1

Consequences

Everybody, sooner or later,
sits down to a banquet of consequences.
– Robert Louis Stevenson

'Watch what you say,' my mother often advised, 'because what you say has consequences.' She was right, and doubly so. There are two senses in which what one says has consequences. One sense, not terribly relevant for present purposes, is captured in the familiar dictum that actions have consequences. To say something is to do something, and doing something is an action. Actions, in turn, are events, and events, as experience tells, have consequences, namely, their causal effects. (Example: a consequence—a causal effect—of your drinking petrol is your being ill, at least other things being equal.) So, in the *causal effects* sense of 'consequences', my mother was perfectly right, but that sense of 'consequence' has little to do with logic.

For present purposes, there is a more relevant sense in which what one says has consequences. What one says, at least in the declarative mode,[1] has *logical consequences*, namely, whatever *logically follows from* what one said, or whatever is *logically implied* by what one said. Suppose, for example, that you're given the following information.

1. Agnes is a cat.
2. All cats are smart.

[1]For purposes of this book, a *declarative sentence* (or a sentence used in the 'declarative mode') is one that is used (successfully or otherwise) to declare or state something about the world. This is hardly a precise definition, but it'll do. (Example. Each of 'You are reading a book', 'Obama is the first black US president', and '1+1=5' are declarative sentences, but sentences such as 'Shut that door!' and 'Do you like Vegemite?' are not declarative, since they fail to declare or state anything about the world.)

A consequence of (1) and (2), taken together, is that Agnes is smart. In other words, *that Agnes is smart* logically follows from (1) and (2); it is implied by (1) and (2), taken together.

1.1 Relations of support

Logical consequence is a relation on sentences of a language, where 'sentence', unless otherwise indicated, is short for 'meaningful, declarative sentence'.[2]

Logical consequence is one among many relations over the sentences of a language. Some of those relations might be called *relations of support*. For example, let $A_1, \ldots, A_n$ and B be arbitrary sentences of some given language—say, English. For some such sentences, the various A_i jointly 'support' B in the following sense.

S1. If all of $A_1, \ldots, A_n$ are true, then B is *probably* true.

Consider, for example, the following sentences.

3. Max took a nap on Day 1.
4. Max took a nap on Day 2.
5. Max took a nap on Day 3.

⋮

n. Max took a nap on Day n (viz., today).

m. Max will take a nap on Day $n+1$ (viz., tomorrow).

On the surface, sentences (3)–(n) support sentence (m) in the sense of (S1): the former, taken together, make the latter more likely. Similarly, (6) supports (7) in the same way.

6. The sun came up every day in the past.
7. The sun will come up tomorrow.

If (6) is true, then (7) is probably true too.

The relation of support given in (S1) is important for empirical science and, in general, for rationally navigating about our world. Clarifying the (S1) notion of 'support' is the job of probability theory (and, relatedly, decision theory), an area beyond the range of this book.

[2] Taking consequence to be relation on sentences simplifies matters a great deal, and sidesteps the issue of so-called 'truth bearers', an ongoing issue in philosophy of logic. For present purposes, simplicity is worth the sidestep.

1.2 Logical consequence: the basic recipe

Logical consequence, the chief topic of logic, is a stricter relation of support than that in (S1). Notice, for example, that while (7) may be very likely true if (6) is true, it is still possible, in some sense, for (6) to be true without (7) being true. After all, the sun might well explode later today.

While (S1) might indicate a strong relation of support between some sentences and another, it doesn't capture the tightest relation of support. Logical consequence, on many standard views, is often thought to be the tightest relation of support over sentences of a language. In order for some sentence B to be a logical consequence of sentences $A_1, \ldots, A_n$, the truth of the latter needs to 'guarantee' the truth of the former, in some suitably strong sense of 'guarantee'.

Throughout this book, we will rely on the following (so-called semantic) account of logical consequence, where $A_1, \ldots, A_n$ and B are arbitrary sentences of some given language (or fragment of a language).

Definition 1. (Logical Consequence) *B is a logical consequence of $A_1, \ldots, A_n$ if and only if there is no case in which $A_1, \ldots, A_n$ are all true but B is not true.*

Notice that the given 'definition' has two parts corresponding to the 'if and only if' construction, namely,

- If B is a logical consequence of $A_1, \ldots, A_n$, then there is no case in which $A_1, \ldots, A_n$ are all true but B is not true.
- If there is no case in which $A_1, \ldots, A_n$ are all true but B is not true, then B is a logical consequence of $A_1, \ldots, A_n$.

Also notable is that the given 'definition' is really just a recipe. In order to get a proper definition, one needs to specify two key ingredients:

- what 'cases' are;
- what it is to be *true in a case.*

Once these ingredients are specified, one gets an account of logical consequence. For example, let $A_1, \ldots, A_n$ and B be declarative sentences of English. If we have a sufficiently precise notion of *possibility* and, in turn, think of 'cases' as such *possibilities*, we can treat 'true in a case' as 'possibly true' and get the following account of logical consequence—call it 'necessary consequence'.

- B is a *(necessary) consequence* of $A_1, \ldots, A_n$ if and only if there is no possibility in which $A_1, \ldots, A_n$ are all true but B is not true. (In other words, B is a consequence of $A_1, \ldots, A_n$ if and only if it is impossible for each given A_i to be true without B being true.)

Presumably, this account has it that, as above, 'Agnes is smart' is a consequence of (1) and (2). After all, presumably, it's not possible for (1) and (2) to be true without 'Agnes is smart' also being true. On the other hand, (7) is not a necessary consequence of (6), since, presumably, it is possible for (6) to be true without (7) being true.

Of course, taking 'cases' to be 'possibilities' requires some specification of what is possible, or at least some class of 'relevant possibilities'. The answer is not always straightforward. Is it possible to travel faster than the speed of light? Well, it's not *physically* possible (i.e., the physical laws prohibit it), but one might acknowledge a broader sense of 'possibility' in which such travel is possible—for example, *coherent* or *imaginable* or the like. If one restricts one's 'cases' to only physical possibilities, one gets a different account of logical consequence from an account that admits of possibilities that go beyond the physical laws.

In subsequent chapters, we will be exploring different logical theories of our language (or fragments of our language). A logical theory of our language (or a fragment thereof) is a theory that specifies the logical consequence relation over that language (or fragment). Some fragments of our language seem to call for some types of 'cases', while other fragments call for other (or additional) types. Subsequent chapters will clarify this point.

1.3 Valid arguments and truth

In general, theses require arguments. Consider the thesis that there are feline gods. Is the thesis true? An argument is required. Why think that there are feline gods? We need to examine the argument—the reasons that purport to 'support' the given thesis.

Arguments, for our purposes, comprise premises and a conclusion. The latter item is the thesis in question; the former purport to 'support' the conclusion. Arguments may be evaluated according to any relation of support (over sentences). An argument might be 'good' relative to some relation of support, but not good by another. For example, the argument from (6) to (7) is a good argument when assessed along the lines of (S1); however, it is not good when assessed in terms of (say) necessary consequence, since, as noted above, (7) is not a necessary consequence of (6).

In some areas of rational inquiry, empirical observation is often sufficient to figure out the truth. Suppose that you want to know whether there's a cat on the table. One reliable method is handy: look at the table and see whether there's a cat on it! Of course, 'real empirical science' is much more complicated than checking out cats, but empirical observation—empirical testing—is nonetheless a critical ingredient.

What about other pursuits for which there is little, if any, opportunity for settling matters by observation? Consider, for example, pure mathematics or philosophy. In such areas, theses cannot be empirically tested, at least in general. How, then, do we figure out the truth in such areas? Argument is the only recourse.

When argument is the only recourse, as in mathematics or (at least much of) philosophy, it makes sense to invoke the strictest relation of support—namely, logical consequence. Traditionally, an argument is said to be *valid*—strictly speaking, *logically valid*—if its conclusion is a logical consequence of its premises. We will follow suit.

Of course, a valid argument needn't be a proof of anything. After all, the 'definition' (or, for now, 'recipe') of *logical consequence* doesn't require that any of the premises be true. Rather, the given account requires only the *absence* of any 'counterexample', where these are defined as follows.

Definition 2. (Counterexample) *A counterexample to an argument is a case in which the premises are true but the conclusion is not true.*

We can say that B is a logical consequence of $A_1, \ldots, A_n$ if and only if there is no counterexample to the argument from (premises) $A_1, \ldots, A_n$ to (conclusion) B. In turn, an argument is *valid* just if there is no counterexample to it.

Accordingly, an argument may be valid—that is, its conclusion be a logical consequence of its premises—even though none of its premises are true. In mathematics and philosophy, validity is a necessary condition on suitable arguments; it is not sufficient. What is sufficient, for such pursuits, is a so-called *sound* argument.

Definition 3. (Sound Argument) *A sound argument is valid and all its premises are true.*

Suppose that, among the 'cases' in our definition of validity (or logical consequence), there is an 'actual case' @ such that A is true-in-@ just if A is true (i.e., actually true). On such an account, every sound argument has a true conclusion. After all, a sound argument, by definition, has all true premises. By supposition, a sentence is true just if true-in-@, and so all premises of a sound argument are true-in-@. But a sound argument, by definition, is also valid, and so, by definition, if its premises are true in a case, then so too is its conclusion. Since, as noted, the premises of any sound argument are true-in-@, so too is its conclusion.

Logic, in the end, serves the pursuit of truth; however, it does not principally concern itself with truth. Instead, logic, as above, has its chief concern with consequence—logical consequence. Logic aims to precisely specify valid arguments. Once the valid arguments are in order, rational inquiry may proceed to discern the sound arguments. For our purposes in this book, we will focus on different accounts of logical consequence, and some of the phenomena that motivate the various accounts.

1.4 Summary, looking ahead, and reading

Summary. Logical consequence is the chief concern of logic. An argument is *valid* just if its conclusion is a logical consequence of its premises. Logical consequence, in this book, will be understood as absence of counterexample, where a *counterexample* is a 'case' in which all the premises are true but the conclusion not true. One of the chief concerns of logic, broadly construed, is to figure out which 'cases' are involved in specifying the consequence relation on a given language (or fragment thereof). In subsequent chapters, we will look at different accounts of logical consequence—different logical theories of our language (or fragments thereof)—and some of the phenomena that have motivated them.

Looking Ahead. The next two chapters are devoted to stage-setting. Chapter 2 discusses features of language that are relevant to logic, and also discusses the general 'model-building' enterprise of formal logic. Chapter 3 briefly—and, for the most part, informally—introduces some useful set-theoretic notions. These two chapters will make subsequent discussion easier.

Further Reading. For related, accessible discussion of logic, see Read 1995, Haack 1978; Haack 1996. (And see the bibliographies therein for a host of other sources!) For a more advanced discussion of the 'recipe' of logical consequence, see Beall and Restall 2005.

Exercises

1. What is an argument?
2. What is a valid argument?
3. What is a sound argument?
4. What is the general 'recipe' for defining logical consequence (or validity)? What are the two key ingredients that one must specify in defining a consequence relation?
5. Consider the 'necessary consequence' relation, which takes cases to be possibilities. Assume, as is reasonable (!), that our actual world is possible—that is, that whatever is true (actually true) is possibly true. Question: on this account of logical consequence, are there any sound arguments that have false conclusions? If so, why? If not, why not?
6. As noted in the chapter, 'if and only if' (which is often abbreviated as 'iff') expresses two conditionals: 'A iff B' expresses both of the following conditionals.[3]
    - If A, then B.
    - If B, then A.

[3]Strictly speaking, what is expressed is the 'conjunction' of the two conditionals, but we leave the notion of *conjunctions* for the next chapter.

For our purposes, a biconditional '*A* iff *B*' is true so long as *A* and *B* are either both true or both false (and such biconditionals are false otherwise). With this in mind, consider the necessary consequence relation. Is the following argument valid (where, here, validity is necessary consequence)? If it is valid—if its conclusion is a necessary consequence of the premises—explain why it is valid. If not, explain why not.

(a) Max is happy if and only if Agnes is sleeping.
(b) Agnes is sleeping.
(c) Therefore, Max is happy.

What about the following argument?

(d) Max is happy if and only if Agnes is sleeping.
(e) Agnes is not sleeping.
(f) Therefore, Max is not happy.

7. Using the 'necessary consequence' account of validity, specify which of the following arguments are valid or invalid. Justify your answer.
   (a) Argument 1.
      i. If Agnes arrived at work on time, then her car worked properly.
      ii. If Agnes's car worked properly, then the car's ignition was not broken.
      iii. The car's ignition was not broken.
      iv. Therefore, Agnes arrived at work on time.
   (b) Argument 2.
      i. Either the sun will rise tomorrow or it will explode tomorrow.
      ii. The sun will not explode tomorrow.
      iii. Therefore, the sun will rise tomorrow.
   (c) Argument 3.
      i. If Max wins the lottery, then Max will be a millionaire.
      ii. Max will not win the lottery.
      iii. Therefore, Max will not be a millionaire.

(d) Argument 4.
   i. If Beetle is an extraterrestrial, then Beetle is not from earth.
   ii. Beetle is an extraterrestrial.
   iii. Therefore, Beetle is not from earth.

Sample answers

Answer 5. On the necessary-consequence sense of 'validity' (the sense in question), an argument is *valid* iff every possibility (e.g., possible circumstance) in which the premises are all true is one in which the conclusion is true. Hence, if the actual world—the 'real' world, the way things really are—counts as a possibility, then it itself cannot be a case in which the premises of a valid argument are true but the conclusion false. But, then, any *sound* argument—that is, a valid argument whose premises are all (actually) true—is one in which the conclusion is true, and so not false.[4]

Answer 6. The argument from (6a) and (6b) to (6c) is valid in the necessary-consequence approach to validity: it is not possible for both of (6a) and (6b) to be true without (6c) being true. After all, recall that (6a) expresses not only that *if Max is happy then Agnes is sleeping*; it also expresses that *if Agnes is sleeping then Max is happy*. Now, consider any possibility (and possible circumstance) in which both (6a) and (6b) are true, that is, a possible circumstance in which not only Agnes is sleeping, but *if Agnes is sleeping (in that circumstance), then Max is happy (in that circumstance)*. Well, then, no matter what possible circumstance we choose, it'll be one in which Max is happy if it's one in which both (6a) and (6b) are true. (Of course, there are, presumably, many possibilities in which neither (6a) nor (6b) are true, but this does not affect the necessary-consequence sense in which the given argument is valid. Why?)

[4] This last step—from *true* to *not false*—is something that some logical theories reject, but these theories are left for later chapters.