Chapter 2

Truth Functions – or Not?

Whether or not the rules of validity are hard-wired into us, we all have pretty strong intuitions about the validity or otherwise of various inferences. There wouldn't be much disagreement, for example, that the following inference is valid: 'She's a woman and a banker; so she's a banker'. Or that the following inference is invalid: 'He's a carpenter; so he's a carpenter and plays baseball'.

But our intuitions can get us into trouble sometimes. What do you think of the following inference? The two premisses occur above the line; the conclusion below it.

The Queen is rich.   The Queen isn't rich.
___
Pigs can fly.

It certainly doesn't seem valid. The wealth of the Queen – great or not – would seem to have no bearing on the aviatory abilities of pigs.

But what do you think about the following two inferences?

The Queen is rich.
___
Either the Queen is rich or pigs can fly.

Either the Queen is rich or pigs can fly.   The Queen isn't rich.

Pigs can fly.

The first of these seems valid. Consider its conclusion. Logicians call sentences like this a *disjunction*; and the clauses on either side of the 'or' are called *disjuncts*. Now, what does it take for a disjunction to be true? Just that one or other of the disjuncts is true. So in any situation where the premiss is true, so is the conclusion. The second inference also seems valid. If one or other of two claims is true and one of these isn't, the other must be.

Now, the trouble is that by putting these two apparently valid inferences together, we get the apparently invalid inference, like this:

The Queen is rich.

Either the Queen is rich or pigs can fly.   The Queen isn't rich.

Pigs can fly.

This can't be right. Chaining valid inferences together in this way can't give you an invalid inference. If all the premises are true in any situation, then so are their conclusions, the conclusions that follow from *these*; and so on, till we reach the final conclusion. What has gone wrong?

To give an orthodox answer to this question, let us focus a bit more on the details. For a start, let's write the sentence 'Pigs can fly' as p, and the sentence 'The Queen is rich' as q. This makes things a bit more compact; but not only that: if you think about it for a moment, you can see that the two particular sentences actually used in the examples above don't have much to do with things; I could have set everything up using pretty much any two sentences; so we can ignore their content. This is what we do in writing the sentences as single letters.

The sentence 'Either the Queen is rich or pigs can fly' now becomes 'Either q or p'. Logicians often write this as $q \vee p$. What of 'The Queen isn't rich'? Let us rewrite this as 'It is not the case that the Queen is rich', pulling the negative particle to the front of the sentence. Hence, the sentence becomes 'It is not the case that *q*'. Logicians often write this as $\neg q$, and call it the *negation* of *q*. While we are at it, what about the sentence 'The Queen is rich *and* pigs can fly', that is, '*q* and *p*'? Logicians often write this as $q \,\&\, p$ and call it the *conjunction* of *q* and *p*, *q* and *p* being the *conjuncts*. With this machinery under our belt, we can write the chain-inference that we met thus:

$$\frac{\dfrac{q}{q \vee p} \quad \neg q}{p}$$

What are we to say about this inference?

Sentences can be true, and sentences can be false. Let us use *T* for truth, and *F* for falsity. After one of the founders of modern logic, the German philosopher/mathematician Gottlob Frege, these are often called *truth values*. Given any old sentence, *a*, what is the connection between the truth value of *a* and that of its negation, $\neg a$? A natural answer is that if one is true, the other is false, and vice versa. Thus, if 'The Queen is rich' is true, 'The Queen isn't rich' is false, and vice versa. We can record this as follows:

> $\neg a$ has the value *T* just if *a* has the value *F*.
> $\neg a$ has the value *F* just if *a* has the value *T*.

Logicians call these the *truth conditions* for negation. If we assume that every sentence is either true or false, but not both, we can depict the conditions in the following table, which logicians call a *truth table*:

a	$\neg a$
T	*F*
F	*T*

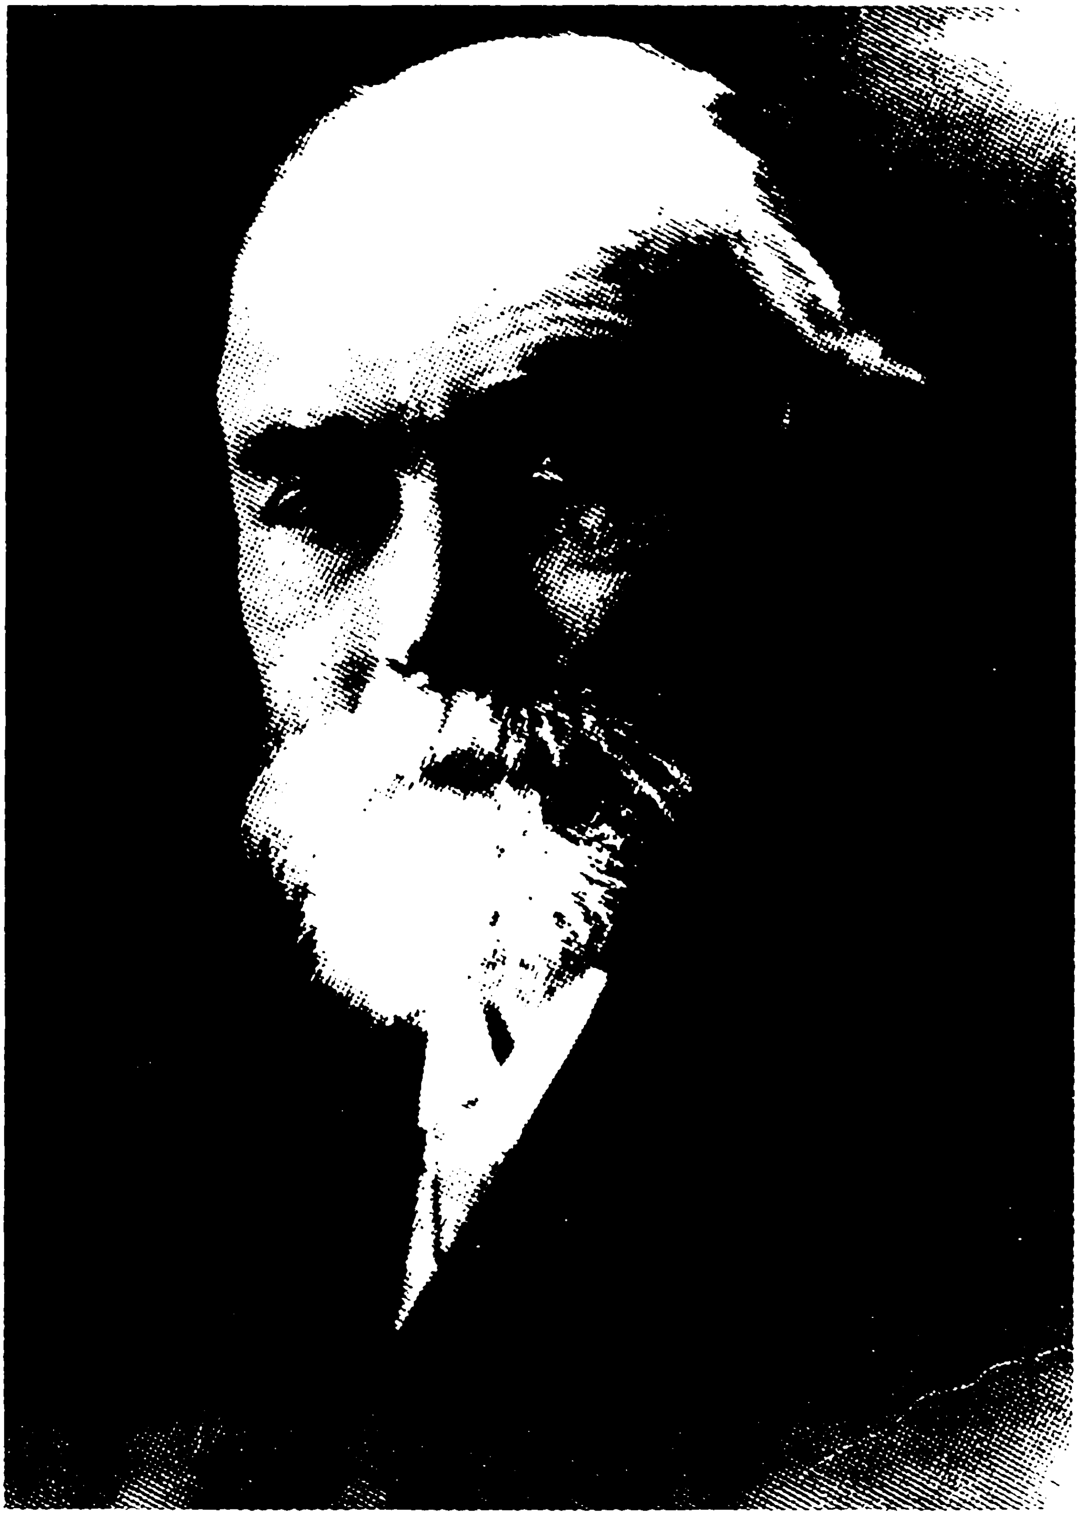

2. Gottlob Frege (1848–1925), one of the founders of modern logic

If a has the truth value given in the column under it, $\neg a$ has the corresponding value to its right.

What of disjunction, V? As I have already noted, a natural assumption is that a disjunction, $a \vee b$, is true if one or other (or maybe both) of a and b are true, and false otherwise. We can record this in the truth conditions for disjunction:

> $a \vee b$ has the value T just if at least one of a and b has the value T.
> $a \vee b$ has the value F just if both of a and b have the value F.

These conditions can be depicted in the following truth table:

a	b	$a \vee b$
T	T	T
T	F	T
F	T	T
F	F	F

Each row - except the first, which is the header - now records a possible combination of the values for a (first column) and b (second column). There are four such possible combinations, and so four rows. For each combination, the corresponding value of $a \vee b$ is given to its right (third column).

Again, while we are about it, what is the connection between the truth values of a and b, and that of a & b? A natural assumption is that a & b is true if both a and b are true, and false otherwise. Thus, for example, 'John is 35 and has brown hair' is true just if 'John is 35' and 'John has brown hair' are both true. We can record this in the truth conditions for conjunction:

> a & b has the value T just if both of a and b have the value T.
> a & b has the value F just if at least one of a and b has the value F.

These conditions can be depicted in the following truth table:

a	b	$a \& b$
T	T	T
T	F	F
F	T	F
F	F	F

Now, how does all this bear on the problem we started with? Let us come back to the question I raised at the end of the last chapter: what is a situation? A natural thought is that whatever a situation is, it determines a truth value for every sentence. So, for example, in one particular situation, it might be true that the Queen is rich and false that pigs can fly. In another it might be false that the Queen is rich, and true that pigs can fly. (Note that these situations may be purely hypothetical!) In other words, a situation determines each relevant sentence to be either T or F. The relevant sentences here do not contain any occurrences of 'and', 'or' or 'not'. Given the basic information about a situation, we can use truth tables to work out the truth values of the sentences that do.

For example, suppose we have the following situation:

$p : T$
$q : F$
$r : T$

(r might be the sentence 'Rhubarb is nutritious', and '$p : T$' means that p is assigned the truth value T, etc.) What is the truth value of, say, $p \& (\neg r \vee q)$? We work out the truth value of this in exactly the same way that we would work out the numerical value of $3 \times (-6 + 2)$ using tables for multiplication and addition. The truth value of r is T. So the truth table for $\neg$ tells us that the truth value of $\neg r$ is *F*. But since the value of *q* is *F*, the truth table for V tells us that the value of $\neg r \vee q$ is *F*. And since the truth value of *p* is *T*, the truth table for & tells us that the value of $p \,\&\, (\neg r \vee q)$ is *F*. In this step-by-step way, we can work out the truth value of any formula containing occurrences of &, V, and $\neg$.

Now, recall from the last chapter that an inference is valid provided that there is no situation which makes all the premisses true, and the conclusion untrue (false). That is, it is valid if there is no way of assigning *T*s and *F*s to the relevant sentences, which results in all the premisses having the value *T* and the conclusion having the value *F*. Consider, for example, the inference that we have already met, $q / q \vee p$. (I write this on a single line to save Oxford University Press money.) The relevant sentences are *q* and *p*. There are four combinations of truth values, and for each of these we can work out the truth values for the premiss and conclusion. We can represent the result as follows:

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

The first two columns give us all the possible combinations of truth values for *q* and *p*. The last two columns give us the corresponding truth values for the premiss and the conclusion. The third column is the same as the first. This is an accident of this example, due to the fact that, in this particular case, the premiss happens to be one of the relevant sentences. The fourth column can be read off from the truth table for disjunction. Given this information, we can see that the inference is valid. For there is no row where the premiss, *q*, is true and the conclusion, $q \vee p$, is not.

What about the inference $q \vee p$, $\neg q / p$? Proceeding in the same way, we obtain:

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

This time, there are five columns, because there are two premisses. The truth values of the premisses and conclusion can be read off from the truth tables for disjunction and negation. And again, there is no row where both of the premisses are true and the conclusion is not. Hence, the inference is valid.

What about the inference with which we started: q, $\neg q / p$? Proceeding as before, we get:

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

Again, the inference is valid; and now we see why. There is no row in which both of the premisses are true and the conclusion is false. Indeed, there is no row in which both of the premisses are true. The conclusion doesn't really matter at all! Sometimes, logicians describe this situation by saying that the inference is *vacuously* valid, just because the premisses could never be true together.

Here, then, is a solution to the problem with which we started. According to this account, our original intuitions about this inference were wrong. After all, people's intuitions can often be misleading. It seems obvious to everyone that the earth is motionless – until they take a course in physics, and find out that it is really hurtling through space. We can even offer an explanation as to why our logical intuitions go wrong. Most of the inferences we meet in practice are not of the vacuous kind. Our intuitions develop in this sort of context, and don't apply generally – just as the habits you build up learning to walk (for example, not to lean to the side) don't always work in other contexts (for example, when you to learn to ride a bike).

We will come back to this matter in a later chapter. But let us end this one with a brief look at the adequacy of the machinery we have used. Things here are not as straightforward as one might have hoped. According to this account, the truth value of a sentence $\neg a$ is completely determined by the truth value of the sentence *a*. In a similar way, the truth values of the sentences $a \vee b$ and *a* & *b* are completely determined by the truth values of *a* and *b*. Logicians call operations that work like this *truth functions*. But there are good reasons to suppose that 'or' and 'and', as they occur in English, are not truth functions – at least, not always.

For example, according to the truth table for &, '*a* and *b*' always has the same truth value as '*b* and *a*': namely, they are both true if *a* and *b* are both true, and false otherwise. But consider the sentences:

1. John hit his head and fell down.
2. John fell down and hit his head.

The first says that John hit his head *and then* fell down. The second says that John fell down *and then* hit his head. Clearly, the first could be true whilst the second is false, and vice versa. Thus, it is not just the truth values of the conjuncts that are important, but which conjunct caused which.

Similar problems beset 'or'. According to the account we had, '*a* or *b*' is true if one or other of *a* and *b* is true. But suppose a friend says:

Either you come now or we will be late;

and so you come. Given the truth table for V, the disjunction is true. But suppose you discover that your friend had been tricking you: you could have left half an hour later and still been on time. Under these circumstances, you would surely say that your friend had lied: what he had said was false. Again, it is not merely the truth values of the disjuncts that are important, but the existence of a connection of a certain kind between them.

I will leave you to think about these matters. The material we have been looking at gives us at least a working account of how certain logical machinery functions; and we will draw on this in succeeding chapters, unless the ideas in those chapters explicitly override it - which they will sometimes.

The machinery in question deals only with certain kinds of inferences: there are many others. We have only just started.

Main Ideas of the Chapter

- **In a situation, a unique truth value (*T* or *F*) is assigned to each relevant sentence.**
- **$\neg a$ is *T* just if *a* is *F*.**
- **$a \vee b$ is *T* just if at least one of *a* and *b* is *T*.**
- **a & b is *T* just if both of *a* and *b* are *T*.**