This logical form will be discussed in more detail in chapter II.2.4.

This form is valid, which therefore is also true for the following argument of the same form:

Didier is a lawyer.
Not a single lawyer understands mathematics.

Didier does not understand mathematics.

This example does not only illustrate the valid logical form but also that 'translating' an argument rendered in natural language into its logical form is not always simple. For instance, the logical form contains the phrase 'is a(n)', whereas this phrase is not found in the argument. And in the conclusion of the logical form we see 'is no', while it says 'does not' in the actual argument.

Exercises

It is not easy to determine what the logical form of an argument is without having defined beforehand which parts of an argument concern its form and which parts relate to its content. Nevertheless, it is often possible to give a correct answer – or at least one that is not wrong – by intuition.

1. Determine the logical form of each of the following arguments.

   a. Someone who is a lawyer can become a judge.
   Petra cannot become a judge.

   Petra is not a lawyer.

   b. If he has won the lottery, he is rich.
   He is rich.

   He has won the lottery.

   c. If you have the flu, you have the symptoms X, Y en Z.
   Jean has the symptoms X, Y and Z.

   Jean has the flu.

2. Which of the following arguments has the same form as that in exercise 1b above?

   a. If it rains, the roofs get wet.
   It is raining.

   Therefore, the roofs are getting wet.

   b. If it rains, the roofs get wet.
   The roofs are getting wet.

   Therefore, it is raining.

c. If it rains, the roofs get wet.
   The roofs are not getting wet.

   Therefore, it is not raining.

3. INTERIM SUMMARY

Logic is the study of good arguments and that usually means the study of deductively *valid* arguments. An argument consists of one or more *premises* and – per definition – one *conclusion*.

Both the premises and the conclusion are *propositions*. A proposition is what is expressed by a *declarative sentence*. Exactly which proposition it is that is expressed by a declarative sentence can depend on the speaker, the place and time at which the sentence is uttered, and possibly on other circumstances.

An argument is *deductively valid* (hereinafter: valid) if it is logically impossible that all its premises are true while its conclusion is not. A valid argument with only true premises guarantees that the conclusion is also true. Such an argument is called *sound*.

Whether an argument is valid depends on its logical form. Arguments which share the same logical form are either all valid or all invalid, depending on whether the form is valid or invalid.

4. WHAT IS LOGIC GOOD FOR?

Logic is first and foremost useful for determining whether a given argument is good. In rough terms, one could say that an argument is good if it is technically sound. That means that all the premises of the argument must be true and that the argument is valid, so that the conclusion follows from the premises. Generally, logic does not assess the truth of the premises,[9] but only whether the conclusion logically follows from the premises – in other words, whether the argument is valid. Where a conclusion does not logically follow from the premises, this does not automatically say that they are not true, but only that they do not support the conclusion. If the argument in question was the only reason to consider the conclusion true, invalidity of that argument means that there is no longer a reason to do so.

Consider, for instance, the following line of argumentation:

> If Guy has committed the murder, he had gunshot residues on his hands.
> Guy had gunshot residues on his hands.
> Hence, Guy has committed the murder.
> A person who has committed murder must be given a prison sentence.
> Therefore, Guy must be given a prison sentence.

The first argument, whose conclusion is that Guy has committed the murder, is invalid, since it has the following invalid form:

[9] If the premises of an argument are interim conclusions, their truth usually cannot be proven on logical grounds alone. However, their credibility – which is quite distinct from truth – is affected where the arguments from which they follow are invalid. And validity *can* be assessed by logic.

If A then B.

B.

Therefore A.

Therefore, the gunshot residues do not constitute evidence – or at least constitute insufficient evidence[10] – for the claim that Guy has committed the murder. If there is no other proof that Guy has committed the murder, the claim that he has done it is not supported. The second conclusion – that Guy must be given a prison sentence – follows from an argument which is valid, but it still lacks support, since the truth of a conclusion is only guaranteed if all the premises are true. And since one of the premises of the second argument is the conclusion of the invalid first one, we cannot be certain of Guy's guilt.

Logic can not only be used to assess the quality of existing arguments, but also to formulate good arguments. The latter is particularly useful when it comes to writing argumentative texts, i.e. texts which seek to convince the reader of a certain position or claim. In essence, such a text is one long argument which consists of several steps.

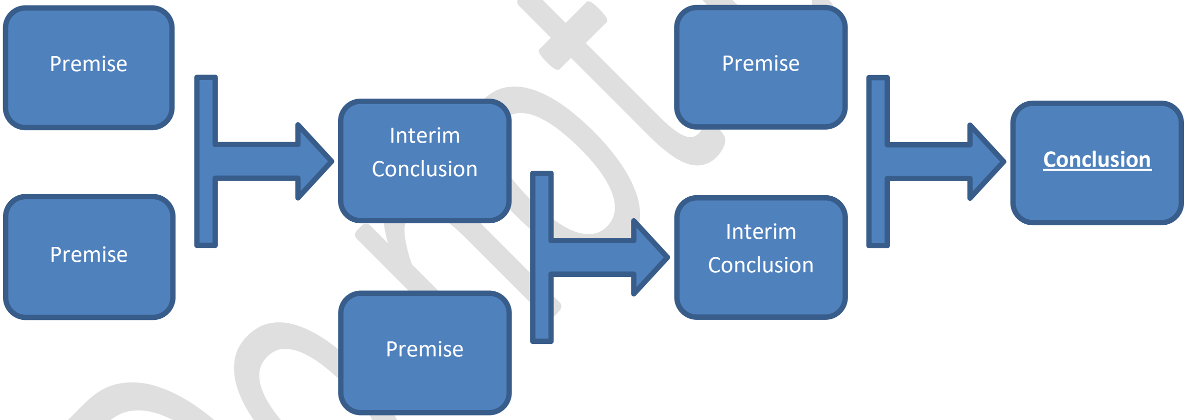

If each of the individual steps of such a line of argumentation constitutes a valid argument, the final conclusion logically follows from the line of argumentation as a whole. This does not guarantee the truth of the conclusion, since one or more of the premises can still be false. But if the line of argumentation is valid, one of the two requirements of a good argumentative text is fulfilled. The other requirement is that the premises are true, but logic has little to say about that.

Not only can logic show which premises are necessary to support the final conclusion, it can also show what is *not* necessary. In a 'stringent' argumentative text, each part of the argumentation is a premise, an interim conclusion, the final conclusion,[11] or something closely connected to one of them, such as an example. Anything which is not a premise, interim conclusion or final conclusion is

[10] The presence of gunshot residues has a certain evidence value, but does not by itself guarantee the truth of the conclusion. The argument therefore does have a certain force but is not deductively valid.
[11] As opposed to a single argument, an argumentative text frequently has more than one conclusion. Where that is the case, the argumentative text consists of more than one line of argumentation.

therefore at first glance irrelevant and may be omitted. *At first glance*, since a readable text – also an argumentative one – also contains elements which enhance its readability. But including such elements in a text requires that there are good reasons for doing so. Where no such reasons exist, we are dealing with superfluous passages which may better be deleted.

5. WHAT FOLLOWS

In the following chapters, two issues will be discussed. The second chapter is going to expound on valid and invalid argumentative forms in so-called 'propositional logic'. Propositional logic deals with arguments like this:

> The weather lady forecasts rain and there was no error in the weather computer.
> If the weather lady forecasts rain and there was no error in the weather computer, it is going to rain.
> --
> It is going to rain.

The third chapter deals with arguments relating to entire classes or categories, such as 'all murderers', or members of such categories, such as Guy, who has committed murder. Therefore, it deals with arguments like the following:

> Guy has committed murder.
> A person who has committed murder must be given a prison sentence.
> --
> Guy must be given a prison sentence.

We will discuss such arguments on the basis of so-called Venn diagrams (overlapping circles). The logic which applies to them is a variant of what is referred to as 'syllogistics'.

II. PROPOSITIONAL LOGIC

Propositional logic is logic based on the meaning of so-called logical operators, such as 'not', 'and', 'or' and 'if … then'. These operators combine elementary propositions into compound propositions. For instance, the 'if … then' operator combines the elementary proposition 'It rains' and 'The roofs get wet' into 'If it rains, then the roofs get wet'.

Given the meaning of 'if … then', the following argument is deductively valid:

IF it rains, THEN the roofs get wet.
It rains.

The roofs are getting wet.

Propositional logic is a formal logic and in discussing it we will encounter a number of symbols and formulae. But first, we are going to look at some informal examples to get an expression of what propositional logic is all about.

1. ELEMENTARY AND COMPOUND PROPOSITIONS

In a certain sense, the proposition 'It is raining and the roofs are getting wet' actually consists of two propositions: 'It is raining' and 'The roofs are getting wet'. As far as propositional logic is concerned, these two propositions cannot be split any further; they are *elementary propositions*. Elementary propositions can be used to create *compound propositions*. Examples of compound propositions are:

1. It is raining AND the roofs are getting wet.
2. It is raining OR the roofs are getting wet.
3. IF it is raining THEN the roofs are getting wet.

and – perhaps surprisingly

4. The roofs are NOT getting wet.

The first three examples illustrate how two elementary propositions can be combined to one compound proposition. The words AND, OR and IF … THEN are called *logical operators*. All three of them are operators which fuse two propositions together to form one proposition.

NOT is also a logical operator, but it does not combine two propositions but transforms one proposition into another, the negation of the first proposition. The proposition 'The roofs are NOT getting wet' is therefore also a compound proposition, but a compound proposition based on only one elementary proposition. In propositional logic, NOT is the only logical operator which works with only one proposition.

In the four examples we have seen compound propositions based on one or two elementary propositions. But also compound propositions can be part of another – more complicated – compound proposition, for example:

5. It is NOT raining, OR the roofs are getting wet.
6. The roofs are covered in plastic OR the roofs get wet IF it rains.

Example 5 illustrates how a NOT-proposition can be combined with an elementary proposition. In example 6, an elementary proposition is combined with an IF … THEN-proposition. In order to let the sentence run more smoothly, the word THEN is omitted. Moreover, the THEN-part now stands before the IF-part. This shows, once again, that the 'translation' from natural language into the language of propositional logic is not always straightforward.

2. LOGICAL OPERATORS – AN INFORMAL MEETING

We will examine the logical operators more closely at a later point, but it is useful first to get an impression of what they are about on the basis of informal examples. The informal operators which are important in this respect are NOT, END, OR and IF … THEN.

2.1 The operator NOT

The operator NOT transforms a proposition into another proposition which is the negation of the original one. 'It is NOT raining' is the negation of 'It is raining'.

This transformation entails that if the proposition 'It is raining' is true, the proposition 'It is NOT raining' is false, and vice versa.

WELL-FORMEDNESS

Since the NOT-operator can be used with every proposition and since a proposition with a negation contained in it is also a proposition, double negations are possible. 'It is NOT the case that it does NOT rain' is a 'well-formed' proposition. That is a proposition which satisfies the requirements of propositional logic for the structure of propositions. This does not mean that the proposition is true, since truth is not the same as well-formedness. A compound proposition is well-formed if the application of the logical operator has led to a new proposition. An example of a proposition which is not well-formed is: 'It is NOT the case that AND it is raining'.

DOUBLE NEGATION

The proposition 'It is NOT the case that it is NOT raining' is true if the proposition 'It is NOT raining' is false. The latter, in turn, is the case if the proposition 'It is raining' is true. In short, the proposition 'It is NOT the case that it is NOT raining' is true if 'It is raining' is true. The double negation cancels itself out. This can be expressed schematically:

'It is raining'	true	false
'It is NOT raining'	false	true
'Het is NIET zo dat het NIET regent'	true	false

Since 'It is NOT the case that it does NOT rain' is true if 'It is raining' is true and is false if 'It is raining' is false, the former proposition can be deduced from the latter and vice versa. It is always possible to add or remove double negations without changing the truth value.

The propositions to which the NOT-operator is applied may also be compound. For instance, the proposition 'It is NOT the case that (Jean is a thief AND a murderer)' is true if and only if the compound proposition 'Jean is a thief AND a murderer' is false.

IFF

In logic it often happens that a proposition is true (or false) if and only if some other proposition is true (or false). The expression ‘if and only if’ is a bit clumsy, and therefore the convention has arisen to replace this expression by the single word ‘iff’. We might therefore also have written that , the proposition ‘It is NOT the case that (Jean is a thief AND a murderer)’ is true *iff* the compound proposition ‘Jean is a thief AND a murderer’ is false.

2.2 The operator AND

The operator AND fuses two propositions together to form a compound proposition. This compound proposition is true iff both of its component parts are true. For instance, the proposition ‘Jean is a thief AND so is Petra’ is true only iff the propositions ‘Jean is a thief’ and ‘Petra is a thief’ are both true.

By the way, it is possible that the two component propositions are about completely different subjects. That is, for instance, the case in the proposition ‘Jean is a thief AND Petra is good at cooking’.

The propositions to which the AND-operator is applied may themselves also be compound. For example, ‘(Jean is a thief OR Jean is innocent) AND (Petra is NOT a thief)’ is a well-formed compound proposition. It is true iff the propositions ‘Jean is a thief OR Jean is innocent’ and ‘Petra is NOT a thief’ are both true.

Since an AND-proposition is true if both of its component parts are, and not otherwise, both of these component parts can be deduced from the AND-proposition in which they are contained. The following deductions are therefore logically valid:

It is raining AND (Petra is NOT a thief)

It is raining.

It is raining AND (Petra is NOT a thief)

Petra is NOT a thief.

Conversely, an AND-proposition can be deduced from the fact that both of its component parts are true:

It is raining.
Petra is NOT a thief.

It is raining AND (Petra is NOT a thief)

2.3 The operator OR

The operator OR fuses two propositions together to form a compound proposition. This compound proposition is true iff at least one of the two component propositions is true. For instance, the proposition ‘Jean is a thief OR Petra is good at cooking’ is true iff

- the proposition ‘Jean is a thief’ is true,
- or the proposition ‘Petra is good at cooking’ is true,
- or both component propositions are true.

As demonstrated by this example, the two component propositions need not have anything to do with each other.

In natural language, the word ‘or’ is also used to express that exactly one of two possibilities is the case, such as, for instance, in the sentence ‘Annie arrives tomorrow or the day after tomorrow’. But this meaning of the word ‘or’, in which two possibilities exclude each other, is not the meaning of the logical operator OR. In section 3.5 we will see that is possible to construct the so-called *exclusive or* by means of the operators NOT, OR and AND.

The proposition to which the OR-operator is applied may also be compound. For instance, the proposition ‘(Jean is a thief OR Jean is innocent) OR (Petra is NOT good at cooking)’ is well-formed. It is true iff the proposition ‘Jean is a thief OR Jean is innocent’ is true, or if the proposition ‘Petra is not good at cooking’ is true, or if both are true.

Since an OR-proposition is only true if at least one of the component propositions is true, the following arguments, which are very similar, are both valid:

```
Jean is a thief OR Petra is good at cooking.
Jean is NOT a thief.
-------------------------------------------------------
Petra is good at cooking.
```

```
Jean is a thief OR Petra is good at cooking.
Petra is NOT good at cooking.
------------------------------------------------------
Jean is a thief.
```

Conversely, an OR-proposition can be deduced from any random proposition:

```
It is raining.
----------------------------------------------
It is raining OR the roofs remain dry.
```

```
It is raining.
------------------------------------------------------------
It is raining OR Pierre is the owner of a chair.
```

From a logical point of view, nothing is wrong with the latter argument, because if the proposition ‘It is raining’ is true, the compound proposition ‘It is raining OR Pierre is the owner of a chair’ must also be true. Of course, in daily life this would be an odd argument. This example has been chosen on purpose to illustrate that not every argument which is logically valid is also a ‘reasonable’ argument in normal usage.

2.4 The operator IF … THEN

The operator IF … THEN fuses two propositions together to form a compound proposition. This compound proposition is true iff either the proposition after 'IF' is false, or the proposition after THEN is true, or both.

This is a strange definition which will be explained in section 3.6. But the reader should realise already at this point that the meaning of IF … THEN as a logical operator is not identical with the meaning of 'if … then' in normal language use. The difference is that, in normal language, 'if … then' implies a material link between the proposition after 'if' and that after 'then', while no such link needs to exist between propositions connected by the IF … THEN-operator.

MODUS PONENDO PONENS

There is, however, also an important similarity between 'if … then' in normal language use and IF … THEN as a logical operator, namely that arguments of the form 'modus ponendo ponens' are valid both in natural language and in propositional logic. These are arguments in which the proposition after THEN is deduced from the compound proposition and the proposition after IF. The following valid argument is an example of this:

IF Marcel has murdered Tina, THEN Marcel is punishable.
Marcel has murdered Tina.

Marcel is punishable.

But for the same reason, the following argument is also valid:

IF Marcel has murdered Tina, THEN Petra is good at cooking.
Marcel has murdered Tina.

Petra is good at cooking.

The Latin name *Modus Ponens*, which refers to the argumentative form *modus ponendo ponens*, is so commonly used that is unavoidable to briefly explain this name. The Latin verb 'ponere' means 'to deposit', 'to lay down' or 'to put'. In arguments of the form *modus ponendo ponens*, the conclusion is put forward (drawn; assumed as true), which is expressed by the word '*ponens*'. The conclusion is drawn *by* assuming the truth of a premise, which is expressed by the word '*ponendo*'. Therefore, in the *modus ponendo ponens* a conclusion is assumed to be true by assuming a premise to be true. The name '*modus ponendo ponens*' is abbreviated to '*Modus Ponens*'.

The propositions to which the IF … THEN-operator is applied may also be compound. For instance, the proposition 'IF (Jean is NOT a thief) THEN (John is punishable OR John has committed another crime)' is a well-formed proposition. It is true iff either the proposition 'Jean is NOT a thief' is false (therefore, if 'Jean is a thief' is true), or if the compound proposition 'John is punishable OR John has committed another crime' is true, or both.

‘STRANGE’ DEDUCTIONS

As explained earlier, an IF … THEN-proposition is true if the proposition after IF is false, or both. Therefore, the following to deductions of an IF … THEN-proposition are both valid:

Marcel has NOT murdered Tina.

IF Marcel has murdered Tina, THEN Marcel is punishable.

Marcel is punishable.

IF Marcel has murdered Tina, THEN Marcel is punishable.

The validity of these two arguments has to do with the fact no material link needs to exist between the proposition after IF and the proposition after THEN. Even the following argument is therefore valid:

Marcel has NOT murdered Tina.

IF Marcel has murdered Tina, THEN Pierre is the owner of a chair.

SUFFICIENT CONDITION

The IF … THEN-operator can clearly lead to strange deductions, but it is a very useful tool to demonstrate that something is a sufficient condition for something else. A *sufficient condition* is a condition which, if it is fulfilled, guarantees that this ‘something else’ is also the case. A nice example is the proposition:

IF Marcel has murdered Tina, THEN Marcel is punishable.

After all, if Marcel has murdered Tina, this guarantees that he is punishable.[12] In section 3.8 we will encounter examples of logical constructions which are akin to the sufficient condition, namely necessary conditions and conditions which are necessary and sufficient at the same time.

MODUS TOLLENDO TOLLENS

Another deduction based on the IF … THEN-operator is the argumentative form known as ‘*modus (tollendo) tollens*’. In arguments of this form, the part after THEN in a premise is denied and from this the conclusion is drawn that the part after IF is also false.

IF Marcel has murdered Tina, THEN Marcel is punishable.
Marcel is NOT punishable.

Marcel has NOT murdered Tina.

[12] Lawyers may object here that it may be the case that Marcel was not of sound mind or a minor and thus is not punishable. That is correct, but then the compound proposition ‘IF Marcel has murdered Tina, THEN Marcel is punishable’ is also not (unconditionally) true.

The name of this argumentative form is derived from the Latin word '*tollere*', which means 'to take away'. This taking away stands for to the denial of a proposition. By denying (*tollendo*) the part after THEN, the part after IF is also denied, which then constitutes the conclusion. Often, the shorter name *Modus Tollens* is used for *modus tollendo tollens*.

An argumentative form akin to Modus Tollens is the so-called 'transposition'. In this form, the sequence of the two component propositions of an IF … THEN-proposition is reversed and both component propositions are negated, and in this way a new IF … THEN-proposition is deduced:

IF Marcel has murdered Tina, THEN Marcel is punishable.

IF Marcel is NOT punishable, THEN Marcel has NOT murdered Tina.

3. FORMALISED PROPOSITIONAL LOGIC

Even though informal examples provide a good impression of how propositional logic works, a number of things can be better understood when looking at the 'official', formal variant. This is what we will do in the following. First, we will devote some attention to the advantages of presenting arguments in a formal language, such as that of propositional logic. Then, we will introduce that language itself in order to subsequently present the various logical operators with the help of so-called 'truth tables'. We will focus on a number of argumentative forms that are widely used.

3.1 The advantage of a formal language

Real arguments, certainly those of lawyers, are formulated in ordinary language. This language may well contain a number of technical terms but this does not take away that an argument is expressed in normal words and sentences. The language of propositional logic is a so-called 'formal language' which consists of abbreviations, symbols and formulae. To most people, this appears complicated. It is therefore legitimate to ask why we need a formal language.

The answer is simple: the formal language of propositional logic is clearly defined, which makes it easy to assess the validity of arguments rendered in this language. It is possible to translate arguments rendered in ordinary language into the formal language of propositional logic and to assess the validity of the result. The finding that the argument in formal language is valid/invalid equally applies to the same argument in ordinary language. To make this clearer, consider the following argument:

If Claude is the owner of this house, he can prohibit Lisette to enter it.
Claude is the owner of this house.

Claude can prohibit Lisette to enter the house.

This argument can be translated into the formal language of propositional logic ('formalised'). It then looks like this:

$E \supset V$
E

V

This argument takes the form Modus Ponens which we have already encountered earlier. Since this argumentative form is valid, the original argument in ordinary language is also valid.

The following argument is another example:

If the Scottish Nationalist Party wins the elections, Scotland will become independent or it will be impossible to form a Scottish government.
The Scottish Nationalist Party is going to win the elections, but it will certainly be possible to form a Scottish government.

Scotland will not become independent.

Assessing the validity of this argument is somewhat more difficult, which illustrates why formalisation can be useful. The downside is, however, that it is also more difficult to formalise this argument. If we try, we should get the following result:

$W \supset (I \vee {\sim}G)$
W

${\sim}I$

At first glance, this does not seem to be a great improvement with regard to the complexity of the argument. But now that the argument is formalised, we can easily apply a method which allows us to check its validity. This method involves so-called *truth tables* and can be programmed into a computer if desired. Applying this method will show that the formalised argument above is invalid (we will see this in section 4.5.). This means that the original argument is equally invalid, *provided that the formal argument is an accurate representation of the original argument*.

THE FUNCTION OF FORMALISATION

This addition, '*provided that the formal argument is an accurate representation of the original argument*', is by no means superfluous, since it not always easy to translate an argument from ordinary language into the formal language of propositional logic. The difficulty may seem to speak against the endeavour of formalising arguments, but there is a positive side to it as well: where formalising an argument is difficult, this is so because the argument is difficult to understand. Consequently, trying to formalise the argument will at least bring to light what the problem is. Once the argument is understood, formalising it and testing the validity of the formal argument is easy. In other words, formalising an argument forces us to do what we would anyway have to do, namely to study and understand the argument well enough in order to assess its validity. Once the formalisation is completed, the most important part of the work is done. This is the second, and perhaps the most important, advantage of a formal language: it makes it necessary to understand an argument well enough to assess its validity.

3.2 The language of propositional logic

The language of propositional logic allows the formulation of a number of well-formed propositions. The propositions can be either elementary or compound.

ELEMENTARY PROPOSITIONS

Elementary propositions are marked by capital letters. For greater clarity, we will use a specific font (`Courier New`) for them. Thus, examples of elementary propositions are `A B C … P Q R`. Each of these letters stands for one proposition which may also be expressed in ordinary language, such as 'John is a thief', 'Lisa is the owner of the house' or 'The civil servant is competent to issue a building permit'.

Often it is important to know which 'ordinary' proposition a certain letter stands for, but not always. In order to assess the validity of an argument, for instance, that is not important. After all, validity only depends on the form of the argument. For which proposition precisely a letter stands does not affect the form and therefore also not the validity of the argument.

An important requirement in propositional logic is that the truth of each *elementary* proposition is, from a logical point of view, independent of the truth of *all* other elementary propositions. For example, whether the proposition 'John is a thief' is true is logically independent of the truth of the proposition 'Lisa owns the house'.

*From a logical perspective, t*he truth value of 'John is a thief' is also independent of 'John is punishable'. This seems less straightforward, but the connection between John being a thief and his being punishable is based on a legal rule which, from a logical point of view, need not necessarily exist. In terms of *logic*, it is therefore possible that John is a thief but is not punishable, or that he is not a thief but is still punishable.[13] This logical independence of the truth values of elementary propositions will be essential in section 3.3, where truth tables will be discussed.

COMPOUND PROPOSITIONS

Compound propositions are propositions which feature a logical operator. The logical operators of propositional logic are:

operator	name	stands for
~	negation	NOT
&	conjunction	AND
∨	disjunction	OR
⊃	material implication	IF … THEN
≡	equivalence	IF and ONLY IF

Except one, all of these operators combine two propositions into a compound proposition. The only exception is the operator ~, which transforms one proposition into another, namely the negation of the original one.

It should be noted that logical operators work with *all* propositions, not only with elementary propositions. Hence, it is possible to combine two very complex compound propositions into one

[13] By the way, both of these constellations are also legally possible.

that is even more complex. For instance, the proposition ~(P ∨ Q) ⊃ (R&~S) is the result of the application of the ⊃ - operator to the propositions ~(P ∨ Q) and R&~S.

This example illustrates, by the way, that parentheses may be used within propositions to indicate which parts of a proposition belong together. There is a difference between the proposition ~(P ∨ Q) and the proposition ~P ∨ Q. If P stands for 'André possesses a car' and Q stands for 'André possesses a bike', then ~(P ∨ Q) means that it is not the case that André has a car or a bike (that is: he has neither), while ~P ∨ Q can in fact mean three things: (1) André has no car, (2) he does have a bike, or (3) André both has a bike and no car.

Exercise

For each of the following formulae, determine whether it is a well-formed proposition.

a. P
b. ~P
c. ~~P
d. q
e. Q~
f. P ⊃ Q
g. (P ⊃ Q)
h. ⊃P&~Q
i. (A ∨ B) ≡ ~(~A&B)

3.3 The ~ - operator

In this section and the following one, the different operators which form part of the formal language of propositional logic will be discussed. First, however, we must consider an important characteristic of propositional logic which is also of importance to the different operators. Logic is about the *form* of arguments and propositional logic, in particular, deals with the *possible truth values* of the propositions which feature as premises and conclusions in arguments. For neither of the two the content of the propositions is particularly important. For the argumentative form it does not matter at all. For the truth value of a proposition it is of course important what it is that a proposition says, but in propositional logic it is irrelevant *why* a proposition is or is not true. The only relevant question is *whether* it is true. In other words, a proposition merely stands for a truth value; how or why it got this truth value is immaterial as far as *logic* is concerned. In fact, logic does not even deal with the actual truth values of propositions but only with their *possible* truth values. After all, the validity of an argument does not depend on whether the propositions which feature in it are true or false.

All of this is still very abstract, but we will soon see what it means for logical practice once we consider the different logical operators.

TRUTH TABLES

Informally, the ~-operator stands for the denial of a proposition, also referred to as *negation*. Formally, it is an operator which affects *one* proposition and which transforms this proposition into a new proposition whose truth value is the opposite of that of the original proposition. Therefore, a true proposition becomes false and a false proposition becomes true.

The meaning of the ~ - operator can be shown in a so-called truth table. A truth table for a compound proposition Π is a table which displays the truth value of Π for each possible combination of truth values of the elementary propositions which are part of Π.[14]

A proposition of the form ~Π has only one elementary proposition, which is Π. This proposition can only have two possible truth values, namely true and false. The truth table for the proposition ~Π therefore contains only two rows with truth values. It looks like this:

1	**Π**	**~Π**
2	true	false
3	false	true

In this truth table, the first column contains the numbers of the rows.

The truth table must be read as follows: if the proposition Π is true, the proposition ~Π is false (row 2) and if the proposition Π is false, the proposition ~Π must be true (row 3).

By means of a truth table it is easy to show how adding a double negation to a proposition results in a truth value equal to that of the original proposition:

1	**Π**	**~Π**	**~~Π**
2	true	false	true
3	false	true	false

The truth values in the fourth column are 'calculated' by applying the definition of negation to the truth values in the third column. The truth values in the second and fourth column are identical, which shows that the truth value of a double negation is always identical to that of the original (non-negated) proposition. *Always*, since it is irrelevant what the content of proposition Π is; the rule applies to all propositions.

3.4 The &- operator

In informal language, the symbol & stands for the word 'and'. The operator is usually called a 'conjunction'. Formally, it is an operator which combines two propositions, elementary or compound, into a new compound proposition. That compound proposition is true iff both combined propositions – the two 'conjuncts' – are true.

TRUTH TABLE

The meaning of the &-operator can again be illustrated by means of a truth table.

A compound proposition of the form Π&Θ has two elementary propositions, Π and Θ. Each of these two propositions can have two possible truth values, true and false. Moreover, these truth values are independent of each other. This means that there are four possible combinations; hence, the truth tables must contain 4+1=5 rows. It should look like this:

[14] Greek capitals, such as Α, Β, Γ, Ο, Π, and Θ are used to signify propositions whose content is irrelevant. Different capitals stand for different propositions.

1	**Π**	**Θ**	**Π&Θ**
2	true	true	true
3	true	false	false
4	false	true	false
5	false	false	false

The table must be read as follows: If Π is true and Θ is also true, then so is the compound proposition Π&Θ (row 2). If Π is false, the proposition Π&Θ is also false (rows 4 and 5). In addition, Π&Θ is false if Θ is false (rows 3 and 5).

It is, again, important to realise that the truth value of Π&Θ depends exclusively on the truth value of the two elementary propositions of which it is a combination, Π and Θ. What Π and Θ stand for, what they mean in ordinary language, is completely irrelevant.

There also does not need to be any material connection between Π and Θ. An extreme example may help to understand this: if Θ stand for the proposition 'Brussels is the capital of Portugal' and Π stands for 'French fries are edible', the compound proposition Π&Θ, rendered in formal language, is true if, and only if, Brussels is the capital of Portugal and French fries are indeed edible.

THE RELATIONSHIP WITH ORDINARY LANGUAGE

The meaning of the &-operator is fully expressed in the truth table. A formal sentence of the form Π&Θ means nothing more and nothing less than that Π and Θ are both true. The word 'and' in ordinary language sometimes means more than that. Consider, for instance, the sentence 'John set the alarm clock and went to sleep'. In ordinary language and in an everyday context, this means that John first set the alarm clock and then went to sleep, and moreover that little time elapsed between the two events. In translating this sentence into the compound proposition Π&Θ, this extra information is lost. But since such implicit extra information is often irrelevant for the validity of an argument, this loss is acceptable. Sometimes, however, the implicit extra information *is* relevant, which may lead to the problem that an argument, translated into formal language, may *seem* valid or invalid while it actually is not. It is therefore good to be aware of this possibility and to watch out for it when formalising arguments rendered in ordinary language. An example of an argument which is valid but whose validity is no longer visible in the formalised version is the following:

John set the alarm clock and went to sleep.
He set the alarm clock at twelve o'clock.

John only went to bed after twelve o'clock.

3.5 The ∨-operator

Informally, the ∨-operator stands for the word 'or'; it is usually called 'disjunction'. In formal language, it is an operator which combines two propositions, elementary or compound, into one compound proposition. The compound proposition is true iff at least one of the of the component propositions, the 'disjuncts', is true.

TRUTH TABLE

The meaning of the ∨-operator can again best be illustrated by means of a truth table.

A proposition of the form $\Pi \vee \Theta$ has two elementary propositions, Π and Θ. Each of these two propositions can have two truth values, true and false. These truth values are independent of each other, so that there are four possible combinations. The truth table looks as follows:

1	Π	Θ	$\Pi \vee \Theta$
2	true	true	true
3	true	false	true
4	false	true	true
5	false	false	false

The table above must be read as follows: if at least one of the propositions Π and Θ is true, the proposition $\Pi \vee \Theta$ is also true (rows 2-4) and only if Π and Θ are both false (row 5), the compound proposition $\Pi \vee \Theta$ is also false.

It is again important to realise that the truth value of the compound proposition $\Pi \vee \Theta$ is determined exclusively by the individual truth values of the component propositions Π and Θ. It is completely irrelevant what Π and Θ stand for; only their truth value is important. An extreme example, similar (but *not* identical, due to the different operator) to the one we used for the &-operator, illustrates this: if Π stands for 'Brussels is the capital of Portugal' and Θ stands for 'French fries are edible', then still $\Pi \vee \Theta$ is only false if Brussels is not the capital of Portugal and French fries are also not edible.

EXCLUSIVE DISJUNCTION

The meaning of the $\vee$-operator is fully expressed in the truth table. A formal sentence of the form $\Pi \vee \Theta$ means nothing more and nothing less than at least one of the propositions Π and Θ is true. This is called the 'inclusive or'.

In ordinary language, the word 'or' is sometimes taken to mean that exactly one of the two disjuncts is true, and not both. That is the 'exclusive or'.

In propositional logic the $\vee$-operator stands for the inclusive or and there is no specific operator to signify the 'exclusive or'. However, with the help of the three operators $\vee$, & and ~, the 'exclusive or' can be simulated. Assume that J stands for the proposition 'John has the ball' and P stands for 'Paul has the ball'. Assume further that John and Paul cannot both have the ball. In that case, the sentence 'John has the ball or Paul has the ball' contains an 'exclusive or'. Tis sentence could then be formalised as $(J \vee P) \& \sim(J \& P)$, which means that at least one of the two, John or Paul, has the ball ($J \vee P$; inclusive or) and that not both of them have the ball: $\sim(J \& P)$.

The following complex truth table illustrates how truth tables can be used in order to check whether an operator is defined well. A compound proposition which expresses an 'exclusive or' must be true if precisely one of the disjuncts is true. Otherwise, i.e. if either both of them are true or both of them are false, the compound proposition is false. For the truth table, this means that in the column which expresses the exclusive disjunction we must find the value 'true' in the rows in which precisely one of the disjuncts is true, and the value 'false' in the rows in which the disjuncts are either both true or both false.

1	2	3	4	5	6	7
2	Π	Θ	$\Pi \vee \Theta$	$\Pi \& \Theta$	$\sim(\Pi \& \Theta)$	$(\Pi \vee \Theta) \& \sim(\Pi \& \Theta)$
3	true	true	true	true	false	false
4	true	false	true	false	true	true
5	false	true	true	false	true	true
6	false	false	false	false	true	false

In order to make it easier to explain, the table above contains an extra row (the first one) with the numbers of the columns.

Columns 2 and 3 show the four possible combinations of truth values of the propositions Π and Θ. Columns 4 and 5 display the truth values of the disjunction and the conjunction of these two propositions, respectively. Column 6 shows the opposite values of column 5, since the proposition represented in it constitutes the negation of the proposition represented in column 5. Column 7 indicates the truth value of the conjunction of the propositions in columns 4 and 6. This last column stands for the exclusive conjunction and – as could be expected – provides the value 'true' in the rows which contain the value 'true' in exactly one of the columns 2 and 3, thus in row 4 and 5.

> A truth table, such as the one above, may initially not be easy to grasp. But learning how to work with truth tables is worth some effort and time, since once you do understand them, you will have understood much of propositional logic.

3.6 The $\supset$-operator

Informally, the $\supset$-operator stand for the 'if … then'-combination. The operator is usually called 'material implication'. In formal language, it is an operator which combines two propositions, elementary or compound, into one compound proposition. This compound proposition is true iff the first component proposition (the '*antecedens*') is false, or the second component proposition (the '*consequens*') is true, or both.

TRUTH TABLE

The meaning of the $\supset$-operator can again be illustrated by means of a truth table.

A proposition of the form $\Pi \supset \Theta$ has two elementary propositions, Π and Θ. Each of these two propositions can have two truth values, true and false. The truth table looks like this:

1	Π	Θ	$\Pi \supset \Theta$
2	true	true	true
3	true	false	false
4	false	true	true
5	false	false	true

The table must be read as follows:

- If Π is false, the proposition $\Pi \supset \Theta$ is true (rows 4 and 5).
- If Θ is true, the proposition $\Pi \supset \Theta$ is also true (rows 2 and 4).
- This entails that the proposition $\Pi \supset \Theta$ is only false if Π is true and Θ is false (row 3)

It follows from this definition that the proposition $\Pi \supset \Theta$ means the same as the proposition $\sim\Pi \vee \Theta$. This can also be shown by means of a truth table. See exercise 3 after section 4.

Once more it is important to realise that the truth value of the compound proposition $\Pi \vee \Theta$ is determined exclusively by the individual truth values of the component propositions Π and Θ. It is completely irrelevant what Π and Θ stand for; only their truth value is important. There does not need to be a material connection between Π and Θ. To illustrate this, let us return to the extreme example of Brussels and French fries: if Π stands for 'Brussels is the capital of Portugal' and Θ stands for 'French fries are edible', then $\Pi \supset \Theta$ is only true if Brussels is not the capital of Portugal or (inclusive!) if French fries are edible.

THE RELATIONSHIP WITH ORDINARY LANGUAGE

The meaning of the $\supset$-operator is fully expressed in the truth table. A sentence of the form $\Pi \supset \Theta$ does not express anything more or less than that either Π is false, or Θ is true (or both). In particular, it does *not* express that the truth of Π is relevant for the truth of Θ, or that Π constitutes a reason for Θ. The meaning of the material implication does not correspond to the meaning of 'if … then'-constructions in ordinary language. But they do have something in common in that they make it possible to deduce the truth of the *consequens* from the truth of the *antecedens*, to make deductions of the types Modus Ponens and Modus Tollens.

We will return to this in section 4.3. There we will also see that it is possible to deduce the truth of a material implication from the falsity of the *antecedens* and from the truth of the *consequens*. See also section 2.4.

3.7 The $\equiv$-operator

Informally, the $\equiv$-operator stands for the 'if, and only if' (iff) combination. The operator is usually called 'equivalence'. In formal language, it is an operator which combines two propositions, elementary or compound, into one compound proposition. The compound proposition is true iff the two component propositions are either both true or both false. Therefore, their truth values must be equal, hence the name 'equivalence'. If the two component propositions do not have the same truth value, the equivalence is false.

TRUTH TABLE

The meaning of the $\equiv$-operator can again be illustrated by means of a truth table.

A proposition of the form $\Pi \equiv \Theta$ has two elementary propositions, Π and Θ. Each of these two propositions can have two truth values, true and false, which are independent of each other. The truth table looks as follows:

1	**Π**	**Θ**	**$\Pi \equiv \Theta$**
2	true	true	true
3	true	false	false
4	false	true	false
5	false	false	true

If Π and Θ are equivalent to each other, this means that Π materially implies Θ and vice versa. That this is the case can be seen in the following truth table:

1	2	3	4	5	6	7
2	Π	Θ	$\Pi \supset \Theta$	$\Theta \supset \Pi$	$(\Pi \supset \Theta)\&(\Theta \supset \Pi)$	$\Pi \equiv \Theta$
3	true	true	true	true	true	true
4	true	false	false	true	false	false
5	false	true	true	false	false	false
6	false	false	true	true	true	true

Column 4 indicates the truth value for the material implication $\Pi\supset\Theta$ and column 5 indicates the truth value for the material implication $\Theta\supset\Pi$. Column 6 indicates the truth values of the conjunction of both implications, with the same result as column 7, which indicates the truth values of the equivalence.

The truth table for the equivalence yields precisely the opposite outcome of that of the exclusive disjunction:

1	2	3	4	5
2	Π	Θ	$(\Pi \vee \Theta)\&\sim(\Pi\&\Theta)$	$\Pi \equiv \Theta$
3	true	true	false	true
4	true	false	true	false
5	false	true	true	false
6	false	false	false	true

This makes it possible to provide an easier definition of the exclusive disjunction, namely to define it as the negation of equivalence: $\sim\Pi\equiv\Theta)$

3.8 Necessary and sufficient conditions

Like the other operators discussed above, equivalence is purely defined in terms of the truth values of the component propositions. It says nothing about any material link between those propositions. But while we bear this in mind, it is useful to devote some attention to the way in which material implication and equivalence can be used to formalise necessary and sufficient conditions.

SUFFICIENT CONDITIONS

The event or fact A is a sufficient condition for event or fact B if whenever A happens or is the case, B also happens or is the case.

For instance, the fact that C is a circle is a sufficient condition for the fact that C is round. This is an example of a *timeless* sufficient condition.

An example in which the sufficient condition precedes that for which it is a condition in time is the breaking of the dam of a reservoir lake, which is a sufficient condition for the flooding of the land downstream.

From a logical point of view it is not necessary that a necessary condition precedes that for which it is a condition in time. A sufficient condition does not have to be a *cause*. For example, the fact that a window is broken is a sufficient condition for the breakability of a window – after all, if something is broken it must have been breakable.

A sufficient condition can be formalised by means of a material implication in which the sufficient conditions operates as the *antecedens* and that for which it is a condition operates as the *consequens*.

> If B stands for 'the dam breaks' and O stands for 'the land downstream is flooded', the B ⊃ O expresses that the breaking of the dam is a sufficient condition for the flooding.
>
> If R stands for the fact that something is breakable and B stands for the fact that it is broken, then B⊃ R expresses that the fact of being broken is a sufficient condition for being breakable.

If A is a sufficient condition for B, then the implication A⊃B is true. That this is the effect can also be shown by means of the truth table for this formula:

1	**A**	**B**	**A ⊃ B**
2	true	true	true
3	true	false	false
4	false	true	true
5	false	false	true

If the formula is true, that is, if we are looking at one of the rows 2, 4 or 5, then B is always true if A is also true. We see this in row 2, in which A is true and B is true as well. (In row 3, A is also true, but A⊃B is false and therefore we have to leave this row out of consideration.)

NECESSARY CONDITIONS

The event or fact A is a necessary condition for event or fact B, if B only happens or is the case if A also happens or is the case.

For instance, the fact that P has the right to dispose of a house is – exceptions aside – a necessary condition for the transfer of ownership of the house by P to another person. The fact that the criminal law defines the type of act H as punishable is a necessary condition for punishing someone because he or she has committed an act of type H.

> The former are both examples of timeless necessary conditions. An examples in which the necessary condition precedes that for which it is a condition in time is that a train must have stopped before it can be boarded (in a decent way).

Also necessary conditions can easily be formalised by means of a material implication, although the implication will then point in the 'wrong' direction. If A is a necessary condition for B, then the following implication is true: $B \supset A$. That this is the case can be illustrated by means of a truth table:

1	B	A	B ⊃ A
2	true	true	true
3	true	false	false
4	false	true	true
5	false	false	true

If the formula is true, that is, if we are looking at one of the rows 2, 4 or 5, then B is only true (in row 2) if A is also true.

THE RELATIONSHIP BETWEEN NECESSARY AND SUFFICIENT CONDITIONS

If a fact F is a *necessary* condition for a fact G, then, conversely, G is a *sufficient* condition for F.

We can easily see this in our earlier example of breaking and breakability. The fact that O is broken is sufficient condition for the breakability of O. Conversely, this means that the breakability of O – the *possibility* of O to break – is a necessary condition for the breaking of O.

Since the right to dispose of a house is a necessary condition for the transfer of ownership of that house, that transfer of ownership is a sufficient condition for the right to dispose. (Be careful: if something is a necessary or sufficient condition for something else, this does not *per se* imply a causal link, a legal consequence, or a certain sequence in time.)

CONDITIONS WHICH ARE NECESSARY AND SUFFICIENT FOR EACH OTHER

Sometimes, two facts go together in the sense that, if one is the case, so is the other, and vice versa. Where this is so, each of the two facts is both a necessary and a sufficient condition for the other. (If this is not immediately clear to you, try to find your own example.)

> An example is the relationship between the fact that Barack Obama is the President of the United States and the fact that Barack Obama is the commander-in-chief of the US armed forces. From the one fact, the other can be deduced, and vice versa.

The propositions which express these facts are therefore equivalent. If K stands for the fact that Barack Obama is President of the United States and O stands for the fact that Barack Obama is the commander-in-chief of the US armed forces, then the equivalence of these two propositions can be expressed by means of the formula $K \equiv O$.

4. LOGICAL VALIDITY

How can propositional logic help us determine whether an argument is valid? The first step in determining this is to realise under which circumstances an argument is logically valid. This is the case where it is *logically impossible* that all premises of the argument are true while its conclusion is false.

But what does 'logically impossible' mean? This question can be answered in terms of propositional logic: something is logically impossible if it is not the case, regardless of what the concrete facts are. There is no combination of one or more facts that makes a proposition true which expresses something that is logically impossible. This means that a compound proposition which expresses something impossible is false in all rows of the truth table. Such a proposition is a contradiction (see section 4.2).

With regard to the validity of an argument, this means that there is no row in a truth table which lists all premises of a valid argument as true but the conclusion as false.

Before we go on to show that this characterisation of logical validity is correct, it is useful first to consider tautologies and contradictions.

4.1 Tautologies

A tautology is a statement which is true, regardless of the concrete facts. This may seem odd at first glance, but some examples will show that tautologies can indeed exist.

One such example is the proposition that bachelors are unmarried. This is by definition the case, and there can be no facts or circumstances which render the proposition 'Bachelors are unmarried' false.

Another example of a tautology is the proposition 'It is raining or it is not raining'. Either it is raining, in which case the proposition is true. Or it is not raining, in which case it is also true. Therefore, the proposition is true regardless of what the facts are.

What does the truth table of a tautology look like? It is difficult to draw up a truth table for the proposition that bachelors are unmarried in the language of propositional logic, therefore we use the proposition 'It is raining or it is not raining' as an example. Assume that R stands for 'It is raining'. The truth table then looks as follows:

1	**R**	**~R**	**R ∨ ~R**
2	true	false	true
3	false	true	true

Rows 2 and 3 represent all possible combinations of facts which are relevant to the truth value of the compound proposition R∨~R. (There are only two possibilities, since there is only a single elementary proposition.) In both cases, the compound proposition is true. Hence, the proposition is true regardless of the facts.

Of course there are generally more than two possible combinations of facts. For instance, it may be raining during a storm or in still air. But since the truth value of the compound proposition R∨~R depends solely on the truth value of R (and on logic) and not, for example, on any information about wind speed, only two sorts of combinations are relevant, namely combinations in which R is true and combinations in which R is false. Both of these combinations are represented in the truth table.

4.2 Contradictions

While tautologies are propositions which are true regardless of what the facts may be, contradictions are propositions which are *false* regardless of what the facts are. Easy examples are the propositions

'This is a square circle' and 'It is raining and it is not raining'. Let us have a look at the truth table for the latter proposition:

1	**R**	**~R**	**R&~R**
2	true	false	false
3	false	true	false

We now see that the compound proposition R&~R is false in all relevant combinations of facts.

4.3 Valid arguments

Let us now look at some examples of valid arguments and the corresponding truth tables.

DOUBLE NEGATION

We start out simple:

> It is raining.
> --
> It is not the case that it is not raining.

The corresponding truth table is constructed by adding one column each for all elementary propositions which feature in the argument and for all premises and for the conclusion. The validity of an argument can be seen in a truth table, namely if the *conclusion is true in all rows in which all premises are true*.

For the simple example above, drawing up the truth table is easy:

1	**R**	**~R**	**~~R**
2	true	false	true
3	false	true	false

All premises (i.e. the only one there is) are only true in row 2 and in this row the conclusion is also true. The argument is therefore valid.

MODUS PONENS

> If Jean is a thief (D), then Jean is punishable (S).
> Jean is a thief.
> ---
> Jean is punishable.

This argument takes the logical form Modus Ponens, which we have seen earlier. The corresponding truth table looks as follows:

1	**2**	**3**	**5**
2	D	S	D ⊃ S
3	true	true	true
4	true	false	false
5	false	true	true
6	false	false	true

The premises, D and D ⊃ S, are only both true in row 3. In that row, the conclusion S is also true; the argument is therefore valid.

Modus Tollens

If Jean is a thief (D), then Jean is punishable (S)
Jean is not punishable.

Jean is not a thief.

The corresponding truth table looks like this:

1	**2**	**3**	**4**	**5**	**6**
2	D	~D	S	~S	D ⊃ S
3	true	false	true	false	true
4	true	false	false	true	false
5	false	true	true	false	true
6	false	true	false	true	true

The premises, ~S and D ⊃ S are only both true in row 6. In that row, the conclusion ~D is also true. Hence, the argument is valid.

A complicated example

Finally, let us consider an example of a complicated argument which really shows that propositional logic is useful. The argument looks as follows:

If Brussels is not located in Belgium (~B) then one of the following two is the case: Brussels is not the capital of Belgium (~H), or Liège is the capital of Wallonia (L).
Brussels is not located in Belgium.
Brussels is the capital of Belgium.

Liège is the capital of Wallonia.

When we formalise this argument, it looks like this:

```
~B ⊃ (~H ∨ L)
~B
H
----------------------------------------
L
```

Let us now see what the corresponding truth table would look like. There are three elementary propositions, B, H and L, three premises and – of course – one conclusion. All of these are represented in the truth table:

1	2	3	4	5	6	7	8
2	B	~B	H	~H	L	~H ∨ L	~B ⊃ (~H ∨ L)
3	true	false	true	false	true	true	true
4	true	false	true	false	false	false	true
5	true	false	false	true	true	true	true
6	true	false	false	true	false	true	true
7	false	true	true	false	true	true	true
8	false	true	true	false	false	false	false
9	false	true	false	true	true	true	true
10	false	true	false	true	false	true	true

Only in row 7 all premises are true. In the first 4 rows, ~B is false. In the remaining rows 5, 6 and 8, H is only true in rows 7 and 8. And in these rows, ~B ⊃ (~H ∨ L) is only true in row 7. Apparently, the three premises are only simultaneously true if B is false, H is true and L is also true.

The conclusion of the argument is L and L is true if all three premises are true (row 7). Therefore, the argument is valid.

4.4 Logical laws

On the basis of truth tables, we can easily show the validity of several logical 'laws'. These laws always state that two formulae are logically the same, that they are equivalent.

$\underline{\Pi\&(\Theta \vee \Sigma) \equiv (\Pi\&\Theta) \vee (\Pi\&\Sigma)}$

The equivalence formula above looks complicated, but it becomes much easier to understand once we fill in simple propositions for the variables.

Assume that Π stands for 'Eddy has a car', Θ for 'Eddy has a horse' and Σ for 'Eddy has a donkey'. The equivalence therefore means that the proposition that Eddy has a car and that he has a horse or a donkey next to that, is equivalent to the proposition that Eddy has a car and a horse, or a car and a donkey. That this equivalence is necessarily true and therefore can be regarded as a 'law' of logic becomes evident in the following truth table:

1	2	3	4	5	6	7	8
2	Π	Θ	Σ	$\Theta \vee \Sigma$	$\Pi\&\Theta$	$\Pi\&\Sigma$	$\Pi\&(\Theta \vee \Sigma) \equiv (\Pi\&\Theta) \vee (\Pi\&\Sigma)$
3	true	true	true	true	true	true	true
4	true	true	false	true	true	false	true
5	true	false	true	true	false	true	true
6	true	false	false	false	false	false	true
7	false	true	true	true	false	false	true
8	false	true	false	true	false	false	true
9	false	false	true	true	false	false	true
10	false	false	false	false	false	false	true

The formula $\Pi\&(\Theta \vee \Sigma)$ is false in rows 7-10, in which Π is false, and moreover in row 6, in which $\Theta \vee \Sigma$ is false. The formula is therefore true in rows 3-5.

The formula $(\Pi\&\Theta) \vee (\Pi\&\Sigma)$ is true in the rows in which either $\Pi\&\Theta$ or $\Pi\&\Sigma$ (or both) is true. This is the case in rows 3-5. The formula is therefore false in rows 6-10.

It appears that the two formulae are true in exactly the same rows and that the equivalence of the two formulae is therefore a tautology. Hence, the formulae are logically the same.

<u>$\Pi \vee (\Theta\&\Sigma) \equiv (\Pi \vee \Theta)\&(\Pi \vee \Sigma)$</u>

Assume again that Π stands for 'Eddy has a car', Θ for 'Eddy has a horse' and Σ for 'Eddy has a donkey'. The equivalence above then means that that the proposition that Eddy has a car or that he has a horse as well as a donkey is the same as the proposition that Eddy has a car or a horse, and a car or a donkey. That this equivalence is necessarily true and therefore can be regarded as a 'law' of logic becomes evident in the following truth table:

1	2	3	4	5	6	7	8
2	Π	Θ	Σ	$\Theta\&\Sigma$	$\Pi \vee \Theta$	$\Pi \vee \Sigma$	$\Pi \vee (\Theta\&\Sigma) \equiv (\Pi \vee \Theta)\&(\Pi \vee \Sigma)$
3	true	true	true	true	true	true	true
4	true	true	false	false	true	true	true
5	true	false	true	false	true	true	true
6	true	false	false	false	true	true	true
7	false	true	true	true	true	true	true
8	false	true	false	false	true	false	true
9	false	false	true	false	false	true	true
10	false	false	false	false	false	false	true

$\underline{\sim(\Pi \vee \Theta) \equiv (\sim\Pi \& \sim\Theta)}$

Assume that Π stands for 'Eddy has a car' and Θ stands for 'Eddy has a horse'. This equivalence then means that the proposition 'It is not the case that Eddy has a car or a horse' is equivalent to the proposition 'Eddy has no car and also no horse'.

The truth table from which the equivalence appears looks as follows:

1	2	3	4	5	6
2	Π	Θ	$\Pi \vee \Theta$	$\sim(\Pi \vee \Theta)$	$\underline{\sim\Pi \& \sim\Theta}$
3	true	true	true	false	false
4	true	false	true	false	false
5	false	true	true	false	false
6	false	false	false	true	true

$\underline{\sim(\Pi \& \Theta) \equiv (\sim\Pi \vee \sim\Theta)}$

Assume that Π stands for 'Eddy has a car' and Θ stands for 'Eddy has a horse'. This equivalence then means that the proposition 'It is not the case that Eddy has a car as well as a horse' is equivalent to 'Eddy has no car or (inclusive 'or'!) Eddy has no horse'.

The truth table from which the equivalence appears looks as follows:

1	2	3	4	5	6
2	Π	Θ	$\Pi \& \Theta$	$\sim(\Pi \& \Theta)$	$\underline{\sim\Pi \vee \sim\Theta}$
3	true	true	true	false	false
4	true	false	false	true	true
5	false	true	false	true	true
6	false	false	false	true	true

4.5 Invalid arguments

Let us now consider two examples of invalid arguments and the corresponding truth tables. The first example is the argument:

> It is raining (R), or somebody is pouring water (G).
> Somebody is pouring water.
> ---
> It is not raining.

In formal language, this becomes:

```
R ∨ G
G
-----
~R
```

When we construct a truth table for this argument, it looks like this:

1	**2**	**3**	**4**	**5**
2	R	~R	G	R ∨ G
3	true	false	true	true
4	true	false	false	true
5	false	true	true	true
6	false	true	false	false

The two premises are both true in rows 3 and 5. In row 5, the conclusion is also true, but in row 3 it is not. If it is raining and (at the same time) somebody pours water, both premises are true but the conclusion is not. It is therefore logically possible that all premises of this argument are true while the conclusion is false. Therefore, the argument is invalid.

In section 3.1 we encountered a somewhat more complicated argument, namely:

If the Scottish Nationalist Party wins the elections, Scotland will become independent or it will be impossible to form a Scottish government.
The Scottish Nationalist Party is going to win the elections, but it will certainly be possible to form a Scottish government.

Scotland will not become independent.

We formalised this argument as follows:

W ⊃ (I ∨ ~G)
W

~I

Let us now see what the corresponding truth table would look like. There are three elementary propositions, three premises and – of course – one conclusion. The truth table therefore becomes:

1	2	3	4	5	6	7	8
2	W	I	G	~I	~G	I ∨ ~G	W ⊃ (I ∨ ~G)
3	true	true	true	false	false	true	true
4	true	true	false	false	true	true	true
5	true	false	true	true	false	false	false
6	true	false	false	true	true	true	true
7	false	true	true	false	false	true	true
8	false	true	false	false	true	true	true
9	false	false	true	true	false	false	true
10	false	false	false	true	true	true	true

W is only true in rows 3-6. In those four rows, G is only true in rows 3 and 5. In the latter two rows, W ⊃ (I ∨ ~G) is only true in row 3. But in this row, the conclusion ~I is also false. The argument is therefore invalid.

Exercises

1. Use truth tables to determine whether the following compound propositions are tautologies, contradictions, or neither:

   a. `A&(B ∨ ~B)`

   b. `~(~C)`

   c. `P&(Q&~Q)`

   d. `P ⊃ ~P`

   e. `(P&(P ⊃ Q)) ⊃ Q`

   f. `((A ∨ B)&C) ∨ ((~C ∨ ~B) ∨ ~A)`

2. Use truth tables to determine whether the following arguments are valid:

a.
```
P ∨ Q
~Q
-----
~P
```

b.
```
R ∨ S
S
-----
~R
```

c.
```
P ∨ R
----------
(~P&~R) ⊃ Q
```

d. (A&~B) ∨ C
~C&~B

C ∨ ~A

e. Either the Labour Party or the Conservative Party wins the elections.

If the Labour Party wins, taxes are raised.

If taxes are not raised, the Conservative Party therefore wins.

f. If the Belgian national team does not win and the French national team does, everyone passes the logic exam.
The Belgian national team wins, since some do not pass the logic exam.

3. Use a truth table to prove that the propositions $\Pi \supset \Theta$ and $\sim\Pi \vee \Theta$ are logically equivalent. (A hint: this will be the case if their equivalence is a tautology).

III. CLASS LOGIC

1. INTRODUCTION

Consider the following argument:

Lawyers do not like mathematics. Nerds do like mathematics. -- Lawyers are not nerds.

If we formalise this argument in the language of propositional logic, this is the result:

```
P
Q
--
R
```

It does not require a truth table to see that this argument is invalid. From a logical point of view, P, Q and R are all elementary propositions and their truth values are therefore independent of each other. It is therefore logically possible that P and Q are both true while R is false. In this light, the argument is invalid. But if we look at the original argument in ordinary language, it does seem valid. How is that possible?

It has to do with the previously mentioned fact (section 1.2) that an argument can have more than one logical form. In order to be valid, it is only necessary that an argument has one form which is valid. Where that is the case, it is irrelevant that the argument also has invalid forms. The form which the above argument has according to propositional logic is invalid, but there is also another sort of logic, predicate logic,[15] and according to this logic the argument does have a valid form. The argument is therefore valid.

It may seem useful to continue with the study of predicate logic at this point, but for the time being this would lead to far. Instead, we will focus on so-called 'class logic', a precursor of predicate logic which studies a specific sort of arguments called 'categorical syllogisms'. Categorical syllogisms are arguments which deal with classes of individuals (categories). The advantage of discussing class logic here is threefold:

1. The validity of categorical syllogism can be visualised by means of a graphic tool, the so-called 'Venn diagrams'.
2. Categorical syllogisms are very suitable for analysing the application of legal rules. The application of a legal rule to an individual case is therefore referred to as 'legal syllogism'.
3. Testing the validity of arguments in class logic is easier than in predicate logic.

This chapter will therefore discuss class logic. We start by introducing a number of basic concepts of this logic in section 2. Section 3 explains what categorical syllogisms are. In section 4 we see how Venn diagrams can be used to test the validity of these syllogisms. Lastly, we discuss a number of

[15] There are many more logics, but leave them out of consideration here.

rules by means of which the validity of categorical syllogisms can easily be assessed. In this context, we also discuss the 16 different forms of the categorical syllogism (also referred to as 'modi').

Before we continue, a word of warning is in order. Syllogistics is an ancient form of logic which was already discussed by Aristotle (384-322 BC) and further developed during the middle ages. What we are going to discuss in the following sections is not completely congruent with the traditional treatment of syllogistics but is influenced by modern predicate logic.[16] Historical correctness is not the aim of this chapter. Instead, it will content itself with explaining how to recognise categorical syllogisms and how to assess their validity.

2. CLASSES AND INDIVIDUALS

The example at the beginning of this chapter dealt with classes of persons: lawyers, people who like mathematics, and nerds. The argument was about how these classes relate to each other. The first premise stated that members of the class of lawyers do not belong to the class of those who like mathematics. The second premise stated that nerds do belong to this class. The conclusion was that lawyers are therefore not nerds.

The logic we are discussing in this chapter is class logic, thus logic in which classes play a special role. A class consists of a number of individuals which all share a certain characteristic. This characteristic defines the class. It may therefore be that being a lawyer is such a defining characteristic, or being a nerd, or having a liking for mathematics. But also cows constitute a class, just as chairs, pencils with broken tips, world travellers who have never been to Nepal or mail boxes which have been painted yesterday. In short, classes can be defined in very different ways, but they always consist of individuals that share a particular characteristic, however complicated or narrowly defined.

It is not *per se* the case that there are in fact any individuals in a class: this, for instance, of the class of winged horses, the class of square circles, or the class of generous scrooges. There are also class which by definition contain only one individual, for example the class of persons identical to Queen Elisabeth II, or the class consisting of the number three (*all* numbers three, but of course there is only one).

We will later see that it is sometimes useful to treat an individual as a class with only one element.

3. VENN DIAGRAMS

The most interesting categorical syllogisms have two premises and deal with three different classes. But there are also arguments which deal with only two classes. The study of those arguments is better suited to showing what class logic is about. They only need a single premise in order to be valid.

3.1 Euler diagrams

Consider the following argument:

Lawyers are not nerds. ------------------------------ Nerds are not lawyers.

[16] In particular, we are not going to claim that classes are not empty.

This is a valid argument. Why this is so becomes apparent once you realise that, according to the premise, the class of lawyers falls entirely outside the class of nerds. This is illustrated by the following diagram:

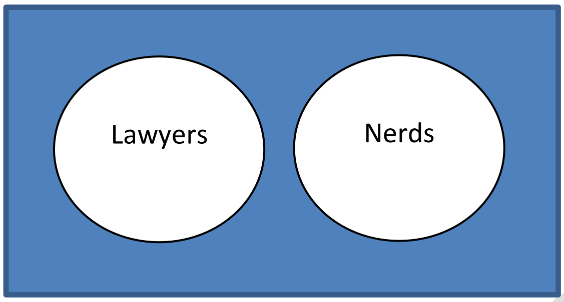

The same diagram can also be used to show that nerds are not lawyers. In fact, the propositions that lawyers are not nerds and that nerds are not lawyers are equivalent: it is impossible that one of them is true while the other one is false. An argument in which one of these two propositions is the only premise and the other one is the conclusion is therefore valid.

The diagram above, in which the circles represent the scope of the propositions, is a so-called Euler diagram, named after the famous Swiss mathematician Leonhard Euler (1707-1783).

Another example of an Euler diagram is the following one, which represents the proposition that all lawyers are intelligent beings.

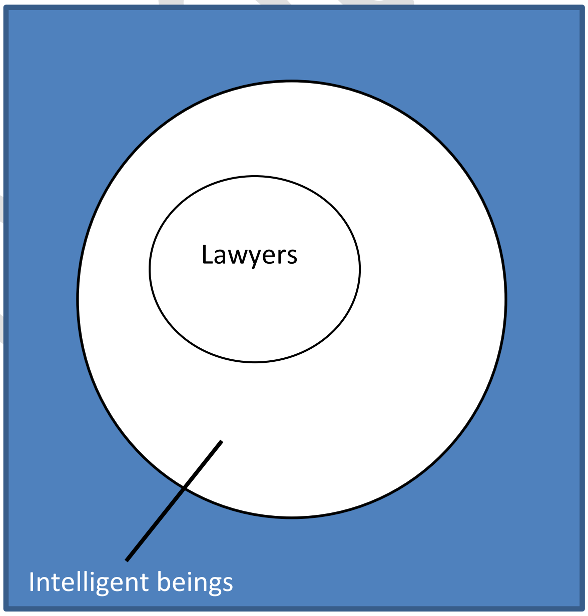

Euler diagrams are useful in that they allow us to 'read' the validity of arguments from them. However, their disadvantage is that in order to draw such a diagram one actually has to know beforehand whether or not an argument is valid.

3.2 Venn diagrams

We will therefore work with a different sort of diagrams, the so-called Venn diagrams. These diagrams are named after the English logician John Archibald Venn (1834-1923). Venn diagrams are slightly less easy to read that Euler diagrams, but their advantage is that they can be drawn even without a full grasp of an argument. Venn diagrams are therefore more suitable as a tool to assess the validity of arguments.

Venn diagrams always consist of intersecting circles. For an argument which deals with two classes, we need two circles which are drawn in the following way:

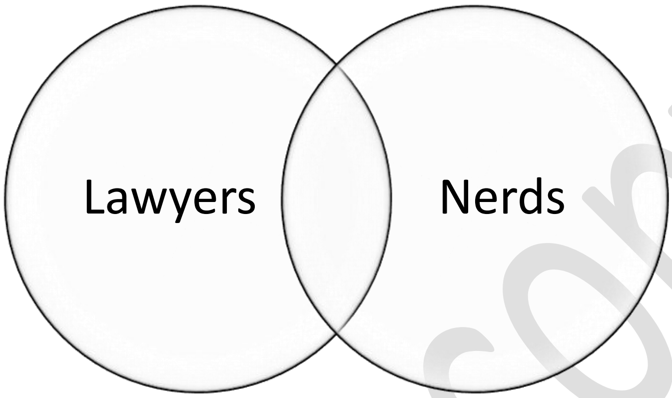

The left circle represents one class, for example that of lawyers. The right circle represents the other class, here that of nerds. The part of the diagram in which the two circles overlap stands for everything which falls into both classes, thus for all lawyers who are also nerds or – which is the same – for all nerds who are also lawyers.

3.3 Classes with an empty intersection

The first premise of our example argument states that lawyers are not nerds. Expressed in the terminology of classes, this means that the area in which the class of lawyers and the class of nerds overlap does not contain any elements. This overlapping area is called the 'intersection' of two classes. If there are no lawyers who are also nerds, this intersection is 'empty'. In a Venn diagram, this is expressed by scratching out the area which stands for the intersection, that is the area where the two circles overlap. This looks as follows:

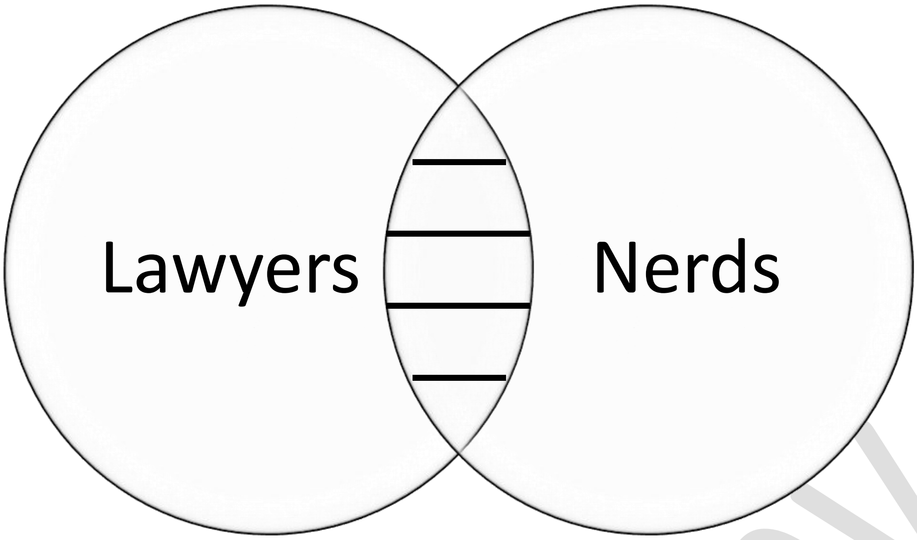

The proposition that no lawyers are nerds would result in precisely the same drawing; that already tells you that the proposition contains the same information as the proposition that lawyers are not nerds. Drawing one proposition into a Venn diagram allows you to 'read off' that another proposition is also true. Therefore it is possible to validly deduce one proposition from another.

> From the proposition that lawyers are not nerds, it can be deduced that nerds are not lawyers, and vice versa.

It is clear that this does not only count for lawyers and nerds. From the proposition that cows are not horses, it can be deduced that horses are not cows, and vice versa. The drawing would be exactly the same; only the labels 'lawyers' and 'nerds' would have to be replaced by 'cows' and 'horses'. By leaving the labels out altogether, we transform the diagram from the representation of one argument into the representation of a *form of argumentation*. The Venn diagram then shows that this form of argumentation is valid.

3.4 Classes with common elements

Consider the following argument:

Some lawyers are nerds. --------------------------------- Some nerds are lawyers.

This is also a valid argument. Why this is so becomes clear once we realise that the class of lawyers contains members (elements) which are also members of the class of nerds. That a class contains at least one element is indicated in Venn diagram by inserting an X into the circle which represents this class. Both of the following arguments indicate that at least one lawyer exists:

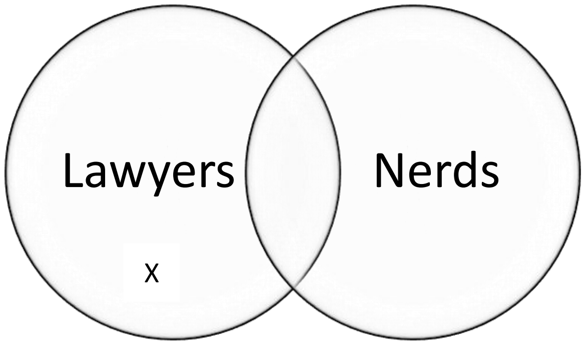

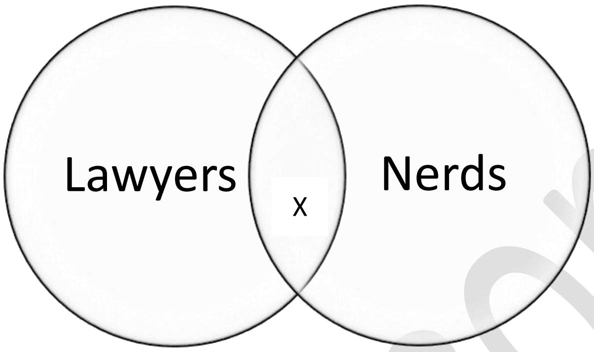

But there is a difference between what the two diagrams say. In the first diagram, the X is located within the circle of lawyers (indicating that at least one lawyer exists) but outside the circle of nerds. According to this diagram, there is therefore at least one lawyer who is not a nerd.

In the second diagram, the X is located in the intersection of the two circles, that is both in the circle of lawyers and in the circle of nerds. Hence, this diagram indicates that there is at least one lawyer who is also a nerd. But the same diagram says that there is at least one nerd who is also a lawyer. Hence, according to this diagram, both propositions mean the same and one can be validly deduced from the other.

Also here, leaving out the labels transforms the diagram from a representation of one argument into the representation of an *argumentative form*. From the fact that there is at least one X which is also Y, it can be deduced that there is at least one Y which is also X.

In fact, both of the two diagrams above contain more information than that there are lawyers. The first diagram indicates that there are lawyers who are not nerds; the second one indicates that there are lawyers who are also nerds. Can a Venn diagram also indicate that there are lawyers without simultaneously saying anything about nerds? Yes, this is possible by placing the X on the borderline of the nerd circle and within the lawyer circle. This signals that the question whether or not the existing lawyer is also a nerd.

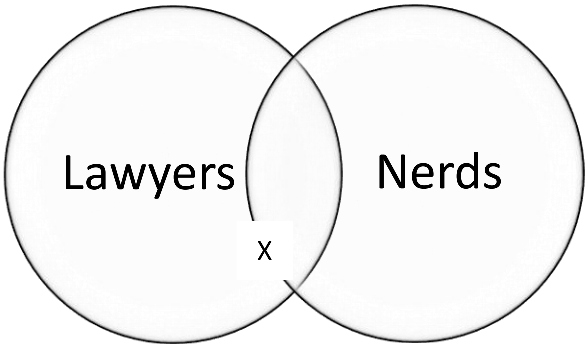

Therefore, the above diagram only shows that there are lawyers; the circle for nerds is actually superfluous and could just as well have been omitted from the diagram.

3.5 Classes which contain each other

Consider the following argument:

All thieves are criminals.

Some criminals are thieves.

Perhaps this argument seems valid at first glance, but it is not.[17] Why it is invalid can easily be illustrated by means of Venn diagrams. Let us start by drawing one for the premise:

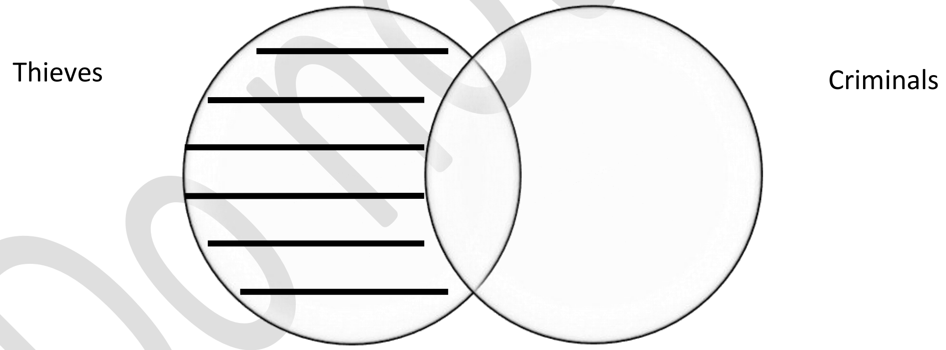

The proposition that all thieves are criminals is represented by crossing out the part of the 'thieves circle' which does not intersect the 'criminals circle'. This indicates that there are no thieves who are non-criminals. In other words, as far as there are thieves – and nothing is said about that yet – those thieves also fall within the category of 'criminals'. Remember: crossing out part of a circle means that this part of the category does not contain any elements. But that does not automatically mean, conversely, that the part of the circle which has not been crossed out does contain any elements. Only Xs indicate that (a part of) a class contains one or more elements.

On the basis of this diagram, there can therefore be two sorts of criminals, namely criminals who are thieves and criminals who are not thieves (for example drunk drivers or murderers). The former

[17] In traditional syllogistics the argument would be valid, since traditional syllogistics assumes that there are no empty classes.

category is represented by the part of the diagram in which the two circles intersect; the second by that part of the 'criminals' circle which does not fall within the 'thieves' circle. None of those two parts is crossed out, so both may possibly contain one or more elements.

However, there are no crosses in this diagram. Hence, we do not know whether the classes contain any elements. That may well be the case, but the premise does not provide us with any information. Therefore, we do not know whether there are criminal thieves, or criminals who are not thieves. In fact, we do not even know whether any criminals exist at all. The diagram does not allow us to draw any conclusions as to the existence of criminal thieves, and therefore we also cannot conclude that some criminals are thieves. The argument is therefore invalid.

We will see the latter conclusion again (in section 5.2) in the form of the rule that from universal premises, which only say something about classes as a whole, no valid conclusion can be drawn with respect to the existence of any elements in those classes.

4 ARGUMENTS WITH TWO PREMISES

4.1 Venn diagrams with three circles

Traditional categorical syllogisms are always arguments with two premises which deal with three classes. The following argument is an example:

All thieves are criminals.
All criminals are punishable.

All thieves are punishable.

This is a valid argument. In order to show that it is valid, we need a Venn diagram with three circles which represent the class of thieves, the class of criminals and the class of punishable persons, respectively. A Venn diagram representing three classes usually looks as follows:

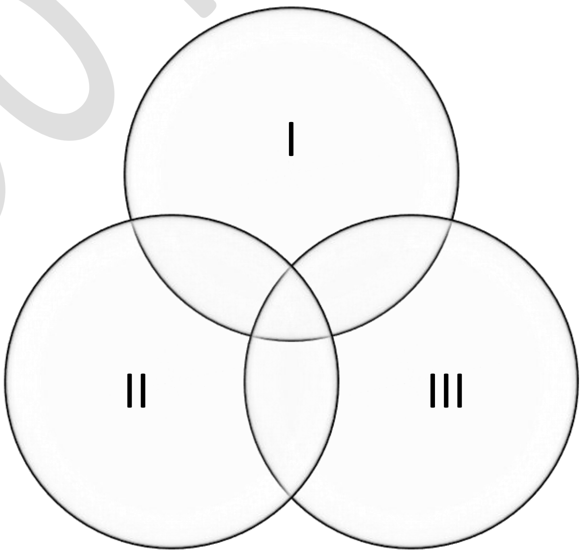

In this diagram, there are three intersecting circles. They are drawn in such a way that all forms of overlap between categories are possible:

- I and II, but not III
- I and III, but not II
- II and III, but not I and
- I, II and III.

4.2 Classes which contain each other

If we want to represent the example argument above, we can start with the representation of the first premise. In principle is does not matter in which order the premises are entered into the diagram, but for reasons we will see later (in section 4.3) it is often useful to start with those premises which say something about the relevant classes as a whole before considering the premises which deal with one or more elements of these classes.

The premise that all thieves are criminals is represented by crossing out the part of the 'thieves' circle which falls outside the 'criminals' circle. The diagram in which the first premise is represented then looks as follows:

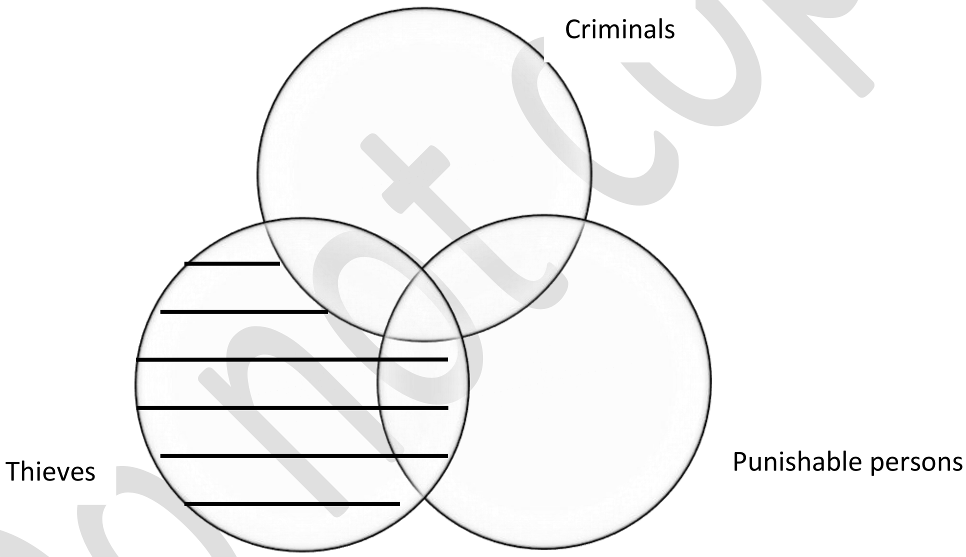

The premise that all criminals are punishable is now added by crossing out that part of the 'criminals' circle which falls outside the 'punishable persons' circle. This is the result:

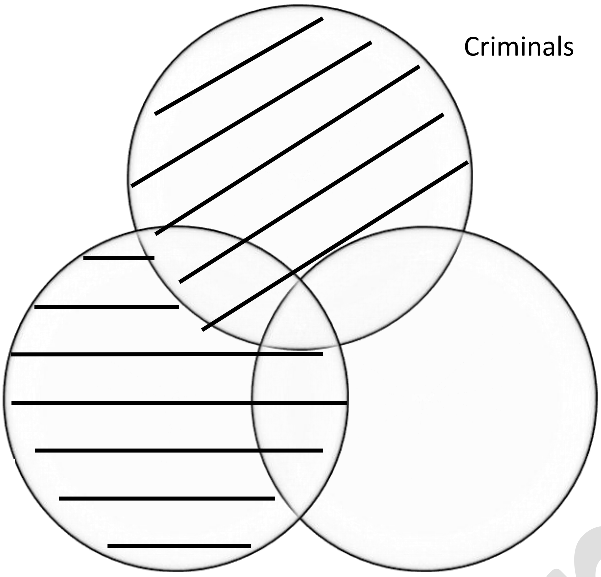

We can now see what sorts of thieves there can be, given the two premises: there can only be thieves who are also criminals and who are also punishable. After all, all parts of the 'thieves' circle which fall outside the area in which all three circles overlap have been crossed out. We can therefore validly conclude that all thieves who exist are punishable.

Whether there actually *are* any thieves is another question. We cannot draw any conclusion about this as the diagram does not contain any Xs.

4.3 Classes which exclude each other and classes with elements

Let us take a look at a new example:

No one who possesses a thing owns that thing. Some people are owners of a thing. -- Some people are not possessors of a thing.

AMBIGUITY

Before we go on to construct a Venn diagram of the argument above, there is an ambiguity in this argument that we must solve first. The second premise says that some people are owners of a thing and the conclusion says that some people are not possessors of *a* thing. The second premise and the conclusion have nothing to do with each other (and hence do not carry any implications with regard to each other) if they are about possession and ownership of *different* things. From the fact that somebody owns a car, nothing can be inferred about whether that person possesses a house. For any chance of a logical relationship between the two, they must concern the same thing. If that is the case, the conclusion should better have been:

Some people are not possessors of *that* thing.

We will assume that the conclusion indeed concerns the same thing as the second premise. In that case, the Venn diagram in which both premises are represented looks as follows:

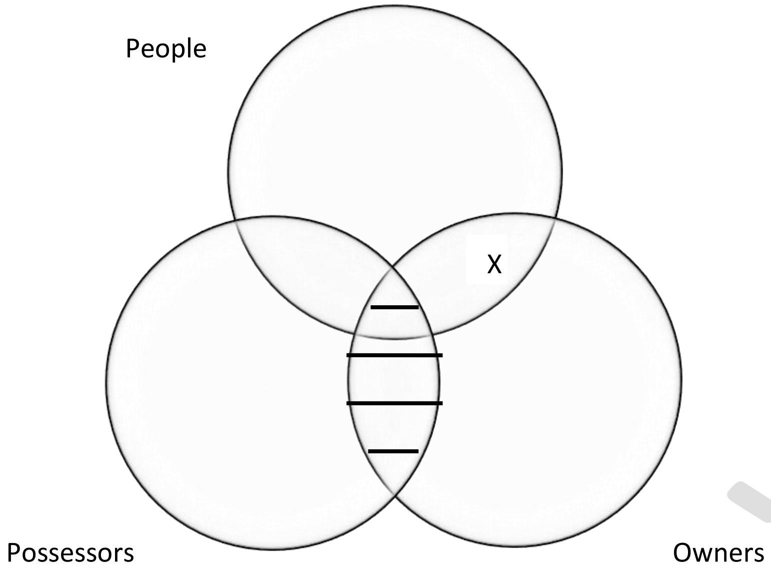

The first premise is represented by crossing out the intersection of possessors and owners. The second premise is represented by placing an X in the intersection of owners and people. Since part of the latter intersection is crossed out, we know that the X cannot be located in that part. Hence, we do not have a free choice as to where in the intersection of owners and people the X must be, we must place it in the area which falls outside the ‘possessors’ circle.

> Since crossing out part of an intersection limits our possibilities for placing the X, it is useful to code universal propositions (i.e. propositions about classes as a whole) first and specific ones (i.e. propositions that deal with one or more individuals within classes) later. We already encountered this phenomenon in section 4.2.

Once both premises are represented in the diagram, we can see that there is an X in the ‘people’ circle which is located outside the ‘possessors’ circle. Therefore, there are people who are not possessors (of *that* particular thing). Accordingly, the conclusion follows logically from the premises.

4.4 Universal, specific and concrete propositions

With regard to classes, we can distinguish at least three sorts of propositions:

1. Universal propositions, which say something about all members of a class. For instance, the proposition ‘All thieves are punishable’ says something about all thieves (but not about all punishable persons).
2. Specific propositions, which say something about some (one or more) members of a class without explicitly identifying *which* members are meant. For example, the proposition ‘Some thieves are punishable’ says something about thieves, but not about all thieves *per se*.
3. Concrete propositions, which say something about one or more members of a class and explicitly identify them. For example, the proposition ‘Luc is a thief’ says something about Luc and also about thieves, but not about all thieves *per se*.

In class logic, we actually do not have suitable means at our disposal to work with concrete propositions. We can say something about all members of a class (universal) or about some of them (specific), but not about concrete things or persons. However, depending on context, a concrete proposition can be treated as if it were a universal or specific one. For instance, a proposition about

Luc can be considered a proposition about *all persons identical to Luc*. That is only one person, but a proposition about a single person is also a proposition about all persons identical to that person (therefore, only one). This way, logic that relates to universal propositions can be applied to concrete propositions. An example will illustrate this:

All owners have the right to dispose. Elise is an owner. -- Elise has the right to dispose.

We treat Elise as if she were a class of her own, the class of all persons identical to Elise. The Venn diagram of this argument then looks as follows:

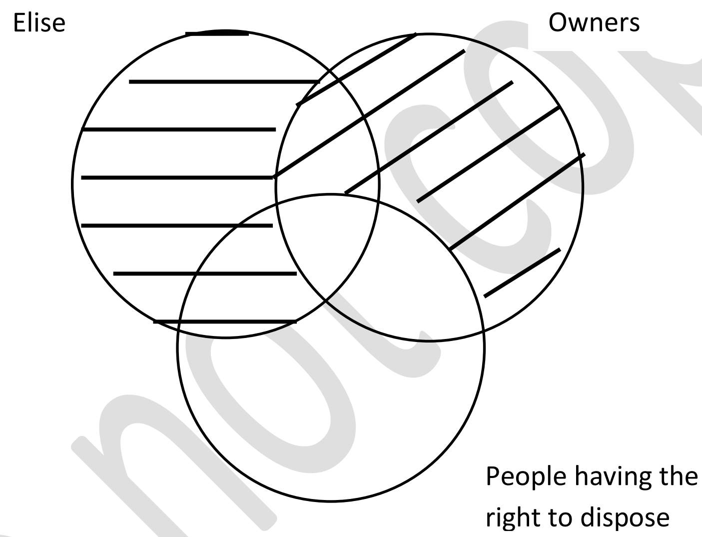

Once the argument about Elise is represented in a Venn diagram, we immediately see that it has the same form as that about thieves who are punishable. That form is valid; therefore the argument is also valid.

The example above shows that it is sometimes useful to treat a concrete proposition as if it were a universal one. In another context, however, it can be useful to treat it as if it were a specific proposition. Consider the following valid argument:

Christie is married. All married persons have a spouse. -- There are persons who have a spouse.

For the purpose of drawing a Venn diagram, we treat the proposition that Christie is married as the proposition that some persons (at least one person) are married. That makes it possible to conclude that some have a spouse. We must put an X in the ‘married persons’ circle, but so far it is not yet clear whether or not it must be located in the intersection. In order to find out, we first code the second premise by crossing out the part of the ‘married persons’ circle which does not overlap with the ‘spouses’ circle. This reveals that the X must be placed in the part of the ‘married persons’ circle which does overlap with the ‘spouses’ circle.

Once again we see why it is useful to represent universal propositions first, before entering specific ones.

The fact that the X must be placed in the intersection, the area in which the ‘married persons’ circle overlaps with the ‘spouses’ circle, immediately shows that the conclusion follows from the premises: there is at least one person with a spouse (namely Christie).

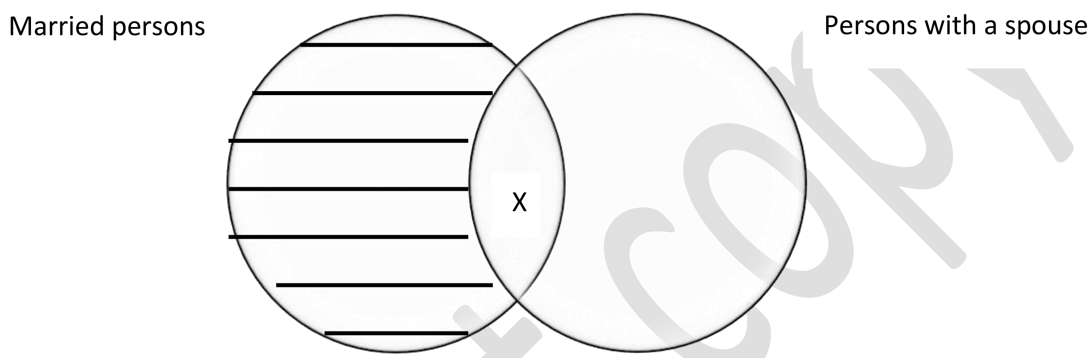

By the way, this is an example of an argument with two premises which deals with only two classes instead of the usual three.

5 RULES FOR THE VALIDITY OF SYLLOGISMS

It is possible to assess the validity of all arguments based on classes and their relationships with each other by means of Venn diagrams. But often it is even easier than that. There is a set of rules by means of which the validity categorical syllogisms can be tested quickly. Moreover, knowing these rules makes it easier to fully understand these syllogisms. In order to understand these rules of validity, several concepts must be known. The following section introduces them.

5.1 The basic structure of categorical syllogisms

The syllogisms we are going to discuss in the following are all arguments with two premises that deal with three classes. The conclusion is a proposition about the relationship of two of these classes with each other; in each of the two premises, one of the classes is connected with a third class, which functions as a nexus or connecting element between the two others. To illustrate this scheme, let us consider the following argument once more:

All thieves are criminals.
All criminals are punishable.

All thieves are punishable.

In the first premise, the class of the thieves is linked to the 'connecting class' of criminals. In the second premise, this connecting class is linked with the class of punishable persons. The conclusion no longer includes the connecting class of criminals and the classes of thieves and punishable persons are linked directly to each other.

The relationship between the two classes which feature in any of the three propositions can take eight different forms. In order to understand them, we should start with an analysis of the conclusion.

SUBJECT TERM AND PREDICATE TERM

In the conclusion of a categorical syllogism, two classes are linked with each other. These classes are indicated by two terms; in our example these are the terms 'thieves' and 'punishable', which stand for the classes of thieves and punishable persons, respectively. Conventionally, the first term ('thieves') is called the *subject term*. The second term is called *predicate term* (the word 'predicate' means the characteristic that is attributed to the subject[18]).

CONFIRMATION AND NEGATION

The linkage between the two classes described in a proposition can be positive or negative. Examples of a positive linkage are:

Louise is a poisoner.
Some possessors are owners.
All police are agents of the executive.

Examples of a negative linkage are:

John does not have a single book.
Some criminals are insane.
No politician always speaks the truth.

UNIVERSALITY AND SPECIFICITY

The propositions in categorical syllogisms are always universal or specific. Strictly speaking, concrete propositions are impossible in categorical syllogisms, but we have already seen how they can be simulated by means of 'virtual' universal or specific propositions.

Examples of universal propositions are:

All police are agents of the executive.
No politician always speaks the truth.

Examples of specific propositions are:

[18] Not in all conclusions a characteristic is attributed to the subject, hence the name is somewhat misleading.

Some possessors are owners.
Some criminals are insane.

Propositions which are both universal and confirmatory (positive), such as ‘All thieves are punishable’, are called a-propositions.

Propositions which are both specific and confirmatory, such are ‘Some owners possess their goods’, are called i-propositions.

Propositions which are both universal and negatory (negative), such as ‘Not a single head of state is also a member of the police force’, are called e-propositions.

Propositions which are both specific and negatory, such as ‘Some drivers do not have a licence’, are called o-propositions.[19]

In the conclusion of a syllogism, the subject term and the predicate term by definition appear in a fixed order: the subject term comes first, the predicate term last. In the two premises, the order of appearance is not fixed. Some premises mention the subject term or the predicate term first, then the middle term (which signifies the connecting class); in other premises the middle term is mentioned first.

Conventionally, the subject term is indicated by means of a capital S, the predicate term by means of a capital P and the middle term by means of a capital M.

MODI (FORMS) OF THE SYLLOGISM

In accordance with the above, the conclusion of a categorical syllogism has one of the following four forms:

SaP (All S are P)

SiP (Some S are P)

SeP (Not a single S is P)

SoP (Some S are not P)

Both of the premises can have eight different forms each, since the sequence of the S- and P-terms and the M-term can differ (which is not possible in the conclusion, where the S-term always comes first). The eight possible forms for premises are the following:

	universal	**specific**
confirmatory	SaM PaM	SiM PiM
	MaS MaP	MiS MiP
negatory	SeM PeM	SoM PoM
	MeS MeP	MoS MoP

[19] The letters a, i, e and o stem from the Latin words ‘*affirmo*’ (I confirm) and ‘*nego*’ (I negate/deny).

Since both premises can take eight different forms and the conclusion can take four different forms, the syllogism as a whole can take 8 x 8 x 4 = 256 different forms (modi). Of these 256, only a small portion is logically valid.

THE STANDARD FORM OF THE SYLLOGISM

There is a conventional standard of formalisation for categorical syllogisms. The conclusion always takes the form S*x*P, in which the *x* is replaced by one of the letters a, e, i or o.

The first premise (also called the '*maior*') conventionally always contains the predicate term and the middle term. The second premise (also called the '*minor*') then contains the subject term and the middle term.

Exercise

1. Formalise the following arguments in the standard form for categorical syllogisms. For example:

All owners have the right to dispose.
Elise is an owner.

Elise has the right to dispose.

S=Elise; P=has right to dispose; M=owner.
The standard form of this argument is: MaP, SaM, therefore SaP.

a. All humans are mortal.
Soldiers are humans.

Soldiers are mortal.

b. There are no superfluous human beings.
Jean-Pierre is a human being.

Jean-Pierre is not superfluous.

c. Some chickens are wheelbarrows.
There is no wheelbarrow which is yellow.

There is no yellow chicken.

d. Not a single fish can fly. Therefore, all flying animals are edible, since fish are edible.

e. Professors don't know the law, since only decent people know the law and professors are decent people.

2. For each of the arguments in exercise 1, use a Venn diagram to assess its validity.

5.2 Rules for valid syllogisms[20]

The concepts that have been introduced in the foregoing section, together with some that will follow in the course of the present one, make it possible to formulate a set of simple rules by means of which the validity of syllogism can be assessed with relative ease.

These rules only apply to:

- syllogisms
- with two premises, whereby
- one premise links the subject of the conclusion to the connecting class (signified by the middle term),
- the other premise links the predicate of the conclusion to the connecting class, and
- all terms are used in one sense only (no ambiguities).

The following is an example in which one term is used ambiguously:

All laws which apply in England have received royal assent. The law of gravity applies in England. --- The law of gravity has received royal assent.

This argument is invalid because the term 'laws' is used ambiguous: in the first premise, it refers to legal rules, while in the second it means laws of nature.

RULE 1: AT LEAST ONE OF THE PREMISES MUST BE CONFIRMATORY

Validity of a categorical syllogism is based on the link between the classes in the conclusion that is constructed by means of the connecting class. If both premises are negatory (e- or o-premises), this merely expresses that both the subject class and the predicate class wholly or partially fall outside the connecting class. However, this does not say anything about the relationship between the subject class and the predicate class, so that from two negatory premises no valid conclusion can be drawn with respect to that relationship.

[20] The rules that are discussed in this section could be formulated thanks to the assistance of Elena Dauzacker.

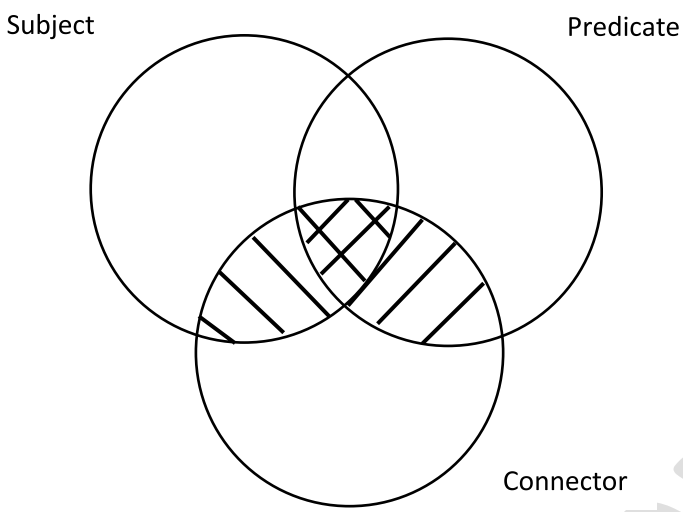

As the above Venn diagram shows, information about the *lack* of overlap between, on the one hand, the subject class and the connecting class and, on the other hand, the predicate class and the connecting class says nothing (in a positive or negative sense) about any possible overlap between the subject class and the predicate class. The diagram only relates to e-premises, but for o-premises the result is essentially the same.

Therefore each categorical syllogism with two negatory premises is invalid.

RULE 2: IN SYLLOGISMS WITH A NEGATORY CONCLUSION, EXACTLY ONE OF THE PREMISES MUST BE NEGATORY

If the conclusion is negatory, it says that the subject class and the predicate class do not overlap at all (an e-conclusion), or that some elements of the subject class are not part of the predicate class (o-conclusion). If both premises of the syllogism are confirmatory, they will – at best – provide information about the existence of (partial) overlap between the subject class and the predicate class.[21] A conclusion about the *lack* of such overlap is never sufficiently supported by such premises. If, however, *both* premises are negatory, nothing follows as all (rule 1). Therefore:

A categorical syllogism with a negatory conclusion can only be valid if exactly one of the premises is negatory.

By the way, this is a necessary condition for validity, but not a sufficient one.

RULE 3: AT LEAST ONE OF THE PREMISES MUST BE UNIVERSAL

Validity of a categorical syllogism is based on the link between the classes in the conclusion that is constructed by means of the connecting class. If both premises are specific (i- or o-premises), they merely say that both the subject class and the predicate class *partially* overlap with the connecting

[21] However, not even this is certain, since it is possible that both premises are specifc. See rule 3.

class. This does, however, not guarantee that the subject class and the predicate class overlap with *the same part* of the connecting lass. A Venn diagram illustrates this:

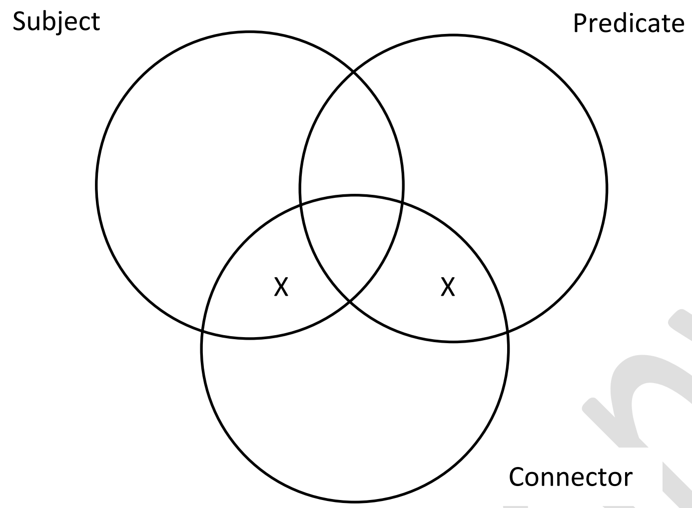

This diagram shows that both the subject class and the predicate class has some overlap with the connecting class, but since this overlap relates to different parts of the connecting class, there is no information about possible overlap between the subject class and the predicate class themselves.

With regard to rule 3 it is important that concrete premises can be rendered as (virtual) universal premises. If at least one premise is concrete, all requirements of this rule are fulfilled.

RULE 4: IN SYLLOGISMS WITH A SPECIFIC CONCLUSION, EXACTLY ONE OF THE PREMISES MUST BE SPECIFIC

If the conclusion is specific, it either says that the subject class and the predicate class share one or more elements (an i-conclusion), or that some elements of the subject class are not part of the predicate class (o-conclusion). If both premises are universal, there can (at best) only be information about the *possibility* of overlap between the subject class and the predicate class, but not about the existence of any elements within that potential overlap.

In terms of Venn diagrams, this can be expressed as follows: universal premises merely provide information about which parts of the circles must be crossed out, but not where Xs must be placed. A specific conclusion, however, tells us precisely that. It says where an X must be placed in the Venn diagram.

Therefore, in order to be able to validly draw a specific conclusion, at least one of the premises must be specific. However, if *both* premises are specific, nothing follows from them (rule 3). Therefore:

> *A categorical syllogism with a specific conclusion can only be valid if exactly one of its premises is specific.*

Again, this is a necessary condition, but not a sufficient one.

RULE 5: FOR A UNIVERSAL CONCLUSION, BOTH PREMISES MUST BE UNIVERSAL

A universal conclusion states that there is a universal connection between the S- and the P-class; it does not say anything about specific elements of these classes. For example, if a conclusion reads that all judges are lawyers, it does not say anything about individual persons. A premise that does say something about individual elements of a class, for instance that there are one or more lawyers or judges, does not contribute to such a universal conclusion. All information about the universal conclusion must then stem from the other premise. However, if that is the case, it would have been possible to draw the conclusion from that latter premise alone, and in that case we are not speaking anymore of a categorical syllogism in the sense discussed here. In other words, if a categorical syllogism is to lead to an universal conclusion, it must have two universal premises.

RULE 6: FOR A CONFIRMATORY CONCLUSION, BOTH PREMISES MUST BE CONFIRMATORY

A confirmatory conclusion indicates that all or some elements of one class are also elements of another class. For example it contains the information that all, or some of, the students stem from Germany. 'Negative' information, for instance that some students do not stem from the Netherlands, cannot contribute to such a conclusion. All information about the confirmatory conclusion must then stem from the other premise. However, if that is the case, it would have been possible to draw the conclusion from that latter premise alone, and in that case we are not speaking anymore of a categorical syllogism in the sense discussed here. In other words, if a categorical syllogism is to lead to an confirmatory conclusion, it must have two confirmatory premises.

RULE 7: THE MIDDLE TERM MUST BE DISTRIBUTED IN PRECISELY ONE OF THE PREMISES

We have already devoted some attention to the crucial role of the connecting class in the context of rules 1 and 3. If both premises merely say that (part of) the subject class and the predicate class fall outside the connecting class (two negatory premises), of if they both merely say that the subject class and the predicate class share or do not share some elements with the connecting class (two specific premises), then the argument is invalid. In either case, the reason for this invalidity is that, on the basis of the information provided by the premises, the connecting class is insufficiently linked with the subject class and the predicate class to be able to connect these two with each other.

Next to the requirements that not more than one of the premises may be negative and not more than one premise may be specific, there is a third requirement: the middle term must be 'distributed' in at least one of the two premises.[22] What does that mean? As we have seen previously, the terms in categorical premises stand for classes. But while they sometimes stand for the entire class, they may also sometimes only stand for several elements or a part of a class. We can best see this by having a closer look at a-, i-, e- and o-propositions separately:

- A proposition of the form SaP (for example 'All thieves are criminals') says something about the entire class S (therefore about all thieves), but not about all elements of class P (about all criminals). After all, there may be criminals who are not thieves and the proposition does not

[22] This requirement is narrowly connected with the requirement that not both of the premises may be specific.

contain any information about those. Since the proposition has something to say about all elements of class S, the S-term in the proposition is distributed. The P-term is not distributed.

- A proposition of the form SiP (for example 'Some banks engage in money laundering') does not say something about the entire class S (banks). Likewise, the proposition has only something to say about some, not all, elements of class P (those who engage in money laundering). This means that in a SiP-proposition neither the S-term nor the P-term is distributed.

- A proposition of the form SeP (for example 'There are no blue elephants') says something about the entire class S (namely that it falls entirely outside class P; not a single elephant is blue) and also about the entire class P (namely that it falls outside class S; not a single thing that is blue is an elephant). This means that in an SeP-proposition both the S-term and the P-term are distributed.

- A proposition of the form SoP (for example 'Some managers do not crave power') does not say something about the entire class S (all managers) but only about some elements of it (namely that they do not crave power). The S-term in this proposition is therefore not distributed. However, the P-term is distributed. After all, the proposition does say something about the entire class P (those who crave power), namely that it does not overlap with some elements of class S (some managers).

In order for the connecting class of a categorical syllogism to connect the S-class and the P-class in such a way that a valid conclusion can be drawn, at least one of the two premises must say something about *all* elements of the connecting class. In other words, the middle term which stands for the connecting class must be distributed in at least one of the two premises. Therefore:

A categorical syllogism can only be valid if the middle term is distributed in at least one of the two premises.

This is a necessary condition, but not a sufficient one. Moreover, the rule can be made even more precise. because there is still a second requirement concerning the distribution of the middle term: it cannot be distributed in both premises.

This requirement is hard to explain. Many arguments that are invalid because the middle term is distributed in both premises are invalid because one of the other rules is violated. They are for instance arguments with two negatory premises. However, sometimes it seems to be a mere 'coincidence' that an arguments which has two distributed middle terms is invalid. An example would be the following argument:

Everyone who benefited from an inheritance is rich.
Some house-owners did not benefit from an inheritance.

Some house-owners are not rich.

RULE 8: A TERM THAT IS DISTRIBUTED IN THE CONCLUSION MUST BE DISTRIBUTED IN THE PREMISE WHICH CONTAINS THIS TERM

If a term is distributed in the conclusion, the conclusion contains information about all elements of the class denoted by that term. Then the premises must also contain information about all elements of that class. This information can only be contained in the premises which uses that term, and the term must therefore be distributed in that premise. In the following example argument this rule is violated:

All honest people are lawyers. All lawyers studied law. -- Everybody who studied law is honest.

The conclusion of this argument says something about everybody who has studied law. The second premises says something about some who studied law, but not necessarily about everybody who did so. Therefore the conclusion contains information that is not contained in the premises, and the argument is therefore invalid.

A COMPLETE TEST

The 8 rules explained above constitute a complete test for the validity of categorical syllogisms. If one of the rules is violated, the syllogism is invalid; if none of the rules is violated, it is valid.

Exercises

1. For each of the following modi, determine whether it is valid. If the modus is invalid, state which rule(s) it violates.

   a. PaM, MaS, therefore SaP

   b. PoM, MaS, therefore SoP

   c. MoP, SeM, therefore SoP

   d. PiM, SoM, therefore SoP

   e. PaM, SiM, therefore SiP

2. Assess the validity of each of the following arguments. If the argument is invalid, state which rule(s) it violates.

   a. The judge may order Jolande to pay compensation, since Jolanda has unlawfully caused damage and the judge may order persons who have unlawfully caused damage to pay compensation.

b. No administrator acts outside her competences. Therefore, some administrative acts are not punishable, since persons who act without being competent are sometimes punishable.

c. Members of the government may not also be mayor of a municipality. Since mayors are elected, elected persons cannot be members of the government.

d. Some civil servants are can be bribed. Therefore, some civil servants are corrupt, since some corrupt people can be bribed.

e. Members of left-wing political parties do not understand economics. Members of right-wing political parties have no sense of social responsibility. Therefore, persons with a sense of social responsibility do not understand economics.