CHAPTER 1

Agenda

> [T]he understanding of "if" is not a narrow academic concern, but a matter of central importance in the understanding of what makes human intelligence special and distinctive.
>
> Jonathan Evans and David Over, *If* (p. 153)

Conditionals are sentences of the forms "If φ, [then] ψ" and "ψ if φ," such as

(1.1) a. If the village is flooded, then the dam must have broken.
b. If Henry had come to the party, Sue would have come too.
c. Paul would have bought the house if it hadn't been so expensive.

One may also classify as conditionals sentences that can be naturally put in the above forms, such as

(1.2) a. They will leave in an hour, unless John changes his mind.
b. No guts, no glory.
c. Give Louis a toy and he'll ruin it.

which can be rephrased as, respectively, (1.3a), (1.3b), and (1.3c):

(1.3) a. If John does not change his mind, they will leave in an hour.
b. If a person lacks courage, there will be no glory for him or her.
c. If one gives Louis a toy, he ruins it.

In "If φ, then ψ," φ is called the "antecedent" and ψ the "consequent."

Conditionals are *special*. They are special for a number of reasons, but probably most conspicuously for the heated controversy that they have generated, and continue to generate. Not that controversy is anything out of the ordinary in philosophy. But even in philosophy, controversies

commonly take place against a shared background of basic assumptions. For instance, while there is ongoing controversy about the concept of knowledge, there is at the same time broad (if not universal) agreement on many core issues surrounding that concept. Few dispute that knowledge is factive; that it requires belief as well as justification; that justified true belief is not sufficient for knowledge, however; that coherence amongst one's beliefs is not enough to elevate these to the status of knowledge; that we can gain knowledge from testimony; and so on and so forth. Not so in the case of conditionals. For almost any claim about conditionals that is not downright trivial, it will be exceedingly hard to find a majority, or even a sizable minority, of philosophers who adhere to it. Nearly every interesting question about conditionals that one can think of must be considered to be still largely open. This is so in spite of centuries of hard work on the topic of conditionals by leading philosophers, and more recently also by linguists, psychologists, and (increasingly) computer scientists.

To get an impression of just how vastly theorists disagree on conditionals, consider a long-standing dispute about the semantics of conditionals. Some philosophers who have thought seriously about conditionals hold that conditionals like

(1.4) If it continues to rain, the match will be canceled.

have the truth conditions of the corresponding material conditionals. That is to say, according to these philosophers, (1.4) and conditionals like it are false if their antecedent is true and their consequent false, and are true in all other cases, so that (1.4) would be logically equivalent to

(1.5) Either it won't continue to rain or the match will be canceled (or both).

Some of the same philosophers hold that conditionals like

(1.6) If it had continued to rain, the match would have been canceled.

are just loose talk and do not possess truth conditions at all. But then there are other philosophers who, having thought equally seriously about the matter, hold that conditionals like (1.4) *lack* truth conditions while conditionals like (1.6) do have truth conditions. (Most of these philosophers think the latter have the truth conditions attributed to them by Stalnaker's

possible worlds semantics, to be encountered later on.) This is not a minor disagreement![1]

On the one hand, you will have suspected that conditionals are special even if you are a novice to this area of research. You may have found this book shouldering ones entitled *If* (Evans and Over [2004]), *Ifs* (Harper, Stalnaker, and Pearce [1981]), *A Philosophical Guide to Conditionals* (Bennett [2003]), and several books with the title *Conditionals* (e.g., Jackson [1987], Jackson (ed.) [1991], Crocco, Fariñas del Cerro, and Herzig (eds.) [1995], Woods [1997]), with no book called *And*, or *Ands*, or *A Philosophical Guide to Conjunctions*, or *Conjunctions* (or *Or*, etc.) in sight.[2]

On the other hand, however, it *is* surprising that conditionals are special. Conditionals may not be quite as common as conjunctions or disjunctions in daily parlance,[3] but they are common enough not to strike us as exotic beasts – not, at least, until we start viewing them from a theoretical perspective.

The question of why conditionals are special, of why there is all this disagreement about them, is not so easy to answer. In metaphysics, there is a lot of controversy about the notion of an object, even though, one would have supposed, we all know perfectly well what objects are, having been surrounded by them for all our lives. And yet it is a heavily debated question whether objects are three-dimensional or four-dimensional entities; and equally whether only the most elementary physical particles are really objects or whether, on the contrary, those particles do not deserve to be called "objects" in the first place; and whether objects are nothing over and above the properties they possess or whether there is some substratum to which those properties adhere;[4] and so on. But at least here it is reasonable to conjecture that much of the controversy has arisen in

[1] A prominent example of the first type of philosopher is Quine [1959]. Prominent examples of the second type are Gibbard [1981] and Bennett [2003].

[2] Though, admittedly, you may have come across *A Natural History of Negation* (Horn [1989]), *The Syntax of Negation* (Haegeman [2005]), and *The Genealogy of Disjunction* (Jennings [1994]).

[3] For what it is worth, comparisons of "if" with "and" and "or" on Google Fight (`http://www.googlefight.com`), which allows for a kind of poor man's version of corpus research, showed "if" to occur about three times more often on webpages and in documents accessible via the Internet than "and," although "if" apparently occurs only about half as often as "or." On the other hand, Google's Ngram Viewer (`http://books.google.com/ngrams`), when asked to search for the occurrence of "if," "and," and "or" in books that were published after 1800 and that are available online, gives a very different picture: then "and" comes out as being used much more frequently than "or," which is used somewhat more frequently than "if"; this outcome is stable over the whole period of time from 1800 onward. (These comparisons were made on January 29, 2013. Thanks to Sylvia Wenmackers for the pointer.)

[4] It is even debated whether there is any real difference between the bundle-of-properties view and the substratum view; see Benovsky [2008].

response to developments in physics, which has continued to reveal truths that tend to sit badly with our ordinary conception of objects (like, in the relatively recent history of physics, the idea that ordinary objects are mostly empty space). No such easy explanation is available for the controversy (or controversies) surrounding conditionals.

An observation made by experimental psychologists may go some way toward explaining why conditionals are special. It has been well documented that not only do young children have difficulty comprehending conditionals, but also that in adults, IQ is a good predictor of how well people comprehend conditionals. Specifically, young children and low-IQ adults tend to interpret conditionals as conjunctions (see Chapter 3). No such difficulties, or differences between cognitively more and less able people, have been found in the interpretation of conjunctions, for example. This suggests that interpreting conditionals is cognitively more demanding than interpreting most other types of sentences. Perhaps interpreting conditionals involves unique mechanisms that science is still to uncover.[5]

A further explanatory factor may be that conditional constructions are put to a great variety of uses (I say more about this below) and that philosophers, being naturally inclined to look for the most general theories, have tended to overgeneralize in the case of conditionals. The diversity of uses to which the word "if" is put is likely to cause trouble even if one resists the inclination. We may try the best we can to disentangle and keep separate the different uses, and yet the risk remains that intuitions related to one type of use contaminate our thinking about other types.

1.1 Key questions

As said, nearly *anything* about conditionals is controversial. This is so even with regard to the most fundamental questions concerning conditionals, which can be grouped into five clusters:

1. The first of these concerns *truth*. Specifically, what are the truth conditions of conditionals? Do conditionals have truth conditions at all? If so, should there be a uniform semantics for all conditionals, or may different types of conditionals have different truth conditions? Or may some types have truth conditions while others lack them?

[5] According to Récanati [2000:57 f], the cognitive complexity of conditionals is due to the fact that their evaluation requires us to go through a process of simulation. This idea is related to the Ramsey test, which we shall come across at various later points in this book.

We are strongly inclined (at least, I am) to think that (1.7a) is true and (1.7b) is false:

(1.7) a. If global warming continues unabated, various species will become extinct.
b. If Lance Armstrong had admitted to doping earlier, he would not have been stripped of his seven Tour de France titles.

But however strong our inclinations, the failure of literally every extant truth-conditional semantics of conditionals to accommodate the various registered facts about conditionals – how we reason with conditionals, when we are willing to accept or assert them and when not, the probabilities we accord to them, and so on – constitutes something of an inductive argument against the idea that conditionals are true or false. It could be argued that, given the right division of labor between semantics and pragmatics, and possibly also given some pragmatic principles still to be discovered, some of our present semantics of conditionals will eventually get the relevant facts right. Even granting this, later on we shall come across considerations that appear to impinge more severely on the viability of truth-conditional semantics of conditionals.

So suppose, for now, that conditionals are not truth-conditional. Then what are we to make of the following apparent facts?

- We seem to believe many conditionals. But to believe something is to believe it to be true!
- We seem to know many conditionals. But knowledge requires truth!
- It seems that there can be evidence in favor of a conditional. But favoring evidence is understood as an indication of truth!
- Conditionals seem meaningful, and we seem to know their meanings. But according to the mainstream in semantics, the meaning of a sentence is given by its truth conditions, and knowing the meaning of a sentence is knowing its truth conditions!

This is merely scratching the surface, but it provides a sense of just how much is at stake in the debate concerning the semantics of conditionals.

2. The second cluster of questions concerns the *probabilities* of conditionals. What *are* these probabilities? Can we reasonably attribute probabilities to conditionals even if conditionals lack truth conditions? How much weight is to be given to actual data about people's judgments of the probabilities of conditionals? Might these judgments be systematically mistaken?

As will be seen, a number of the main semantics of conditionals make predictions about the probabilities of conditionals that not only clash with pretheoretic intuition but are also inconsistent with the outcomes of various recent experiments investigating people's probability assignments. What is more, one of the few issues about conditionals on which philosophers do agree is a claim concerning what these probabilities can *not* be: the probability of a conditional cannot, on formal grounds, be the conditional probability of the consequent given the antecedent, or so it is believed. Surprisingly, the aforementioned experiments appear to show that people do judge the probability of a conditional to be the corresponding conditional probability. According to mainstream thinking, these data are to be explained away, possibly as being due to some systematic bias in people's probability judgments, or to an error in the setup of the experiments, or to people's mistaking probability for something else. We shall see that it may be possible to take the data at face value after all.

3. Questions in the third cluster concern the *use conditions* – the acceptability and assertability conditions – of conditionals. What are these? That is, when are we warranted to accept a conditional, when to assert it? Is there a uniform account of use conditions for all types of conditionals? And what is the relation – if any – between the truth conditions of conditionals (if such there be) and their use conditions? What is the relation – if any – between the probabilities of conditionals and their use conditions?

Here again we find some agreement among philosophers. The agreement is on the degrees of acceptability of so-called simple conditionals, that is, conditionals whose antecedent and consequent are not themselves conditional in form. According to virtually all who have thought about the matter, a simple conditional is acceptable to the extent that its consequent is probable given its antecedent. But recent experimental results show this claim to be descriptively inadequate. That might just be another instance of people failing to obey a norm of rationality, on a par perhaps with the base rate fallacy (Kahneman and Tversky [1973]) and the conjunction fallacy (Tversky and Kahneman [1983]). That response to the data is hard to maintain, however, in the absence of any arguments to the effect that we *ought* to judge a conditional acceptable to the degree to which we believe its consequent, given its antecedent. There are also a fair number of philosophers who agree on the categorical acceptability conditions of conditionals: a conditional is categorically acceptable iff (if and only if) its

consequent is highly probable given its antecedent. Again, it will be seen that this proposal is undermined by linguistic intuition as well as, probably more perniciously, experimental data. What are we to put in the place of these flawed proposals?

And what are we to say about the assertability of conditionals? Most philosophers nowadays hold that whether an assertion is warranted depends on the asserter's epistemic position vis-à-vis what he or she asserts. The main divide is over how strong that epistemic position must be: whether the asserter must know what he or she has asserted or whether something weaker, like rational acceptability, suffices.[6] Those holding that rational acceptability warrants assertion will have no further work to do concerning the assertability of conditionals once the (categorical) acceptability conditions of conditionals have been specified. Insofar as specifying the conditions under which conditionals are known poses no difficulties over and above those associated with specifying their acceptability conditions, the same will be true for those who hold that assertion requires knowledge.

4. Then there are questions regarding the *logic* of conditionals. Which principles for conditional reasoning are valid? For instance, is "if" transitive, in the sense that "If φ, then χ" follows from "If φ, then ψ" and "If ψ, then χ"? Many instances of this schema strike us as embodying impeccable reasoning. For instance, the following argument is seemingly unassailable:

> If Jim passes the exam, his parents will buy him a car.
> If Jim's parents buy him a car, he will go to Italy on vacation this summer.
> ∴ If Jim passes the exam, he will go to Italy on vacation this summer.

However, philosophers have come up with counterexamples to the putative transitivity of the conditional, like this one from Cooper [1978:183]:

> If Brown dies, Jones will win the election on Wednesday.
> If Jones wins the election on Wednesday, Brown will retire.
> ?? ∴ If Brown dies, Brown will retire.

Similar counterexamples have been discovered for quite a number of argument forms that have the same prima facie plausibility as the one considered here.

[6] For defenses of the claim that assertion requires knowledge, see, e.g., Williamson [1996], [2000, Ch. 11], Adler [2002], DeRose [2002], and Turri [2011]. For defenses of the claim that something like rational acceptability suffices for assertion, see, e.g., Douven [2006], [2009], and Lackey [2007].

But for an even greater number of argument forms involving conditionals, no counterexamples are known to exist. Given, then, that certain inferences involving conditionals seem valid, do we have another argument against the view that conditionals do not have truth conditions? After all, logical validity is defined in terms of the preservation of *truth*. How can an inference involving one or more conditionals be said to preserve truth if conditionals are not truth-conditional? As Adams has shown, however, validity need not be defined in terms of truth-preservation but can be defined in terms of the preservation of some other property – like, in his own proposal, high probability: according to Adams, an argument is valid iff whenever its premises are highly probable so is its conclusion.[7] Later in this book, we look at conditional inferences and define validity in terms of acceptability-preservation.[8]

5. And finally there are questions related to *belief change*. How ought one respond doxastically to the acceptance of a conditional? That is, how ought one to revise one's beliefs, or degrees of belief, upon receiving conditional information? Specifically, are traditional mechanisms of belief change (such as Bayes' rule) also adequate for handling the receipt of conditional information, or does this kind of information require its own update rule or rules?

Unlike most of the previous questions, the questions in this cluster are not so much sources of continuing controversy among philosophers. Rather, they mark blind spots; they have been mostly overlooked, willingly or not. This cannot be because these questions appear not to be very pressing. No one will want to deny that we change our beliefs, or degrees of belief, upon receiving conditional information. For instance, upon being informed by a colleague that Millie will not pass the exam if she does not work harder, I may drop my earlier belief that Millie is among our brightest students, or lower considerably the degree to which I believe this. But however ordinary and natural such changes of belief appear, it will be seen that the familiar mechanisms of belief change have great difficulty handling conditional information.

[7] To be more precise, on his account an argument is valid iff the improbability of the conclusion (that is, 1 minus the probability of the conclusion) does not exceed the sum of the improbabilities of the premises. Adams [1966], [1975] proves that all arguments that are classically valid are also valid in his sense; see also Adams and Levine [1975].

[8] Field [2009] argues that it has been a mistake to think that logic is primarily concerned with identifying those argument forms that are truth-preserving. Rather, logic is concerned with identifying *good* arguments, where, on Field's notion of goodness, this brings logic close to being concerned with acceptability-preservation in our sense. I will not press this point in the following.

1.2 Scope of the book

Much time and effort has been spent trying to come up with a semantics of conditionals. This has led to an abundance of different proposals, a small number of which have at times enjoyed some popularity, or even widespread popularity, among the general philosophical or psychological readership. However, at least among specialists, there is currently no account of the truth conditions of conditionals that is favored by more than a tiny minority. Except for the question of the probabilities of conditionals – which, as will be seen, is intimately related to semantic questions – none of the epistemological questions concerning conditionals has received nearly as much attention as the cluster of questions concerning the truth conditions of conditionals. Perhaps this is just an oversight on the part of epistemologists. Or they may have thought that conditionals are unremarkable, epistemologically speaking. That thought would be understandable, for at least prima facie it seems that conditionals can be believed, asserted, known, disbelieved, doubted, and learned, in the same way that, for instance, conjunctions and disjunctions can be believed, asserted, and so on. The thought is wrong, however. For instance, as just mentioned, and as will be seen in detail later on, none of the known mechanisms for incorporating new information into one's belief system seems to yield materially adequate results when applied to conditionals.

Some epistemologists may also have thought it wiser to leave the epistemology of conditionals for a later date, when (it is hoped) we will have a firmer grip on their semantics. However, given that so little progress has been made with regard to stating a satisfactory such semantics, the time may be ripe for a sort of reverse-engineering approach, one that addresses the epistemological questions first and then, armed with new insights and data about the epistemology of conditionals, tackles the semantics with a fresh look, wholly or partly as a kind of optimization problem: which semantics does best in explaining why we assign to conditionals the probabilities we do, why we accept and assert certain conditionals but not others, why we change our beliefs in certain ways and not others upon the receipt of conditional information, and so on? There is no shortage of ideas as to what the truth conditions of conditionals might be. A clearer view of how conditionals function in our epistemic lives may help us choose among those ideas by asking which of them offers the best prospects for explaining the epistemic functioning of conditionals. At a minimum, it may enable us to eliminate some candidate semantics that otherwise (e.g., on grounds of simplicity) look theoretically appealing. Indeed, some of

the clearest evidence against the already briefly mentioned material conditional account will be seen to come from work on the probabilities of conditionals.

This book deals primarily with the epistemology of conditionals. As intimated, some epistemologists may have refrained from considering epistemological questions concerning conditionals because it has seemed to them that these questions cannot be properly addressed without having a working semantics of conditionals at hand. That is not so, as I hope to show. Much can be said about the probabilities of conditionals, about their use conditions, as well as about updating on conditionals, that is independent of what the truth conditions of conditionals are, and independent even of whether conditionals have truth conditions to begin with. Moreover, while we cannot say anything about which inferential principles involving conditionals preserve truth without a semantics of conditionals, given an account of the acceptability conditions of conditionals, we *can* ask which inferential principles preserve acceptability. So, we may for instance ask whether "If φ, then ψ and χ" is acceptable if both "If φ, then ψ" and "If φ, then χ" are acceptable. In this sense, we can – and will – consider the logic of conditionals.

The scope of the book is actually still more specific: it will deal with the epistemology of *indicative* conditionals, and then mainly *normal* ones, and only *simple* ones. It was already stated in Section 1.1 that simple conditionals are conditionals whose antecedent and consequent are not themselves conditional in form. I will explain the other two terms.

It is common practice to group conditionals into two major types, *indicative* conditionals and *subjunctive* conditionals. I alluded to this distinction when, on page 2, I tried to give an impression of the kind of disagreement one finds in the literature on conditionals. We have also come across examples of both types; for instance, (1.1a), (1.4), and (1.7a) are indicative conditionals, and (1.1b), (1.6), and (1.7b) are subjunctive conditionals.

Indicative and subjunctive conditionals are usually distinguished on the basis of the grammatical mood of their antecedent. If the antecedent is in the indicative mood, the sentence is an indicative conditional; if the antecedent is in the subjunctive mood, the sentence is a subjunctive conditional.[9] Given that not all languages have a subjunctive mood,

[9] See, e.g., Bennett [2003:10]; see Gibbard [1981, Sect. 4] for a more precise grammatical characterization of the distinction. For recent evidence from neuroscience relating to the indicative–subjunctive distinction, see Kulakova et al. [2013].